

Fusion 3.1 Documentation

2017-12-07

Table of Contents

- User Manual 1
- 1. Getting Started..... 2
 - 1.1. Quick start 2
 - 1.2. Workflow contexts 5
 - 1.3. Fusion Concepts 7
 - 1.3.1. Collections..... 8
 - Primary and Auxiliary Collections 8
 - System Collections 9
 - Collection Configuration Properties..... 9
 - Collection Profiles..... 10
 - 1.3.2. Fusion Components 11
 - Fusion processes..... 11
 - The Fusion UI..... 13
 - UI service 14
 - Connectors 15
 - Parsers..... 21
 - REST API Services..... 26
 - Solr 27
 - Apache Spark..... 28
 - ZooKeeper 29
 - bin/fusion..... 33
 - 1.3.3. Deployment Types 35
 - Deployment goals..... 35
 - 1.4. Fusion Installation and Upgrades..... 38
 - 1.4.1. Installation..... 38
 - 1.4.2. Upgrades..... 38
 - 1.4.3. Related Topics 38
 - 1.4.4. System Requirements 39
 - Supported Operating Systems 39
 - Hardware requirements..... 39
 - Java 40
 - Cluster Requirements 41
 - Recommended HTTP Clients 42
 - 1.4.5. New Installation 43
 - Install Fusion on a Single Node..... 44
 - Integrate Fusion with an Existing Solr Deployment 48
 - 1.4.6. Upgrade Fusion..... 50
 - Per-version instruction sets 50
 - Upgrade overview: migrating data, configurations, and customizations..... 50
 - Version Incompatibilities 52
 - Upgrade Fusion 3.0.0 to 3.0.1..... 53
 - Upgrade Fusion 2.1.4 or 2.4.x to Fusion 3.0.0..... 55

Upgrade Fusion 2.1 or 2.2 to Fusion 2.4	61
Upgrade Fusion 1.2 to Fusion 2.4	67
1.4.7. bin/fusion	77
Fusion Agent Process	77
Start Fusion	77
Check the status of Fusion	77
Stop Fusion	77
Troubleshooting	78
1.4.8. Directories and Logs	79
Directories	79
Logfiles	80
1.4.9. Configuration Files	82
Logging configuration files	82
1.4.10. Default Ports	84
Default Ports	84
ZooKeeper Port Configuration	85
1.4.11. Checking System State	86
Full System Check	86
Solr Health Check	86
REST API Services Health Check	87
Connectors Health Check	87
1.4.12. Migrating Fusion Data	89
Migrating collections and related configurations	89
Migrating application configuration data	96
1.5. The Fusion Workflow	101
1.5.1. The Index Workbench	101
1.5.2. The Query Workbench	102
1.5.3. Developing your search application	102
1.6. Explore Objects	103
1.6.1. Object Types	103
1.6.2. Use Object Explorer	103
2. Search	105
2.1. Datasources	106
2.1.1. Index Workbench	107
Datasources	107
Parser Configuration	108
Index Pipelines	109
2.1.2. Index Pipelines	111
Collection-specific Pipelines	112
Pre-configured Pipelines	112
Fusion PipelineDocument Objects	115
Index Pipeline Stages	119
Index Profiles	121
Entity Extraction	122

Blob Storage	124
Time-Based Partitioning	126
Custom JavaScript Stages For Index Pipelines	132
2.1.3. Other Ingestion Methods	137
Importing Data with Pig	138
Importing Data with Hive	143
Pushing Documents to a Pipeline	148
2.2. Query Pipelines	150
2.2.1. Default Query Pipelines	150
2.2.2. Custom Query Pipelines	151
2.2.3. Query Workbench	152
Query pipelines and pipeline stages	152
Boosting and blocking	156
Compare mode	160
Editing parameters	161
Formatting the search results	161
Step-by-step Query Workbench workflow	168
2.2.4. Fusion Query Request Objects	169
The Request Java API	169
JSON representation of a Request object	169
2.2.5. Fusion Query Response Objects	172
2.2.6. Faceting	173
Field faceting	173
Range faceting	173
Faceting concepts	173
Further Reading	174
2.2.7. Search Query Pipeline Stages	175
Configuring query pipeline stages	175
Conditional query processing	175
Reference topics	175
Setup	175
Results relevancy	175
Fetch data	176
Troubleshooting	176
Advanced	176
Other	176
2.2.8. Query Profiles	177
Query Profiles in the UI	177
2.2.9. Using Query Pipelines with SolrJ	178
Authentication with SolrJ	178
Example	178
2.2.10. Search Query Reporting	180
Available Reports	180
2.2.11. Custom JavaScript Stages for Query Pipelines	182

JavaScript Query Stage Global Variables	182
JavaScript Query Stage Examples.	183
Debugging and Troubleshooting	183
The JavaScript Engine Used by Fusion	183
2.2.12. Solr Query Language Cheat Sheet	185
Wildcardcards	185
Common query parameters.	185
Query examples	185
2.3. Signals and Aggregations	186
2.3.1. Signals	186
2.3.2. Aggregations	186
2.3.3. The cold start problem	187
2.3.4. Signals	188
Enabling and disabling signals	188
Signal document structure.	188
The default signals index pipeline	190
Removing signals	190
Video tutorial.	190
2.3.5. Aggregations	191
Aggregation Pipelines	191
Aggregation Functions	191
Aggregation properties.	191
EventMinerAggregator	192
Aggregation job configuration	193
EventMinerAggregator configuration parameters	193
Example Configuration.	193
Creating Aggregation Jobs	195
Aggregator Functions	199
Aggregator Scripting	211
2.4. Recommendations and Boosting	214
2.4.1. Recommendation methods	214
2.4.2. Using the Recommendations API	215
2.4.3. Enable Spark-based Collaborative Recommendations	215
2.4.4. Disable Spark-based Collaborative Recommendations	217
2.4.5. Content-based Filtering.	219
2.4.6. Collaborative Filtering	220
Parts of the Process	220
Use Signals to Report Data to Fusion.	220
Aggregate Signal Data	220
Compute the Documents to Recommend	220
Recommend Documents at Query Time	221
2.4.7. Item-Item Similarities	222
2.4.8. Use Multiple Recommendation/Boosting Stages	223
Examples of Using Multiple Collaborative Recommendation Stages	223

2.4.9. Step by Step	224
Step 1: Create a Collection to Hold the Data	224
Step 2: Enable Recommendations	224
Step 3: Collect User-Item Interaction Data	226
Step 4: Configure Recommender Jobs	226
Step 5: Run a Job to Aggregate Signal Data	227
Step 6: Run a Job to Create Collaborative Recommendations	227
Step 7: Set Up Query Pipelines	227
Step 8: Look Up Precomputed Recommendations at Query Time	228
2.5. Search Applications	229
2.5.1. Autocomplete	230
2.5.2. DateTime Processing	231
DateUtils - Uniform API for Date Parsing and Formatting	231
Supported Formats	231
DateUtils Usage	232
2.5.3. Stopwords Files	234
2.5.4. Synonyms Files	235
Synonyms Manager	235
3. Analytics	237
3.1. Get Started with Fusion Dashboards	238
3.1.1. Dashboard Controls	238
3.1.2. Create the Raw Signals Collection	239
3.1.3. Create the Dashboard	241
Configure the Histogram Panel	242
Configure the Table Panel	242
3.1.4. Save the Dashboard	243
3.1.5. Conclusion	243
3.2. Log Analytics Dashboards	245
3.2.1. Search query controls	245
3.2.2. Display controls	245
3.3. Signals Dashboards	246
3.3.1. Signals Analytics Components	247
Search query controls	247
Display controls	247
3.4. Fusion Metrics Dashboard	248
3.5. Search Analytics Dashboard	249
3.6. Non-Time Series Dashboards	250
3.6.1. Search query controls	250
3.6.2. Display controls	250
3.7. Dashboard Layouts	251
3.7.1. Configuring Rows	251
3.7.2. Configuring Panels	251
3.7.3. Nested Panel Layout using Column Panels	251
3.8. Input Panels	252

3.8.1. Query Syntax	254
Rules for the Simplest Cases	254
Examples	254
3.8.2. Inspect a Panel Query	255
3.8.3. Query Panel	256
3.8.4. Timepicker Panel	257
3.9. Display Panels	258
3.9.1. What Data is Displayed	258
3.9.2. Layout Panels	258
3.9.3. Textual Information Panels	258
3.9.4. Graphical Visualization Panels	259
3.9.5. Bettermap Panel	266
3.9.6. Filtering Panel	267
3.9.7. Heatmap Panel	268
3.9.8. Histogram Panel	269
Configuration Options	269
3.9.9. Hits Panel	271
3.9.10. Map Panel	272
3.9.11. Table Panel	273
Configuration Options	273
3.9.12. Terms Panel	274
3.9.13. Text Panel	275
3.9.14. Ticker Panel	276
4. Dev Ops	277
4.1. System Metrics	278
4.1.1. Types of Metrics Collected	278
4.1.2. Metrics of Particular Interest	278
Slow Web Service Calls	278
System Memory	279
Query and Index Pipeline Stage Metrics	280
Web Service Endpoint Metrics	282
Solr Request Metrics	282
4.1.3. Changing Metric Collection Frequency	283
4.2. Messaging and Alerting	284
4.2.1. Supported Message Types	284
4.2.2. Messaging Service Configuration	284
Enabling Messaging Services	284
String Templates	285
4.2.3. Triggering Messages and Alerts	285
Messaging and Alerting Pipeline Stages	285
Setting Properties Upstream	285
4.3. System Usage Monitor	287
4.3.1. Information Collected by the Usage Monitor	287
4.3.2. How Data is Sent	287

4.3.3. How to Opt-Out	288
4.4. Schedules	289
4.4.1. Scheduler job definitions	289
Time properties	289
Call properties	290
4.4.2. How to define a scheduler job	290
Defining a scheduler job in the Fusion UI	290
API Examples	291
4.5. bin/fusion	293
4.5.1. Fusion Agent Process	293
4.5.2. Start Fusion	293
4.5.3. Check the status of Fusion	293
4.5.4. Stop Fusion	293
4.5.5. Troubleshooting	294
4.5.6. Directories and Logs	295
Directories	295
Logfiles	296
4.5.7. Default Ports	298
Default Ports	298
ZooKeeper Port Configuration	299
4.5.8. Checking System State	300
Full System Check	300
Solr Health Check	300
REST API Services Health Check	301
Connectors Health Check	301
4.5.9. Migrating Fusion Data	303
Migrating collections and related configurations	303
Migrating application configuration data	310
4.6. Blob Storage	315
4.6.1. Blob Types	315
4.6.2. The Blob manager	315
4.7. Jobs and Schedules	317
4.7.1. Job types	317
4.7.2. The Jobs manager	317
4.7.3. The Scheduler	321
4.7.4. Datasource Jobs	322
4.7.5. Spark Jobs	323
Spark job subtypes	323
Spark job configuration	324
4.7.6. Tasks	325
Task subtypes	325
Task configuration	325
4.8. Access Control	326
4.8.1. User Authentication and Authorization	326

4.8.2. User Account Administration	326
4.8.3. Video tutorial.....	326
4.8.4. Users	327
Add Users	327
Manage Users in the Fusion UI	327
Manage Users via HTTP Requests to the Users API	327
User Information	327
4.8.5. Roles.....	329
Where You Specify Roles	329
Default Roles	329
Role Information	331
Manage Roles.....	331
4.8.6. Permissions	332
Specify UI Permissions	332
Specify API Permissions	332
Example Permissions Set	333
4.8.7. Security Realms.....	334
Security Realm Types	334
Manage Security Realms	335
4.8.8. Configuring Fusion for SSO	336
4.8.9. Configuring Fusion for LDAP.....	338
Bind	340
Search	340
Kerberos	341
Basic LDAP Concepts and Terminology	344
4.8.10. Configuring Fusion for Kerberos	348
Configuring the Kerberos client	349
The Service Principal Keytab file	350
Configuring Fusion Authentication for Kerberos Realm	351
Kerberos/SPNEGO HTTP Authentication	352
Testing and Troubleshooting.....	353
References and Tutorials	354
4.8.11. Configuring Fusion for SAML	355
Fusion Configuration for a SAML Realm	355
SAML Authority Identify Provider Configuration for Fusion	356
Example SAML Realm Configuration	356
References	357
4.8.12. User Access Request Params	358
Per-Request Authentication.....	358
Session Cookies	358
4.8.13. Configuring Fusion for SSL	360
Installing an SSL certificate	360
Enabling SSL in the Fusion UI.....	360
Disabling HTTP	362

Configuring Fusion for SSL Solr/SolrCloud	363
Generating a self-signed certificate	364
References and tutorials	367
4.8.14. Web Authentication Cookie FAQs	368
5. Spark and Machine Learning	369
5.1. Further Reading	369
5.2. Spark Concepts and Terminology	370
5.2.1. Spark Fusion Terminology	370
5.2.2. Spark Fusion Directories	371
5.3. Spark Getting Started	372
5.3.1. Starting Spark Master and Worker Processes	372
5.3.2. Running a Job in the Spark Shell	373
5.3.3. The Spark Shell UI	374
5.3.4. Spark Job Tuning	376
5.4. Spark Driver Processes	379
5.4.1. Default Driver	379
5.4.2. Scripted Driver	381
5.4.3. Spark Drivers in a Multi-node Cluster	381
5.5. Spark Configuration	382
5.5.1. Spark Master / Worker Resource Allocation	382
Number of Cores Allocated	382
Memory Allocation	383
Cores per Driver Allocation	384
5.5.2. Ports used by Spark in Fusion	385
5.5.3. Directories and Temporary Files	385
Shaded jar file	385
Temporary work directories	385
5.5.4. Connection Configurations for an SSL-enabled Solr cluster	386
5.6. Scaling Spark Aggregations	388
5.7. Spark Troubleshooting	389
5.7.1. Job Hung in Waiting Status	389
5.7.2. Lost Executor Due to Heartbeat Timeout	389
5.7.3. Spark Master won't start on EC2	389
5.8. Machine Learning Models in Fusion	390
5.8.1. Training Models	390
5.8.2. Model Prediction	390
5.8.3. Model Checking	391
5.8.4. Metadata file "spark-mllib.json"	391
Reference Manual	393
6. Connectors Configuration Reference	394
6.1. List of connectors	394
6.2. Installing a connector	398
6.2.1. Installing a connector using the Fusion UI	398
6.2.2. Installing a connector using the API	399

6.3. Updating a connector	400
6.4. Deleting a connector	400
6.4.1. Deleting a connector using the Fusion UI.....	400
6.4.2. Deleting a connector using the REST API	401
6.5. Alfresco Connector and Datasource Configuration	403
6.5.1. Configuration.....	403
Connector-specific Properties	403
Security Trimming.....	404
Link Discovery	404
Limit Documents	404
Crawl Performance	406
Dedupe	408
Recrawl Rules	409
Crawl History	410
Field Mapping.....	411
ConnectorDb Configuration.....	413
General Configuration	414
6.6. Azure Connector and Datasource Configuration	415
6.6.1. Configuration.....	415
Field Mapping.....	417
ConnectorDb Configuration.....	418
General Configuration	419
6.7. Box.com Connector and Datasource Configuration.....	420
6.7.1. How the Box Connector Works	420
6.7.2. Overview of Steps.....	420
6.7.3. Set Up Box.....	421
Step 1: Sign Up for a Box Developer Account	421
Step 2: Enable 2-Step Verification	421
Step 3: Create a Box App that Fusion Can Use to Crawl the Box files	422
Step 4: Configure Your App to Use a Box Service Account	422
6.7.4. Set Up Fusion.....	423
Step 5: Install Fusion's Box Connector	423
Step 6: Create Datasources.....	423
6.7.5. Crawl a Box Data Repository	425
Step 7: Crawl the Fusion Datasources	425
6.7.6. Box Authorization, Access, and Refresh Tokens	425
Authentication Using JWT	425
6.7.7. Authentication Using OAuth 2.0	426
6.7.8. Configuration.....	427
Connector-specific Properties	427
Limit Documents	431
Security Trimming.....	432
Crawl Performance	433
Dedupe	434

Recrawl Rules	435
Crawl History	436
Field Mapping	437
ConnectorDb Configuration	439
General Configuration	440
6.8. Couchbase Connector and Datasource Configuration	441
6.8.1. Indexing and Commits	441
6.8.2. Splitting Couchbase Documents	441
6.8.3. Field Mapping with Couchbase	442
Example	442
6.8.4. Configuration	444
Field Mapping	446
ConnectorDb Configuration	448
General Configuration	449
6.9. Dropbox Connector and Datasource Configuration	450
6.9.1. Dropbox Authentication	450
6.9.2. Configuration	450
Connector-specific Properties	450
Limit Documents	451
Crawl Performance	453
Dedupe	454
Recrawl Rules	455
Crawl History	456
Field Mapping	457
ConnectorDb Configuration	459
General Configuration	460
6.10. Drupal 7.x Connector and Datasource Configuration	461
6.10.1. Configuration	461
Connector-specific Properties	461
Drupal URL Aliases	463
Limit Documents	464
Crawl Performance	465
Dedupe	467
Recrawl Rules	468
Crawl History	469
Field Mapping	470
ConnectorDb Configuration	472
General Configuration	473
6.11. FTP Connector and Datasource Configuration	474
6.11.1. Configuration	474
Field Mapping	477
ConnectorDb Configuration	479
General Configuration	480
6.12. GitHub Connector and Datasource Configuration	481

6.12.1. Configuration	481
Connector-specific Properties	481
Limit Documents	483
Crawl Performance	484
Dedupe	485
Recrawl Rules	486
Crawl History	487
Field Mapping	488
ConnectorDb Configuration	490
General Configuration	491
6.13. Google Drive Connector and Datasource Configuration	492
6.13.1. Google Drive authentication	492
Authentication for access to site-wide documents	492
Authentication for access to per-user documents	496
6.13.2. Configuration	499
Connector-specific Properties	500
Limit Documents	502
Security Trimming	503
Crawl Performance	504
Dedupe	505
Recrawl Rules	506
Crawl History	507
Field Mapping	508
ConnectorDb Configuration	511
General Configuration	512
6.14. HDFS Connector and Datasource Configuration	513
6.14.1. Configuration	513
Field Mapping	517
ConnectorDb Configuration	518
General Configuration	519
6.15. Hadoop Connector and Datasource Configuration	520
6.15.1. Hadoop Installation and Configuration	521
Permission Issues	521
Configuration for a Kerberos Hadoop Cluster	521
6.15.2. Configuration	522
Field Mapping	526
ConnectorDb Configuration	527
General Configuration	528
6.16. JDBC Connector and Datasource Configuration	529
6.16.1. SQL queries for document retrieval	529
Delta queries	529
Nested queries	529
6.16.2. Uploading a JDBC driver	529
6.16.3. Indexing binary data	532

6.16.4. Troubleshooting	532
6.16.5. Configuration	533
Field Mapping	535
ConnectorDb Configuration	537
General Configuration	538
6.17. JIRA Connector and Datasource Configuration	539
6.17.1. Configuration	539
Connector-specific Properties	539
Security Trimming	542
Limit Documents	542
Crawl Performance	544
Dedupe	545
Recrawl Rules	546
Crawl History	547
Field Mapping	548
ConnectorDb Configuration	550
General Configuration	551
6.18. Javascript Connector and Datasource Configuration	552
6.18.1. The JavaScript Program	552
6.18.2. Examples	553
Return content as a java.lang.String	553
Return content as a byte array	553
Return content as a JavaScript array	553
Return content as a JavaScript map	553
Leverage the Fetcher	553
6.18.3. Configuration	554
Connector-specific Properties	554
Link Discovery	554
Limit Documents	555
Crawl Performance	556
Dedupe	557
Recrawl Rules	558
Crawl History	559
Field Mapping	560
ConnectorDb Configuration	562
General Configuration	563
6.19. Jive Connector and Datasource Configuration	564
6.19.1. Updating the Jive connector	564
Removing the Old Jive Connector	564
Adding the New Jive Connector	564
Installing the Jive Add-on	565
6.19.2. Security Trimming of Results	565
6.19.3. Configuration	566
Field Mapping	571

ConnectorDb Configuration	573
General Configuration	574
6.20. Local Filesystem Connector and Datasource Configuration	575
6.20.1. Configuration	575
Connector-specific Properties	575
Limit Documents	575
Crawl Performance	577
Dedupe	578
Recrawl Rules	579
Crawl History	580
Field Mapping	581
ConnectorDb Configuration	583
General Configuration	584
6.21. MongoDB Datasource and Connector Configuration	585
6.21.1. Configuration	585
Field Mapping	587
ConnectorDb Configuration	589
General Configuration	590
6.22. Solr Push Endpoint Datasource and Configuration	591
6.22.1. Configuration	591
Field Mapping	592
ConnectorDb Configuration	593
General Configuration	594
6.23. S3 Connector and Datasource Configuration	595
6.23.1. Bucket Permissions	595
6.23.2. Configuration	595
Field Mapping	599
ConnectorDb Configuration	601
General Configuration	602
6.24. Salesforce Connector and Datasource Configuration	603
6.24.1. Security Trimming	603
6.24.2. Creating a Salesforce Connected App	603
6.24.3. Configuration	604
Field Mapping	605
ConnectorDb Configuration	607
General Configuration	608
6.25. ServiceNow Connector and Datasource Configuration	609
6.25.1. Configuration	609
Field Mapping	611
ConnectorDb Configuration	613
General Configuration	614
6.26. SharePoint Connector and Datasource Configuration	615
6.26.1. SharePoint Permissions	615
Required Permissions	616

Troubleshooting Permission Issues	617
6.26.2. Configuration	617
Connector-specific Properties	617
Authentication	619
Security Trimming.....	620
Limit Documents	620
Dedupe	622
Recrawl Rules	622
Crawl Performance	624
Crawl History	625
Field Mapping.....	625
ConnectorDb Configuration.....	627
General Configuration	628
6.27. Solr Connector and Datasource Configuration.....	629
6.27.1. Limitations.....	629
6.27.2. Configuration	629
Field Mapping.....	632
ConnectorDb Configuration.....	633
General Configuration	634
6.28. SolrXML Connector and Datasource Configuration.....	635
6.28.1. The SolrXML Format	635
6.28.2. Configuration	635
Field Mapping.....	637
ConnectorDb Configuration.....	639
General Configuration	640
6.29. Subversion Connector and Datasource Configuration	641
6.29.1. Configuration	641
Connector-specific Properties	641
Limit Documents	642
Crawl Performance	643
Dedupe	644
Recrawl Rules	645
Crawl History	646
Field Mapping.....	647
ConnectorDb Configuration.....	649
General Configuration	650
6.30. Twitter Search Connector and Datasource Configuration	651
6.30.1. Registering for Twitter Credentials	651
6.30.2. Configuration	651
Field Mapping.....	653
ConnectorDb Configuration.....	655
General Configuration	656
6.31. Twitter Stream Connector and Datasource Configuration.....	657
6.31.1. Registering for Twitter Credentials	657

6.31.2. Configuration	657
Field Mapping	659
ConnectorDb Configuration	661
General Configuration	662
6.32. Web Connector and Datasource Configuration	663
6.32.1. Crawling JavaScript Web sites	663
6.32.2. Limiting Crawl Scope	663
6.32.3. Extracting Content from Pages	663
6.32.4. Sitemap Processing	664
6.32.5. Website Authentication	664
Form-based Authentication	664
Complex Form-based Authentication	665
Basic and Digest Authentication	665
NTLM Authentication	666
6.32.6. Configuration	666
Connector-specific Properties	666
Crawl Authorization	679
Link Discovery	684
Limit Documents	685
Crawler ID	686
Document Parsing	687
Crawl Performance	689
Dedupe	691
Javascript Evaluation	692
Recrawl Rules	693
Crawl History	695
Field Mapping	696
ConnectorDb Configuration	698
General Configuration	699
6.33. Websphere Connector and Datasource Configuration	700
6.34. Windows Share Connector and Datasource Configuration	701
6.34.1. Access Control Lists (ACLs)	701
6.34.2. Configuration	702
Field Mapping	706
ConnectorDb Configuration	708
General Configuration	709
6.35. Zendesk Connector and Datasource Configuration	710
6.35.1. Authorization	710
6.35.2. Required Configuration Properties	710
6.35.3. Configuration Properties	710
Field Mapping	712
ConnectorDb Configuration	713
General Configuration	714
7. Parsers	715

7.1. Built-in parsing stages	715
7.1.1. HTML parser stage	715
7.1.2. XML parser stage	716
7.1.3. CSV parser stage	716
7.1.4. JSON parser stage	716
7.1.5. Text parser stage	716
7.1.6. Archive parser stage	717
7.1.7. Apache Tika parser stage	717
7.1.8. Fallback parser stage	717
7.2. Configuring parsers	717
7.2.1. Parser configuration in the Fusion UI	718
7.2.2. Parser configuration in the REST API	718
7.3. Parser index pipeline stage	719
7.4. Apache Tika Parser Stage	720
7.4.1. Configuration	720
Global configuration	720
Apache Tika parser stage configuration	720
7.5. Archive Parser Stage	722
7.5.1. Configuration	722
Global configuration	722
Archive parser stage configuration	722
7.6. CSV Parser Stage	724
7.6.1. Configuration	724
Global configuration	724
CSV parser stage configuration	724
7.7. Fallback Parser Stage	727
7.7.1. Configuration	727
Global configuration	727
Fallback parser stage configuration	727
7.8. HTML Parser Stage	729
7.8.1. Configuration	729
Global configuration	729
HTML parser stage configuration	729
7.9. JSON Parser Stage	732
7.9.1. Configuration	732
Global configuration	732
JSON parser stage configuration	732
7.10. Text Parser Stage	734
7.10.1. Configuration	734
Global configuration	734
Text parser stage configuration	734
7.11. XML Parser Stage	736
7.11.1. Configuration	736
Global configuration	736

- XML parser stage configuration 737
- 8. Pipeline Stages Reference 739
 - 8.1. Pipeline Stage Properties 739
 - 8.2. Conditional Processing 739
 - 8.3. Index Pipeline Stages 740
 - 8.3.1. Document transformation 740
 - 8.3.2. Document filtering and enrichment 740
 - 8.3.3. Field transformation 740
 - 8.3.4. Natural language processing 740
 - 8.3.5. Indexing 741
 - 8.3.6. Troubleshooting 741
 - 8.3.7. Advanced 741
 - 8.4. Apache Tika Parser Index Stage 742
 - 8.4.1. Configuration 742
 - Configuration Properties 742
 - 8.5. Call Pipeline Index Stage 745
 - 8.5.1. Configuration 745
 - Configuration Properties 745
 - 8.6. Date Parsing Index Stage 746
 - 8.6.1. Timestamp splitting options 746
 - Example: splitLocal parsing 747
 - Example: splitUTC parsing 747
 - 8.6.2. Configuration 748
 - Configuration Properties 748
 - 8.7. Detect Sentences Index Stage 750
 - 8.7.1. Sentence Detection in a NLP Pipeline 750
 - 8.7.2. Stage Setup 750
 - 8.7.3. Configuration 751
 - Configuration Properties 751
 - 8.8. Exclude Documents Index Stage 752
 - 8.8.1. Examples 752
 - 8.8.2. Configuration 753
 - Configuration Properties 753
 - 8.9. Exclusion Filter Index Stage 754
 - 8.9.1. Uploading an exclusion list 754
 - 8.9.2. Example 754
 - 8.9.3. Configuration 754
 - Configuration Properties 755
 - 8.10. Field Mapping Index Stage 756
 - 8.10.1. Field Mapping Stage Properties 756
 - Mapping Rules and Unmapped Rules 756
 - 8.10.2. Field Mapping Behavior 757
 - 8.10.3. Examples 758
 - 8.10.4. Configuration 759

Configuration Properties	759
8.11. Field Parser Index Stage	762
8.11.1. Configuration	762
Configuration Properties	762
8.12. Filter Short Fields Index Stage	764
8.12.1. Example Stage Specification	764
8.12.2. Configuration	764
Configuration Properties	764
8.13. Find and Replace Index Stage	766
8.13.1. Configuration	766
Configuration Properties	766
8.14. Format Signals Index Stage	768
8.14.1. Date/time parsing and formatting	768
8.14.2. Example Stage Specification	768
8.14.3. Configuration	768
Configuration Properties	768
8.15. Gazetteer Lookup Extraction Index Stage	770
8.15.1. Gazetteers and OpenNLP Tools	770
8.15.2. Uploading Lookup Lists to Fusion Blob Store	770
8.15.3. Name Lookup Example	770
8.15.4. Configuration	771
Configuration Properties	771
8.16. Include Documents Index Stage	772
8.16.1. Examples	772
8.16.2. Configuration	773
Configuration Properties	773
8.17. JavaScript Index Stage	775
8.17.1. Examples	775
8.17.2. Configuration	777
Configuration Properties	777
8.18. JDBC Index Stage	778
8.18.1. Example	778
8.18.2. Configuration	778
Configuration Properties	778
8.19. Logging Index Stage	781
8.19.1. Configuration	781
Configuration Properties	781
8.20. Machine Learning Index Stage	782
8.20.1. Configuration	782
Configuration Properties	782
8.21. OpenNLP NER Extraction Index Stage	784
8.21.1. Example Specification	784
8.21.2. Configuration	785
Configuration Properties	785

8.22. Regex Field Extraction Index Stage	787
8.22.1. Example Stage Specification	787
8.22.2. Configuration	787
Configuration Properties	787
8.23. Regex Field Filter Index Stage	790
8.23.1. Example Stage Specification	790
8.23.2. Configuration	790
Configuration Properties	790
8.24. Regex Field Replacement Index Stage	792
8.24.1. Configuration	792
Configuration Properties	792
8.25. Detect Sentences Index Stage	795
8.25.1. Sentence Detection in a NLP Pipeline	795
8.25.2. Stage Setup	795
8.25.3. Configuration	796
Configuration Properties	796
8.26. Resolve Multivalued Fields Index Stage	797
8.26.1. Configuration	797
Configuration Properties	797
8.27. REST Query Index Stage	799
8.27.1. Call Parameters	799
8.27.2. Mapping Rules	800
8.27.3. Examples	800
8.27.4. Configuration	801
Configuration Properties	802
8.28. Send PagerDuty Message Index Stage	804
8.28.1. Enabling PagerDuty Messaging	804
8.28.2. Configuration	804
Configuration Properties	804
8.29. Send Slack Message Index Stage	808
8.29.1. Enabling Slack Messaging	808
8.29.2. Configuration	808
Configuration Properties	808
8.30. Send SMTP Email Index Stage	810
8.30.1. Enabling Email Messaging	810
8.30.2. Configuration	810
Configuration Properties	810
8.31. Set Property Index Stage	813
8.31.1. Configuration	813
Configuration Properties	813
8.32. Solr Dynamic Field Name Mapping	815
8.32.1. Configuration	815
Configuration Properties	815
8.33. Solr Indexer Stage	817

8.33.1. Configuration	817
Configuration Properties	817
8.34. Solr Partial Update Indexer Stage	819
8.34.1. Example Stage Specification	819
8.34.2. Types of Update Operations.....	819
8.34.3. Document Identifier Field	820
8.34.4. Optimistic Concurrency	820
8.34.5. Performance Considerations.....	820
8.34.6. Solr Date Formats.....	820
8.34.7. Configuration	821
Configuration Properties	821
8.35. Tag Part-of-Speech Index Stage.....	826
8.35.1. Part-of-speech Tagging in a NLP Pipeline.....	826
8.35.2. Stage Setup.....	826
8.35.3. Configuration	827
Configuration Properties	827
8.36. Update Experiment Stage.....	828
8.36.1. Configuration	828
Configuration Properties	828
8.37. Write Log Message Index Stage	829
8.37.1. Configuration	829
Configuration Properties	829
8.38. XML Transformation Index Stage	831
8.38.1. Pipeline Configuration	831
8.38.2. XML Transforms.....	831
8.38.3. Example Stage Specification	832
8.38.4. Configuration	834
Configuration Properties	834
9. Query Pipeline Stages	836
9.1. Setup	836
9.2. Results relevancy	836
9.3. Fetch data.....	836
9.4. Troubleshooting	836
9.5. Advanced	837
9.6. Other	837
9.7. Active Directory Security Trimming Stage.....	838
9.7.1. Example Stage Setup	838
9.7.2. Configuration.....	838
Configuration Properties	838
9.8. Additional Query Parameters Stage.....	841
9.8.1. Example Stage Specification	841
9.8.2. Configuration.....	841
Configuration Properties	841
9.9. Analytics Catalog Stage	843

9.9.1. Configuration	843
Configuration Properties	843
9.10. Block Documents Stage	845
9.10.1. Configuration	845
Configuration Properties	845
9.11. Boost Documents Stage	847
9.11.1. Configuration	847
Configuration Properties	847
9.12. Boost with Signals Stage	849
9.12.1. Configuration	849
Configuration Properties	849
9.13. Run Query Pipeline Stage	852
9.13.1. Configuration	852
Configuration Properties	852
9.14. Experiment Query Parameters Stage	853
9.14.1. Configuration	853
Configuration Properties	853
9.15. Facets Stage	854
9.15.1. Configuration	854
Configuration Properties	854
9.16. JavaScript Stage	857
9.16.1. Configuration	857
Configuration Properties	857
9.17. JDBC Lookup Stage	858
9.17.1. Example	858
9.17.2. Configuration	858
Configuration Properties	859
9.18. Landing Pages Stage	861
9.18.1. Configuration	861
Configuration Properties	861
9.19. Logging Stage	863
9.19.1. Configuration	863
Configuration Properties	863
9.20. Machine Learning Stage	864
9.20.1. Configuration	864
Configuration Properties	864
9.21. Parameterized Boosting Stage	866
9.21.1. Configuration	866
Configuration Properties	866
9.22. Parameterized Faceting Stage	868
9.22.1. Configuration	868
Configuration Properties	868
9.23. Query Fields Stage	873
9.23.1. Configuration	873

Configuration Properties	873
9.24. Recommend Items for Item Stage	875
9.24.1. Prerequisites	875
9.24.2. Configuration	875
Configuration Properties	875
9.25. Recommend Items for User Stage	879
9.25.1. Prerequisites	879
9.25.2. Configuration	879
Configuration Properties	879
9.26. Recommend More Like This Stage	883
9.26.1. Configuration	883
Configuration Properties	883
9.27. Query RPC Stage	886
9.27.1. Configuration	886
Configuration Properties	886
9.28. Retrieve Stored Parameters Stage	890
9.28.1. Configuration	890
Configuration Properties	890
9.29. Return Query Parameters Stage	891
9.29.1. Configuration	891
Configuration Properties	891
9.30. Rollup Aggregation Stage	892
9.30.1. Configuration	892
Configuration Properties	892
9.31. Security Trimming Stage	894
9.31.1. Configuration	894
Configuration Properties	894
9.32. Send PagerDuty Message Stage	896
9.32.1. Enabling PagerDuty Messaging	896
9.32.2. Configuration	896
Configuration Properties	896
9.33. Send SMTP Email Stage	900
9.33.1. Enabling Email Messaging	900
9.33.2. Configuration	900
Configuration Properties	900
9.34. Send Slack Message Stage	903
9.34.1. Enabling Slack Messaging	903
9.34.2. Configuration	903
Configuration Properties	903
9.35. Solr Query Stage	905
9.35.1. Configuration	905
Configuration Properties	905
9.36. Write Log Message Stage	906
9.36.1. Configuration	906

- Configuration Properties 906
- 10. REST API Reference 908
 - 10.1. Listing all Fusion component services 908
 - 10.2. Fusion Components REST API Reference Pages 908
 - 10.3. Authentication and Authorization APIs 910
 - 10.3.1. Realms API 911
 - Create, Update, Delete or List Realms 911
 - 10.3.2. Roles API 913
 - Create, Update or Delete Roles 913
 - Role Specification 913
 - Examples 914
 - 10.3.3. Sessions API 915
 - Create a Session 915
 - 10.3.4. User API 917
 - Create, Update, Delete or List Users 917
 - 10.4. Blob Store API 920
 - 10.4.1. Blob Types 920
 - 10.4.2. Examples 920
 - 10.5. Catalog API 923
 - 10.5.1. Intra-shard splits 923
 - 10.5.2. Body attributes 923
 - 10.5.3. Configuration options 924
 - 10.5.4. Examples 928
 - 10.6. Collections API 932
 - 10.6.1. Examples 932
 - 10.7. Collection Features API 935
 - 10.7.1. Examples 935
 - 10.8. Connector APIs 937
 - 10.8.1. Connector Datasources API 938
 - List a Datasource or All Datasources 938
 - Create or Update a Datasource 939
 - Delete a Datasource or All Datasources 941
 - 10.8.2. Connector History API 943
 - Get or Delete the History for a Datasource 943
 - 10.8.3. Connector JDBC API 946
 - Upload a Driver 946
 - List Available Drivers 947
 - Delete a Driver 947
 - 10.8.4. Connector Jobs API 949
 - Start, Stop or Check Status of a Job 949
 - Get Detailed Job Statistics 950
 - 10.8.5. Connector Plugins API 952
 - List Connectors 952
 - List Connector Types 952

Drop ConnectorDB State	954
10.8.6. Connector Status API	956
Reload, Quit or Get Status of Connector Registry	956
10.8.7. Connectors Crawl Database API	957
Get Statistics for a Datasource Database or Drop Database	957
Get Table Statistics or Drop the Table	958
Get or Delete Table Items	959
10.9. Experiments API (experimental)	961
10.9.1. Experiment life cycle	961
10.9.2. Experiment types	962
A/B testing and multi-armed bandits	962
IR quality metrics based on signals	964
Arbitrary aggregations	965
Machine learning pipelines	965
10.9.3. Integration with query pipelines and index pipelines	965
Experiment Query Stage	965
Experiment Update Stage	965
10.10. Groups API	966
10.11. Index Pipelines API	967
10.11.1. Examples	967
10.12. Index Profiles API	979
Examples	979
10.13. Index Stages API	981
10.13.1. Examples	981
10.14. Jobs API	986
10.14.1. Job runtime configuration	986
Job types	986
Job triggers	987
10.14.2. Examples	988
Get all datasource jobs	988
Start a datasource job now	988
Get the job history for a datasource	989
Define a trigger for a datasource job	989
10.15. Links API	990
10.15.1. Link types and link reversibility	990
10.15.2. Examples	992
10.16. Messaging API	994
10.16.1. Attributes	994
10.16.2. Examples	994
10.17. Objects API	997
10.17.1. Object export and import	997
Supported objects	997
Exporting and importing collections	997
Encrypted passwords	998

Import policies	998
Filtering on export	998
Exporting linked objects	999
10.17.2. Validation	999
10.17.3. Status messages	1000
10.17.4. Examples	1000
10.18. Parsers API	1002
10.19. Query Pipelines API	1003
10.19.1. Query Pipeline Definition Properties	1003
10.19.2. Examples	1003
10.20. Query Profiles API	1007
10.20.1. Examples	1007
10.21. Query Stages API	1009
10.21.1. Examples	1009
10.22. Recommendations API	1012
10.22.1. Recommendation types	1012
10.22.2. Output	1012
10.22.3. Examples	1013
itemsForQuery	1013
queriesForItem	1014
10.23. Reporting API	1015
10.23.1. Report Configuration Information	1015
10.23.2. Examples	1015
10.24. Scheduler API	1022
10.24.1. Schedule Definition Properties	1022
10.24.2. Examples	1023
10.25. Search Cluster API	1028
10.25.1. Search Cluster Definition Properties	1028
10.25.2. Examples	1028
10.26. Signals API	1031
10.26.1. Signal document structure	1031
10.26.2. Examples	1033
10.27. Signals Aggregator API	1035
10.27.1. Signals Aggregator Definitions Properties	1035
10.27.2. Examples	1037
10.28. Solr API	1041
10.28.1. Examples	1041
10.29. Solr Configuration API	1043
10.29.1. Solr Configuration Definition Properties	1043
10.29.2. ZooKeeper Collection Configuration Definition Properties	1044
10.29.3. Examples	1044
10.30. SolrAdmin API	1046
10.30.1. Example	1046
10.31. Spark Jobs API	1047

10.31.1. Spark job subtypes	1047
10.31.2. Spark Configuration Properties	1048
10.31.3. Fusion Configuration Properties	1049
10.32. Stopwords API	1051
10.32.1. Examples	1051
10.33. Synonyms API	1053
10.33.1. Examples	1053
10.34. System Admin APIs	1055
10.34.1. Configurations API	1056
Examples	1056
10.34.2. History API	1057
Examples	1057
10.34.3. Nodes API	1059
Examples	1059
10.34.4. System API	1061
Examples	1061
10.34.5. Usage API	1065
Examples	1065
10.35. Tasks API	1066
10.36. Taxonomy API	1067
10.36.1. Taxonomy structure	1067
10.36.2. JSON example of a taxonomy	1067
10.37. ZooKeeper Import/Export API	1069
10.37.1. ZooKeeper	1069
10.37.2. Utility script zkImportExport.sh	1069
Examples	1070
10.37.3. Fusion REST API service ZKImportExport	1071

User Manual

Chapter 1. Getting Started

These topics explain how to get started with Fusion:


- Fusion concepts - Learn how Fusion works.
- Fusion installation - Plan, deploy, or upgrade your Fusion system.
- Fusion scripts - Start and stop Fusion.
- Fusion workflow - End to end workflow for ingesting data and developing search applications.

1.1. Quick start

For a quick Fusion installation on your Unix host:

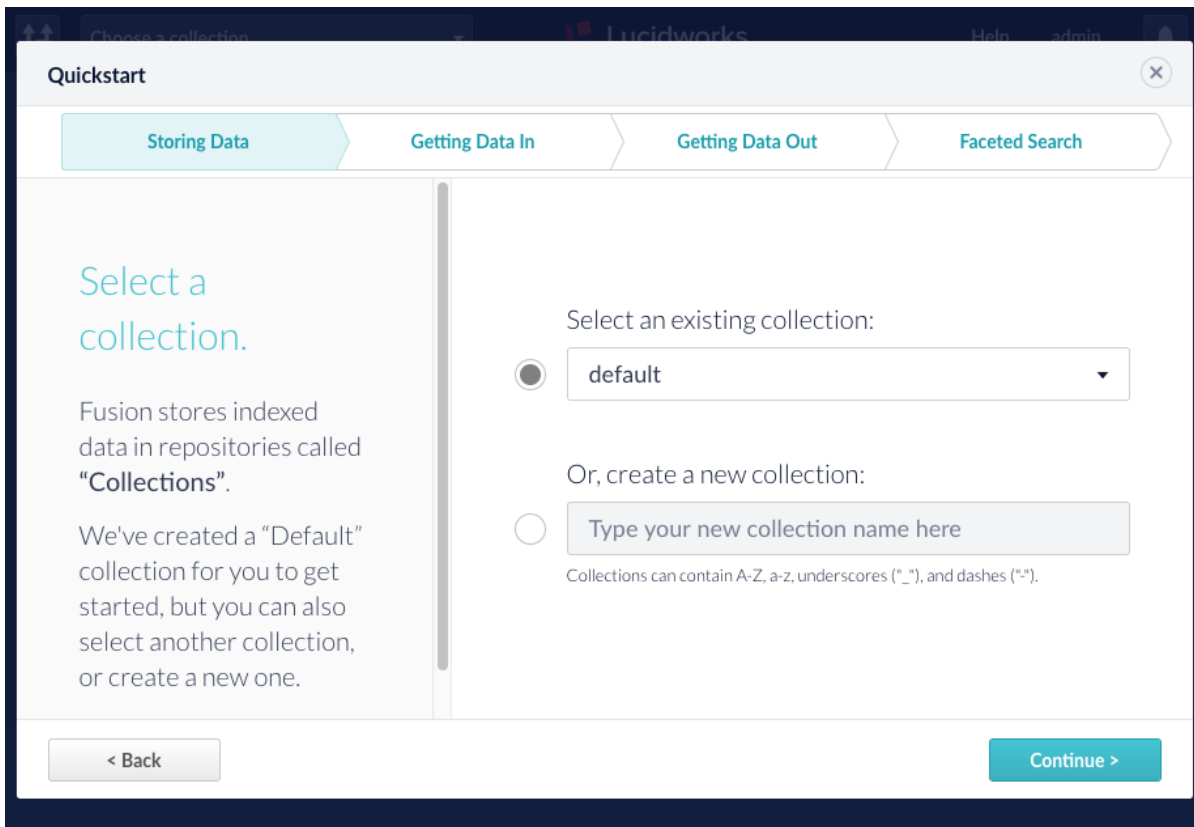
1. [Download Fusion](#).
2. `tar -xf fusion-3.1.x.tar.gz`
3. `fusion/3.1.x/bin/fusion start`
4. Go to <http://localhost:8764/>.

The first time you log in, Fusion automatically launches the Quickstart interface:



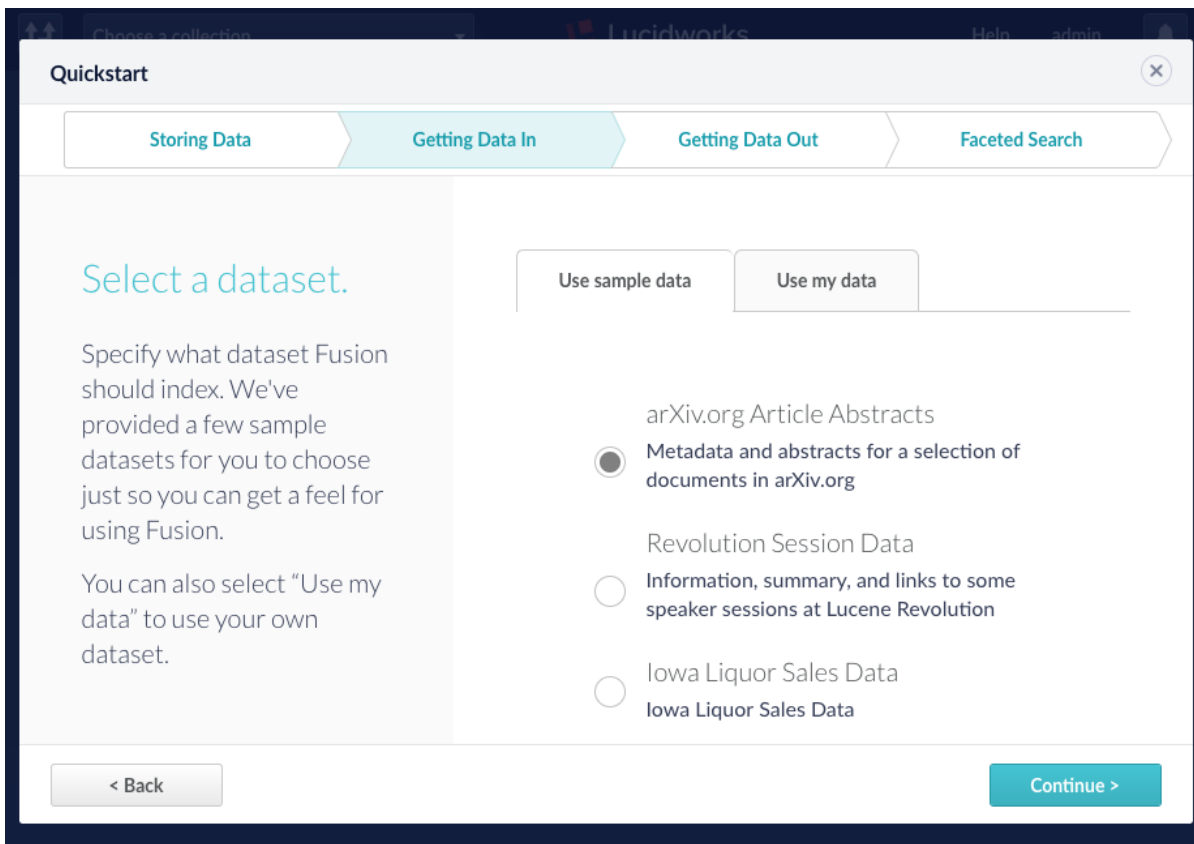
5. Click **Continue**.

On the **Storing Data** screen, you can select or create a collection for your quickstart data. A collection called "default" is created automatically for you to use:




6. Click **Continue**.

On the **Getting Data In** screen, you can either select one of the built-in sample datasets or click **Use my data** to upload your own:



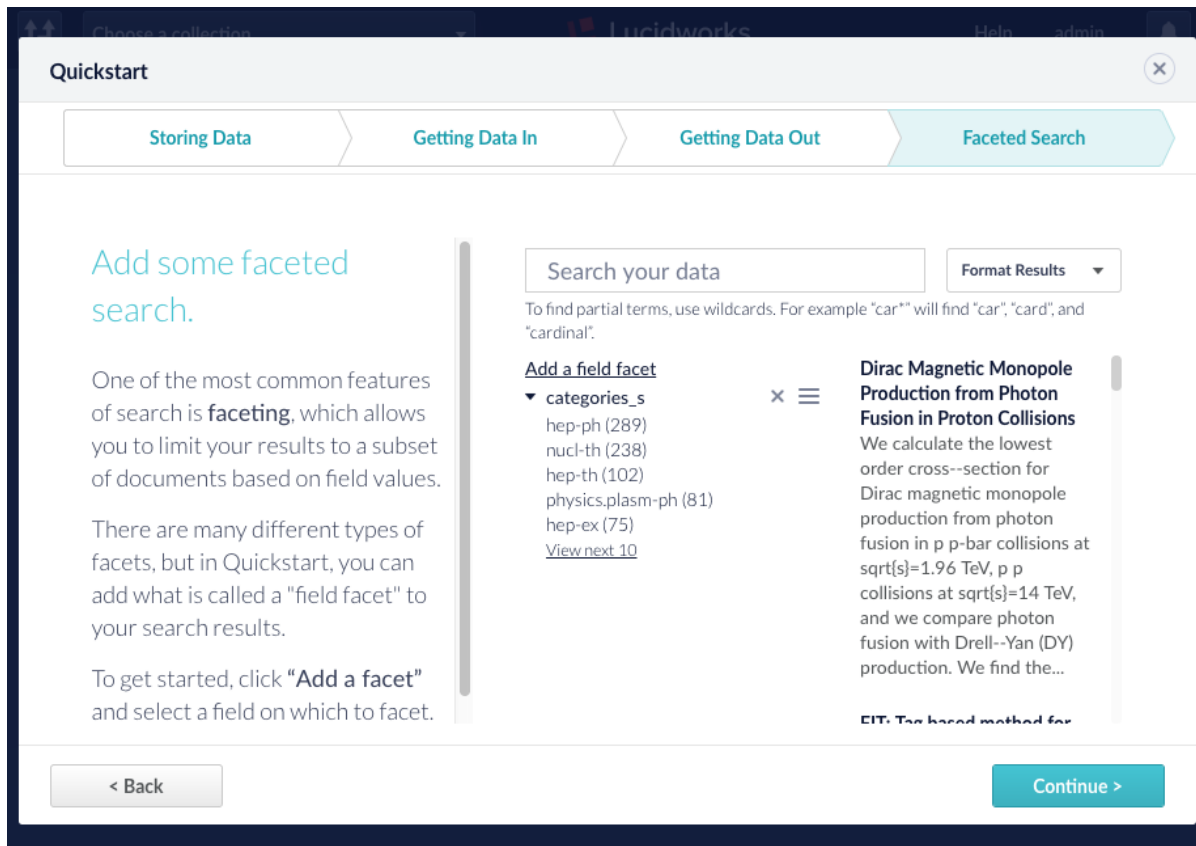
7. Click **Continue**.

On the **Getting Data Out** screen, you can see all search results and enter your own search queries to test the indexed dataset. You can also select the display fields or view the parsed documents:



8. Click **Continue**.

On the **Faceted Search** screen, you can select field to use as search facets, then experiment with the search results:




9. Click **Continue**.

From here, you can open the Index Workbench to modify the index pipeline, or open the Query Workbench to modify the query pipeline. The workbenches are essential tools in the Fusion workflow.

At any time, you can open the Quickstart interface by navigating to the Launcher (described below) and clicking **Try the Quickstart**. It provides a simple way to quickly ingest a dataset before refining the index or query pipelines.

1.2. Workflow contexts

Whenever you log in to Fusion (and it's not the first login), the Launcher lets you choose a context for your workflow:



- The **Search** context provides all the tools you need to get your data into Fusion and develop a search-ready back end for your application.
- The **Analytics** context gives you access to Fusion's dashboard application, where you can customize your view into system metrics and search activity.
- Use the **Devops** context to administrate and monitor the Fusion system.

1.3. Fusion Concepts

Fusion is a search application development platform whose core is the Solr open-source search engine. These core concepts can help you get started:

- A Collection is a Solr dataset that is managed by Fusion. It may include multiple datasources, but all of its data is queried as a single set.
- Fusion's components work together on a single host or in a distributed environment.
- A flexible range of deployment types provide scalable, high-availability search.

1.3.1. Collections

Fusion collections are Solr collections managed by Fusion. A Solr collection is a distributed index defined by a named configuration stored in ZooKeeper, with these properties:

- Number of shards
Documents are distributed across this number of partitions.
- Document routing strategy
How documents are assigned to shards.
- Replication factor
How many copies of each document in the collection.
- Replica placement strategy
Where to place replicas in the cluster.

When you first install Fusion, a collection called "default" is created automatically. You can view the simplest collection configuration by using the Collections API endpoint at <http://localhost:8765/api/v1/collections/default/>, if you haven't modified the default collection yet.

Solr is the underlying engine which indexes, stores, and searches your data. Fusion manages Solr collections, manipulates data and queries before passing them to Solr, and provides analytics and monitoring features.

If your data is already stored in a Solr instance or cluster, you can manage this collection in Fusion by creating a Fusion collection which is configured to import the existing Solr collection. See Installation with an existing Solr instance or cluster.

Primary and Auxiliary Collections

In Fusion, the "Primary" collection is the collection which contains your application data, that is, the set of documents over which search and indexing happens. Fusion registers the collection name and information about the Solr cluster that manages this collection.

Note	All collection names should be considered to be case-insensitive, even though Fusion preserves case in referring to these collections.
------	--

If your application uses Fusion's signals, analytics, or monitoring services, then Fusion will create a set of auxiliary collections in which to store signals, query, and other logfiles. Naming conventions relate auxiliary collections with the primary collection. Auxiliary collections have the same base name as the name of the primary collection plus a suffix which indicates the kind of auxiliary collection, e.g., the suffix for a query logs auxiliary collection is "_logs" so that for a primary collection named "COLL", Fusion creates an auxiliary collection named "COLL_logs". These auxiliary collections include:

- A search query logs collection, suffix "_logs".
- A pair of associated collections for signals and aggregated signals, suffixes "_signals", "_signals_aggr" respectively.

Note	Do not create primary collections with names that end in suffix "_logs", "_signals", or "_signals_aggr". Such names can only be used for Fusion auxiliary collections, which are created and managed by Fusion directly.
------	--

Fusion maintains a set of Solr collections which store Fusion's own logfiles and other internal information. These are called System Collections, described below.

Note	Do not create primary collections named "logs", or which begin with "system_". These names are reserved for Fusion system collections.
------	--

Fusion uses ZooKeeper to register information about all collections, and the Fusion components and services related to a collection. The Fusion components associated with a collection include:

- Datasources
- Pipelines
- Profiles
- Signals and aggregations
- Analytics Dashboards

System Collections

Fusion uses *system collections* for internal purposes:

- **logs** indexes the log messages from the Fusion API services.
- **audit_logs** indexes all HTTP requests sent to the Fusion API services.
- **system_banana** stores configurations used by Fusion dashboards.
- **system_blobs** stores blobs in Solr. This is used to store model files for the NLP components and other binary files used by Fusion components.
- **system_messages** is used by Fusion's Messaging Services.
- **system_metrics** stores information about the running process itself, such as the amount of memory in the system, the average response time for services, Solr heap size, etc. The data is polled at regular intervals according to the internal configuration variable: `com.lucidworks.apollo.metrics.poll.seconds`. This collection doesn't appear until after the first set of metrics are collected.

Collection Configuration Properties

Collections have three configurable properties which are set to default values in the Fusion UI. They can be configured as appropriate for your application by creating the collection using the Fusion API service Collections API.

Property	Description
signals	Property <code>signals</code> determines whether or not to create an auxiliary collections "_signals" and "_signals_aggr". When creating a collection in the Fusion UI, this property defaults to true . When creating a collection using Fusion's API services, this property defaults to false .
searchLogs	Property <code>searchLogs</code> determines whether or not to create an auxiliary search query logs collection with suffix "_logs". When creating a collection in the Fusion UI, this property defaults to true . When creating a collection using Fusion's API services, this property defaults to false .
dynamicSchema	Property <code>dynamicSchema</code> always defaults to false . When <code>dynamicSchema</code> is true , Fusion and Solr use schemaless mode to administer search and indexing over that collection.

Signals are events with timestamps that can be used to improve search results. For more information about signals in Fusion, see the section [Signals](#).

Search logs data is used for Search Query Reporting. The set of reports available includes most popular documents, queries that generated less than a minimum number of results, and search histograms.

The name schemaless mode is misleading: Solr always uses a schema when managing a collection. In schemaless mode, if a document contains a field not currently in the Solr schema, Solr processes the field value to determine what the field type should be defined as, and then adds a new field to the schema with the field name and field type. This behavior may be convenient during preliminary application development, but is rarely appropriate in a production environment, therefore the default is false.

Collection Profiles

Profiles are used to create pipeline aliases for a specific collection. In Fusion, index and query pipelines are not connected to a specific collection by default so that pipeline can be created once and re-used in several collections. This complicates the way that pipelines are used with collections. Profiles provide a shortcut.

- Index Profiles work with index pipelines for getting content into the system.
- Query Profiles work with query pipelines for user queries.

1.3.2. Fusion Components

The Fusion platform is comprised of a set of Java programs, each of which runs in its own JVM. Apache ZooKeeper provides the shared, synchronized data store for all user and application configuration information.

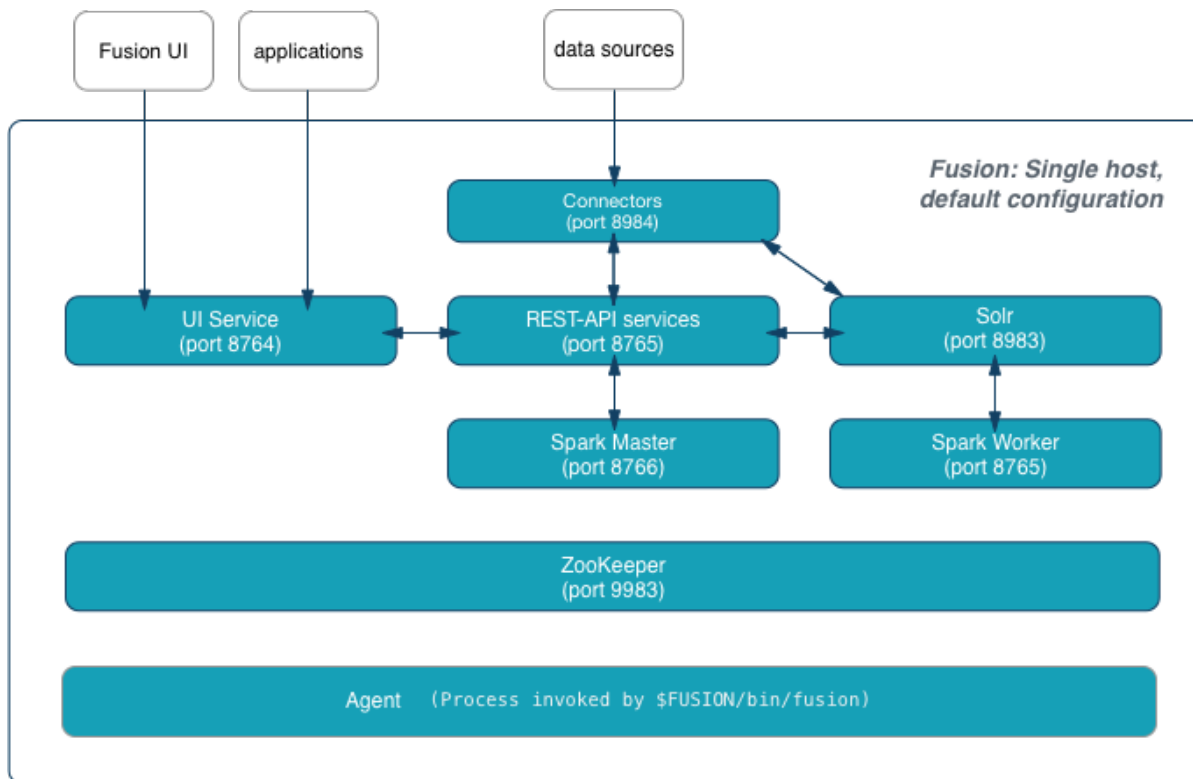
You can adjust the set of Fusion components running on each node to meet processing requirements.

Every Fusion node in a deployment runs the Fusion API Services process. Beyond that, the set of processes running on a particular Fusion node depends on the processing and throughput needs of the search application.

- Running Solr on all Fusion nodes scales out document storage as well as providing data replication. (Alternatively, you can use an external SolrCloud cluster to store Fusion collections, see Integrating Fusion with an Existing Solr Deployment.)
- Running Fusion Connectors on multiple nodes provides high throughput for indexing and updates, e.g., for applications that run analytics over live data streams such as logfile indexing or mobile tracking devices.
- Running the Fusion UI on two or more nodes provides failover for Fusion’s authentication proxy.
- Running Apache Spark on multiple nodes provides processing power for applications that aggregate clicks and other signals or use Fusion machine learning components.

Fusion processes

This diagram shows the full set of Fusion processes that run on a single node and the default ports used by each, with arrows representing the flow of HTTP requests between components for document search and indexing:



The inputs to this diagram represent:

- Users working directly in the Fusion UI, whether for developing and refining search applications, viewing analytics dashboards, or performing system administration tasks. Fusion’s UI component relays all requests to the API Services component.

- Search queries, which originate from the search application, are sent to the Fusion UI for authentication. The Fusion UI sends the requests to the Fusion API Services component, which invokes a query pipeline to build out the raw query and send the resulting query to Solr.
- Fusion datasources ingest data that will be indexed into a Solr collection. A datasource sends this raw data to Fusion's connector services. A connector invokes an index pipeline to extract, transform, and otherwise enrich the raw data, and then sends the resulting document to Solr for indexing.

Apache Spark carries out signal processing and aggregations. The Apache Spark master distributes tasks across one or more worker processes.

Apache ZooKeeper is included in this diagram because all Fusion processes across all nodes in a Fusion deployment communicate with the ZooKeeper cluster (also called an ensemble) at the socket layer via [ZooKeeper's Java API](#).

The Fusion Agent process is the server process that starts, stops, and monitors all Fusion components running on the node.

See these topics for details about each component:

- Fusion UI
- UI Service
- Connectors
- REST API Services
- Solr
- Spark
- ZooKeeper
- bin/fusion

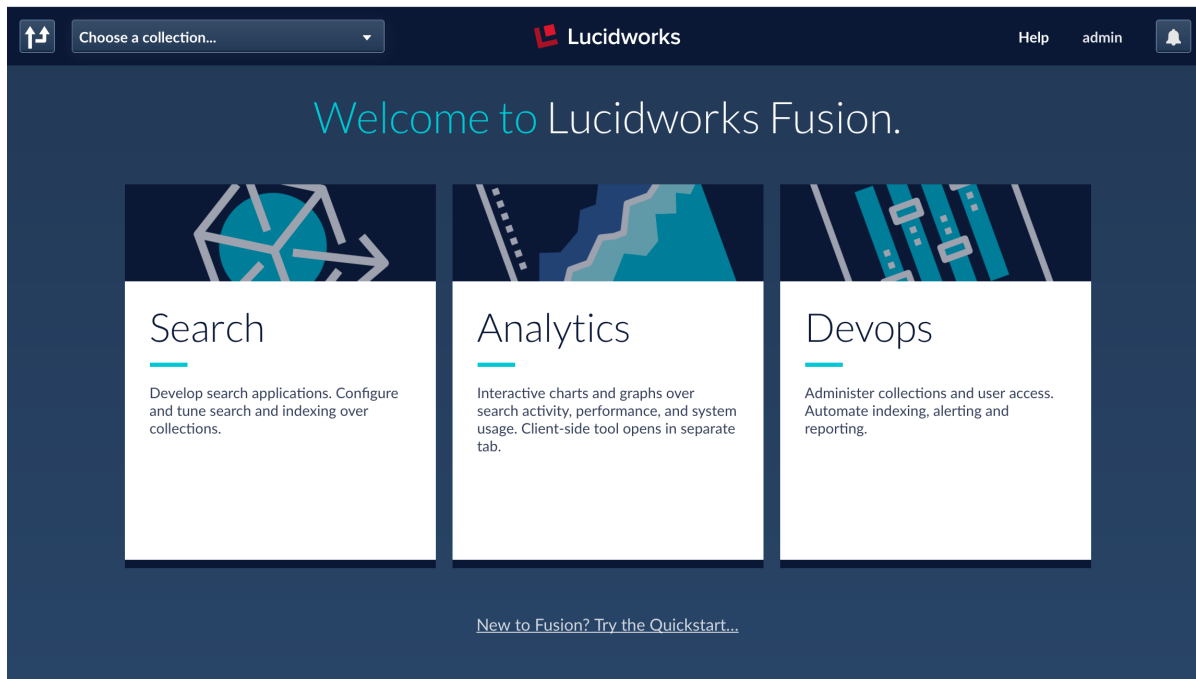
The Fusion UI

The Fusion UI provides a graphical interface to Fusion's REST APIs. To access it, point your browser to <http://localhost:8764/>, or the host on which the UI Service is running.

The first time you log in to the UI, you'll be prompted to create an admin account.

Note	Save your admin credentials. There is no password recovery system in Fusion.
------	--

Whenever you log in to Fusion (and it's not the first login), the Launcher lets you choose a context for your workflow:



- The **Search** context provides all the tools you need to get your data into Fusion and develop a search-ready back end for your application.
- The **Analytics** context gives you access to Fusion's dashboard application, where you can customize your view into system metrics and search activity.
- Use the **Devops** context to administrate and monitor the Fusion system.

Since the Fusion UI is just an interface to the REST API service, you can also use the APIs to do almost anything that the UI can do.

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

UI service

The UI service is the interface between your applications and the REST API services. All HTTP requests from external clients must pass through the UI service, because it hosts the auth proxy. It also hosts the Web-based Fusion UI.

The auth proxy

The auth proxy generates the session cookie required by all REST API services. The REST API services accept no commands without this cookie. The auth proxy requires valid username/password credentials with every request, like this:

```
curl -u <user>:<pwd> localhost:8764/api/apollo/<servicename>
```

UI configuration

The default port is 8764, configurable as `ui.port` in `fusion/3.1.x/conf/fusion.properties`.

UI Web service configuration is in `fusion/3.1.x/apps/jetty/ui`.

[This blog post](#) explains how to secure the UI using SSL.

UI logs

UI service log files are in `fusion/3.1.x/var/log/ui`.

Connectors

Connectors are the built-in mechanism for pulling your data into Fusion. Fusion comes with a wide variety of connectors, each specialized for a particular data type. When you add a datasource to a collection, you specify the connector to use for ingesting data. See Connector Types for a complete list of connectors, with links to configuration reference information for each one.

Fusion comes with a standard set of built-in connectors:

- Local Filesystem connector
- File Upload connector
- JDBC connector
- Web connector

Built-in connectors are in [fusion/3.1.x/apps/connectors/bootstrap-plugins/](#).

Additional connectors are available for download at <http://lucidworks.com/connectors/>. You can look in [fusion/3.1.x/apps/connectors/plugins/](#) to see which additional connectors are currently installed.

Installing a connector

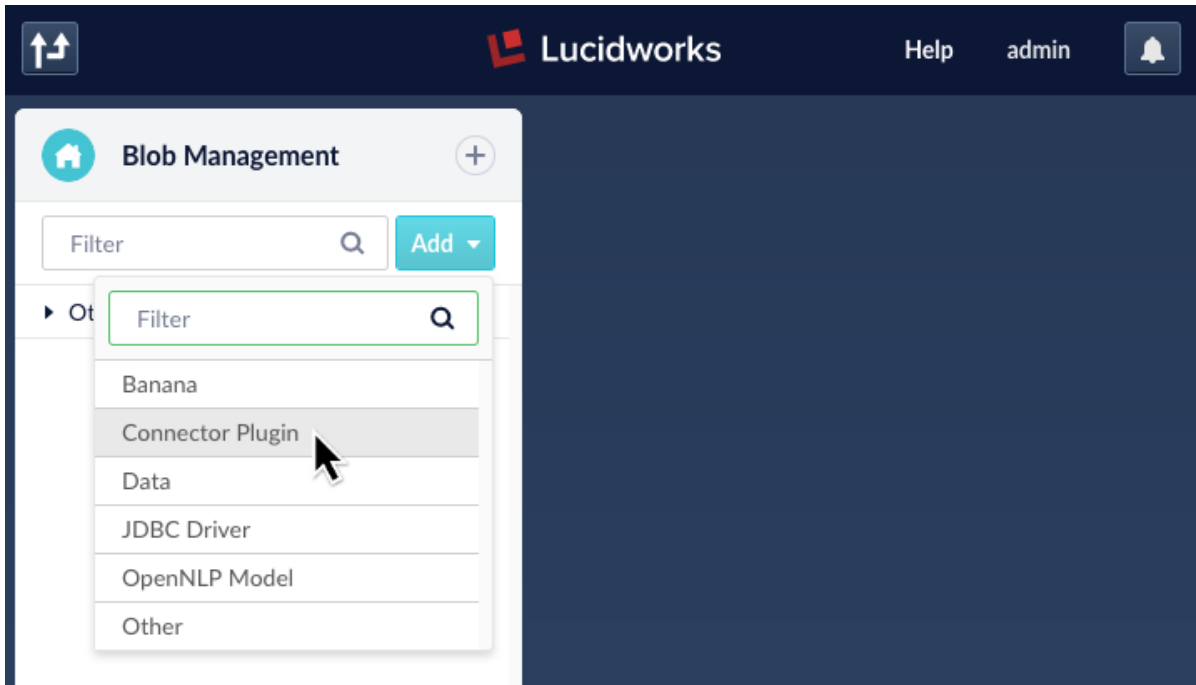
Connectors are installed by uploading them to the blob store. You can do this using the Fusion UI or the Blob Store API. Both methods are explained below.

Installing a connector using the Fusion UI

1. Download the connector zip file from <http://lucidworks.com/connectors/>.

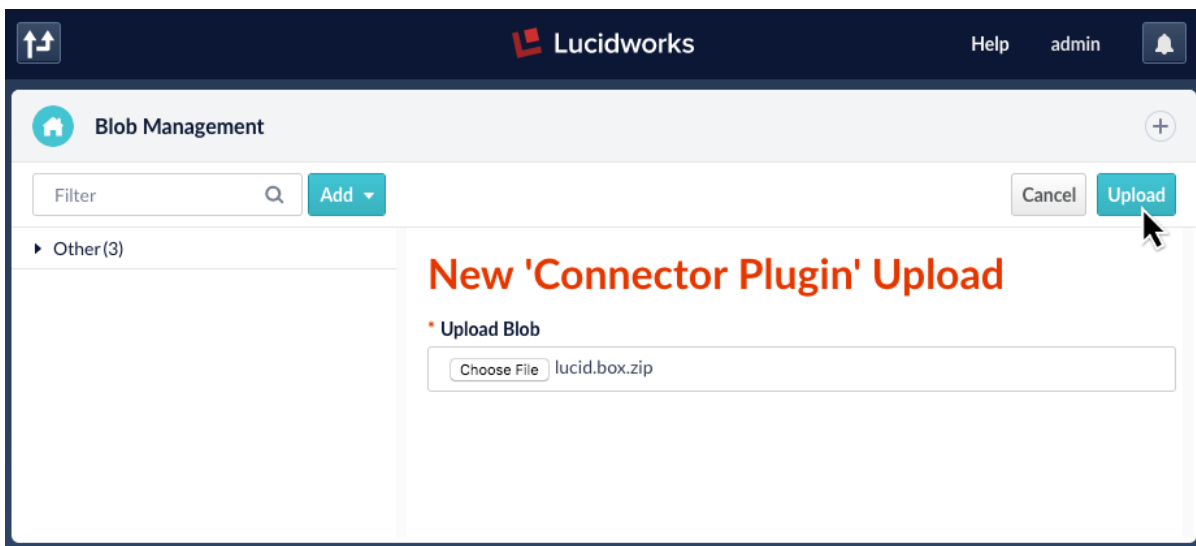
Do not expand the archive; Fusion consumes it as-is.

2. In the Fusion UI, navigate to **DevOps > Blobs**.
3. Click **Add**.
4. Select **Connector Plugin**.



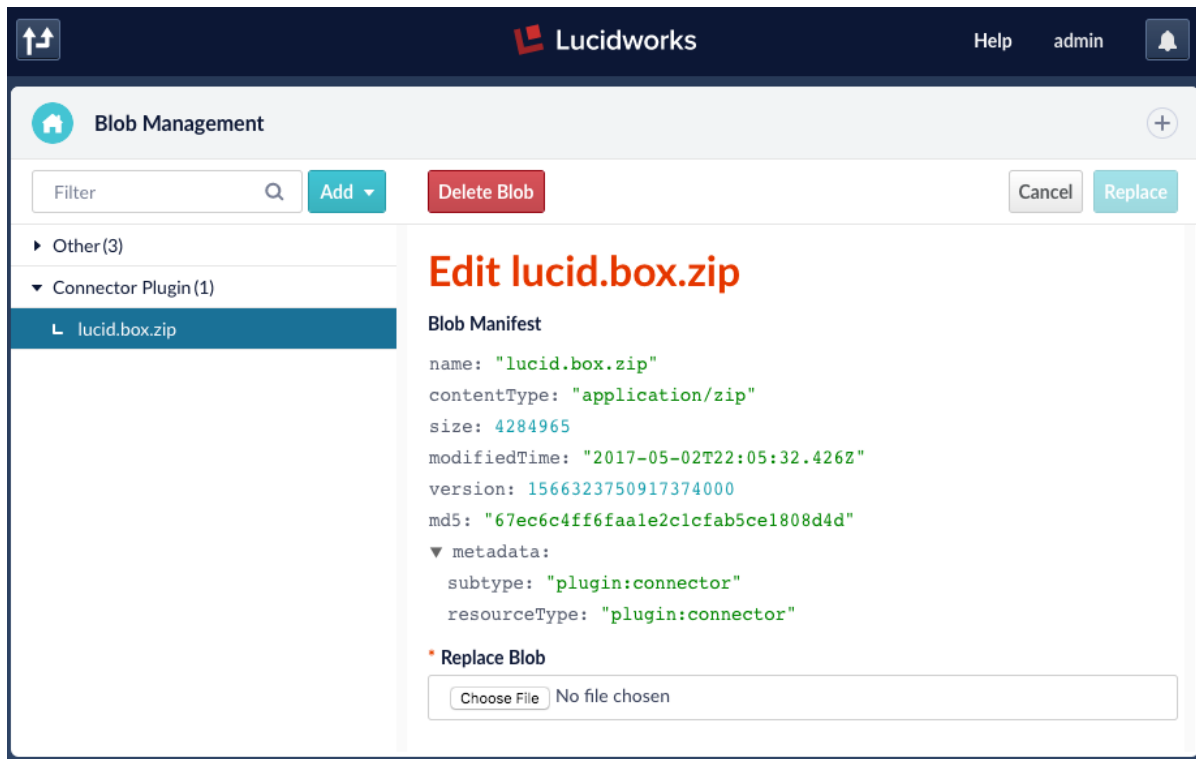
The "New 'Connector Plugin' Upload" panel appears.

5. Click **Choose File** and select the downloaded zip file from your file system.



6. Click **Upload**.

The new connector's blob manifest appears.



From this screen you can also delete or replace the connector.

Installing a connector using the API

1. Download the connector zip file from <http://lucidworks.com/connectors/>.

Do not expand the archive; Fusion consumes it as-is.

2. Upload the connector zip file to Fusion's blob store.

Specify an arbitrary blob ID, and a `resourceType` value of `plugin:connector`, as in this example:

```
curl -H 'content-type:application/zip' -X PUT
'localhost:8765/api/v1/blobs/myplugin?resourceType=plugin:connector' --data-binary @myplugin.zip
```

Fusion automatically publishes the event to the cluster, and the listeners perform the connector installation process on each node.

Tip	If the blob ID is identical to an existing one, the old connector will be uninstalled and the new connector will be installed in its place. To get the list of existing blob IDs, run: <code>curl -u user:pass localhost:8764/api/apollo/blobs</code>
-----	---

3. Look in `fusion/3.1.x/apps/connectors/plugins/` to verify that the new connector is installed.

Updating a connector

On Unix, you can update a connector by simply uploading the new one. Fusion overwrites the old one, and no restart is needed.

On Windows, a different procedure is needed:

How to update a Fusion connector on Windows

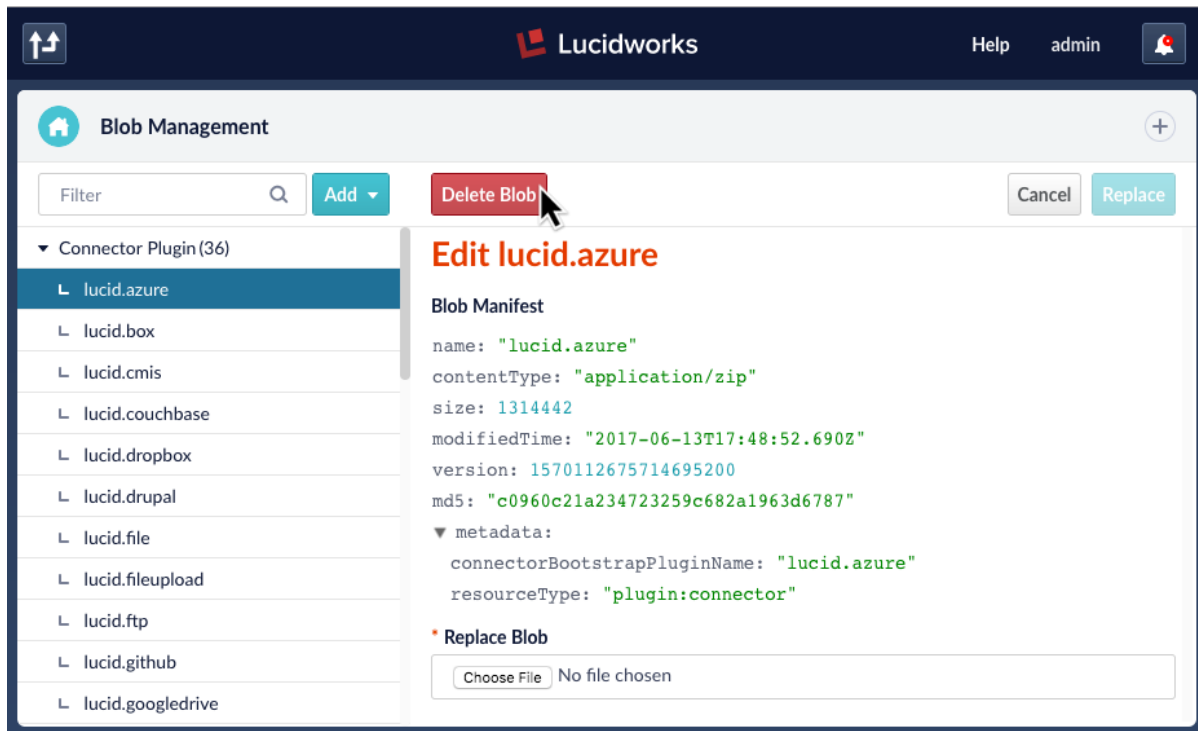
1. Delete the old connector, as explained below.
2. Restart Fusion.
3. Upload the new connector.

Deleting a connector

You can delete a connector using the Fusion UI or the Blob Store API.

Deleting a connector using the Fusion UI

1. In the Fusion UI, navigate to **DevOps > Blobs**.
2. Under **Connector Plugin**, select the connector to delete.
3. Click **Delete Blob**.



Fusion prompts you to confirm that you want to delete the blob.

4. Click **Yes, Delete**.

The connector disappears from the blob list.

Deleting a connector using the REST API

1. Get the list of blobs of the connector plugin type:

```
curl -u user:pass http://localhost:8764/api/apollo/blobs?resouType=plugin:connector
```

2. Locate the connector you want to delete, and copy its ID.

For example, the Jive connector ID is `lucid.jive`:

```
{
  "name" : "lucid.jive",
  "contentType" : "application/zip",
  "size" : 125302,
  "modifiedTime" : "2017-06-13T17:49:20.171Z",
  "version" : 1570112704530612224,
  "md5" : "7032bf2c038bb2d1e27aee82c056c0fb",
  "metadata" : {
    "connectorBootstrapPluginName" : "lucid.jive",
    "resourceType" : "plugin:connector"
  }
}
```

1. Delete the connector as follows:

```
curl -u user:pass -X DELETE http://localhost:8764/api/apollo/blobs/<id>
```

For example

```
curl -u user:pass -X DELETE http://localhost:8764/api/apollo/blobs/lucid.jive
```

A null response indicates success. You can verify that the connector is deleted like this:

```
curl -u user:pass http://localhost:8764/api/apollo/blobs | grep lucid.jive
```

Connector configuration

When you add a datasource to a collection, you select a connector and configure it. There are two ways to do this:

- Using the API
- Using the UI

Configuring Connectors Using The API

You can create or update a datasource with the Connector Datasources API, specifying the connector, its properties, and their values.

Example: *Create and configure a datasource to index Solr-formatted XML files*

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id":"SolrXML",
"connector":"lucid.solrxml", "type":"solrxml", "properties":{"path":"/Applications/solr-
4.10.2/example/exampledocs", "generate_unique_key":false, "collection":"MyCollection}}'
http://localhost:8764/api/apollo/connectors/datasources
```

See the Connectors and Datasources Reference for details about configuration options.

Tip	Be sure to include the <code>collection</code> property; otherwise the datasource will not be available in the Fusion UI.
-----	---

Example: *Change the `max_docs` value for the above datasource*

```
curl -u user:pass -X PUT -H 'Content-type: application/json' -d '{"id":"SolrXML", "connector":"lucid.solrxml", "type":"solrxml", "properties":{"path":"/Applications/solr-4.10.2/example/exampledocs", "max_docs":10}}' http://localhost:8764/api/apollo/connectors/datasources/SolrXML
```

Configuring Connectors Using The Fusion UI

- To create and configure a new datasource and its connector:
 1. Click **Applications > Collections**.
 2. Select a Collection, or click **Add a Collection** to create a new one.
 3. Click **Add a Datasource**.
 4. Select a datasource type; these correspond to Fusion's Connectors.
 5. Edit the configuration fields in the datasource panel that appears.

See the Connectors and Datasources Reference for details about configuration options.

6. Click **Save**.
- To change the connector configuration for an existing datasource:
 1. Navigate to **Applications > Collections**.
 2. Click your collection name.
 3. Click the datasource you want to change.
 4. Edit the configuration fields as needed.
 5. Click **Save**.

Connector logs

You can find connector logs in `fusion/3.1.x/var/log/connectors`.

Parsers

Parsers were introduced in Fusion 3.0 to provide more fine-grained configuration for inbound data. Parsers are configured in stages, much like index pipelines and query pipelines. They can include conditional parsing and nested parsing, and can be configured via the Fusion UI or the Parsers API.

Connectors receive the inbound data, convert it into a byte stream, and send the byte stream through the configured parsing stages. The stream moves through the parser stage by stage until it has been successfully parsed, then proceeds to the index pipeline.

Each parsing stage evaluates whether the inbound stream matches the stage's default media type or filename extension. The first stage that finds a match can output one or both of the following:

- Zero or more pipeline documents for consumption by the index pipeline
- Zero or more new input streams for re-parsing

This recursive approach is useful for containers (zip or tar files, for example). The output of the container parsing may be another container or a stream of uncompressed content which requires its own parsing.

There are a few static fields that impact the overall configuration and are accessible whenever you have selected the parser in the Index Workbench:

- Document ID Source Field
- Enable Automatic Media Type Detection
- Maximum Recursion Depth

Built-in parsing stages

These stages are available for configuration:

- HTML parser stage
- XML parser stage
- CSV parser stage
- JSON parser stage
- Text parser stage
- Archive parser stage
- Apache Tika parser stage
- Fallback parser stage

Datasources which use connectors that retrieve fixed-structure content (like Twitter or Jira) have hard-coded parsers and do not expose any configurable parser options.

HTML parser stage

This parser stage processes the following HTML elements:

- `<title>`
- `<body>` (with tags removed)

- `<meta>`
- `<a>` and `<link>`

Additionally, you can configure [JSoup selectors](#) to extract specific HTML and CSS elements from a document and map them to PipelineDocument fields. For example, you could use this to process navigational `DIV` elements one way, then process content-ful `DIV` elements another way.

See HTML parser stage for configuration details.

Note	The HTML Transformation index pipeline stage is deprecated in favor of this parser stage.
------	---

XML parser stage

The XML parser stage parses whole XML documents by default, but it can also be configured to parse only specific nodes without loading the entire document into memory. It can also split an XML document into multiple documents. XPATH-like expressions are used to select specific nodes to parse, such as `/posts/row` or `/posts/record`. Nested XML elements are flattened.

CSV parser stage

This parser breaks down incoming CSV files into the most efficient components for Fusion to index. It produces one new document per row from the CSV input, excluding comment rows and header rows.

See CSV parser stage for configuration details.

JSON parser stage

JSON parsing converts JSON content from a single document field into one or more new documents. This parser uses Solr's [JsonRecordReader](#) to split JSON into sub-documents.

See JSON parser stage for configuration details.

Text parser stage

The Plain Text parser can split a text file by lines or consume it into a single document.

Options for treatment of this filetype include:

- Plain Text Parser Fields
- Number of header rows to skip
- Split on line end or not
- Comment character
- Skip empty lines
- Charset

See Text parser stage for configuration details.

Archive parser stage

The Archive parser stage can parse the majority of common archive and compressed file formats. They are parsed into

their constituent documents, which can then be parsed further or sent straight to the index pipeline. The following archive formats are supported:

- tar
- zip
- jar
- 7z
- ar
- arj
- Unix dump
- cpio

See Archive parser stage for configuration details.

Apache Tika parser stage

Apache Tika is a versatile parser that supports many types of unstructured document formats, such as HTML, PDF, Microsoft Office documents, OpenOffice, RTF, audio, video, images, and more. A complete list of supported formats is available at <http://tika.apache.org/>.

See Apache Tika parser stage for configuration details.

Fallback parser stage

The Fallback parser stage is useful for processing data that Fusion does not have a specified parsing process for. Fallback does not technically parse data, since it does not know what to do with it, it simply copies the raw bytes into a Solr document. If your Fusion parser stage configuration encounters data it does not know how to parse, such as someone's proprietary data file format, it will copy it as-is, whereas if it encounters recognizable data in more common file types, such as PDFs, Fusion will parse the text and metadata using Tika.

The Fallback parser acts as the final stage that attempts to parse any documents that haven't been parsed already. When the correct parsing stage lands on the data, it executes accordingly.

See Fallback parser stage for configuration details.

Configuring parsers

When you configure a datasource, you can use the Index Workbench or the Parsers API to create a parser. A parser consists of an ordered list of parser stages, some global parser parameters, and the stage-specific parameters. You can re-order the stages list by dragging them up or down in the Index Workbench.

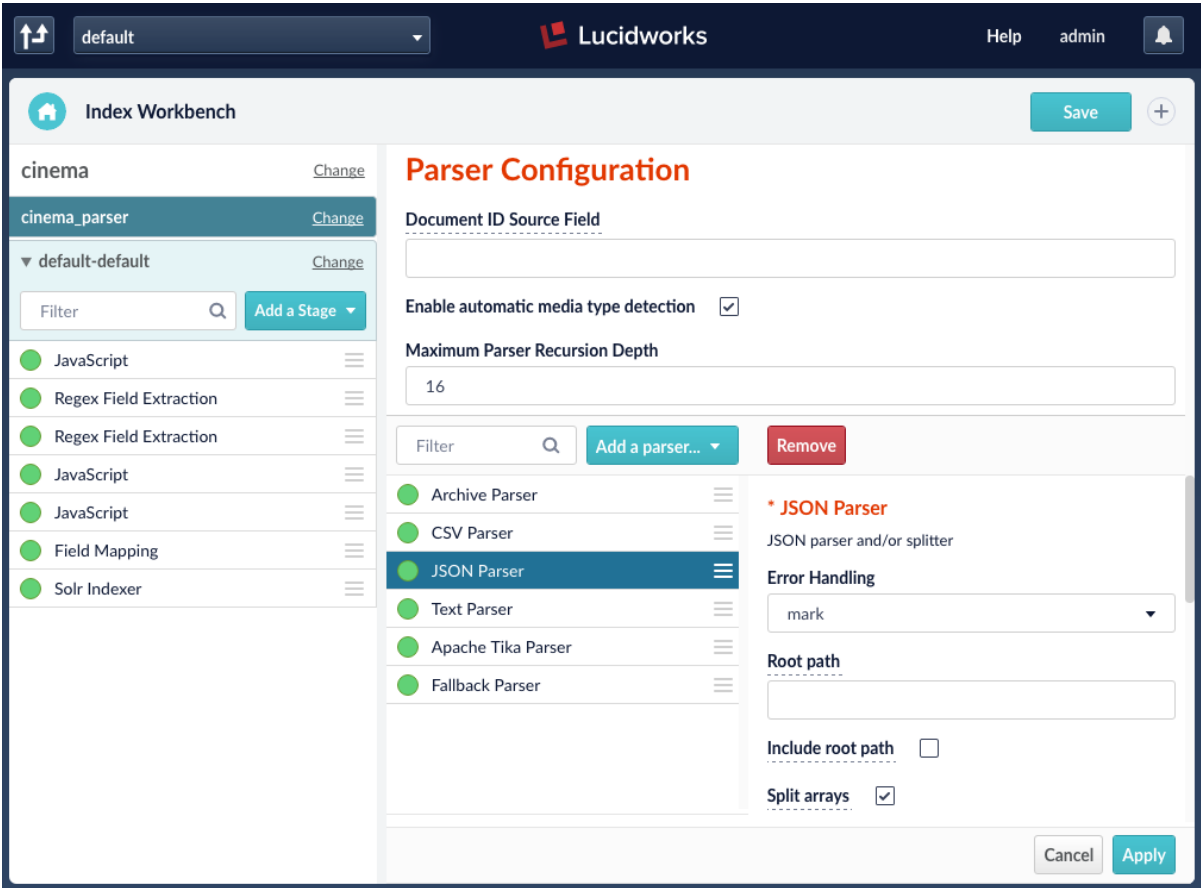
Any parser stage can be added to the same parser multiple times if different configuration options are needed for different stages. Datasources with fixed-structure data will also be parsed by Fusion, but with default settings that do not need to be customized.

There is no limit to the number of stages that can be included in a parser. The order in which they run is also completely flexible and can be linear or recursive. When the end of the parsing sequence is reached, a default parser stage automatically attempts to parse anything that has not yet been matched.

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Parser configuration in the Fusion UI

In the Search context, select a collection, then navigate to **Home > Index Workbench** and click the parser, usually called "*datasource-name* _parser". Clicking a specific stage opens its configuration panel.



Parser configuration in the REST API

The Parsers API provides a programmatic interface for viewing, creating, and modifying parsers, as well as sending documents directly to a parser.

- To get all currently-defined parsers: <http://localhost:8765/api/v1/parsers/>
- To get the parser schema: http://localhost:8765/api/v1/parsers/_schema

Here's a very simple parser example, for parsing JSON input:

```
{ "id": "simple-json",
  "type": "json",
    arbitrary parser-specific options here.
  "prettyfy": false
}
```

The example below shows a parser that can parse JSON input, as well as JSON that is inside zip, tar, or gzip containers, or any combination (such as .tar.gz). The order of the stages begins with the outermost containers and ends with the innermost content.

```
{ "id": "default-json",
  "type": "composite",
  "parsers": [
    { "id" : "zip-parser",
      "type" : "zip" },
    { "type" : "gz" },
    { "type" : "tar" },
    { "id": "json-parser",
      "type": "json",
      "prettify": false
    }
  ]
}
```

ID is optional, just as in pipeline stages. Many parser stages require no configuration other than **type**.

Parser index pipeline stage

The parsers themselves only parse whole documents. Parsing of content embedded in fields is performed separately by the Parser Index Pipeline Stage. This stage identifies the field or context that requires parsing, the appropriate parser to use, and what to do with the parsed content.

REST API Services

Fusion includes many REST APIs, all described in detail in the REST API Reference. Almost every aspect of Fusion can be controlled via the APIs.

The Fusion UI provides another administrative interface to Fusion, using a subset of the REST APIs.

API Service configuration

The Web service for the APIs is configured in `fusion/3.1.x/apps/jetty/api`.

API Service logs

API log files are in `fusion/3.1.x/var/log/api`.

Solr

Solr is the search platform that powers Fusion. There are multiple aspects to Fusion's use of Solr:

- Fusion components manage Solr search and indexing and provide analytics over these collections. Fusion's analytics components depend on aggregations over information which is stored in a Solr collection.
- Fusion collections are all Solr collections.
- Application data is stored as one or more Solr collections.
- Fusion's own logs are stored as Solr collections.
- A few Fusion service APIs use Solr as a backing store, notably Parameter Sets.

Solr configuration

Fusion requires that Solr run with [SolrCloud](#) enabled.

Configuration for Solr's Web service is in `fusion/3.1.x/apps/jetty/solr`.

Solr logs

Solr log files are in `fusion/3.1.x/var/log/solr`.

Accessing the Solr UI

With Fusion installed out of the box, you can still access the Solr UI at <http://localhost:8983/solr/>.

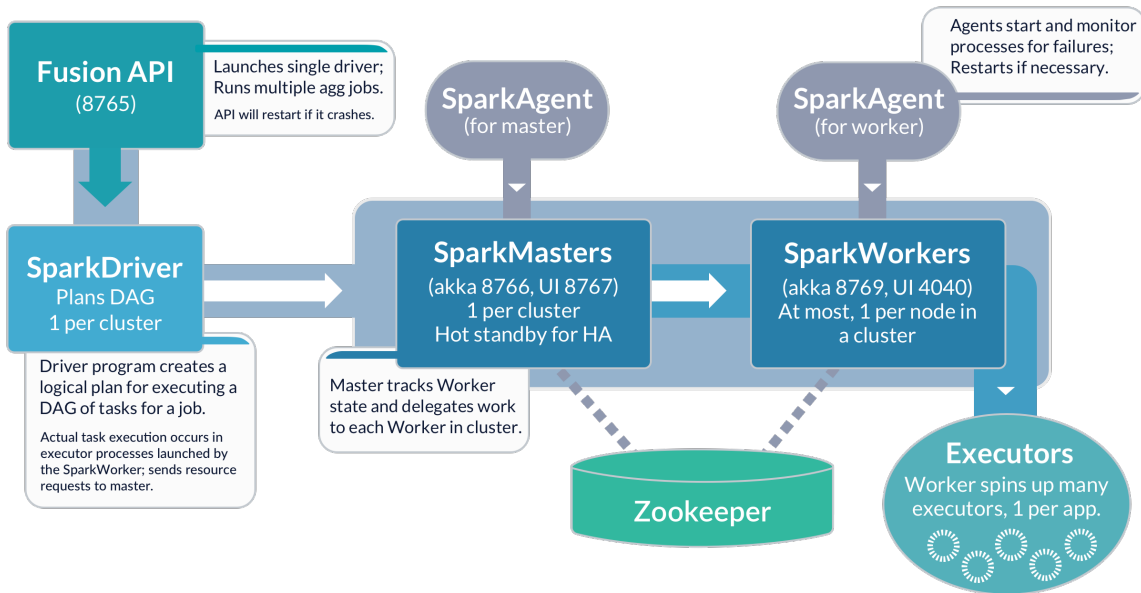
Solr documentation

Solr documentation and additional resources are available at <http://lucene.apache.org/solr/resources.html>.

You can also find plenty of Solr tips and technical discussions in our [knowledge base](#), [blog](#), and [webinars](#). Lucidworks also maintains a search interface to Solr's community discussions at searchhub.org.

Apache Spark

Apache Spark is a fast and general execution engine for large-scale data processing jobs that can be decomposed into stepwise tasks which are distributed across a cluster of networked computers. Spark provides faster processing and better fault-tolerance than previous MapReduce implementations. The following schematic shows the Spark components available from Fuson:



Fusion 2.1 introduced the use of a Spark cluster for all signal aggregation processes. This use of Spark is managed entirely by Fusion and is hidden from view of a Fusion application developer or end user.

As of Fusion 2.4, the Spark Jobs API provides a set of REST API endpoints for configuring and running Spark jobs via Fusion.

The section Spark and Machine Learning covers how to use Fusion's Spark cluster to run your own Spark jobs for custom aggregations and Machine Learning tasks.

ZooKeeper

[Apache ZooKeeper](#) is a distributed configuration service, synchronization service, and naming registry.

Fusion uses ZooKeeper to configure and manage all Fusion components in a single Fusion deployment, therefore a ZooKeeper service must always be running as part of the Fusion deployment. For high availability, this should be an external 3-node ZooKeeper cluster. All Fusion Java components communicate with ZooKeeper using the ZooKeeper API.

For ZooKeeper installation instructions, see the [ZooKeeper documentation](#).

You can find ZooKeeper's logs at `fusion/3.1.x/var/log/zookeeper`.

ZooKeeper Terminology

- **znode:** ZooKeeper data is organized into a hierarchal name space of data nodes called znodes. A znode can have data associated with it as well as child znodes. The data in a znode is stored in a binary format, but it is possible to import, export, and view this information as JSON data. Paths to znodes are always expressed as canonical, absolute, slash-separated paths; there are no relative reference.
- **ephemeral nodes:** An ephemeral node is a znode which exists only for the duration of an active session. When the session ends the znode is deleted. An ephemeral znode cannot have children.
- **server:** A ZooKeeper service consists of one or more machines; each machine is a server which runs in its own JVM and listens on its own set of ports. For testing, you can run several ZooKeeper servers at once on a single workstation by configuring the ports for each server.
- **quorum:** A quorum is a set of ZooKeeper servers. It must be an odd number. For most deployments, only 3 servers are required.
- **client:** A client is any host or process which uses a ZooKeeper service.

See the official [ZooKeeper documentation](#) for details about using and managing a ZooKeeper service.

Fusion ZooKeeper Nodes

Fusion configuration data is stored in ZooKeeper under two znodes:

- Node `lucid` stores all application-specific configurations, including collection, datasource, pipeline, signals, aggregations, and associated scheduling, jobs, and metrics.
- Node `lucid-apollo-admin` stores all access control information, including all users, groups, roles, and realms.

The Solr Admin tool provides a ZooKeeper node browser tool. In the case of the Fusion default developer deployment, the Fusion runs scripts are configured to run the instances of both Solr and ZooKeeper which are included with the Fusion distribution, and therefore we take a fresh installation of a Fusion developer instance and use the embedded Solr's Admin tool to explore how Fusion's configurations are managed in ZooKeeper.

On initial install, the "lucid" znode contains the set of default configurations used by Fusion's services:

version 0
 aversion 0
 children_count 20
 ctime Wed Jun 01 17:55:55 EDT 2016 (1464818155808)
 cversion 20
 czxid 35
 ephemeralOwner 0
 mtime Wed Jun 01 17:55:55 EDT 2016 (1464818155808)
 mzxid 35
 pzxid 811
 dataLength 0

Node "/live_nodes" has no utf8 Content

The "lucid-apollo-admin" znode contains the set of nodes used by Fusion's access control services:


version 0
 aversion 0
 children_count 0
 ctime Wed Jun 01 17:56:18 EDT 2016 (1464818178276)
 cversion 0
 czxid 209
 ephemeralOwner 0
 mtime Wed Jun 01 17:56:18 EDT 2016 (1464818178276)
 mzxid 209
 pzxid 209
 dataLength 0

Node "/lucid-apollo-admin/user" has no utf8 Content

Documentation Issue Tracker IRC Channel Community forum

In the above screenshot, the ZooKeeper node browser is browsing the contents of znode "lucid-apollo-admin/users" which is empty. The Fusion distribution ships without any user accounts. The initial user added to Fusion is the Fusion native realm "admin" user. This entry is only created on initial startup via the Fusion UI "set admin password" panel. Once you submit the admin password, the admin user account is created. Until Fusion contains at least the admin user account, you cannot use the system, because all Fusion requests require proper authorization.

Once the admin password is set, and you have created one or more Fusion collections and have populated them by running one or more datasources, these collections, datasources, pipelines, and other application configuration settings are stored under the "lucid" znode:



In the above screenshot, the ZooKeeper node browser is browsing the contents of znode "lucid/connectors/datasources/ds1". This datasource was used to populate a Fusion collection with documents retrieved via a webcrawl. Note that in the initial screenshot for znode "lucid", there is no "connectors" node at all.

The "lucid-apollo-admin" znode now contains one user accounts for user "admin":

Solr Admin

localhost:8983/solr/#/~cloud?view=tree

Solr

- Dashboard
- Logging
- Cloud
 - Tree
 - Graph
 - Graph (Radial)
 - Dump
- Core Admin
- Java Properties
- Thread Dump

Core Selector

File Tree:

- /
- /aliases.json
- /clusterstate.json
- /collections
- /configs
- /live_nodes
- /lucid
- /lucid-apollo-admin
 - collection-pref
 - proxy-startup-lock
 - realm-config
 - role
 - search-app-pref
 - sys
 - user
 - overseer
- /overseer_elect
- /zookeeper

version	0
aversion	0
children_count	0
ctime	Wed Jun 01 18:46:14 EDT 2016 (1464821174094)
cversion	0
czxid	3630
ephemeralOwner	0
mtime	Wed Jun 01 18:46:14 EDT 2016 (1464821174094)
mzxid	3630
pzxid	3630
dataLength	241

```
{ "created-at": "2016-06-01T22:46:14Z", "permissions": [], "role-names": [ "admin" ],
```

bin/fusion

For every server in a Fusion deployment, the script `fusion/3.1.x/bin/fusion` is used to start, stop, and check the status of the Fusion instance running on that server.

Fusion Agent Process

The Fusion agent process makes sure that all Fusion processes start up and shut down correctly. It prevents problems that can arise by trying to start Fusion on a server where it is already running.

Start Fusion

Run the script `fusion/3.1.x/bin/fusion` with the argument `start`:

```
$ cd /path/to/fusion/3.1.x
$ ./bin/fusion start
Starting zookeeper.
Successfully started zookeeper on port 9983 (process ID 77295)
Starting solr.
Successfully started solr on port 8983 (process ID 77297)
Starting api.....
Successfully started api on port 8765 (process ID 77301)
Starting connectors.....
Successfully started connectors on port 8984 (process ID 77388)
Starting ui....
Successfully started ui on port 8764 (process ID 77469)
```

Check the status of Fusion

Run the script `fusion/3.1.x/bin/fusion` with the argument `status`:

```
$ cd /path/to/fusion/3.1.x
$ ./bin/fusion status
zookeeper is running on port 9983 (process ID 77295)
solr is running on port 8983 (process ID 77297)
api is running on port 8765 (process ID 77301)
ui is running on port 8764 (process ID 77469)
connectors is running on port 8984 (process ID 77388)
```

Stop Fusion

Run the script `fusion/3.1.x/bin/fusion` with the argument `stop`:

```
$ cd /path/to/fusion/3.1.x
$ ./bin/fusion stop
Successfully stopped ui (process ID 41524)
Successfully stopped connectors (process ID 41328)
Successfully stopped api (process ID 41159)
Successfully stopped solr (process ID 41153)
Successfully stopped zookeeper (process ID 41151)
```

Troubleshooting

The [Java Virtual Machine Process Status Tool](#) utility at `/usr/bin/jps` is useful for reporting on all Fusion processes reported by script `fusion/3.1.x/bin/fusion`:

```
$ jps
77294 AgentMain
77295 zookeeper-path-1475182112123.jar
77297 start.jar
77301 start.jar
77388 start.jar
77469 start.jar
79455 Jps
```

The process `zookeeper-path-1475182112123.jar` is the ZooKeeper process used by Fusion. The 4 `start.jar` processes are Fusion's Solr, API Services, Connectors, and UI.

If the `path/to/fusion/3.1.x/bin/fusion` script doesn't run, or if it fails to start all services, see the Troubleshooting topic or the [knowledge base](#) for help.

1.3.3. Deployment Types

The Fusion platform is designed to support enterprise search applications at any scale. You can deploy Fusion across multiple nodes in order to store large amounts of data or to achieve high processing throughput or both.

Deployment goals

- **Demo, trial, and development deployments** – The simplest possible architecture is the one you get out of the box, by unpacking the tar/zip file and running `fusion/3.1.x/bin/fusion start`, so that all components are running on a single host in their default configurations.

You can quickly install and run Fusion on a computer (even on your laptop) to explore its features and work with sample data. See [Getting Started for Quickstart instructions](#). This diagram illustrates a single-node Fusion deployment:

Computer

Fusion

Solr

ZooKeeper

- **Production deployments** – Fusion is designed for flexible, distributed deployment. Any of its components can be distributed across your network, and some can be clustered. A production deployment usually requires multiple Fusion nodes, each of which runs some or all Fusion services.

If you already have SolrCloud clusters managing your data, you can integrate them into a Fusion deployment using ZooKeeper.

1.4. Fusion Installation and Upgrades

1.4.1. Installation

You can install Fusion:

- **On a single computer** – Install Fusion on a single computer for purposes of development, evaluations, or demonstrations.

For information, see [Install Fusion on a Single Computer](#).

- **On multiple servers** (as a *Fusion cluster*) – For a number of reasons such as performance, scaling, and availability, you *must* install Fusion on multiple servers for a production deployment.

For information, see [Install a Fusion Cluster](#).

These procedures are for *initial* installations of Fusion.

1.4.2. Upgrades

If you already have Fusion installed and want to upgrade to a later release, see the [upgrade documentation](#).

1.4.3. Related Topics

Additional documentation that pertains to installation includes:

- [System Requirements](#) – Hardware and software needed to run Fusion in both single-node (trial) and multi-node (production) deployments
- [Deployment Types](#) – An explanation of deployment goals and cluster arrangements for Fusion, Solr, and ZooKeeper
- [Components](#) – An explanation of the components that comprise Fusion
- [Start and Stop Fusion](#) – An explanation of the components that comprise Fusion and of deployment types
- [Troubleshoot When Installing Fusion](#) – An explanation of how to troubleshoot difficulties that occur when installing or upgrading Fusion
- [Directories and Logs](#) – What's in the Fusion home directory (`/path/to/fusion/3.1.x`), including default log file locations
- [Default Ports](#) – Default Fusion server ports and how to change them
- [Checking System State](#) – How to check the default Fusion distribution
- [Integrating with existing Solr instances](#) – How to use existing Solr instances to store Fusion collections

After installation, proceed to these topics:

- [Security](#) – How to secure a Fusion deployment
- [Administration](#) – Information about configuration and monitoring

1.4.4. System Requirements

Requirements for Fusion installation are detailed below.

Important	Lucidworks recommends <i>not</i> virus scanning the <code>fusion/data</code> folder. Virus scanning can cause slow performance, and it can cause downtime if it quarantines an index file identified as a possible virus.
-----------	---

Supported Operating Systems

Supported for production use of Fusion services

Operating system:

- Windows Server 2012, 2012 R2, and 2016 (x64 only)

Note	Windows systems must have the Microsoft Visual C++ 2010 SP1 package installed.
------	--

- Linux 64-bit (x86_64/amd64 only) with 2.6 or later kernel

Note	RedHat based Linux, including CentOS, must be 6.6.x or later, or else the bug fix https://rhn.redhat.com/errata/RHSA-2013-1605.html must have been applied. This fix remediates a RedHat bug that causes Fusion to hang. See http://www.infoq.com/news/2015/05/redhat-futex for more information.
------	--

JVM:

- Oracle JRE or JDK 1.8, 64-bit (x64 only)
- OpenJDK JRE or JDK 1.8, 64-bit (x64/x86_64 only)

Also supported for trial and development use of Fusion services

Operating system:

- Windows 7, 8, 8.1, and 10 64-bit (x64 only)
- Mac OS X 10.8 or later

Hardware requirements

Fusion and Solr nodes

Here are some *minimum* recommendations for different Fusion deployments. These are for the nodes that run Fusion (including ones that *also* run Solr and/or ZooKeeper, if those cluster arrangements are used), as well as for external (non-Fusion) nodes for SolrCloud nodes:

Deployment type	Memory	CPU
Development/Testing	16 GB	4 cores

Deployment type	Memory	CPU
Small Production*	32 GB	8 cores
Large Production**	64+ GB	8+ cores

- Small production environments will have 2+ nodes with these specs, on which both Fusion and Solr are installed. (3+ nodes are recommended.)
 - Large production environments will have 3+ nodes with these specs. Additional nodes are needed for an external SolrCloud cluster, if that is used. An external SolrCloud cluster is only necessary for Solr collections with very large numbers of documents (for example, 100 million documents or more), though a Fusion deployment can use it regardless. On these nodes, Fusion and Solr are installed if no external SolrCloud cluster is used. If an external SolrCloud cluster is used, these nodes might or might not have Solr installed.

ZooKeeper nodes

Here are some *minimum* recommendations for the nodes that run ZooKeeper. These nodes can be more lightweight than the nodes that run Fusion and/or Solr. With the exception of small development/testing deployments, ZooKeeper nodes should only run ZooKeeper (not Fusion or Solr). Because ZooKeeper synchronizes all its operations to disk, we recommend using a disk with high throughput and low latency for your ZooKeeper nodes.

Deployment type	Memory	CPU
Development/Testing	4 GB	4 cores
Small Production	4+ GB	4+ cores
Large Production	4+ GB	4+ cores

Fusion components

You can run Fusion components on different nodes. Different Fusion components require different amounts of resources. Below are the *minimum* recommended memory requirements; consult Lucidworks for specific recommendations tailored to your unique use case, data load, and production needs.

Fusion component	Minimum RAM
Connectors	2 GB
API service	1 GB
Fusion UI	512 MB
Solr	2 GB
Spark master	512 MB
Spark worker	1 GB

Java

Fusion is a Java-based application, and thus requires a pre-installed JDK.

Fusion runs on [JDK 1.8](#).

Fusion's JavaScript pipeline stages are application-specific custom JavaScript programs. The first time a pipeline is run,

the JDK compiles the JavaScript using the Nashorn JavaScript engine.

Fusion scripts execute both the `java` and `javac` commands, which are usually aliases for the current Java installation. To check which version of Java is invoked, run the following commands from a shell or terminal window:

```
java -version
javac -version
echo $JAVA_HOME // Unix
echo %JAVA_HOME% // Windows
```

Cluster Requirements

Supported Solr and SolrCloud Versions

Fusion includes an embedded instance of Solr; see the release history to find out which Solr version is included in each Fusion release.

If your search requirements are *very large* (for example, 100 million documents or more), we recommend that you use an external SolrCloud cluster. (You can use an external SolrCloud cluster regardless.)

Optionally, you can use *both* embedded Solr instances *and* an external SolrCloud cluster. In this case, only store logs in the embedded Solr instances. Store data in the external SolrCloud cluster.

Solr 4.10 and higher are supported.

If you want to know whether other versions of Solr might be compatible with Fusion, you can [ask Lucidworks](#).

If you use an external SolrCloud cluster, it has ZooKeeper bundled with it. However, Apache [recommends that you not use the bundled ZooKeeper in production](#). Instead, create an external ZooKeeper cluster (external to both Fusion and SolrCloud) or use the ZooKeeper embedded with Fusion, depending on the circumstance. For more information, see Supported ZooKeeper Versions and Cluster Requirements.

Note	If you decide to set up an external SolrCloud cluster, check the requirements for this as well in the Solr documentation .
------	--

We strongly recommend that you use Network Time Protocol (NTP) on a SolrCloud cluster to ensure that nodes use synchronized time. While this is not strictly required, reasoning about log contents and database entries becomes impossible without it. Information and instructions on how to install and run NTP are available at www.ntp.org.

Supported ZooKeeper Versions

Fusion includes an embedded instance of ZooKeeper; see the release history to find out which version of ZooKeeper is included in each Fusion release. For demo, trial, and some testing deployments, you can use the embedded ZooKeeper. *For production deployments, we recommend using an external ZooKeeper cluster.*

ZooKeeper 3.4.6 is supported for external ZooKeeper clusters.

A ZooKeeper cluster must have 3+ nodes configured as an ensemble, and it must have an odd number of nodes. We recommend three ZooKeeper nodes in a ZooKeeper cluster/ensemble. Only for very large SolrCloud clusters (50+ nodes) might you need more ZooKeeper nodes.

If you want to know whether other versions of ZooKeeper might be compatible with Fusion, you can [ask Lucidworks](#).

Recommended HTTP Clients

The Fusion API can be accessed from any HTTP client, and allows you to build user interfaces and applications that work with any browser. However, the Fusion Administration UI, Dashboards, and built-in Search UI are supported only with:

- Chrome latest version
- Firefox latest version and latest ESR
- Internet Explorer 11

1.4.5. New Installation

Install new instances of Fusion:

- Install Fusion on a Single Node

Start here for the simplest deployment, which is useful for evaluation, demo, and development purposes. You can install Fusion on a server, desktop, or laptop.

- Install a Fusion Cluster (Introduction)

For production deployments, install Fusion on multiple servers (a *Fusion cluster*). Start with this introduction.

- Install a Fusion Cluster (Unix)

Install a Fusion Cluster on Unix nodes. These instructions also cover relevant aspects of installing the SolrCloud and ZooKeeper clusters.

- Install a Fusion Cluster (Windows)

Install a Fusion cluster on Windows nodes. These instructions also cover relevant aspects of installing the SolrCloud and ZooKeeper clusters.

- Integrate Fusion with an Existing Solr Deployment

Use these instructions when installing Fusion in an existing Solr environment.

Install Fusion on a Single Node

Note	These instructions are for an initial installation of Fusion on a single node (computer). To install Fusion on multiple nodes (a <i>Fusion cluster</i>), see Install a Fusion Cluster. If you already have a version of Fusion installed and want to upgrade it, see the Fusion upgrade instructions.
------	--

Out of the box, Fusion uses the instances of Solr, ZooKeeper, and Spark that are included in the Fusion distribution. See the release history to find out which versions of Solr, Spark, and ZooKeeper are included in each Fusion release.

To use Fusion with your existing Solr installations, see Integrating with existing Solr instances.

Ports

To run Fusion as a single-server installation, the following ports should be available and not used by other applications or services:

- 8764
- 8765
- 8983
- 8984

See Default Ports if you need to modify the default Fusion ports before starting the application.

Unix installation

Fusion for Unix is distributed as a gzipped tar file.

To install Fusion on Linux or Mac

1. [Download the Fusion tar/zip file](#) for the latest version of Fusion and move it to where you would like it to reside in your filesystem (if you would like to use Upstart for process management, you must install Fusion in `/opt/lucidworks`).
2. Become the user that will run Fusion.

Important	Do not run Fusion as the root user.
-----------	-------------------------------------

3. Change your working directory to the directory in which you placed the `fusion-version.x.tar.gz` file, for example:

```
$ cd /opt/lucidworks
```

4. Unpack the archive with `tar -xf` (or `tar -xvf`), for example:

```
$ tar -xf fusion-version.x.tar.gz
```

The resulting directory is named `fusion/3.1.x`. You can rename this if you wish. This directory is considered your Fusion home directory. See Directories and Logs for the contents of the `fusion/3.1.x` directory.

Starting Fusion

All Fusion start scripts must be executed by a user who has permissions to read and write to the directories where Fusion is installed. These scripts don't need to be run as root (or sudo), nor should they be. Use a suitable ID, or create a new one, and then ensure that it owns the directory where Fusion resides, (e.g. `/opt/lucidworks`).

To start all required services

```
./bin/fusion start
```

For information about starting individual services, see Start and Stop Fusion.

Running Fusion In The Foreground

To run Fusion or any of its services in the foreground, use the `run` command-line argument in place of `start`.

Stopping Fusion

To stop Fusion or any of its services, use the `stop` command-line argument in place of `start`.

Ubuntu Upstart Scripts

Under Ubuntu 12.04 LTS or newer we support Upstart for process management. This requires Fusion to be installed in the `/opt/lucidworks/` directory.

To configure upstart, run the following commands:

```
$ cd init/upstart
$ sudo bash install.sh
```

If this complains with `no JAVA_HOME set`, replace `sudo` with `sudo -E`. Then you can use the `service` command to control the server:

```
$ sudo service fusion-solr start
$ sudo service fusion-api start
$ sudo service fusion-connectors start
$ sudo service fusion-ui start
```

and similarly use `stop` and `status`.

Logfiles for Fusion services are found in directories under `fusion/3.1.x/var/log`.

Upstart log files for each service are in the `/var/log/upstart` directory.

For convenience, you can start/stop all services with Upstart using:

```
$ sudo bash start.sh
$ sudo bash stop.sh
```

Windows installation

Fusion for Windows is distributed as a compressed zip file. To unpack the Fusion zip file on Windows, you can use a

native compression utility or the freely available [7zip](#) file archiver. Visit the [7zip download page](#) for the latest version.

To install Fusion on Windows

1. Download the zip file for the latest version of Fusion and move it to where you would like Fusion to reside in your filesystem. It will appear as a compressed folder.
2. Unpack the archive. In most cases, you need only right-click and choose "Extract all...". If you don't see this option, check that you have permissions to extract folders on your system.

The resulting directory is named `fusion\3.1.x`. This directory is considered your Fusion home directory. See Directories and Logs for the contents of the Fusion home directory.

To install Fusion as a set of Windows services

1. Run `bin/install-services.cmd`.
2. Enter the name of the windows user that is used to launch this service.

Remember the username is `COMPUTERNAME\username` or `DOMAIN\username` (if your computer is part of a Windows domain).

3. Enter the user's password.
4. Enter the path to the directory containing the JDK to use for running the services.

Starting Fusion

All Fusion start scripts must be executed by a user who has permissions to read and write to the directories where Fusion is installed. Ensure that such a user owns the directory where Fusion resides.

To start all required services

- `bin\fusion.cmd start`
- `bin\start-services.cmd` (Start all Fusion services as Windows services)

For information about starting individual services, see Start and Stop Fusion.

Stopping Fusion

To stop Fusion or any of its services, use the `stop` command-line argument in place of `start`.

Installation with an existing Solr instance or cluster

Fusion supports Solr versions 4.4 and higher. Solr 4.6.0 or 4.7.0 are not supported, as they contain severe bugs that will impact the ability of Fusion to work with your Solr system.

If installing Fusion to work with an existing Solr instance, either in SolrCloud mode or standalone, you should install Fusion as described above. You should start each of the services as described above.

Once Fusion installation is complete, you can register your existing Solr installation with Fusion to be able to use the two systems together. For details on how to do that, see the section Search Clusters.

Troubleshooting

For information about problems you might encounter when installing Fusion, and solutions, see Troubleshoot When

Installing Fusion.

Integrate Fusion with an Existing Solr Deployment

If you have already implemented Solr as a standalone instance or as a SolrCloud cluster, you can add Fusion to your existing Solr deployment and import your Solr collections into Fusion. Each Fusion collection can import one Solr collection.

- *If your existing Solr instance is running in SolrCloud mode*, you can use Fusion's UI to modify configuration files (such as schema.xml or solrconfig.xml) and create Solr collections.
- *If your existing Solr instance is running in standalone mode*, you can still connect it to Fusion. Fusion can send documents to a standalone Solr instance and query the instance. But you won't be able to use Fusion's UI to create Solr collections (Solr cores) or to modify Solr configuration files.

Prerequisites

- You have already installed Fusion.
- You have already installed Solr, which must meet these Solr requirements.
- You have already installed ZooKeeper, which must meet these ZooKeeper requirements.

Note	We recommend that you create an <i>external</i> ZooKeeper cluster (external to both Fusion and SolrCloud).
------	--

- Your Solr deployment must contain one or more collections (cores).
- In SolrCloud mode, Solr must be configured to use ZooKeeper.

Configure Fusion to use an existing Solr deployment

Use the Fusion UI or the Fusion API to integrate Fusion with an existing Solr deployment.

Use the Fusion UI

1. Create a Fusion search cluster:
 - a. In the Fusion UI, navigate to **System > Solr Clusters** and click **New Solr Cluster**.
 - b. Enter this information:
 - A cluster ID of your choice
 - Whether SolrCloud is enabled
 - The connect string (to tell Fusion how to connect to the SolrCloud cluster or Solr instance)
 - For SolrCloud, this is the ZooKeeper connect string.
 - For a standalone Solr instance, this is the URL of the Solr instance.
 - c. Verify that the connection is working by clicking **Cores** in the new cluster and inspecting the contents.
2. Create a Fusion collection that points to your Solr cluster and collection:
 - a. In the UI, navigate to **Collections** and click **Add a Collection**.
 - b. Enter a name for the new collection.
 - c. Click **Advanced**.
 - d. Select your SolrCloud cluster or Solr instance from the dropdown.
 - e. Enter the name of the Solr collection to import.

Use the Fusion API

Use the Search Cluster API to create a Solr cluster.

Then use the Collections API to create and configure a collection.

Sending Documents to Solr through Fusion

You can use the Fusion connectors to crawl documents and index them to your existing Solr deployment.

1. Follow the steps above to create and configure a search cluster and a collection that points to Solr.
2. Define an index pipeline that ends with a Solr Indexer stage that sends the documents to Solr.
3. Use one of these methods to ingest your data:
 - In the collection that points to your Solr collection, define a datasource using the connector of choice.
 - Send prepared documents directly to the index pipeline for processing. See [Pushing Documents to a Pipeline](#).
 - It's also possible to use a different indexing process besides a connector, such as a script that sends documents through the index pipeline.

When documents are sent to Solr, a buffering `solrServer` is used. Buffering the updates reduces the number of HTTP requests made from Fusion to Solr, which can significantly affect processing time. For example, when processing simple documents, you should always try to buffer as many documents as possible to increase throughput. When processing complex documents, you should use small batch sizes. You should only turn buffering off if you are using an older version of Solr and you want Fusion to catch and document indexing errors.

Querying Solr via Fusion requests

Indexed documents are stored in Solr indexes. You can query for these documents by using query pipelines. The query pipelines let you define your query parameters – such as how many records to return, the fields you'd like, how to structure facets, and so on. You also have the ability to add JavaScript to the response processing, and define landing pages or specific boost levels depending on the user's query. See [Query Pipelines](#).

If you prefer, you can also use the Solr API and SolrAdmin API to query Solr directly.

1.4.6. Upgrade Fusion

After you have a Fusion-based search application running, at some point it might be necessary to upgrade to a later version of Fusion. Your goal is to transfer over all of your data together with all configurations and customizations necessary to support your applications.

In this topic, we:

- Provide links to upgrade instructions for specific versions of Fusion.
- Give a general overview of the procedure for upgrading.
- Present notes about version incompatibilities.

Tip	See the release history to find out what's new, including which versions of Solr, Spark, and ZooKeeper are bundled with each Fusion release.
-----	--

Per-version instruction sets

To upgrade to a later version of Fusion from an existing installation requires transferring over all configurations and data from your existing Fusion installation to the new version.

Perform the steps in the appropriate section:

- Upgrade Fusion 2.1 to Fusion 2.4 – Instructions for upgrading Fusion versions 2.1.x and 2.2.1 to Fusion 2.4.
- Upgrade Fusion 1.2 to Fusion 2.4 – Instructions for upgrading Fusion versions 1.2.3 through 1.2.8 and version 1.4.0 to the latest version of Fusion 2.4.
- Upgrade Fusion 2.4 to Fusion 3.0 – Instructions for upgrading Fusion version 2.4.x to Fusion 3.0.0.
- Upgrade Fusion 3.0.0 to 3.0.1 – Instructions for upgrading Fusion version 3.0.0 to 3.0.1.
- Upgrade Fusion 3.x to a more recent release – Instructions for upgrading Fusion version 3.x to a more recent release. This procedure uses a migrator that migrates all Fusion objects in a seamless upgrade.

Upgrade overview: migrating data, configurations, and customizations

The Upgrade Process

The upgrade process leaves the current Fusion deployment in place while a new Fusion deployment is installed and configured. All of the upgrade operations copy information from the current Fusion over to the new Fusion. This provides a rollback option should the upgrade procedure encounter problems.

The current Fusion configurations must remain as-is during the upgrade process. In order to capture indexing job history, no indexing jobs should be running. If the new Fusion installation is being installed onto the same server that the current Fusion installation is running on, you must either run only one version at a time or else change the Fusion component server ports so that all components are using unique ports for both the current and new versions.

ZooKeeper

Migration consists of the following steps:

- Copy the ZooKeeper data nodes which contain Fusion configuration information from the FUSION-CURRENT ZooKeeper instance to the FUSION-NEW ZooKeeper instance

- Rewrite Fusion datasource and pipeline configurations, working against the FUSION-NEW ZooKeeper instance

Important	Because some Fusion configurations have changed, ZooKeeper data must be rewritten accordingly using scripts available in the public GitHub repository: https://github.com/LucidWorks/fusion-upgrade-scripts .
-----------	--

Solr

Fusion-based search applications store your data in Solr. If your data is stored in an external Solr cluster, and if you aren't upgrading your Solr cluster, then you don't need to migrate your Solr data at all; you just need to configure the new Fusion deployment to use your external Solr cluster.

Fusion uses Solr as a data store for server logs, search logs, and binary components such as jar files and compiled models used by Fusion pipelines. The Fusion distribution includes a complete Solr server and is configured to use this embedded Solr instance by default. If the current Fusion deployment uses the embedded Solr for its system collections, then you must copy all of these collections over to the embedded Solr instance included with the new Fusion distribution.

Connector services data: crawldb, database drivers

The directory `fusion/3.1.x/data/connectors/lucid.jdbc` contains third party JDBC driver files that have been registered with Fusion in order to run the JDBC connector.

The directory `fusion/3.1.x/data/connectors/crawldb` is managed by Fusion's connector service. Fusion datasources that walk over websites, filesystems, or similar repository use the crawldb to store information about files visited during the crawl; this allows incremental updates and avoids data re-indexing. In current versions of Fusion, the default location is the Fusion directory `fusion/3.1.x/data/connectors/crawldb`.

Important	The crawldb data format was changed, therefore for upgrades from Fusion 1.2.x to the latest Fusion 2.1, these must be processed using a reformatting program that is available in the public upgrade scripts repository: https://github.com/LucidWorks/fusion-upgrade-scripts .
-----------	--

Pipeline services data: models used by pipeline stages

The directory `fusion/3.1.x/data/connectors/lucid.jdbc` contains third party JDBC driver files that have been registered with Fusion in order to run the JDBC connector.

Customized settings for Fusion run commands and configuration scripts

The scripts used to start, stop, and restart Fusion and its components are found in the top-level "bin" directory of the Fusion distribution. As of Fusion 2.0, the configuration scripts previously found in the "bin" directory were put in their own top-level directory called "conf".

Important	If you have customized the settings in these files and wish to carry these settings over to the new Fusion deployment the only way to do so is to edit the command and configuration scripts in the new Fusion deployment by hand. You cannot copy over the old configuration files because they may contain commands or settings which are no longer valid.
-----------	--

Custom Fusion connector plugins and pipeline stages

Custom Java components written for one version of Fusion must be re-compiled and installed anew for the latest version of Fusion as Fusion's Java API may have changed.

Version Incompatibilities

There are no version incompatibilities between Fusion 2.4 and later versions.

Upgrade Fusion 3.0.0 to 3.0.1

The steps below describe how to upgrade Fusion 3.0 to a more recent release. The `kazoo` Python library must be installed before you begin.

Note	A migration tool will be available soon for upgrading to Fusion 3.1 from versions 3.0.x
------	---

How to upgrade Fusion 3.0.0 to 3.0.1

1. [Download the latest version of Fusion.](#)
2. Extract the zip file to the same parent `fusion` folder as your earlier 3.0 release.

For example, if Fusion 3.0.0 is installed in `/opt/fusion/3.0.0`, then place the downloaded `fusion-3.0.1.tar.gz` file in the `/opt/` folder and run:

```
tar -xf fusion-3.0.1.tar.gz
```

3. Set the `FUSION_OLD_HOME` environment variable to the full path of the version-numbered folder below the `fusion` directory for the older version of Fusion:

```
export FUSION_OLD_HOME=/path/to/fusion/3.0.0
```

4. Set the `FUSION_HOME` environment variable to the full path of the version-numbered folder below the `fusion` directory for the new version of Fusion:

```
export FUSION_HOME=/path/to/fusion/3.0.1
```

5. Copy data from the old Fusion instance to the new instance:

```
cp -R $FUSION_OLD_HOME/data/ $FUSION_HOME/
```

6. Copy configurations from the old Fusion instance to the new instance:

```
mv $FUSION_HOME/conf/ $FUSION_HOME/conf_backup/  
cp -R $FUSION_OLD_HOME/conf/ $FUSION_HOME/
```

If you are running distributed Fusion, then perform this step on all the Fusion nodes.

Note	If you are running external Solr or Zookeeper, then modify the <code>group.default</code> property in <code>fusion.properties</code> to reflect the services that should be started when the <code>bin/fusion</code> script is executed.
------	--

7. Update links to point to the new Fusion instance:

```
cd $FUSION_HOME/..
unlink latest
ln -s 3.0.1 latest
```

8. Start the Zookeeper server which will be used by the new Fusion installation.

If you are using the Zookeeper bundled within Fusion, that would be:

```
cd $FUSION_HOME
./bin/zookeeper start
```

9. Clone the fusion-upgrade-scripts repository, if you haven't already:

```
git clone https://github.com/lucidworks/fusion-upgrade-scripts
```

10. Run the upgrade script:

```
cd fusion-upgrade-scripts/src
python upgrade-3.0.x.py
```

11. Run the new version of Fusion (all services defined in `fusion.properties`) and validate via the Fusion UI:

```
cd $FUSION_HOME
./bin/fusion start
```

Tip	Before accessing the Fusion UI, clear your browser's cache. Otherwise, you may inadvertently access a cached version of the old Fusion UI and see inconsistent behavior.
-----	--

Upgrade Fusion 2.1.4 or 2.4.x to Fusion 3.0.0

This article describes how to upgrade from Fusion 2.1.4 or 2.4.x to Fusion 3.0.0 for Linux and Windows.

Note	To upgrade to Fusion 3.0.0 from an earlier version of Fusion, first follow the steps to upgrade from 1.2 to 2.4; then follow the steps below.
------	---

Linux

The steps in this section describe how to upgrade from Fusion 2.1.4 or 2.4.x to Fusion 3.0 on Linux.

Prerequisites

Ensure that these prerequisites are met before upgrading Fusion.

Note	If you have multi-node Fusion, this software is required on all nodes.
------	--

The [upgrade scripts](#) require Python 2.7 or later, and that the following Python libraries are already installed:

- [kazoo](#)
- [requests](#)

Steps to upgrade from Fusion 2.4.x to 3.0.0 on Linux

These steps describe how to upgrade from Fusion 2.1.4 or 2.4.x to Fusion 3.0.0 on Linux.

Note	If you have a multi-node installation and unless otherwise stated, you must perform each step on each node as you go. For example, if you have 3 nodes, perform step 1 on node 1, node 2, and then node 3. Next, perform step 2 on node 1, node 2, and then node 3. Run steps that are only needed on a single node on your first node; not on the other nodes.
------	---

1. (On all nodes) [Download Fusion 3.0.0](#).
2. (On all nodes) Extract Fusion to a folder. Here we'll call it `fusion-new`, as an example:

```
mkdir fusion-new
tar -C fusion-new --strip-components=1 -xf fusion-3.0.0.tar.gz
```

3. (On all nodes) Set the `FUSION_OLD_HOME` environment variable to the location of the old version of Fusion (2.x.y):

```
export FUSION_OLD_HOME=/path/to/fusion/2.x.y
```

4. (On all nodes) Set the `FUSION_HOME` environment variable to the location of the new version of Fusion (3.0):

```
export FUSION_HOME=fusion-new/3.0.0
```

5. (On all nodes) Copy data from the older fusion instance to the new fusion instance:

```
cp -R $FUSION_OLD_HOME/data/* $FUSION_HOME/data/
```

6. (On all nodes) Clone the fusion-upgrade-scripts repository, if you haven't already:

```
git clone https://github.com/lucidworks/fusion-upgrade-scripts
```

7. (On all nodes) Upgrade the configuration. This upgrades the customized properties in `$FUSION_OLD_HOME/conf/config.sh` to the new properties file in 3.0.0 (`$FUSION_HOME/conf/fusion.properties`).

```
cd fusion-upgrade-scripts/src
python upgrade-to-3.0.py --upgrade config
```

Alternatively, if all of the Fusion nodes use the same configuration, then you can just run this on one node, and then copy the new `fusion.properties` file to all of the other nodes.

8. (On all nodes) Check the upgrade configuration in `$FUSION_HOME/conf/config.sh` and ensure that all of the modified properties in `config.sh` are reflected. Pay particular attention to the ZK connection strings if your existing Fusion installation is connected to an external ZooKeeper cluster.
9. (On all nodes) If you are running an external Solr or Zookeeper, then modify the `group.default` property in `fusion.properties` to reflect the services that should be started when the `bin/fusion` script is executed.
10. (On all nodes) Start the ZooKeeper service that the new Fusion installation will use. If you are using the ZooKeeper bundled within Fusion, that would be:

```
cd $FUSION_HOME
./bin/zookeeper start
```

11. (On all nodes) Copy your ZooKeeper configuration file over from the old installation to the new one:

```
cp $FUSION_OLD_HOME\conf\zookeeper\zoo.cfg $FUSION_NEW_HOME\conf\zookeeper\zoo.cfg
```

12. (For multi-node Fusion, only run this step on the first node) Upgrade ZooKeeper:

```
cd fusion-upgrade-scripts-internal/src
python upgrade-to-3.0.py --upgrade zk
```

13. (On all nodes) Run Fusion 3.0.0 (start all services defined in `fusion.properties`):

```
cd $FUSION_HOME
./bin/fusion start
```

Ensure that all services have started on all nodes before proceeding.

14. (For multi-node Fusion, only run this step on the first node) Upgrade your custom dashboards, which are stored in the

Solr collection `system_banana` (after the upgrade, they are stored in the Solr collection `system_blobs`):

```
cd fusion-upgrade-scripts-internal/src
python upgrade-to-3.0.py --upgrade banana
```

15. (On all nodes) Use the Fusion UI to validate the deployment.

Tip

Before opening the Fusion UI, clear your browser's cache. Otherwise, you might inadvertently access a cached version of the old Fusion UI and see inconsistent behavior.

Windows

The steps in this section describe how to upgrade from Fusion 2.1.4 or 2.4.x to Fusion 3.0 on Windows.

Prerequisites

Ensure that these prerequisites are met before upgrading Fusion.

Note

If you have multi-node Fusion, perform these steps on all nodes.

1. Install python 2.7.13 from here: <https://www.python.org/ftp/python/2.7.13/python-2.7.13.amd64.msi>
2. Open a normal Command Prompt window and set a `PYTHON_HOME` environment variable to the Python installation directory you just created:

```
set PYTHON_HOME=C:\path\to\python\directory
```

Example:

```
set PYTHON_HOME=C:\Python27
```

Keep this command prompt window open.

3. Install the `kazoo` and `requests` python packages

```
"%PYTHON_HOME%\Scripts\pip" install kazoo
"%PYTHON_HOME%\Scripts\pip" install requests
```

Steps to upgrade from Fusion 2.4.x to 3.0.0 on Windows

These steps describe how to upgrade from Fusion 2.1.4 or 2.4.x to Fusion 3.0.0 on Windows.

Note	If you have a multi-node installation and unless otherwise stated, you must perform each step on each node as you go. For example, if you have 3 nodes, perform step 1 on node 1, node 2, and then node 3. Next, perform step 2 on node 1, node 2, and then node 3. Run steps that are only needed on a single node on your first node; not on the other nodes.
------	---

1. (On all nodes) Create a `fusion-new` folder somewhere on your computer, and download and extract [Fusion 3.0.0](#) to it.
2. (On all nodes) Set the `FUSION_HOME` environment variable to the full path of the `fusion\3.0.0` directory:

Example:

```
SET FUSION_HOME=C:\lucidworks\fusion-3.0.0\fusion\3.0.0
```

At this point, `%FUSION_HOME%\bin\fusion.cmd` should point to the Fusion 3.0.0 `fusion.cmd` file.

3. (On all nodes) Set the `FUSION_OLD_HOME` environment variable to the full path of the old Fusion home directory; for example:

```
SET FUSION_OLD_HOME=C:\lucidworks\fusion
```

At this point, `%FUSION_OLD_HOME%\bin\fusion.cmd` should point to the Fusion 2.x.y `fusion.cmd` file.

4. (On all nodes) Stop the old Fusion services:

```
cd "%FUSION_OLD_HOME%"  
"%FUSION_OLD_HOME%\bin\fusion" stop
```

5. (On all nodes) In Windows Explorer, copy the `%FUSION_OLD_HOME%\data` directory on top of the `%FUSION_HOME%\data` directory; respond with **Overwrite All** when prompted.
6. (On all nodes) Edit `%FUSION_HOME%\conf\fusion.properties` and apply all custom JVM settings that you had applied to your `%FUSION_OLD_HOME%\conf\config.cmd` and `%FUSION_OLD_HOME%\bin*.cmd` files into the new `jvmOptions` properties. If you have multi-node Fusion, make sure you uncomment `default.zk.connect` and set it to what you have `FUSION_ZK` set to.

For example, many customers customize the max heap size for `%FUSION_OLD_HOME%\bin\connectors.cmd`. Here is an example for someone who changed it to use a max heap size of 4 GB:

```
set "JAVA_OPTIONS=-Djava.net.preferIPv4Stack=true -Xmx4g -Dapple.awt.UIElement=true"
```

If this were the case, you would need to use this entry in `%FUSION_HOME%\conf\fusion.properties` for the `connectors.jvmOptions` variable:

```
connectors.jvmOptions=-Xmx4g -Xss256k -Dcom.lucidworks.connectors.pipelines.embedded=false
```

7. (On all nodes) Copy your ZooKeeper configuration file over from the old installation to the new one, that is, copy

`%FUSION_OLD_HOME%\conf\zookeeper\zoo.cfg` to `%FUSION_NEW_HOME%\conf\zookeeper\zoo.cfg`.

8. (On all nodes) Start the new ZooKeeper service:

```
"%FUSION_HOME%\bin\zookeeper" start
```

Ensure that ZooKeeper services have started on all nodes before proceeding.

9. (For multi-node Fusion, only run this step on the first node) Download the `fusion-upgrade-scripts` project from <https://github.com/lucidworks/fusion-upgrade-scripts>. Click **Download** > **Download Zip**.

Extract this zip file somewhere on your local computer, for example, to `C:\lucidworks`. When the zip file is extracted, you will have the directory `C:\lucidworks\fusion-upgrade-scripts-master`.

10. (For multi-node Fusion, only run this step on the first node) Run the ZooKeeper upgrade script:

```
cd "C:\lucidworks\fusion-upgrade-scripts-master\src"  
"%PYTHON_HOME%\python" upgrade-to-3.0.py --upgrade zk
```

Example output:

```
C:\lucidworks\fusion-upgrade-scripts-master\src>"%PYTHON_HOME%\python" upgrade-to-3.0.py --upgrade zk  
2017-10-04 08:45:25,517 - root - load_fusion_3x_config.py:36 - INFO - Creating config file using agent  
2017-10-04 08:45:27,563 - root - load_fusion_3x_config.py:45 - INFO - Generated config file at path  
'ui.config.json'  
2017-10-04 08:45:27,563 - root - zookeeper_client.py:12 - INFO - Starting zookeeper client  
2017-10-04 08:45:27,565 - kazoo.client - connection.py:599 - INFO - Connecting to 127.0.0.1:9983  
2017-10-04 08:45:27,596 - kazoo.client - client.py:465 - INFO - Zookeeper connection established, state:  
CONNECTED  
2017-10-04 08:45:27,612 - root - upgrade-to-3.0.py:64 - INFO - Migrating from fusion version '2.4.5' to  
'3.0.0'  
2017-10-04 08:45:27,612 - root - upgrade-to-3.0.py:67 - INFO - Copying znodes from old fusion paths to new  
paths  
2017-10-04 08:45:27,612 - root - znodes_migration.py:15 - INFO - Migrating Solr data to new ZK namespace  
/lwfusion/3.0.0/solr  
2017-10-04 08:45:47,562 - root - znodes_migration.py:21 - INFO - Migrating api data to new ZK namespace  
lwfusion/3.0.0/core  
2017-10-04 08:45:58,421 - root - znodes_migration.py:28 - INFO - Migrating proxy data to new ZK namespace  
lwfusion/3.0.0/proxy  
2017-10-04 08:45:59,338 - root - upgrade-to-3.0.py:69 - INFO - Migration from old znode paths to new paths  
complete  
2017-10-04 08:45:59,352 - root - api_pojo_migrator.py:16 - INFO - Updating search-cluster payload at path  
'lwfusion/3.0.0/core/search-clusters/default'  
2017-10-04 08:45:59,405 - root - proxy_pojo_migrator.py:15 - INFO - Updating init-meta payload at path  
'lwfusion/3.0.0/proxy/sys/init-meta'  
2017-10-04 08:45:59,490 - root - resource_manager.py:17 - INFO - Loading file from path  
C:\lucidworks\fusion-upgrade-scripts-master\src\utils\../resources/migrators.json  
2017-10-04 08:45:59,490 - root - connectors_migrator.py:126 - INFO - Trying to migrate datasource:  
test2352352, type: jdbc  
2017-10-04 08:45:59,490 - root - upgrade-to-3.0.py:78 - INFO - Performing splitter migrator  
2017-10-04 08:45:59,523 - kazoo.client - connection.py:566 - INFO - Closing connection to 127.0.0.1:9983  
2017-10-04 08:45:59,523 - kazoo.client - client.py:469 - INFO - Zookeeper session lost, state: CLOSED
```

11. (On all nodes) Run Fusion 3.0.0 (start all services defined in `fusion.properties`), and wait for all services to start:

```
"%FUSION_HOME%\bin\fusion" start
```

12. (For multi-node Fusion, only run this step on the first node) Upgrade your custom dashboards, which are stored in the Solr collection `system_banana` (after the upgrade, they are stored in the Solr collection `system_blobs`):

```
cd "C:\lucidworks\fusion-upgrade-scripts-master\src"  
"%PYTHON_HOME%\python" upgrade-to-3.0.py --upgrade banana
```

If you get an error containing `Unknown collection system_banana`, you can safely ignore it. That just means that you had no dashboards to migrate in the first place.

13. (On all nodes) Use the Fusion UI to validate the deployment.

Tip

Before opening the Fusion UI, clear your browser's cache. Otherwise, you might inadvertently access a cached version of the old Fusion UI and see inconsistent behavior.

Upgrade Fusion 2.1 or 2.2 to Fusion 2.4

This instruction set is valid for all Fusion 2.1 releases as well as Fusion 2.2.0.

Note	<p>Fusion 2.4 introduces changes to the configuration properties for some Fusion datasources. To update these configurations, we have provided a program which can be downloaded from: https://github.com/LucidWorks/fusion-upgrade-scripts.</p> <p>Once you have migrated all Fusion configurations from the current Fusion 2.1 ZooKeeper service to the new Fusion 2.4 ZooKeeper service, you must run this script against the new ZooKeeper service, see Migrate ZooKeeper data</p>
------	---

The upgrade process leaves the current Fusion deployment in place while a new Fusion deployment is installed and configured. All of the upgrade operations copy information from the current Fusion over to the new Fusion. This provides a rollback option should the upgrade procedure encounter problems.

The current Fusion configurations must remain as-is during the upgrade process. In order to capture indexing job history, no indexing jobs should be running. If the new Fusion installation is being installed onto the same server that the current Fusion installation is running on, you must either run only one version at a time or else change the Fusion component server ports so that all components are using unique ports for both the current and new versions.

Terminology

These instructions use the following names to refer to the directories involved in the upgrade procedure:

- FUSION_HOME: Absolute pathname to the top-level directory of the Fusion distribution
- FUSION-CURRENT: Name of the FUSION_HOME directory for the current Fusion version, e.g. "/opt/lucidworks/fusion-2.1.2"
- FUSION-NEW: Name of the directory of the upgrade Fusion distribution during the upgrade process, e.g. "/opt/lucidworks/fusion-2.4.1"
- INSTALL-DIR: Directory where the new Fusion version will be installed, e.g. "opt/lucidworks" All scripts and commands in the upgrade instruction set are carried out from this directory.
- FUSION-UPGRADE-SCRIPTS: Full path to the directory that contains the upgrade scripts from <https://github.com/LucidWorks/fusion-upgrade-scripts>.

Requirements

- File-system permissions: the user running the upgrade scripts and commands must have read/write/execute (rwx) permissions on directory INSTALL-DIR.
- Download but *do not* unpack a copy of the FUSION-NEW distribution. The compressed Fusion distribution requires approximately 1.7 GB disk space. All supported version are available from [Lucidworks Fusion Get Started](#) page.
- Disk space requirements: the INSTALL-DIR must be on a disk partition which has enough free space for the complete FUSION-NEW installation, that is, there must be at least as much free space as the size of the FUSION-CURRENT directory. On a Unix system, the following commands can be used:

- `du -sh fusion` - total size of FUSION-CURRENT.
- `df -kH` - amount of free space on all file-systems.
- Download a copy of the Fusion upgrade scripts from the GitHub repository <https://github.com/LucidWorks/fusion-upgrade-scripts>. These upgrade scripts run under Python 2.7. They have been tested with version 2.7.10. If this version of Python isn't available, you should use Python's `virtualenv`. If you don't have permissions to install packages, you can use `python` to install `virtualenv` and then from your `virtualenv` python environment, you can install your own versions of these packages.

Note	These scripts require the environment variable <code>FUSION_OLD_HOME</code> which should be set to the location of the current Fusion installation, i.e., the existing 1.2 or 2.1 install.
------	--

- Upgrades from 2.1 to 2.4 use the script `src/upgrade-ds-2.1-to-2.4.py`. This script requires the python package `kazoo` which is a ZooKeeper client.
- Upgrades from 1.2 to 2.4 use two scripts: `src/upgrade-ds-1.2-to-2.4.py` and `bin/download_upload_ds.py`. These scripts require python packages `kazoo` and `requests`, which is an HTTP request handler.

Procedure

Unpack FUSION-NEW

- **Current working directory must be `INSTALL-DIR`**

The commands in this section assume that your current working directory is `INSTALL-DIR` (e.g., `opt/lucidworks`), therefore `cd` to this directory before continuing.

- **Avoid directory name conflicts between `FUSION-CURRENT` and `FUSION-NEW`**

By default, the Fusion distribution unpacks into a directory named "fusion". If the `INSTALL-DIR` is the directory which contains the `FUSION-CURRENT` directory and if the `FUSION-CURRENT` directory is named "fusion", then you must create a new directory with a different name into which to unpack the Fusion distribution. For example, if your `INSTALL-DIR` is `/opt/lucidworks` and your `FUSION-CURRENT` directory is `/opt/lucidworks/fusion`, then you should create a directory named "fusion-new" and unpack the contents of the distribution here:

```
> mkdir fusion-new
> tar -C fusion-new --strip-components=1 -xf fusion-2.4.1.tar.gz
```

If you are working on a Windows machine, the zipfile unzips into a folder named "fusion-2.4.1" which contains a folder named "fusion". Rename folder "fusion" to "fusion-new" and move it into folder `INSTALL-DIR`.

Customize `FUSION-NEW` configuration files and run scripts

The Fusion run scripts in the `FUSION_HOME/bin` directory start and stop Fusion and its component services. The Fusion configuration files `FUSION_HOME/conf` define environment variables used by the Fusion run scripts. The configuration and run scripts for the `FUSION-NEW` installation must be edited by hand, you cannot copy over existing scripts from the current installation.

The Fusion configuration scripts **might** need to be updated if you have changed default settings. These scripts **will** need to be updated for deployments that:

- Use an external ZooKeeper cluster as Fusion's ZooKeeper service

- Use an external Solr cluster to manage Fusion’s system collections
- Run on non-standard ports
- Have been configured to run over SSL

To facilitate the task of identifying changes made to the current installation, the FUSION-UPGRADE-SCRIPTS repository contains a directory "reference-files" which contains copies of the contents of these directories for all Fusion releases. To identify changes, use the Unix `diff` command with the `-r` flag; e.g., if FUSION-CURRENT is 2.1.1, then these diff commands will report the set of changed files and the changes that were made:

```
> diff -r FUSION-CURRENT/bin FUSION-UPGRADE-SCRIPTS/reference-files/bin-2.1.1
> diff -r FUSION-CURRENT/conf FUSION-UPGRADE-SCRIPTS/reference-files/conf-2.1.1
```

A copy of Fusion is installed on every node in a Fusion deployment. Depending on the role that node plays in the deployment, the configuration settings and run scripts are customized accordingly. Therefore, if you are running a multi-node Fusion deployment this configuration step will be carried out for each node in the cluster.

Copy local data stores in directory `FUSION-CURRENT/data`

The directory `FUSION_HOME/data` contains the on-disk data stores managed directly or indirectly by Fusion services.

- `FUSION_HOME/data/connectors` contains data required by Fusion connectors.
 - `FUSION_HOME/data/connectors/lucid.jdbc` contains third-party JDBC driver files. If your application uses a JDBC connector, you must copy this information over to every server on which will this connector will run.
 - `FUSION_HOME/data/connectors/crawlDb` contains information on the filed visited during a crawl. (Preserving crawlDb history may not be possible if there are multiple different servers running Fusion connectors services.)
- `FUSION_HOME/data/nlp` contains data used by Fusion NLP pipeline stages. If you are using Fusion’s NLP components for sentence detection, part-of-speech tagging, and named entity detection, you must copy over the model files stored under this directory.
- `FUSION_HOME/data/solr` contains the backing store for Fusion’s embedded Solr (developer deployment only).
- `FUSION_HOME/data/zookeeper` contains the backing store for Fusion’s embedded ZooKeeper (developer deployment only).

If `FUSION_CURRENT` and `FUSION_NEW` are installed on the same server, you can copy a subset of these directories using the Unix "cp" command, e.g.:

```
> cp -R FUSION-CURRENT/data/connectors/lucid.jdbc FUSION-NEW/data/connectors
> cp -R FUSION-CURRENT/data/connectors/crawlDb FUSION-NEW/data/connectors
> cp -R FUSION-CURRENT/data/nlp FUSION-NEW/data/
```

If `FUSION_CURRENT` and `FUSION_NEW` are on different servers, use the Unix `rsync` utility.

Migrate ZooKeeper and Solr for single-node Fusion deployment

If you are running a single-node Fusion deployment and using both the embedded ZooKeeper and the embedded Solr that ships with this distribution, then you must copy over both the configurations and data.

To copy the ZooKeeper configuration:

```
> cp -R FUSION-CURRENT/data/zookeeper FUSION-NEW/data
```

To copy the Solr data:

```
> cp -R FUSION-CURRENT/data/solr FUSION-NEW/data
```

If the Solr collections are very large this may take a while.

Migrate Fusion configurations between ZooKeeper instances

Migration consists of the following steps:

- Copy the ZooKeeper data nodes which contain Fusion configuration information from the FUSION-CURRENT ZooKeeper instance to the FUSION-NEW ZooKeeper instance

Fusion's utility script `zkImportExport.sh` is used to copy ZooKeeper data between ZooKeeper clusters. This script is included with all Fusion distributions in the top-level directory named `scripts`.

- Rewrite Fusion datasource configurations

Fusion 2.4 changed and standardized the configuration properties used by several datasources. The public GitHub repository <https://github.com/LucidWorks/fusion-upgrade-scripts> contains a python script `src/upgrade-ds-2.1-to-2.4.py` which rewrites these properties.

Copying ZooKeeper data nodes

Note	This step is not necessary if you are doing an in-place upgrade of a single-node Fusion deployment; the copy command described in procedure single-node Fusion ZooKeeper data (above) is sufficient.
------	--

Fusion configurations are stored in Fusion's ZooKeeper instance under two top-level znodes:

- Node `lucid` stores all application-specific configurations, including collection, datasource, pipeline, signals, aggregations, and associated scheduling, jobs, and metrics.
- Node `lucid-apollo-admin` stores all access control information, including all users, groups, roles, and realms.

Fusion's utility script `zkImportExport.sh` is used to migrate ZooKeeper data between ZooKeeper clusters. Migrating configuration information from one deployment to another requires running this script twice:

- The first invocation runs the script in "export" mode, in order to get the set of configurations to be migrated as a JSON dump file.
- The second invocation runs the script in "import" or "update" mode, in order to sent this configuration set to the other Fusion deployment.

When running this script against a Fusion deployment, it is advisable to stop all Fusion services except for Fusion's ZooKeeper service.

Exporting Fusion configurations from FUSION-CURRENT ZooKeeper Service

The ZooKeeper service for FUSION-CURRENT must be running. Either stop all other Fusion services or otherwise ensure that no changes to Fusion configurations take place during this procedure. If you are upgrading from a Fusion 1.2 installation which uses Fusion's embedded Solr service and the ZooKeeper service included with that Solr installation, then starting just the Solr service will start the ZooKeeper service as well. If you are upgrading from a Fusion 2 installation, you can start just the ZooKeeper service via the script "zookeeper" in the `$FUSION_HOME/bin` directory.

The zkImportExport.sh script arguments are:

- `-cmd export` - This is the command parameter which specifies the mode in which to run this program.
- `-zkhost <FUSION_CURRENT ZK>` - The ZooKeeper connect string is the list of all servers,ports for the FUSION_CURRENT ZooKeeper cluster. For example, if running a single-node Fusion developer deployment with embedded ZooKeeper, the connect string is `localhost:9983`. If you have an external 3-node ZooKeeper cluster running on servers "zk1.acme.com", "zk2.acme.com", "zk3.acme.com", all listening on port 2181, then the connect string is `zk1.acme.com:2181,zk2.acme.com:2181,zk3.acme.com:2181`
- `-filename <path/to/JSON/dump/file>` - The name of the JSON dump file to save to.
- `-path <start znode>`
 - To migrate all ZooKeeper data, the path is `"/`.
 - To migrate only the Fusion services configurations, the path is `"/lucid"`. Migrating just the "lucid" node between the ZooKeeper services used by different Fusion deployments results in deployments which contain the same applications but not the same user databases.
 - To migrate the Fusion users, groups, roles, and realms information, the path is `"/lucid-apollo-admin"`.

Example of exporting Fusion configurations for znode `"/lucid"` from a local single-node ZooKeeper service:

```
> $FUSION_HOME/scripts/zkImportExport.sh -zkhost localhost:9983 -cmd export -path /lucid -filename znode_lucid_dump.json
```

Importing ZooKeeper data into FUSION-NEW

ZooKeeper service for FUSION-NEW must be running.

To import configurations, run the zkImportExport.sh script, this time with arguments:

- command; must be `import`
- ZooKeeper connect string for the FUSION-NEW Zookeeper cluster
- Location of JSON dump file.

This command will fail if the "lucid" znode in this Fusion installation contains configuration definitions which are in conflict with the exported data.

Example of importing exported data from previous step into FUSION_NEW ZooKeeper running on test server 'test.acme.com':

```
> $FUSION_HOME/scripts/zkImportExport.sh -zkhost test.acme.com:9983 -cmd import -filename znode_lucid_dump.json
```

Note that the above command will fail if there is conflict between existing znode structures or contents between the ZooKeeper service and the dump file.

Rewrite datasource configurations for Fusion 2.4

Once all Fusion configurations have been uploaded to the FUSION-NEW ZooKeeper service and while that service is running, you can run the Python programs [upgrade-ds-2.1-to-2.4.py](#) or [upgrade-ds-1.2-to-2.4.py](#) to update these configurations.

Note	<p>These programs require:</p> <ul style="list-style-type: none">• The environment variable "FUSION_HOME" must be set to the FUSION-NEW directory.• The environment variable "FUSION_OLD_HOME" must be set to the FUSION-CURRENT directory.• Python version 2.7, preferably version 2.7.10.• Package: kazoo - a ZooKeeper client <p>The Python virtualenv tool can be used to install the correct Python version and required package.</p>
------	---

Set environment variable "FUSION_HOME" to the full path of the FUSION-NEW directory, e.g.:

```
> export FUSION_HOME=/Users/demo/test_upgrade/fusion_2_4_1
```

Run this program with arguments: "--datasources all"

If your current Fusion version is 1.2, run:

```
> python upgrade-ds-1.2-to-2.4.py --datasources all
```

If your current Fusion is version 2, run:

```
> python upgrade-ds-2.1-to-2.4.py --datasources all
```

If a datasource wouldn't have a valid implementation, the application will print a log message on console and continue with the next datasource.

Troubleshooting the upgrade

- Clear your browser cache after starting the UI in the new Fusion instance
- The Fusion 2.4 [Index Pipeline Simulator](#) can be used to verify that the existing set of datasource configurations work as expected.

Upgrade Fusion 1.2 to Fusion 2.4

These instructions are valid for Fusion 1.2.3 releases through Fusion 1.2.8.

Note	<p>Several changes have been made to Fusion configurations stored in ZooKeeper:</p> <ul style="list-style-type: none">• Fusion 2.1 introduced enhanced security for Fusion datasource passwords which are stored in ZooKeeper as part of datasource and pipeline stage configuration properties.• Fusion 2.4 introduces changes to the configuration properties for some Fusion datasources. <p>To update these configurations, we have provided two python scripts which can be downloaded from: https://github.com/LucidWorks/fusion-upgrade-scripts.</p> <p>Once you have migrated all Fusion configurations from the current Fusion 1.2.x ZooKeeper service to the new Fusion 2.4 ZooKeeper service, you must run both of these scripts against the new ZooKeeper service. This procedure is covered in detail in section Migrate ZooKeeper data</p>
------	--

The upgrade process leaves the current Fusion deployment in place while a new Fusion deployment is installed and configured. All of the upgrade operations copy information from the current Fusion over to the new Fusion. This provides a rollback option should the upgrade procedure encounter problems.

The current Fusion configurations must remain as-is during the upgrade process. In order to capture indexing job history, no indexing jobs should be running. If the new Fusion installation is being installed onto the same server that the current Fusion installation is running on, you must either run only one version at a time or else change the Fusion component server ports so that all components are using unique ports for both the current and new versions.

Terminology

These instructions use the following names to refer to the directories involved in the upgrade procedure:

- FUSION_HOME: Absolute pathname to the top-level directory of the Fusion distribution
- FUSION-CURRENT: Name of the FUSION_HOME directory for the current Fusion version, e.g. "/opt/lucidworks/fusion-2.1.2"
- FUSION-NEW: Name of the directory of the upgrade Fusion distribution during the upgrade process, e.g. "/opt/lucidworks/fusion-2.4.1"
- INSTALL-DIR: Directory where the new Fusion version will be installed, e.g. "opt/lucidworks" All scripts and commands in the upgrade instruction set are carried out from this directory.
- FUSION-UPGRADE-SCRIPTS: Full path to the directory that contains the upgrade scripts from <https://github.com/LucidWorks/fusion-upgrade-scripts>.

Requirements

- File-system permissions: the user running the upgrade scripts and commands must have read/write/execute (rwx) permissions on directory INSTALL-DIR.
- Download but *do not* unpack a copy of the FUSION-NEW distribution. The compressed Fusion distribution requires approximately 1.7 GB disk space. All supported version are available from [Lucidworks Fusion Get Started](#) page.
- Disk space requirements: the INSTALL-DIR must be on a disk partition which has enough free space for the complete FUSION-NEW installation, that is, there must be at least as much free space as the size of the FUSION-CURRENT directory. On a Unix system, the following commands can be used:
 - `du -sh fusion` - total size of FUSION-CURRENT.
 - `df -kH` - amount of free space on all file-systems.
- Download a copy of the Fusion upgrade scripts from the GitHub repository <https://github.com/LucidWorks/fusion-upgrade-scripts>. These upgrade scripts run under Python 2.7. They have been tested with version 2.7.10. If this version of Python isn't available, you should use Python's `virtualenv`. If you don't have permissions to install packages, you can use `python` to install `virtualenv` and then from your `virtualenv` python environment, you can install your own versions of these packages.

Note

These scripts require the environment variable `FUSION_OLD_HOME` which should be set to the location of the current Fusion installation, i.e., the existing 1.2 or 2.1 install.

- Upgrades from 2.1 to 2.4 use the script `src/upgrade-ds-2.1-to-2.4.py`. This script requires the python package `kazoo` which is a ZooKeeper client.
- Upgrades from 1.2 to 2.4 use two scripts: `src/upgrade-ds-1.2-to-2.4.py` and `bin/download_upload_ds.py`. These scripts require python packages `kazoo` and `requests`, which is an HTTP request handler.

Procedure

Unpack FUSION-NEW

- **Current working directory must be INSTALL-DIR**
The commands in this section assume that your current working directory is INSTALL-DIR (e.g., "/opt/lucidworks"), therefore `cd` to this directory before continuing.
- **Avoid directory name conflicts between FUSION-CURRENT and FUSION-NEW**
By default, the Fusion distribution unpacks into a directory named "fusion". If the INSTALL-DIR is the directory which contains the FUSION-CURRENT directory and if the FUSION-CURRENT directory is named "fusion", then you must create a new directory with a different name into which to unpack the Fusion distribution. For example, if your INSTALL-DIR is "/opt/lucidworks" and your FUSION-CURRENT directory is "/opt/lucidworks/fusion", then you should create a directory named "fusion-new" and unpack the contents of the distribution here:

```
> mkdir fusion-new
> tar -C fusion-new --strip-components=1 -xf fusion-2.4.1.tar.gz
```

If you are working on a Windows machine, the zipfile unzips into a folder named "fusion-2.4.1" which contains a folder named "fusion". Rename folder "fusion" to "fusion-new" and move it into folder INSTALL-DIR.

Customize FUSION-NEW configuration files and run scripts

The Fusion run scripts in the `FUSION_HOME/bin` directory start and stop Fusion and its component services. The Fusion configuration files `FUSION_HOME/conf` define environment variables used by the Fusion run scripts. The configuration and run scripts for the FUSION-NEW installation must be edited by hand, you cannot copy over existing scripts from the current installation.

The Fusion configuration scripts **might** need to be updated if you have changed default settings. These scripts **will** need to be updated for deployments that:

- Use an external ZooKeeper cluster as Fusion's ZooKeeper service
- Use an external Solr cluster to manage Fusion's system collections
- Run on non-standard ports
- Have been configured to run over SSL

To facilitate the task of identifying changes made to the current installation, the FUSION-UPGRADE-SCRIPTS repository contains a directory "reference-files" which contains copies of the contents of these directories for all Fusion releases. To identify changes, use the Unix `diff` command with the `-r` flag; e.g., if FUSION-CURRENT is 2.1.1, then these diff commands will report the set of changed files and the changes that were made:

```
> diff -r FUSION-CURRENT/bin FUSION-UPGRADE-SCRIPTS/reference-files/bin-2.1.1
> diff -r FUSION-CURRENT/conf FUSION-UPGRADE-SCRIPTS/reference-files/conf-2.1.1
```

A copy of Fusion is installed on every node in a Fusion deployment. Depending on the role that node plays in the deployment, the configuration settings and run scripts are customized accordingly. Therefore, if you are running a multi-node Fusion deployment this configuration step will be carried out for each node in the cluster.

In Fusion 1.2, the `FUSION_HOME/bin` directory contains both the Fusion run scripts and the helper scripts which define common settings and environment variables. In Fusion 2.1, the configuration files `config.sh` and `config.cmd` have been moved to directory `FUSION_HOME/conf`.

Checking a 1.2 installation against the reference scripts for that release requires only a single diff command:

```
> diff -r FUSION-CURRENT/bin FUSION-UPGRADE-SCRIPTS/reference-files/bin-1.2.3
```

If either the "config.sh" or "config.cmd" files have changed, the corresponding files for the Fusion 2 release will be in directory `FUSION_HOME/conf`.

Copy local data stores in the directory FUSION-CURRENT/data

The directory `FUSION_HOME/data` contains the on-disk data stores managed directly or indirectly by Fusion services.

- `FUSION_HOME/data/connectors` contains data required by Fusion connectors.
 - `FUSION_HOME/data/connectors/lucid.jdbc` contains third-party JDBC driver files. If your application uses a JDBC connector, you must copy this information over to every server on which will this connector will run.
 - `FUSION_HOME/data/connectors/crawlddb` contains information on the files visited during a crawl. (Preserving crawlddb history may not be possible if there are multiple different servers running Fusion connectors services.)
- `FUSION_HOME/data/nlp` contains data used by Fusion NLP pipeline stages. If you are using Fusion's NLP components

for sentence detection, part-of-speech tagging, and named entity detection, you must copy over the model files stored under this directory.

- `FUSION_HOME/data/solr` contains the backing store for Fusion's embedded Solr (developer deployment only).
- `FUSION_HOME/data/zookeeper` contains the backing store for Fusion's embedded ZooKeeper (developer deployment only).

If `FUSION_CURRENT` and `FUSION_NEW` are installed on the same server, you can copy a subset of these directories using the Unix "cp" command, e.g.:

```
> cp -R FUSION-CURRENT/data/connectors/lucid.jdbc FUSION-NEW/data/connectors
> cp -R FUSION-CURRENT/data/connectors/crawldb FUSION-NEW/data/connectors
> cp -R FUSION-CURRENT/data/nlp FUSION-NEW/data/
```

If `FUSION_CURRENT` and `FUSION_NEW` are on different servers, use the Unix `rsync` utility.

Migrate ZooKeeper and Solr for single-node Fusion deployment

If you are running a single-node Fusion deployment and using both the embedded ZooKeeper and the embedded Solr that ships with this distribution, then you must copy over both the configurations and data.

To copy the ZooKeeper configuration:

```
> mkdir -p FUSION-NEW/data/zookeeper
> cp -R FUSION-CURRENT/solr/zoo_data/* FUSION-NEW/data/zookeeper
```

To check your work: compare the directories `FUSION-CURRENT/solr/zoo_data/` and `FUSION-NEW/data/zookeeper` using the `diff` command. This command succeeds silently when the contents are the same.

```
> diff -r FUSION-CURRENT/solr/zoo_data FUSION-NEW/data/zookeeper
```

To copy the Solr data:

```
> find FUSION-CURRENT/solr -maxdepth 1 -mindepth 1 | grep -v -E "zoo*" | while read f ; do cp -R $f FUSION-NEW/data/solr/; done
```

If the Solr collections are very large this may take a while.

You can use the `diff` command to check your work. The copy command excluded ZooKeeper config data, therefore you should see the following output:

```
> diff -r FUSION-CURRENT/solr FUSION-NEW/data/solr
Only in FUSION-NEW/data/solr: configsets
Only in FUSION-CURRENT/solr: zoo.cfg
Only in FUSION-CURRENT/solr: zoo_data
```

Migrate Fusion configurations between ZooKeeper instances

Migration consists of three steps:

- Copy the ZooKeeper data nodes which contain Fusion configuration information from the FUSION-CURRENT ZooKeeper instance to the FUSION-NEW ZooKeeper instance

Fusion’s utility script `zkImportExport.sh` is used to copy ZooKeeper data between ZooKeeper clusters. This script is included with all Fusion distributions in the top-level directory named `scripts`.

- Rewrite Fusion datasource configurations

Fusion 2.4 changed and standardized the configuration properties used by several datasources. The public GitHub repository <https://github.com/LucidWorks/fusion-upgrade-scripts> contains a python script `src/upgrade-ds-1.2-to-2.4.py` which rewrites these properties.

- Rewrite stored password information used by Fusion datasources and pipelines.

Fusion 2 encrypts all passwords use by datasources and pipelines to access password-protected data repositories. The public GitHub repository <https://github.com/LucidWorks/fusion-upgrade-scripts>. contains a a python script `bin/download_upload_ds_pipelines.py` used to edit the stored password information.

Copying ZooKeeper data nodes

Note	This step is not necessary if you are doing an in-place upgrade of a single-node Fusion deployment; the copy command described in procedure single-node Fusion ZooKeeper data (above) is sufficient.
------	--

Fusion configurations are stored in Fusion’s ZooKeeper instance under two top-level znodes:

- Node `lucid` stores all application-specific configurations, including collection, datasource, pipeline, signals, aggregations, and associated scheduling, jobs, and metrics.
- Node `lucid-apollo-admin` stores all access control information, including all users, groups, roles, and realms.

Fusion’s utility script `zkImportExport.sh` is used to migrate ZooKeeper data between ZooKeeper clusters. Migrating configuration information from one deployment to another requires running this script twice:

- The first invocation runs the script in "export" mode, in order to get the set of configurations to be migrated as a JSON dump file.
- The second invocation runs the script in "import" or "update" mode, in order to sent this configuration set to the other Fusion deployment.

When running this script against a Fusion deployment, it is advisable to stop all Fusion services except for Fusion’s ZooKeeper service.

Exporting Fusion configurations from FUSION-CURRENT ZooKeeper Service

The ZooKeeper service for FUSION-CURRENT must be running. Either stop all other Fusion services or otherwise ensure that no changes to Fusion configurations take place during this procedure. If you are upgrading from a Fusion 1.2 installation which uses Fusion’s embedded Solr service and the ZooKeeper service included with that Solr installation, then starting just the Solr service will start the ZooKeeper service as well. If you are upgrading from a Fusion 2 installation, you can start just the ZooKeeper service via the script "zookeeper" in the `$FUSION_HOME/bin` directory.

The `zkImportExport.sh` script arguments are:

- `-cmd export` - This is the command parameter which specifies the mode in which to run this program.
- `-zkhost <FUSION_CURRENT_ZK>` - The ZooKeeper connect string is the list of all servers,ports for the FUSION_CURRENT ZooKeeper cluster. For example, if running a single-node Fusion developer deployment with embedded ZooKeeper, the connect string is `localhost:9983`. If you have an external 3-node ZooKeeper cluster running on servers "zk1.acme.com", "zk2.acme.com", "zk3.acme.com", all listening on port 2181, then the connect string is `zk1.acme.com:2181,zk2.acme.com:2181,zk3.acme.com:2181`
- `-filename <path/to/JSON/dump/file>` - The name of the JSON dump file to save to.
- `-path <start znode>`
 - To migrate all ZooKeeper data, the path is `"/`.
 - To migrate only the Fusion services configurations, the path is `"/lucid`". Migrating just the "lucid" node between the ZooKeeper services used by different Fusion deployments results in deployments which contain the same applications but not the same user databases.
 - To migrate the Fusion users, groups, roles, and realms information, the path is `"/lucid-apollo-admin`".

Example of exporting Fusion configurations for znode `"/lucid`" from a local single-node ZooKeeper service:

```
> $FUSION_HOME/scripts/zkImportExport.sh -zkhost localhost:9983 -cmd export -path /lucid -filename
znode_lucid_dump.json
```

Importing ZooKeeper data into FUSION-NEW

ZooKeeper service for FUSION-NEW must be running.

To import configurations, run the `zkImportExport.sh` script, this time with arguments:

- command; must be `import`
- ZooKeeper connect string for the FUSION-NEW Zookeeper cluster
- Location of JSON dump file.

This command will fail if the "lucid" znode in this Fusion installation contains configuration definitions which are in conflict with the exported data.

Example of importing exported data from previous step into FUSION_NEW ZooKeeper running on test server 'test.acme.com':

```
> $FUSION_HOME/scripts/zkImportExport.sh -zkhost test.acme.com:9983 -cmd import -filename
znode_lucid_dump.json
```

Note that the above command will fail if there is conflict between existing znode structures or contents between the ZooKeeper service and the dump file.

Rewrite datasource configurations for Fusion 2.4

Once all Fusion configurations have been uploaded to the FUSION-NEW ZooKeeper service and while that service is running, you can run the Python programs [upgrade-ds-2.1-to-2.4.py](#) or [upgrade-ds-1.2-to-2.4.py](#) to update these configurations.

Note	<p>These programs require:</p> <ul style="list-style-type: none">• The environment variable "FUSION_HOME" must be set to the FUSION-NEW directory.• The environment variable "FUSION_OLD_HOME" must be set to the FUSION-CURRENT directory.• Python version 2.7, preferably version 2.7.10.• Package: kazoo - a ZooKeeper client <p>The Python virtualenv tool can be used to install the correct Python version and required package.</p>
------	---

Set environment variable "FUSION_HOME" to the full path of the FUSION-NEW directory, e.g.:

```
> export FUSION_HOME=/Users/demo/test_upgrade/fusion_2_4_1
```

Run this program with arguments: "--datasources all"

If your current Fusion version is 1.2, run:

```
> python upgrade-ds-1.2-to-2.4.py --datasources all
```

If your current Fusion is version 2, run:

```
> python upgrade-ds-2.1-to-2.4.py --datasources all
```

If a datasource wouldn't have a valid implementation, the application will print a log message on console and continue with the next datasource.

Rewrite stored password information used by Fusion datasources and pipelines

Once you have migrated all Fusion configurations to the FUSION_NEW ZooKeeper service, you must update the migrated datasource configurations by running the script [download_upload_ds_pipelines.py](#) against the FUSION_NEW zookeeper in order to rewrite any stored datasource passwords that are specified as part of the configuration for a datasource or pipeline.

Note	<p>The script <code>bin/download_upload_ds_pipelines.py</code> requires:</p> <ul style="list-style-type: none">• Python version 2.7, preferably version 2.7.10.• Package: <code>kazoo</code> - a ZooKeeper client• Package: <code>requests</code> - an HTTP request handler• Environment variable <code>FUSION_OLD_HOME</code> set to location of Fusion 1.2 home. <p>The Python <code>virtualenv</code> tool can be used to install the correct Python version and required packages.</p>
------	---

The rewrite process consists of a download step which exports the ZooKeeper configuration information and an upload step which rewrites the information and then imports it back into ZooKeeper.

This script uses the following arguments and values:

- "--zk-connect": the ZooKeeper server:port for FUSION-NEW
- "--action": either "download" or "upload".
- "--fusion-url": URL of Fusion API service to upload configurations to
- "--fusion-username": name of Fusion user with admin privileges; the script will prompt for username's password.

Download configurations from ZooKeeper

No services for FUSION-NEW should be running, except for ZooKeeper. If your Fusion installation uses an external ZooKeeper, then this must be running. If your Fusion installation uses an embedded ZooKeeper, then you must have copied the ZooKeeper data from FUSION-CURRENT to FUSION-NEW.

Start the ZooKeeper service:

```
> FUSION-NEW/bin/zookeeper start
```

Run the script to download the configurations.

```
> python FUSION-UPGRADE-SCRIPTS/bin/download_upload_ds_pipelines.py \  
--zk-connect localhost:9983 --action download
```

To check your work, check that directory "fusion_upgrade_2.1" was created and that it contains definitions for all datasources and pipelines. Do not remove this directory until you have successfully completed the upload step.

If you are running embedded ZooKeeper, shut it down again:

```
> FUSION-NEW/bin/zookeeper stop
```

Upload configurations to the Fusion API service

Start FUSION-NEW:

```
> FUSION-NEW/bin/fusion start
```

Once it is running, run the script in upload mode to propagate the configurations in directory "fusion_upgrade_2.1".

At this point in the migration process, the FUSION-NEW ZooKeeper information is the same as the FUSION-CURRENT Zookeeper information; therefore the password for the admin user is the same.

To upload data to the Fusion API services, you must supply the admin username and password as arguments to the script:

- "--fusion-username": name of Fusion user with admin privileges
- "--fusion-password": password for Fusion user

```
> FUSION-NEW/bin/fusion start
> python FUSION-UPGRADE-SCRIPTS/bin/download_upload_ds_pipelines.py \
--zk-connect localhost:9983 --action upload --fusion-url http://localhost:8764/api \
--fusion-username <admin>
```

Copy and convert the crawldb

The Fusion "crawldb" records the results of running datasource jobs. This information must be copied from FUSION-CURRENT to FUSION-NEW and the data format must be converted to the format used in Fusion 2.1 via the conversion utility [com.lucidworks.fusion-crawldb-migrator-0.1.0.jar](#).

Copy the Fusion "crawldb" directory:

```
> cp -R FUSION-CURRENT/data/connectors/crawldb FUSION-NEW/data/connectors/
```

The crawldb data format changed in Fusion 2.1, therefore to upgrade to 2.1, the crawldb data must be converted with the the conversion utility [com.lucidworks.fusion-crawldb-migrator-0.1.0.jar](#).

The [anda-v1-to-v2](#) command allows Fusion 1.2.x connector DBs to be updated to the new v2.x format. It requires:

- A Fusion pre 2.1 installation (FUSION-CURRENT)
- A Fusion 2.1 or later installation (FUSION-NEW).
 - All FUSION-CURRENT datasource configurations must have been migrated to FUSION-NEW (see Migrate ZooKeeper data)
 - All FUSION-CURRENT crawldb files must have been copied over to the FUSION-NEW deployment.

If the FUSION-NEW installation is not currently running, start it:

```
> FUSION-NEW/bin/fusion start
```

The [anda-v1-to-v2](#) takes the following arguments:

- path-to-FUSION-CURRENT
- path-to-FUSION-NEW

- the `-z` flag specifies the ZooKeeper server:port for FUSION-NEW

The command to run this utility from the INSTALL-DIR is:

```
> java -jar FUSION-UPGRADE-SCRIPTS/bin/com.lucidworks.fusion-crawldb-migrator-0.1.0.jar anda-v1-v2 fusion
fusion-new -z localhost:9983
```

Once the task successfully completes, **the last few lines of logging show the output directory of the new DB files**. The output must be copied over to FUSION-NEW. To do this, remove the existing `lucid.anda` db directories, then copy the new `lucid.anda` directories generated from this utility into that same location:

```
> rm -Rf FUSION-NEW/data/connectors/crawldb/lucid.anda/*
> mv ${path-printed-from-command-output} FUSION-NEW/data/connectors/crawldb/lucid.anda/
```

This completes the upgrade process.

Troubleshooting the upgrade

- Clear your browser cache after starting the UI in the new Fusion instance
- The Fusion 2.4 [Index Pipeline Simulator](#) can be used to verify that the existing set of datasource configurations work as expected.

1.4.7. bin/fusion

For every server in a Fusion deployment, the script `fusion/3.1.x/bin/fusion` is used to start, stop, and check the status of the Fusion instance running on that server.

Fusion Agent Process

The Fusion agent process makes sure that all Fusion processes start up and shut down correctly. It prevents problems that can arise by trying to start Fusion on a server where it is already running.

Start Fusion

Run the script `fusion/3.1.x/bin/fusion` with the argument `start`:

```
$ cd /path/to/fusion/3.1.x
$ ./bin/fusion start
Starting zookeeper.
Successfully started zookeeper on port 9983 (process ID 77295)
Starting solr.
Successfully started solr on port 8983 (process ID 77297)
Starting api.....
Successfully started api on port 8765 (process ID 77301)
Starting connectors.....
Successfully started connectors on port 8984 (process ID 77388)
Starting ui....
Successfully started ui on port 8764 (process ID 77469)
```

Check the status of Fusion

Run the script `fusion/3.1.x/bin/fusion` with the argument `status`:

```
$ cd /path/to/fusion/3.1.x
$ ./bin/fusion status
zookeeper is running on port 9983 (process ID 77295)
solr is running on port 8983 (process ID 77297)
api is running on port 8765 (process ID 77301)
ui is running on port 8764 (process ID 77469)
connectors is running on port 8984 (process ID 77388)
```

Stop Fusion

Run the script `fusion/3.1.x/bin/fusion` with the argument `stop`:

```
$ cd /path/to/fusion/3.1.x
$ ./bin/fusion stop
Successfully stopped ui (process ID 41524)
Successfully stopped connectors (process ID 41328)
Successfully stopped api (process ID 41159)
Successfully stopped solr (process ID 41153)
Successfully stopped zookeeper (process ID 41151)
```

Troubleshooting

The [Java Virtual Machine Process Status Tool](#) utility at `/usr/bin/jps` is useful for reporting on all Fusion processes reported by script `fusion/3.1.x/bin/fusion`:

```
$ jps
77294 AgentMain
77295 zookeeper-path-1475182112123.jar
77297 start.jar
77301 start.jar
77388 start.jar
77469 start.jar
79455 Jps
```

The process `zookeeper-path-1475182112123.jar` is the ZooKeeper process used by Fusion. The 4 `start.jar` processes are Fusion's Solr, API Services, Connectors, and UI.

If the `path/to/fusion/3.1.x/bin/fusion` script doesn't run, or if it fails to start all services, see the Troubleshooting topic or the [knowledge base](#) for help.

1.4.8. Directories and Logs

Directories

The directory where the Fusion files go for a specific version of Fusion is the *Fusion home directory*. The Fusion home directory is a version-numbered directory (for example, `3.1.0`) below the directory `fusion`. This installation strategy lets you install multiple versions of Fusion and switch between them. The Fusion home directory is the directory `fusion`.

The directories found in the Fusion home directory `fusion/3.1.x` are:

Name	Description
<code>apps</code>	Fusion components 3rd-party distributions used by Fusion, including jar files and plugins
<code>bin</code>	Master script to run Fusion, and per-component run scripts
<code>conf</code>	Configuration files for Fusion and ZooKeeper that contain parameters settings tuned for common use cases
<code>data</code>	Default location of data stores used by Fusion apps
<code>docs</code>	License information
<code>examples</code>	Fusion signals example
<code>init</code>	<code>systemd</code> and <code>upstart</code> scripts and configurations for Linux
<code>scripts</code>	Developer utilities, including diagnostic scripts, for Linux and Windows. See <code>scripts/diag/linux/README</code> and <code>scripts/diag/win64/README.txt</code> for details.
<code>var</code>	Logfiles and system files created by Fusion components, as well as <code>.pid</code> files for each running process

Symbolic Links on UNIX

To simplify access to the latest version of Fusion and to files in the `bin`, `conf`, and `var` directories, Fusion creates a symbolic link `latest` to the latest version and symbolic links `bin`, `conf`, and `var` to `latest/bin`, `latest/conf`, and `latest/var` respectively.

For example, if `latest` is `3.1.0`, then instead of entering this command to change to the `bin` directory:

```
$ cd /path/to/fusion/3.1.0/bin
```

You could just type:

```
$ cd /path/to/fusion/bin
```

To avoid possible confusion in the documentation, we spell out the path below the `fusion` directory.

From the `fusion` directory, you can view the symbolic links by typing:

```
$ ls -l . | grep "\->"
```

To change the version of Fusion to which the symbolic links refer, unlink and relink `latest`, for example:

```
$ cd /path/to/fusion
$ unlink latest
$ ln -s 3.1.2 latest
```

Logfiles

Logfiles are found in directories under `fusion/3.1.x/var/log`. Because the Fusion components run in separate JVMs, each component has its own set of logfiles and files that monitor all garbage-collection events for that process.

Name	Description
<code>api</code>	Fusion REST API services logging and error messages. This log shows the result of service requests submitted to the REST API directly via HTTP and indirectly via the Fusion UI.
<code>connectors</code>	Fusion connector services logging and error messages. Fusion index pipeline logging stages write to this file.
<code>solr</code>	Messages from Solr
<code>spark-master</code>	Spark-master logs
<code>spark-worker</code>	Spark-worker logs
<code>ui</code>	Fusion UI messages
<code>zookeeper</code>	ZooKeeper messages

Every component logs all messages to a logfile named `<component>.log`. For example, the full path to the logfile for the connectors services is:

```
fusion/3.1.x/var/log/connectors/connectors.log
```

In addition to component logfiles, every component maintains a set of garbage-collection logfiles which are used for resource tuning. The garbage-collection logfiles are named `gc_<YYYYMMDD>_<PID>.log.<CT>`. In addition, the current garbage-collection logfile has suffix `.current`.

The Fusion REST API, UI, connectors services, and Solr all run inside a Jetty server. The Jetty server logs are also written to that component's logfile directory. The Jetty server logs are named:

- `jetty-YYYY_MM_DD.request.log`
- `jetty-YYYY_MM_DD.stderrout.log`

Fusion uses the [Apache Log4j 2](#) logging framework with Jetty to log each of the Fusion components. Logging is configured via an xml configuration file named `log4j2.xml`. Log levels, frequencies, and log rotation policy can be configured by changing these configuration files:

API service	<code>fusion/3.1.x/conf/api-log4j2.xml</code>
-------------	---

connectors	<code>fusion/3.1.x/conf/connectors-log4j2.xml</code>
Solr	<code>fusion/3.1.x/conf/solr-log4j2.xml</code>
Spark	<code>fusion/3.1.x/conf/spark-driver-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-master-agent-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-master-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-worker-agent-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-worker-log4j2.xml</code>
UI	<code>fusion/3.1.x/conf/ui-log4j2.xml</code>
ZooKeeper	<code>fusion/3.1.x/conf/zk-log4j2.xml</code>

The [Log4j2 Configuration](#) guide provides documentation and examples of all logging configuration options.

1.4.9. Configuration Files

Fusion configuration files are stored in `fusion/3.1.x/conf`. The complete contents of this directory are as follows:

<code>fusion.properties</code>	Fusion's main configuration file, which defines the common environment variables used by the Fusion run scripts.
<code>hive-site.xml</code>	Configuration for Fusion's Serializer/Deserializer (SerDe) for Hive.
<code>zookeeper/commons-logging.properties</code> <code>zookeeper/zoo.cfg</code>	ZooKeeper configuration files.
<code>agent-log4j2.xml</code> <code>api-log4j2.xml</code> <code>connectors-log4j2.xml</code> <code>solr-log4j2.xml</code> <code>spark-driver-log4j2.xml</code> <code>spark-master-agent-log4j2.xml</code> <code>spark-master-log4j2.xml</code> <code>spark-worker-agent-log4j2.xml</code> <code>spark-worker-log4j2.xml</code> <code>sql-agent-log4j2.xml</code> <code>sql-log4j2.xml</code> <code>ui-log4j2.xml</code> <code>zk-log4j2.xml</code> <code>zookeeper/log4j2.xml</code> <code>zookeeper/log4j.properties</code>	Logging configuration files. Fusion uses the Apache Log4j 2 logging framework with Jetty. Log levels, frequencies, and log rotation policy can be configured by changing these configuration files. See the Log4j2 Configuration guide.

Logging configuration files

Fusion uses the [Apache Log4j 2](#) logging framework with Jetty to log each of the Fusion components. Logging is configured via an xml configuration file named `log4j2.xml`. Log levels, frequencies, and log rotation policy can be configured by changing these configuration files:

API service	<code>fusion/3.1.x/conf/api-log4j2.xml</code>
connectors	<code>fusion/3.1.x/conf/connectors-log4j2.xml</code>
Solr	<code>fusion/3.1.x/conf/solr-log4j2.xml</code>

Spark	<code>fusion/3.1.x/conf/spark-driver-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-master-agent-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-master-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-worker-agent-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-worker-log4j2.xml</code>
UI	<code>fusion/3.1.x/conf/ui-log4j2.xml</code>
ZooKeeper	<code>fusion/3.1.x/conf/zk-log4j2.xml</code>

The [Log4j2 Configuration](#) guide provides documentation and examples of all logging configuration options.

1.4.10. Default Ports

Fusion services run in their own JVM and listen for requests on a number of ports. Environment variables, set in a common configuration file, are used to specify the port a service uses. To change the port(s) a service uses, you must change the settings in the configuration file.

Default Ports

The default ports for the Fusion services are as follows:

Port	Service
8764	Fusion UI This service includes the Fusion Authorization Proxy
8765	Fusion API Services
8766	Spark Master
8769	Spark Worker
8984	Connectors Services
8983	Solr This is the embedded Solr instance included in the Fusion distribution.
9983	ZooKeeper The embedded ZooKeeper used by Fusion services. It corresponds to the ZooKeeper <code>clientPort</code> which is defined in file <code>fusion/3.1.x/conf/zookeeper/zoo.cfg</code> .
8766	Apache Spark master REST port
8767	Apache Spark master UI
8082 and 8770	Apache Spark worker UI
4040	Apache Spark driver UI
8769	Apache Spark worker listening port
7337	Shuffle port for Apache Spark worker
8600-8616	Akka ports used between Spark driver, master, workers and API See aka documentation
47100-48099	Apache Ignite TCP communication port range (used by API, Connectors and UI Proxy)

Port	Service
48100-48199	Apache Ignite shared memory port range (used by API, Connectors and UI Proxy)
49200-49299	Apache Ignite discovery port range (used by API, Connectors and UI Proxy)
51500-52000	Executor port, driver port, block manager port, file server port You will need to set these manually in config if needed. Otherwise you can ignore these.

Important	In a production environment, do not expose port 8765 to users. Using your firewall software or the Jetty configuration of the API server, make it accessible only to the auth proxy service and the connectors service.
-----------	---

Port settings are defined in the `fusion.properties` file.

Jetty is used to run Solr, the Fusion UI, API, and connectors services. For each of these services, Jetty runs the service on the assigned port and listens on a second port for shutdown requests. Therefore, `fusion.properties` defines pairs of ports for components running on Jetty, e.g.:

```
api.port = 8765
api.stopPort = 7765
```

ZooKeeper Port Configuration

The ZooKeeper ports are defined both in the `fusion.properties` file and in the zookeeper configuration file, `zoo.cfg`, in the zookeeper subdirectory, path `fusion/3.1.x/conf/zookeeper/zoo.cfg`.

The definition in

`fusion.properties` is:

```
zookeeper.port = 9983
```

The definition in `zoo.cfg` is:

```
clientPort=9983
```

Important	If you change the zookeeper port and are running the embedded zookeeper, the port definitions must match!
-----------	---

1.4.11. Checking System State

As described in the section Default Ports, Fusion runs several components as separate JVMs running on different ports. Each of the components is capable of reporting its status. The proxy component reports status for all of the other components.

Full System Check

To see if each component has been started, a simple API call to the proxy (running on port 8764 by default) will return the status of each component of the system.

```
curl http://localhost:8764/api
```

The response should look similar to the following. If 'ping' is true for each service, all of the system components are running.

```
{
  "version": "0.9.0-SNAPSHOT-jenkins.build.105+git.sha.b425e2a",
  "enabledRealms": [
    "native"
  ],
  "initMeta": {
    "version": "0.9.0-SNAPSHOT-jenkins.build.105+git.sha.b425e2a",
    "initializedAt": "2014-10-06T17:43:31Z",
    "createdAt": "2014-10-06T17:43:31Z"
  },
  "startTime": "2014-10-06T18:38:08Z",
  "status": {
    "connectors": {
      "ping": true
    },
    "apollo": {
      "ping": true
    },
    "apolloZk": {
      "ping": true
    },
    "db": {
      "ping": true
    }
  }
}
```

Solr Health Check

The Fusion UI and API services are not accessible if ZooKeeper and Solr are not in healthy state. A Solr health check can be performed with a ping request.

```
curl http://localhost:8983/solr/admin/ping
```

The response will be a standard Solr XML response, similar to the following.

```
<?xml version="1.0" encoding="UTF-8"?>
<response>
  <lst name="responseHeader">
    <int name="status">0</int>
    <int name="QTime">6</int>
    <lst name="params">
      <str name="df">text</str>
      <str name="echoParams">all</str>
      <str name="rows">10</str>
      <str name="echoParams">all</str>
      <str name="q">solrpingquery</str>
      <str name="distrib">>false</str>
    </lst>
  </lst>
  <str name="status">OK</str>
</response>
```

The 'status' should be "OK" if Solr is running properly.

REST API Services Health Check

All of the Fusion API backend services (except Connectors and the UI) are started at port 8765 when the run.sh script is executed. The Fusion UI depends on all these services.

If all the services are started without any issues, then the below ping request should return response 'ok'.

```
curl http://localhost:8765/api/v1
```

As an additional check, you can also query the `system/status` endpoint, which should return a response 'started'.

```
curl http://localhost:8765/api/v1/system/status
```

The response would look like:

```
{
  "status" : "started"
}
```

Connectors Health Check

The Connectors health check can be performed by a ping request to port 8984. Similar to the previous ping request, the returned response is 'ok' if the service starts successfully.

```
curl http://localhost:8984/connectors/v1
```

As an additional check, you can also query the `system/status` endpoint, which should return a response 'started'.

```
curl http://localhost:8984/connectors/v1/system/status
```

The response would look like:

```
{  
  "status" : "started"  
}
```

1.4.12. Migrating Fusion Data

The instructions in this topic can be used to migrate Fusion data from development environments into testing and production environments, or to back up data and restore it after an incident of data loss.

- Collections and related configurations can be migrated using the Objects API and the Fusion UI (import only). Fusion objects include all your searchable data, plus pipelines, aggregations, and other configurations on which your collections depend.
- Application configuration data includes

Migrating collections and related configurations

Fusion allows you to export objects from one Fusion instance and import them into another. The data that you can migrate includes collections and all collection-related configurations.

Exporting can only be performed using the Objects API. Importing can be performed using the API or the UI.

Object export and import

Collections and encrypted values are treated specially; details are provided below. During import, conflicts are resolved according to the specified import policy.

For objects other than collections, no implicit filtering is performed; all objects are included by default. However, on export you can filter by type and ID.

Supported objects

Fusion lets you export and import these types of objects:

- `collection`
- `index-pipeline`
- `query-pipeline`
- `search-cluster`
- `datasource`
- `banana`
- `parser`
- `group`
- `link`
- `task`
- `job`
- `spark`

Exporting and importing collections

Collections are processed with these dependent objects:

- features
- index profiles
- query profiles

Datasources, parser configurations, and pipeline configurations are not included when collections are exported or imported. These must be exported and imported explicitly.

Only user-created collections are included by default. Certain types of collections are excluded:

- the "default" collection
- collections whose type is not DATA
- collections whose names start with "system_"
- "Secondary" collections, that is, collections created by features

Instead, create the same features on the target system; this automatically creates the corresponding secondary collections.

You can override these exclusions by specifying a collection, like this:

```
http://localhost:8764/api/apollo/objects/export?collection.ids=default
```

Encrypted passwords

Some objects, such as datasources and pipelines, include encrypted passwords for accessing remote data.

- On export, these encrypted values are replaced with `${secret.n.nameOfProperty}`.
- On import, the original, plaintext passwords must be provided in a JSON map:

```
{"secret.1.bindPassword" : "abc", "secret.2.bindPassword" : "def"}
```

The file must be supplied as multipart form data.

Note	Variables that do not start with <code>secret.</code> are ignored.
------	--

Import policies

On import, the `importPolicy` parameter is required. It specifies what to do if any object in the import list already exists on the target system:

<code>abort</code>	If there are conflicts, then import nothing.
<code>merge</code>	If there are conflicts, then skip the conflicting objects.
<code>overwrite</code>	If there are conflicts, then overwrite or delete/create the conflicting objects on the target system.

Filtering on export

On export, there are two ways to specify the objects to include:

- by type

You can specify a list of object types to export all objects of those types. Valid values:

- `collection`
- `index-pipeline`
- `query-pipeline`
- `search-cluster`
- `datasource`
- `banana`
- `parser`
- `group`
- `link`
- `task`
- `job`
- `spark`
- by type and ID

The `type.ids` parameter lets you list the IDs to match for the specified object type.

The `type` and `type.ids` parameters can be combined as needed.

Exporting linked objects

Related Fusion objects are linked. You can view linked objects using the Links API or the Object Explorer.

When exporting a specific Fusion object, you can also export its linked objects without specifying each one individually. To export all objects linked to the specified object, include the `deep="true"` query parameter in your request. See the example below. When `deep` is "true", Fusion follows these link types:

- `DependsOn`
- `HasPart`
- `RelatesTo`

Validation

Objects are validated before import. If any objects fail validation, the whole import request is rejected. A separate endpoint is available for validating objects without importing them.

Validation includes checking whether an object already exists on the target system and whether the user is authorized to create or modify the object.

For collection objects, the following special validation is performed:

- We check the `searchClusterId` of each collection and verify that a cluster with this ID exists on the target system or in the import file (error).
- We check that features, index profiles, and query profiles belong only to the collections specified in the import file (error).
- We check that a feature exists on the target system for each feature in the import file (error).
- We check for index profiles or query profiles that do not exist on the target system or in the import file (warning).

For **job** objects, which contain schedule configurations, Fusion only imports them if their associated **task**, **datasource**, or **spark** objects are also present, either on the target host or in the import file.

Status messages

Validation completed with no errors	The validation method was called and no errors found, though there may be warnings.
Validation found errors	The validation was called and errors found. Validation does not stop on the first error, so the complete list of errors is reported.
Validation was not completed because of system error	The validation was interrupted by system error.
Import was not performed because validation errors exist	The import method was called, but import didn't start because of validation errors.
Import was not performed because of input data error	The import method was called, but import didn't start, because Fusion could not find a substitution for one of the secret values in objects in import.
Import was not completed because of system error	The validation found no errors and import started, but it was interrupted by system error.
Import was completed	Validation found no errors and import finished successfully.

How to export Fusion objects

Exporting can only be performed using the Objects API.

You can select all objects, or limit the operation to specific object types or IDs. In addition to export endpoints, a validation endpoint is provided for troubleshooting.

Note	By default, system-created collections are not exported.
------	--

Some example requests are shown below. For complete reference information about object export endpoints, see the Objects API.

Export all objects

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export
```

Export all datasources

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export?type=datasource
```

Export a specific datasource and its linked objects

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export?export?datasource.ids=movies_csv-ml-movies&deep=true
```


Export all datasources and pipelines, plus two specific parsing configurations

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export?type=datasource,index-pipeline,query-pipeline&parser.ids=cinema_parser,metafiles_parser
```

How to import Fusion objects

Objects can be imported using the REST API or the Fusion UI.

Importing objects with the REST API

Some example requests are shown below. For complete reference information about object export endpoints, see the Objects API.

Import objects from a file and stop if there are conflicts

```
curl -u user:pass -H "Content-Type:multipart/form-data" -X POST -F
'importData=@/Users/admin/Fusion/export.json'
http://localhost:8764/api/apollo/objects/import?importPolicy=abort
```

Import objects, substitute the password variables, and merge any conflicts

```
curl -u user:pass -H "Content-Type:multipart/form-data" -X POST -F
'importData=@/Users/admin/Fusion/export.json' -F 'variableValues=@password_file.json'
http://localhost:8764/api/apollo/objects/import?importPolicy=merge
```


Note

`password_file.json` must contain plaintext passwords.

Importing objects with the Fusion UI

How to import objects using the UI

- 1.



In the upper left, click the


Launcher button and select **Devops**.

2. In the Home panel, click **Import Fusion Objects**.

The Import Fusion Objects window opens.

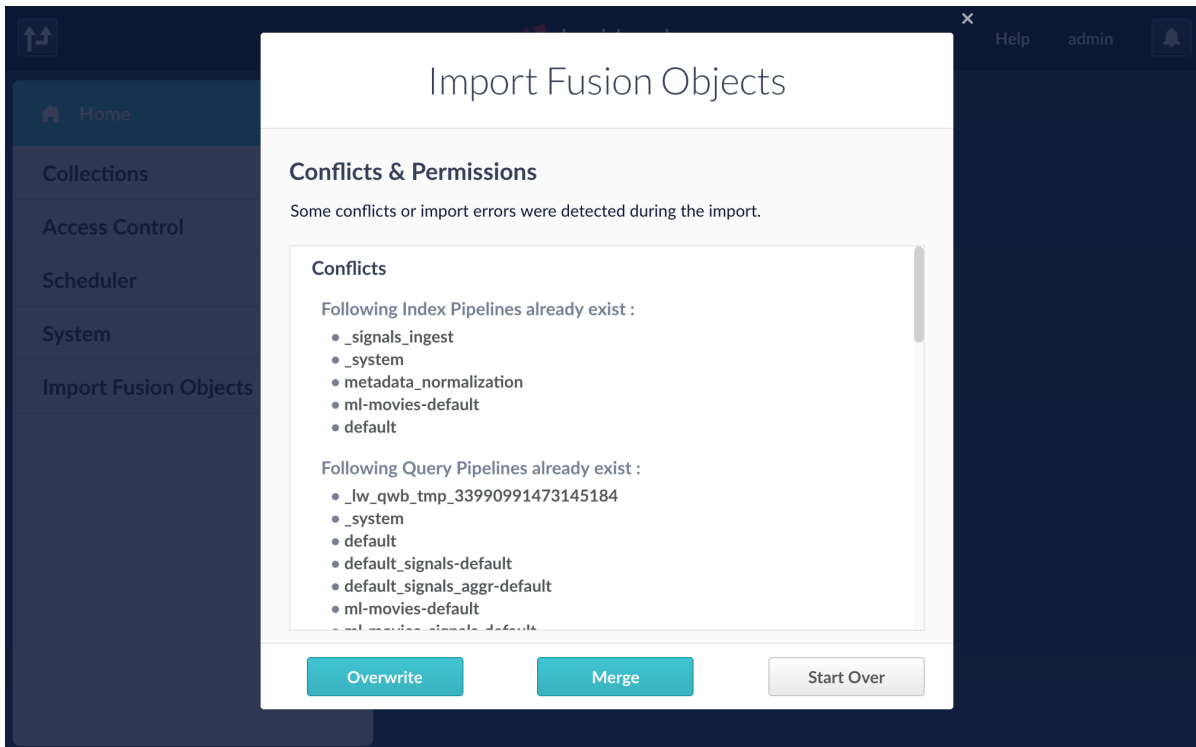
3. Select the data file from your local filesystem.

If you are importing passwords, also select the JSON file that maps variables to plaintext passwords.



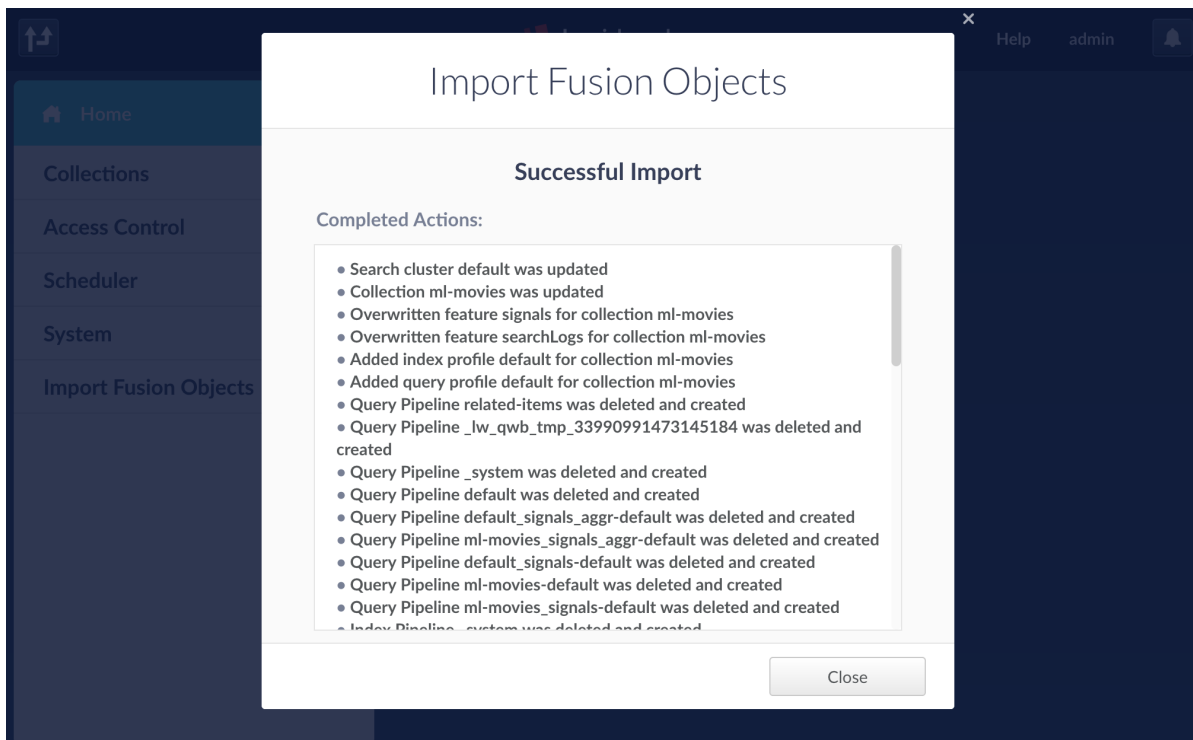
4. Click **Import**.

If there are conflicts, Fusion prompts you to specify an import policy:



- Click **Overwrite** to overwrite the objects on the target system with the ones in the import file.
- Click **Merge** to skip all conflicting objects and import only the non-conflicting objects.
- Click **Start Over** to abort the import.

Fusion confirms that the import was successful:



5. Click **Close** to close the Import Fusion Objects window.

Migrating application configuration data

ZooKeeper configuration data is used to coordinate a distributed Fusion deployment. Additionally, certain Fusion components have configuration data that can be migrated between Fusion instances.

Migrating ZooKeeper data

Migration consists of the following steps:

- Copy the ZooKeeper data nodes which contain Fusion configuration information from the FUSION-CURRENT ZooKeeper instance to the FUSION-NEW ZooKeeper instance
- Rewrite Fusion datasource and pipeline configurations, working against the FUSION-NEW ZooKeeper instance

From ZooKeeper to JSON file

To export configurations from an existing Fusion deployment, the script `zkImportExport.sh` requires parameters:

- `-cmd export` - this is the command parameter which specifies the mode in which to run this program.
- `-zkhost <connect string>` - the ZooKeeper connect string is the list of all servers,ports for the FUSION_CURRENT ZooKeeper cluster. For example, if running a single-node Fusion developer deployment with embedded ZooKeeper, the connect string is `localhost:9983`. If you have an external 3-node ZooKeeper cluster running on servers "zk1.acme.com", "zk2.acme.com", "zk3.acme.com", all listening on port 2181, then the connect string is `zk1.acme.com:2181,zk2.acme.com:2181,zk3.acme.com:2181`
- `-filename <path/to/JSON/dump/file>` - the name of the JSON dump file to save to.
- `-path <start znode>`
 - To migrate Fusion configurations for all applications, the path is `/lucid`. Migrating just the "lucid" node between the ZooKeeper services used by different Fusion deployments results in deployments which contain the same applications but not the same user databases.
 - To migrate the Fusion users, groups, roles, and realms information, the path is `/lucid-apollo-admin`.
 - To migrate all ZooKeeper data, the path is `/`.

Example: export from local developer deployment to file "znode_lucid_dump.json"

```
> {fusion_path}/scripts/zkImportExport.sh -zkhost localhost:9983 -cmd export -path /lucid -filename znode_lucid_dump.json
```

The command produces the following terminal outputs:

```

2016-06-01T19:48:12,512 - INFO [main:URLConfigurationSource@125] - URLs to be used as dynamic configuration
source: [jar:file:/Users/demo/tmp5/fusion/apps/jetty/api/webapps/api/WEB-INF/lib/lucid-base-spark-
2.2.0.jar!/config.properties]
2016-06-01T19:48:12,878 - INFO [main:DynamicPropertyFactory@281] - DynamicPropertyFactory is initialized with
configuration sources: com.netflix.config.ConcurrentCompositeConfiguration@5bf22f18
2016-06-01T19:48:12,961 - INFO [main:CloseableRegistry@45] - Registering a new closeable:
org.apache.curator.frameworkimps.CuratorFrameworkImpl@32fe9d0a
2016-06-01T19:48:12,961 - INFO [main:CuratorFrameworkImpl@234] - Starting
2016-06-01T19:48:12,974 - INFO [main:Environment@100] - Client environment:zookeeper.version=3.4.6-1569965,
built on 02/20/2014 09:09 GMT
2016-06-01T19:48:12,974 - INFO [main:Environment@100] - Client environment:host.name=10.0.1.16
2016-06-01T19:48:12,974 - INFO [main:Environment@100] - Client environment:java.version=1.8.0_25
2016-06-01T19:48:12,974 - INFO [main:Environment@100] - Client environment:java.vendor=Oracle Corporation
2016-06-01T19:48:12,975 - INFO [main:Environment@100] - Client
environment:java.home=/Library/Java/JavaVirtualMachines/jdk1.8.0_25.jdk/Contents/Home/jre
2016-06-01T19:48:12,975 - INFO [main:Environment@100] - Client
environment:java.class.path=./fusion/scripts/.. .. ( rest of path omitted )
2016-06-01T19:48:12,976 - INFO [main:Environment@100] - Client
environment:java.library.path=/Users/demo/Library/Java/Extensions: ... ( rest of path omitted )
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client
environment:java.io.tmpdir=/var/folders/jq/ms_hc8f9269f4h8k4b691d74000gp/T/
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:java.compiler=<NA>
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:os.name=Mac OS X
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:os.arch=x86_64
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:os.version=10.10.5
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:user.name=demo
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:user.home=/Users/demo
2016-06-01T19:48:12,978 - INFO [main:Environment@100] - Client environment:user.dir=/Users/demo/tmp5
2016-06-01T19:48:12,978 - INFO [main:ZooKeeper@438] - Initiating client connection,
connectString=localhost:9983 sessionTimeout=60000 watcher=org.apache.curator.ConnectionState@138fe6ec
2016-06-01T19:48:18,070 - INFO [main-SendThread(fe80:0:0:0:0:0:1%1:9983):ClientCnxn$SendThread@975] -
Opening socket connection to server fe80:0:0:0:0:0:1%1/fe80:0:0:0:0:0:1%1:9983. Will not attempt to
authenticate using SASL (unknown error)
2016-06-01T19:48:18,111 - INFO [main-SendThread(fe80:0:0:0:0:0:1%1:9983):ClientCnxn$SendThread@852] -
Socket connection established to fe80:0:0:0:0:0:1%1/fe80:0:0:0:0:0:1%1:9983, initiating session
2016-06-01T19:48:18,118 - INFO [main-SendThread(fe80:0:0:0:0:0:1%1:9983):ClientCnxn$SendThread@1235] -
Session establishment complete on server fe80:0:0:0:0:0:1%1/fe80:0:0:0:0:0:1%1:9983, sessionId =
0x1550df6b0180017, negotiated timeout = 40000
2016-06-01T19:48:18,121 - INFO [main-EventThread:ConnectionStateManager@228] - State change: CONNECTED
2016-06-01T19:48:18,367 - INFO [main:ZKImportExportCli@198] - Data written to file
'/Users/demo/tmp5/znode_lucid_dump.json'
2016-06-01T19:48:18,370 - INFO [main:ZooKeeper@684] - Session: 0x1550df6b0180017 closed
2016-06-01T19:48:18,370 - INFO [main-EventThread:ClientCnxn$EventThread@512] - EventThread shut down

```

The resulting JSON output file contains the znode hierarchy for znode "lucid", with ZooKeeper binary data:

exported configurations should be uploaded using the script command argument `-cmd update`.

update command example:

```
> {fusion_path}/scripts/zkImportExport.sh -zkhost localhost:9983 -cmd update -path /lucid -filename
znode_lucid_dump.json
```

To verify, start all Fusion services and log in to the new Fusion installation and verify that this installation contains the same set of collections and datasources as the existing collection, and that all Fusion pipelines and stages match those of the existing Fusion installation.

Existing application, existing Fusion deployment

When migrating an existing application to a Fusion deployment which is already running a version of that application, the exported configurations should be uploaded using the script command argument `-cmd update --overwrite`.

update --overwrite command example:

```
> {fusion_path}/scripts/zkImportExport.sh -zkhost localhost:9983 -cmd update --override -path /lucid -filename
znode_lucid_dump.json
```

To verify, start all Fusion services and log in to the new Fusion installation and verify that this installation contains the same set of collections and datasources as the existing collection, and that all Fusion pipelines and stages match those of the existing Fusion installation.

Caveats

- All datasource configurations are copied over as is. If the set of repositories used to populate the collections changes according to deployment environment, then these datasources will need to be updated accordingly.
- The import export script is only guaranteed to work between Fusion deployments running the same Fusion version. The should work across all releases for the same Major.minor version of Fusion, e.g. you should be able to migrate between versions 2.4.1 and 2.4.2. If the set of configurations needed for an application have the same structure and properties across two different versions, these scripts *might* work.

Migrating Fusion component configuration data

The directory `fusion/3.1.x/data` contains the on-disk data stores managed directly or indirectly by Fusion services.

- `fusion/3.1.x/data/connectors` contains data required by Fusion connectors.
 - `fusion/3.1.x/data/connectors/lucid.jdbc` contains third-party JDBC driver files. If your application uses a JDBC connector, you must copy this information over to every server on which will this connector will run.
 - `fusion/3.1.x/data/connectors/crawldb` contains information on the filed visited during a crawl. (Preserving crawldb history may not be possible if there are multiple different servers running Fusion connectors services.)
- `fusion/3.1.x/data/nlp` contains data used by Fusion NLP pipeline stages. If you are using Fusion's NLP components for sentence detection, part-of-speech tagging, and named entity detection, you must copy over the model files stored under this directory.
- `fusion/3.1.x/data/solr` contains the backing store for Fusion's embedded Solr (developer deployment only).
- `fusion/3.1.x/data/zookeeper` contains the backing store for Fusion's embedded ZooKeeper (developer deployment

only).

When migrating these directories, no Fusion services which may change the contents should be running. The choice of which directories to migrate and the utilities used to do the migration are entirely dependent upon the platform, environment, and deployment configurations.

1.5. The Fusion Workflow

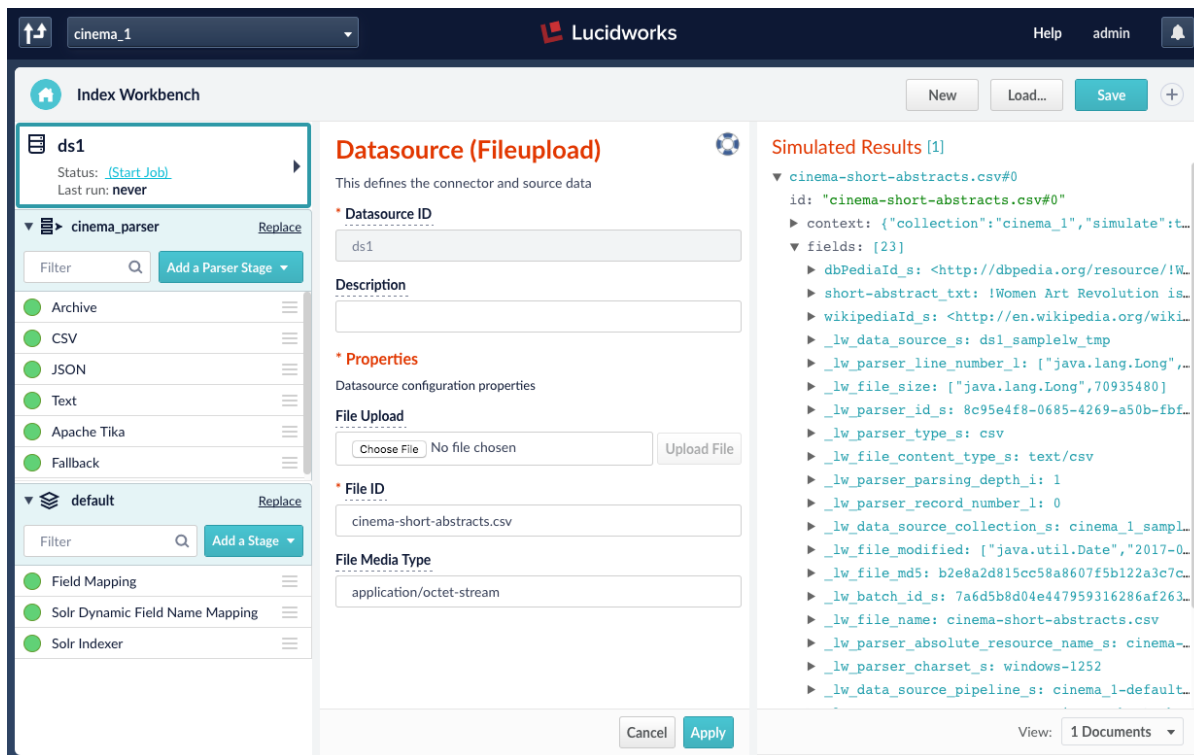
Here's a high-level overview of the Fusion workflow for search developers:

1. Use the Index Workbench to develop your datasources.
2. Use the Query Workbench to develop useful search results.
3. Develop your search application.

To get started, log in to the Fusion UI and click **Search**.

1.5.1. The Index Workbench

Start here to get your data into Fusion in a format that's optimized for your search application.



The screenshot shows the Lucidworks Index Workbench interface. The top navigation bar includes 'cinema_1', the Lucidworks logo, and 'Help admin'. The main interface is divided into three panels:

- Index Workbench:** A sidebar on the left with a tree view showing 'ds1' (Status: Start Job, Last run: never) and two folders: 'cinema_parser' and 'default'. Each folder has a 'Replace' button and a list of components like Archive, CSV, JSON, Text, Apache Tika, Fallback, Field Mapping, Solr Dynamic Field Name Mapping, and Solr Indexer.
- Datasource (Fileupload):** The central configuration panel. It includes a 'Datasource ID' field (ds1), a 'Description' field, a 'Properties' section for 'Datasource configuration properties', a 'File Upload' section with a 'Choose File' button and 'Upload File' button, a 'File ID' field (cinema-short-abstracts.csv), and a 'File Media Type' field (application/octet-stream). 'Cancel' and 'Apply' buttons are at the bottom.
- Simulated Results [1]:** A panel on the right displaying a JSON document structure for 'cinema-short-abstracts.csv#0'. It shows fields like 'id', 'context', and a list of 23 fields including 'dbPediaId_s', 'short-abstract_txt', 'wikipediaId_s', and various Lucidworks internal fields like '_lw_data_source_s', '_lw_parser_line_number_l', '_lw_file_size', etc.

Select what to configure: datasource, parser, or index pipeline.

Configure the selected component's options.

View simulated results to see how your configuration will affect your data.

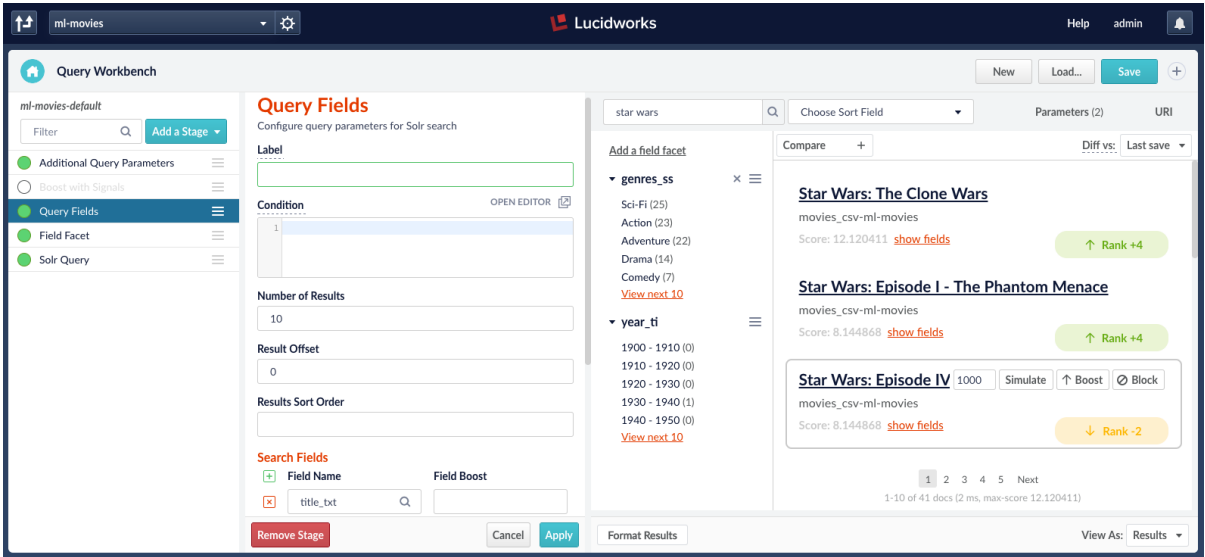
The Index Workbench is a workspace for configuring datasources, parsers, and index pipelines. Here you can configure Fusion to manipulate incoming data, transform its fields, conditionally include or exclude documents, send alerts, and more.

Simulated results are displayed and updated on the fly so you can verify that your configuration will index the data in the format you need for your search application. No indexing occurs until you're satisfied with your configuration.

See The Index Workbench for more details. The Index Pipeline Stages Reference provides complete information about the wide variety of features available in index pipelines.

1.5.2. The Query Workbench

Once your data is indexed, use the Query Workbench to configure how search results are returned.



Select or re-order the stages in your query pipeline.

Configure the selected query pipeline stage.

Configure faceting.

Preview your search results, boost/block results, and view fields.

Here you configure query pipelines to manipulate incoming queries and the results that Fusion returns. You can configure relevancy, faceting, security trimming, external lookups, alerting, and more.

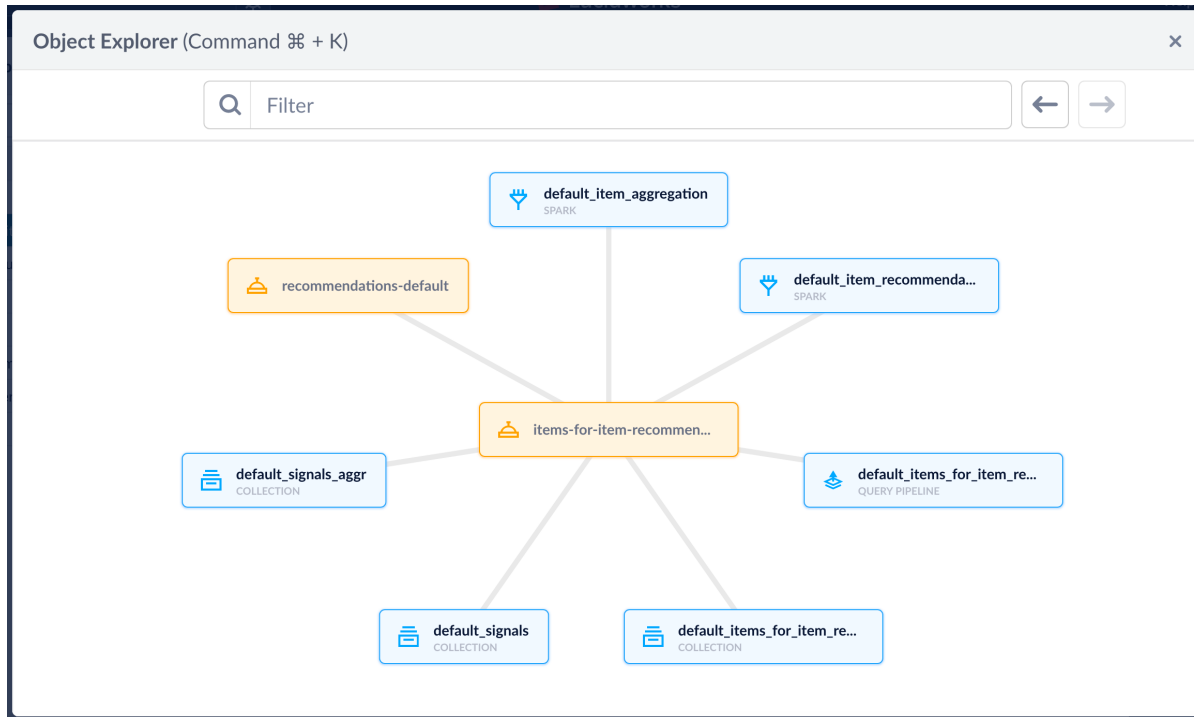
See The Query Workbench for details about how to use this tool. The Query Pipeline Stages Reference provides complete information about the wide variety of features available in query pipelines.

1.5.3. Developing your search application

Ultimately, Fusion is the back end for your own search applications. See Search Applications for details about how to develop a front-end interface, or adapt your existing application, for querying your Fusion collections and displaying search results.

1.6. Explore Objects

Object Explorer lets you easily find Fusion objects across an entire Fusion deployment and navigate between objects. When you select an object, Object Explorer displays all related objects, for example:












You can perform some management tasks inside Object Explorer. For example, you can schedule jobs and tasks from Object Explorer. For other management tasks, Object Explorer takes you to where you need to be.


Note: Which objects you can explore depends on the permissions that your security realm and user definition grant you. As the user `admin`, you can explore all objects supported by Object Explorer.

1.6.1. Object Types

















You can explore these types of objects with Object Explorer:

 Spark jobs including aggregation jobs	 Collection	 Datasource
 Index pipeline	 Parser	 Query pipeline
 Recommendations group	 System group	 User group

1.6.2. Use Object Explorer

- **Open Object Explorer** – On a PC, press Ctrl+K. On a Mac, press Command +K. Alternatively, click **Devops** > **Home**  > **Object Explorer**. Object Explorer opens and displays either a list of all objects or the list of objects that match your most recent search in Object Explorer, whether during this Fusion session or a previous session.
- **Browse all objects** – Open Object Explorer. Object Explorer displays a list of all objects. If you don't see a list of objects, press Enter in the **Filter** box.
- **Search for objects** – Type any part of an object's name in the **Filter** box. Object Explorer displays a list of matching

objects.

- **Display related objects** (with a mouse) – With a list of objects present, click an object to select it. Object Explorer centers the object and displays all related objects.
- **Display related objects** (with Tab keys) – With a list of objects present, tab from object to object. With an object highlighted, press Enter to select it. Object Explorer centers the object and displays all related objects.
- **Move backward and forward through a history of your selections** – After you make a series of object selections, you can move backward and forward through that history. To the right of the **Filter** box, click Back  and Forward , or use the left arrow  and right arrow  keys.
- **Edit an object's name** – For some objects, for example, Group objects, clicking  lets you edit the object's name.
- **Manage an object** – For most objects that display the  icon when you hover over the object,  takes you to the part of the Fusion UI where you can manage the object.
- **Schedule a job or task** – Hover over an object that displays the Schedule  icon when you hover. Click Schedule . From the scheduling dialog, you can click  to open the Scheduler.
- **View existing groups** – A *group* is a means of tagging objects with a shared, arbitrary identifier. Hover over an object that displays the Add To Group  icon when you hover. Click Add To Group . View the list of groups. To close the popup dialog, click anywhere on Object Explorer's background.
- **Add an object to a group** – Hover over an object that displays the Add To Group  icon when you hover. Click Add To Group . Choose an existing group, or enter the name of a new group.
- **Remove an object from a group** – Select the object you want to remove from a group. Object Explorer centers the object. Hover over the group from which you want to remove the object, and then click Remove From Group . Alternatively, select the group and then hover over the object. If this is the last object removed from a group, Object Explorer deletes the group.
- **Close Object Explorer** – On a PC or Mac, press Esc or click Close . Alternatively on a Mac, you can press Command +K.

Chapter 2. Search

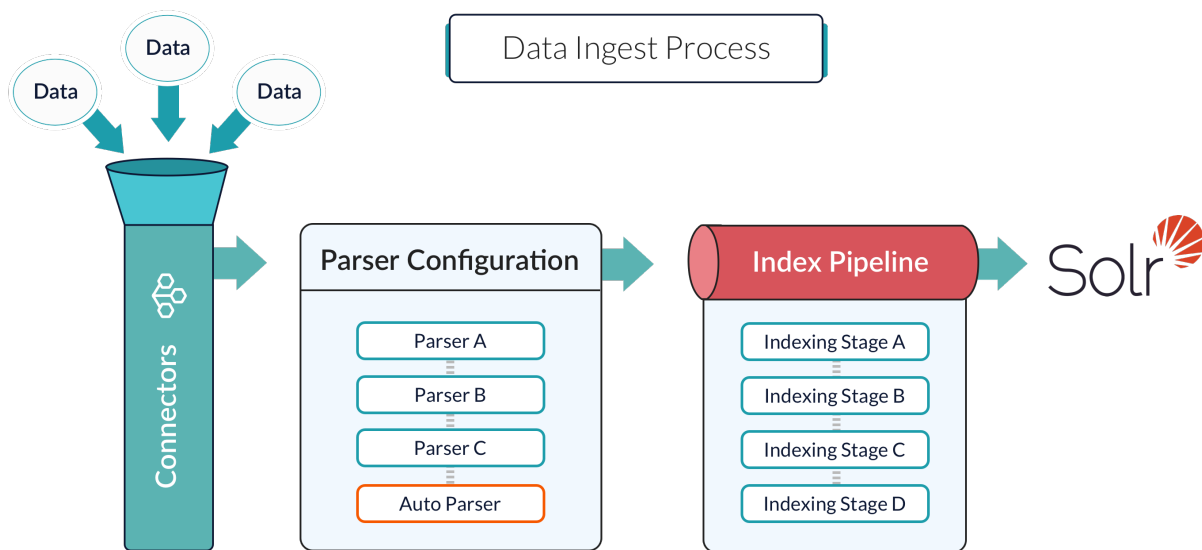
- Datasources are the configurations that import and index data into a collection.
- Query pipelines filter, transform, and augment Solr queries and responses in order to return all and only the most relevant search results.
- Signals and aggregations are ways to collect and compile data for analysis or to boost search results in the future.
- Search applications are the front-end interfaces that you build on top of Fusion using its REST API.

2.1. Datasources

A collection includes one or more datasources. A datasource is a configuration that manages the import and indexing of data into the collection.

The Index Workbench provides a development environment for creating, configuring, and testing a datasource configuration. Every datasource configuration includes the following:

- Connector configuration, specifying the source and format of the incoming data.
- Parser configuration, describing a series of conditional parsing stages to transform the incoming data into PipelineDocument objects.
- Index pipeline configuration, consisting of stages that transform PipelineDocument objects into Solr documents to be indexed.

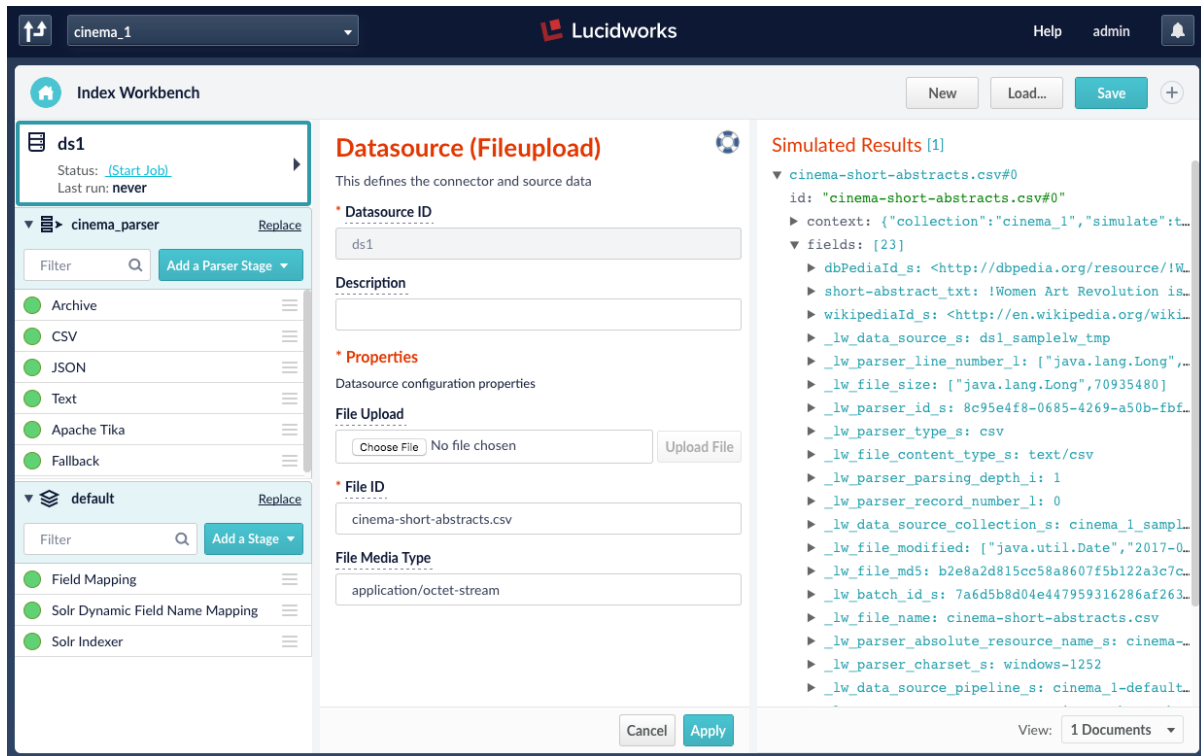


Collections and datasources can also be managed through the REST API.

In some cases it may make sense to bypass the connectors and use other ingest methods for your data.

2.1.1. Index Workbench

The Index Workbench is a powerful tool that combines key aspects of the the data indexing configuration process into one user-friendly panel. It guides the user through the workflow for configuring datasources, parsers and index pipelines.



Select what to configure: datasource, parser, and/or your index pipeline

Configure the selected component's options

View simulated results to see how your configuration will affect your data

Datasources

To set up a new datasource, you have several options. You can paste in a website URL for Fusion to crawl, or you can navigate to a locally saved file through the File Finder dialog. This quicker setup for commonly used Web and Local Filesystem datasources saves steps in the process and helps newer users get started quickly. Alternatively, the Datasource selection dropdown allows you to quickly select and navigate to any configurable datasource option that Fusion has. If there are existing datasources that you have already configured and saved in the Collection, they can be quickly accessed from this pane as well.

Add A New Datasource

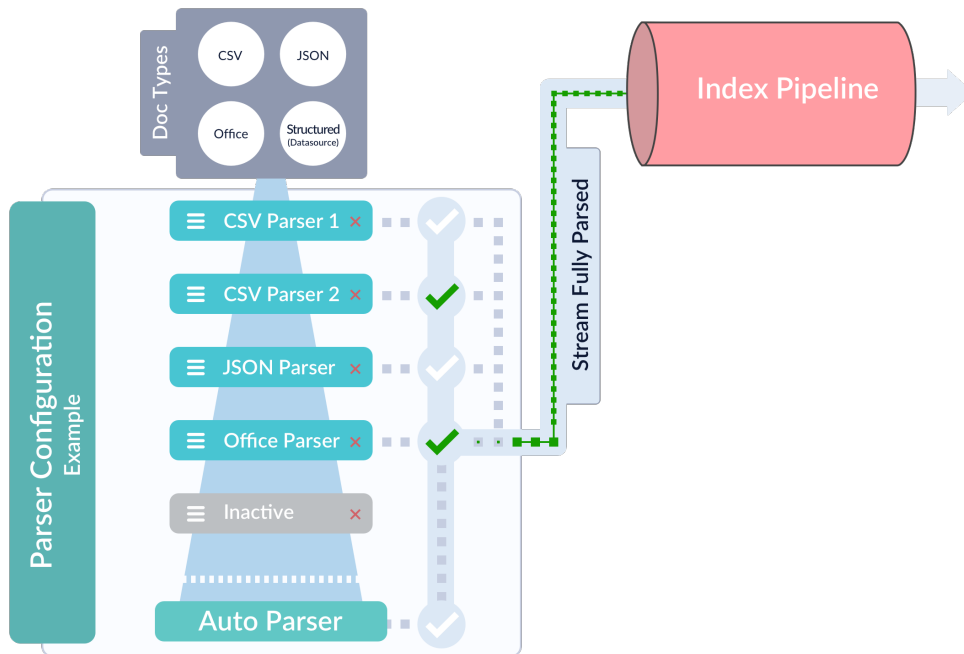
Point to a URL: Paste a URL for web crawling

Or, upload a file: Pull data from a locally saved file

Or, if you know the location of your data: Choose from the full list of datasources

Parser Configuration

In Fusion 3.0, parsers have been introduced as their own configurable component of the indexing workflow. They allow greater flexibility and specificity when parsing inbound data.



A parser consists of an ordered list of parser stages that is completely customizable. The same parser stage can be added to a given configuration multiple times if the different specified settings within those stages best suits the parsing of the data. There is no limit to the number of parser stages that can be included in a parser, and the order in which they run is also completely flexible. In a parser, after all of the doctype-specific parser stages have run, the Tika and Fallback stages are useful catch-all stages that can attempt to parse anything that has not yet been matched. Tika is used for parsing many types of unstructured documents like PDFs, DOCX, and many more. If all of the other stages in the parser

fail to completely parse the data, the Fallback stage can copy the raw bytes directly to Solr.

Index Pipelines

An Index Pipeline transforms incoming data into a document suitable for indexing by Solr via a series of modularized operations called stages. Fusion provides a variety of specialized index stages to index data effectively. Stages can be selected, configured, and enabled or disabled in the Index Pipeline section of the Index Workbench.

The screenshot displays the 'Index Workbench' interface. At the top, there is a header with a home icon and the text 'Index Workbench'. Below this, a card for 'ds1' is shown with a status of '(Start Job)' and 'Last run: never'. A navigation bar contains 'ds1' and a 'Replace' button. Below the navigation bar, a card for 'cinema_1-default' is visible, also with a 'Replace' button. A search bar labeled 'Filter' and an 'Add a Stage' button are present. A dropdown menu is open, showing a search bar and a list of stage options: Natural Language Processing, Detect Sentences, Gazetteer Lookup Extraction, OpenNLP NER Extraction, Tag Part-of-Speech, Indexing, Solr Indexer, Solr Partial Update Indexer, Troubleshooting, Logging, and Send PagerDuty Message. To the right, a 'Set Pipeline ID' section shows a 'Pipeline ID' field with the value 'cinema_1-default'.

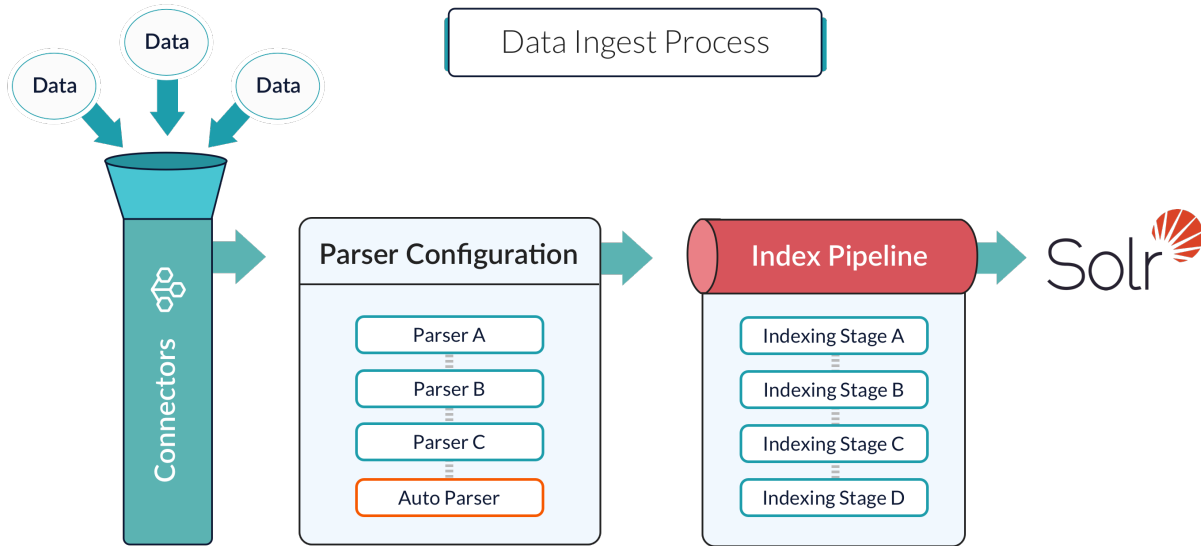
Once you finish configuring a datasource using the Index Workbench, you can move on to setting up queries using the

Query Workbench, which provides a similar workflow for configuring and previewing search results.

2.1.2. Index Pipelines

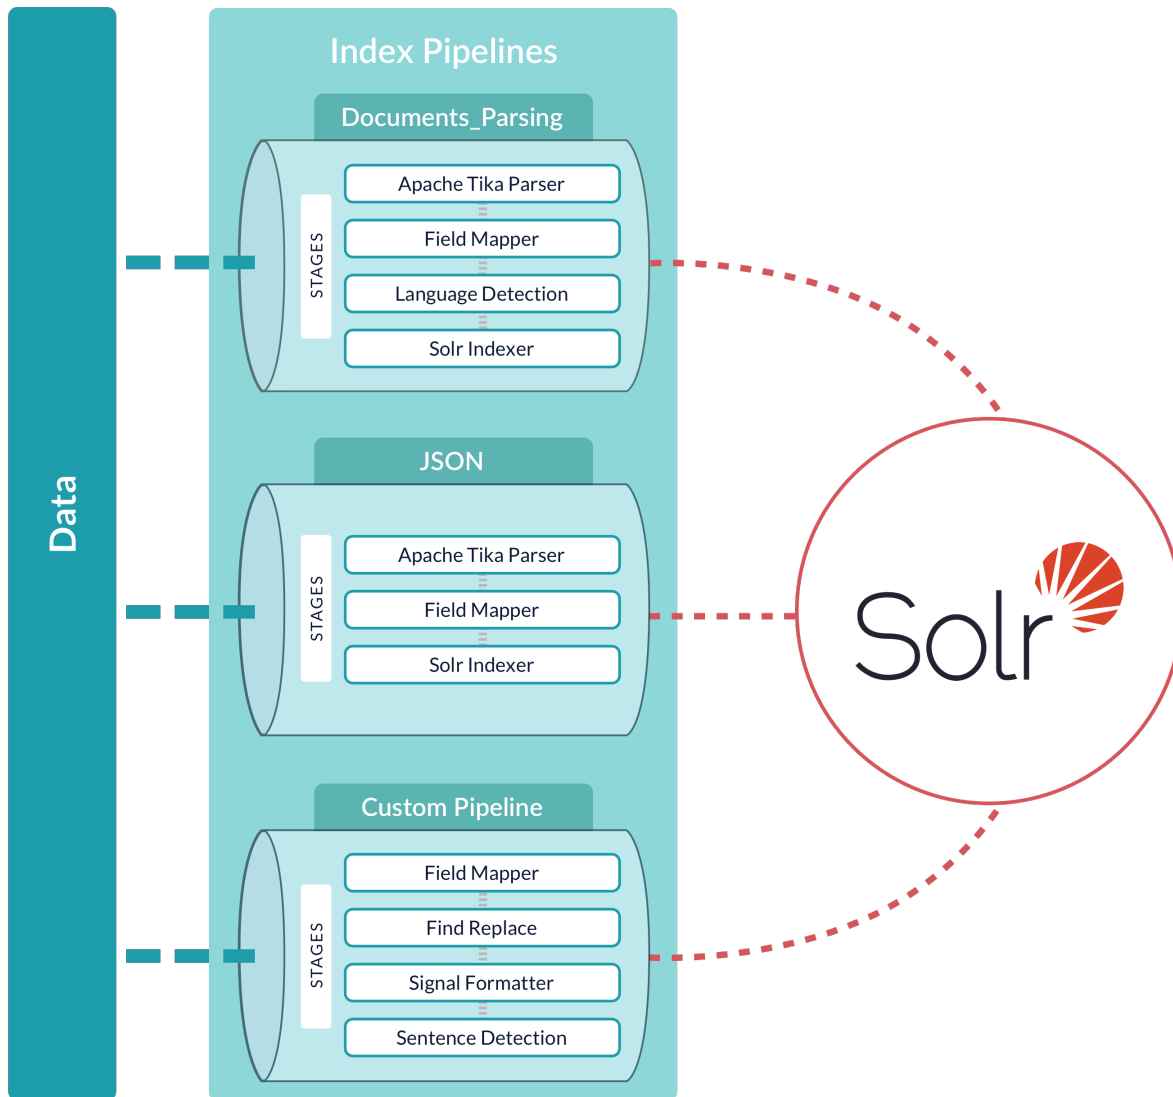
Index pipelines transform incoming data into PipelineDocument objects for indexing by Fusion's Solr core. An index pipeline consists of a series of configurable index pipeline stages, each performing a different transformation on the data before passing the result to the next stage in the pipeline. The final stage is the Solr Indexer stage, which transforms the PipelineDocument into a Solr document and submits it to Solr for indexing in a specific Collection.

Each configured datasource has an associated index pipeline and uses a connector to fetch data to parse and then input into the index pipeline.



Alternatively, documents can be submitted directly to an Index Pipeline via the REST API; see Pushing Documents to a Pipeline.

A pipeline can be re-used across multiple collections. Fusion provides a set of built-in pipelines. You can use the Index Workbench or the REST API to develop custom index pipelines to suit any datasource or application.




Collection-specific Pipelines

When a Fusion collection is created using the Fusion UI, a pair of index and query pipelines are created to that pipeline, where the pipeline name is the collection name with the suffix "-default". This pipeline consists of a Field Mapping index stage

Although default pipelines are created when a Fusion collection is created, they are not deleted when the collection is deleted. This is due to the fact that pipelines can be used across collections, therefore a named pipeline, although originally associated with a collection, may be used by several collections.

Pre-configured Pipelines

Fusion includes several pre-configured pipelines which provide out-of-the-box processing capabilities and/or a starting point for customization. There are also a set of named pipelines which are used by Fusion services for logging, signal processing, and signal aggregation.



General Purpose Pipelines

- **CSV** - a pipeline for handling tabular data from CSV files, using these stages:
 1. CSV Parsing index stage
 2. Field Mapping index stage (with no pre-defined mappings)
 3. Solr Indexer stage
- **Default_Data** - a pipeline for processing general key-value data, i.e., data which has already been parsed into key-value pairs.
 1. Field Mapping index stage
 2. Solr Indexer stage.
- **Discard** (Fusion 2.0) / **conn_noop**** - a pipeline used for testing datasource configurations which has no defined stages.
- **Documents_Parsing** (Fusion 2.0) - a pipeline used to parse and index documents.
 1. Apache Tika Parser index stage, which can process a wide variety of common document formats and is able to extract metadata as well.
 2. Field Mapping index stage, which has mapping rules for common document elements.
 3. Detect Language index stage
 4. Solr Indexer stage
- **Documents_Parsing_debug_logging** (Fusion 2.0) - this pipeline is an augmented version of the Documents_Parsing pipeline where a logging stage has been added before every processing stage.

- **JSON** - a pipeline for handling JSON data.
 1. JSON Parsing index stage.
 2. Field Mapping index stage
 3. Solr Indexer stage
- **Source_Code** - a pipeline for extracting source code from Git and SVN repositories.
 1. Apache Tika Parser index stage, to extract raw data.
 2. Field Mapping index stage to map elements to proper Solr field names.
 3. Solr Indexer stage

Legacy Pipelines

- **conn_solr** - a pipeline used to parse and index documents. The initial stage is a Tika Parser index stage. The next stage is a Field Mapper index stage which has mapping rules for common document elements. The final stage is a Solr Indexer stage.
- **default** - a pipeline which consists of just a Solr Indexer stage, used to push documents which have been completely parsed and have appropriately named fields to Solr for indexing.

Internal Use Pipelines

- **_aggregation_rollup** - also a pipeline which consists of a single Solr Indexer stage which sends aggregations to Solr.
- **_signals_ingest** - a pipeline used to index raw signal data. It has three stages, a Format Signals stage, a Field Mapping stage and a Solr Indexer stage to index the raw signal events.
- **_system_metrics** - a pipeline which consists of a single Solr Indexer stage which sends internal information to the Fusion system_metrics collection.

Fusion PipelineDocument Objects

A PipelineDocument organizes the contents of each document submitted to the pipeline, document-level metadata, and processing commands into a list of fields where each field has a string name, a value, an associated metadata object and a list of annotations. A Solr Indexer stage transforms a PipelineDocument into a Solr document and submits it to Solr for indexing.

The PipelineDocument Java Object

Under the Fusion hood, a PipelineDocument is a Java object, see the PipelineDocument javadocs.

JSON representation of a PipelineDocument

The JSON representation of a PipelineDocument object has four fields:

- **id** : value is a string identifier.
- **commands** : value is a list of processing commands for the index (optional)
- **metadata** : value is single object containing a name : value pair (optional)
- **fields**: value is a list of field objects, where a field object consists of four fields:
 - **name** : value is a string containing the field name
 - **value** : value is a string containing the field value
 - **metadata** : value is a single object containing a name : value pair (optional)
 - **annotations** : value is a list of annotations (optional)

Pipeline stages add, remove, and update the fields of the PipelineDocument. The Solr Indexer stage transforms the list of PipelineDocument fields into a set of Solr document fields.

The commands field can be used to issue a commit at the end of document processing or to delete documents based on documents that match an included query.

If a pipeline includes a logging stage, the PipelineDocument will be pretty-printed to the Fusion connectors logfile (default location `fusion/3.1.x/var/log/connectors/connectors.log`). To see how this works, we set up a pipeline consisting of an initial logging stage, followed by a Apache Tika Parser stage, followed by another logging stage, followed by a Field Mapping stage.

We define a datasource named "email" configured with lucid.anda filesystem connector that submits the contents of a file named "test_email.eml" to the pipeline.

The initial logging stage from the connectors logfile is shown below. The raw bytes from the file are encoded as a BASE64 string in field "`raw_content`". After the initial logging stage (before Tika Parsing), the PipelineDocument object is:


```
{ "id" : "/Users/mitzimorris/tmp/test_email.eml",
  "fields" : [ {
    "name" : "parsing_time_1",
    "value" : [ "java.lang.Long", 157 ],
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "subject",
    "value" : "this is the subject of email message",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "dcterms:created",
    "value" : "2015-02-28T03:13:41Z",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, ...
}
```

The following is an example of an empty PipelineDocument that issues a commit command on the index:

```
{ "fields" : [ ],
  "metadata" : { },
  "commands" : [ {
    "name" : "commit",
    "params" : { }
  } ]
}
```

Submitting PipelineDocuments Directly to a Pipeline via the REST API

PipelineDocuments can be submitted to a pipeline as a POST request to the Fusion REST API path:

[/api/apollo/index-pipelines/<id>/collections/<collectionName>/index](#)

where <id> is the name of an specific pipeline and <collectionName> is the name of a specific collection. The content type header to use for this format is:

[application/vnd.lucidworks-document](#)

Example

Send two documents to collection named "docs" using the "conn_solr" pipeline:

```
curl -u user:pass -X POST -H "Content-Type: application/vnd.lucidworks-document" -d '[{"id":"myDoc1",
"fields":[{"name":"title", "value":"My first document"}, {"name":"body", "value":"This is a simple
document."}], {"id":"myDoc2", "fields":[{"name":"title", "value":"My second document"}, {"name":"body",
"value":"This is another simple document."}]}' http://localhost:8764/api/apollo/index-
pipelines/conn_solr/collections/docs/index
```

Limitations of using the index-pipelines REST API

All fields of a pipeline document sent to a pipeline as an HTTP POST request are treated as plain text strings. If you encode data in a binary format as a BASE64-encoded string, you must then add a stage to decode that data before the Tika Parser stage, else it will be parsed as plain text by Tika.

Index Pipeline Stages

An Index Pipeline takes content and transforms it into a document suitable for indexing by Solr via a series of modular operations called stages. The objects sent from stage to stage are PipelineDocument objects. Fusion provides many specialized index stages as well as a JavaScript Index stage that allows for custom processing via a JavaScript program. The general outline of the Extract/Transform/Load processing performed by an index pipeline is:

- Raw content is parsed into one or more PipelineDocument objects.
- Any number of intermediate stages operate on the document fields directly, or, in the case of specialized NLP tools, add annotations to a document.
- Finally, the PipelineDocument is sent to Solr for indexing.

A pipeline stage definition associates a unique ID with a set of properties. Pipeline definitions are stored in ZooKeeper for reuse across pipelines and search applications. The Fusion UI provides stage-specific panels used to define and configure each pipeline stage. Alternatively, JSON can be used to specify the sequence of pipeline stages and registered via the Fusion REST API. Some stages require additional resources, e.g., text files which contain lists of names, synonyms, places, or binary files which NLP language models. These resources can be uploaded via the Fusion UI or the REST API.

Available index pipeline stages are listed below:

Document transformation

- Apache Tika Parser
- CSV Parsing
- HTML Transformation
- JSON Parsing
- XML Transformation

Document filtering and enrichment

- Detect Language
- Exclude Documents
- Format Signals
- Include Documents
- JDBC Lookup
- REST Query

Field transformation

- Date Parsing
- Field Mapping
- Filter Short Fields
- Find and Replace
- Regex Field Extraction
- Regex Field Filter

- Regex Field Replacement
- Resolve Multivalued Fields
- Solr Dynamic Field Name Mapping

Natural language processing

- Detect Sentences
- Gazetteer Lookup Extraction
- OpenNLP NER Extraction
- Tag Part-of-Speech

Indexing

- Solr Indexer
- Solr Partial Update Indexer

Troubleshooting

- Logging
- Send PagerDuty Message
- Send SMTP Email
- Send Slack Message
- Write Log Message

Advanced

- Call Pipeline
- Exclusion Filter
- Javascript
- Machine Learning
- Set Property
- Update Experiment

Index Profiles

Index profiles allow your applications to send documents for indexing to a consistent endpoint (the profile alias) and change the backend index pipeline as needed. The profile is also a simple way to use one pipeline for multiple collections without any one collection "owning" the pipeline.

Associating a profile with a pipeline is simply a mapping, and the profile is considered an alias for the pipeline. The mapping can be managed in the UI or with the Index Profiles API.

The Profiles tab shows the index profiles on the left and the query profiles on the right. Hover over the name of a profile, and an **edit** button will appear to allow you to change the pipeline the profile is mapped to. Hover over the name of a pipeline, and you will be able to jump to edit that pipeline.

Click **Add Profile** to add a profile. The next screen will show a form allowing you to define the profile name and either select an existing pipeline or create a new pipeline with the name you choose. Click **Create** to save the new profile.

Entity Extraction

Fusion includes extensive entity extraction capabilities. Entity extraction is configured as an index pipeline stage and there are several stage types to correspond to the different types of entity extraction you'd like to perform on your documents. The different types are described in more detail below.

Many of the entity extraction capabilities require models or lookup files, and we have provided a number of these by default. You can find the files in [fusion/3.1.x/data/nlp/](#) but in order to use them in an index pipeline stage, you will need to load them to Solr using the Blob Store API.

Loading Models and Lookup Files to Solr

To load the files, you simply need to make a PUT request with the Blob Store API, as in this example:

```
curl -u admin:pass -X PUT --data-binary @data/nlp/models/en-sent.bin -H 'Content-type: application/octet-stream' http://localhost:8764/api/apollo/blobs/sentenceModel.bin
```

Note that the endpoint is the ID of the file. You will use the ID you've assigned with the endpoint in the index pipeline definition to indicate the model or lookup file to use. If the ID is omitted from the request (which is possible with a POST request), a random ID will be assigned and it will be difficult to tell one stored blob from another.

Entity Extraction Capabilities

Lookup Lists

The lookup lists are the most numerous of the available files. Many of these are simple lists that you will want to add values to. However, some may be robust enough for your needs.

The available lists are found in [fusion/3.1.x/data/nlp/gazetteer](#).

To use a lookup list-based entity extraction, you would configure a Gazetteer Lookup Extraction index stage as part of your pipeline.

OpenNLP Models

We have also included entity extraction models from the OpenNLP project. These models are based on news articles and may not be suitable for all entity extraction needs. The Fusion-supplied models are located in [fusion/3.1.x/data/nlp/models](#).

If you have created your own model for your data, you can load it to the blob store and use it in the nlp index stage as described above.

To use an OpenNLP-based entity extraction, you would configure a OpenNLP NER Extraction index stage as part of your pipeline.

Regular Expression Extraction

The regular expression extraction allows using a regular expression to find entities in documents that should be extracted. The extracted entities will then be copied to a field defined by you.

To use regular expression extraction, you would configure a Regex Field Extraction index stage as part of your pipeline.

Regex Field Filter

The Regex Field Filter stage allows you to remove a field based on a regular expression. This removes the entire field from the document, there is not yet an option for removing only specific entities found in the field or for excluding entire documents based on values found in a field.

To use Regex Field Filter, you would configure a Regex Field Filter index stage as part of your pipeline.

Exclusion Lists

An exclusion list is a list of known items that should be removed from a document. This removes the entire field from the document, there is not yet an option for removing only specific entities found in the field or for excluding entire documents based on values found in a field.

Note that the list must be loaded to Solr using the Blob Store API before it can be used with a Fusion index pipeline.

To use exclusion list filtering, you would configure a Exclusion Filter Index Stage as part of your pipeline.

Filter Short Fields

The Filter Short Fields stage removes entities that are equal to or smaller than a defined character limit.

To use the Filter Short Fields stage, you would configure a Filter Short Fields index stage as part of your pipeline.

Blob Storage

Fusion accepts large binary objects (blobs) for upload, and stores them in Solr. Blob uploads are used to install models, lookup lists, JDBC drivers, connectors, and more.

Blob Types

A `resourceType` query parameter can be used to specify the a blob type. For example, specify `plugin:connector` when uploading a connector, like this:

```
curl -H 'content-type:application/zip' -X PUT
'localhost:8765/api/v1/blobs/myplugin?resourceType=plugin:connector' --data-binary @myplugin.zip
```

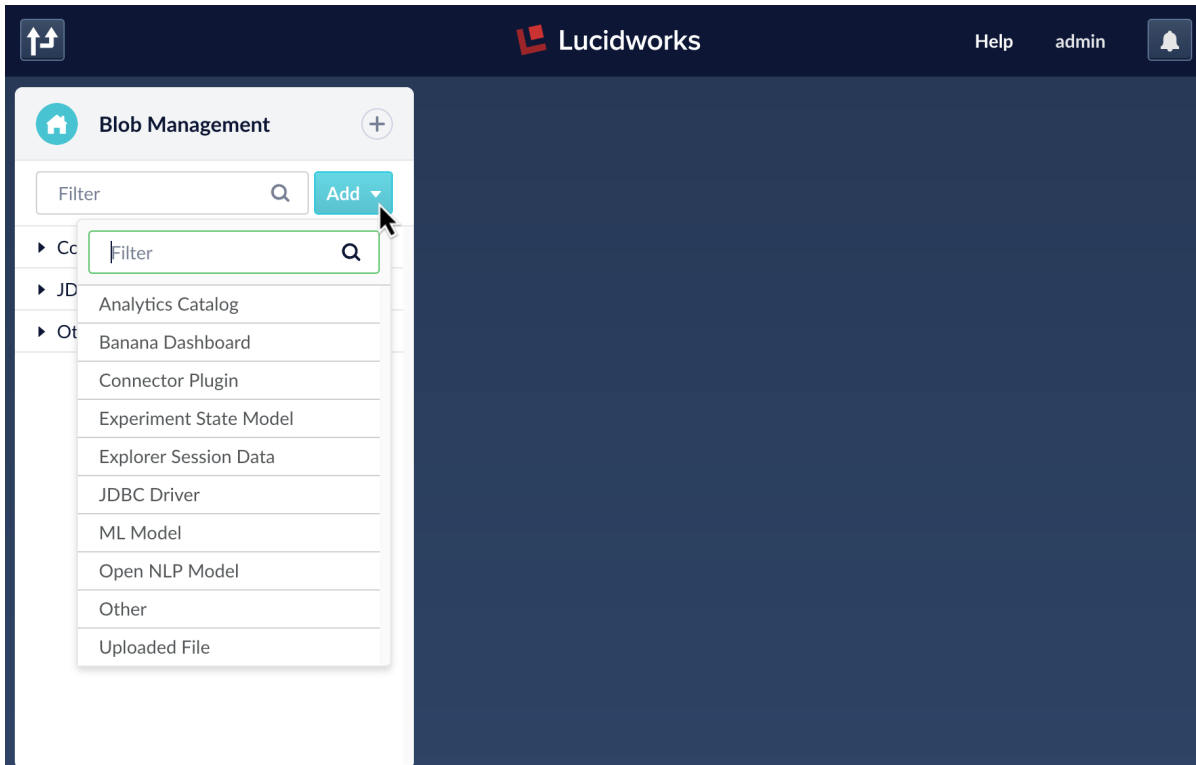
The complete list of valid values for `resourceType` is below:

Type	Description
<code>catalog</code>	An analytics catalog
<code>driver:jdbc</code>	A JDBC driver
<code>plugin:connector</code>	A connector plugin
<code>model:ml-model</code>	A machine learning model
<code>model:open-nlp</code>	An OpenNLP model
<code>file-upload</code>	Any uploaded file, such as from the Quickstart or the Index Workbench.
<code>banana</code>	A Banana dashboard
<code>other</code>	A blob of unknown type If no <code>resourceType</code> is specified on upload, "other" is assigned by default.

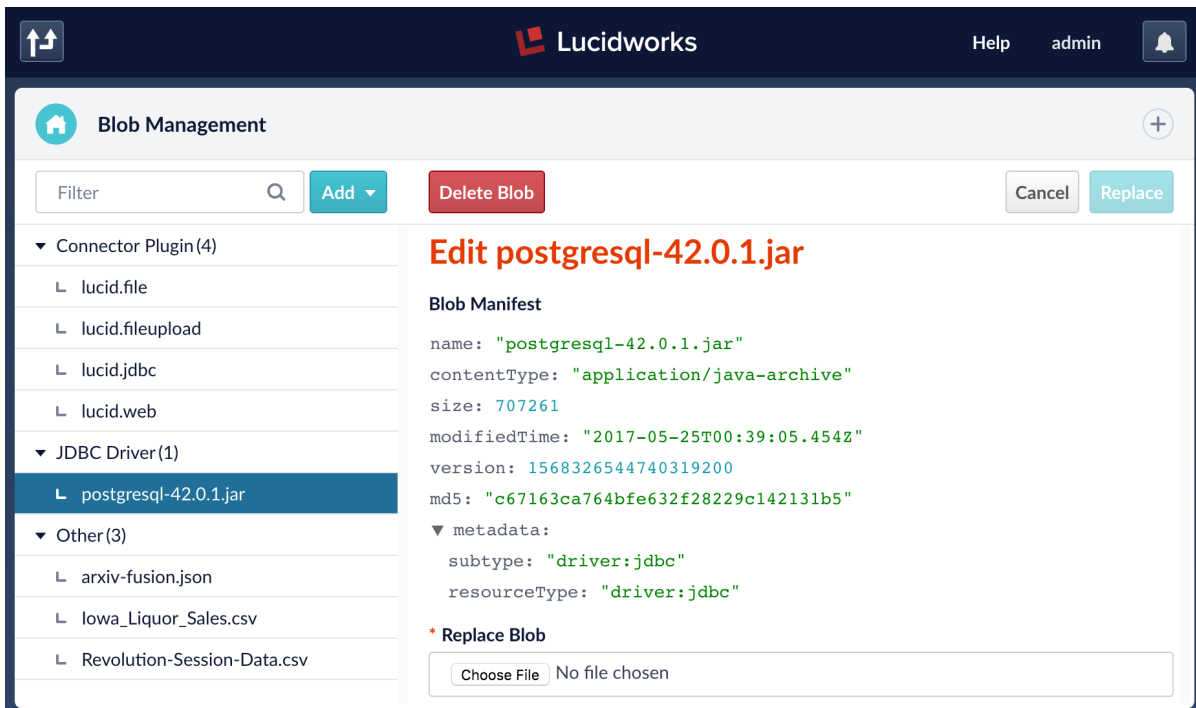
The Blob manager

In addition to the Blob Store API, the Fusion UI provides an interface to the blob store, at **DevOps > Blobs**.

- Click **Add** to upload a new blob:




- Select an uploaded blob to view, replace, or delete it:



Time-Based Partitioning

A Fusion collection can be configured to map to multiple Solr collections, known as partitions in this context, where each partition contains data from a specific time range. An example is time-based partitioning for logs:



Once a collection is configured for time-base partitioning, Fusion automatically ages out old partitions and creates new ones, using the configured partition sizes, expiration intervals, and so on. No manual maintenance is needed.

This feature is not enabled by default. Enable it for each collection using the Collection Features API.

Note	Fusion cannot retroactively partition data that has already been indexed. It can only perform time-based partitioning on incoming data.
------	---

Enabling time-based partitioning

- In the UI, you can only enable time-based partitioning for *new* collections.
- In the API, you can only enable time-based partitioning for *existing* collections.

Enablement using the Fusion UI

1. While creating a collection, click **Advanced**.
2. Scroll down to "Time Series Partitioning".
3. Click **Enable**.

4. Save the collection.

Currently, you cannot use the UI to enable time-based partitioning for an existing collection.

Enablement using the API

Use the Collection Features API to enable time-based partitioning for an existing collection.

Enable time-based partitioning using the default configuration:

```
curl -X PUT -H 'Content-type: application/json' -d '{"enabled": true}'
http://localhost:8765/api/v1/collections/<collection>/features/partitionByTime
```

No response is returned.

Submit an empty request to the same endpoint to verify that time-based partitioning is enabled:

```
curl -X GET http://localhost:8765/api/v1/collections/<collection>/features/partitionByTime
```

Response:

```
{
  "name" : "partitionByTime",
  "collectionId" : "<collection>",
  "params" : { },
  "enabled" : true
}
```

To change the configuration, see the options and examples below.

Configuration options

When time series indexing is enabled for a collection, you can configure these options using the UI or the Collections API. None are required.

UI Label, API Name	Description
Timestamp Field Name <code>timestampFieldName</code>	The name of the field from which to read timestamps. The default is "timestamp".
Partition Time Period <code>timePeriod</code>	The time range for each partition. The default is one day.
Max Active Partitions <code>maxActivePartitions</code>	The number of partitions to keep active.
Delete Expired Partitions <code>deleteExpired</code>	"True" to automatically delete partitions that fall outside of the <code>maxActivePartitions</code> window, at intervals of <code>scheduleIntervalMinutes</code> . The default is "false".
Preemptive Create Enabled <code>preemptiveCreateEnabled</code>	"True" (the default) to create partitions in advance.
Schedule Interval <code>scheduleIntervalMinutes</code>	The interval, in minutes, at which to perform background maintenance, including preemptively creating partitions (<code>preemptiveCreateEnabled</code>) and deleting expired partitions (<code>deleteExpired</code>). The default is five minutes.
Partition Num Shards <code>numShards</code>	The number of shards per partition. The default is the value configured for the main Fusion collection.

UI Label, API Name	Description
Partition Replication Factor <code>replicationFactor</code>	The number of copies to keep, per partition. The default is the value configured for the main Fusion collection.
Partition Config Name <code>configName</code>	The name of the Solr configuration set to be applied to new partitions; the default is the configuration used by the primary collection.

Examples

Create a new collection called "TimeSeries1":

```
curl -X PUT -H 'Content-type: application/json' -d '{
  "solrParams": {
    "numShards": 1,
    "replicationFactor": 1
  }
}' http://localhost:8765/api/v1/collections/TimeSeries1
```

Enable and configure time-based partitioning for the "TimeSeries1" collection:

```
curl -X PUT -H 'Content-type: application/json' -d '{
  "enabled": true,
  "timestampFieldName": "ts",
  "timePeriod": "5MINUTES",
  "scheduleIntervalMinutes": 1,
  "preemptiveCreateEnabled": false,
  "maxActivePartitions": 4,
  "deleteExpired": true
}' http://localhost:8765/api/v1/collections/TimeSeries1/features/partitionByTime
```

Verify that time-based partitioning is enabled:

```
curl -X GET http://localhost:8765/api/v1/collections/TimeSeries1/features/partitionByTime
```

Import some sample data into this collection:

```
curl -X POST -H "Content-type:application/vnd.lucidworks-document" -d '[
  {
    "id": "1",
    "fields": [
      {
        "name": "ts",
        "value": "2016-02-24T00:00:01Z"
      },
      {
        "name": "partition_s",
        "value": "eventsim_2016_02_24_00_00"
      }
    ]
  }
]
```

```

    }
  ]
},
{
  "id": "2",
  "fields": [
    {
      "name": "ts",
      "value": "2016-02-24T00:05:01Z"
    },
    {
      "name": "partition_s",
      "value": "eventsim_2016_02_24_00_05"
    }
  ]
},
{
  "id": "3",
  "fields": [
    {
      "name": "ts",
      "value": "2016-02-24T00:10:01Z"
    },
    {
      "name": "partition_s",
      "value": "eventsim_2016_02_24_00_10"
    }
  ]
},
{
  "id": "4",
  "fields": [
    {
      "name": "ts",
      "value": "2016-02-24T00:15:01Z"
    },
    {
      "name": "partition_s",
      "value": "eventsim_2016_02_24_00_15"
    }
  ]
},
{
  "id": "5",
  "fields": [
    {
      "name": "ts",
      "value": "2016-02-24T00:20:01Z"
    },
    {
      "name": "partition_s",
      "value": "eventsim_2016_02_24_00_20"
    }
  ]
}
]
}' http://localhost:8765/api/v1/index-pipelines/TimeSeries1-default/collections/TimeSeries1/index

```

View the Solr configuration for this collection:

```
curl -X GET "http://localhost:8765/api/v1/query-pipelines/TimeSeries1-  
default/collections/TimeSeries1/select?q=*:*"
```

The response includes a list of active Solr collections that correspond to this Fusion collection:

```
<str  
name="collection">TimeSeries1_2016_02_24_00_05,TimeSeries1_2016_02_24_00_10,TimeSeries1_2016_02_24_00_15,TimeS  
eries1_2016_02_24_00_20</str>
```

Custom JavaScript Stages For Index Pipelines

The JavaScript Index stage allows you to write a custom processing logic using JavaScript to manipulate Pipeline Documents and the index pipeline context. which will be compiled by the JDK into Java bytecode that is executed by the Fusion pipeline. The first time that the pipeline is run, Fusion compiles the JavaScript program into Java bytecode using the JDK's JavaScript engine.

For a JavaScript Index stage, the JavaScript code must return either: a single document or array of documents; or the null value or an empty array. In the latter case, no further processing is possible, which means that the document will not be indexed or updated.

JavaScript Index Stage Global Variables

[JavaScript](#) is a lightweight scripting language. The JavaScript in a JavaScript stage is standard [ECMAScript](#). What a JavaScript program can do depends on the container in which it runs. For a JavaScript Index stage, the container is a Fusion index pipeline. The following global pipeline variables are available:

Name	Type	Description
<code>doc</code>	PipelineDocument	The contents of each document submitted to the pipeline. See: PipelineDocument Objects for a complete description of this object.
<code>ctx</code>	Context	A reference to the container which holds a map over the pipeline properties. Used to update or modify this information for downstream pipeline stages.
<code>collection</code>	String	The name of the Fusion collection being indexed or queried.
<code>solrServer</code>	BufferingSolrServer	The Solr server instance that manages the pipeline's default Fusion collection. All indexing and query requests are done by calls to methods on this object. See SolrClient for details.
<code>solrServerFactory</code>	SolrClientFactory	The SolrCluster server used for lookups by collection name which returns a Solr server instance for a that collection, e.g. <pre>var productsSolr = solrServerFactory.getSolrServer("products");</pre>

Note	The now-deprecated global variable "_context" refers to the same object as "ctx".
------	---

The JavaScript in a JavaScript Index stage must return either a single document or an array of documents. This can be accomplished by either:

- a series of statements where the final statement evaluates to a document or array of documents
- a function which returns a document or an array of documents

As of Fusion 2.4, all pipeline variables referenced in the body of the JavaScript function must be passed in as arguments to the function. E.g., in order to access the PipelineDocument in global variable 'doc', the JavaScript function **must** be written as:

```
function doWork(doc) {
  // do some work ...
  return doc;
}
```

The allowed set of function declarations are:

```
function doWork(doc) { ... return doc; }
function doWork(doc, ctx) { ... return doc; }
function doWork(doc, ctx, collection) { ... return doc; }
function doWork(doc, ctx, collection, solrServer) { ... return doc; }
function doWork(doc, ctx, collection, solrServer, solrServerFactory) { ... return doc; }
```

The order of these arguments is according to the (estimated) frequency of use. The assumption is that most processing only requires access to the document object itself, and the next-most frequent type of processing requires only the document and read-only access of some context parameters. If you need to reference the solrServerFactory global variable, you must use the 5-arg function declaration.

In order to use other functions in your JavaScript program, you can define and use them, as long as the final statement in the program returns a document or documents.

Global variable `logger`

The global variable named `logger` writes messages to the logfile of the server running the pipeline. This variable is truly global and doesn't need to be declared as part of the function parameter list.

Since Fusion's connectors service does the index pipeline processing, these log messages go into the logfile: `fusion/3.1.x/var/log/connector/connector.log`. There are 5 methods available, which each take either a single argument (the string message to log) or two arguments (the string message and an exception to log). The five methods are, "debug", "info", "warn", and "error".

JavaScript Index Stage Examples

Add a field to a document

```
function (doc) {
  doc.addField('some-new-field', 'some-value');
  return doc;
}
```

Join two fields

The following example conjoins separate latitude and longitude fields into a single geo-coordinate field, whose field name follows Solr schema conventions and ends in "_p". It also removes the original latitude and longitude fields from the document.

```
function(doc) {
  var value = "";
  if (doc.hasField("myGeo_Lat") && doc.hasField("myGeo_Long")) {
    value = doc.getFirstFieldValue("myGeo_Lat") + "," + doc.getFirstFieldValue("myGeo_Long");
    doc.addField("myGeo_p", value);
    doc.removeFields("myGeo_Lat");
    doc.removeFields("myGeo_Long");
    logger.debug("conjoined Lat, Long: " + value);
  }
  return doc;
}
```

Return an array of documents

```
function (doc) {
  var subjects = doc.getFieldValues("subjects");
  var id = doc.getId();
  var newDocs = [];
  for (i = 0; i < subjects.size(); i++) {
    var pd = new com.lucidworks.apollo.common.pipeline.PipelineDocument(id+'-'+i );
    pd.addField('subject', subjects.get(i));
    newDocs.push( pd );
  }
  return newDocs;
}
```

Parse a JSON-escaped string into a JSON object

While it's simpler to use a JSON Parsing index stage, the following code example shows you how to parse a JSON-escaped string representation into a JSON object.

This code parses a JSON object into an array of attributes, and then find the attribute "tags" which has as its value a list of strings. Each item in the list is added to a multi-valued document field named "tag_ss".

```

var imports = new JavaImporter(Packages.sun.org.mozilla.javascript.internal.json.JsonParser);
function(doc) {
  with (imports) {
    myData = JSON.parse(doc.getFirstFieldValue('body'));
    logger.info("parsed object");
    for (var index in myData) {
      var entity = myData[index];
      if (index == "tags") {
        for (var i=0; i<entity.length;i++) {
          var tag = entity[i][0];
          doc.addField("tag_ss",tag);
        }
      }
    }
  }
  doc.removeFields("body");
  return doc;
}

```

Do a lookup on another Fusion collection

```

function doWork(doc, ctx, collection, solrServer, solrServerFactory) {
  var imports = new JavaImporter(
    org.apache.solr.client.solrj.SolrQuery,
    org.apache.solr.client.solrj.util.ClientUtils);
  with(imports) {
    var sku = doc.getFirstFieldValue("sku");
    if (!doc.hasField("mentions")) {
      var mentions = ""
      var productsSolr = solrServerFactory.getSolrServer("products");
      if( productsSolr != null ){
        var q = "sku:"+sku;
        var query = new SolrQuery();
        query.setRows(100);
        query.setQuery(q);
        var res = productsSolr.query(query);
        mentions = res.getResults().size();
        doc.addField("mentions",mentions);
      }
    }
  }
  return doc;
}

```

Reject a document

If the function returns `null` or an empty array, it will not be indexed or updated into Fusion.

```

function doWork(doc) {
  if (!doc.hasField("required_field")) {
    return null;
  }
  return doc;
}

```

Debugging and Troubleshooting

To debug a JavaScript Index stage you can:

- Check the Fusion api server logs for compilation errors.
- Check the Fusion connectors server logs for runtime processing errors.
- Use the `logger` object for print debugging (in the Fusion connectors logfile).
- Use the Pipeline Preview tool (not available in Fusion 2.0, 2.1, or 2.2).

The JavaScript Engine Used by Fusion

The JavaScript engine used by Fusion is the Nashorn engine from Oracle. See [The Nashorn Java API](#) for details.

Upgrading to the latest Nashorn engine

The default version of the Nashorn engine used by Fusion versions 2.4.1 and earlier is the `nashorn-0.1-jdk7.jar` which contains many bugs that have since been fixed in the official JDK 1.8 version. In order to use the latest version of the Nashorn engine, you must:

- Have an up-to-date version of Java 8 installed.
- Remove the `nashorn-0.1-jdk7.jar` from the Fusion classpaths:
 - `cd fusion/3.1.x`
 - `find . -name "nashorn-0.1-jdk7.jar" -print -exec rm -i {} \;`

Creating and accessing Java types

The following information is taken from Oracle's JavaScript programming guide section 3, [Using Java From Scripts](#).

To create script objects that access and reference Java types from Javascript use the `Java.type()` function:

```
var ArrayList = Java.type("java.util.ArrayList");
var a = new ArrayList;
```

2.1.3. Other Ingestion Methods

Usually, the simplest way to get data into Fusion is through its connectors. However, in some cases it makes sense to use other methods:

- Importing directly into Solr

Fusion can read any data that has been imported into its Solr core. If you're already a proficient Solr user, you may already have mechanisms in place for doing this. You can continue importing your data directly into Solr, then use Fusion to read and manage it.

The Fusion package also includes tools for direct import into Solr:

- a Hive Serializer/Deserializer
- a set of Pig functions
- Pushing to an index pipeline

You can bypass the connectors and parsers to push documents directly to an index pipeline, using the Index-Pipelines REST API. You'll need to understand Fusion PipelineDocument Objects first.

Importing Data with Pig

You can use Pig to import data into Fusion, using the `lucidworks-pig-functions-2.2.6.jar` file found in `fusion/3.1.x/apps/connectors/resources/lucid.hadoop/jobs`.

Available Functions

The Pig functions included in the `lucidworks-pig-functions-2.2.6.jar` are three UserDefined Functions (UDF) and two Store functions. These functions are:

- `com/lucidworks/hadoop/pig/SolrStoreFunc.class`
- `com/lucidworks/hadoop/pig/FusionIndexPipelinesStoreFunc.class`
- `com/lucidworks/hadoop/pig/EpochToCalendar.class`
- `com/lucidworks/hadoop/pig/Extract.class`
- `com/lucidworks/hadoop/pig/Histogram.class`

Using The Functions

Register the Functions

There are two approaches to using functions in Pig: **REGISTER** them in the script, or load them with your Pig command line request.

If using **REGISTER**, the Pig function jars must be put in HDFS in order to be used by your Pig script. It can be located anywhere in HDFS; you can either supply the path in your script or use a variable and define the variable with `-p` property definition.

The example below uses the second approach, loading the jars with the `-Dpig.additional.jars` system property when launching the script. With this approach, the jars can be located anywhere on the machine where the script will be run.

Indexing Data to Fusion

When indexing data to Fusion, there are several parameters to pass with your script in order to output data to Fusion for indexing.

These parameters can be made into variables in the script, with the proper values passed on the command line when the script is initiated. The example script below shows how to do this for Solr. The theory is the same for Fusion, only the parameter names would change as appropriate:

`fusion.endpoints`

The full URL to the index pipeline in Fusion. The URL should include the pipeline name and the collection data will be indexed to.

`fusion.fail.on.error`

If `true`, when an error is encountered, such as if a row could not be parsed, indexing will stop. This is `false` by default.

`fusion.buffer.timeoutms`

The amount of time, in milliseconds, to buffer documents before sending them to Fusion. The default is 1000. Documents will be sent to Fusion when either this value or `fusion.batchSize` is met.

`fusion.batchSize`

The number of documents to batch before sending the batch to Fusion. The default is 500. Documents will be sent to

Fusion when either this value or `fusion.buffer.timeoutms` is met.

`fusion.realm`

This is used with `fusion.user` and `fusion.password` to authenticate to Fusion for indexing data. Two options are supported, `KERBEROS` or `NATIVE`. + Kerberos authentication is supported with the additional definition of a JAAS file. The properties `java.security.auth.login.config` and `fusion.jaas.appname` are used to define the location of the JAAS file and the section of the file to use. These are described in more detail below. + Native authentication uses a Fusion-defined username and password. This user must exist in Fusion, and have the proper permissions to index documents.

`fusion.user`

The Fusion username or Kerberos principal to use for authentication to Fusion. + If a Fusion username is used (`'fusion.realm' = 'NATIVE'`), the `fusion.password` must also be supplied.

`fusion.pass`

This property is not shown in the example above. The password for the `fusion.user` when the `fusion.realm` is `NATIVE`.

Indexing to a Kerberized Fusion Installation

When Fusion is secured with Kerberos, Pig scripts must include the full path to a JAAS file that includes the service principal and the path to a keytab file that will be used to index the output of the script to Fusion.

Additionally, a Kerberos ticket must be obtained on the server for the principal using `kinit`.

`java.security.auth.login.config`

This property defines the path to a JAAS file that contains a service principal and keytab location for a user who is authorized to write to Fusion. + The JAAS configuration file **must** be copied to the same path on every node where a Node Manager is running (i.e., every node where map/reduce tasks are executed). Here is a sample section of a JAAS file: +

```
Client { (1)
  com.sun.security.auth.module.Krb5LoginModule required
  useKeyTab=true
  keyTab="/data/fusion-indexer.keytab" (2)
  storeKey=true
  useTicketCache=true
  debug=true
  principal="fusion-indexer@FUSIONSERVER.COM"; (3)
};
```

- +
 1. The name of this section of the JAAS file. This name will be used with the `'fusion.jaas.appname'` parameter.
 2. The location of the keytab file.
 3. The service principal name. This should be a different principal than the one used for Fusion, but must have access to both Fusion and Pig. This name is used with the `'fusion.user'` parameter described above.
- `'fusion.jaas.appname'::`
Used only when indexing to or reading from Fusion when it is secured with Kerberos.
- +

This property provides the name of the section in the JAAS file that includes the correct service principal and keytab path.

Sample CSV Script

The following Pig script will take a simple CSV file and index it to Solr.

```
set solr.zkhost '$zkHost';
set solr.collection '$collection'; (1)

A = load '$csv' using PigStorage(',') as
(id_s:chararray,city_s:chararray,country_s:chararray,code_s:chararray,code2_s:chararray,latitude_s:chararray,lon
gitude_s:chararray,flag_s:chararray); (2)
--dump A;
B = FOREACH A GENERATE $0 as id, 'city_s', $1, 'country_s', $2, 'code_s', $3, 'code2_s', $4, 'latitude_s', $5,
'longitude_s', $6, 'flag_s', $7; (3)

ok = store B into 'SOLR' using com.lucidworks.hadoop.pig.SolrStoreFunc(); (4)
```

This relatively simple script is doing several things that help to understand how the Solr Pig functions work.

1. This and the line above define parameters that are needed by `SolrStoreFunc` to know where Solr is. `SolrStoreFunc` needs the properties `solr.zkhost` and `solr.collection`, and these lines are mapping the `zkhost` and `collection` parameters we will pass when invoking Pig to the required properties.
2. Load the CSV file, the path and name we will pass with the `csv` parameter. We also define the field names for each column in CSV file, and their types.
3. For each item in the CSV file, generate a document id from the first field (`$0`) and then define each field name and value in `name, value` pairs.
4. Load the documents into Solr, using the `SolrStoreFunc`. While we don't need to define the location of Solr here, the function will use the `zkhost` and `collection` properties that we will pass when we invoke our Pig script.

Warning

When using `SolrStoreFunc`, the document ID **must** be the first field.

When we want to run this script, we invoke Pig and define several parameters we have referenced in the script with the `-p` option, such as in this command:

```
./bin/pig -Dpig.additional.jars=/path/to/lucidworks-pig-functions-2.2.6.jar -p
csv=/path/to/my/csv/airports.dat -p zkHost=zknode1:2181,zknode2:2181,zknode3:2181/solr -p
collection=myCollection ~/myScripts/index-csv.pig
```

The parameters to pass are:

csv

The path and name of the CSV file we want to process.

zkhost

The ZooKeeper connection string for a SolrCloud cluster, in the form of `zkhost1:port,zkhost2:port,zkhost3:port/chroot`. In the script, we mapped this to the `solr.zkhost` property, which is required by the `SolrStoreFunc` to know where to send the output documents.

collection

The Solr collection to index into. In the script, we mapped this to the `solr.collection` property, which is required by

the `SolrStoreFunc` to know the Solr collection the documents should be indexed to.

Tip

The `zkhost` parameter above is only used if you are indexing to a SolrCloud cluster, which uses ZooKeeper to route indexing and query requests.

If, however, you are not using SolrCloud, you can use the `solrUrl` parameter, which takes the location of a standalone Solr instance, in the form of `http://host:port/solr`.

In the script, you would change the line that maps `solr.zkhost` to the `zkhost` property to map `solr.server.url` to the `solrUrl` property. For example:

```
`set solr.server.url '$solrUrl';`
```

How to Contribute

1. Fork this repo i.e. `<username|organization>/hadoop-solr`, following the [fork a repo](#) tutorial. Then, clone the forked repo on your local machine:

```
$ git clone https://github.com/<username|organization>/hadoop-solr.git
```

2. Configure remotes with the [configuring remotes](#) tutorial.
3. Create a new branch:

```
$ git checkout -b new_branch  
$ git push origin new_branch
```

Use the [creating branches](#) tutorial to create the branch from GitHub UI if you prefer.

4. Develop on `new_branch` branch only, **do not merge `new_branch` to your master**. Commit changes to `new_branch` as often as you like:

```
$ git add <filename>  
$ git commit -m 'commit message'
```

5. Push your changes to GitHub.

```
$ git push origin new_branch
```

6. Repeat the commit & push steps until your development is complete.
7. Before submitting a pull request, fetch upstream changes that were done by other contributors:

```
$ git fetch upstream
```

8. And update master locally:

```
$ git checkout master  
$ git pull upstream master
```

9. Merge master branch into `new_branch` in order to avoid conflicts:

```
$ git checkout new_branch  
$ git merge master
```

10. If conflicts happen, use the [resolving merge conflicts](#) tutorial to fix them:

11. Push master changes to `new_branch` branch

```
$ git push origin new_branch
```

12. Add jUnits, as appropriate, to test your changes.

13. When all testing is done, use the [create a pull request](#) tutorial to submit your change to the repo.

Note

Please be sure that your pull request sends only your changes, and no others. Check it using the command:

```
git diff new_branch upstream/master
```

Importing Data with Hive

Fusion 3.1 has 4 different versions of the Hadoop connector, and they are no longer included in the out-of-box Fusion. You must download them from [the connectors downloads table](#). Notice how there are 4 different connectors now:

- hadoop-apache2
- hadoop-cloudera
- hadoop-hortonworks
- hadoop-mapR

Once you have taken the plugin zip file(s) you desire and installed the connector(s) into Fusion, you will now find the Serializer/Deserializer (SerDe) for Hive located at `lucidworks-hive-serde-2.2.7.jar` in `fusion/3.1.x/apps/connectors/plugins/lucid.hadoop-SUBTYPE/assets`, where `SUBTYPE` is either `apache2`, `cloudera`, `hortonworks` or `mapR` depending on what one you want to use.

Features

- Index Hive table data to Solr.
- Read Solr index data to a Hive table.
- Kerberos support for securing communication between Hive and Solr.
- As of v2.2.4 of the SerDe, integration with [Lucidworks Fusion](#) is supported.
 - Fusion's index pipelines can be used to index data to Fusion.
 - Fusion's query pipelines can be used to query Fusion's Solr instance for data to insert into a Hive table.

Install the SerDe Jar to Hive

In order for Hive to work with Solr, the Hive SerDe jar must be added as a plugin to Hive.

From a Hive prompt, use the `ADD JAR` command and reference the path and filename of the SerDe jar for your Hive version.

```
hive> ADD JAR lucidworks-hive-serde-2.2.7.jar;
```

This can also be done in your Hive command to create the table, as in the example below.

Indexing Data to Fusion

If you use Lucidworks Fusion, you can index data from Hive to Solr via Fusion's index pipelines. These pipelines allow you several options for further transforming your data.

Tip	<p>If you are using Fusion v3.0.x, you already have the Hive SerDe in Fusion's <code>./apps/connectors/resources/lucid.hadoop/jobs</code> directory. The SerDe jar that supports Fusion is v2.2.4 or higher. This was released with Fusion 3.0.</p> <p>If you are using Fusion 3.1.x, you will need to download the Hive SerDe from http://lucidworks.com/connectors/. Choose the proper Hadoop distribution and the resulting .zip file will include the Hive SerDe.</p> <p>A 2.2.4 or higher jar built from this repository will also work with Fusion 2.4.x releases.</p>
-----	--

This is an example Hive command to create an external table to index documents in Fusion and to query the table later.

```
hive> CREATE EXTERNAL TABLE fusion (id string, field1 string, field2 int)
  STORED BY 'com.lucidworks.hadoop.hive.FusionStorageHandler'
  LOCATION '/tmp/fusion'
  TBLPROPERTIES('fusion.endpoints' = 'http://localhost:8764/api/apollo/index-
pipelines/<pipeline>/collections/<collection>/index',
                'fusion.fail.on.error' = 'false',
                'fusion.buffer.timeoutms' = '1000',
                'fusion.batchSize' = '500',
                'fusion.realm' = 'KERBEROS',
                'fusion.user' = 'fusion-indexer@FUSIONSERVER.COM',
                'java.security.auth.login.config' = '/path/to/JAAS/file',
                'fusion.jaas.appname' = 'FusionClient',
                'fusion.query.endpoints' = 'http://localhost:8764/api/apollo/query-pipelines/pipeline-
id/collections/collection-id',
                'fusion.query' = '*:*');
```

In this example, we have created an external table named "fusion", and defined a custom storage handler (`STORED BY 'com.lucidworks.hadoop.hive.FusionStorageHandler'`).

The `LOCATION` indicates the location in HDFS where the table data will be stored. In this example, we have chosen to use `/tmp/fusion`.

In the section `TBLPROPERTIES`, we define several properties for Fusion so the data can be indexed to the right Fusion installation and collection:

`fusion.endpoints`

The full URL to the index pipeline in Fusion. The URL should include the pipeline name and the collection data will be indexed to.

`fusion.fail.on.error`

If `true`, when an error is encountered, such as if a row could not be parsed, indexing will stop. This is `false` by default.

`fusion.buffer.timeoutms`

The amount of time, in milliseconds, to buffer documents before sending them to Fusion. The default is 1000.

Documents will be sent to Fusion when either this value or `fusion.batchSize` is met.

`fusion.batchSize`

The number of documents to batch before sending the batch to Fusion. The default is 500. Documents will be sent to Fusion when either this value or `fusion.buffer.timeoutms` is met.

`fusion.realm`

This is used with `fusion.user` and `fusion.password` to authenticate to Fusion for indexing data. Two options are supported, `KERBEROS` or `NATIVE`. + Kerberos authentication is supported with the additional definition of a JAAS file. The properties `java.security.auth.login.config` and `fusion.jaas.appname` are used to define the location of the JAAS file and the section of the file to use. + Native authentication uses a Fusion-defined username and password. This user must exist in Fusion, and have the proper permissions to index documents.

`fusion.user`

The Fusion username or Kerberos principal to use for authentication to Fusion. If a Fusion username is used (`'fusion.realm' = 'NATIVE'`), the `fusion.password` must also be supplied.

`fusion.password`

This property is not shown in the example above. The password for the `fusion.user` when the `fusion.realm` is `NATIVE`.

`java.security.auth.login.config`

This property defines the path to a JAAS file that contains a service principal and keytab location for a user who is authorized to read from and write to Fusion and Hive. + The JAAS configuration file **must** be copied to the same path on every node where a Node Manager is running (i.e., every node where map/reduce tasks are executed). Here is a sample section of a JAAS file: +

```
Client { (1)
  com.sun.security.auth.module.Krb5LoginModule required
  useKeyTab=true
  keyTab="/data/fusion-indexer.keytab" (2)
  storeKey=true
  useTicketCache=true
  debug=true
  principal="fusion-indexer@FUSIONSERVER.COM"; (3)
};
```

+
 1. The name of this section of the JAAS file. This name will be used with the `'fusion.jaas.appname'` parameter.
 2. The location of the keytab file.
 3. The service principal name. This should be a different principal than the one used for Fusion, but must have access to both Fusion and Hive. This name is used with the `'fusion.user'` parameter described above.

`'fusion.jaas.appname'::`
 Used only when indexing to or reading from Fusion when it is secured with Kerberos.

+
 This property provides the name of the section in the JAAS file that includes the correct service principal and keytab path.

`'fusion.query.endpoints'::`
 The full URL to a query pipeline in Fusion. The URL should include the pipeline name and the collection data will be read from. You should also specify the request handler to be used.

+
 If you do not intend to query your Fusion data from Hive, you can skip this parameter.

`'fusion.query'::`
 The query to run in Fusion to select records to be read into Hive. This is `'*:*'` by default, which selects all records in the index.

+
 If you do not intend to query your Fusion data from Hive, you can skip this parameter.

Query and Insert Data to Hive

Once the table is configured, any syntactically correct Hive query will be able to query the index.

For example, to select three fields named "id", "field1", and "field2" from the "solr" table, you would use a query such as:

```
hive> SELECT id, field1, field2 FROM solr;
```

Replace the table name as appropriate to use this example with your data.

To join data from tables, you can make a request such as:

```
hive> SELECT id, field1, field2 FROM solr left
      JOIN sometable right
      WHERE left.id = right.id;
```

And finally, to insert data to a table, simply use the Solr table as the target for the Hive INSERT statement, such as:

```
hive> INSERT INTO solr
      SELECT id, field1, field2 FROM sometable;
```

Example Indexing Hive to Solr

Solr includes a small number of sample documents for use when getting started. One of these is a CSV file containing book metadata. This file is found in your Solr installation, at `$SOLR_HOME/example/exampledocs/books.csv`.

Using the sample `books.csv` file, we can see a detailed example of creating a table, loading data to it, and indexing that data to Solr.

```

CREATE TABLE books (id STRING, cat STRING, title STRING, price FLOAT, in_stock BOOLEAN, author STRING, series
STRING, seq INT, genre STRING) ROW FORMAT DELIMITED FIELDS TERMINATED BY ','; (1)

LOAD DATA LOCAL INPATH '/solr/example/exampledocs/books.csv' OVERWRITE INTO TABLE books; (2)

ADD JAR {packageUser}-hive-serde-{connectorVersion}.jar; (3)

CREATE EXTERNAL TABLE solr (id STRING, cat_s STRING, title_s STRING, price_f FLOAT, in_stock_b BOOLEAN,
author_s STRING, series_s STRING, seq_i INT, genre_s STRING) (4)
  STORED BY 'com.lucidworks.hadoop.hive.LWStorageHandler' (5)
  LOCATION '/tmp/solr' (6)
  TBLPROPERTIES('solr.zkhost' = 'zknode1:2181,zknode2:2181,zknode3:2181/solr',
                'solr.collection' = 'gettingstarted',
                'solr.query' = '*:*'), (7)
                'lww.jaas.file' = '/data/jaas-client.conf'); (8)

INSERT OVERWRITE TABLE solr SELECT b.* FROM books b;

```

1. Define the table `books`, and provide the field names and field types that will make up the table.
2. Load the data from the `books.csv` file.
3. Add the `lucidworks-hive-serde-2.2.7.jar` file to Hive. Note the jar name shown here omits the version information which will be included in the jar file you have. If you are using Hive 0.13, you must also use a jar specifically built for 0.13.
4. Create an external table named `solr`, and provide the field names and field types that will make up the table. These will be the same field names as in your local Hive table, so we can index all of the same data to Solr.
5. Define the custom storage handler provided by the `lucidworks-hive-serde-2.2.7.jar`.
6. Define storage location in HDFS.
7. The query to run in Solr to read records from Solr for use in Hive.
8. Define the location of Solr (or ZooKeeper if using SolrCloud), the collection in Solr to index the data to, and the query to use when reading the table. This example also refers to a JAAS configuration file that will be used to authenticate to the Kerberized Solr cluster.

Pushing Documents to a Pipeline

Documents can be sent directly to an index pipeline using the Index-Pipelines REST API. The request path is:

```
/api/apollo/index-pipelines/<id>/collections/<collectionName>/index
```

where *<id>* is the name of a specific pipeline and *<collectionName>* is the name of a specific collection.

These requests are sent as a POST request. The request header specifies the format of the contents of the request body.

To send in a streaming list of JSON documents to the index pipeline you can send the JSON file which holds these objects to the API listed above with "application/json" as the content type. If your JSON file is a list/array of many items, the pipeline will operate in a streaming way and index the docs as necessary.

Example:

The JSON document called "myJsonDoc.json" holds 4.3M entries. Send the document to the index pipeline with the following command:

```
curl -u user:password -X POST -H 'Content-Type: application/json' -d@myJsonDoc.json  
"http://localhost:8764/api/apollo/index-pipelines/<id>/collections/<collectionName>/index"
```

Documents can be created on the fly using the PipelineDocument JSON notation. See Fusion PipelineDocument Objects for details and an example of how to do this.

Indexing PDFs and MS Office Documents

If you can access the filesystem in which the PDFs or MS Office documents reside, you can index these documents using properly configured datasource with the appropriate connector for that filesystem type. See Connectors and Datasources Reference for a list of all Fusion connectors.

If, however, there are obstacles to using the connectors, it may be simpler to index these types of documents with an index pipeline. The pipelines can only be used with REST API calls, and there is complete documentation in the section Index Pipelines API.

When sending the documents, it's important to set the content type header properly for the content being sent. This is not a complete list, but here are some frequently used content types:

- PDF documents: `application/pdf`
- MS Office:
 - .docx: `application/vnd.openxmlformats-officedocument.wordprocessingml.document`
 - .xlsx: `application/vnd.openxmlformats-officedocument.spreadsheetml.sheet``
 - .pptx: `application/vnd.vnd.openxmlformats-officedocument.presentationml.presentation``
 - More types: http://filext.com/faq/office_mime_types.php
- Text: `text/json`, `text/xml`, `text/csv`, etc.

Examples

Index a PDF document through the 'conn_solr' index pipeline to a collection named 'docs'. The pre-configured 'conn_solr' pipeline includes stages to parse documents with Tika, map fields, and index the documents to Solr (in that

order).

```
curl -u user:pass -X POST -H "Content-Type: application/pdf" --data-binary @/solr/core/src/test-files/mailling_lists.pdf http://localhost:8764/api/apollo/index-pipelines/conn_solr/collections/docs/index
```

Index one of the example Solr XML documents found in the `./example/exampledocs` directory of Solr. For this example to work well, the default pipeline was modified to include a field mapping stage in addition to indexing the documents to Solr. In this example, the custom pipeline is named 'docs-default' and the collection is 'docs'.

```
curl -u user:pass -X POST -H "Content-Type: text/xml" --data-binary @/Applications/solr-4.10.0/example/exampledocs/hd.xml http://localhost:8764/api/apollo/index-pipelines/docs-default/collections/docs/index
```

Indexing CSV Files

In the usual case, to index a CSV (or TSV) file, the file is split into records, one per row, and each row is indexed as a separate document. Datasources which use crawlers that are based on either the lucid.anda or lucid.fs framework can do the CSV splitting as part of the connector process.

Alternatively, the index pipeline can include a CSV Parsing stage.

2.2. Query Pipelines

A Query Pipeline transforms a set of inputs into a Solr query request and it can execute requests and manipulate the Solr response as well, via a set of modularized operations called Query Stages. The objects sent from stage to stage are Request objects and Response objects.

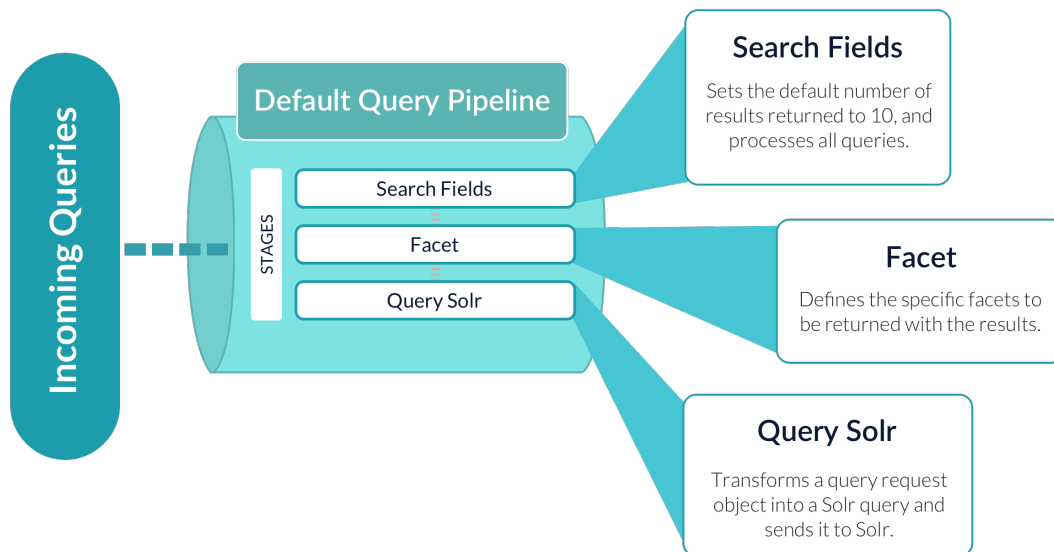
Fusion stores pipeline names and definitions, allowing a pipeline to be reused across applications. Pipeline definitions can be modified, so that as an application evolves, the pipelines used by that application can evolve accordingly. During application development, the Fusion UI can be used to develop and debug a Query Pipeline.

The available stage types allow setting specific parameters for the query, such as the number of results to return or the query parser to use. You can also define facets and recommendations to be returned with the results. If Access Control Lists (ACLs) are in use, you can apply a security-trimming stage to apply user access restrictions to the results.

For details about the available REST APIs, see [Query Pipelines API](#) and [Query Stages API](#).

2.2.1. Default Query Pipelines

Fusion ships with one default query pipeline named 'default'.



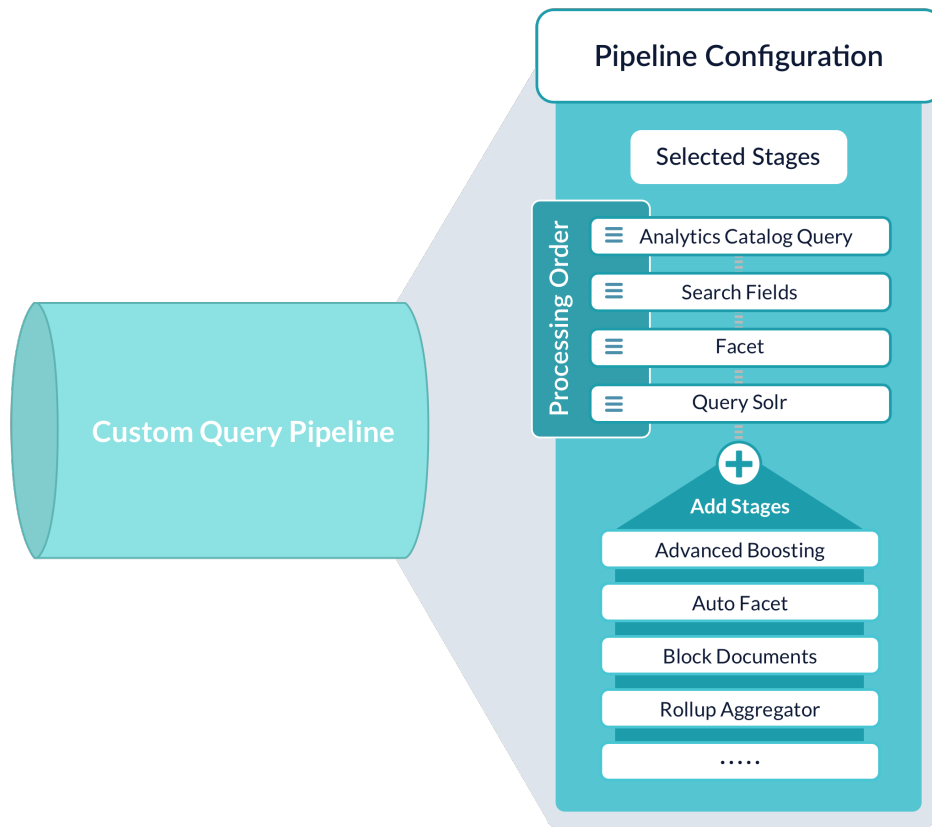
This pipeline has the following pre-configured stages:

- A Boost with Signals stage which boosts based on signals (if signals are collected and sent to the app).
- A Query Fields stage which sets the default number of results returned to 10. The property "skip" is false and the property "condition" is empty, so that all queries are processed.
- A Facets stage. No facet fields are specified. As configured, this stage has no effect on the pipeline.
- A Solr Query stage which sends the full query request to Solr.

When a collection is created, a default query pipeline for the collection is created. The pipeline name is the collection name with '-default' appended. For example, collection *foo* will have default query pipeline *foo-default*. This pipeline has the same configuration as the pipeline named 'default'.

2.2.2. Custom Query Pipelines

Using the Query Workbench or the REST API, you can develop custom pipelines to suit any search application. Start with any of Fusion's built-in query pipelines, then add, remove, and re-order the pipeline stages as needed to produce the appropriate query results.

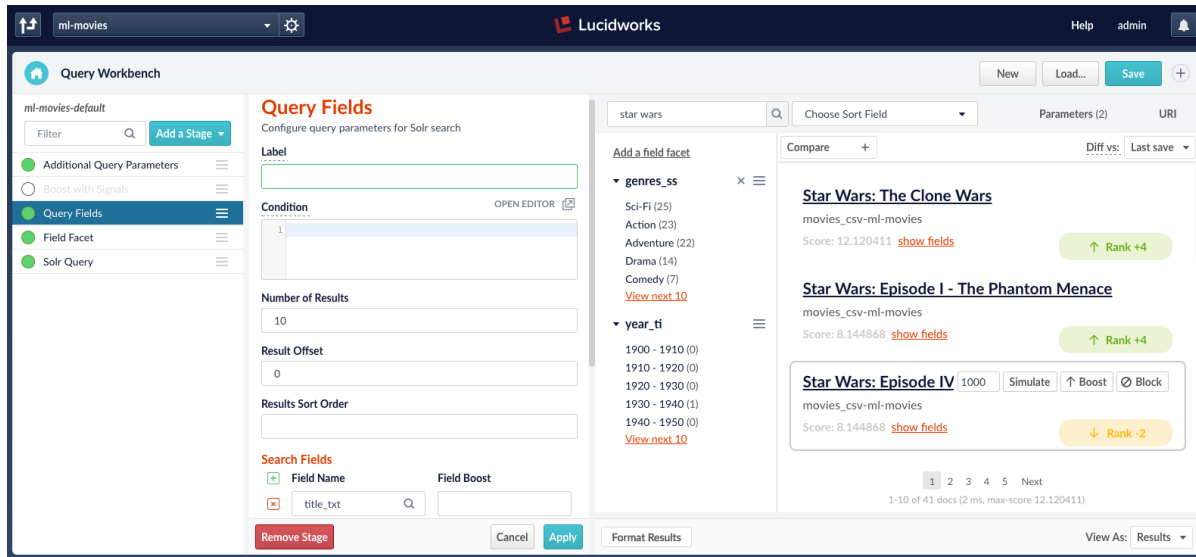


2.2.3. Query Workbench

The Query Workbench is where you edit query pipeline stages, results parameters, and more, then preview the impacts of the changes in real time before saving them. Relevancy tools such as boosting and blocking are accessible the Query Workbench and can be altered directly from the Workbench. The Query Workbench streamlines the process of querying indexed results and fine-tuning Query Pipelines to surface the results that best satisfy the end user's needs.

Beginning in Fusion 3.0, the Query Workbench replaces the Search UI that was in earlier versions.

See the Query Language Cheat Sheet for help constructing queries.



Select or re-order the stages in your query pipeline.

Configure the selected query pipeline stage.

Configure faceting.


Preview your search results, boost/block results, and view fields.

With this sandbox paradigm, you can experiment with query pipelines without consequence because Fusion uses a copy of a pipeline to simulate the new results, and it is not permanently modified in the collection until you click **Save**.


Query pipelines and pipeline stages

Query pipelines work by processing search requests and returning an ordered list of matching documents. Each pipeline consists of a series of query stages that can be added, ordered, and modified using the Query Workbench.

Select any stage in the pipeline to open its configuration panel:



Adding a pipeline stage

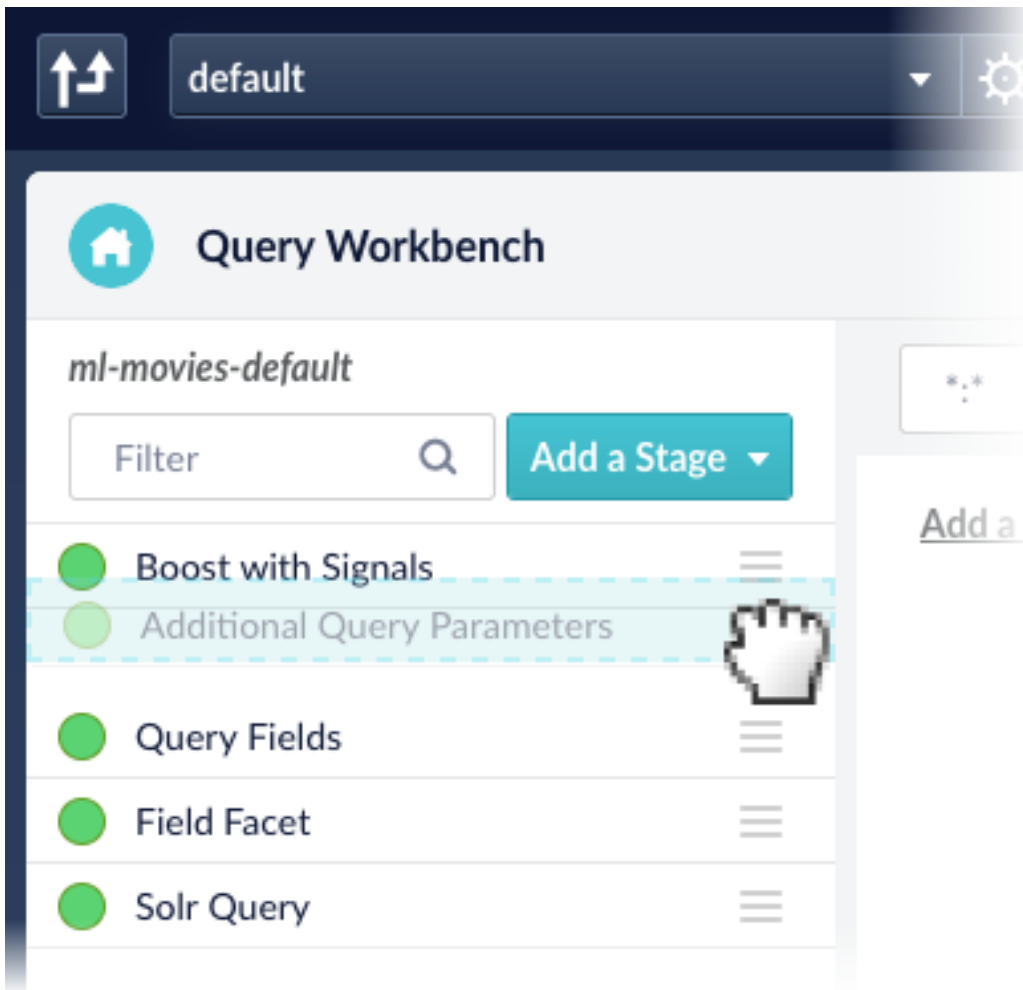


Click **Add a**

Stage to add query pipeline stages that can perform query setup, results relevancy, troubleshooting, and more.

Each pipeline stage definition associates a unique ID with a set of properties. The Solr Query stage is the only pipeline stage that is required for querying processes to complete, and therefore exists in every query pipeline. It is always the last stage in a series.

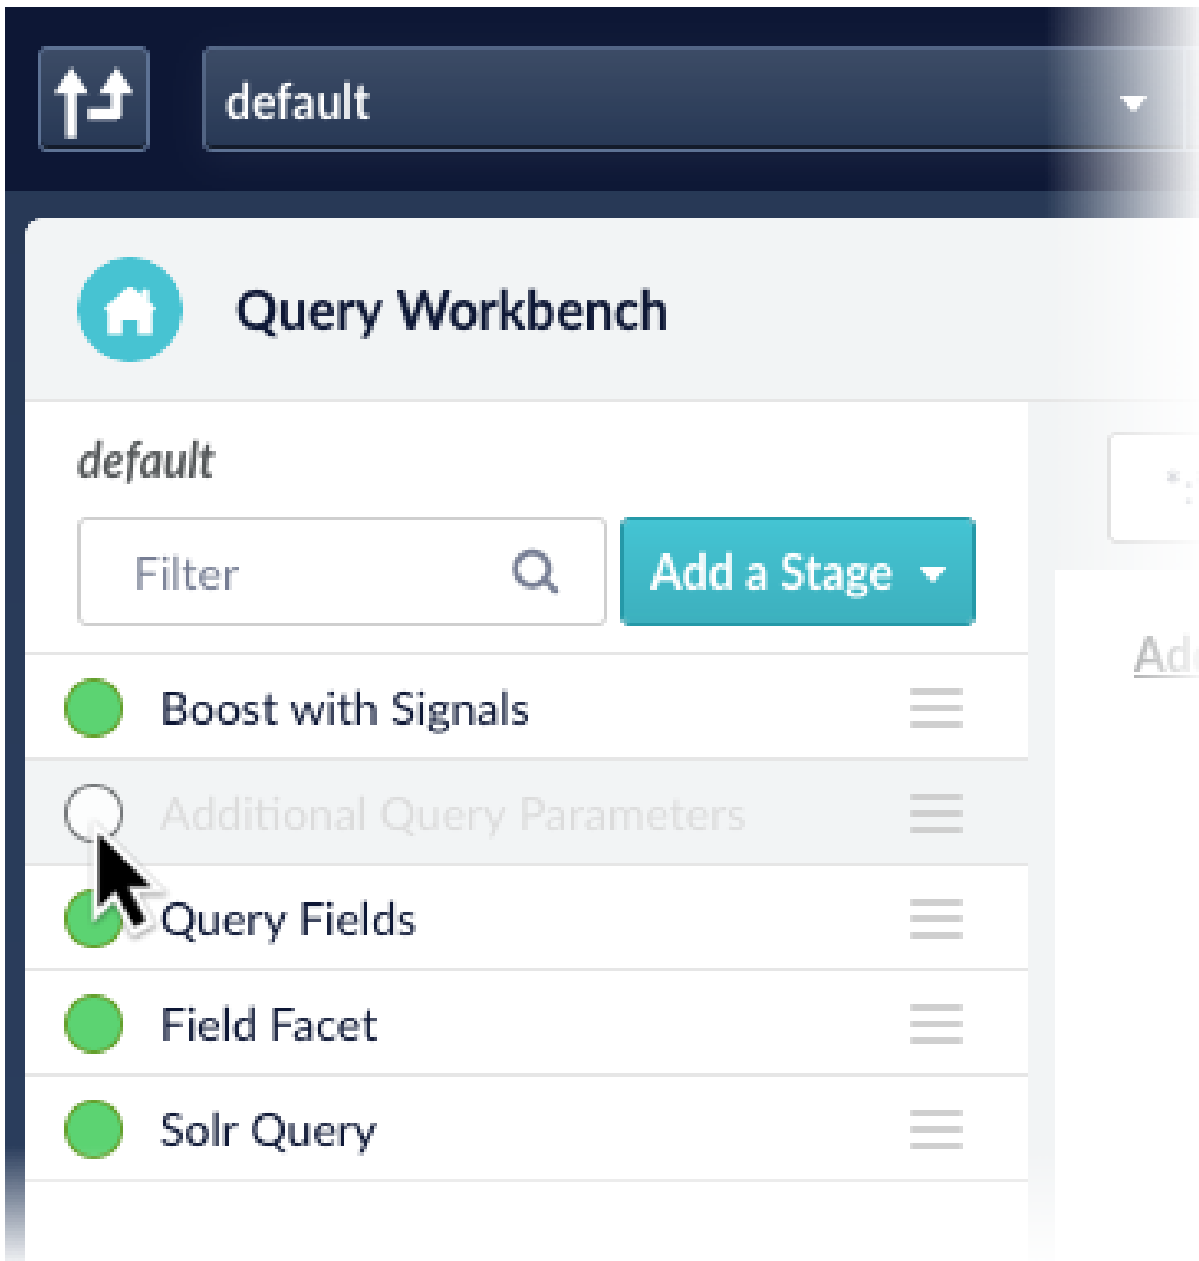
Re-ordering pipeline stages



The order of the pipeline stages matters, because the output from one stage becomes the input to the next stage. For example, the `width=600%` must always come last in the sequence, so that data is indexed only after it has been processed by all other stages. Putting this stage first in the sequence means that subsequent stages have no effect on the indexed data.

Drag any stage in the pipeline to move it up or down in the sequence of stages. The preview panel automatically updates the search results to reflect the output of the new sequence.

Enabling and disabling pipeline stages



By default, every stage in a query pipeline is enabled. While working with a query pipeline, it can be helpful to disable stages without removing them completely. This allows you to preserve a stage's configuration while observing how the search results change in its absence. You can re-enable the stage at any time. When you save a query pipeline, the enabled/disabled state of each stage is also saved.

Click the circle next to any stage in order to enable or disable it.

Boosting and blocking

As you search your data and inspect the results, you can manipulate the rankings of individual documents. Boosting a document raises its ranking, while blocking a document removes it from search results.

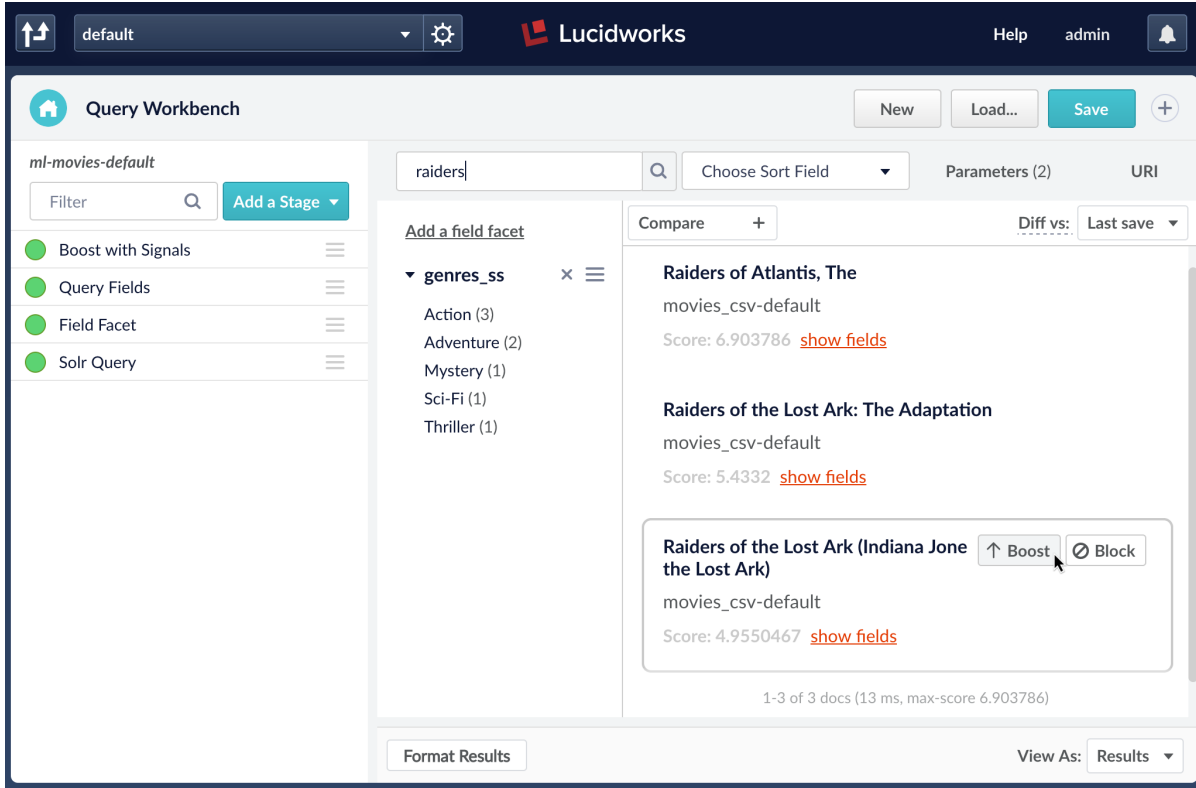
Boosting and blocking affects the results for specific queries. For example, when you search for "citizen" and boost the document for "Citizen Kane", your end users will see that title boosted to the top when they also search for "citizen", but not necessarily when they search for "kane".

The Query Workbench provides convenient buttons for instantly blocking or boosting the documents that you see in the

preview panel. Clicking **Boost** automatically adds the Boost Documents stage to your query pipeline, with a boost rule that matches your query and the document you clicked. Likewise, clicking **Block** adds and configures the Block Documents stage.


To boost a document

1. In the preview panel, hover over the document you want to boost.




2. Click **Boost**.

The preview panel automatically updates the rankings of the search results, and tags the differences:




Notice that the Boost Documents stage is now in our pipeline. You can click this stage to view the boost rules you've added. This is also where you can remove boost rules:




To block a document

1. In the preview panel, hover over the document you want to block.

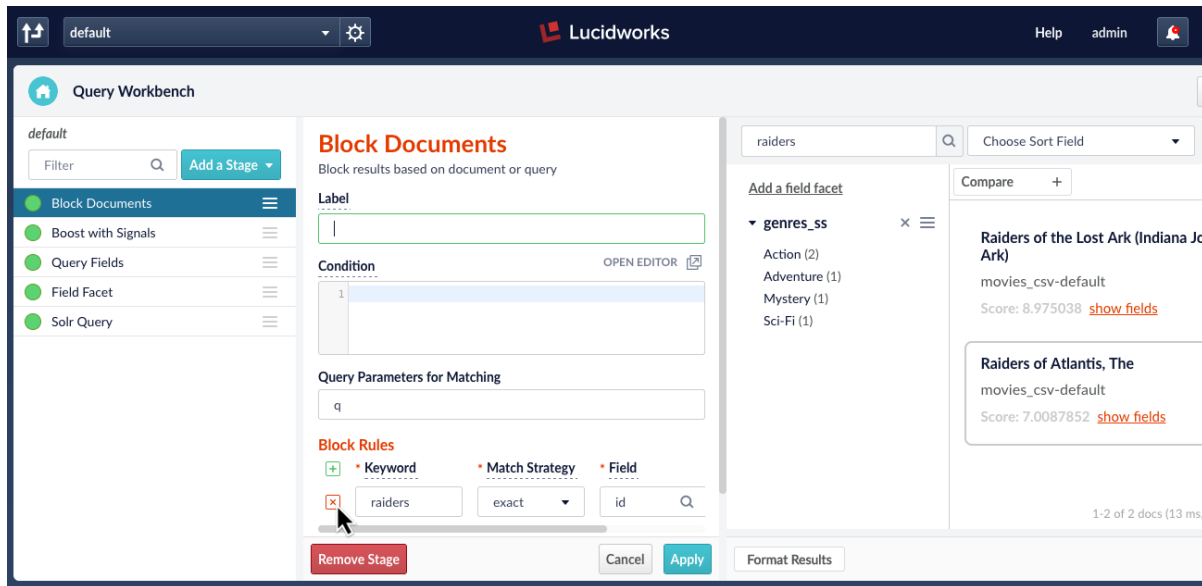


2. Click **Block**

The preview panel automatically updates the rankings of the search results, and tags the differences.



Notice that the Block Documents stage is now in our pipeline. You can click this stage to view the block rules you've added. This is also where you can remove block rules:



Compare mode

Compare mode allows you to compare search results using two different query pipelines. On the right is the working pipeline, which you can edit. On the left is another pipeline, selected from all existing pipelines.

search results side by side while modifying pipelines and boosting/blocking results in real time. In compare mode, you can view the differences between search results from the original query pipeline and the copy being modified in the Query Workbench.

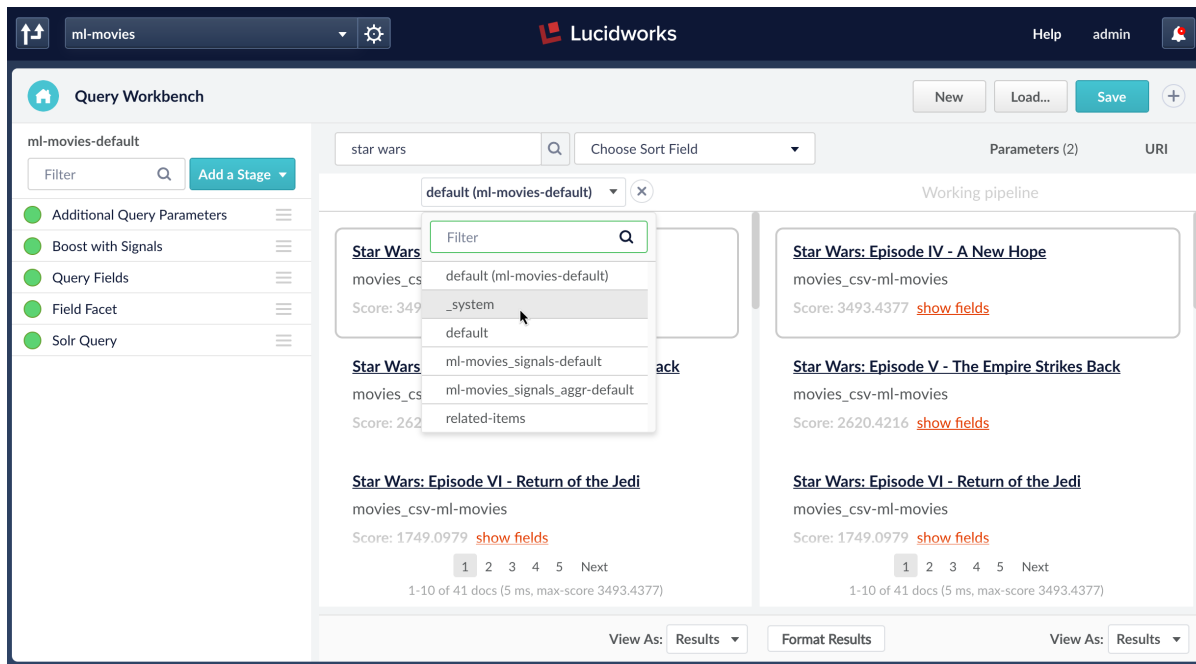
To use compare mode

1. Click **Compare**.

Another preview panel opens. In this view, you can compare results from one query pipeline side by side with another query pipeline.

2. In the left panel, select a pipeline to compare to your working pipeline.

Now you can see how the search results differ between the two pipelines:




Tip

When you click a document in one panel, the other panel automatically scrolls to the same document.

To exit compare mode, click the Close  icon.

Editing parameters

The Query Workbench allows you to edit the search parameters using these controls:



Formatting the search results

Ultimately, your search application will format the search results that end users see. The Query Workbench provides some formatting options for the preview panel.

At the bottom of the screen, click **Format Results** to configure how results are displayed in the QWB:

General Documents Facets ×

Control the number of results displayed, if highlighting information is utilized, and signals.

Set results per page

Display highlighting?

Signals

Show signal generators

Send click signals

View As: ▼

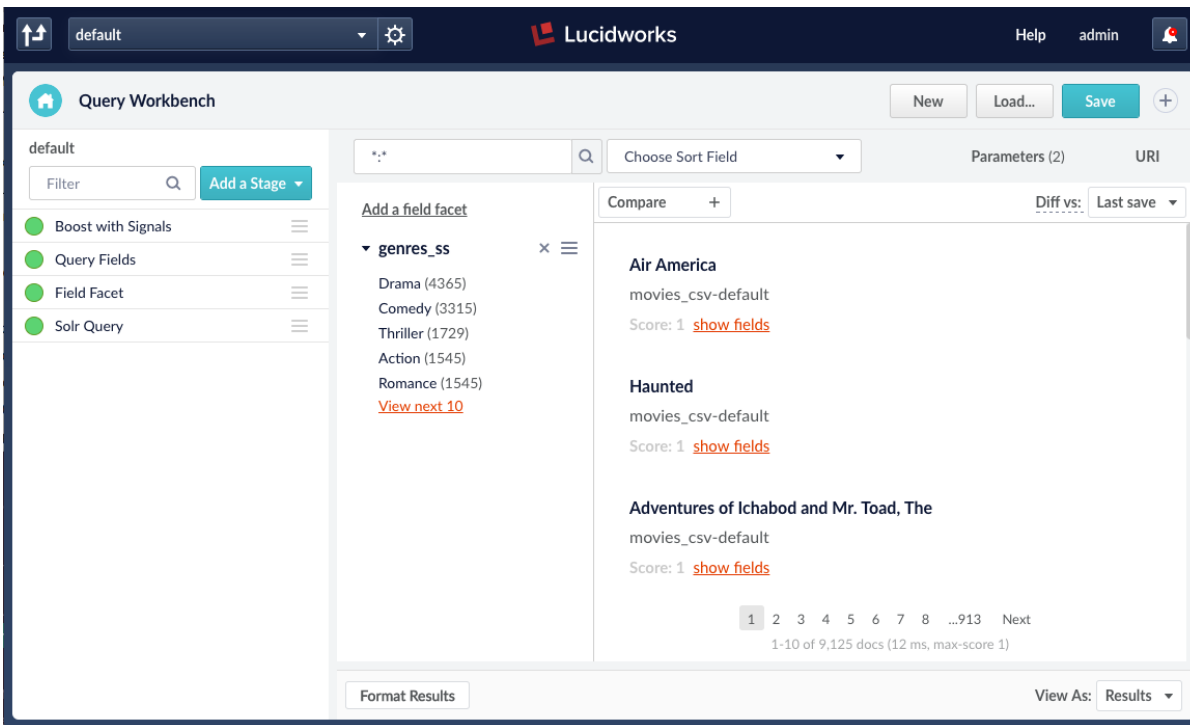
- The **General** tab configures pagination and highlighting, as well as enabling synthetic signals.
- The **Documents** tab configures the primary and secondary fields to display for every document.

These options only affect how the Query Workbench displays results; they have no effect on how your search application displays them.

Selecting the fields to display

While the Query Fields stage configures which fields are matched against incoming queries, the search application itself determines which fields to display in the search results. In this case, the Query Workbench is our search application.

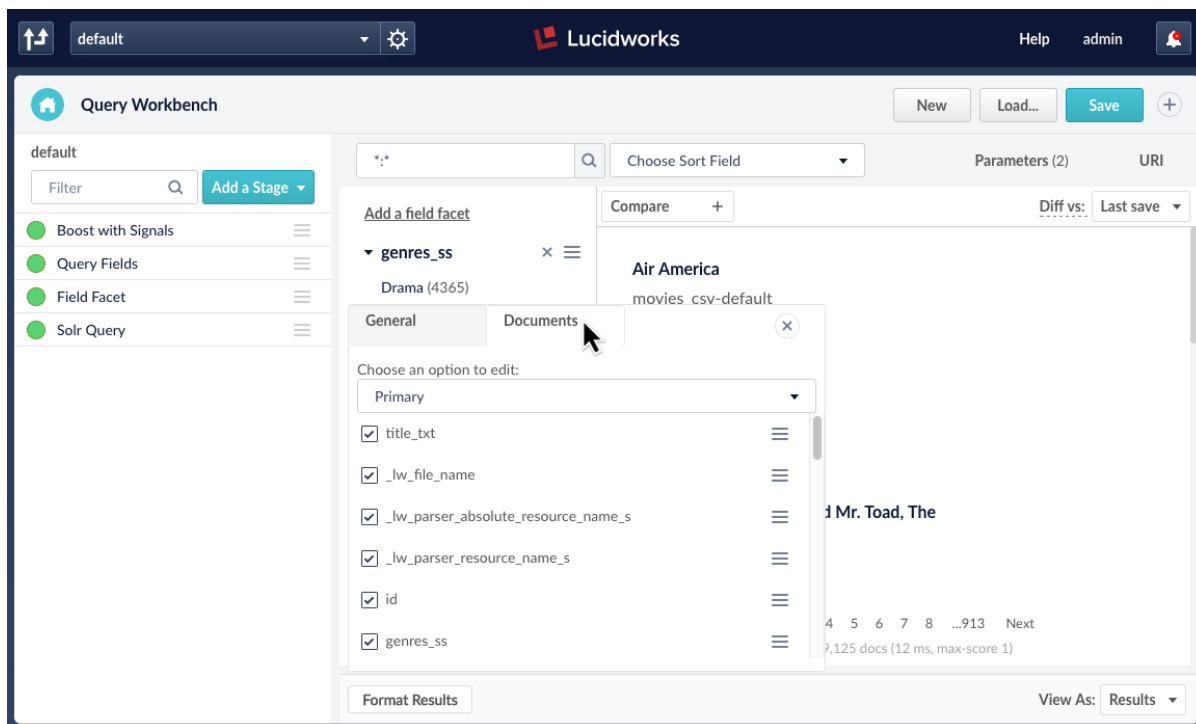
For example, the search results below include the `title_txt` field and the `_lw_data_source_s` field for movie documents:



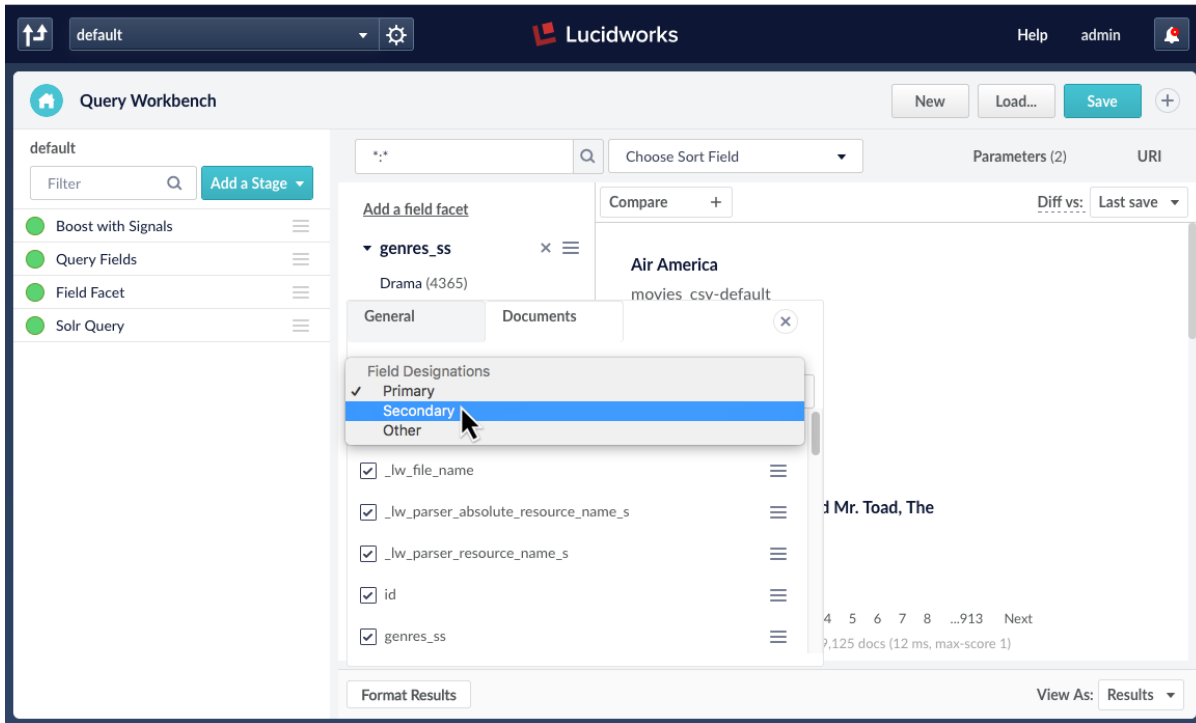
The `_lw_data_source_s` field is an internal field added by Fusion during ingest; our end users will not find it useful. The procedure below shows how to replace it with a more relevant field.

To select the display fields


1. In the lower left, click **Format Results**.
2. Click the **Documents** tab.



3. From the drop-down list, select **Secondary**.




4. Scroll down to the `year_ti` field.
5. Drag the `year_ti` field to the top of the list.



6. Close the Format Results panel by clicking Close (✕).

The preview panel automatically updates, and now we see the year of release for each movie title in our search results:



Configuring highlighting

Search results can be displayed with the search terms highlighted, by adding these [Solr query parameters](#) to the Additional Query Parameters stage of the query pipeline:

- `hl=true`
- `hl.fl=*`

By default, the Query Workbench ignores these parameters when rendering search results. To view highlighted search results in the preview panel of the Query Workbench, you must configure the parameters above *and* enable the **Display highlighting** option.

To enable highlighting in the preview panel

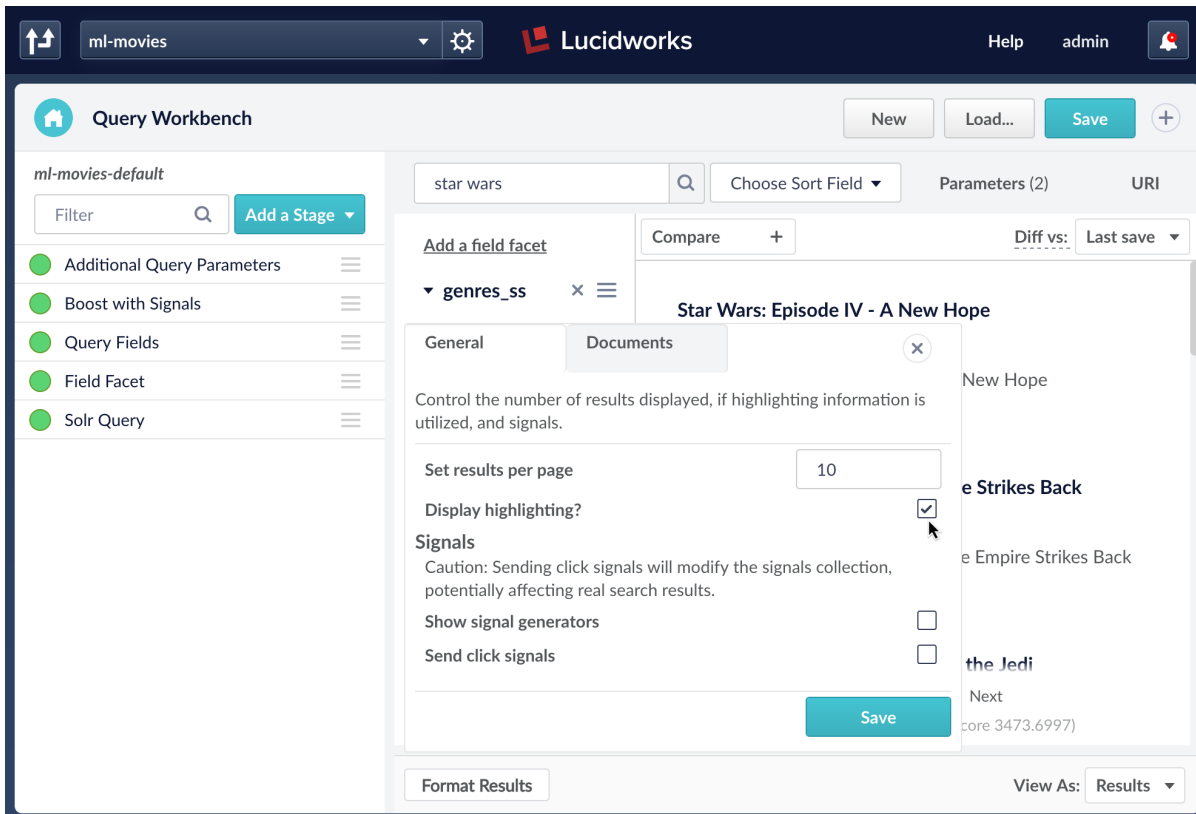
1. Add the Additional Query Parameters stage to your query pipeline, or select it if it is already there.
2. Under **Parameters and Values**, add the following query parameters and values:

<input type="checkbox"/>	hl	true	append ▼
<input type="checkbox"/>	hl.fl	*	append ▼

3. Click **Apply**.

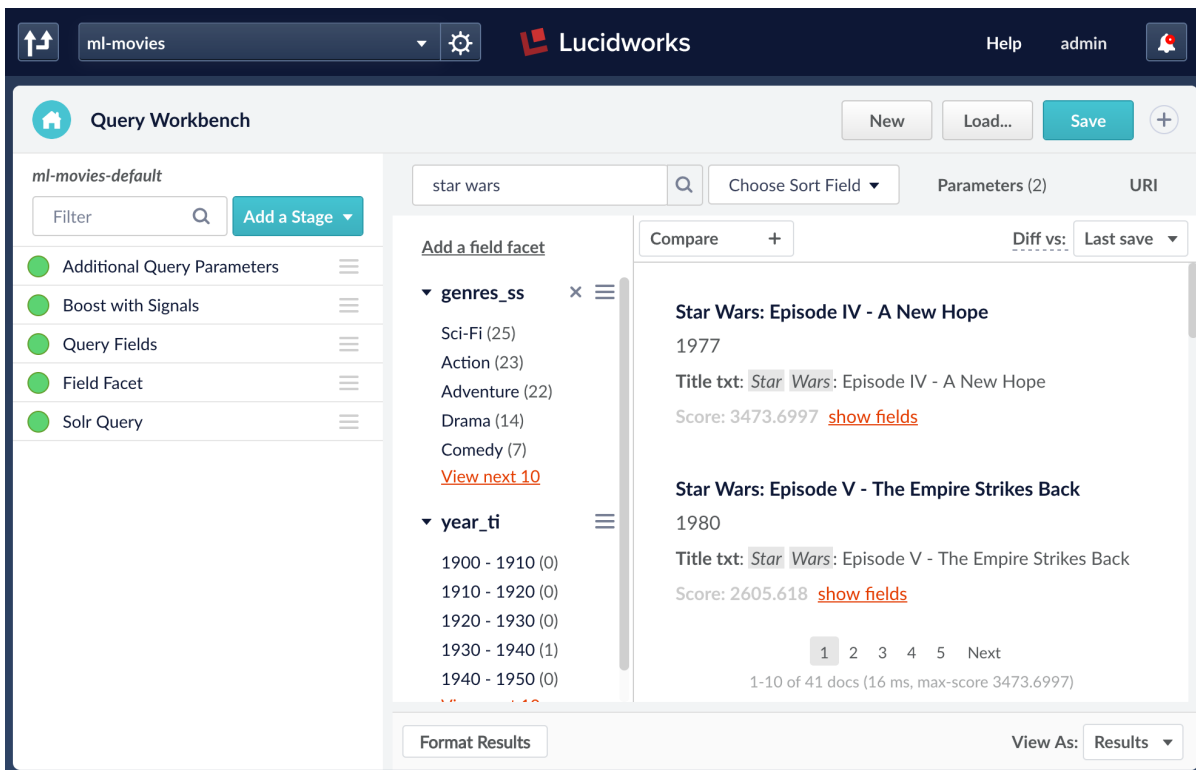
At this point, no highlighting appears in the preview panel.

4. At the bottom of the window, click **Format Results**.
5. Select **Display highlighting**.



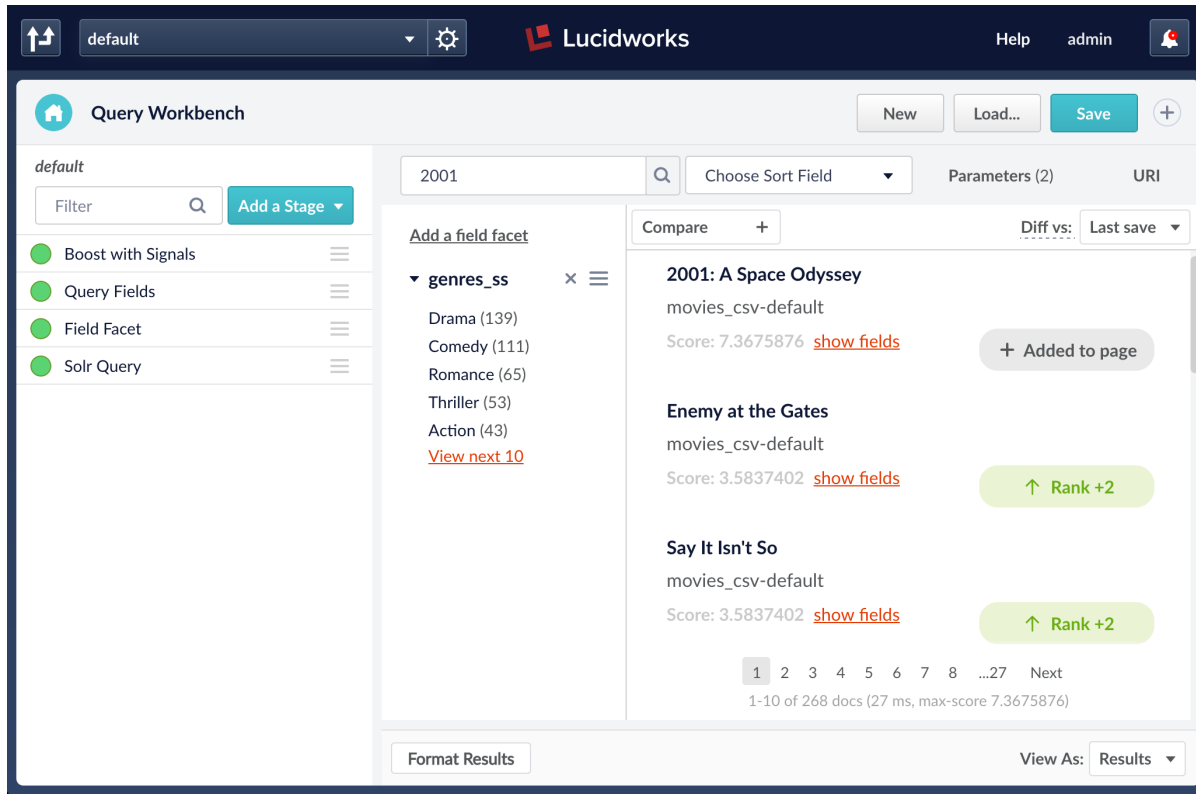
6. Click **Save**.

Now our search results include highlighted search terms:



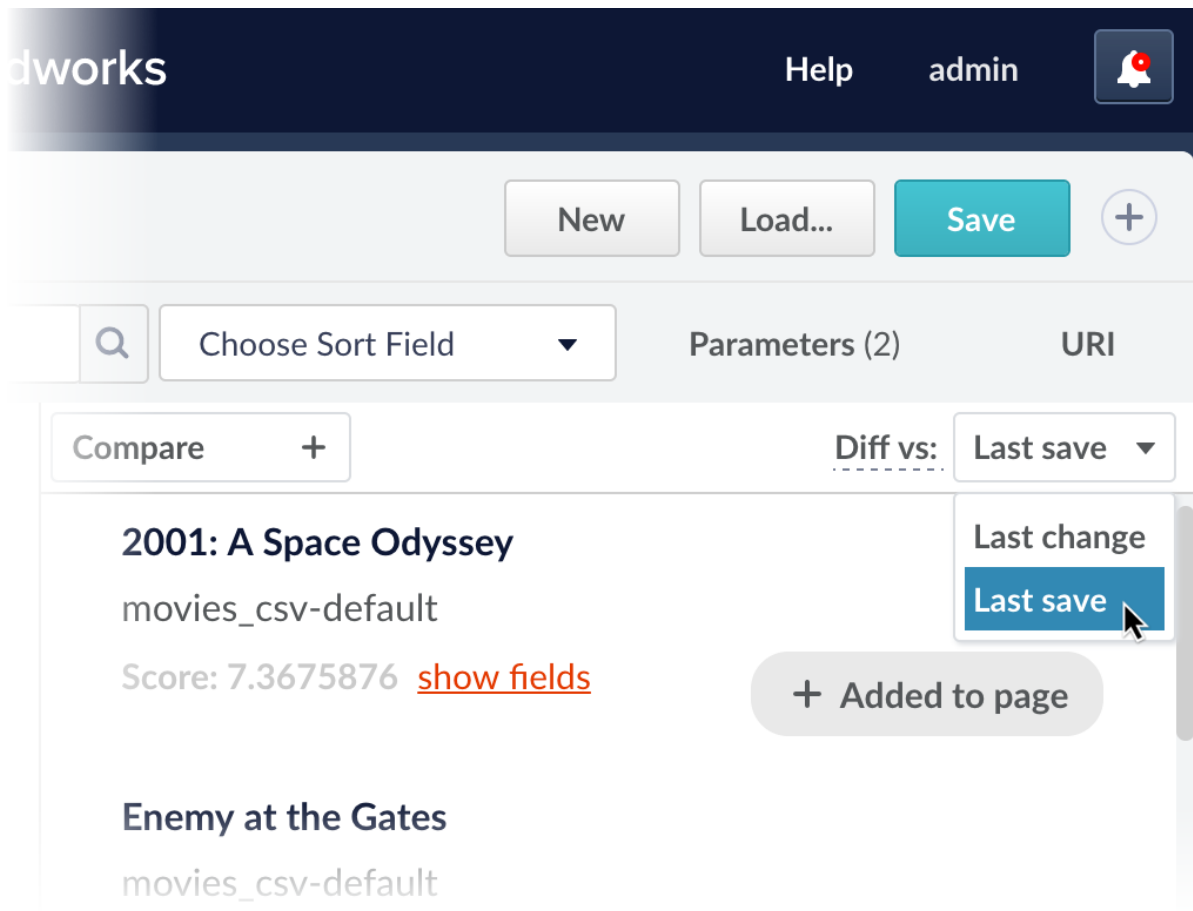
Configuring differences

As you modify the working pipeline, the Query Workbench marks the search results whose score is changed by your modifications.



The screenshot displays the Lucidworks Query Workbench interface. At the top, there is a navigation bar with a home icon, a dropdown menu set to 'default', a settings gear icon, the Lucidworks logo, and 'Help' and 'admin' links. Below this, the main interface is titled 'Query Workbench' and includes buttons for 'New', 'Load...', and 'Save'. The search bar contains the query '2001'. To the right of the search bar are options for 'Choose Sort Field', 'Parameters (2)', and 'URI'. A 'Diff vs:' dropdown menu is set to 'Last save'. The left sidebar shows a list of stages: 'Boost with Signals', 'Query Fields', 'Field Facet', and 'Solr Query'. The main content area shows a facet for 'genres_ss' with counts for Drama (139), Comedy (111), Romance (65), Thriller (53), and Action (43). Below the facet, three search results are displayed: '2001: A Space Odyssey' (Score: 7.3675876), 'Enemy at the Gates' (Score: 3.5837402), and 'Say It Isn't So' (Score: 3.5837402). The '2001: A Space Odyssey' result has a '+ Added to page' button. The other two results have '+ Rank +2' buttons. At the bottom, there is a 'Format Results' button and a 'View As: Results' dropdown menu.

You can select whether these differences are based on the last saved version of the pipeline or the last change that was made. Click the **Diff vs** menu in the upper right:



Step-by-step Query Workbench workflow

1. Use the Index Workbench to set up datasource(s).
2. Run a crawl on the data.
3. Open the Query Workbench.
4. Select a query pipeline to use or create a new, blank pipeline (the Default pipeline is created automatically and can be used as-is or with modification).
5. Modify the pipeline stages and set the order in which the queries will run.
6. Examine the live results in the Search panel.
7. Modify any pertinent stage settings to fine tune your results.
8. Click **Save** to commit the new settings to the pipeline

After a custom Query Pipeline configuration is saved, it becomes available throughout Fusion and can be re-used and modified.

2.2.4. Fusion Query Request Objects

A Fusion Query Request organizes the contents of each query submitted to a Fusion query pipeline. A Request object is comprised of the HTTP method, a set of HTTP headers, and a set of query parameters.

The Request Java API

Under the Fusion hood, a Request is a Java object. The link to the public API javadoc page is: [Request](#)

JSON representation of a Request object

The JSON representation of a Request object has three fields:

- **httpMethod**: value is a string, one of the defined HTTP methods (verbs), usually GET.
- **headers** : value is an object consisting of the set of defined header fields on the request.
- **params** : value is an object consisting of the set of query parameter names and corresponding parameter values.

The httpMethod and headers reflect the initial pipeline request. The Fusion Request object public API doesn't provide methods to update these fields.

The Fusion Request object public API provides methods to get and set the param fields and their values.

Query pipeline stages add, remove, and update the values in the params field. To see how the query parameters are affected by query pipeline stage, we insert a Logging Query stage at the very beginning of a Query Pipeline and capture the request object printed to the Fusion log file fusion/3.1.x/logs/api/api.log. The example query pipeline used here is called "medsamp-default." This pipeline is composed of the following sequence of stages:

- Logging Query
- Query Fields
- Logging Query
- Facet
- Logging Query
- Solr Query

Using the Admin Search tool, we submit the query "oxygen". The initial Logging Query stage shows the Request object. The "params" field takes an object consisting of a set of param names, param values. The param name "q" takes as its value the list of query terms, which in this case, is a list consisting of the word "oxygen":

```

{
  "headers" : {
    "Accept-Collections" : [ "*" ],
    "Fusion-User-Auth-Realm" : [ "native" ],
    "Host" : [ "172.16.2.12:8765" ],
    "Content-Length" : [ "0" ],
    "Fusion-Session-Id" : [ "e7085302-f5c5-4401-ba79-a4ead8a9f65a" ],
    "Accept-Encoding" : [ "gzip, deflate" ],
    "User-Agent" : [ "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_9_5) AppleWebKit/537.36 (KHTML, like Gecko)
Chrome/41.0.2272.89 Safari/537.36" ],
    "Via" : [ "1.0 apollo-proxy" ],
    "Fusion-User-Role" : [ "admin" ],
    "Accept" : [ "application/json, text/plain, */*" ],
    "Content-Type" : [ "application/json; charset=UTF-8" ],
    "Fusion-User-Name" : [ "admin" ],
    "Fusion-User-Id" : [ "a4096f7e-2de8-4910-beb4-93f30a3a5eb6" ]
  },
  "params" : {
    "echoParams" : [ "all" ],
    "start" : [ "0" ],
    "lw.pipelineId" : [ "medsamp-default" ],
    "q" : [ "oxygen" ],
    "json.nl" : [ "arrarr" ],
    "wt" : [ "json" ]
  },
  "httpMethod" : "GET"
}

```

The Query Fields stage updates the Request object, and the second Logging Query stage reflects these changes. Here we have omitted the "headers", as they are unchanged (and unchangable). The Query Fields stage added params "fl" (fields list - document fields returned), "qf" (query fields - the fields used for search), and "defType" (the default scoring type) to the params set.

```

{
  "headers" : { ... }
  "params" : {
    "fl" : [ "article-title_txt", "article-vernacular-title_txt", "article-author-lastname_txt", "mesh-
heading_ss", "pmid_txt", "article-language_s", "article-abstract_txt", "article-author-affiliation_txt" ],
    "echoParams" : [ "all" ],
    "start" : [ "0" ],
    "lw.pipelineId" : [ "medsamp-default" ],
    "q" : [ "oxygen" ],
    "qf" : [ "article-title_txt", "article-abstract_txt" ],
    "json.nl" : [ "arrarr" ],
    "wt" : [ "json" ],
    "rows" : [ "10" ],
    "defType" : [ "edismax" ]
  },
  "httpMethod" : "GET"
}

```

The Facets stage further updates the Request object, and the third Logging Query stage reflects these changes. The params set now includes param "facet".

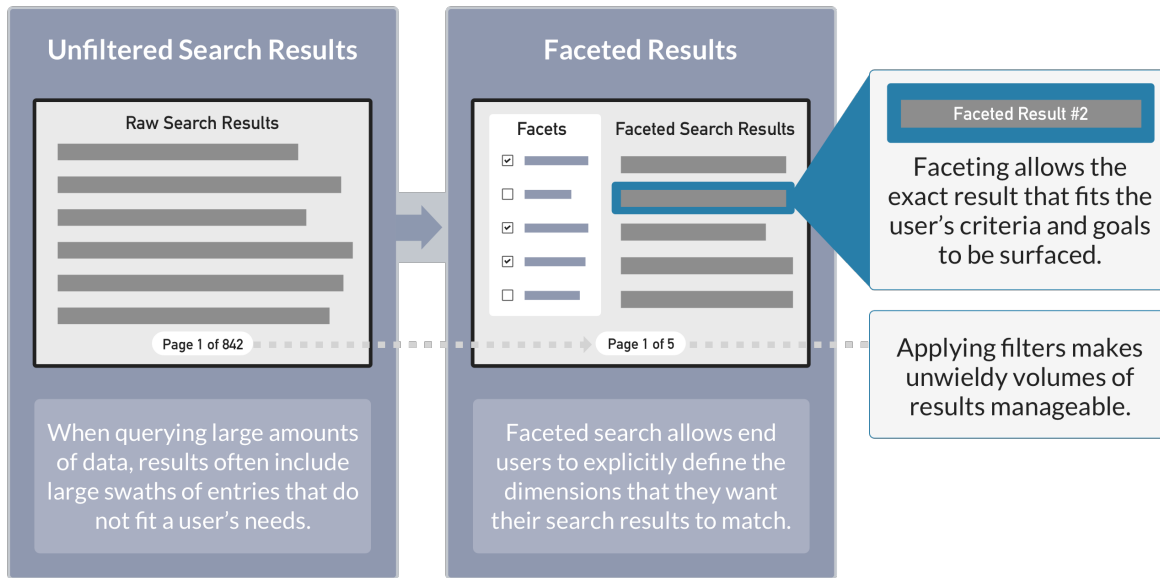
```
{
  "headers" : { ... }
  "params" : {
    "facet" : [ "true" ],
    "fl" : [ "article-title_txt", "article-vernacular-title_txt", "article-author-lastname_txt", "mesh-
heading_ss", "pmid_txt", "article-language_s", "article-abstract_txt", "article-author-affiliation_txt" ],
    "echoParams" : [ "all" ],
    "start" : [ "0" ],
    "lw.pipelineId" : [ "medsamp-default" ],
    "q" : [ "oxygen" ],
    "qf" : [ "article-title_txt", "article-abstract_txt" ],
    "json.nl" : [ "arrarr" ],
    "wt" : [ "json" ],
    "rows" : [ "10" ],
    "defType" : [ "edismax" ]
  },
  "httpMethod" : "GET"
}
```

2.2.5. Fusion Query Response Objects

A Fusion query Response object contains the Solr response as well as a set of HTTP headers. It is used to improve the search experience by refining, expanding, filtering, or otherwise modifying the Solr response.

Under the Fusion hood, a Response is a Java object. The link to the public API javadoc page is: [Response](#).

2.2.6. Faceting



Faceting is the name given to a set of computed counts over a search result which are returned together with the documents which match the search query. Facets are most often used to create additional navigational controls on the search results page or panel which allow users to expand and restrict their search criteria in a natural way, without having to construct complicated queries. For example, popular e-commerce facets include product category, price range, availability, and user ratings.

Fusion leverages [Solr's Faceting](#) search components.

Field faceting

In Solr the most straightforward kind of faceting is field faceting, in which Solr's FacetComponent computes the top values for a field and returns the list of those values along with a count of the subset of documents in the search results which match that term. Field faceting works best over fields which contain a single label or set of labels from a finite, controlled lexicon such as product category. Fusion's Facet Query Stage can be used to configure field faceting as part of the search query pipeline.

Range faceting

Range facets are used for fields which contain date or number values. Values can be grouped into ranges by specifying additional query parameters.

To configure range faceting, use the Additional Query Parameters Stage to specify [Solr range faceting parameters](#).

Faceting concepts

Key Facet Concepts:

Term

A specific value from a field.

Limit

The maximum number of terms to be returned.

Offset

The number of top facet values to skip in the response (just like paging through search results and choosing an offset of 51 to start on page 2 when showing 50 results per page).

Sort

The order in which to list facet values: `count` ordering is by documents per term, descending, and `index` ordering is sorted on term values themselves.

Missing

The number of documents in the results set which have no value for the facet field.

Choice of facet method

(advanced) Specify Solr algorithm used to calculate facet counts. (See [Facet Method Configuration](#) for details). One of:
+

- `enum` - small number of distinct categories
- `fc` ("field cache") - many different values in the field, each document has low number of values, multi-valued field
- `fcs` ("single value string fields") - good for rapidly changing indexes

Further Reading

<https://lucidworks.com/blog/2014/10/03/pivot-facets-inside-and-out/>

<https://lucidworks.com/blog/2015/01/29/you-got-stats-in-my-facets/>

<https://lucidworks.com/blog/2016/08/12/pivoting-to-the-query-using-pivot-facets-to-build-a-multi-field-suggester/>

2.2.7. Search Query Pipeline Stages

A query pipeline is made up of a series of query stages that process incoming search queries.

A pipeline stage definition associates a unique ID with a set of properties. These definitions are registered with the Fusion API service and stored in ZooKeeper for re-use across pipelines and search applications.

Fusion includes a number of specialized query stages as well as a JavaScript stage that allows advanced processing via a JavaScript program.

Configuring query pipeline stages

- In the Fusion UI, the Query Workbench provides an environment for configuring the stages in a query pipeline.
- The Query Stages API is used to create, list, update, or delete query stages using JSON. See also the Query Pipelines API.

Conditional query processing

Query Pipeline stages are used to modify Request objects and Response objects. Each stage can include a conditional JavaScript expression (the 'condition' property in its configuration) that can access these objects.

For example, this condition first checks that the property "fusion-user-name" is specified in the Request object, then checks for a particular value:

```
request.hasParam("fusion-user-name") && request.getFirstParam("fusion-user-name").equals("SuperUser");
```

Reference topics

See these reference topics for complete details about each query pipeline stage:

Setup

- Active Directory Security Trimming
- Field Facet
- More Like This
- Query Fields
- Security Trimming

Results relevancy

- Block Documents
- Boost Documents
- Landing Pages
- Parameterized Boosting
- Recommend More Like This stage
- Boost with Signals stage
- Recommend Items for User stage

- Recommend Items for Item stage

Fetch data

- JDBC Lookup
- REST Query
- Solr Query
- Solr Subquery

Troubleshooting

- Logging
- Send PagerDuty Message
- Send Slack Message
- Send SMTP Email
- Write Log Message

Advanced

- Additional Query Parameters
- Javascript
- Retrieve Stored Parameters

Other

- Analytics Catalog Query
- Call Pipeline
- Experiment Query
- Machine Learning
- Parameterized Faceting
- Return Query Parameters
- Rollup Aggregation

2.2.8. Query Profiles

Query profiles allow you to consistently point your search application at a static endpoint, but give you the flexibility to change the actual query pipeline being used.

For example, an e-commerce site may want to create a query pipeline to support a month-long promotion. Once the pipeline is configured, it can be easily enabled by changing the profile in use by the front-end application to use the new pipeline.

A profile can be created or modified with the UI or with a REST API. The UI is described below, the REST API is described in the section Query Profiles API.

Query Profiles in the UI

The Profiles tab shows the index profiles on the left and the query profiles on the right. Hover over the name of a profile, and an **edit** button will appear to allow you to change the pipeline the profile is mapped to. Hover over the name of a pipeline, and you will be able to jump to edit that pipeline.

Click **Add Profile** to add a profile. The next screen will show a form allowing you to define the profile name and either select an existing pipeline or create a new pipeline with the name you choose. Click **Create** to save the new profile.

2.2.9. Using Query Pipelines with SolrJ

Fusion Pipelines can be used in conjunction with a SolrJ client, allowing you to use the power of Fusion pipelines with an existing Solr installation.

Authentication with SolrJ

Fusion user authentication and authorization is carried out by the Fusion UI service. For details on how Fusion handles authentication and authorization, please see Access Control.

When using SolrJ, however, there are two approaches that can be used: basic authentication and passing credentials in the URL. Once the authentication has occurred, the roles that have been assigned to the user provide the authorization.

Basic Authentication

The Basic authentication approach looks very similar to the session-based approach. However, some classes are changed.

```
HttpClient client = HttpClientBuilder.create().useSystemProperties()
    .addInterceptorLast(new PreEmptiveBasicAuthenticator(user, password))
    .build();
HttpSolrServer server = new HttpSolrServer(url, client);
// ...
public static class PreEmptiveBasicAuthenticator implements HttpRequestInterceptor {
    private final UsernamePasswordCredentials credentials;
    public PreEmptiveBasicAuthenticator(String user, String pass) {
        credentials = new UsernamePasswordCredentials(user, pass);
    }
    public void process(HttpRequest request, HttpContext context) throws HttpException, IOException {
        request.addHeader(BasicScheme.authenticate(credentials, "US-ASCII", false));
    }
}
```

URL Credentials

The URL credential approach provides the ability to pass the authentication properties in the URL, as in this example:

```
String url = "http://user:pass@localhost:8764/api/apollo/solr/demo";
HttpSolrServer server = new HttpSolrServer(url);
```

Example

The example below demonstrates the use of query profiles and query pipelines for querying a collection named 'test' with a query pipeline named 'default', using the basic authentication approach.

```

package com.lucidworks.apollo.testQueryPipeline;

import org.apache.http.HttpRequestInterceptor;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.impl.client.DefaultHttpClient;
import org.apache.solr.client.solrj.SolrServerException;
import org.apache.solr.client.solrj.impl.HttpSolrServer;
import org.apache.solr.client.solrj.impl.XMLResponseParser;
import org.apache.solr.client.solrj.response.QueryResponse;
import org.apache.solr.common.params.ModifiableSolrParams;
public class TestQPSolrJ {
    public static void main(String[] args) throws SolrServerException {
        /**
         * Request URL points to Fusion UI which uses authentication proxy.
         * Example uses basic authentication
         * URL is Apollo query-pipelines endpoint:
         * http://{hostname}:{port_name}/api/apollo/query-
pipelines/{pipeline_name}/collections/{collection_name}
         */
        String url = "http://localhost:8764/api/apollo/query-pipelines/default/collections/demo";
        String user = "admin";
        String password = "password123";
        HttpClient client = HttpClientBuilder.create().useSystemProperties()
            .addInterceptorLast(new PreEmptiveBasicAuthenticator(user, password))
            .build();
        HttpSolrServer solrServer = new HttpSolrServer(url, client);
        solrServer.setParser(new XMLResponseParser());
        ModifiableSolrParams solrParams = new ModifiableSolrParams();
        solrParams.add("q", "*:*");
        QueryResponse out = solrServer.query(solrParams);
        System.out.println("QTime is " + out.getQTime());
        System.out.println("RH is " + out.getResponseHeader().toString());
    }
    public static class PreEmptiveBasicAuthenticator implements HttpRequestInterceptor {
        private final UsernamePasswordCredentials credentials;
        public PreEmptiveBasicAuthenticator(String user, String pass) {
            credentials = new UsernamePasswordCredentials(user, pass);
        }
        public void process(HttpRequest request, HttpContext context) throws HttpException, IOException {
            request.addHeader(BasicScheme.authenticate(credentials, "US-ASCII", false));
        }
    }
}

```

For more information about SolrJ, see the Apache Solr Reference Guide section [Using SolrJ](#).

2.2.10. Search Query Reporting

Fusion includes several reports that can help you gain insight into the performance of your search application and the behavior of your users.

In order to use these reports, you must have the 'searchLogs' feature enabled for a collection. This will create a parallel collection named '<collection>_logs'. When requests for report data is sent to the main collection, the data is pulled from the *_logs collection.

Available Reports

lessThanN

This report provides a way to discover queries that returned less than a defined number of results.

In addition to defining the number of results for inclusion in the report, you can also limit by a date range.

topQueries

This report shows the queries that were performed most often. It reports the user's entire query.

While it does not take a minimum number parameter, it can be limited by a date range.

topN

The TopN report finds the most popular terms. This is distinct from the topQueries report because it reports on terms and not the entire query.

Possible fields are those that appear when investigating a particular item:

- `collection_s`: the collection used for the query
- `id`: the document ID
- `q_txt`: the query text
- `q_s`:
- `qtime_l`: the length of time the query portion of the response
- `totaltime_l`: the total length of time for the response
- `numdocs_l`: the number of documents found for the query
- `timestamp_dt`: the timestamp of the query
- `req_facet_ss`: if facets were requested
- `req_echoParams_ss`: the setting for "echoParams" during the query
- `req_q_ss`: the request itself
- `req_facet.field_ss`: the facet fields requested
- `req_rows_ss`: the number of rows requested
- `version`: The document version at the time of the query

topClicked

This report shows the items that were clicked most often. It requires that Signals have been enabled for a collection, that

click events have been recorded, and that the click signals have been aggregated. See the section on Signals for more information.

In addition to accepting a number to define how many items to return, it can also be limited by date range.

histo

This report provides a histogram that matches the query parameters. A field is defined as the basis for the data, and then a range of values and an interval are provided.

The available fields are:

- `qtime_l`: the length of time the query portion of the response
- `totaltime_l`: the total length of time for the response
- `numdocs_l`: the number of documents found for the query
- `timestamp_dt`: the timestamp of the query
- `req_rows_ss`: the number of rows requested

dateHisto

This report provides a histogram of queries based on a provided date range.

2.2.11. Custom JavaScript Stages for Query Pipelines

The JavaScript Query stage allows you to write a custom processing logic using JavaScript to manipulate search requests and responses. The first time that the pipeline is run, Fusion compiles the JavaScript program into Java bytecode using the JDK's JavaScript engine.

The JavaScript Query stage allows you to run JavaScript functions over search requests and responses by manipulating variables called "request" and "response" which are Request and Response objects, respectively.

JavaScript Query Stage Global Variables

[JavaScript](#) is a lightweight scripting language. The JavaScript in a JavaScript stage is standard [ECMAScript](#). What a JavaScript program can do depends on the container in which it runs. For a JavaScript Query stage, the container is a Fusion query pipeline. The following global pipeline variables are available:

Name	Type	Description
<code>request</code>	Request	The Solr query information.
<code>response</code>	Response	The Solr response information.
<code>params</code>	Context	A reference to the container which holds a map over the pipeline properties. Used to update or modify this information for downstream pipeline stages.
<code>ctx</code>	Context	A reference to the container which holds a map over the pipeline properties. Used to update or modify this information for downstream pipeline stages.
<code>_context</code>	Context	Another reference to the same object as <code>ctx</code> , included because some stages use this name instead.
<code>collection</code>	String	The name of the Fusion collection being indexed or queried.
<code>solrServer</code>	BufferingSolrServer	The Solr server instance that manages the pipeline's default Fusion collection. All indexing and query requests are done by calls to methods on this object. See SolrClient for details.

Name	Type	Description
<code>solrServerFactory</code>	<code>SolrClientFactory</code>	The <code>SolrCluster</code> server used for lookups by collection name which returns a Solr server instance for a that collection, e.g. <pre>var productsSolr = solrServerFactory.getSolrServer("products");</pre>

Global variable `logger`

The global variable named `logger` writes messages to the logfile of the server running the pipeline. Since Fusion's api service does the query pipeline processing, these log messages go into the logfile: `fusion/3.1.x/var/log/api/api.log`. There are 5 methods available, which each take either a single argument (the string message to log) or two arguments (the string message and an exception to log). The five methods are, "debug", "info", "warn", and "error".

JavaScript Query Stage Examples

Add a parameter to the query request

```
request.addParam("foo", "bar");
```

Add a parameter to the query response

This example contains a simple JavaScript function which adds information to the query response. Repeated calls to this function build out the response.

```
function add_to_response(key, list) {
  if (list.length > 0) {
    response.initialEntity.appendStringList(key, Java.to(list, Java.type('java.util.List')));
  }
}
add_to_response('banners', ctx.getProperty('banners'));
add_to_response('landing-pages', ctx.getProperty('redirects'));
```

Debugging and Troubleshooting

To debug a JavaScript Index stage you can:

- Check the Fusion api server logs for errors.
- Use the `logger` object for print debugging (in the Fusion api logfile).
- Use the Pipeline Preview tool (only available in Fusion 1.x)

The JavaScript Engine Used by Fusion

The JavaScript engine used by Fusion is the Nashorn engine from Oracle. See [The Nashorn Java API](#) for details.

Upgrading to the latest Nashorn engine

The default version of the Nashorn engine used by Fusion versions 2.4.1 and earlier is the `nashorn-0.1-jdk7.jar` which

contains many bugs that have since been fixed in the official JDK 1.8 version. In order to use the latest version of the Nashorn engine, you must:

- Have an up-to-date version of Java 8 installed.
- Remove the nashorn-0.1-jdk7.jar from the Fusion classpaths:
 - `cd fusion/3.1.x`
 - `find . -name "nashorn-0.1-jdk7.jar" -print -exec rm -i {} \;`

Creating and accessing Java types

The following information is taken from Oracle's JavaScript programming guide section 3, [Using Java From Scripts](#).

To create script objects that access and reference Java types from Javascript use the `Java.type()` function:

```
var ArrayList = Java.type("java.util.ArrayList");
var a = new ArrayList;
```

2.2.12. Solr Query Language Cheat Sheet

This cheat sheet is a quick reference to the Solr query language. Use this syntax when querying Fusion via the Query Pipelines API

There are two ways to query a Fusion collection using the parameters below:

- Enter query parameters in the Query Workbench or the Quickstart.
- Append query parameters to the `/query-pipelines/{id}/collections/{collection}/{handler}` endpoint.

Wildcards and regular expressions are supported.

Wildcards

- `?` – Single-character wildcard
- `*` – Multi-character wildcard

Common query parameters

<code>q</code>	Query parameters. The full Solr query, using Lucene query syntax .
<code>fq</code>	Filter query. A query string that limits the query results without influencing their scores.
<code>sort</code>	Sort field/direction. The field on which to sort, followed by a space and direction (<code>desc</code> or <code>asc</code>). You can specify multiple sort fields like this: <code>sort=title asc,year desc</code>
<code>rows</code>	Max results per page. This set the "page size" for paginated search results.
<code>start</code>	Pagination offset. The number of results to skip, for pagination purposes.
<code>fl</code>	Field List. The list of fields to return in the query results.
<code>df</code>	Default field. Used to configure the <code>q</code> and <code>fq</code> parameters. If not specified, the default field is <code>text</code> .
<code>wt</code>	Response writer. Select the response format by specifying one of Solr's response writers .

See also the [Solr facet parameters documentation](#).

Query examples

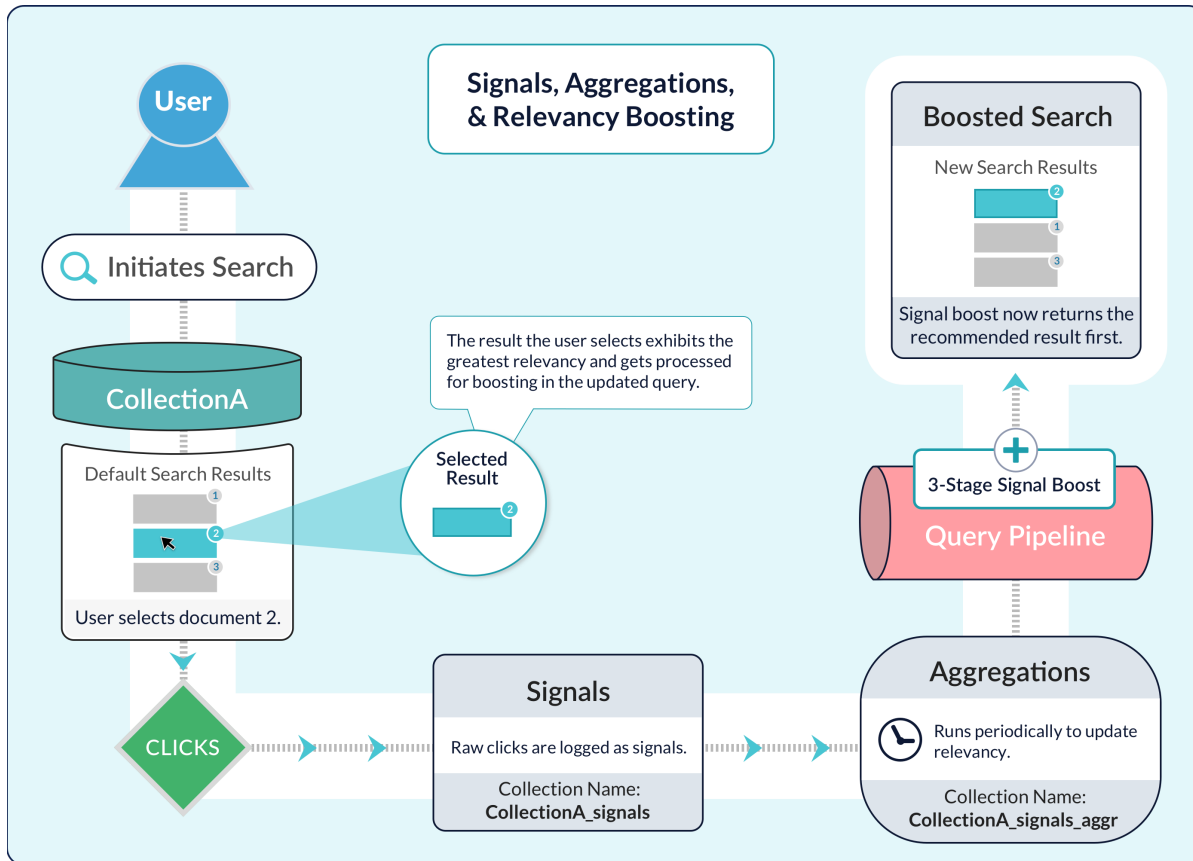
Search only the `title` field in the "docs" collection for the term "solr", and format the results as JSON:

```
http://localhost:8764/api/apollo/query-pipelines/default/collections/docs/select?q=solr&fl=title&wt=json
```

2.3. Signals and Aggregations

In addition to the basic search experience enabled through query pipelines, Fusion provides ways to develop an enhanced search experience for your end users and provide useful data for your analytics team. The primary mechanisms for doing this are signals and aggregations.

By collecting signals and aggregating them, you compile a body of data that allows you to develop a sophisticated search experience, with rich search results for your end users, based on past user behavior.



Signals and aggregated signals are stored each in their own collection. These collections are associated with a primary collection, so that a collection named "products" will have two related collections: "products_signals" and "products_signals_aggr". By default, when using the UI to create a collection, a "signals" and "aggregated signals" collection are also created.

2.3.1. Signals

Signals are events that are collected for analysis or to enhance the search experience for end users. Common types of signal events include clicks, purchases, downloads, ratings, and so on.

Generally, signals are only useful after aggregation.

2.3.2. Aggregations

Aggregations are processed signals. An *aggregator* reads the raw signals and returns interesting summaries, ranging from simple sums to sophisticated statistical functions.

Crucially, it must be possible to relate the documents in an aggregated signals collection to documents in the primary

collection, in order to use the aggregated signals for recommendations and/or boosting of searches over the primary collection.

2.3.3. The cold start problem

The "cold start" problem means it is hard to personalize the search experience when insufficient signals have been aggregated. For example, it is hard to offer recommendations to users who have never visited before, or for queries that have never been issued before, or for items that have been recently introduced into the system.

Fusion provides solutions for this problem using its query pipelines. A query pipeline that includes stages for blocking, boosting, or recommending based on signals can also include stages that provide fallbacks. In the case where there is not enough data to provide specialized blocking, boosting, or recommendations, the pipeline can return a simpler set of search results using Solr's normal relevancy calculation.

A common solution to the cold start problem is to sort or boost on a certain field to provide pseudo-recommendations when more specific recommendations are not available. For example, you can sort on the `sales_rank` field to recommend the most popular products, or boost on the `date_added` field to recommend the newest items.

2.3.4. Signals

A *signal* is a recorded event related to one or more documents in a collection. Signals can record any kind of event that is interesting to your organization. Queries and clicks are the most common types of signals, as they are useful for tracking what users search for and what actions they take.

Signals are indexed in a secondary collection which is linked to the primary collection by the naming convention `<primarycollectionname>_signals`. So, if your main collection is named `products`, the associated signals collection is named `products_signals`. The signals collection is created automatically when signals are enabled for the primary collection.

Signals are indexed just like ordinary documents. The signals collection can be searched like any other collection, for example to retrieve a user's search history or last viewed items.

Signals are most useful when they are *aggregated* into a set of summaries that can be used to enrich the search experience through recommendations and boosting. Like the signals collection, a `<primarycollectionname>_signals_agg` collection is created automatically when signals are enabled for a primary collection. An aggregation job is also created automatically, and scheduled to run every two minutes.

Enabling and disabling signals

Using the Fusion UI, when you create a collection, signals are enabled and a signals collection created by default.

Using the API, the `/collections/{collection}/features/{feature}` endpoint enables or disables signals for any collection:

Check whether signals are enabled for a collection

```
curl -u user:pass http://localhost:8765/api/v1/collections/<collection-name>/features/signals
```

Enable signals for a collection

```
curl -u user:pass -X PUT -H "Content-type: application/json" -d '{"enabled" : true}'  
http://localhost:8765/api/v1/collections/<collection-name>/features/signals
```

Disable signals for a collection

```
curl -u user:pass -X PUT -H "Content-type: application/json" -d '{"enabled" : false}'  
http://localhost:8765/api/v1/collections/<collection-name>/features/signals
```

Signal document structure

A raw signal is stored as a Solr document with the following fields, which are derived from the raw signal as follows:

Field	Description
<code>id</code> <i>Optional</i>	The signal ID. If no ID is supplied, one will be automatically generated.

Field	Description
<p>type <i>Required</i></p>	<p>The signal type that is being sent. This value is used during aggregation to filter events of the same type. Types can be mixed in aggregation jobs, if needed.</p> <p>The type can consist of any string you choose. For consistency, always send events of the same type with the same type value.</p> <p>During indexing, type values will be moved to a field named <code>type_s</code>.</p>
<p>params <i>Optional</i></p>	<p>The params allow flexible definition of the fields you care about and will use later for signal aggregation:</p> <ul style="list-style-type: none"> • docId – A unique document ID <p>This is stored in the Solr raw signal document as field <code>doc_id_s</code>.</p> • userId – A unique user ID <p>This is stored in the Solr raw signal document as field <code>user_id_s</code>.</p> • query – A query string; for example, a user’s search <p>This is copied to the Solr raw signal document as both fields <code>query_s</code> and <code>query_t</code>. Some cleanup occurs to convert the string to lowercase, decode URL encoding, and replace white space with single space characters. The original query is saved in field <code>query_orig_s</code>.</p> • filterQueries – A list of strings, such as filters on the search query <p>This is copied to the Solr raw signal document as both <code>filters_s</code> and <code>filters_orig_ss</code>.</p> • collection – The primary collection name • weight – A float value representing the relative weight of this signal <p>This is saved in the field <code>weight_d</code>.</p> • count – A positive integer value representing the incremented count of signals <p>This is saved in the field <code>count_i</code>.</p>

Field	Description
<code>timestamp</code>	<p>The timestamp of the signal event.</p> <ul style="list-style-type: none"> • When using the Signals API, this property is optional; it defaults to the current server time. • When using the Signal Formatter index stage, one of the following fields must be present: <code>timestamp</code>, <code>timestamp_tdt</code>, <code>timestamp_dt</code>, or <code>epoch</code>.

Here is the JSON representation of one click signal, taken from an example dataset of synthetic clickstream data:

```
{ "params": {
  "docId": "2125233",
  "filterQueries": ["cat00000","abcat0100000", "abcat0101000", "abcat0101001"],
  "query": "Televisiões Panasonic 50 pulgadas" }
"type":"click",
"timestamp": "2011-09-01T23:44:52.533000Z",
}
```

The default signals index pipeline

When indexing signals, a default index pipeline named `_signals_ingest` will be used unless you specify a different index pipeline.

The `_signals_ingest` pipeline has three stages:

1. Format Signals stage
2. Field Mapping stage
3. Solr Indexer stage

If you prefer different options in the signals index pipeline, you can pass a query parameter when indexing signals that contains the name of your custom index pipeline.

If you create a custom pipeline, it must include a Field Mapping stage and a Solr Indexer stage (see Index Pipeline Stages for more details), which sends the documents to Solr. Additionally, the Solr Indexer stage must have the `enforce_schema` property set to "true".

Removing signals

The aggregator includes an option to delete signals after they have been processed. If, however, you have chosen not to remove signals during aggregation, you can also run a "delete" query in Solr to delete documents from the signals collection.

Video tutorial

This video tutorial explains how to boost searches using click signals and aggregations:

2.3.5. Aggregations

Signals are most useful when they are aggregated into a set of summaries that can be used to enrich the search experience through recommendations and boosting. Aggregation jobs are a subtype of Spark jobs.

When signals are enabled for a "primary" collection, a `<primarycollectionname>_signals` collection and a `<primarycollectionname>_signals_aggr` collection are created automatically.

Aggregation Pipelines

Aggregated events are indexed, and use a default pipeline named "aggr_rollup". This pipeline contains one stage, a Solr Indexer stage to index the aggregated events.

You can create your own custom index pipeline to process aggregated events differently if you choose.

Aggregation Functions

The section Aggregator Functions documents the available set of aggregation functions.

Custom aggregation functions can be defined via a JavaScript stage. The options described in Aggregator Scripting provide more detail on the objects available for scripts.

Aggregation properties

The aggregation process is specified by an aggregation type consisting of the following list of properties:

Name	Description
<code>id</code>	Aggregation ID
<code>groupingFields</code>	List of signal field names
<code>signalTypes</code>	List of signal types
<code>aggregator</code>	Symbolic name of the aggregator implementation
<code>selectQuery</code>	Query string, default :
<code>sort</code>	Ordering of aggregated signals
<code>timeRange</code>	String specifying time range, e.g., <code>[* TO NOW]</code>
<code>outputPipeline</code>	Pipeline ID for processing aggregated events
<code>outputCollection</code>	Output collection name
<code>rollupPipeline</code>	Rollup pipeline ID
<code>rollupAggregator</code>	Name of the aggregator implementation used for rollups
<code>sourceRemove</code>	Boolean, default is false
<code>sourceCatchup</code>	Boolean, default is true
<code>outputRollup</code>	Boolean, default is true
<code>aggregates</code>	List of aggregation functions

Name	Description
params	Arbitrary parameters to be used by specific aggregator implementations

EventMinerAggregator

This aggregator type produces synthetic co-occurrence documents for predefined fields, based on co-occurring data in events in the same session.

This aggregator assumes that events have the following fields:

- `timestamp_tdt` - event timestamp
- `user_id_s` - user ID
- `query_s` - (optional) query that is related to this event
- `doc_id_s` - (optional) document ID that is related to this event

A session is then defined as a series of events created by a given `user_id_s` with timestamps falling within a session timeout limit (which is configurable via aggregator parameters).

Example document produced by this aggregator:

```

{id": "29d2c95e48154a4ebdc03158b0dd7875-25170",
  "entity_type_s": "doc_id_s",
  "entity_id_s": "2548405",
  "co_occurring_docIds_ss": [
    "2938114"
  ],
  "co_occurring_docIds_counts_is": [2],
  "in_queries_ss": [
    "tennis",
    "nicktoons kinect",
    "virtua tennis 4"
  ],
  "in_queries_counts_is": [3, 1, 1],
  "in_session_ids_ss": [
    "f2de89f97e1638614795d3b03c50d5b1",
    "eae6242c6b58526d2a039d4bd95a85b6",
    "d87d264a528fd7ba162c20209bd3ca8a"
  ],
  "in_session_ids_counts_is": [1, 1, 1],
  "in_user_ids_ss": [
    "4c79f9ebcf7d50ba5d25a3fca0343929ff05c822",
    "2943dea5408b6f947bbd42aa28f9d8006bfb366a",
    "3fa43872d2275d5e6463ff2de1d95dca51d0003f"
  ],
  "in_user_ids_counts_is": [1, 1, 1]
}

```

This example illustrates the following:

- This is a synthetic co-occurrence document for an entity of type `doc_id_s`, with entity ID "2548405" (that is, a document with this ID). Other similar documents are produced for sessions, queries and users as the primary entity ID.

- The `in_queries_ss` field contains queries that led to this document (in this case, since events were click-throughs, these are the queries that resulted in clicks on a link to this document).
- The `related_docIds_ss` field shows documents that users also clicked within the same search session (thus implicitly indicating that they are related to this one).
- The `in_user_ids_ss` and `in_session_ids_ss` fields contains user IDs and session IDs respectively, where the click events for this document originated from.

This document can be viewed as a row in an $N \times N$ co-occurrence matrix, with the dimensions being (in this case): `doc_id_s`, `user_id_s`, `query_s`. Alternatively, it can be viewed as a vertex of a given type (e.g. `doc_id_s`) in a graph, with the fields representing edges to other vertices of different types.

The EventMinerAggregator accumulates this co-occurrence data from events in sessions, per each entity type, and then periodically outputs co-occurrence documents when its internal LRU cache becomes full. It also flushes all remaining entries from the cache at the end of a job.

This helps the aggregator to limit the total memory budget, while keeping a reasonably long context of co-occurring entities.

The size of this internal cache is adjustable via a configuration parameter. A side-effect of this approach is that there may be multiple partial co-occurrence documents produced if related entities occur in a longer context than the LRU cache is able to handle - consumers of the aggregated data should be prepared to roll-up these multiple partial documents.

Aggregation job configuration

The groupingFields should use just `user_id_s`, and optionally the "sort" parameter should be set to `timestamp_tdt asc` - this way the sessionization process will work most efficiently. On the other hand, sorting by timestamp requires more work on the Solr-side, so it may be omitted, with the possible side-effect that there will be additional partial documents created.

EventMinerAggregator configuration parameters

These parameters are passed in the "params" property of the aggregator configuration:

- `maxSessionTime`: Optional integer, default value: 300. This specifies the maximum time (in seconds) to use for the definition of a session i.e. a series of actions by the same user in a given time period. Normally you want to keep this value fairly small, as events that occur close together in time are more likely to be related than those further apart.
- `maxElementsPerField`: Optional integer, default value: 10. Configures the maximum number of elements to store in each field of the aggregated documents.
- `maxCacheSize`: Optional integer, default value: 25000. This controls the maximum size of the internal LRU ("Least recently used") cache. Depending on the volume of data being processed, smaller values will result in more partial documents being created, while larger values will lead to a higher memory usage during the aggregation run.

Example Configuration

```
{
  "id": "event-miner-aggregation",
  "groupingFields": [
    "user_id_s"
  ],
  "signalTypes": [ ],
  "selectQuery": ":*:*",
  "sort": "timestamp_tdt asc",
  "timeRange": "[* TO NOW]",
  "aggregator": "em",
  "sourceRemove": false,
  "sourceCatchup": false,
  "outputRollup": false,
  "aggregates": [ ],
  "params": {
    "maxSessionTime": 600,
    "maxElementsPerField": 20
    "maxCacheSize": 10000
  }
}
```

Notes:

- The input documents will be grouped together based on their `user_id_s` values.
- The `selectQuery` is set to `.*:.*`, which means match all values in all fields when returning the initial set of input records for processing.
- The `timeRange` specifies all records with a timestamp up to NOW, where the latter will be set by the system to the time the aggregation job starts.
- The value for aggregator is set to `em`, which is a short label for the "EventMinerAggregator".
- The definition of the settings in the `params` section is as defined earlier.

Creating Aggregation Jobs

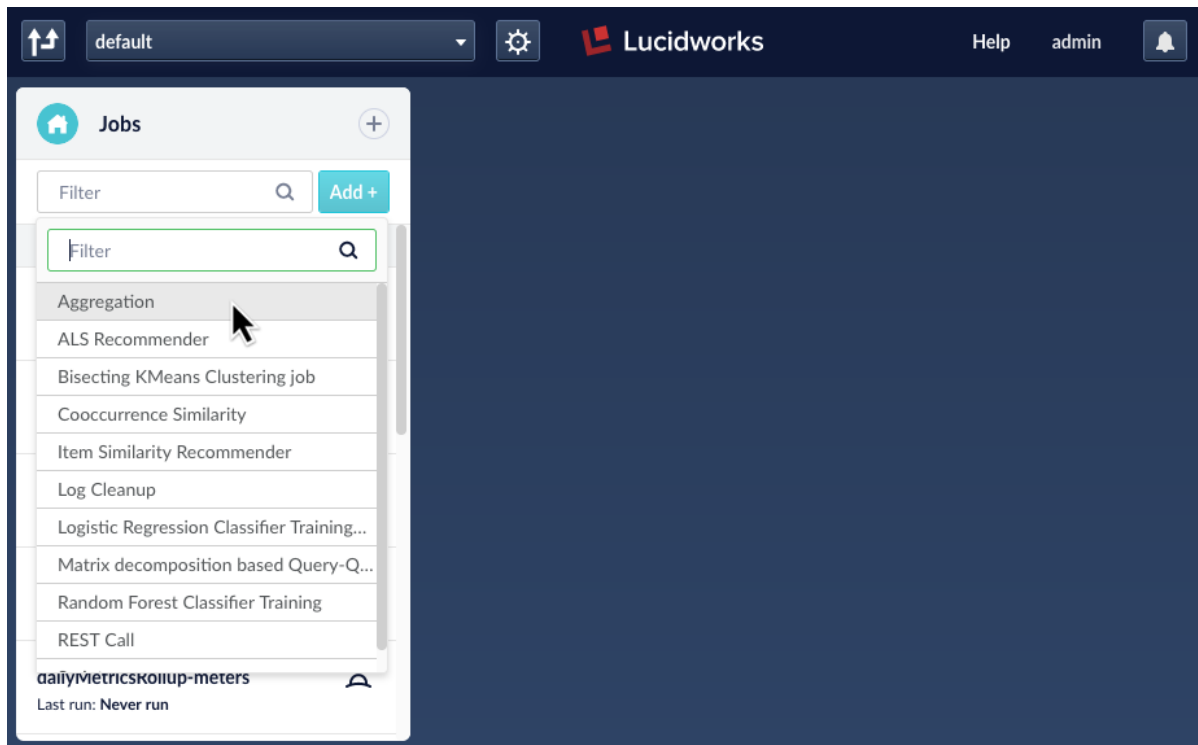
Aggregations are created automatically whenever you enable signals or recommendations. This topic explains how to create or modify aggregations individually. You can do this using the Fusion UI or the Jobs API.

Note	As of Fusion 3.1, the Aggregator API is deprecated in favor of the Jobs API.
------	--

Creating an aggregation job using the Fusion UI

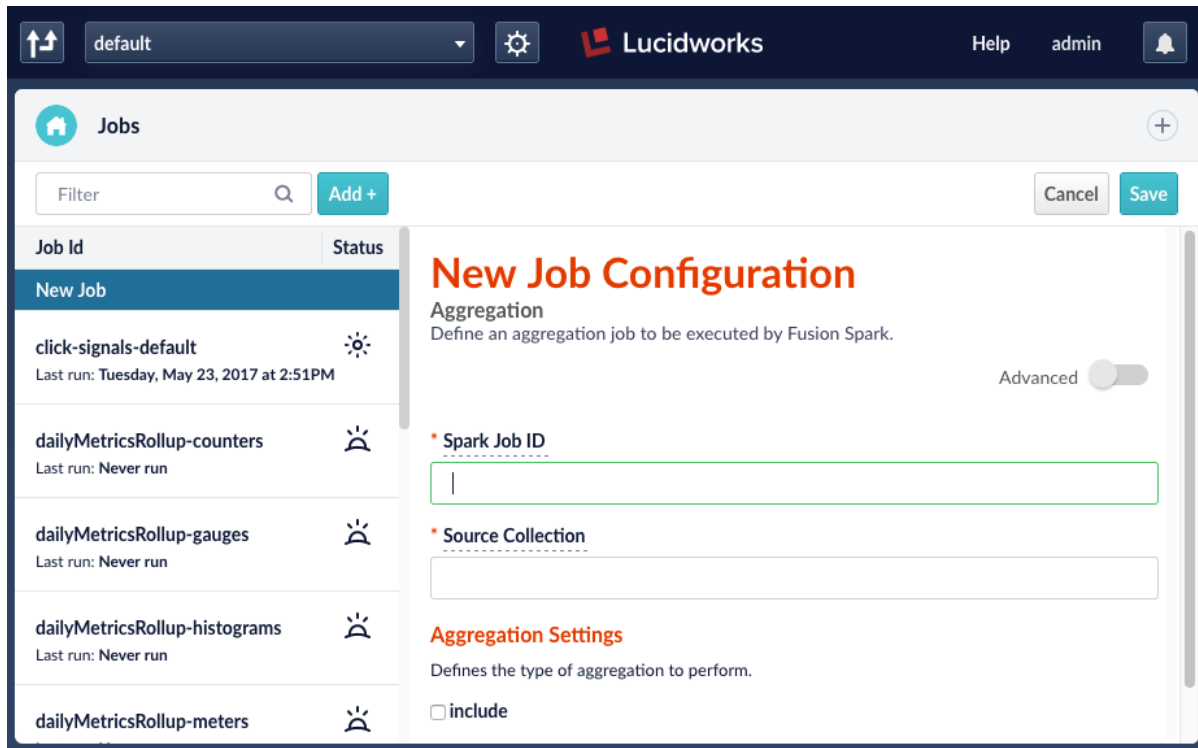
An aggregation is a type of job. Aggregation jobs can be created or modified at **Search > Jobs** in the Fusion UI.

1. Navigate to **Search > Jobs**.
2. Click **Add**.



3. Select **Aggregation**.

The New Job Configuration panel appears.



4. Enter an arbitrary Spark job ID.
5. Enter the name of the signals collection to be aggregated.

Note	Be sure to specify the signals collection (usually <code><primarycollectionname>_signals</code>), not the primary (<code><primarycollectionname></code>) collection.
------	---

6. Under Aggregation Settings, click **include**.
7. Configure the aggregation parameters as needed.

See Aggregation configuration parameters below for descriptions.

8. Click **Save**.

The new aggregation job appears in the jobs list. Now you can run it or schedule it.

Aggregation configuration parameters

<code>groupingFields</code>	An array of strings specifying the fields to group on.
<code>signalTypes</code>	The signal types. If not set then any signal type is selected.
<code>selectQuery</code>	The query to select the desired signals. If not set then <code>:</code> will be used, or equivalent.

<code>sort</code>	The criteria to sort on within a group. If not set then sort order is by ID, ascending.
<code>timeRange</code>	The time range to select signals on.
<code>outputPipeline</code>	What pipeline to use to process the output. If not set then <code>_system</code> pipeline will be used.
<code>rollupPipeline</code>	Pipeline to use for processing results of roll-up. This is by default the same index pipeline used for processing the aggregation results.
<code>rollupAggregator</code>	The aggregator to use when rolling up. If not set then the same aggregator will be used for roll-up.
<code>outputCollection</code>	The collection to write the aggregates to on output. This property is required if the selected output/rollup pipeline requires it (the default pipeline does). A special value of <code>-</code> disables the output.
<code>aggregator</code>	Aggregator implementation to use. This is either one of the symbolic names (<code>simple</code> , <code>click</code> , <code>em</code>) or a fully-qualified class name of a class extending <code>EventAggregator</code> . If not set then <code>'simple'</code> is used.
<code>sourceRemove</code>	If true, the processed source signals will be removed after aggregation. Default is false.
<code>sourceCatchup</code>	If true, only aggregate the signals since the last time the job was successfully run. If there is a record of such previous run then this overrides the starting time of time range set in <code>timeRange</code> property.
<code>outputRollup</code>	Roll-up current results with all previous results for this aggregation id, which are available in <code>outputCollection</code> .

<p>aggregates</p>	<p>List of functions defining how to aggregate events with results. Aggregation functions have these properties:</p> <ul style="list-style-type: none"> • type The function type defining how to aggregate events with results. • sourceFields The fields that the function will read from. • targetField The field that the function will write to. • mapper When true the function will be used in map phase only. • parameters Other parameters specific to individual functions.
<p>statsFields</p>	<p>List of numeric fields in results for which to compute overall statistics.</p>
<p>parameters</p>	<p>Other aggregation parameters (such as start / aggregate / finish scripts, cache size, and so on).</p>

Aggregator Functions

Aggregator Functions provide many ways to customize signals aggregations. These functions execute a specified operation on data coming from source event fields and accumulate the new value in a target field of the aggregated result.

Functions are implemented in a aggregator job definition, as a list within the `aggregates` property. Each function definition includes the function type, source fields, target fields, and additional parameters as needed for the function type. Specifically, each function takes the following properties (unless otherwise noted); additional parameters are noted in the function descriptions below.

- `type`: the function type.
- `sourceFields`: the list of fields from source events. Data will be retrieved from these fields as inputs to the function.
- `targetField`: the name of the target field where the aggregated result will be stored.
- `params`: any additional parameters for the specific function type, as described below.

The "sourceFields" and "targetField" field names in function specifications can be optionally prefixed with "event:" or "result:". If there are no prefixes the sourceFields take values from the current event being aggregated, and the targetField takes (or updates) the value in the current partial aggregated result. With these prefixes values can be processed and e.g. the original event can be updated, or event fields can be considered taking into account the accumulated values in the result.

Examples:

Override default input field source:

```
"sourceField": "result:tweet_split_ss"
```

Override default target field source:

```
"targetField": "event:tweet_split_ss"
```

Arithmetic Functions

Arithmetic functions operate on all valid numeric values (including string fields that are parseable into double numbers) from source fields and compute a single result to the target field.

sum

A sum of numeric values, as a double number.

Example:

```
{
  "type": "sum",
  "sourceFields": [ "count_i" ],
  "targetField": "sum_count_d"
}
```

sumOfLogs

A sum of natural logs of numeric value, as a double number.

Example:

```
{
  "type": "sumOfLogs",
  "sourceFields": [ "script_d" ],
  "targetField": "script_sum_logs_d"
}
```

sumOfSquares

A sum of squares of numeric value, as a double number.

Example:

```
{
  "type" : "sumOfSquares",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "sumOfSquares_position_d"
}
```

count

A count of source values, as a long number.

Example:

```
{
  "type": "count",
  "sourceFields": [ "id" ],
  "targetField": "count_d"
}
```

geoMean

A geometric mean of numeric values, as a double number.

Example:

```
{
  "type" : "geoMean",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "geoMean_position_d"
}
```

mean

An arithmetic mean of numeric values, as a double number.

Example:

```
{
  "type" : "mean",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "mean_position_d"
}
```

min

The minimum numeric value.

Example:

```
{
  "type" : "min",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "min_position_d"
}
```

max

The maximum numeric value.

```
{
  "type" : "max",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "max_position_d"
}
```

decay_sum

A sum of time-based exponentially decayed numeric values. The difference between the aggregationTime and the event time will be decayed using an exponential function with a half-life equaling 30 days.

This function has some additional properties:

- **halfLife**: the number of seconds for the half-life decay function.
- **timestampField**: the name of the field that contains the source event's timestamp. By default, this is `timestamp_dt`.
- **defaultWeight**: the weight of an event if all values from source fields are missing. The default is 0.1f, and this is expressed as a float.

Example:

```
{
  "type" : "decay_sum",
  "sourceFields" : [ "weight_d" ],
  "targetField" : "decay_sum_weight_d",
  "params" : { }
}
```

String Functions

String functions operate all values from source fields treated as strings.

cat

A concatenation of string values.

This function has some additional properties:

- **separator**: the character to use as a delimiter between values. The default is a single space.
- **maxLength**: the maximum length of the concatenated values (including separators). When this limit is exceeded, additional values are discarded. The default value is 10485760 characters (10 * 1024 * 1024).
- **maxValueCount**: the maximum number of values to concatenate. Any values collected after this limit are discarded. The default is 100.

Example:

```
{
  "type" : "cat",
  "sourceFields" : [ "user_id_s" ],
  "targetField" : "cat_user_id_txt",
  "params" : { }
}
```

ucat

A concatenation of unique string values.

This function has some additional properties:

- **separator**: the character to use as a delimiter between values. The default is a single space.
- **maxLength**: the maximum length of of the concatenated values (including separators). When this limit is exceeded, additional values are discarded. The default value is 10485760 characters (10 * 1024 * 1024).
- **maxValueCount**: the maximum number of values to concatenate. Any values collected after this limit are discarded. The default is 100.

Example:

```
{
  "type" : "ucat",
  "sourceFields" : [ "user_id_s" ],
  "targetField" : "ucat_user_id_txt",
  "params" : { }
}
```

split

A simple regex-based string splitting function.

The following function params are supported:

- `regex` - (string, required) a regular expression used for splitting.
- `lower` - (boolean, optional, false by default) after the `regex` has been applied the resulting parts are optionally lower-cased (using US locale).

Example:

```
{
  "type": "split",
  "sourceFields": [ "query_s" ],
  "targetField": "event:query_split",
  "params": {
    "regex": "\\s+",
    "lower": true
  }
}
```

In the example above, the raw signal event field "query_s" is first split on whitespace and then lower-cased, and the result is put back into the raw signal event field "query_split".

replace

A simple regex-based string replace. The `java.util.regex.Pattern` syntax is supported for the regex matching and replacement.

The following function params are supported:

- `regex` - (string, required) a pattern to match.
- `replace` - (string, required) replacement.

Example:

```
{
  "type": "replace",
  "sourceFields": [ "query_split" ],
  "targetField": "event:query_split_clean",
  "params": {
    "regex": "\\P{Alpha}+",
    "replace": "_"
  }
}
```

In the example above, this function takes the "query_split" values and replaces all non-alphabetic characters with underscores, and stores the result in the event's field "query_split_clean". As an extended example, this function would follow after the example **split** function and would add the field "query_split_clean" to the raw signal event. The "query_split_clean" field could be aggregated via other aggregation functions.

Collection Functions

Collection functions simply collect values from the source fields and add them as multiple values to the target field.

discard

This function discards all values from source fields and the target field. This modifies the source event and any in-progress aggregation result. This creates side-effects for subsequent functions, so should be used with care.

Example:

```
{
  "type" : "discard",
  "sourceFields" : [ "user_id_s" ],
  "targetField" : "collect_user_id_ss",
  "params" : { }
}
```

collect

Collect values from source fields.

This function has one additional property, 'maxValueCount', which defines the number of fields to collect from source fields. Any fields collected after this limit are discarded. The default is 100.

Example:

```
{
  "type" : "collect",
  "sourceFields" : [ "user_id_s" ],
  "targetField" : "collect_user_id_ss",
  "params" : { }
}
```

ucollectCollect unique values from source fields.

This function has one additional property, 'maxValueCount', which defines the number of fields to collect from source fields. Any fields collected after this limit are discarded. The default is 100.

Example:

```
{
  "type" : "ucollect",
  "sourceFields" : [ "user_id_s" ],
  "targetField" : "unique_user_id_ss",
  "params" : { }
}
```

Statistical Functions

Statistical functions compute scalar and matrix statistics. When the function has multiple results, such as for matrix or vector results, the data is stored in multiple fields.

varianceThe square of standard deviation of numeric values, as a double number.

Example:


```
{
  "type" : "covariance",
  "sourceFields" : [ "params.position_s", "position_rnd_1", "position_rnd_2" ],
  "targetField" : "cov_position_d",
  "params" : { }
}
```

stddev

The standard deviation of numeric values, as a double number.

Example:

```
{
  "type" : "stddev",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "stddev_position_d",
  "params" : { }
}
```

cardinality

An estimate of the number of unique elements in the set of values from source fields (which are treated as strings). This uses the HyperLogLog implementation.

This function has one additional property, 'error', which defines the acceptable probability of error from real value, specifically the standard deviation from real results. Smaller values cause exponentially higher RAM consumption during processing. For example, the default, 0.1, uses ~8Kb of RAM, while tests have shown 0.0001 uses ~64Mb.

Example:

```
{
  "type" : "cardinality",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "cardinality_position_l",
  "params" : { }
}
```

skewness

The measure of asymmetry of the distribution around its mean. This function is performed on numeric values and is expressed as a double number.

Example:

```
{
  "type" : "skewness",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "skewness_position_d",
  "params" : { }
}
```

kurtosis

The adjusted Pearson's kurtosis of numeric values, expressed as a double. This provides a comparison of the shape of the distribution to that of the normal distribution.

Example:

```
{
  "type" : "kurtosis",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "kurtosis_position_d",
  "params" : { }
}
```

quantiles

The quantiles of numeric values, stored as a double number in 0-N.targetField, or as a list of values in the target field (depending on the 'multivalued' property, described below). This implementation uses the T-Digest structure.

This function has the following additional properties:

- **quantiles**: the number of quantiles. The default is 10.
- **multiValued**: when true, all quantiles will be stored as multiple values in the target field. If false, then multiple values will be created in the format '0.targetField' to 'N.targetField'.

Example:

```
{
  "type" : "quantiles",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "quantiles_position_ss",
  "params" : {
    "multiValued" : true
  }
}
```

topk

An estimate of the top-K elements in the source fields and their frequency. The result is stored in three multi-valued fields, each with the same number of values. The three fields are:

- **counts.targetField**: integer counts (frequencies) of elements.
- **values.targetField**: elements.
- **errors.targetField**: estimation errors.

This function has one additional property, 'k', which is the number of elements to report. The default is 10.

Example:

```

{
  "type" : "topk",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "topk_position_ss",
  "params" : { }
}

```

covariance

A covariance matrix of numeric values from $N > 1$ source fields, with no smoothing. Missing or invalid values are treated as 0.0. A row of missing values is ignored. The resulting covariance matrix is stored in $N * (N - 1)$ fields following the naming pattern 'sourceField1.sourceField2.targetField'.

If source fields contain multiple values, only the first value from each source field will be used.

This implementation runs in a constant and small memory budget.

Example:

```

{
  "type" : "covariance",
  "sourceFields" : [ "params.position_s", "position_rnd_1", "position_rnd_2" ],
  "targetField" : "cov_position_d",
  "params" : { }
}

```

correlation

A correlation matrix of numeric values from $N > 1$ source fields. This implementation is based on the covariance function. The resulting correlation matrix is stored in $N * (N - 1)$ fields following the naming pattern 'sourceField1.sourceField2.targetField'.

Example:

```

{
  "type" : "correlation",
  "sourceFields" : [ "params.position_s", "position_rnd_1", "position_rnd_2" ],
  "targetField" : "corr_position_d",
  "params" : { }
}

```

histogram

An approximate histogram of values and their counts in source fields, using the T-Digest algorithm. Results are stored as corresponding multiple values in 'means.targetField' (for double values) and 'counts.targetField' (for integer values).

Example:

```
{
  "type" : "histogram",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "histogram_position_ss",
  "params" : { }
}
```

sigmoid

This function uses hyperbolic tangent (tanh) to limit the impact of source values according to an s-shaped curve. The following parameters control the shape of the curve:

- **weight** - controls the range of values. Default weight is 1.0, which means that the sigmoid function values will range between (-1, 1). E.g. weight = 2.0 means that values will range between (-2, 2).
- **intercept** - sets the constant shift of function values. Default is 0, which means that $\text{sigmoid}(0) = 0$ and $\text{sigmoid}(\text{Inf}) = 1$. E.g. intercept = 2.0 means that $\text{sigmoid}(0) = 2.0$ and $\text{sigmoid}(\text{Inf}) = 3.0$.
- **slope** - this parameter controls the slope of the function, i.e. how quickly it reaches saturation. Default value is 1.0. E.g. slope = 2 will cause the function to saturate quickly, slope = 0.1 will cause the function to saturate for larger values of source.
- **final** - boolean, default is true. This controls how the sigmoid is applied to the source value. First, all numeric values from source fields are summed. Then, if final = false the current sum is passed to the sigmoid function and added to the previous total. If final = true then the current sum is added to the total and the sigmoid function is applied only at the end of the aggregation.

Example:

```
{
  "type" : "sigmoid",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "sigmoid_position_ss",
  "params" : {
    "weight" : 2.0,
    "intercept" : 10.0,
    "slope" : 0.5,
    "final" : true
  }
}
```

Logical Functions

when

A logical function where processing will continue only if this function evaluates to true.

This function takes one additional property, 'expr', which is a JavaScript expression that must evaluate to a Boolean true/false. This property takes the same objects as the 'expr' function, described above. If this property is missing, the function will evaluate to true when any sourceField or targetField is present.

Example:

```
{
  "type" : "when",
  "sourceFields" : [ "params.position_s" ],
  "params" : {
    "expr" : "parseFloat(params_position_s) < 3"
  }
}
```

unless

A logical function where processing will continue only if this function evaluates to false.

This function takes one additional property, 'expr', which is a JavaScript expression that must evaluate to a Boolean true/false. This property takes the same objects as the 'expr' function, described above. If this property is missing, the function will evaluate to false when any sourceField or targetField is present.

Example:

```
{
  "type" : "unless",
  "sourceFields" : [ "params.position_s" ],
  "params" : {
    "expr" : "parseFloat(params_position_s) > 1"
  }
}
```

Scripting Functions

script

A scripted function. Scripts are evaluated as snippets, not as a function, and are expected to operate directly on the source event and the result. Their final values are discarded, since snippets in JavaScript are treated as expressions that evaluate to a specific value.

This function ignores the sourceFields and targetFields properties. Instead, the snippets are passed the following properties:

- startScript: the script defined is executed when the aggregation for the next unique tuple starts.
- aggregateScript: the script defined is executed for each source event.
- finishScript: the script is defined when all events for the current tuple have been processed and the result is about to be returned.

Example:

```

{
  "type" : "script",
  "sourceFields" : [ ],
  "params" : {
    "aggregateScript" : "result.addField('script_event_id_ss', event.getFieldValue('id'));"
  }
}, {
  "type" : "script",
  "sourceFields" : [ ],
  "params" : {
    "aggregateScript" : "event.addField('position_rnd_1', event.getFieldValue('params.position_s') + 1.0 -
Math.random());event.addField('position_rnd_2', event.getFieldValue('params.position_s') + 5.0 - Math.random()
* 10.0);"
  }
}

```

expr

A script expression. The script is evaluated as a snippet, and its final value is assigned to the targetField.

This function has only one additional property, 'expr', which contains the script expression.

Example:

```

{
  "type" : "expr",
  "sourceFields" : [ "query_s", "filters_s" ],
  "targetField" : "expr_s",
  "params" : {
    "expr" : "v = ''; if (value != null) v = value + ' '; v + query_s + '_' + filters_s"
  }
}, {
  "type" : "expr",
  "sourceFields" : [ "params.position_s" ],
  "targetField" : "expr_d",
  "params" : {
    "expr" : "v = 0; if (value != null) v = parseFloat(value); v + parseFloat(params_position_s)"
  }
}

```

Special Functions

noop

A function that does nothing (is non-operational). This is a fallback function when invalid function parameters or execution errors are encountered.

Example:

```

{
  "type" : "noop"
}

```

Aggregator Scripting

The aggregation jobs that process signals can be customized with using JavaScript. There are several options for scripts, and each option will be executed at a different point of the aggregation process. The options available at each stage of the process will vary, as explained for each option below.

Scripts are run after the main logic of the class they are customizing. This allows overriding default behavior of the class if needed.

The scripts are defined in the aggregation job using the Signals Aggregator API, with the 'params' property. Here is an example of declaring a script in an aggregator definition, using the the specialFields script option:

```
{
  "id" : "r1",
  "signalTypes" : [ "click" ],
  "selectQuery" : "*:*",
  "timeRange" : "[* TO NOW]",
  "params" : {
    "specialFields" : "unless_pos_gt_1_ss,when_pos_lt_3_ss"
  }
}
```

Note:

In many cases, the scripts defined will be executed many times during the aggregation job (i.e., for every event). For this reason, it's good practice to keep the scripts as simple as possible to avoid a negative impact on system performance. The `initScript` option includes a `"_context"` object that can be used for storing values that may require lengthy initialization or heavy computation.

initScript

A JavaScript defined with this option is executed when the signal aggregator instance (i.e., the specific aggregator job) is initialized. The following objects are available to the script:

- `logger`: an SLF4J Logger object.
- `aggregator`: the aggregator instance.
- `initArgs`: the initiation arguments.
- `_context`: the current scripting context. This can be used for storing small objects between executions of other scripted methods.

startScript

A script defined with this option is executed when a new tuple is about to be aggregated. All of the objects available to `initScript` are available to `startScript`, plus:

- `type`: the aggregation type, which is a string. Currently only the 'click' type is supported.
- `aggregationTime`: the reference point from which the aggregation is calculated, which is expressed in epoch time, an integer.
- `currentTuple`: a map of field names and values for the current tuple being aggregated.

aggregateScript

A script defined with this option is executed when a new event is being processed for the current tuple. All of the objects available to `initScript` and `startScript` are available, plus:

- `event`: the current event for aggregation. This is a `SolrDocument`.
- `result`: the aggregated result so far. This will also contain the original tuple fields. This is a `SolrDocument`.

If this script is present, it overrides the default logic for processing events. This means that the script must completely process the events as desired; it's not possible to build on existing rules. Note also that defining an `aggregateScript` will override any options defined as `specialFields`, described below.

It's possible to emit more than one result of aggregation for any given group of source events. This may be invoked in scripts, like the following snippet:

```
doc = $.prepareResult();
$.emit(doc);
```

"\$" is a reference to the current instance of aggregation function. The `prepareResult` method finishes calculations of some of the more complex functions (e.g. `topK`, `percentiles`, `correlation`, etc) and updates the result `PipelineDocument` (note: after this function is called the current "result" document is discarded, and a new `PipelineDocument` will be created to hold results of aggregating the following events, and the returned document can't be used for incremental calculations). An example use for this functionality would be to extract the month part of the date from a set of events which are sorted by timestamps, in order to produce aggregated results for every month within the current tuple defined by `groupingFields`.

finishScript

A script defined with this option is executed when all of the events for the current tuple have been processed and it's time to return the aggregated result. All of the objects available to `initScript`, `startScript` and `aggregator script` are available, plus:

- `result`: the final aggregated result. This is a `SolrDocument`.

specialFields

A script defined with this option uses a comma-separated, a whitespace-separated, or a JSON list of field names that are exempt from the default processing logic. These fields will **not** be processed in any way, which means they will not be included in the aggregated result.

If an `aggregatorScript` has been defined, it will be used instead of this option.

The default processing logic is as follows:

- skip any fields declared in `specialFields`;
- skip the event ID field (`id`);
- if the field value is a `Number`, then sum up all values as a `Double`;
- if field name ends with `'_s'` or `'_dt'`, retain only the first value and discard all other values (these dynamic fields are single-value only);
- otherwise add all values as-is to the result.

halfLife

This option allows defining a time period, in milliseconds, for the half-life decay formula. This formula is used when determining boost values for clicked documents: documents that have not been clicked in a longer period of time will not receive as high of a boost as documents that have been clicked more recently.

The default value is equivalent to 30 days (i.e., 2,592,000,000 ms).

weightScript

A script defined with this option is used when weighting the current event. It must evaluate to a numeric value, but has the following additional objects available:

- **event**: the current event for aggregation. This is a SolrDocument.
- **result**: the aggregated result so far. This will also contain the original tuple fields. This is a SolrDocument.
- **eventFlag**: the flag that indicates if the event is the result of a previous aggregation ("aggr") or is a new event ("event").
- **eventTime**: the timestamp of the event.
- **eventWeight**: the initial event weight, expressed as a float.
- **defaultWeight**: the default weight (if the script fails or the eventWeight is not entered properly), expressed as a float.
- **position**: the click position. This is 0-based, or 0 if not available. The data is retrieved from the 'params_position_s' field of the event.

2.4. Recommendations and Boosting

Apps can call Fusion to return documents for a user to interact with in various contexts. These include:

- **Search contexts:** A user provides a query and performs a search.
- **Non-search contexts:** A user *doesn't* perform a search, but rather lands on app page, for example, one that lists products or displays details for a single product. There is a query, but it comes from the app, not from the user.

The primary goal for the user and the business is *relevance*. Display the items that are most useful for the user, or for which the user is most likely to take some action, for example, to make a purchase.

In a basic search, a user enters search terms and Fusion uses an algorithm to rank the documents that it returns to the user. The matching and ranking depend solely on a match between the search terms and the contents of the documents (and possibly document metadata).

Recommendations are documents that a recommender stage in a query pipeline chooses to recommend to a user based on some criteria. But they aren't recommendations for a user yet. Whether the documents are recommended to the user depends on the remainder of the query pipeline.

Boosts are changes in numeric scores associated with documents that match a query. Changing the scores can move documents up or down in the query response. Positive boosts can bring documents into the response that wouldn't have been included without boosting. Negative boosts (which are actually fractional) can move documents out of the response.

Item similarities also play a role in some approaches for recommendations and boosting. When a user searches for an item, a search application can recommend and/or boost similar items.

2.4.1. Recommendation methods

There are two ways to get recommendations from Fusion: the Recommendations API or recommender stages for query pipelines.

Below is a more detailed comparison of the two methods:

	Recommendations API	Recommender query stages
Customization	Limited customization	Highly configurable

<p>"Cold start"</p>	<p>When no pertinent signals exist, no recommendations are returned.</p>	<p>When no pertinent signals exist, the query's normal search results are returned (if recommendations are applied as boosts). If recommendations are returned as the result, then no recommendations are returned.</p> <p>For the Recommend Items for User and Recommend Items for Items stages, the optional configuration setting Estimate Recent Results (enabled by default) estimates results since the last run of the recommender jobs.</p>
<p>Recommendation types</p>	<p>The Recommendations API provides the following kinds of recommendations:</p> <ul style="list-style-type: none"> • A list of items recommended for the given query (the <code>itemsForQuery</code> endpoint). • A list of queries associated with the given item, using aggregated clicks (the <code>queriesForItem</code> endpoint). 	<p>The query stages provide different kinds of recommendations and/or boosting:</p> <p>Content-based filtering:</p> <ul style="list-style-type: none"> • Boost Documents stage • Parameterized Boosting stage • Recommend More Like This stage <p>Collaborative filtering:</p> <ul style="list-style-type: none"> • Boost with Signals stage • Recommend Items for User stage • Recommend Items for Item stage

2.4.2. Using the Recommendations API

In order to use the REST API Recommendations service to get recommendations for items in some collection, that collection must have associated signals and aggregated-signals system collections. How good recommendations are depends on how well the information in the signals and aggregated signals collections, which is derived from observed user behavior, matches user behavior going forward.

See the the Recommendations API for details.


Note: Recommendations based on collaborative filtering aren't available through the Recommendations API. Instead, use the query-pipeline stages Recommend Items for User and Recommend Items for Item.

2.4.3. Enable Spark-based Collaborative Recommendations

Enable Recommendations in the Fusion UI enables Spark-based collaborative recommendations, that is, items-for-user and items-for-item recommendations (and users-for-item recommendations if you enable those). You don't need to give

this command to set up the recommendations infrastructure for items-for-query collaborative recommendations (Boost with Signals) or for content-based filtering.

To enable collaborative recommendations:

1. Log in to Fusion as the admin user.
2. From the Collections drop-down list, choose the collection for which you want to enable collaborative recommendations.
3. Click **Search**.
4. Click **Settings**  **> Enable Recommendations**.

Enabling recommendations creates these objects:

- **Collections:**

Object	Description
<code><collection>_signals</code>	Collection to hold user interactions. If the signals feature is enabled for this collection (which it is by default), then this collection will already exist.
<code><collection>_signals_aggr</code>	Collection to hold aggregated user interactions. If the signals feature is enabled for this collection (which it is by default), then this will already exist. The default aggregation groups together all signals of type click by the joint key (<code>user_id_s</code> , <code>doc_id_s</code>) from the <code><collection>_signals</code> collection, and performs a time-decayed count of these results to generate an estimate of the "implicit preference" that <code>user_id_s</code> has for <code>doc_id_s</code> .
<code><collection>_items_for_item_recommendations</code>	Collection to hold generated item-item similarities (by default 10 per item). No <code>user_id_s</code> data is present. A Recommend Items for Item query-pipeline stage can use the similarities to return item recommendations. For example, a query in which <code>doc_id_s = docA</code> would return an ordered list of other <code>doc_id_s</code> values for documents that are similar to document <code>docA</code> , along with the similarities. For example: <code>[("docB", 0.83), ("docC", 0.55), ("docD", 0.43), ..., ("docK", 0.22)]</code> .
<code><collection>_items_for_user_recommendations</code>	Collection to hold recommended items for a user. By default the job creates 10 recommendations per user.

- **Jobs:**

Object	Description
<code><collection>_item_aggregation</code>	Aggregates user-item interactions to generate weights. This aggregates user-item interactions and calculates a weight based on the recency of the interaction (so that more recent interactions have more impact on recommendations).
<code><collection>_item_recommendations</code>	Runs a Spark job to train an ALS-based (Alternating Least Squares) recommendation model, and then uses that model to generate recommendations. By default, this job generates both items-for-user recommendations and items-for-item recommendations. The job stores the results in the <code><collection>_items_for_user_recommendations</code> and <code><collection>_items_for_item_recommendations</code> collections.

- **Job schedules:**

Object	Description
Schedule for the job <code><collection>_item_aggregation</code>	Runs the aggregation job once a day.
Schedule for the job <code><collection>_item_recommendations</code>	Runs the recommender job after successful completion of the aggregation job.

- **Query pipelines:**

Object	Description
<code><collection>_items_for_user_recommendations</code>	Query pipeline to generate recommendations of items for a user
<code><collection>_items_for_item_recommendations</code>	Query pipeline to generate recommendations of items similar to an item

2.4.4. Disable Spark-based Collaborative Recommendations

Disable Recommendations in the Fusion UI disables Spark-based collaborative recommendations (items-for-user, items-for-item, and users-for-item).

When you disable recommendations, Fusion disables these schedules:


- Schedule for the job `<collection>_item_aggregation`
- Schedule for the job `<collection>_item_recommendations`

This stops the production of new collaborative recommendations. Query pipeline stages continue to return recommendations based on existing data. No data is deleted when you disable recommendations.

Disabling recommendations has no effect on items-for-query recommendations (Boost with Signals) or on the recommendations infrastructure for content-based filtering.

To disable collaborative recommendations:

1. Log in to Fusion as the admin user.

2. From the Collections drop-down list, choose the collection for which you want to disable collaborative recommendations.
3. Click **Search**.
4. Click **Settings**  **> Disable Recommendations**.

2.4.5. Content-based Filtering

Content-based filtering uses data about a user's search results, browsing history, and/or purchase history to determine which content to serve to the user. The filtering is non-collaborative, that is, related data from other users isn't used.

These query-pipeline stages use content-based filtering:

- Boost Documents stage
- Parameterized Boosting stage
- Recommend More Like This stage

2.4.6. Collaborative Filtering

Collaborative filtering lets your app take advantage of knowledge about the behavior of many individuals, in order to return the most useful documents to a user in a search or non-search context. It makes *serendipitous discovery* possible—a user is presented with items that other users deem relevant, for example, socks when buying shoes.

Parts of the Process

Collaborative filtering consists of these parts:

- **Use signals to collect data** – Collaborative filtering relies on user-engagement data (for example, clicks or purchases) that the app passes to Fusion as signals. Different kinds of collaborative filtering require different data.
- **Aggregate signal data** – Aggregate signal data so that the data can be used to compute recommendations.
- **Compute which documents to recommend** – Calculating recommended items using data from many users regarding many items can be time consuming. The calculations are done periodically by running jobs.
- **Recommend the documents at query time** – Document recommendations happen in a query pipeline, as boosts or as responses (document IDs). At query time, the documents to recommend are known, so the operation is fast.

Use Signals to Report Data to Fusion

Use signals to report user activity to Fusion. Signals contain data of interest. Signals generated from user-item interactions contain a `type` field that identifies the type of interaction, for example, a click or a purchase. They also contain these `params` fields:

- `query` – Either a search query or a context-query. This is used for items-for-query (Boost with Signals) recommendations, and might be used for items-for-user, items-for-item, and users-for-item recommendations.
- `userId` – A user ID. This can be anything that is associated with the user. A persistent identifier is needed, but the persistence can be limited. For example, a cookie can provide the user ID. The longer duration the persistence is, the better the recommendations. This is used for items-for-user, items-for-item, and users-for-item recommendations.
- `doc_Id` – Document ID. This is used for all types of collaborative recommendations.

Heterogeneous signals are OK. Signals that lack `query` won't contribute to items-for-query recommendations. Signals that lack `userId` won't contribute to items-for-user, items-for-item, or users-for-item recommendations.

Aggregate Signal Data

An aggregation job runs to aggregate signal data. These are the jobs:

- **Boost with Signals** – `<collection>_click_signals_aggregation`. The job places aggregated data in the collection `<collection>_signals_aggr`.
- **Recommend Items for User** – `<collection>_item_aggregation`. The job places aggregated data in the collection `<collection>_signals_aggr`.
- **Recommend Items for Item** – `<collection>_item_aggregation`. The job places aggregated data in the collection `<collection>_signals_aggr`.

Compute the Documents to Recommend

For items-for-query recommendations (Boost with Signals), the `<collection>_click_signals_aggregation` job produces the information needed to recommend documents at query time.

For the other collaborative recommendations, Fusion runs the `<collection>_item_recommendation` job automatically upon successful completion of the `<collection>_item_aggregation` job. The `<collection>_item_recommendation` job places recommendations in these collection:

`<collection>_items_for_user_recommendations` `<collection>_items_for_item_recommendations`

- **Items-for-user recommendations** (*produced by default*) – If you don't need items-for-user recommendations, delete the collection name `<collection>_items_for_user_recommendations` from the Items-for-users Recommendation text box.
- **Items-for-item recommendations** (*produced by default*) – If you don't need items-for-item recommendations (or item-item similarities), delete the collection name `<collection>_items_for_item_recommendations` from the Item-to-item Similarity Collection text box.
- **Users-for-item recommendations** (*not produced by default*)– To produce users-for-item recommendations, add the collection name `<collection>_users_for_item_recommendations` in the Users-for-items Recommendation Collection text box.

The `<collection>_item_recommendations` job can produce items-for-user, items-for-item, and/or users-for-item recommendations in a single (periodic) run. You don't need to set up multiple jobs.

Recommend Documents at Query Time

An app calls Fusion at query time to obtain documents to recommend. The query pipeline must contain a recommender pipeline stage, as well as other pipeline stages as needed.

Query-pipeline Stages

These query-pipeline stages provide collaborative filtering:

- Boost with Signals stage
- Recommend Items for User stage
- Recommend Items for Item stage

Built-in Query Pipelines

Fusion creates these query pipelines when you enable recommendations:


- `<collection>_items_for_user_recommendations` – Returns items-for-user recommendations.
- `<collection>_items_for_item_recommendations` – Returns items-for-item recommendations.

2.4.7. Item-Item Similarities

Item-item similarities are similarities between items. Fusion computes similarities based on user-item preferences. You can use item-item similarities for whatever you want. The most common use is as input data for Spark-based collaborative filtering.

By default, the `<collection>_item_recommendations` job is configured to produce item-item similarities and item-item recommendations.

To produce item-item similarities and item-item recommendations:

1. From the Fusion Launcher, click **Search** > **Home**  > **Jobs**.
2. Click the `<collection>_items_for_user_recommendations` job.
3. In the Item-to-item Similarity Collection text box, enter `<collection>_items_for_item_recommendations`.
4. Click **Save** to save the job configuration.

2.4.8. Use Multiple Recommendation/Boosting Stages

Fusion provides three query-pipeline stages for content-based filtering and three query-pipeline stages for collaborative filtering. All of these stages recommend documents.

Content-based filtering:

- Boost Documents stage
- Parameterized Boosting stage
- Recommend More Like This stage

Collaborative filtering:

- Boost with Signals stage
- Recommend Items for User stage
- Recommend Items for Item stage

Examples of Using Multiple Collaborative Recommendation Stages

In some cases, using multiple recommendation stages in the query pipeline might make sense. The `<collection>_items_for_user_recommendations` and `<collection>_items_for_item_recommendations` query pipelines are examples. In addition to a Recommend Items for User or Recommend Items for Item stages, each query pipeline also has a Boost with Signals stage. Thus, recommendations based on a Spark collaborative filtering computation are combined with recommendations based on the association of the query with items.

A Hybrid Recommender System

You can combine content-based filtering and collaborative filtering in the same query pipeline. A hybrid recommender system can leverage the capabilities and strengths of both approaches.

2.4.9. Step by Step


Fusion can leverage user behavior to generate recommendations and boosts based on the behavior of similar users. This is called collaborative filtering. The default objects and schedules that are created when you enable Spark-based collaborative recommendations can produce these types of recommendations:

- **Items for User** – Recommend items based on which items similar users have interacted with.
- **Items for Item** – Recommend items based on the similarity of the other items to the specified item. Similarity is based on user-item interactions.
- **Users for Item** – Recommend users for an item, for example to target advertising.

Here we give a step-by-step walkthrough of how to implement Items for User and Items for Item recommendations.

Step 1: Create a Collection to Hold the Data

Create a collection to contain the input data (signals about user-item interactions) and the data that Fusion generates to provide collaborative recommendations.


1. Click **Devops** > **Home**  > **Collections** > **New**.
2. Name the collection and click **Save Collection**.

Step 2: Enable Recommendations

Create the objects and schedules that Fusion needs to process the data for items-for-user, items-for-item, and users-for-item collaborative recommendations.

Enable Recommendations in the Fusion UI enables Spark-based collaborative recommendations, that is, items-for-user and items-for-item recommendations (and users-for-item recommendations if you enable those). You don't need to give this command to set up the recommendations infrastructure for items-for-query collaborative recommendations (Boost with Signals) or for content-based filtering.

To enable collaborative recommendations:

1. Log in to Fusion as the admin user.
2. From the **Collections** drop-down list, choose the collection for which you want to enable collaborative recommendations.
3. Click **Search**.
4. Click **Settings**  > **Enable Recommendations**.

Enabling recommendations creates these objects:

- **Collections:**

Object	Description
<code><collection>_signals</code>	Collection to hold user interactions. If the signals feature is enabled for this collection (which it is by default), then this collection will already exist.

Object	Description
<code><collection>_signals_aggr</code>	<p>Collection to hold aggregated user interactions. If the signals feature is enabled for this collection (which it is by default), then this will already exist.</p> <p>The default aggregation groups together all signals of type click by the joint key (<code>user_id_s</code>, <code>doc_id_s</code>) from the <code><collection>_signals</code> collection, and performs a time-decayed count of these results to generate an estimate of the "implicit preference" that <code>user_id_s</code> has for <code>doc_id_s</code>.</p>
<code><collection>_items_for_item_recommendations</code>	<p>Collection to hold generated item-item similarities (by default 10 per item). No <code>user_id_s</code> data is present. A Recommend Items for Item query-pipeline stage can use the similarities to return item recommendations. For example, a query in which <code>doc_id_s = docA</code> would return an ordered list of other <code>doc_id_s</code> values for documents that are similar to document <code>docA</code>, along with the similarities. For example: <code>[("docB", 0.83), ("docC", 0.55), ("docD", 0.43), ..., ("docK", 0.22)]</code>.</p>
<code><collection>_items_for_user_recommendations</code>	<p>Collection to hold recommended items for a user. By default the job creates 10 recommendations per user.</p>

- **Jobs:**

Object	Description
<code><collection>_item_aggregation</code>	<p>Aggregates user-item interactions to generate weights. This aggregates user-item interactions and calculates a weight based on the recency of the interaction (so that more recent interactions have more impact on recommendations).</p>
<code><collection>_item_recommendations</code>	<p>Runs a Spark job to train an ALS-based (Alternating Least Squares) recommendation model, and then uses that model to generate recommendations. By default, this job generates both items-for-user recommendations and items-for-item recommendations. The job stores the results in the <code><collection>_items_for_user_recommendations</code> and <code><collection>_items_for_item_recommendations</code> collections.</p>

- **Job schedules:**

Object	Description
Schedule for the job <code><collection>_item_aggregation</code>	Runs the aggregation job once a day.
Schedule for the job <code><collection>_item_recommendations</code>	Runs the recommender job after successful completion of the aggregation job.

- **Query pipelines:**

Object	Description
<code><collection>_items_for_user_recommendations</code>	Query pipeline to generate recommendations of items for a user
<code><collection>_items_for_item_recommendations</code>	Query pipeline to generate recommendations of items similar to an item

Step 3: Collect User-Item Interaction Data

Both items-for-user and items-for-item recommendations are based on user-item interactions, that is, on some interactions that users have with items. For example, users might look at the items, click on the items, or buy the items. User-item interactions often reflect user-item preferences.

Your app must collect user-item interaction data and use signals to get the data to Fusion. Data collection is an ongoing process. The more data, the better; and current data provides the best recommendations.

If you have collected user-item interaction data previously, you can load that data into Fusion. This is one way to address the cold-start problem—without user-item interaction data, the algorithms lack the data they need to make recommendations.


Use signals to report user activity to Fusion. Signals contain data of interest. Signals generated from user-item interactions contain a `type` field that identifies the type of interaction, for example, a click or a purchase. They also contain these `params` fields:

- `query` – Either a search query or a context-query. This is used for items-for-query (Boost with Signals) recommendations, and might be used for items-for-user, items-for-item, and users-for-item recommendations.
- `userId` – A user ID. This can be anything that is associated with the user. A persistent identifier is needed, but the persistence can be limited. For example, a cookie can provide the user ID. The longer duration the persistence is, the better the recommendations. This is used for items-for-user, items-for-item, and users-for-item recommendations.
- `doc_id` – Document ID. This is used for all types of collaborative recommendations.

Heterogeneous signals are OK. Signals that lack `query` won't contribute to items-for-query recommendations. Signals that lack `userId` won't contribute to items-for-user, items-for-item, or users-for-item recommendations.

Step 4: Configure Recommender Jobs

The jobs `<collection>_item_aggregation` and `<collection>_item_recommendations` have default configurations that produce both items-for-user and items-for-item recommendations as boosts.

Configure jobs in **Search** > > **Home**  > **Jobs**. Click the job you want to configure. Click **Save** to save the configuration.

Job `<collection>_item_aggregation`

You shouldn't need to configure the `<collection>_item_aggregation` job.

Job `<collection>_item_recommendations`

Options for job configuration are:


- **Training Data Sampling Fraction** – A value between 0 and 1. Specify a decimal fraction to use that fraction of the data when training. The job will run more quickly and use less memory. Consider reducing this value during testing. For production environments, it should be set to 1.

- **Produce items-for-user recommendations and item-item similarities** (*produced by default*) – Fusion produces item-item similarities when it produces item-item recommendations. You get both or neither. To get both, enter `<collection>_items_for_item_recommendations` in the Item-to-item Similarity Collection text box. To get neither, delete the collection name `<collection>_items_for_user_recommendations` from the Items-for-users Recommendation text box.
- **Produce items-for-item recommendations** (*produced by default*) – If you don't need items-for-item recommendations (or item-item similarities), delete the collection name `<collection>_items_for_item_recommendations` from the Item-to-item Similarity Collection text box.
- **Produce users-for-item recommendations** (*not produced by default*) – To produce users-for-item recommendations, add the collection name `<collection>_users_for_item_recommendations` in the Users-for-items Recommendation Collection text box.
- **Compute more or fewer recommendations** – The default numbers for items-for-user, items-for-item, and users-for-item recommendations are all 10. Increasing the number *doesn't* have a large cost.
- **Recommender Rank**: You might need to tune this (make it bigger). There is a memory cost to doing so.
- **Grid Search Width** (Advanced tab) – Set to 1 to have Spark estimate the best parameters. After one run with the value set to 1, you can set it back to 0.
- **Implicit Preferences** – If checked, Fusion uses user actions to guess user preferences. Uncheck if you have user preference data (for example, ratings).

Step 5: Run a Job to Aggregate Signal Data


Aggregate user-item interaction data to create an aggregation document for each user-item pair. The aggregated data is equivalent to a user-item interaction matrix, which is the input for the ALS Recommender Spark job.

Enabling collaborative recommendations creates the `<collection>_item_aggregation` job and schedules it to run once a day.

You can also run the `<collection>_item_aggregation` job manually. Navigate to **Devops** > **Home**  > **Scheduler**. Click the `<collection>_item_aggregation` job, and then click **Start**. Click **Cancel** to exit scheduling for the job.

Step 6: Run a Job to Create Collaborative Recommendations

Note: By default, this job is scheduled to run automatically upon successful completion of the `<collection>_item_aggregation` job.

To run the `<collection>_item_recommendations` job manually, navigate to **Devops** > **Home**  > **Scheduler**. Click the `<collection>_item_recommendations` job, and then click **Start**. Click **Cancel** to exit scheduling for the job.

This job populates several collections:

- `<collection>_items_for_item_recommendations` – N recommended items for each item
- `<collection>_items_for_user_recommendations` – N recommended items for each user
- `recommender_models` (*not produced by default*) – Weights for each user-item pair. This collection can contain multiple models for different data collections, which are tagged by `model_id`. By default, subsequent runs of a collaborative-filtering recommender job reuse a previously generated model, if possible, for better performance.

Step 7: Set Up Query Pipelines

By default, enabling collaborative recommendations sets up two query pipelines that an app can use to obtain

recommended items. Alternatively, you can add corresponding query-pipeline stages to other query pipelines.

Use Built-in Query Pipelines for Collaborative Recommendations

When you enable collaborative recommendations, Fusion creates two query pipelines that you can use for retrieving the recommendations:

- `<collection>_items_for_user_recommendations` – This pipeline retrieves items to recommend for the user specified by the `userId` query parameter, based on which items similar users interact with. As its first stage, it has a Recommend Items for User stage. By default, results from that stage are applied as boosts.
- `<collection>_items_for_item_recommendations` – This pipeline retrieves items to recommend for the item specified by the `docId` query parameter, based on the similarity of the other items to the specified item. Similarity is based on user-item interactions, for example, users who clicked on item A also tended to click on item Z. As its first stage, it has a Recommend Items for Item stage. By default, results from this stage are applied as boosts.

Add Collaborative-Recommendation Stages to Other Query Pipelines

As an alternative to using the built-in query pipelines for collaborative recommendations, you can add stages for collaborative recommendations to other query pipelines:

- **Recommend Items for User** – Add this stage to recommend items for users, based on what items similar users interact with. By default, results from that stage are applied as boosts.
- **Recommend Items for Item** – Add this stage to recommend items based on the similarity of the other items to the specified item. Similarity is based on user-item interactions, for example, users who clicked on item A also tended to click on item Z. By default, results from this stage are applied as boosts.

Include Required Stages

Depending on the source of the query pipeline and the default stages, you might need to add or remove query-pipeline stages. These are examples:

Use Items-for-User and Items-for-Query Recommendations to Boost Items that Match a Query

Use the pipeline stages present in the built-in pipeline `<collection>_items_for_user_recommendations` – Recommend Items for User, Boost with Signals, Query Fields, Field Facet, and Solr Query.

Use Items-for-Item and Items-for-Query Recommendations to Boost Items that Match a Query

Use the pipeline stages present in the built-in pipeline `<collection>_items_for_item_recommendations` – Recommend Items for User, Boost with Signals, Query Fields, Field Facet, and Solr Query.

Step 8: Look Up Precomputed Recommendations at Query Time

Get collaborative recommendations for a user at query time. Query the pipelines that you set up in Step 7.

2.5. Search Applications

Ultimately, Fusion is the back end for your own search applications.

- Your application uses Fusion’s REST API to interact with the Fusion system. The REST API supports all the features available in the Fusion UI. At a minimum, your application will employ the `/query-pipelines/{id}/collections/{collection}/{handler}` endpoint to query Fusion collections.
- *Recommendations* are a way to use aggregations to enhance the search experience. Based on the current search, or signals collected previously, Fusion can return results that are relevant in the end user’s current context.
- Certain front-end features require some Fusion configuration:
 - Autocomplete
 - Faceting
 - Stopwords
 - Synonyms
 - DateTime processing

2.5.1. Autocomplete

Autocomplete, also known as auto-suggest, look-ahead, or type-ahead, is a feature that displays a list of common search terms that begin with, or are contained by, the query string. For example, the query "search technology" might result in an autocomplete list with results "search technology rocks", "i love search technology", and/or "this is a search technology company", depending on how the autocomplete has been configured. Like most other search components, Fusion leverages Solr's [autosuggest](#) component. Assuming auto-suggest is configured in the Solr config, it is sufficient to enable it in Fusion in order to use the component.

How to configure the suggest component:

1. Choose a suggester implementation from the Solr Suggester documentation and define it in the solrconfig.xml file.

In this example, a suggester called mySuggester is created in solrconfig.xml using the AnalyzingInfixLookupFactory and defining 'description' as the field to be suggested on:

```
<searchComponent name="suggest" class="solr.SuggestComponent">
  <lst name="suggester">
    <str name="name">mySuggester</str>
    <str name="lookupImpl">AnalyzingInfixLookupFactory</str>
    <str name="dictionaryImpl">DocumentDictionaryFactory</str>
    <str name="field">description</str>
    <str name="suggestAnalyzerFieldType">string</str>
    <str name="buildOnStartup">>false</str>
  </lst>
</searchComponent>
```

1. Ensure that the field used for autosuggest is properly defined in the schema.

```
<field name="description" type="text_general" stored="true" indexed="true" multivalued="false" />
```

1. Once modified, the custom configuration files can be uploaded to Zookeeper using [Zookeeper's command line interface](#).

The zkcli command `putfile` can be used to replace an existing ZK configuration file. Here `CONFIG_NAME` will likely represent the name of the collection, otherwise the name of a custom configuration that has been defined.

```
{fusion_path}/apps/solr-dist/server/scripts/cloud-scripts/zkcli.sh -z localhost:2181 -cmd putfile
/configs/CONFIG_NAME/solrconfig.xml /Desktop/solrconfig.xml
```

1. Enabling autosuggest in Fusion requires adding the "suggest" request handler to the AllowedRequestHandlers in the "Solr Query" stage of the query pipeline. The suggester should then be built using `.../suggest?suggest.build=true`.

Ex: <http://localhost:8764/api/apollo/query-pipelines/PIPELINE/collections/COLLECTION/suggest?suggest.buld=true>

1. Start searching!

<http://localhost:8764/api/apollo/query-pipelines/PIPELINE/collections/COLLECTION/suggest?suggest.q=ca...>

2.5.2. DateTime Processing

DateUtils - Uniform API for Date Parsing and Formatting

Date and time processing is difficult due to the complexity of rules (and exceptions from rules) specific to historical changes, calendar systems, locales, time zones, and DST rules (daylight saving time).

Fusion uses `com.lucidworks.apollo.common.util.DateUtils` for date / time parsing and formatting, and developers are strongly encouraged to use this class for parsing date formats.

The `DateUtils` class uses the [Joda Time](#) library, which was the basis for the `java.util.time` API in JDK 8 but is compatible with earlier versions of Java. It add support for parsing abbreviated time zone names (e.g. PST). Although use of abbreviated time zone names has been deprecated because many of them are ambiguous, they are still in wide use. It also supports parsing full time zone names in any letter case (Joda Time accepts only canonical mixed-case names, e.g. "America/New_York"). Robust identification of time zone is helpful for reasoning about time intervals, because time zone rules cover phenomena like DST changes with gap and overlap hours, leap seconds, administrative changes to offsets, etc, etc, for which a simple offset from UTC is insufficient.

Supported Formats

Supported formats for date / time parsing can be divided into three disjoint groups:

- ISO 8601 formats ("Solr formats"), represented as `yyyy-MM-dd'T'HH:mm:ss.SSS'Z'`. The milliseconds part is optional, with precision varying between 0-3 digits.
- Fusion "zoned format" - a modification of a common Internet format that specifies day of week and uses full names for time zones, represented as `EEE yyyy-MM-dd HH:mm:ss.SSS ZZZ` in US locale. This format is useful for storing and easy processing of date/time with zone, as it's unambiguous, exact and easy to parse.
- Epoch formats - represented in Java API as objects of `java.util.Date` or the number of seconds, or the number of milliseconds since epoch (either as a number or as a String). These formats by definition don't contain any time zone information, and if permitted they are treated as absolute instants in the default time zone (see below).
- Common global formats - i.e. formats that explicitly specify the timezone.
- Common local formats - i.e. formats that don't specify the timezone.

The definition of formatting symbols used below can be found in the Joda Time [DateTimeFormat](#) documentation.

ISO 8601 / Solr Formats

See <https://cwiki.apache.org/confluence/display/solr/Working+with+Dates>

```
"yyyy-MM-dd'T'HH:mm:ss'Z'", // Solr format without milliseconds
"yyyy-MM-dd'T'HH:mm:ss.SSS'Z'", // standard Solr format, with literal "Z" at the end
"yyyy-MM-dd'T'HH:mm:ss.SS'Z'", // standard Solr format, with literal "Z" at the end
"yyyy-MM-dd'T'HH:mm:ss.S'Z'" // standard Solr format, with literal "Z" at the end
```

Common Global Formats

```

"EEE yyyy-MM-dd HH:mm:ss.SSS zzz",
"yyyy-MM-dd'T'HH:mm:ss.SSSZ", // with numeric +-HHmm timezone at the end
"yyyy-MM-dd'T'HH:mm:ss.SSSZ", // with numeric +-HH:mm timezone at the end
"yyyy-MM-dd'T'HH:mm:ss.SSSz", // with symbolic XXX timezone at the end
"yyyy-MM-dd'T'HH:mm:ssz", // with symbolic XXX timezone at the end
"yyyy-MM-dd'T'HH:mm:ssZ", // with offset
"EEE MMM d HH:mm:ss z yyyy",
"EEE MMM d HH:mm:ss Z yyyy",
"EEE MMM d HH:mm:ss z yyyy",
"EEE MMM d HH:mm:ss.SSS z yyyy",
"EEE, dd MMM yyyy HH:mm:ss zzz", // RFC 1123, with either short or full time zone
"EEEE, dd-MMM-yy HH:mm:ss zzz", // RFC 1036
"yyyy-MM-dd HH:mm:ss Z",
"yyyy-MM-dd HH:mm:ss ZZ",
"yyyy-MM-dd HH:mm:ss z",
"yyyy-MM-dd HH:mm:ss.SSS Z",
"yyyy-MM-dd HH:mm:ss.SSS ZZ",
"yyyy-MM-dd HH:mm:ss.SSS z",
"yyyy-MM-dd HH:mm:ss zzz", // with full time zone (e.g. America/New_York)
"yyyy-MM-dd'T'HH:mm:ss'GMT'Z", // with literal "GMT" and offset
"yyyy-MM-dd'T'HH:mm:ss.SSS'GMT'Z", // with literal "GMT" and offset
"yyyy-MM-dd'T'HH:mm:ss'UTC'Z", // with literal "UTC" and offset
"yyyy-MM-dd'T'HH:mm:ss.SSS'UTC'Z", // with literal "UTC" and offset
"yyyy-MM-dd HH:mm:ss 'UTC'Z",
"yyyy-MM-dd HH:mm:ss.SSS 'UTC'Z",
"yyyy-MM-dd HH:mm:ss 'GMT'Z",
"yyyy-MM-dd HH:mm:ss.SSS 'GMT'Z"

```

Note: Joda-Time cannot parse ZZZ when zone name is in incorrect case (e.g. "EUROPE/WARSAW" will fail), for this reason we use zzz which accepts any letter case.

Common Local Formats

```

"yyyy-MM-dd'T'HH:mm:ss",
"yyyy-MM-dd",
"yyyy-MM-dd hh:mm:ss",
"yyyy-MM-dd HH:mm:ss",
"yyyy-MM-dd HH:mm:ss.SSS",
"EEE MMM d HH:mm:ss yyyy", // ANSI C
"EEE MMM d HH:mm:ss.SSS yyyy" // ANSI C

```

Missing parts of the data are filled in with defaults - e.g. time zone is filled in with the default time zone (see below), missing time data is filled with 00:00:00.000

DateUtils Usage

Instances of DateUtils can be created with the following arguments:

- requireTimezone - (boolean, required) when this argument is true then only patterns with timezone component are accepted (unless custom patterns are provided). Epoch formats are not accepted. If false then any recognizable pattern is used.
- formats - (list of strings, optional) a list of formats to use instead of the built-in ones.

- locale - (string, optional) locale to use for parsing, or null for the platform default locale.
- defaultTimeZone - (string, optional) time zone name to use as default (when using local formats), or null for the platform default time zone. Time zone names can be provided in short or long form, or as a fixed offset [-+]HH:mm. Further information is provided in the Javadoc of the class. The default instance of DateUtils uses common global formats (with ISO8601 and epoch formats), requireTimezone == true, locale "en-US" and time zone "UTC".

2.5.3. Stopwords Files

Fusion collections are Solr collections which are managed by Fusion. Solr itself manages a set of [resources](#) for a collection. Stopword lists are one such resource.

The Fusion UI provides a Stopwords manager tool which is reached from the "Stopwords" tab in the "Configuration" section of the collection home panel. Clicking on this tool allows you to view the contents of all configured stopwords files in an editable browser.

2.5.4. Synonyms Files

Fusion collections are Solr collections which are managed by Fusion. Solr itself manages a set of [resources](#) for a collection. These resources include stopword lists and synonyms.

There are two ways of specifying synonyms: as a set of interchangeable tokens, or as a mapping from one token to one or more other tokens. A set of interchangeable tokens is specified as a series of comma-separated words:

```
Television, Televisions, TV, TVs
```

This will be used to expand any document or query by expanding any single token in this set to all tokens in this set.

A mapping from one token to one or more other tokens is specified as follows:

```
pixima => pixma
```

In this case, tokens on the left hand side of the rule will be replaced by the tokens on the right hand side.

The "Configuration" section of the collection home panel contains options for managing each of these resources from inside Fusion.

You can view this file as a text file from the "Solr Config" panel:

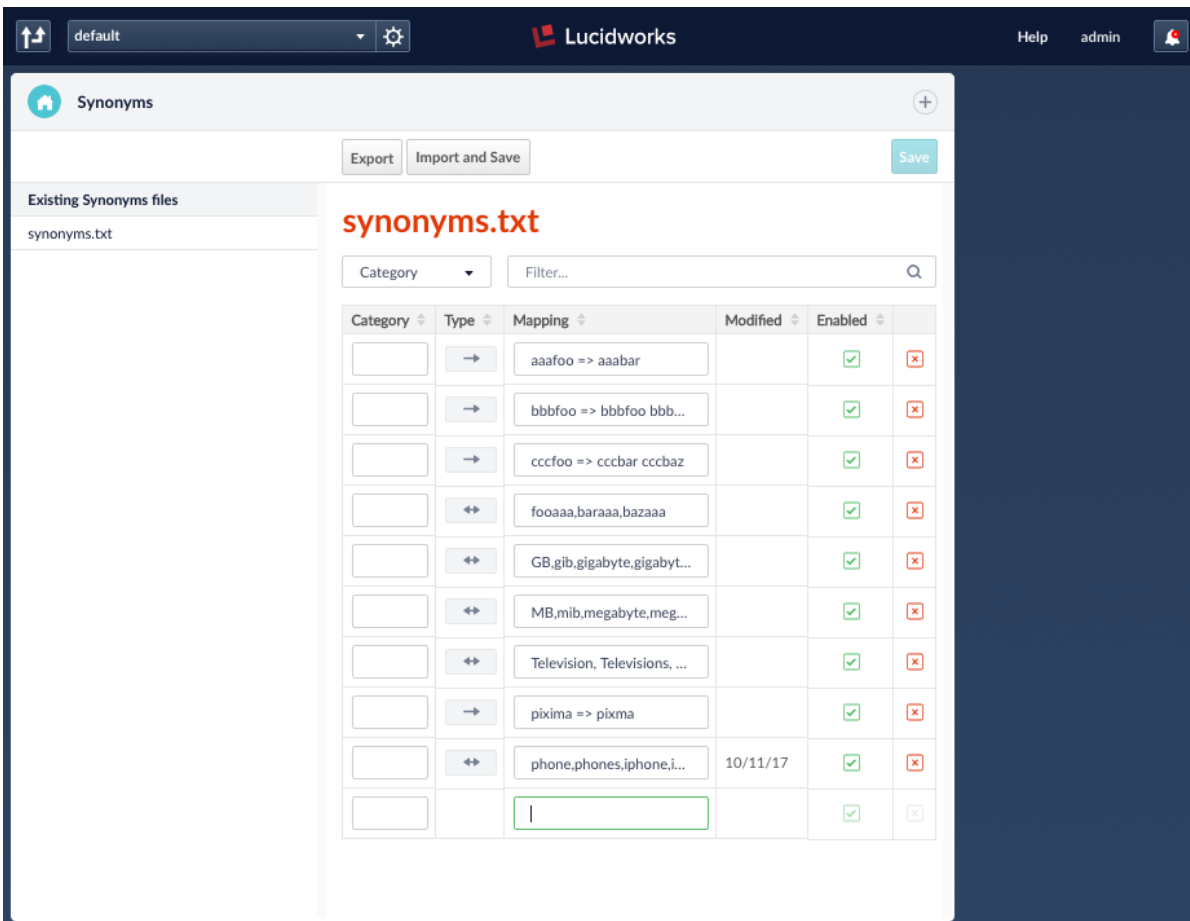
```
#some test synonym mappings unlikely to appear in real input text
aaafoo => aaabar
bbbfoo => bbbfoo bbbbar
cccfoo => cccbar cccbaz
fooaaa,baraaa,bazaaa

# Some synonym groups specific to this example
GB,gib,gigabyte,gigabytes
MB,mib,megabyte,megabytes
Television, Televisions, TV, TVs
#notice we use "gib" instead of "GiB" so any WordDelimiterFilter coming
#after us won't split it into two words.

# Synonym mappings can be used for spelling correction too
pixima => pixma
```

Synonyms Manager

The Fusion UI provides a Synonyms manager tool which is reached from the "Synonyms" tab in the "Configuration" section of the collection home panel. Once opened, the Synonyms manager panel will display the synonyms.txt as a series of editable objects, 1 per line.



Moving the cursor out of the editor field will add the item, clicking "Save" in the upper right corner will apply it to the current collection immediately, storing the config file in the correct Fusion folder. Choosing "Export" will download a copy via the browser to the local user computer. "Import and Save" reads a local file into Fusion.

An additional feature of the Synonym manager in the Fusion UI is that label terms can be used to restrict the display of synonyms to a browsable subset by adding a "category" label to the synonym set or mapping rule. This category label is a convenience for working with data in the Fusion UI only. The information is not saved in the synonyms file used by Solr.

Chapter 3. Analytics

Dashboards are Fusion's built-in analytics tool. These sections explain how to use, create, and manage dashboards:

- [About Dashboards](#)
- [Use Dashboards](#)
- [Create Dashboards](#)
- [Input Panels](#)
- [Display Panels](#)
- [Manage Dashboards](#)

To use third-party data analytics tools with Fusion, see the [Catalog API](#).

3.1. Get Started with Fusion Dashboards

In this tutorial we build an analytics dashboard over signal data.

The Fusion distribution includes a directory named "fusion/3.1.x/examples/signals", where "fusion/3.1.x" is the top-level directory of the Fusion distribution, which contains:


- The data file `signals.json`, a synthetic dataset of 20,000 click-signal events, based on a set of Best Buy query logs from 2011. This file contains a list of JSON objects, where each object contains information about a search query, a set of search categories, and the item ultimately clicked on.
- The script file `signals.sh`, which loads the raw signal data and then runs aggregations.

First we load the raw signal data into a Fusion collection. Then we create a new time-series dashboard and populate it with input and display panels.

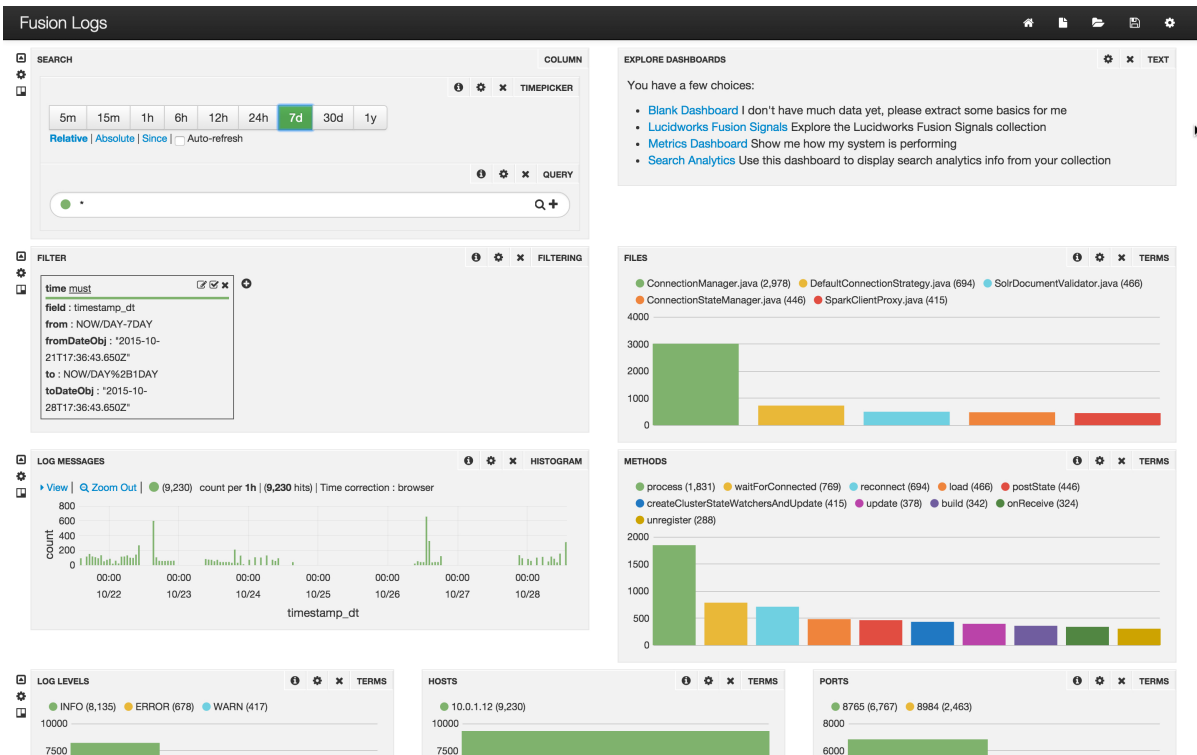
3.1.1. Dashboard Controls

Here's how to interact with the Dashboard interface.


You can launch the Dashboards tool from the Fusion UI, via the Fusion Applications menu:



When launched from the Fusion UI, the Dashboards tool appears in a new tab titled Banana 3. The title changes to the name of the dashboard. The initial dashboard is the built-in Fusion Logs dashboard:



The left side of the top menu bar displays the Dashboard title. The right side of the top menu contains a set of controls:



- The home icon is the "Go to Default Dashboard" control. The initial default dashboard is the "Fusion Logs" dashboard.
- The sheet of paper icon opens the "Create Dashboard" dialog.
- The folder icon opens the "Load Dashboard" dialog.
- The diskette icon opens the "Save Dashboard" dialog.
- The gear icon opens the "Configure Dashboard" dialog for the current dashboard, which has tabs:
 - "General" : title and page style
 - "Rows" : for adding, deleting and arranging the rows
 - "Controls" : loading, save, and export options
 - "Solr" : configure Solr server and set global query parameters.

3.1.2. Create the Raw Signals Collection

In this example, we create a Log Analytics Dashboard using a synthetic log dataset. This file contains a list of JSON objects, where each object contains information about a search query, a set of search categories, and the item ultimately clicked on.

The first JSON object in this dataset is:

```
"timestamp": "2011-09-01T23:44:52.533Z",
"params": {
  "query": "Televisiones Panasonic 50 pulgadas",
  "docId": "2125233",
  "filterQueries": [
    "cat00000",
    "abcat0100000",
    "abcat0101000",
    "abcat0101001"
  ]
},
"type": "click"
```

The top-level attributes of this object are:

- **type** - A required field for all signals. In the example dataset, all signals are of type click.
- **timestamp** - The time at which this click was logged.
- **params** - This attribute contains a set of key-value pairs for search-query event information.

In this dataset, the information captured includes the free-text search query entered by the user, the document id of the item clicked on, and the set of Best Buy site categories that the search was restricted to. These are codes for categories and sub-categories such as "Electronics" or "Televisions".

The example script `signals.sh` loads the raw signal via a POST request to the Fusion REST API endpoint: `/api/apollo/signals/<collectionName>` where `<collectionName>` is the name of the primary collection itself.

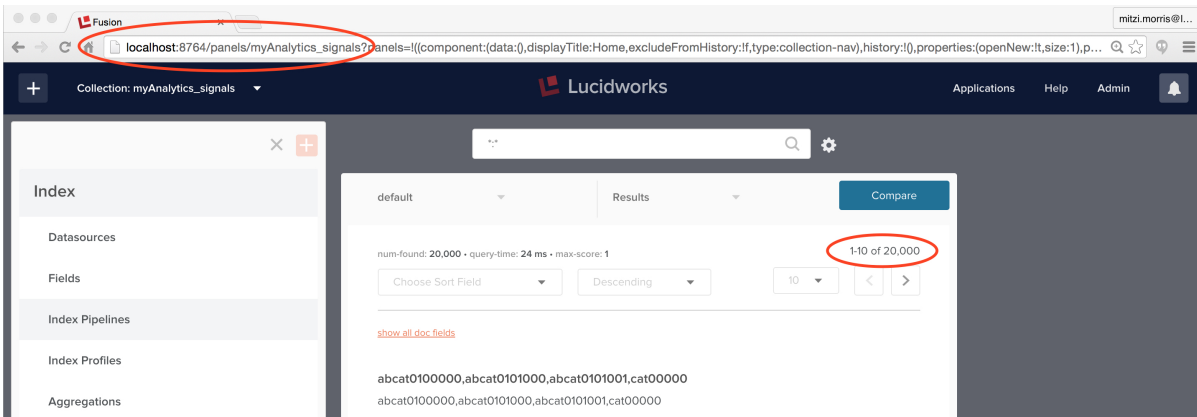
I create the primary collection "myAnalytics" from the Fusion UI. By default, this collection has feature "signals" enabled, therefore I can load the raw signal data into Fusion collection "myAnalytics_signals" by sending a POST request to the endpoint: `/api/apollo/signals/myAnalytics`

To run this command on a local Fusion installation using all default ports, the request is:

```
curl -u <user>:<password> -X POST\  
-H 'Content-type:application/json\  
http://localhost:8764/api/apollo/signals/myAnalytics?commit=true\  
--data-binary @<some-path>/signals.json
```

This command succeeds silently.

To verify that the collection "myAnalytics_signals" exists, view the collection details in the Fusion UI via URL: "localhost:8764/panels/myAnalytics_signals".




3.1.3. Create the Dashboard

Time-series dashboards show trends over time by using the timestamp field to aggregate query results. To create a time-series dashboard over the collection "myAnalytics_signals", from the Dashboards application, use the "New" control to open the "Create Dashboard" dialog and choose the option "new time-series dashboard"



The next step is configuring the dashboard settings:



The newly created dashboard is called "New Time Series Dashboard". It contains 4 rows: * Row 1 contains the input controls and a display panel showing total hits. * Row 2 displays the time range specified in the search query * Row 3 contains a visual display of hits per date-time interval * Row 4 contains a table over all documents in the result set.

In the example above, the dashboard is populated with results for the following Solr query parameters:

```
q=*:*&
fq=timestamp_tdt:[NOW/MINUTE-15MINUTE TO NOW/MINUTE+1MINUTE]&
wt=json&rows=0
```

The example signals data was collected during the months of August through October 2011. After the "Time Range" panel is configured with this date range, the new Solr query parameters are:


```
q=*:*&
fq=timestamp_tdt:[2011-08-01T04:00:00.000Z TO 2011-11-01T03:59:59.000Z]&
wt=json&rows=0
```

This query returns all 20,000 raw signals and the panel is updated accordingly.

Configure the Histogram Panel

Solr facets provide the counts displayed by the Histogram. The facet parameters are derived from the TimePicker selection and the controls on the Histogram display panel. These controls are used to select the type of count display and the time interval over which counts are aggregated.

After this interval is set to 1-day, the dashboard shows the daily activity for the timespan:



Clicking on the "info" icon on the Histogram controls shows the Solr query parameters:

```
q=*:*&
wt=json&rows=0&
fq=timestamp_tdt:[2011-08-01T04:00:00.000Z TO 2011-11-01T03:59:59.000Z]&
facet=true&
facet.range=timestamp_tdt&
facet.range.start=2011-08-01T04:00:00.000Z&
facet.range.end=2011-11-01T03:59:59.000Z&
facet.range.gap=+1DAY
```

Configure the Table Panel

The table panel displays the documents in the results set, 1 per row. The default display shows all document fields. The table "Fields" control is used to select the fields. The table header row column headers have arrows which allow shifting the column position left or right.

After selecting a subset of the document fields and ordering the display accordingly, the table panel shows the essential raw signal document fields:




TABLE PANEL

0 to 100 of 500 available for paging

Fields

- _text_
- _version_
- count_l
- doc_id_s
- filters_orig_ss
- filters_s
- flag_s
- id
- query_orig_s
- query_s
- query_t
- timestamp_tdt
- type_s
- tz_timestamp_txt


doc_id_s	query_orig_s	timestamp_tdt	filters_orig_ss
2009324	Sharp	2011-09-05T12:25:37.42Z	abcat0100000,abcat0101000,abcat0101001,cat00000
1517163	nook	2011-08-24T12:56:58.91Z	abcat0500000,cat00000,pcmcat193100050014,pcmcat223300050025
2877134	rca	2011-10-25T07:19:51.69Z	abcat0100000,abcat0101000,abcat0101005,cat00000
2416092	Flat screen tvs	2011-09-07T15:54:47.95Z	cat00000,pcmcat139900050002,pcmcat143200050016
2264036	Blue tooth headphones	2011-09-23T12:40:20.87Z	abcat0800000,abcat0811002,cat00000,pcmcat171900050028
2740208	memory card	2011-10-25T22:55:58.68Z	abcat0500000,abcat0504001,cat00000,pcmcat186100050006
2584273	AC power cord	2011-09-11T12:48:44.13Z	abcat0500000,abcat0515000,abcat0515012,cat00000,pcmcat138100050040
1230537	Zagg iPhone	2011-10-18T17:21:33.91Z	abcat0800000,abcat0811002,cat00000,pcmcat171900050031,pcmcat201900050009
3168067	Watch The Throne	2011-09-04T10:55:20.42Z	abcat0600000,cat00000,cat02001,cat02012,cat02713
7997055	Remote control extender	2011-10-28T16:26:29.22Z	abcat0107039,abcat0200000,abcat0208000,cat00000,pcmcat224000050003
1988047	3ds	2011-09-23T22:14:08.96Z	abcat0700000,abcat0707000,abcat0707001,cat00000
1260591	Joy stick	2011-09-19T21:17:48.55Z	abcat0515020,abcat0515022,abcat0700000,abcat0715000,cat00000
3425255	Zookeeper	2011-09-06T17:14:53.23Z	abcat0600000,cat00000,cat02015
9755322	Macbook	2011-09-26T15:12:48.53Z	abcat0500000,abcat0502000,cat00000,pcmcat247400050001
1432551	Ipad	2011-10-27T21:18:51.26Z	abcat0200000,abcat0204000,cat00000,pcmcat144700050004
9695793	Cable management	2011-10-05T10:10:35.55Z	abcat0500000,abcat0507000,cat00000,pcmcat230800050017
1092403	Driver ps3	2011-09-23T18:23:42.24Z	abcat0700000,abcat0703000,abcat0703002,cat00000
3271964	nunch	2011-10-07T08:11:41.56Z	abcat0200000,abcat0302005,cat00000,pcmcat223000050008

3.1.4. Save the Dashboard

The work done to create and configure this new time series dashboard is all work that has been done client-side. In order to save this dashboard for future sessions, it must be saved back to Fusion.

Before saving the dashboard, we rename it to "myAnalytics_signals dashboard" using the "Dashboard Settings" controls, which are accessed by clicking the gear icon on the top menu bar.

The top menu bar diskette icon opens the Save dialog. The option "Solr" saves the configured dashboard to Fusion.



The top menu bar folder icon opens the Load dialog. From this controls you can either reload or delete a saved dashboard configuration.

When the dashboard configuration is reloaded, the configured Solr query or queries will be rerun and the panel displays will be populated with the results.

3.1.5. Conclusion

In this section we have demonstrated how to create, configure, and save a new dashboard. The essential principles are:

- Dashboards are client-side application.
- Input panels create and submit queries to Fusion

- Dashboard displays are highly configurable.
- Configurations can be saved on the server.

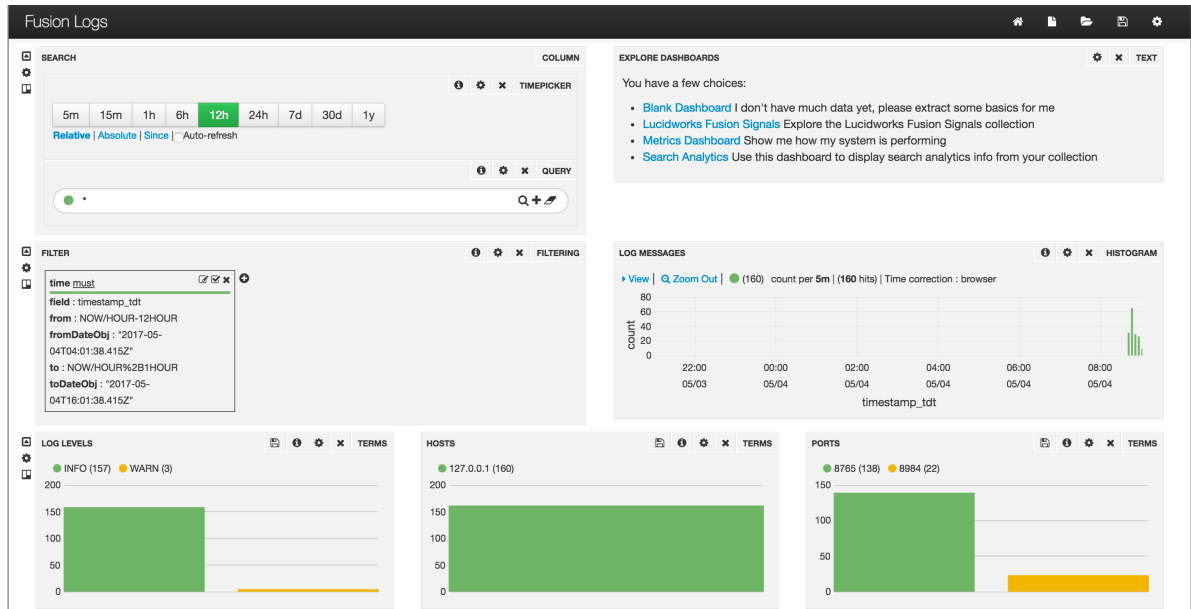
For more on this topic, see the Lucidworks blog:

<https://lucidworks.com/blog/2015/04/20/noob-notes-log-analytics-fusion/>

3.2. Log Analytics Dashboards

You can build a log analytics dashboard for any Fusion collection that contains one or more logfiles. Each logfile entry contains a timestamp plus some amount of additional event information. The time field must be a Solr trie-date field (with field suffix "_tdt").

This is an example of a log analytics dashboard. It is the default dashboard that appears when you click **Analytics** on the Fusion Launcher. Note the navigation to other dashboards in the upper right corner.



The more information and structure that can be extracted from the logfile and modeled in a Solr document, the more possibilities for analytics and visualizations. At a minimum log data must include:

- a timestamp in an allowed standard date/time format
- text message(s) which be unstructured or semi-structured

3.2.1. Search query controls

- query
- timepicker
- dateRange

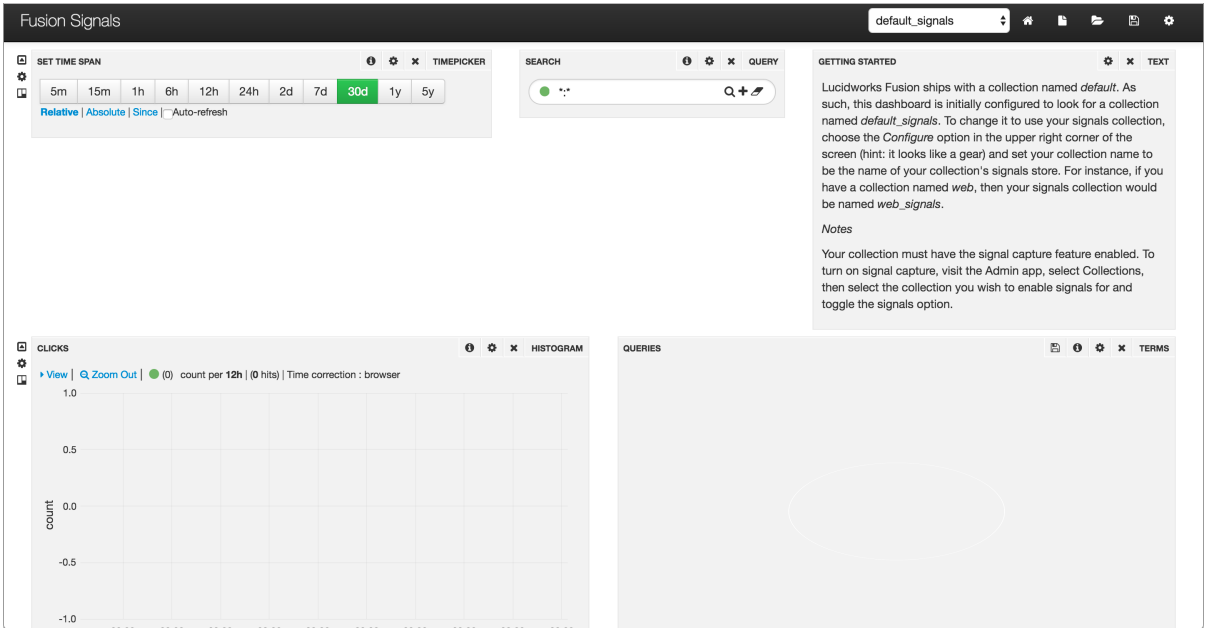
3.2.2. Display controls

At a minimum, a log analytics dashboard contains a histogram.

3.3. Signals Dashboards

Signals dashboards are a type of time-series dashboard that you can use to monitor signals collections, using the signal timestamp as the time field.

This is an example of a Fusion Signals dashboard:



The time field must be a Solr trie-date field (with field suffix "_tdt").

The default dashboard layout for a time-series dashboard is:

Row	Title	Type	Description
1	Time Window	Timepicker	input control which selects timestamp range, faceting granularity
1	Search	Query	keyword search terms
1	Total Hits	Hits	number of results for Solr query
2	-	Filtering	displays Solr params corresponding to Timepicker selections
3	Event Counts	Histogram	binned results as bar chart where X-axis is timeline and Y-axis is signals per date-time interval
4	Table Panel	Table	documents in results set, displayed 1 document per row. Table panel has controls over which fields to display, order in which fields are displayed

3.3.1. Signals Analytics Components

Search query controls

- query
- timepicker
- dateRange


Display controls

At a minimum, a signals analytics dashboard contains a histogram. If the signal contains discrete labels, additional display panels should be added to drill down on the signal contents.

3.4. Fusion Metrics Dashboard

The Fusion Metrics dashboard displays key Fusion metrics from the Fusion collection `system_metrics`.

This is the Fusion Metrics dashboard:



3.5. Search Analytics Dashboard

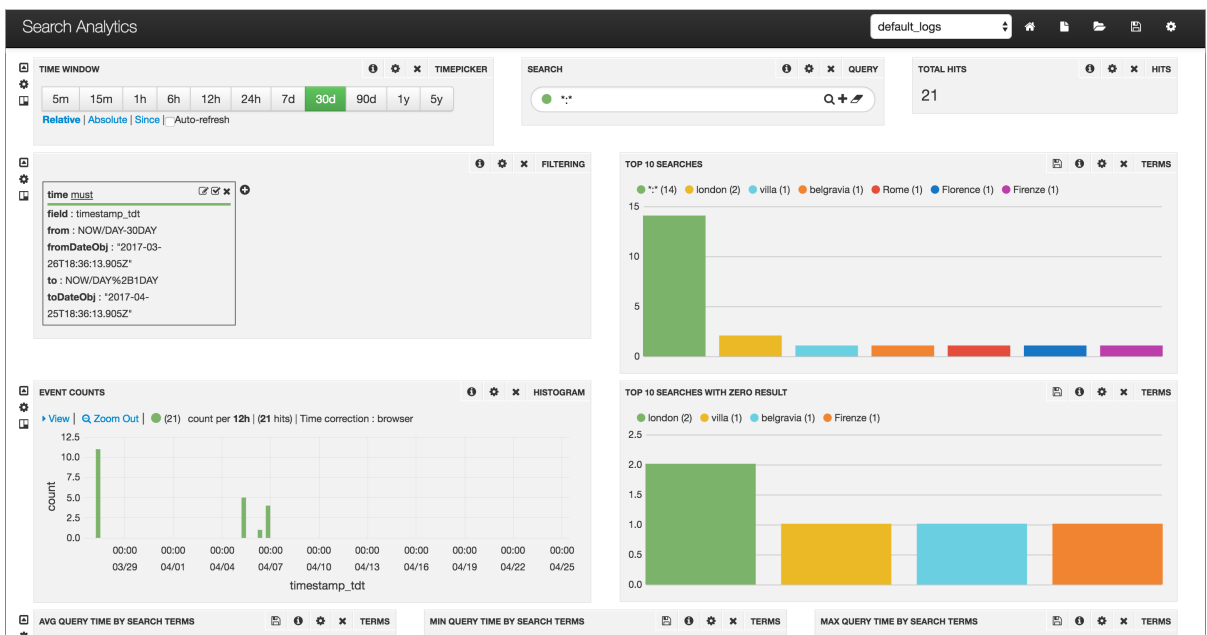
Use the built-in Search Analytics dashboard to view search analytics derived from collection logs.

Every Fusion collection has an associated logs collection named `<collection-name>_logs`. On the Search Analytics dashboard, select a logs collection (`collection_logs`) from the drop-down menu at the top of the page to view search analytics derived from that collection's logs.

You can open the Search Analytics dashboard from a link on the default dashboard (Fusion Logs). The Search Analytics dashboard `lucidworks-searchanalytics.json` is also available here:

```
/path/to//fusion/3.1.x/apps/jetty/ui/webapps/root/WEB-INF/classes/public/banana/app/dashboards`
```

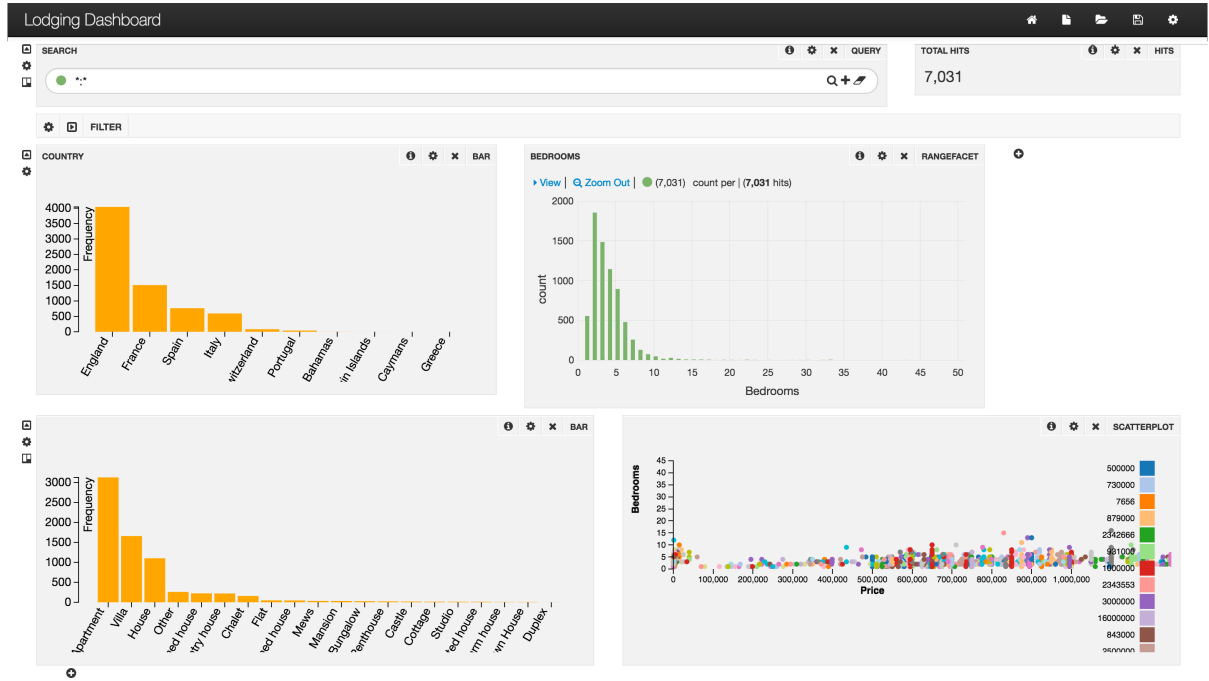
This is an example of a Search Analytics dashboard:



3.6. Non-Time Series Dashboards

Fusion provides the option of creating a non time-series dashboard over the contents of a Fusion collection. For non-time series data, the "filtering" widget controls the faceting.

This is an example of a non time-series dashboard for a collection that contains data about lodgings:



3.6.1. Search query controls

- query
- filtering

3.6.2. Display controls

If no faceting or quantitative information is present, a table can be used to drill-down on the documents in the Fusion collection.

3.7. Dashboard Layouts

Dashboard layouts are controlled by the Dashboard Settings menu. The dashboard setting menu is toggled open and close.

Note	Always remember to click "Save" to save your work.
------	--

3.7.1. Configuring Rows

A Dashboard consists of one or more rows, each of which contains one or more panels.

The Dashboard display contains a set of control icons on each row allowing you to hide/unhide, position, and delete that row. To add a row, enter a title for the new row, the height of the row, if it is editable, and then click **Create Row**. After the row is created, it will appear in the table of rows, and you can click the up or down arrows to arrange the new row with the existing rows.

3.7.2. Configuring Panels

To add a panel to a row, click the + icon to the right of the last panel in a row; if this icon does not appear, the row is full.

You can also click the 'gear' icon to the left of the first panel. This will bring up the row configuration popup, where the last tab is 'Add Panel'.

When adding a panel, the configuration screen varies depending on the required properties for each panel type. From the Add Panel tab, all of the available properties of each type are displayed, allowing you to define the panel name, the data properties, and any queries that should be used to limit the data used in the panel.

When editing a panel, however, the view is split between three tabs: General, for name and size configuration; Panel, for the data properties; and Queries, for defining queries.

3.7.3. Nested Panel Layout using Column Panels

A column panel is a container panel for other display panels. The properties of a column panel allow you to define the panels you would like included.

Choose the type of panel, and then define the properties according to the type chosen. When you have finished configuring the included panel, click Create Panel.



Below the Add Panel to Column area, the configured panels will be shown as a list, in the order they will appear on the dashboard. From this list, you can change the height of each panel, remove panels, change the display order, or temporarily hide panels.

3.8. Input Panels

Input panels let you control which data output panels display. The values you specify in input panels become parts of the Solr queries that output panels use to obtain data.

You can use these input panels:

- **Query Panel** – Enter a free-form query (one or more query terms) in a search bar. Add additional search bars to a query panel to submit separate queries. Some visualization panels (for example, Rangefacet) keep the data separate so you can compare it. This is an example of a Query panel with two search boxes. The parameters for query strings are `Primary_Specialty:Urology` and `Primary_Specialty:Psychiatry`. Here is the Query Panel and the resulting Rangefacet panel:



- **Time Picker Panel** – Apply a time range to time-series data. The time range can be:
 - Relative – A time range starting now and reaching backward in time, for example, the last 15 minutes or one hour

TIME WINDOW ⓘ ⚙️ ✕ TIMEPICKER

5m 15m 1h 6h 12h 24h 7d 30d 90d **1y** 5y

[Relative](#) | [Absolute](#) | [Since](#) Auto-refresh

- Absolute – A specific time-and-date range, for example, from 06/01/2017 00:00:00 to 06/30/17 23:59:59

TIME WINDOW ⓘ ⚙️ ✕ TIMEPICKER

06/29/2017 19:27:46 to 08/08/2017 14:16:21 ✓

[Relative](#) | [Absolute](#) | [Since](#)

- Since – A time range since a specific date and time, for example, since 01/01/2017 00:00:00.

- **Facet** – A Facet panel can facet any data field.

ADDITIONAL FILTERS ⓘ ⚙️ ✕ FACET

Filter


City ▼

Medical_School_Name ▼

Gender ▲


f (813306) ✕

- **Filtering** – A Filtering panel lets you apply field-based filters to the hits returned by the query. The filters apply to all display panels.



Important: You must use a Filtering panel somewhere on your dashboard, so that all panels work correctly when you interact with data.

3.8.1. Query Syntax

Enter a search term or phrase in the search box of the Query panel. Autocompletion provides a list of possibly related prior searches. Finish typing your search query or select an autocompleted query, and then click Search .

Enter queries in a Query panel using [Apache Lucene Query Parser syntax](#). You enter the parameter for the query string (for example, *Susan* or *Gender:F*), not the query string (for example, *q=Susan* or *q=Gender:F*).

Rules for the Simplest Cases

Here are some syntax rules for the simplest cases:

- A single term is either *field:value* or *value*. With *value*, the search is over all fields.
- For an exact-case match, you must specify the field name.
- Surround a term that contains spaces in *double quotation marks* (" ").
- You can't specify *:value* or *:"value"* to search over all fields; that syntax doesn't work.
- To retrieve all records, use the search term *:*.
- For OR logic, enter *OR* between the terms, or just use spaces. For AND logic, enter *AND* between the terms.

Examples

These are examples of the query syntax:

Goal	Syntax and example
Single term in any field; no blanks in term	<i>term</i> (matches any case)
Single term in any field; blanks in term	<i>"term"</i> (matches any case)
Multiple terms, each in any field; with OR logic; no blanks in terms	<i>term1 term2 ...</i>
Multiple terms, each in any field; with OR logic; blanks in terms	<i>"term1" "term2" ...</i> (matches any case)

Goal	Syntax and example
Multiple terms, each in any field; with AND logic (in the same record); no blanks in terms	<code>term1 AND term2 ...</code>
Multiple terms, each in any field; with AND logic (in the same record); blanks in terms	<code>"term1" AND "term2" ...</code> (matches any case)
Single term in a specific field; no blanks in term	<code>field:term</code> (matches any case)
Single term in a specific field; blanks in term	<code>"field:term"</code> (matches <i>exact</i> case)
Multiple terms, each in a specific field; with OR logic; no blanks in terms	<code>field1:term1 field2:term2 ...</code>
Multiple terms, each in a specific field; with OR logic; blanks in terms	<code>"field1:term1" "field2:term2" ...</code> (matches <i>exact</i> case)
Multiple terms, each in each in a specific field; with AND logic (in the same record); no blanks in terms	<code>field1:term1 AND field2:term2 ...</code>
Multiple terms, each in a specific field; with AND logic (in the same record); blanks in terms	<code>"field1:term1" AND "field2:term2" ...</code> (matches <i>exact</i> case)

For more information about the query syntax, see [Standard Query Parser Parameters](#).

Tip You can use a Text panel to advise the user regarding the syntax of search terms. Also, if you want the dashboard user to explore subsets of the data, use a Filtering panel or a Facet panel to achieve that. Don't expect users to enter complex search expressions. For example, add the field `gender` as a facet, instead of expecting the user to add `AND gender:male` to a search expression.

3.8.2. Inspect a Panel Query

You can't inspect the panel query in the Query panel. You can inspect the panel query in other panels, for example, in a Histogram or Heatmap panel. You can see the contributions that the different parts of the query make. In this example, there are no global query parameters.

Query:

```
start_station_name:"Broadway and E 14 St" AND gender:Male
```

Panel query for a Histogram panel:

The part of the query from the Query panel is the first part, from `q=` through `Male`.

```
q=start_station_name%3A%22Broadway%20and%20E%2014%20St%22%20AND%20gender%3AMale&wt=json&rows=0&fq=start_time:[2014-01-28T20:16:59.000Z%20T0%202014-03-15T05:36:55.000Z]&facet=true&facet.range=start_time&facet.range.start=2014-01-28T20:16:59.000Z&facet.range.end=2014-03-15T05:36:55.000Z&facet.range.gap=%2B12HOUR
```

3.8.3. Query Panel

The query panel provides a search box to allow real-time filtering of data.

It is a best practice to include one of this type of panel on your dashboard. With this panel, you can add, remove, label, pin and color queries.

There are no specific properties for a query panel.

3.8.4. Timepicker Panel

The timepicker panel controls the time range filters. This control is should be included on log and signal analytics dashboards, and any other dashboard based on time-series data.

The configuration properties are:

- **Default Mode:** The options are **relative**, which provides a series of relative timeframes (such as 30 days ago, 1 year, etc.); **absolute**, where you define the start and end dates; or **since**, where you define only the starting date, with the current date assumed.
- **Time Field:** The field to use for time-based data.
- **Relative Time Options:** When the mode is set to relative, you can provide a comma-separated list of relative time options, such as "5m,1h,2d". If you use the default range, you should set the panel to span at least 6, to prevent the time selections from overrunning the edges of the panel.
- **Default Timespan:** The time option that should be selected as a default.
- **Auto-refresh:** When the mode is set to either relative or since, you may want your dashboard to automatically refresh with the latest data. These options allow you to configure auto-refresh:
 - **Enable:** Select to enable auto-refresh.
 - **Interval:** The interval, in seconds, to refresh.
 - **Minimum Interval:** The minimum interval, in seconds. to refresh.

3.9. Display Panels

Display panels on a dashboard display information about the data in a single collection. Fusion has these types of display panels:

- **Layout** – Use these display panels to organize the panels on a dashboard.
- **Textual Information** – Use these panels to display textual information. For example, use a Table panel to displays the values of fields in a set of records.
- **Graphical Visualization** – Use these panels to help users to visualize data. For example, display category data in a bar chart and geographical data on a map.

3.9.1. What Data is Displayed

These things determine which data Fusion displays in a display panel:

- Where the query is sent (to a Fusion query pipeline or to Solr)
- The collection that contains the data. You can specify this when you create a dashboard, or permit the user to choose the collection.
- Global query parameters (optional)
- Input panel configuration settings (possibly including a panel query)
- A user's selection of the collection (if permitted)
- A user's query
- A user's interactions with the data

3.9.2. Layout Panels

These layout panels can help you organize a dashboard:

- **Column** – Lay out panels in a column within a row or part of a row. (Instead of the usual left-to-right layout within rows.)
- **Text** – Add text to a dashboard, for example, to instruct the user regarding how to use the dashboard, or to describe the data.

3.9.3. Textual Information Panels


These panels display textual information:

- **Doc Viewer** – Display documents as single pages of information, and let the user page through documents. Paging through documents doesn't affect the data displayed in other panels.
- **Full Text Search** – Provide full text search capability.

Note: Make a Full Text Search panel wide. Start with Span = 12. Depending on the content, you might be able to make it narrower. If you make a Full Text Search panel too narrow, content and controls can be inaccessible.

- **Hits** – A Hits panel displays statistical information about the hits returned by the query, including all filters that are applied to the query. The default information is the count (number of hits). You can display the count, minimum, mean, maximum, sum, standard deviation, the sum of squares, and/or the number of hits that lack a value for a

field.




- **Table** – Display data in a table.

Note: Make a Table panel wide. Start with Span = 12. Depending on the content, you might be able to make it narrower. If you make a Table panel too narrow, content and controls can be inaccessible.


3.9.4. Graphical Visualization Panels

Most display panels are graphical visualization panels. They visually present the data requested in the input panel(s), subject to configuration of the display panel and to any global filtering. Panels for visualizing information graphically are:


- **Bar** – Graph data as frequencies in a bar chart.




- **Better Map** – Display geolocated points in clustered groups on a map. As you zoom out, points are clustered into fewer groups. As you zoom in, points are clustered into fewer groups (and at some point are not clustered). The Better Map panel *doesn't* use the geospatial search capabilities of Solr. It transfers more data than the Map panel and generally requires more computation, while showing less data. If you have a time filter, the panel will show the most recent points in your search, up to your defined limit. This panel is best used after filtering the results through other queries and filter queries, or when you want to inspect a recent sample of points.




- **Heat Map** – Display a heat map, that is, a graphical representation of data along two facet axes.




- **Histogram** – Display a histogram. When used in `count` mode, this is a bucketed chart of the current query, including all applied time and non-time filters. When used in `values` mode, the histogram plots the value of a specific field over time, and lets the user group the values based on the values of a second field.




- **Map** – Display a map of shaded regions using a field that contains a 2-letter country code or U.S. state code. Regions with more hits are shaded darker. Instead of a count, you can choose to shade regions based on the minimum, maximum, mean, or sum. The map panel uses facets, so it is important that you set field values to the appropriate 2-letter codes at index time.




- **Multi-series** – Graph multiple fields of the same type that vary with time in the same graph. The X-axis is used for the time series field. Values in the other fields are graphed along the Y-axis.
- **Range Facet** – Display a histogram of a numeric field. The Range Facet panel is similar to the time series histogram. It lets you select ranges and zooming in/out to the desired numeric range. Range selections in the panel are reflected across the entire dashboard. With multiple search boxes in a Query panel, the Rangefacet panel can display multiple data sets, as shown in this example:



- **Scatter Plot** – Display a scatter plot between two variables or four variables.



- **Sunburst** – Display a sunburst plot based on facets.




- **Tag Cloud** – Display a tag cloud of the top N words in a field.



TERMS

START STATION



3.9.5. Bettermap Panel

The bettermap panel displays geographic points in clustered groups on a map.

This panel type does not use the terms facet and it does query sequentially. This means that it transfers more data and is generally heavier to compute, while showing less actual data

If you have a time filter, it will attempt to show to most recent points in your search, up to your defined limit.

3.9.6. Filtering Panel

A Filtering Panel shows the ranges applied to the query facets. It is a feedback/information panel that shows the limits applied to the data. A Filtering Panel should always be included on an analytics dashboard as an aid to interpreting/understanding the other visualizations on the dashboard.

3.9.7. Heatmap Panel

The heatmap panel provides a heat map for representing pivot facet counts.

The Panel tab defines what data is displayed in the panel and how it is displayed.

Row Field: The field from Solr that will provide data for rows.

Column Field: The field from Solr that will provide data for columns.

Rows Limit: The maximum number of rows to display.

Heatmap Color: The base color for the heatmap. The intensity of the color in any cell is proportional to the pivot facet count for that cell.

Transposed

3.9.8. Histogram Panel

The histogram panel provides a binned display of queries per time interval, using Solr's range facets for data.

Configuration Options

Mode

The value for the y-axis. The options are **count** or **values**.

When choosing values, you must select a field to use as the basis for the values. This field must have a numeric field type. You can also select a Group By Field, which cannot be a multi-valued field, and creates multiple charts.

Time_field

The value for the x-axis. This is the field to use for display of the time information for the histogram.

Chart Settings

There are several chart settings.

- **Bars:** Enable to show bars on the chart.
- **Lines:** Enable to show lines on the chart. When enabled a few more settings will be available:
 - **Line Fill:** The fill area, from 0-10.
 - **Line Width:** The width of the lines, in pixels.
 - **Smooth:** Enable to remove 0 values from lines.
- **Points:** Enable to show points on the chart.
- **Stack:** Enable to stack multiple series together.
- **Percent:** Enable to show the stack as a percentage of the total (only displayed when stack is enabled).
- **Legend:** Enable to show the legend.
- **xAxis:** Enable to display the x-axis values.
- **yAxis:** Enable to display the y-axis values.
- **Time correction:** If time correction should be applied to use the browser's timezone. Select 'utc' to always display times in UTC.
- **Selectable:**
- **Zoom Links:** Enable to allow users to zoom out in time.
- **View Options:** Enable to show an options section.
- **Auto-interval:** Enable to allow the chart to automatically scale the intervals.
- **Resolution:** When Auto-interval is enabled, a best effort will be made to show this number of bars.

Tooltip Settings

The tooltip settings control the display of data when users hover over a line or bar on the chart.

- **Stacked Values:** When using stacked values, this defines if the data be displayed as cumulative, or as individual values.

- **Display Query:** If an alias is set, it will be shown in the tooltip. If no alias is set, enable this to show the entire query.

3.9.9. Hits Panel

The hits panel shows the total hits for the current query input to a query panel.

The properties allow providing a title for the panel, the style of results, and the font size.

3.9.10. Map Panel

A map panel displays a map of shaded regions using any field that contains a 2-letter country or US state code. Regions with more hits are shaded darker.

This uses the Solr terms facet, so it is important that you set it to the correct field.

The following properties are defined for a map panel:

Field: the Solr field that has the 2-letter codes that will be used for the location data.

Max: a maximum number of countries to plot. The default is 100.

Map: the style of map to display. If your data spans the world, you can choose the world map. If your data is focused on Europe or the US instead, you can choose the Europe or US maps respectively.

Mode: the approach to summarizing the data. You can chose count, mean, maximum, minimum or sum. In order to use any mode other than count, the field type must be numeric.

Decimal Points: the number of digits to display after a decimal points.

3.9.11. Table Panel

The table panel allows you to create a table of field values from the Solr index that match the filters applied to the dashboard. While only the fields selected are displayed, clicking on any entry expands to show all fields of the document.

You can also export the documents, if needed, to CSV, XML or JSON by clicking the "Export" icon in the upper right of the panel.

In the display, you can add new fields on the fly, but clicking a field name from the list in the left side of the panel.

Configuration Options

Add Column

Enter the field(s) you would like to see in the table. You can enter as many fields as you'd like, but too many columns may require you to scroll horizontally to see all of your data.

Click the '+' button to add the field, and you should see it listed in the **Columns** section to the right. Click on a column to remove it from the list.

Options

The display options allow defining how the table is displayed.

- **Header:** Enable to see a header row in the table.
- **Sorting:** Enable to be able to sort the data in the table.
- **Sort:** Choose the field to sort on. Also choose the sort order by selecting the up or down carat to the right of the column name.
- **Font Size:** Choose the font size for the display of the data.
- **Trim Factor:** If the data is too long to fit in the column at your desired width, you can configure the point at which the data will be "trimmed" from display.

Paging

The paging options allow control over how to deal with large amounts of data that may be easier to work with on separate pages.

- **Show Controls:** Enable to show page forward and back options.
- **Overflow:** If the data expands beyond the height of the window, you can choose to **scroll** through the records, or enable **overflow** to expand the size of the panel to fit all of the data for that page.
- **Per Page:** The number of items to display on each page.
- **Page Limit:** The number of pages to display.
- **Pageable:** This will automatically update based on the values of Per Page and Page Limit, to show the total number of items that will be displayed in the table.

3.9.12. Terms Panel

The terms panel displays the results of a Solr facet as a pie chart, bar chart, or a table.

A statistics field can be displayed as min/max/mean/sum, faceted by the Solr facet field, also displayed either as a pie chart, bar chart or a table.

A terms panel takes several properties, described below.

- **Field:** The field to use as the basis of the facets that are used for display.
- **Length:** The maximum number of terms to display.
- **Order:** The sort order of the facets.
- **Style:** The style of chart, either **bar**, **pie** or **table**.
- **Legend:** If you choose bar or pie as the style, you can then choose no legend, or to display it above or below the data.
- **Font Size:** If you choose table, you will be given the option to define the font size.
- **Missing:** Enable to display missing values.
- **Other:**
 - **Donut:** If you choose pie chart as the style, you can choose to display the chart as a donut, with an empty circle in the middle.
 - **Tilt:** If you choose pie chart as the style, you can choose to tilt the chart as an added effect.
- **Labels:** Enable to show labels in the chart for your data.
- **Mode:** The mode for the data. Choose **count**, **mean**, **min**, **max**, or **sum**. If choosing any mode other than count, the Stats Field selected must be a numeric field.
- **Stats Field:** If you choose any mode other than count, you must then specify the field to use for statistics. This field must be a numeric field.
- **Display Precision:** Choose the number of digits to display after a decimal point, as appropriate.

3.9.13. Text Panel

A text panel displays static text, which can be either unformatted plain text or formatted using either [Markdown](#) or HTML.

To configure a text panel, first choose the text format mode, then enter the panel contents. For plain text displays, the font size can be specified.

3.9.14. Ticker Panel

The ticker panel provides a stock-ticker style representation of how queries are moving over time.

When configuring a ticker panel, there is one primary property, "Time Ago", which defines the point in time to use as the basis for comparison.

For example, if the time is 1:10pm, your time picker was set to "Last 10m", and the Time Ago parameter was set to '1d', the panel would show how much the current query results have changed since 1:00 to 1:10pm yesterday.

Chapter 4. Dev Ops

Performance monitoring tools and reports are covered in System Metrics.

Tools for creating and scheduling jobs are covered in Jobs and Schedules.

Fusion provides system messages and notifications via Messaging and Alerting.

4.1. System Metrics

By default, collection of system metrics is disabled. When it is enabled, Fusion continuously indexes system and Solr metrics to the system collection `system_metrics`. Collection of system metrics is enabled using the Configurations API, like this:

```
curl -u user:pass -H 'Content-type:application/json' -X PUT -d 'true'
"http://localhost:8764/api/apollo/configurations/com.lucidworks.apollo.metrics.enabled"
```

There are around 600 different metrics available. In this topic we've highlighted a few that are likely to be the most useful or interesting to you.

The `/system/metrics` endpoint of the System API lists all the metrics that the system is currently collecting. Metrics are returned for the current instance only; Fusion instances do not aggregate metrics between nodes.

4.1.1. Types of Metrics Collected

There are several types of metrics:

- **Gauges:** These are single values, valid for the point in time at which the metrics are collected.
- **Counters:** These are values that are incremented or decremented over time.
- **Meters:** These measure the rate of events over time. They include a mean rate, as well as a 1-, 5- and 15-minute moving average. Most of these moving averages are exponentially weighted, so that more recent values contribute more heavily than older values; exceptions to this rule have the word "unweighted" in their name.
- **Histograms:** These measure the distribution of values. They will report the minimum, maximum, mean, and the values at the 50th, 75th, 95th, 98th, 99th, and 99.9th percentiles.
- **Timers:** A timer is a meter combined with a histogram; it measures the length of time that a particular operation takes (both mean duration and moving averages) as well as the distribution of those durations.

Many of the metrics are for internal use by the system. However, Fusion may ask for a dump of the metrics data (using the System API endpoint) to help diagnose performance issues. Some metrics are also subject to change pending performance tuning and additional testing.

4.1.2. Metrics of Particular Interest

Slow Web Service Calls

For each web service endpoint in the system, the system keeps a list of the last several requests whose request time has been in the 99th percentile – that is, examples of the top 1% of slow requests for that endpoint. These are recorded as `com.lucidworks.apollo.resources.serviceName.methodName.weighted.slow.examples`, where `serviceName` is the name of the service and `methodName` is the name of a valid method for that service.

This information might be helpful when diagnosing performance issues. Here is an example of the 5 slowest calls to the `getCollectionMetrics` method of the `CollectionResource` service:

```

"com.lucidworks.apollo.resources.CollectionResource.getCollectionMetrics.weighted.slow.examples" : {
  "value" : [ {
    "requestUri" : "http://localhost:8765/api/v1/collections/lws5_metrics/stats",
    "queryParams" : { },
    "userPrincipal" : null,
    "method" : "GET",
    "cookies" : { }
  }, {
    "requestUri" : "http://localhost:8765/api/v1/collections/logs/stats",
    "queryParams" : { },
    "userPrincipal" : null,
    "method" : "GET",
    "cookies" : { }
  }, {
    "requestUri" : "http://localhost:8765/api/v1/collections/logs/stats",
    "queryParams" : { },
    "userPrincipal" : null,
    "method" : "GET",
    "cookies" : { }
  }, {
    "requestUri" : "http://localhost:8765/api/v1/collections/lws5_metrics/stats",
    "queryParams" : { },
    "userPrincipal" : null,
    "method" : "GET",
    "cookies" : { }
  }, {
    "requestUri" : "http://localhost:8765/api/v1/collections/lws5_metrics/stats",
    "queryParams" : { },
    "userPrincipal" : null,
    "method" : "GET",
    "cookies" : { }
  } ]
}

```

System Memory

There are several memory-related metrics reported:

- **mem.heap.used**: the current amount of heap memory, in bytes, used by the system.
- **mem.heap.max**: the maximum amount of heap memory, in bytes, that the system could use.
- **mem.heap.usage**: the percentage (0 - 1.0) of available heap memory that the system is currently using (this is equal to $\text{mem.heap.used} / \text{mem.heap.max}$).
- **mem.non-heap.used**: the current amount of non-heap memory (also called "off-heap memory"), in bytes, used by the system.
- **mem.non-heap.max**: the maximum amount of non-heap memory, in bytes, that the system could use.
- **mem.non-heap.usage**: the percentage (0 - 1.0) of available non-heap memory that the system is currently using (this is equal to $\text{mem.non-heap.used} / \text{mem.non-heap.max}$).
- **mem.total.used**: the current total amount of memory (heap plus non-heap), in bytes, used by the system.
- **mem.total.max**: the maximum amount of total memory (heap plus non-heap), in bytes, that the system could use.

Here is an example of **mem.heap.used**:

```

{
  "version" : "3.0.0",
  "gauges" : {
    "mem.heap.used" : {
      "value" : 94783360
    }
  },
  "counters" : { },
  "histograms" : { },
  "meters" : { },
  "timers" : { }
}

```

Query and Index Pipeline Stage Metrics

For each query pipeline and index pipeline stage, Fusion collects aggregate performance metrics for successful executions and for errors. All executions for each stage are stored in a metric named `stages.stageType.stageName.process`, where `stageType` is the type of stage, and `stageName` is the name of a specific stage.

Here is an example of a request to get the performance metrics for an index pipeline stage named 'solr-default' (`stages.solr-index.solr-default.process`), which is included with Fusion:

```

{"version" : "3.0.0",
 "gauges" : { },
 "counters" : { },
 "histograms" : { },
 "meters" : { },
 "timers" : {
  "stages.solr-index.solr-default.process" : {
    "count" : 109195,
    "max" : 0.128585,
    "mean" : 0.004011065175097276,
    "min" : 0.0022500000000000003,
    "p50" : 0.0030645000000000004,
    "p75" : 0.0033495,
    "p95" : 0.005410449999999992,
    "p98" : 0.014195759999999965,
    "p99" : 0.02462230000000001,
    "p999" : 0.12850243700000002,
    "stddev" : 0.007408363728123277,
    "m15_rate" : 11.957732876922531,
    "m1_rate" : 8.784289947811962,
    "m5_rate" : 9.037172472578138,
    "mean_rate" : 9.214233776748047,
    "duration_units" : "seconds",
    "rate_units" : "calls/second"
  }
}
}

```

This shows the number of uses of the stage ("count"), the maximum and minimum times, the mean, the 50th, 75th, 95th, 98th, 99th, and 99.9th percentiles (p50, p75, etc.), and the mean rates over 1-, 5- and 15-minute intervals ('m1_rate', etc.). In this case, the pipeline has been used 109,195 times, with a mean rate of 9.214 events per second, with only .003 events in the 50th percentile.

Metrics for successful completions of stages are stored in metrics named `stages.index.stageType.stage.stageName.ok` or `stages.query.stageType.stage.stageName.ok`, depending on if the stage is part of an index pipeline or a query pipeline. Here is an example of the mean rates for successful runs of the 'solr-default' index pipeline stage (`stages.index.solr-index.stage.solr-default.ok`):

```
{
  "version" : "3.0.0",
  "gauges" : { },
  "counters" : { },
  "histograms" : { },
  "meters" : {
    "stages.index.solr-index.stage.solr-default.ok" : {
      "count" : 110855,
      "m15_rate" : 5.270163206842968,
      "m1_rate" : 8.485969925086419,
      "m5_rate" : 8.06785229981572,
      "mean_rate" : 9.18230056255745,
      "units" : "events/second"
    }
  },
  "timers" : { }
}
```

This shows the number of uses of the stage ("count") and the mean rates over 1-, 5- and 15-minute intervals ('m1_rate', etc.). From the above, we can see that the solr-default stage has been executed 110,855 times, with a mean rate of 9.18 events per second.

If you prefer to see the metrics for the entire stage type, you can omit the stage name entirely, and simply get metrics for the stage type. This takes the form of `stages.index.stageType.ok` (for an index pipeline) or `stages.query.stageName.ok` (for a query pipeline). Here is an example, using the solr-index stage type:

```
{
  "version" : "3.0.0",
  "gauges" : { },
  "counters" : { },
  "histograms" : { },
  "meters" : {
    "stages.index.solr-index.ok" : {
      "count" : 116425,
      "m15_rate" : 6.178851947720613,
      "m1_rate" : 8.814380052133192,
      "m5_rate" : 8.585203640734829,
      "mean_rate" : 9.19499774409566,
      "units" : "events/second"
    }
  },
  "timers" : { }
}
```

In this example, we see that the solr-index stage has been successfully run 116,425 times, with a mean rate of 9.19 events per second.

Web Service Endpoint Metrics

For each web service endpoint, Fusion keeps a timer recording the duration and rate of requests. The duration is calculated using an exponentially-weighted moving average with a heavy bias toward measurements from the last 5 minutes.

These metrics have names in the form: `com.lucidworks.apollo.resources.serviceName.methodName.weighted.timer`, or for a specific example, `com.lucidworks.apollo.resources.CollectionResource.getCollectionMetrics.weighted.timer`:

```
"com.lucidworks.apollo.resources.CollectionResource.getCollectionMetrics.weighted.timer" : {
  "count" : 2624,
  "max" : 0.134712,
  "mean" : 0.031589107976653694,
  "min" : 0.022424000000000003,
  "p50" : 0.028440000000000003,
  "p75" : 0.036908,
  "p95" : 0.044644449999999995,
  "p98" : 0.05026944,
  "p99" : 0.054440510000000004,
  "p999" : 0.134693411,
  "stddev" : 0.00936497282768644,
  "m15_rate" : 0.07113433590025664,
  "m1_rate" : 0.06387037028343223,
  "m5_rate" : 0.06218407166715861,
  "mean_rate" : 0.0663172057583814,
  "duration_units" : "seconds",
  "rate_units" : "calls/second"
}
```

Solr Request Metrics

The system keeps track of the performance of requests to each Solr server that it communicates with.

The metrics have names in the form `solr.solrIdentifier.requestType`. The `solrIdentifier` is the address of the Solr instance, and the `requestType` can be 'get-requests', 'post-requests' or 'put-requests'.

This example shows get-requests to a Solr instance that is found on '10.0.1.8' and port 8983:

```

{
  "version" : "3.0.0",
  "gauges" : { },
  "counters" : { },
  "histograms" : { },
  "meters" : { },
  "timers" : {
    "solr.10.0.1.8-8983.get-requests" : {
      "count" : 3170,
      "max" : 0.873981,
      "mean" : 0.2451200904669261,
      "min" : 0.001678,
      "p50" : 0.318176,
      "p75" : 0.48169550000000005,
      "p95" : 0.53017705,
      "p98" : 0.5617982399999999,
      "p99" : 0.62812218000000003,
      "p999" : 0.8710894970000004,
      "stddev" : 0.2448979377578966,
      "m15_rate" : 0.02059326561557774,
      "m1_rate" : 0.03249432457272969,
      "m5_rate" : 0.030788223074952624,
      "mean_rate" : 0.033875616252208286,
      "duration_units" : "seconds",
      "rate_units" : "calls/second"
    }
  }
}

```

From this we can see that there have been 3,170 GET requests to that Solr instance, and the mean response rate is .03 requests per second.

4.1.3. Changing Metric Collection Frequency

The default frequency to collect metrics is 60 seconds. Since the metrics are stored in a system collection (and a Solr instance), the data can grow to be quite large over time. If you do not need metrics collection to happen as frequently (perhaps during initial implementation), you can change the frequency by modifying the `com.lucidworks.apollo.metrics.poll.seconds` configuration parameter with the Configurations API.

For example:

```

curl -u user:pass -X PUT -H 'Content-type: application/json' -d '600'
http://localhost:8764/api/apollo/configurations/com.lucidworks.apollo.metrics.poll.seconds

```

To disable metrics, you could set the `com.lucidworks.apollo.metrics.poll.seconds` parameter to '-1'.

```

curl -u user:pass -X PUT -H 'Content-type: application/json' -d '-1'
http://localhost:8764/api/apollo/configurations/com.lucidworks.apollo.metrics.poll.seconds

```

4.2. Messaging and Alerting

Fusion's messaging services provide implementations to send messages and alerts to any application or device capable of displaying the supported message types. Read a primer on Fusion's messaging services on our [blog](#).

4.2.1. Supported Message Types

Fusion supports these types of messages and alerts:

- logging

This service logs any message sent to it in the configured logger. You can use it in an index pipeline or a query pipeline.

- Slack

[Slack](#) is a team messaging service with document integration and a focus on collaborative communication. See the Slack Index Stage and the Slack Query Stage.

- SMTP

Email, via the Simple Mail Transfer Protocol. See the Email Index Stage and the Email Query Stage.

- PagerDuty

Alerting and monitoring. See the PagerDuty Index Stage and the PagerDuty Query Stage.

4.2.2. Messaging Service Configuration

The Message Services as a whole can be configured via the Configurations API with these attributes:

Attribute	Description
<code>rateLimit</code>	The time, in milliseconds, to wait between sending messages on a per-second basis. This does not synchronize throttling between requests.
<code>storeAllMessages</code>	Boolean flag that indicates whether messages should be indexed and stored. By default, only scheduled messages are stored, as they need to be retrieved by the scheduler at a later time. Storing all messages can be useful for auditing the system, but it will have an impact on the system storage requirements.

Enabling Messaging Services

The logging service is enabled by default, but Slack, email, and PagerDuty messaging services must be explicitly enabled.

You can do this through the UI at **Applications > System > Messaging Services**, or through the Messaging API.

To see which messaging services are currently enabled:

```
curl -u <user>:<pass> http://localhost:8764/api/apollo/messaging/service
```


String Templates

String templates are libraries used for structured text generation outputs. They are a powerful way of doing variable substitution into a provided template using values contained in documents, requests, and contexts. String templates are made available in the Messaging Service System setup, where users can fill in these portions with document or query values from the working system.

See Messaging Services Templates for details.

4.2.3. Triggering Messages and Alerts

The messaging services can be invoked in several ways:

- Via the Scheduler API to send messages at designated intervals.
- Via the Messaging API.
- Through an index or query pipeline; see Messaging and Alerting Pipeline Stages below.

Note	By default, only scheduled messages are stored. To configure Fusion to store all messages, see Messaging Service Configuration above. The default collection for message storage is system_messages, which is created on startup.
------	---

Messaging and Alerting Pipeline Stages

The Messaging Service supports these pipeline stages:

- Email Message Query Stage
- Email Message Index Stage
- PagerDuty Message Query Stage
- PagerDuty Message Index Stage
- Slack Message Query Stage
- Slack Message Index Stage
- SetPropertyIndex Stage

The pipeline stages above send messages and alerts when specific conditions are met. Conditions can be specified using regular expressions, database lookups, and more. Any upstream stage can affect how Fusion behaves when a match occurs, so pay special attention to the order in which stages occur.

Additionally, the Set Property Index Stage allows conditions to be specified before messages are sent.

Setting Properties Upstream

Fusion includes two index pipeline stages that are useful for setting properties on indexed documents so that you can evaluate those properties in one of the messaging stages, either downstream in the index pipeline or in a query pipeline:

- The Set Property Index Stage can be used to set a property on a document, or a context, by evaluating one or more simple conditions. It is an index-only, conditional stage that allows the setting of properties without the use of JavaScript.

- The JavaScript Index Stage provides a more sophisticated means of setting properties upstream from the messaging stages.

4.3. System Usage Monitor

The System Usage Monitor is a voluntary program to allow users to anonymously send basic information about their system to Lucidworks. We use this information to analyze the types of systems in use by our customers and how they are used so we can improve our product.

At no point does the system collect information that could identify you, your organization or the specific documents indexed. Only minimal data is sent about the type of content indexed. Our website has more information about our [privacy policy](#).

4.3.1. Information Collected by the Usage Monitor

The System Usage Monitor collects the following information:

- uuid - a randomly generated identifier per Fusion cluster.
- System information: the operating system, version and java version will also be reported.
- nodes - the number of Apollo nodes in use. This is calculated from Jetty processes on the same node or on a different node.
- Solr statistics:
 - Number of Solr nodes, total number of collections and number of Fusion collections.
 - Number of documents and the number that are regular documents and the number that are signals.
 - Number of total search requests and the number that are document search requests and the number that are signal search requests.
 - Total time to execute all search requests, to execute only document search requests, and to execute signal search requests.
- Aggregations:
 - Number of aggregation runs and how long they took (in ms).
 - Number of signals processed.
 - Number of aggregated signals.
- Recommendations:
 - Number of recommendation requests for each type of recommendation.
 - Total time to execute each type of recommendation request.

You can see the data that will be sent to Fusion with the Usage API and also in the UI by going to the Systems tab, then 'Heartbeat data'. The UI and the REST API will only report the data currently scheduled to be sent so is not a complete picture of all data collected.

4.3.2. How Data is Sent

Data is sent to Fusion once per week, and also whenever the system is restarted.

When Fusion is started, the System Usage Monitor will transmit data about your system to a server hosted by Lucidworks with two HTTP requests. The first request contains system-level information and if that is successful, the second request will send system-specific information.

The information is sent via an encrypted POST request to <https://heartbeat.lucidworks.io>. Each request includes a unique identifier, which is anonymous and can't be used to identify the sender. The IP that sent the request is not stored with the request.

4.3.3. How to Opt-Out

By default, the usage monitor is enabled in your system. If you would like to opt-out of sending this data to Lucidworks, you can disable the usage monitor. There are two ways to enable or disable the usage monitor:

1. Go to the Heartbeat Data page in the Fusion UI (System → Heartbeat Data), and deselect the "Report Heartbeat" option.
2. Use the Configurations API and send a PUT request as follows:

```
curl -u user:pass -H 'Content-type: application/json' -X PUT -d '"false"'
http://localhost:8764/api/apollo/configurations/usageMonitor
```

4.4. Schedules

Schedules in Fusion allow you to execute any Fusion service, Solr request, or other HTTP request on a defined timetable.

For example, you could schedule a Solr query to run at a specified time every day, or you could define a datasource to be re-crawled once a week. The schedules service does not execute any business logic; the service at the specified endpoint must provide this.

The Fusion scheduler is fault-tolerant and distributed across the nodes of your cluster. Several instances of the scheduler service can run on different nodes, but only one of them at a time executes and modifies schedules. This instance is elected as the schedule "leader", which occurs in ZooKeeper in a similar way to how SolrCloud node leaders are elected. The instances that are not the leader are on standby in case the leader goes down. The schedule job definitions are also kept in ZooKeeper, which allows them to be restored to any node whenever needed.

4.4.1. Scheduler job definitions

When defining a job with the scheduler service, there are two main aspects to configuration:

- The time properties that define when the job will run and how often it will repeat
- The call properties defining the call that will be executed

Time properties

These properties define the start time, end time, and repeat interval, if any.

- `startTime` defines when the job should first run.
- `endTime` defines when the job should no longer run.

The `endTime` does not stop a running job; instead it has the same effect as setting the entire schedule to "inactive" at a certain date.

- `interval` is an integer.
 - `millisecond` or `ms`
 - `second` or `sec`
 - `minute` or `min`
 - `hour` or `hr`
 - `day`
 - `week`
 - `month`

The interval can be "0", in which case the scheduled job only runs once. When the interval is higher than "0", then `repeatUnit` must also be defined.

- `repeatUnit` defines the unit of time to use in conjunction with the `interval`. The allowed values are:

These values are case-insensitive, meaning they can be entered in upper or lower case as you prefer.

Call properties

The call properties are where the actual task of the schedule is defined.

- `uri`

This can take several forms:

- An HTTP or HTTPS request: `<protocol>://<path>`
- A Solr request: `solr://<collection>/...`

For example, you could periodically issue a commit request to Solr. Or you could periodically run a query against a specific collection

- A Fusion service request: `service://<serviceName>/<path>`

The services available are stored in ZooKeeper. You can find them in the Admin UI under the "System" tab, or with a REST API call to the `/introspect` endpoint

- `method` is the HTTP method to use.
- `header` contains any additional required headers.
- `queryParams` contains any additional query parameters.

For Solr requests, `queryParams` may be any valid query parameter for the specified URI.

- `entity` is the request body, if any.

4.4.2. How to define a scheduler job

There are two ways to define a scheduler job:

- In the Fusion UI, at **DevOps > Scheduler**.

See below for instructions.

- Using the Scheduler API.

API examples are provided below.

Defining a scheduler job in the Fusion UI

1. Navigate to **DevOps > Scheduler**.
2. Click **Add a Schedule**.
3. Enter the parameters for the scheduler job:
 - **Schedule Name** – Any arbitrary string (required)
 - **Service** – The endpoint and method for the service to run (required)

Select the protocol:

- `http://` or `https://`

- `solr://{collection}/...`

A SolrCloud request.

- `service://{serviceName}/{path}`

A load-balanced Fusion service request.

- **Start Time** – The date and time at which to begin running the first instance of this job
- **End Time** – The date and time after which this job will be disabled
- **Run Once** – To run the job at regular intervals, uncheck this option.
- **Interval** – The interval at which to repeat this job
- **Active?** – To disable the job, uncheck this option.

Now your configuration should look something like this:

The screenshot shows a configuration form for a scheduler job. It includes fields for Service, path, method, Start Time, End Time, Run Once checkbox, Interval, and Active? checkbox. There are also Cancel and Save buttons.

In this example, the scheduler job crawls MyDataSource every hour. If you want this to happen every hour *on the hour*, you can set the start time to 11:00:00 (or any other hour). In the case of a crawl job like this one, you can check the job history by navigating to **Applications > Collections > CollectionName > Datasources > DatasourceName > Job History**.

4. Click **Save**.

API Examples

Each of these examples shows setting a schedule for an action in the system, using the Scheduler API. To see the results of a job, you will likely need to query the History API.

Issue a commit every 10 seconds

```
{"creatorType":"human", "creatorId":"me", "repeatUnit":"SECOND", "interval":10, "active":true, "callParams":{"uri":"solr://myCollection/update", "method":"GET", "queryParams":{"stream.body":"<commit/>"}}
```

In this example, we've defined the `callParams` with a URI for Solr that calls a collection named 'myCollection' and the 'update' `updateHandler`. The method is GET. The queryParams define the commit call for Solr. For timing, we've defined the job to run every 10 seconds.

Run a datasource every 20 minutes

```
{"creatorType":"human", "creatorId":"me", "repeatUnit":"MINUTE", "interval":20, "active":true, "callParams":{"uri":"service://connectors/jobs/TwitterSearch", "method":"POST"}}
```

In this example, we've defined the callParams with a URI for Fusion that calls the TwitterSearch datasource job. The method is POST, which is the method to use when starting a crawl. There aren't any other properties needed to define the task. For timing, we've defined the job to run every 20 minutes.

Remove signals older than 1 month

```
{"creatorType":"human", "creatorId":"me", "repeatUnit":"MONTH", "interval":1, "active":true, "callParams":{"uri":"solr://myCollection_signals/update", "queryParams":"stream.body=<delete><query>timestamp_dt:[* TO NOW-1MONTH]</query></delete>", "method":"GET"}}
```

In this example, we're again calling Solr's 'update' updateHandler with a collection named 'myCollection_signals', which is the default location for signals. This time we've also defined queryParams to delete documents that match a date query that finds all documents older than 1 month old. For timing, we've set this to run once a month.

4.5. bin/fusion

For every server in a Fusion deployment, the script `fusion/3.1.x/bin/fusion` is used to start, stop, and check the status of the Fusion instance running on that server.

4.5.1. Fusion Agent Process

The Fusion agent process makes sure that all Fusion processes start up and shut down correctly. It prevents problems that can arise by trying to start Fusion on a server where it is already running.

4.5.2. Start Fusion

Run the script `fusion/3.1.x/bin/fusion` with the argument `start`:

```
$ cd /path/to/fusion/3.1.x
$ ./bin/fusion start
Starting zookeeper.
Successfully started zookeeper on port 9983 (process ID 77295)
Starting solr.
Successfully started solr on port 8983 (process ID 77297)
Starting api.....
Successfully started api on port 8765 (process ID 77301)
Starting connectors.....
Successfully started connectors on port 8984 (process ID 77388)
Starting ui....
Successfully started ui on port 8764 (process ID 77469)
```

4.5.3. Check the status of Fusion

Run the script `fusion/3.1.x/bin/fusion` with the argument `status`:

```
$ cd /path/to/fusion/3.1.x
$ ./bin/fusion status
zookeeper is running on port 9983 (process ID 77295)
solr is running on port 8983 (process ID 77297)
api is running on port 8765 (process ID 77301)
ui is running on port 8764 (process ID 77469)
connectors is running on port 8984 (process ID 77388)
```

4.5.4. Stop Fusion

Run the script `fusion/3.1.x/bin/fusion` with the argument `stop`:

```
$ cd /path/to/fusion/3.1.x
$ ./bin/fusion stop
Successfully stopped ui (process ID 41524)
Successfully stopped connectors (process ID 41328)
Successfully stopped api (process ID 41159)
Successfully stopped solr (process ID 41153)
Successfully stopped zookeeper (process ID 41151)
```

4.5.5. Troubleshooting

The [Java Virtual Machine Process Status Tool](#) utility at `/usr/bin/jps` is useful for reporting on all Fusion processes reported by script `fusion/3.1.x/bin/fusion`:

```
$ jps
77294 AgentMain
77295 zookeeper-path-1475182112123.jar
77297 start.jar
77301 start.jar
77388 start.jar
77469 start.jar
79455 Jps
```

The process `zookeeper-path-1475182112123.jar` is the ZooKeeper process used by Fusion. The 4 `start.jar` processes are Fusion's Solr, API Services, Connectors, and UI.

If the `path/to/fusion/3.1.x/bin/fusion` script doesn't run, or if it fails to start all services, see the Troubleshooting topic or the [knowledge base](#) for help.

4.5.6. Directories and Logs

Directories

The directory where the Fusion files go for a specific version of Fusion is the *Fusion home directory*. The Fusion home directory is a version-numbered directory (for example, `3.1.0`) below the directory `fusion`. This installation strategy lets you install multiple versions of Fusion and switch between them. The Fusion home directory is the directory `fusion`.

The directories found in the Fusion home directory `fusion/3.1.x` are:

Name	Description
<code>apps</code>	Fusion components 3rd-party distributions used by Fusion, including jar files and plugins
<code>bin</code>	Master script to run Fusion, and per-component run scripts
<code>conf</code>	Configuration files for Fusion and ZooKeeper that contain parameters settings tuned for common use cases
<code>data</code>	Default location of data stores used by Fusion apps
<code>docs</code>	License information
<code>examples</code>	Fusion signals example
<code>init</code>	<code>systemd</code> and <code>upstart</code> scripts and configurations for Linux
<code>scripts</code>	Developer utilities, including diagnostic scripts, for Linux and Windows. See <code>scripts/diag/linux/README</code> and <code>scripts/diag/win64/README.txt</code> for details.
<code>var</code>	Logfiles and system files created by Fusion components, as well as <code>.pid</code> files for each running process

Symbolic Links on UNIX

To simplify access to the latest version of Fusion and to files in the `bin`, `conf`, and `var` directories, Fusion creates a symbolic link `latest` to the latest version and symbolic links `bin`, `conf`, and `var` to `latest/bin`, `latest/conf`, and `latest/var` respectively.

For example, if `latest` is `3.1.0`, then instead of entering this command to change to the `bin` directory:

```
$ cd /path/to/fusion/3.1.0/bin
```

You could just type:

```
$ cd /path/to/fusion/bin
```

To avoid possible confusion in the documentation, we spell out the path below the `fusion` directory.

From the `fusion` directory, you can view the symbolic links by typing:

```
$ ls -l . | grep "\->"
```

To change the version of Fusion to which the symbolic links refer, unlink and relink `latest`, for example:

```
$ cd /path/to/fusion
$ unlink latest
$ ln -s 3.1.2 latest
```

Logfiles

Logfiles are found in directories under `fusion/3.1.x/var/log`. Because the Fusion components run in separate JVMs, each component has its own set of logfiles and files that monitor all garbage-collection events for that process.

Name	Description
<code>api</code>	Fusion REST API services logging and error messages. This log shows the result of service requests submitted to the REST API directly via HTTP and indirectly via the Fusion UI.
<code>connectors</code>	Fusion connector services logging and error messages. Fusion index pipeline logging stages write to this file.
<code>solr</code>	Messages from Solr
<code>spark-master</code>	Spark-master logs
<code>spark-worker</code>	Spark-worker logs
<code>ui</code>	Fusion UI messages
<code>zookeeper</code>	ZooKeeper messages

Every component logs all messages to a logfile named `<component>.log`. For example, the full path to the logfile for the connectors services is:

```
fusion/3.1.x/var/log/connectors/connectors.log
```

In addition to component logfiles, every component maintains a set of garbage-collection logfiles which are used for resource tuning. The garbage-collection logfiles are named `gc_<YYYYMMDD>_<PID>.log.<CT>`. In addition, the current garbage-collection logfile has suffix `.current`.

The Fusion REST API, UI, connectors services, and Solr all run inside a Jetty server. The Jetty server logs are also written to that component's logfile directory. The Jetty server logs are named:

- `jetty-YYYY_MM_DD.request.log`
- `jetty-YYYY_MM_DD.stderrout.log`

Fusion uses the [Apache Log4j 2](#) logging framework with Jetty to log each of the Fusion components. Logging is configured via an xml configuration file named `log4j2.xml`. Log levels, frequencies, and log rotation policy can be configured by changing these configuration files:

API service	<code>fusion/3.1.x/conf/api-log4j2.xml</code>
-------------	---

connectors	<code>fusion/3.1.x/conf/connectors-log4j2.xml</code>
Solr	<code>fusion/3.1.x/conf/solr-log4j2.xml</code>
Spark	<code>fusion/3.1.x/conf/spark-driver-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-master-agent-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-master-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-worker-agent-log4j2.xml</code> <code>fusion/3.1.x/conf/spark-worker-log4j2.xml</code>
UI	<code>fusion/3.1.x/conf/ui-log4j2.xml</code>
ZooKeeper	<code>fusion/3.1.x/conf/zk-log4j2.xml</code>

The [Log4j2 Configuration](#) guide provides documentation and examples of all logging configuration options.

4.5.7. Default Ports

Fusion services run in their own JVM and listen for requests on a number of ports. Environment variables, set in a common configuration file, are used to specify the port a service uses. To change the port(s) a service uses, you must change the settings in the configuration file.

Default Ports

The default ports for the Fusion services are as follows:

Port	Service
8764	Fusion UI This service includes the Fusion Authorization Proxy
8765	Fusion API Services
8766	Spark Master
8769	Spark Worker
8984	Connectors Services
8983	Solr This is the embedded Solr instance included in the Fusion distribution.
9983	ZooKeeper The embedded ZooKeeper used by Fusion services. It corresponds to the ZooKeeper <code>clientPort</code> which is defined in file <code>fusion/3.1.x/conf/zookeeper/zoo.cfg</code> .
8766	Apache Spark master REST port
8767	Apache Spark master UI
8082 and 8770	Apache Spark worker UI
4040	Apache Spark driver UI
8769	Apache Spark worker listening port
7337	Shuffle port for Apache Spark worker
8600-8616	Akka ports used between Spark driver, master, workers and API See aka documentation
47100-48099	Apache Ignite TCP communication port range (used by API, Connectors and UI Proxy)

Port	Service
48100-48199	Apache Ignite shared memory port range (used by API, Connectors and UI Proxy)
49200-49299	Apache Ignite discovery port range (used by API, Connectors and UI Proxy)
51500-52000	Executor port, driver port, block manager port, file server port You will need to set these manually in config if needed. Otherwise you can ignore these.

Important	In a production environment, do not expose port 8765 to users. Using your firewall software or the Jetty configuration of the API server, make it accessible only to the auth proxy service and the connectors service.
-----------	---

Port settings are defined in the `fusion.properties` file.

Jetty is used to run Solr, the Fusion UI, API, and connectors services. For each of these services, Jetty runs the service on the assigned port and listens on a second port for shutdown requests. Therefore, `fusion.properties` defines pairs of ports for components running on Jetty, e.g.:

```
api.port = 8765
api.stopPort = 7765
```

ZooKeeper Port Configuration

The ZooKeeper ports are defined both in the `fusion.properties` file and in the zookeeper configuration file, `zoo.cfg`, in the zookeeper subdirectory, path `fusion/3.1.x/conf/zookeeper/zoo.cfg`.

The definition in

`fusion.properties` is:

```
zookeeper.port = 9983
```

The definition in `zoo.cfg` is:

```
clientPort=9983
```

Important	If you change the zookeeper port and are running the embedded zookeeper, the port definitions must match!
-----------	---

4.5.8. Checking System State

As described in the section Default Ports, Fusion runs several components as separate JVMs running on different ports. Each of the components is capable of reporting its status. The proxy component reports status for all of the other components.

Full System Check

To see if each component has been started, a simple API call to the proxy (running on port 8764 by default) will return the status of each component of the system.

```
curl http://localhost:8764/api
```

The response should look similar to the following. If 'ping' is true for each service, all of the system components are running.

```
{
  "version": "0.9.0-SNAPSHOT-jenkins.build.105+git.sha.b425e2a",
  "enabledRealms": [
    "native"
  ],
  "initMeta": {
    "version": "0.9.0-SNAPSHOT-jenkins.build.105+git.sha.b425e2a",
    "initializedAt": "2014-10-06T17:43:31Z",
    "createdAt": "2014-10-06T17:43:31Z"
  },
  "startTime": "2014-10-06T18:38:08Z",
  "status": {
    "connectors": {
      "ping": true
    },
    "apollo": {
      "ping": true
    },
    "apolloZk": {
      "ping": true
    },
    "db": {
      "ping": true
    }
  }
}
```

Solr Health Check

The Fusion UI and API services are not accessible if ZooKeeper and Solr are not in healthy state. A Solr health check can be performed with a ping request.

```
curl http://localhost:8983/solr/admin/ping
```

The response will be a standard Solr XML response, similar to the following.


```
<?xml version="1.0" encoding="UTF-8"?>
<response>
  <lst name="responseHeader">
    <int name="status">0</int>
    <int name="QTime">6</int>
    <lst name="params">
      <str name="df">text</str>
      <str name="echoParams">all</str>
      <str name="rows">10</str>
      <str name="echoParams">all</str>
      <str name="q">solrpingquery</str>
      <str name="distrib">>false</str>
    </lst>
  </lst>
  <str name="status">OK</str>
</response>
```

The 'status' should be "OK" if Solr is running properly.

REST API Services Health Check

All of the Fusion API backend services (except Connectors and the UI) are started at port 8765 when the run.sh script is executed. The Fusion UI depends on all these services.

If all the services are started without any issues, then the below ping request should return response 'ok'.

```
curl http://localhost:8765/api/v1
```

As an additional check, you can also query the `system/status` endpoint, which should return a response 'started'.

```
curl http://localhost:8765/api/v1/system/status
```

The response would look like:

```
{
  "status" : "started"
}
```

Connectors Health Check

The Connectors health check can be performed by a ping request to port 8984. Similar to the previous ping request, the returned response is 'ok' if the service starts successfully.

```
curl http://localhost:8984/connectors/v1
```

As an additional check, you can also query the `system/status` endpoint, which should return a response 'started'.

```
curl http://localhost:8984/connectors/v1/system/status
```

The response would look like:

```
{  
  "status" : "started"  
}
```

4.5.9. Migrating Fusion Data

The instructions in this topic can be used to migrate Fusion data from development environments into testing and production environments, or to back up data and restore it after an incident of data loss.

- Collections and related configurations can be migrated using the Objects API and the Fusion UI (import only). Fusion objects include all your searchable data, plus pipelines, aggregations, and other configurations on which your collections depend.
- Application configuration data includes

Migrating collections and related configurations

Fusion allows you to export objects from one Fusion instance and import them into another. The data that you can migrate includes collections and all collection-related configurations.

Exporting can only be performed using the Objects API. Importing can be performed using the API or the UI.

Object export and import

Collections and encrypted values are treated specially; details are provided below. During import, conflicts are resolved according to the specified import policy.

For objects other than collections, no implicit filtering is performed; all objects are included by default. However, on export you can filter by type and ID.

Supported objects

Fusion lets you export and import these types of objects:

- `collection`
- `index-pipeline`
- `query-pipeline`
- `search-cluster`
- `datasource`
- `banana`
- `parser`
- `group`
- `link`
- `task`
- `job`
- `spark`

Exporting and importing collections

Collections are processed with these dependent objects:

- features
- index profiles
- query profiles

Datasources, parser configurations, and pipeline configurations are not included when collections are exported or imported. These must be exported and imported explicitly.

Only user-created collections are included by default. Certain types of collections are excluded:

- the "default" collection
- collections whose type is not DATA
- collections whose names start with "system_"
- "Secondary" collections, that is, collections created by features

Instead, create the same features on the target system; this automatically creates the corresponding secondary collections.

You can override these exclusions by specifying a collection, like this:

```
http://localhost:8764/api/apollo/objects/export?collection.ids=default
```

Encrypted passwords

Some objects, such as datasources and pipelines, include encrypted passwords for accessing remote data.

- On export, these encrypted values are replaced with `${secret.n.nameOfProperty}`.
- On import, the original, plaintext passwords must be provided in a JSON map:

```
{"secret.1.bindPassword" : "abc", "secret.2.bindPassword" : "def"}
```

The file must be supplied as multipart form data.

Note	Variables that do not start with <code>secret.</code> are ignored.
------	--

Import policies

On import, the `importPolicy` parameter is required. It specifies what to do if any object in the import list already exists on the target system:

<code>abort</code>	If there are conflicts, then import nothing.
<code>merge</code>	If there are conflicts, then skip the conflicting objects.
<code>overwrite</code>	If there are conflicts, then overwrite or delete/create the conflicting objects on the target system.

Filtering on export

On export, there are two ways to specify the objects to include:

- by type

You can specify a list of object types to export all objects of those types. Valid values:

- `collection`
- `index-pipeline`
- `query-pipeline`
- `search-cluster`
- `datasource`
- `banana`
- `parser`
- `group`
- `link`
- `task`
- `job`
- `spark`
- by type and ID

The `type.ids` parameter lets you list the IDs to match for the specified object type.

The `type` and `type.ids` parameters can be combined as needed.

Exporting linked objects

Related Fusion objects are linked. You can view linked objects using the Links API or the Object Explorer.

When exporting a specific Fusion object, you can also export its linked objects without specifying each one individually. To export all objects linked to the specified object, include the `deep="true"` query parameter in your request. See the example below. When `deep` is "true", Fusion follows these link types:

- `DependsOn`
- `HasPart`
- `RelatesTo`

Validation

Objects are validated before import. If any objects fail validation, the whole import request is rejected. A separate endpoint is available for validating objects without importing them.

Validation includes checking whether an object already exists on the target system and whether the user is authorized to create or modify the object.

For collection objects, the following special validation is performed:

- We check the `searchClusterId` of each collection and verify that a cluster with this ID exists on the target system or in the import file (error).
- We check that features, index profiles, and query profiles belong only to the collections specified in the import file (error).
- We check that a feature exists on the target system for each feature in the import file (error).
- We check for index profiles or query profiles that do not exist on the target system or in the import file (warning).

For **job** objects, which contain schedule configurations, Fusion only imports them if their associated **task**, **datasource**, or **spark** objects are also present, either on the target host or in the import file.

Status messages

Validation completed with no errors	The validation method was called and no errors found, though there may be warnings.
Validation found errors	The validation was called and errors found. Validation does not stop on the first error, so the complete list of errors is reported.
Validation was not completed because of system error	The validation was interrupted by system error.
Import was not performed because validation errors exist	The import method was called, but import didn't start because of validation errors.
Import was not performed because of input data error	The import method was called, but import didn't start, because Fusion could not find a substitution for one of the secret values in objects in import.
Import was not completed because of system error	The validation found no errors and import started, but it was interrupted by system error.
Import was completed	Validation found no errors and import finished successfully.

How to export Fusion objects

Exporting can only be performed using the Objects API.

You can select all objects, or limit the operation to specific object types or IDs. In addition to export endpoints, a validation endpoint is provided for troubleshooting.

Note	By default, system-created collections are not exported.
------	--

Some example requests are shown below. For complete reference information about object export endpoints, see the Objects API.

Export all objects

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export
```

Export all datasources

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export?type=datasource
```

Export a specific datasource and its linked objects

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export?export?datasource.ids=movies_csv-ml-movies&deep=true
```

Export all datasources and pipelines, plus two specific parsing configurations

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export?type=datasource,index-pipeline,query-pipeline&parser.ids=cinema_parser,metafiles_parser
```

How to import Fusion objects

Objects can be imported using the REST API or the Fusion UI.

Importing objects with the REST API

Some example requests are shown below. For complete reference information about object export endpoints, see the Objects API.

Import objects from a file and stop if there are conflicts

```
curl -u user:pass -H "Content-Type:multipart/form-data" -X POST -F
'importData=@/Users/admin/Fusion/export.json'
http://localhost:8764/api/apollo/objects/import?importPolicy=abort
```

Import objects, substitute the password variables, and merge any conflicts

```
curl -u user:pass -H "Content-Type:multipart/form-data" -X POST -F
'importData=@/Users/admin/Fusion/export.json' -F 'variableValues=@password_file.json'
http://localhost:8764/api/apollo/objects/import?importPolicy=merge
```


Note

`password_file.json` must contain plaintext passwords.

Importing objects with the Fusion UI

How to import objects using the UI

- 1.



In the upper left, click the


Launcher button and select **Devops**.

2. In the Home panel, click **Import Fusion Objects**.

The Import Fusion Objects window opens.

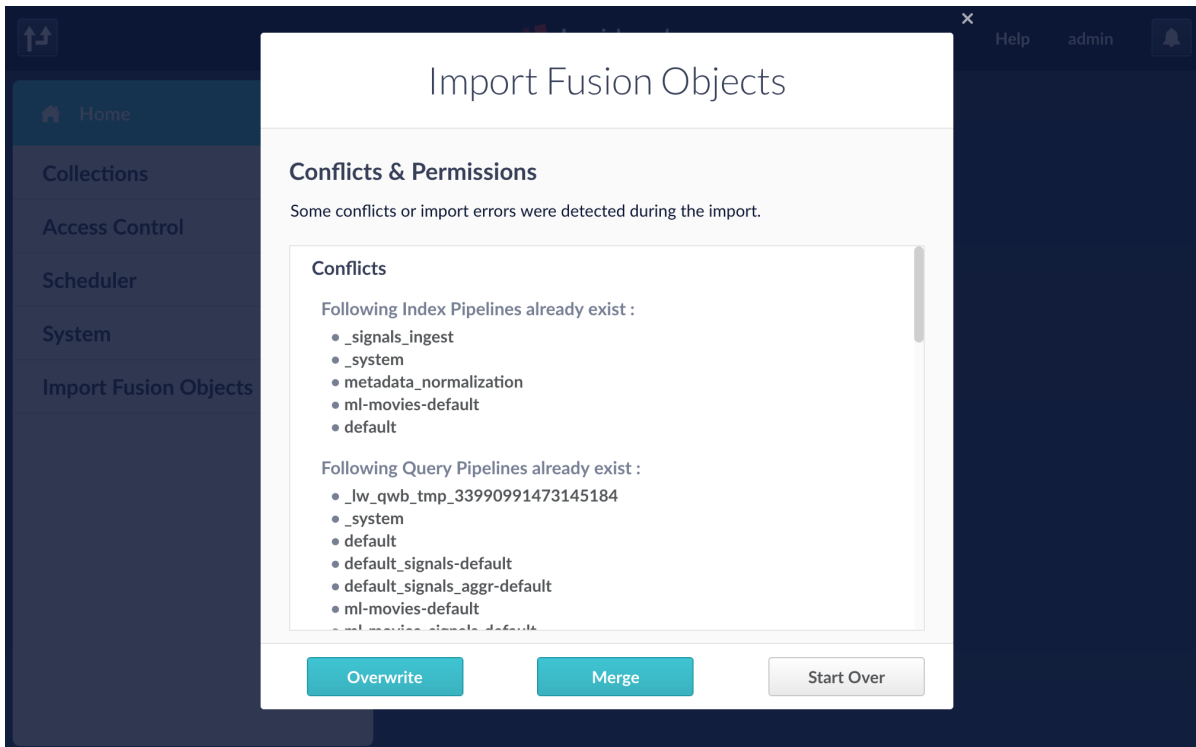
3. Select the data file from your local filesystem.

If you are importing passwords, also select the JSON file that maps variables to plaintext passwords.



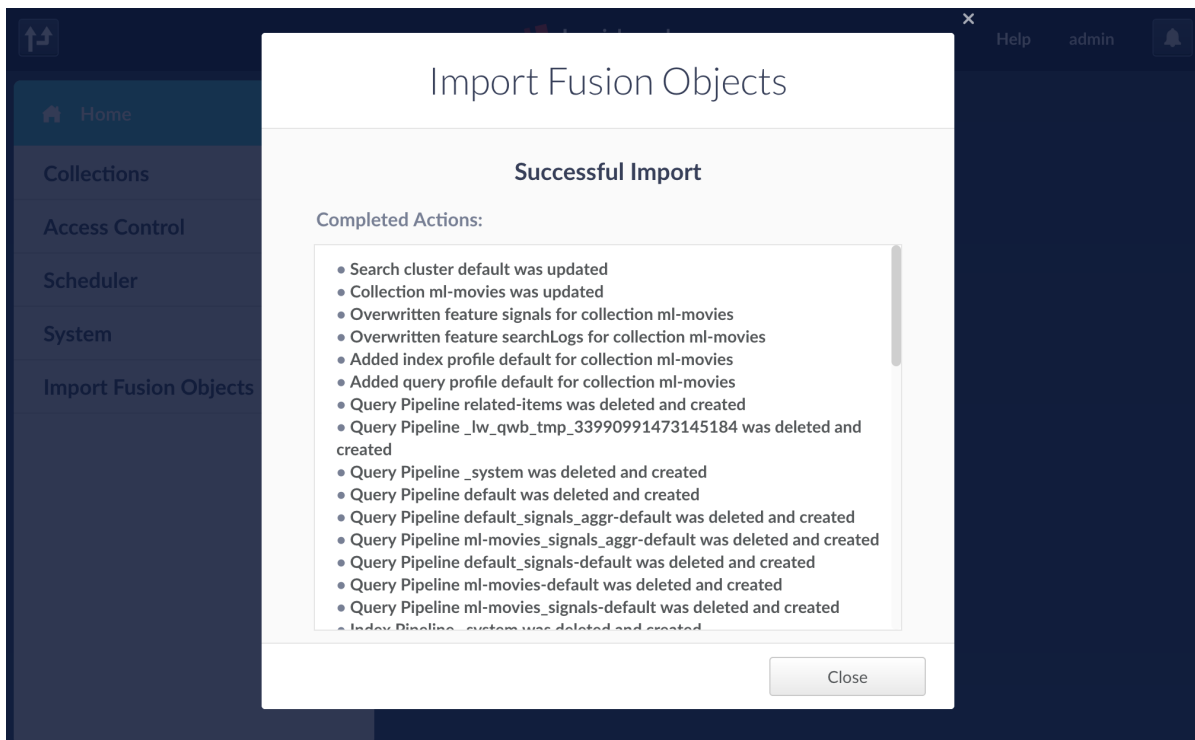
4. Click **Import**.

If there are conflicts, Fusion prompts you to specify an import policy:



- Click **Overwrite** to overwrite the objects on the target system with the ones in the import file.
- Click **Merge** to skip all conflicting objects and import only the non-conflicting objects.
- Click **Start Over** to abort the import.

Fusion confirms that the import was successful:



5. Click **Close** to close the Import Fusion Objects window.

Migrating application configuration data

ZooKeeper configuration data is used to coordinate a distributed Fusion deployment. Additionally, certain Fusion components have configuration data that can be migrated between Fusion instances.

Migrating ZooKeeper data

Migration consists of the following steps:

- Copy the ZooKeeper data nodes which contain Fusion configuration information from the FUSION-CURRENT ZooKeeper instance to the FUSION-NEW ZooKeeper instance
- Rewrite Fusion datasource and pipeline configurations, working against the FUSION-NEW ZooKeeper instance

From ZooKeeper to JSON file

To export configurations from an existing Fusion deployment, the script `zkImportExport.sh` requires parameters:

- `-cmd export` - this is the command parameter which specifies the mode in which to run this program.
- `-zkhost <connect string>` - the ZooKeeper connect string is the list of all servers,ports for the FUSION_CURRENT ZooKeeper cluster. For example, if running a single-node Fusion developer deployment with embedded ZooKeeper, the connect string is `localhost:9983`. If you have an external 3-node ZooKeeper cluster running on servers "zk1.acme.com", "zk2.acme.com", "zk3.acme.com", all listening on port 2181, then the connect string is `zk1.acme.com:2181,zk2.acme.com:2181,zk3.acme.com:2181`
- `-filename <path/to/JSON/dump/file>` - the name of the JSON dump file to save to.
- `-path <start znode>`
 - To migrate Fusion configurations for all applications, the path is `/lucid`. Migrating just the "lucid" node between the ZooKeeper services used by different Fusion deployments results in deployments which contain the same applications but not the same user databases.
 - To migrate the Fusion users, groups, roles, and realms information, the path is `/lucid-apollo-admin`.
 - To migrate all ZooKeeper data, the path is `/`.

Example: export from local developer deployment to file "znode_lucid_dump.json"

```
> {fusion_path}/scripts/zkImportExport.sh -zkhost localhost:9983 -cmd export -path /lucid -filename znode_lucid_dump.json
```

The command products the following terminal outputs:

```

2016-06-01T19:48:12,512 - INFO [main:URLConfigurationSource@125] - URLs to be used as dynamic configuration
source: [jar:file:/Users/demo/tmp5/fusion/apps/jetty/api/webapps/api/WEB-INF/lib/lucid-base-spark-
2.2.0.jar!/config.properties]
2016-06-01T19:48:12,878 - INFO [main:DynamicPropertyFactory@281] - DynamicPropertyFactory is initialized with
configuration sources: com.netflix.config.ConcurrentCompositeConfiguration@5bf22f18
2016-06-01T19:48:12,961 - INFO [main:CloseableRegistry@45] - Registering a new closeable:
org.apache.curator.frameworkimps.CuratorFrameworkImpl@32fe9d0a
2016-06-01T19:48:12,961 - INFO [main:CuratorFrameworkImpl@234] - Starting
2016-06-01T19:48:12,974 - INFO [main:Environment@100] - Client environment:zookeeper.version=3.4.6-1569965,
built on 02/20/2014 09:09 GMT
2016-06-01T19:48:12,974 - INFO [main:Environment@100] - Client environment:host.name=10.0.1.16
2016-06-01T19:48:12,974 - INFO [main:Environment@100] - Client environment:java.version=1.8.0_25
2016-06-01T19:48:12,974 - INFO [main:Environment@100] - Client environment:java.vendor=Oracle Corporation
2016-06-01T19:48:12,975 - INFO [main:Environment@100] - Client
environment:java.home=/Library/Java/JavaVirtualMachines/jdk1.8.0_25.jdk/Contents/Home/jre
2016-06-01T19:48:12,975 - INFO [main:Environment@100] - Client
environment:java.class.path=./fusion/scripts/.. .. ( rest of path omitted )
2016-06-01T19:48:12,976 - INFO [main:Environment@100] - Client
environment:java.library.path=/Users/demo/Library/Java/Extensions: ... ( rest of path omitted )
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client
environment:java.io.tmpdir=/var/folders/jq/ms_hc8f9269f4h8k4b691d740000gp/T/
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:java.compiler=<NA>
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:os.name=Mac OS X
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:os.arch=x86_64
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:os.version=10.10.5
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:user.name=demo
2016-06-01T19:48:12,977 - INFO [main:Environment@100] - Client environment:user.home=/Users/demo
2016-06-01T19:48:12,978 - INFO [main:Environment@100] - Client environment:user.dir=/Users/demo/tmp5
2016-06-01T19:48:12,978 - INFO [main:ZooKeeper@438] - Initiating client connection,
connectString=localhost:9983 sessionTimeout=60000 watcher=org.apache.curator.ConnectionState@138fe6ec
2016-06-01T19:48:18,070 - INFO [main-SendThread(fe80:0:0:0:0:0:1%1:9983):ClientCnxn$SendThread@975] -
Opening socket connection to server fe80:0:0:0:0:0:1%1/fe80:0:0:0:0:0:1%1:9983. Will not attempt to
authenticate using SASL (unknown error)
2016-06-01T19:48:18,111 - INFO [main-SendThread(fe80:0:0:0:0:0:1%1:9983):ClientCnxn$SendThread@852] -
Socket connection established to fe80:0:0:0:0:0:1%1/fe80:0:0:0:0:0:1%1:9983, initiating session
2016-06-01T19:48:18,118 - INFO [main-SendThread(fe80:0:0:0:0:0:1%1:9983):ClientCnxn$SendThread@1235] -
Session establishment complete on server fe80:0:0:0:0:0:1%1/fe80:0:0:0:0:0:1%1:9983, sessionId =
0x1550df6b0180017, negotiated timeout = 40000
2016-06-01T19:48:18,121 - INFO [main-EventThread:ConnectionStateManager@228] - State change: CONNECTED
2016-06-01T19:48:18,367 - INFO [main:ZKImportExportCli@198] - Data written to file
'/Users/demo/tmp5/znode_lucid_dump.json'
2016-06-01T19:48:18,370 - INFO [main:ZooKeeper@684] - Session: 0x1550df6b0180017 closed
2016-06-01T19:48:18,370 - INFO [main-EventThread:ClientCnxn$EventThread@512] - EventThread shut down

```

The resulting JSON output file contains the znode hierarchy for znode "lucid", with ZooKeeper binary data:

exported configurations should be uploaded using the script command argument `-cmd update`.

update command example:

```
> {fusion_path}/scripts/zkImportExport.sh -zkhost localhost:9983 -cmd update -path /lucid -filename
znode_lucid_dump.json
```

To verify, start all Fusion services and log in to the new Fusion installation and verify that this installation contains the same set of collections and datasources as the existing collection, and that all Fusion pipelines and stages match those of the existing Fusion installation.

Existing application, existing Fusion deployment

When migrating an existing application to a Fusion deployment which is already running a version of that application, the exported configurations should be uploaded using the script command argument `-cmd update --overwrite`.

update --overwrite command example:

```
> {fusion_path}/scripts/zkImportExport.sh -zkhost localhost:9983 -cmd update --override -path /lucid -filename
znode_lucid_dump.json
```

To verify, start all Fusion services and log in to the new Fusion installation and verify that this installation contains the same set of collections and datasources as the existing collection, and that all Fusion pipelines and stages match those of the existing Fusion installation.

Caveats

- All datasource configurations are copied over as is. If the set of repositories used to populate the collections changes according to deployment environment, then these datasources will need to be updated accordingly.
- The import export script is only guaranteed to work between Fusion deployments running the same Fusion version. The should work across all releases for the same Major.minor version of Fusion, e.g. you should be able to migrate between versions 2.4.1 and 2.4.2. If the set of configurations needed for an application have the same structure and properties across two different versions, these scripts *might* work.

Migrating Fusion component configuration data

The directory `fusion/3.1.x/data` contains the on-disk data stores managed directly or indirectly by Fusion services.

- `fusion/3.1.x/data/connectors` contains data required by Fusion connectors.
 - `fusion/3.1.x/data/connectors/lucid.jdbc` contains third-party JDBC driver files. If your application uses a JDBC connector, you must copy this information over to every server on which will this connector will run.
 - `fusion/3.1.x/data/connectors/crawlddb` contains information on the filed visited during a crawl. (Preserving crawlddb history may not be possible if there are multiple different servers running Fusion connectors services.)
- `fusion/3.1.x/data/nlp` contains data used by Fusion NLP pipeline stages. If you are using Fusion's NLP components for sentence detection, part-of-speech tagging, and named entity detection, you must copy over the model files stored under this directory.
- `fusion/3.1.x/data/solr` contains the backing store for Fusion's embedded Solr (developer deployment only).
- `fusion/3.1.x/data/zookeeper` contains the backing store for Fusion's embedded ZooKeeper (developer deployment

only).

When migrating these directories, no Fusion services which may change the contents should be running. The choice of which directories to migrate and the utilities used to do the migration are entirely dependent upon the platform, environment, and deployment configurations.

4.6. Blob Storage

Fusion accepts large binary objects (blobs) for upload, and stores them in Solr. Blob uploads are used to install models, lookup lists, JDBC drivers, connectors, and more.

4.6.1. Blob Types

A `resourceType` query parameter can be used to specify the a blob type. For example, specify `plugin:connector` when uploading a connector, like this:

```
curl -H 'content-type:application/zip' -X PUT
'localhost:8765/api/v1/blobs/myplugin?resourceType=plugin:connector' --data-binary @myplugin.zip
```


The complete list of valid values for `resourceType` is below:

Type	Description
<code>catalog</code>	An analytics catalog
<code>driver:jdbc</code>	A JDBC driver
<code>plugin:connector</code>	A connector plugin
<code>model:ml-model</code>	A machine learning model
<code>model:open-nlp</code>	An OpenNLP model
<code>file-upload</code>	Any uploaded file, such as from the Quickstart or the Index Workbench.
<code>banana</code>	A Banana dashboard
<code>other</code>	A blob of unknown type If no <code>resourceType</code> is specified on upload, "other" is assigned by default.

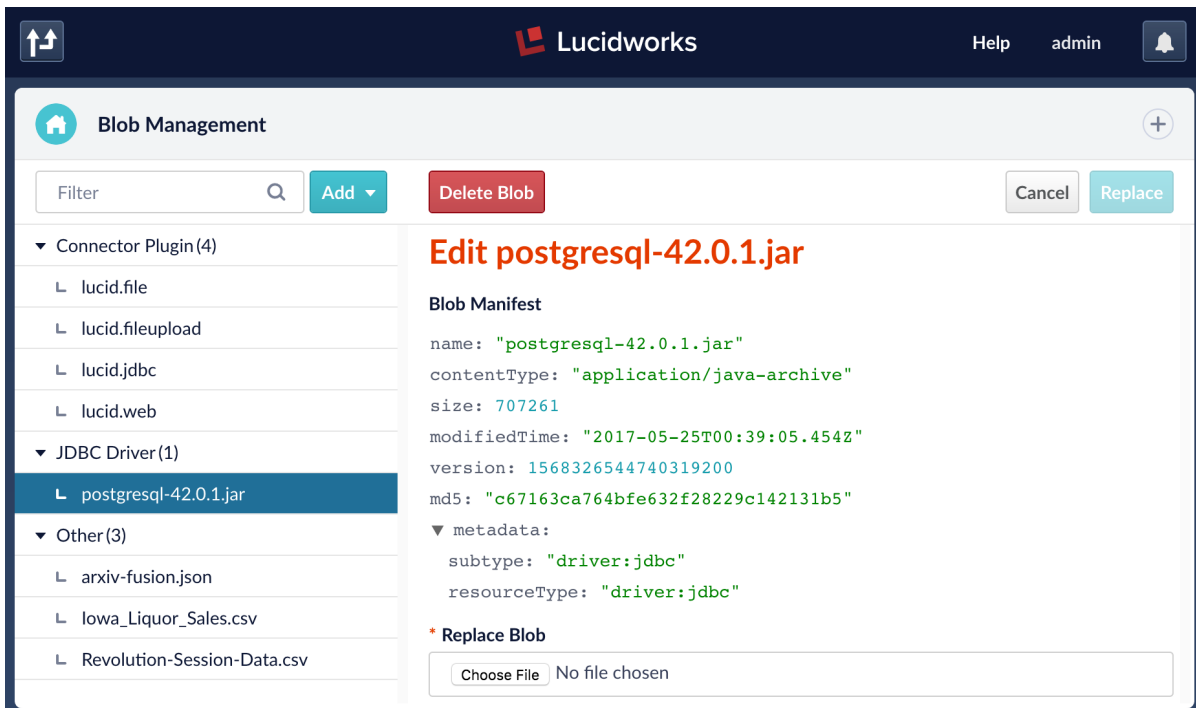
4.6.2. The Blob manager

In addition to the Blob Store API, the Fusion UI provides an interface to the blob store, at **DevOps > Blobs**.

- Click **Add** to upload a new blob:



- Select an uploaded blob to view, replace, or delete it:



4.7. Jobs and Schedules

The Jobs API allows you to view jobs and their run histories, configure their schedules, and control them directly.

In the Fusion UI, jobs are accessible in the Search context, while schedules are in the DevOps context.

- A job is a runnable Fusion object.

Datasources, Spark jobs, and tasks are all jobs.

- A schedule defines when Fusion will perform a job.

A job can be scheduled using cron notation, ISO-8601 interval definitions, or using triggers that depend on the results of other jobs.

4.7.1. Job types

Each job type is a type of Fusion object that you can run or schedule to be run.

<code>datasource</code>	<p>A job to ingest data according to the specified datasource configuration, such as <code>datasource:movie-db</code>. Datasources are created using the Connector Datasources API or the Fusion UI.</p> <p>See Datasource Jobs.</p>
<code>spark</code>	<p>A Spark job to process data, such as <code>spark:dailyMetricsRollup-counters</code>. Spark jobs are created using the Spark Jobs API or the Fusion UI.</p> <p>See Spark Jobs.</p>
<code>task</code>	<p>A job to perform an HTTP call or log cleanup, such as <code>task:delete-old-system-logs</code>. Tasks are created using the Tasks API or the Fusion UI.</p> <p>See Tasks.</p>

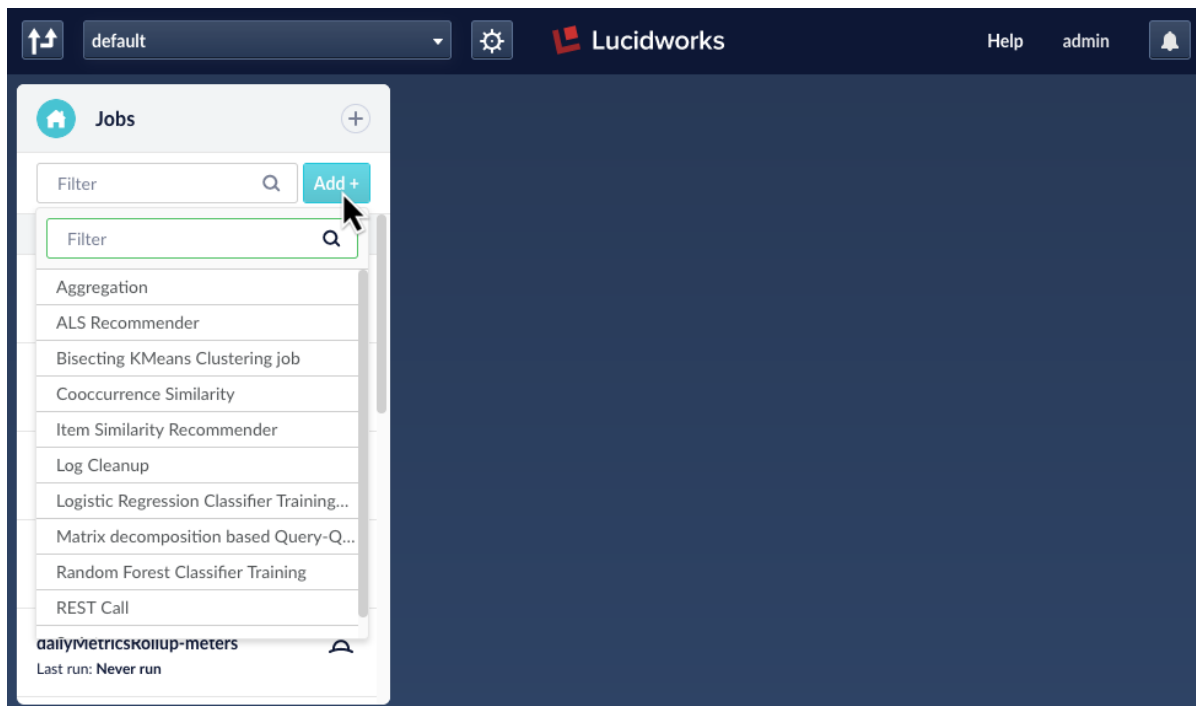
4.7.2. The Jobs manager

The Jobs manager, available in the Fusion UI at **Search** > **Jobs**, provides a simple interface for viewing and scheduling jobs. Tasks and Spark jobs can also be created here.

Datasource jobs cannot be created, run, or scheduled here. They must be created on the Datasources page, in the Index Workbench, or in the Quickstart. To run a datasource once, go to **Search** > **Datasources** or **DevOps** > **Scheduler**. To schedule a datasource job, go to **DevOps** > **Scheduler**.

To create a new job

1. Click **Add** and select the job type.

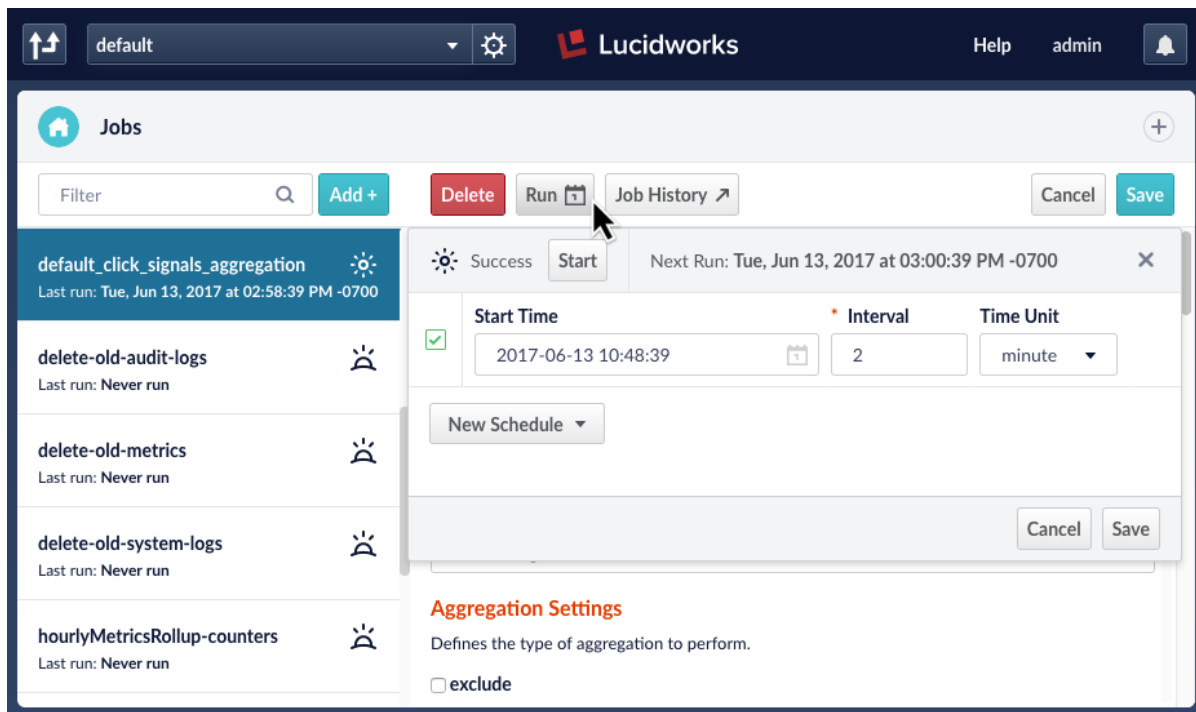


The New Job Configuration panel appears. This panel is different for each job type. See the job types above for details.

2. Configure the new job as needed.
3. Click **Save**.

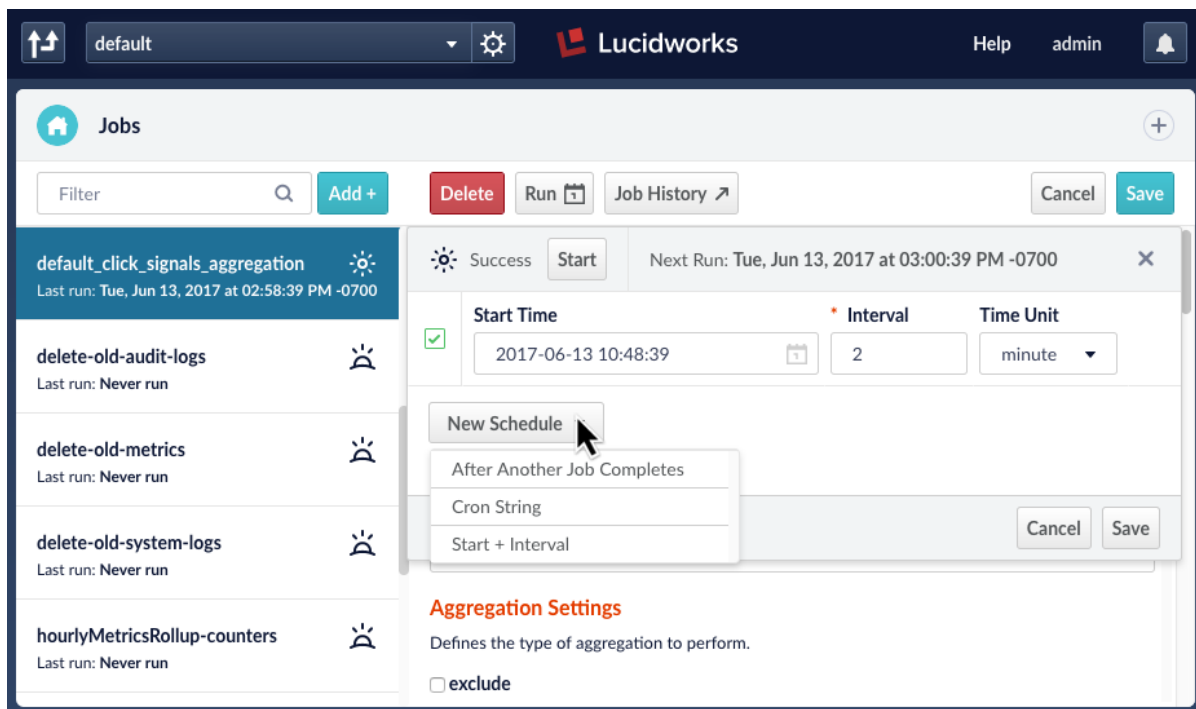
To run a job once

1. Select the job from the job list.
2. Click **Run**.
3. Click **Start**.



To schedule a job

1. Select the job from the job list.
2. Click **Run**.
3. Click **New Schedule**.



4. Select a job trigger:
 - **After Another Job Completes**

Enter the job ID and job result that trigger this one.

The screenshot shows a scheduling configuration window. At the top, there is a sun icon, the text 'Never Run', a 'Start' button, and 'Next Run: N/A'. Below this, there are two dropdown menus: 'Another Job' with the value 'task:delete-old-system-logs' and 'Another Job Status' with the value 'on_success'. A 'New Schedule' button is located below the dropdowns. At the bottom right, there are 'Cancel' and 'Save' buttons.

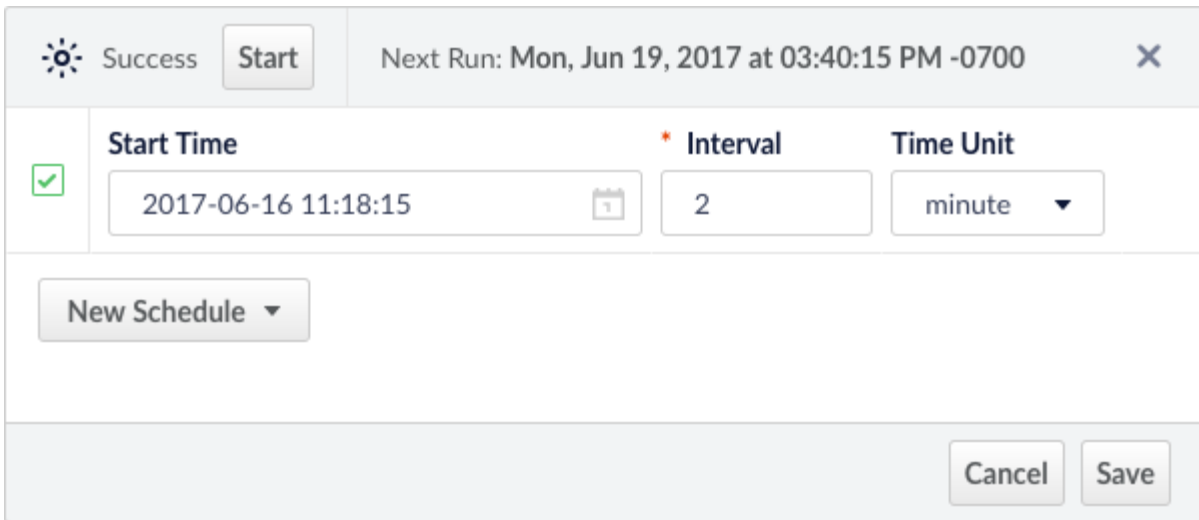
- **Cron String**

Enter a crontab expression.

The screenshot shows a scheduling configuration window. At the top, there is a sun icon, the text 'Never Run', a 'Start' button, and 'Next Run: Sun, Jun 25, 2017 at 04:20:00 AM -0700'. Below this, there is a text field labeled 'Crontab Expression' containing the value '0 20 4 ? * SUN'. A 'New Schedule' button is located below the text field. At the bottom right, there are 'Cancel' and 'Save' buttons.

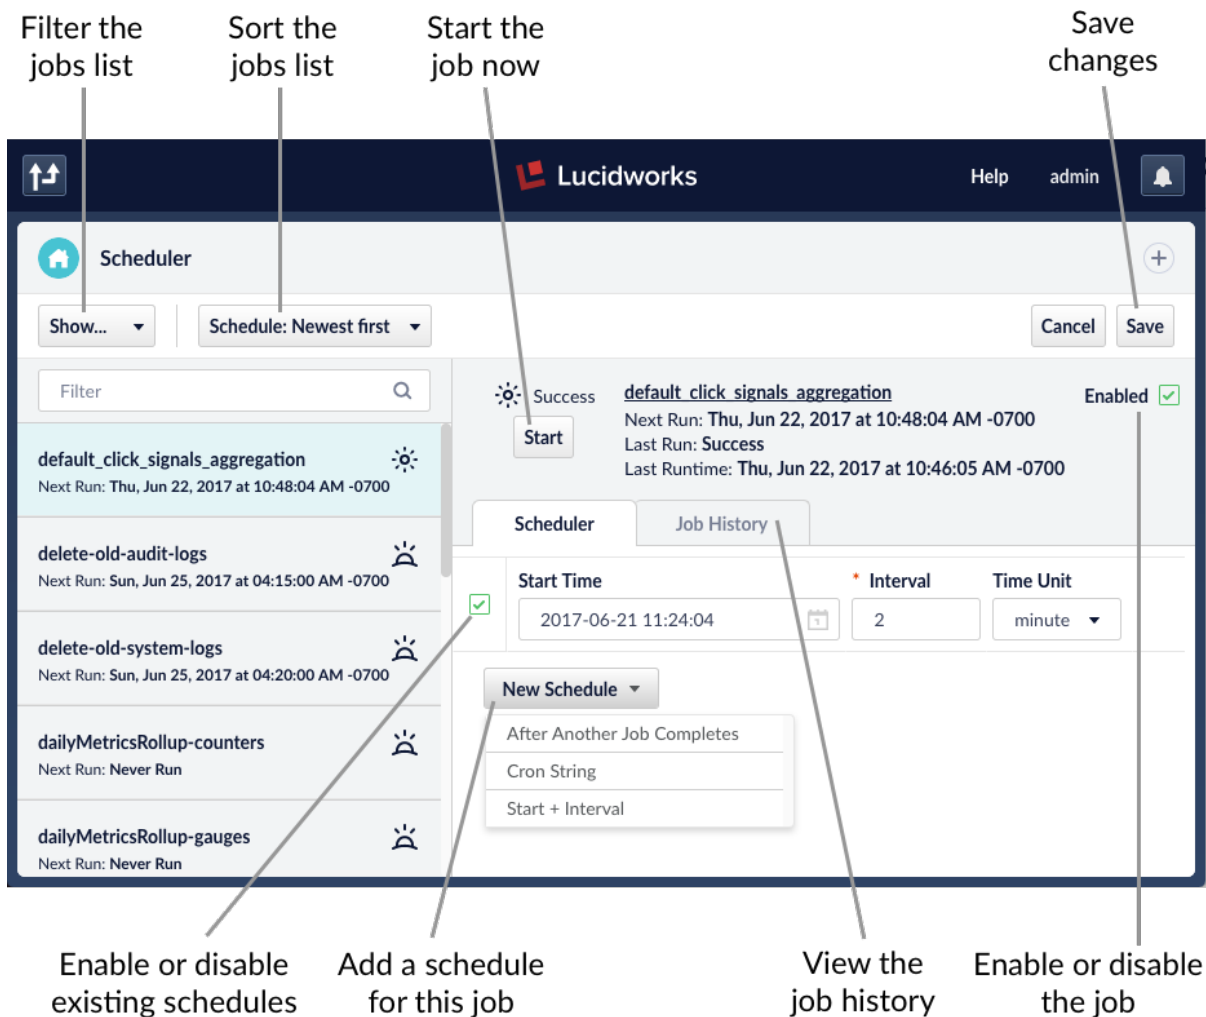
- **Start + Interval**

Enter a start date/time, an interval, and the interval units.



4.7.3. The Scheduler

The Scheduler, available in the Fusion UI at **DevOps > Scheduler**, provides an interface for scheduling jobs. Jobs cannot be created here; to create new schedulable objects, use the Jobs manager, or the REST API for each job type.



4.7.4. Datasource Jobs

A datasource is a configuration that manages the import and indexing of data into the collection.

Datasources are created using the Index Workbench, or at **Search > Datasources**, or using the Connector Datasources API.

A datasource is also a job that can run on demand or on a schedule.

Scheduling a datasource job in the Fusion UI

1. Create a datasource.
2. Navigate to **DevOps > Scheduler**.
3. Select the datasource from the job list.
4. Click **New Schedule**.
5. Select and configure a trigger:
 - **After Another Job Completes**
Enter the job ID and job result that trigger this one.
 - **Cron String**
Enter a crontab expression.
 - **Start + Interval**
Enter a start date/time, an interval, and the interval units.

4.7.5. Spark Jobs

Apache Spark can power a wide variety of data analysis jobs. In Fusion, Spark jobs are especially useful for generating recommendations.

Spark job subtypes

For the Spark job type, the available subtypes are listed below.

ALS Recommender	Train a collaborative filtering matrix decomposition recommender using SparkML's Alternating Least Squares (ALS) to batch-compute user recommendations and item similarities.
Aggregation	Define an aggregation job to be executed by Fusion Spark.
Co-occurrence Similarity	Compute a mutual-information item similarity model.
Random Forest Classifier Training	Train a random forest classifier for text classification.
Script	Run a custom Scala script as a Fusion Job.
Matrix Decomposition-Based Query-Query Similarity Job	Train a collaborative filtering matrix decomposition recommender using SparkML's Alternating Least Squares (ALS) to batch-compute query-query similarities.
Bisecting KMeans Clustering Job	Train a bisecting KMeans clustering model .
Logistic Regression Classifier Training Job	Train a regularized logistic regression model for text classification.
Item Similarity Recommender	Compute user recommendations based on pre-computed item similarity model.
Levenshtein	Compare the items in a collection and produces possible spelling mistakes based on the Levenshtein edit distance.
Collection Analysis	Produce statistics about the types of documents in a collection and their lengths.
Statistically Interesting Phrases (SIP)	Output statistically interesting phrases in a collection, that is, phrases that occur more frequently or less frequently than expected.
Doc Clustering	An end-to-end document clustering job that preprocesses documents, separates out extreme-length documents and other outliers, automatically selects the number of clusters, and extracts keyword labels for clusters. You can choose between Bisecting KMeans and KMeans clustering methods, and between TFIDF and word2vec vectorization methods.
Outlier Detection	Find groups of outliers for the entire set of documents in the collection.
Cluster Labeling	Attach keyword labels to documents that have already been assigned to groups.

Spark job configuration

Spark jobs can be created and modified using the Fusion UI or the Spark Jobs API. They can be scheduled using the Fusion UI or the Jobs API.

To see the complete list of configuration parameters for all Spark job subtypes, use the `/spark/schema` endpoint:

```
curl -u user:pass http://localhost:8764/api/apollo/spark/schema
```


4.7.6. Tasks

Tasks are a flexible job type that can be used to clean up old logs or run any REST call. They are equivalent to "scheduler jobs" in pre-3.1 releases of Fusion.

Task subtypes

The "task" job type has the following subtypes:

Log Cleanup	Delete old log messages from system logs collection.
REST Call	Run an arbitrary REST/HTTP/Solr command.

Task configuration

Tasks can be created and modified using the Fusion UI or the Tasks API. They can be scheduled using the Fusion UI or the Jobs API.

To see the complete list of configuration parameters for all task subtypes, use the `/tasks/_schema` endpoint:

```
curl -u user:pass http://localhost:8764/api/apollo/tasks/_schema
```

4.8. Access Control

4.8.1. User Authentication and Authorization

Fusion provides application security by restricting access to known users via a two-stage process consisting of:

- Authentication - users must sign on using a username and password.
- Authorization - each username is associated with one or more permissions which specify the Fusion UI components and REST API requests that user has access to. Permissions can be restricted to specific endpoints and path parameters. Roles are named sets of permissions which provide access to a specific function.

The access control component runs in the same process as the Fusion UI. It referred to as the "auth proxy" because it handles authentication and authorization for all requests to the Fusion REST API services.

All requests to Fusion must be authenticated, as described in section User Access Request Params.

4.8.2. User Account Administration

A Fusion Security Realm encapsulates a user database together with specific authentication and authorization mechanisms. This information is stored in [ZooKeeper](#) so that it is always available to all Fusion components across the deployment.

Fusion's native security realm manages both authentication and authorization directly. All user information is stored in ZooKeeper: usernames, passwords, roles, and permissions. Stored passwords are encrypted using [bcrypt](#), the strongest possible encryption algorithm available to all JDKs. Authentication consists of a password-hash comparison between the login password and the encrypted password. The native realm is the home of the Fusion admin user and is the default realm type.

Fusion can be configured to use the host domain's security mechanism for user administration. The following configurations are possible:

- LDAP - Fusion stores a local user record in ZooKeeper. Authentication is performed by the LDAP server. LDAP group membership can be used to assign Fusion permissions.
- Kerberos - Fusion stores a local user record in ZooKeeper. [SPNEGO](#) is used for authentication via Kerberos.
- Kerberos authentication, LDAP authorization - Fusion stores a local user record in ZooKeeper. [SPNEGO](#) is used for authentication via Kerberos. LDAP group membership can be used to assign Fusion permissions.
- SAML - Fusion stores a local user record in ZooKeeper. The SAML 2.0 protocol is used to provide web-browser single sign-on.

See also:

- Authentication and Authorization APIs

4.8.3. Video tutorial

4.8.4. Users

All Fusion requests must come from a registered user.

Add Users

The first user who logs in becomes the user `admin`.

There are two approaches for adding users:

- **Manual** – Add users manually to a security realm that doesn't auto-create users.
- **Automatic** – For a security realm that uses an external authentication provider, Fusion can add users automatically. When creating the security realm, check **auto-create users**. Fusion creates a user the first time someone logs into Fusion.

When you add a new user manually, you must provide a unique username and valid password. All other information is optional. However, unless either roles or permissions are specified (or both), this user won't be able to do anything in Fusion.

If you specify API permissions in a user definition, those permissions *override* corresponding permissions defined in the user's roles. See Permissions for more information about how permissions supplied by multiple roles and by user definitions combine.

Manage Users in the Fusion UI

Only Fusion users with administrative privileges (for example, those who are assigned the built-in role `admin`) can manage users.

Manage users in the Fusion UI. Click **Devops > Access Control > Users**.

Manage Users via HTTP Requests to the Users API

See page Users API.

User Information

Fusion stores user information in [Apache ZooKeeper](#).

Each User entry in ZooKeeper contains the following:

- `id`– A globally unique user ID (UUID), created by Fusion based on username, realm-name
- `realm-name`– The Fusion security realm name; the default is "native".
- `username`– The username string, which is unique within the specified security realm
- `permissions`– List of permissions that have been explicitly assigned to the user in the Fusion UI (in **Devops > Access Control**)
- `role-names`– List of roles assigned to the user in the Fusion UI (in **Devops > Access Control**)
- `created-at`– Timestamp; created by Fusion
- `updated-at`– Timestamp for the last edit; created by Fusion

The following JSON shows the ZooKeeper record for the Fusion admin user:

```
{
  "id": "57f539d2-3f53-4011-ad6f-257a3f00fc6b",
  "username": "admin",
  "realm-name": "native"
  "password-hash": "$2a$08$3I82um1XLPShQIW6ngj.Or06DOVgDLGohGmCB9GC0yRtvy5Nfkn6",
  "permissions": [],
  "role-names": ["admin"],
  "created-at": "2016-01-28T00:00:18Z"
}
```

The following JSON shows the ZooKeeper record for a user entry managed by Fusion:

```
{
  "id": "ae9b345a-79e2-4e6d-8620-e6ed4ed2cc16",
  "username": "firstname.lastname",
  "realm-name": "lwLDAP",
  "permissions": [{"path": "collections/**", "methods": ["GET"]}],
  "role-names": [],
  "created-at": "2016-04-01T21:17:36Z"
  "updated-at": "2016-04-01T21:42:15Z",
}
```

4.8.5. Roles

Roles are named sets of permissions that encapsulate the permissions needed for different kinds of users. Permissions grant users access to subsets of Fusion functionality. A role can specify UI permissions, API permissions, or both:

- UI permissions grant users access to parts of the Fusion UI
- API permissions grant users access to specific API commands for specific REST API endpoints.

See Permissions for information about how permissions supplied by multiple roles and by user definitions combine.

Where You Specify Roles

You can specify which roles to apply for a user in one or more of these places:

- **Security realm (directly)** – Under the heading **Roles**, specify the roles to always apply to all users in the security realm.
- **Security realm (from a group/role mapping)** – Security realms of types `ldap` and `trusted-http` can provide a list of groups to which the user belongs. The security realm can map the group names to role names.
- **User definition** – A user definition can specify roles for the user. These roles don't override the other roles. They are added to the other roles.

Default Roles

At initial startup, Fusion creates a set of default roles for common types of users.

admin

The admin role is the the equivalent to the Unix `root` or superuser. It allows full access to all Fusion services:

```
GET,POST,PUT,DELETE,PATCH,HEAD:/**
```

developer

The developer role has all the read/write permissions required for building and running applications.

```

GET,POST,PUT:/system/**
GET,POST,PUT,DELETE,HEAD:/stopwords/**
GET,POST,PUT:/usage/**
GET:/features/**
GET,POST,PUT,DELETE,HEAD:/blobs/**
GET,POST,PUT,DELETE,HEAD:/scheduler/**
GET,POST,PUT,DELETE,HEAD:/aggregator/**
GET,POST,PUT,DELETE,HEAD:/experiments
GET:/introspect/**
PUT:/usage/**
GET,POST,PUT,DELETE,HEAD:/index-stages/**
GET,POST,PUT,DELETE,HEAD:/messaging/**
GET,POST,PUT,DELETE,HEAD:/catalog
GET,POST,PUT,DELETE,HEAD:/parsers/**
GET,POST,PUT:/recommend/**
GET,POST,PUT,DELETE,HEAD:/history/**
GET,POST:/dynamicSchema/**
GET,POST,PUT,DELETE,HEAD:/solr/**
GET,POST,PUT:/signals/**
GET,POST,PUT:/searchLogs/**
GET,POST,PUT,DELETE,HEAD:/query-pipelines/**
GET,POST,PUT:/configurations/**
GET:/suggestions/**
GET,POST,PUT,DELETE,HEAD:/searchCluster/**
GET,POST,PUT,DELETE,HEAD:/index-pipelines/**
GET,POST,PUT,DELETE,HEAD:/spark/**
GET,POST,PUT,DELETE,HEAD:/query-stages/**
GET,POST,PUT,DELETE,HEAD:/prefs/apps/search/*
GET:/nodes/**
GET,POST,PUT,DELETE,HEAD:/solrAdmin/**
GET,POST,PUT:/synonyms/**
GET,POST,PUT,DELETE,HEAD,OPTIONS:/collections/**
GET,POST,PUT,DELETE,HEAD:/connectors/**
GET,POST,PUT:/templates/**
PATCH:/users/{id}:id=#ID
GET,POST,PUT:/registration/**
GET,POST,PUT:/objects/**

```

Note

The permission `PATCH:/users/{id}:id=#ID` uses the variable value `#ID` as a placeholder for the currently logged-in user ID. It is included so the Fusion UI "change password" feature is available to native realm users.

search

The search role has read-only query and write-only signal API access to the Fusion "default" collection. These permissions are required for search applications.

```

POST:/signals/default
GET:/collections/default/query-profiles/default/select
GET:/query-pipelines/default/collections/default/select
PATCH:/users/{id}:id=#ID

```

Note	The permission <code>PATCH:/users/{id}:id=#ID</code> uses the variable value <code>#ID</code> as a placeholder for the currently logged-in user ID. It is included so the Fusion UI "change password" feature is available to native realm users.
------	---

Role Information

Fusion stores role information in [Apache ZooKeeper](#). Each role in a ZooKeeper entry contains the following:

- `id`– ID string, created by Fusion
- `name`– Role name string
- `desc`– Text description; optional
- `permissions`– A list of Fusion permission specifications
- `ui-permissions`– A list of names of Fusion UI components
- `created-at`– Timestamp; created by Fusion
- `updated-at`– Timestamp for last edit; created by Fusion

Manage Roles

Only Fusion users with admin privileges can manage roles.

Restricting access to a subset of Fusion’s functionality requires several narrowly defined permissions. Path variables can be used to designate specific collections. As an example, it’s possible to define a role which allows read-only access to Fusion dashboards for a specific collection:

- `GET:/solr/{id}/*:id=test` – Read-only access to the collection "test"
- `GET:/solr/{id}/admin/luke:id=test` – Also read-only access
- `GET:/solr/system_banana/*` – Read-only access to dashboards
- `GET:/collections/system_banana` – Read-only access to the collection where dashboard definitions are stored

Manage Roles in the Fusion UI

Manage roles in the Fusion UI. Click **Devops > Access Control > Roles**.

To create a new role from the Fusion Admin UI, first you choose a unique role name, then edit the set of permissions. Specify API permissions one per line in the **Permissions** input box. There is a separate list of checkboxes which allow access to the Fusion UI components. If users who are assigned this role require access to the Fusion UI, then you must specify UI permissions in addition to REST API permissions.

Manage Roles via HTTP Requests to the Roles API

See page Roles API.

4.8.6. Permissions

Permissions determine what a user can do in Fusion. There are two kinds of permissions:

- **UI permissions** – Control which parts of the Fusion UI a user can access. These parts show up in menus and the user can view them. But the ability to *use* the functionality depends on API permissions. (The UI uses the API.)
- **API permissions** – Control which requests a user can submit to which REST API endpoints.

Fusion uses permissions for authorization as follows:

- UI permissions are positive (permission needs to be given) and additive (the user has the sum of all specified permissions. This is true of roles specified in a user definition, roles specified in a security realm, and roles determined dynamically based on groups in an LDAP authentication provider.
- API permissions specified in roles are positive (permission needs to be given) and additive (the user has the sum of all specified permissions; that is, for a specific endpoint, the most permissive permissions are used). This is true of roles specified in a user definition, roles specified in a security realm, and roles determined dynamically based on groups in an LDAP authentication provider.
- API permissions specified in the role(s) but not in the user definition are used.
- If an API permission for a specific endpoint is specified in both one or more roles *and* in the user definition, then the permissions in the user definition are used, *overriding* the permissions in the role(s). Use permissions in user definitions to give specific users permissions that are less permissive than the permissions for their role(s). Alternatively, you could define less permissive roles.

Specify UI Permissions

Specify UI permissions in roles.

Specify API Permissions

A Fusion API permission denotes an allowed request to a Fusion REST API endpoint or endpoints. A permissions specification consists of:

- HTTP request method or methods allowed. Multiple HTTP methods are separated by commas.
- REST API services endpoint, which can contain wildcards or named variables. All calls to the REST API start with "api/apollo", followed by the service name and any methods and parameters. The permissions specification includes everything following "api/apollo". The endpoint can include wildcards.

Wildcards make it easy to grant broad access to Fusion services. The wildcard symbol '*' matches all possible values for a single path segment and two wildcards match all possible values for any number of path segments. Granting access to a subset of Fusion's functionality requires a list of narrowly defined permissions.

A path segment can be a named variable enclosed in curly braces: {*variable-name*}. Variables are used when a wildcard would be too permissive and a single path segment too restrictive.

- Optionally, the allowed values for any named variables in the endpoint. The variable specification component specifies the restricted value or values for all named variables in the path. Each specification consists of the variable name, followed by "=" (the equals sign), followed by one or more values which are separated by commas. If the endpoint specification has multiple variable, the semi-colon character ";" is used as the separator between parameter specifications.

Permissions specifications are coded up as a string using the colon character ":" as the separator between the permission elements.

Here are some examples of permissions specifications:

- `GET:/query-pipelines//collections//select` – Search access to any Fusion collection.
- `GET,PUT:/collections/Collection345/synonyms/**` – Permission to edit synonyms for the collection named "Collection345"
- `GET:/collections/{id}:id=Collection345,Collection346` – Read access to collections named "Collection345" and "Collection346"

In ZooKeeper, both User and Roles entries contain a list of Permission specifications. A Permission entry has three attributes: "methods", "path", and "params".

Example Permissions Set

Wildcards make it easy to give wide access to Fusion services. The permissions for the admin user can be written in a single line:

```
GET,POST,PUT,DELETE,PATCH,HEAD:/**
```

To restrict access to a single collection and a single Fusion facility requires a set of narrowly defined permissions. For example, the following set of permissions allows a user to run Fusion's analytics dashboards over a collection named "test":

- `GET:/solr/{id}/*:id=test`– Read-only access to collection named "test"
- `GET:/solr/{id}/admin/luke:id=test`– Dashboards require read-only access to Solr utility luke to compile collection metrics.
- `GET:/solr/system_banana/*`– Read-only access to dashboards
- `GET:/collections/system_banana`– Read-only access to the collection where dashboard definitions are stored

Read-only access to the dashboard definitions collection means that the user cannot save the configured dashboard back to Fusion.

The following JSON shows how Fusion stores this list of permissions in ZooKeeper:

```
"permissions":[
  {"params": {"id":["mdb1"]},
   "path":"/solr/{id}/*",
   "methods":["GET"]},
  {"params":{"id":["mdb1"]},
   "path":"/solr/{id}/admin/luke",
   "methods":["GET"]},
  {"path":"/solr/system_banana/*",
   "methods":["GET"]},
  {"path":"/collections/system_banana",
   "methods":["GET"]}
]
```

4.8.7. Security Realms

Fusion uses *security realms* to authenticate users of the Fusion UI. Each user has an assigned security realm, which the user must choose when logging in. Choosing a different realm results in an authentication failure.

A security realm also provides a list of roles:

- The list always includes the role(s) that are specified in the security realm.
- (Optional) The security realm can reference one or more Fusion roles and/or get groups to which the user belongs from an external directory service that is the authentication provider. Fusion maps the group names to role names and adds these roles to the user's list of roles.

Note: Fusion doesn't use permissions from the LDAP for authorization of UI access or API requests. Fusion only obtains group names from the LDAP (optionally), which you can configure the security realm to map to role names.

Requests to the Fusion REST API must specify a security realm for per-request authentication, unless a session cookie is used (which contains information about the security realm).

Fusion authorizes requested operations based on API permissions specified for the user and for the user's role(s). Fusion considers the role(s) specified in the user definition and in the security realm. Fusion creates a list of roles when a session is created, that is, when a user logs in or when the Sessions REST API creates a session. Authorization based on permissions is at request time.

You can define multiple security realms for a Fusion instance. This lets you give different sets of users different levels of access to specific Fusion collections.

Security Realm Types

When you create a security realm, you can choose among the following security realm types:

Native

Fusion has a single preconfigured security realm named *native*. The admin user is in the native realm, and is the default realm. The native realm also provides a fallback mechanism in case of LDAP server or communication failure.

This realm is required to bootstrap Fusion. Because all requests to Fusion require authentication and authorization, on initial startup you must access the Fusion UI to set the admin password. After Fusion has a valid admin password, it creates the admin account in the Fusion native realm.

For the native realm, Fusion manages all authentication and permissions information directly.

You can create Fusion user accounts and manage them using either the Fusion UI or the User API.

Stored passwords are encrypted using [bcrypt](#), the strongest possible encryption algorithm available to all JDKs.

SSO Trusted HTTP

The "SSO Trusted HTTP" realm type ([trusted-http](#) in the REST API) is useful in single sign-on (SSO) environments.

If the SSO environment contains groups that make sense regarding partitioning Fusion functionality for Fusion users (that is, giving Fusion users different UI and API permissions based on the SSO groups), then you can configure an SSO Trusted HTTP security realm to return a list of group names, and then map the groups to Fusion roles in the security realm definition.

See [Configuring Fusion for SSO](#).

Kerberos

In the case where a host domain uses Kerberos for authentication and LDAP for authorization, Fusion can be configured to do the same, by configuring a realm of type "LDAP" and then specifying Kerberos as the authentication mechanism.

Fusion stores a local user record in ZooKeeper and a mapping to the Kerberos [principal](#).

[SPNEGO](#) is used for authentication via Kerberos.

See [Configuring Fusion for Kerberos](#).

LDAP

You can use an LDAP as an authentication provider for Fusion. If the LDAP contains groups that make sense regarding partitioning Fusion functionality for Fusion users (that is, giving Fusion users different UI and API permissions based on the LDAP-group memberships of LDAP users), then you can configure an LDAP security realm to search for LDAP groups and to map the LDAP groups to Fusion roles.

Fusion stores a local user record in ZooKeeper, and authentication is performed by the LDAP server. User accounts can be managed by Fusion, or created automatically, in which case the Fusion user ID maps directly to the LDAP [Distinguished Name](#) (DN). Fusion permissions can be assigned automatically based on LDAP group membership.

See [Configuring Fusion for LDAP](#).

SAML

Fusion stores a local user record in ZooKeeper and the URL and information about the SAML Identity Provider. The SAML 2.0 protocol is used to provide web browser single sign-on.

See [Configuring Fusion for SAML](#).

Manage Security Realms

Only Fusion users with admin privileges can manage security realms. There are two ways to manage security realms:

In the Fusion UI

Navigate to **Devops > Access Control > Security Realms**.

Using the Realms API

Use the <http://localhost:8764/api/realm-configs/> endpoint to manage security realms. See the Realms API reference for details. In production environments, use port 8765.

4.8.8. Configuring Fusion for SSO

The "SSO Trusted HTTP" realm type (`trusted-http` in the REST API) is useful in single sign-on (SSO) environments.

If SSO is already set up in your environment, user identities and group information can be sent to Fusion through HTTP headers (`REMOTE_USER`, for example). The SSO Trusted HTTP realm type provides the configuration options for integrating this into Fusion's authentication systems. It also supports allowing access to only a set of known client IPs, and mapping groups to Fusion roles.

Use the Realms API to configure this realm type:

```
curl -u user:pass -H 'content-type:application/json' -X POST :3000/api/realm-configs -d @./realm-config.json
```

Below is a sample configuration:

```
{
  "id": "test-id",
  "enabled": true,
  "name": "sso-test",
  "realmType": "trusted-http",
  "config": {
    "identityKey": "REMOTE_USER",
    "groups": {
      "key": "GROUPS",
      "delimiter": "|",
      "roleMapping": [
        ["a", "admin"],
        ["b", "foo"]
      ]
    },
    "allowedIps": ["127.0.0.1", "0:0:0:0:0:0:1", "localhost"]
  }
}
```

<code>identityKey</code>	<p>The name of an HTTP headers entry. If this key is found in the headers map, it used as the identity of the client (username, for example).</p> <p>The <code>X-FORWARDED-FOR</code> header is inspected for this client IP first; the value is split on comma, and the first entry is taken. This would normally be used in cases where the client was forwarded to Fusion through one or more external proxy servers. If the <code>X-FORWARDED-FOR</code> header is not present in the request, the <code>REMOTE-ADDR</code> header value is used instead.</p>
<code>groups</code>	<p>Configuration keys for auth groups:</p> <ul style="list-style-type: none">• <code>key</code> + The name of an HTTP header, used as the source of group names.• <code>delimiter</code> + The character used to split the value (defaults to comma).• <code>roleMapping</code> + A set of 2-tuples, used for mapping the external group values to Fusion Roles.

`allowedIps`

Allow access to only a set of known client IPs. When this property is defined and the client IP is not included in it, the realm logic return a 401.

4.8.9. Configuring Fusion for LDAP

You can create security realms that use external LDAP servers for authentication. Optionally, Fusion can search in the LDAP for groups to which a user belongs, and then map those groups to Fusion roles. Fusion performs authorization using permissions stored in Fusion users and Fusion roles.

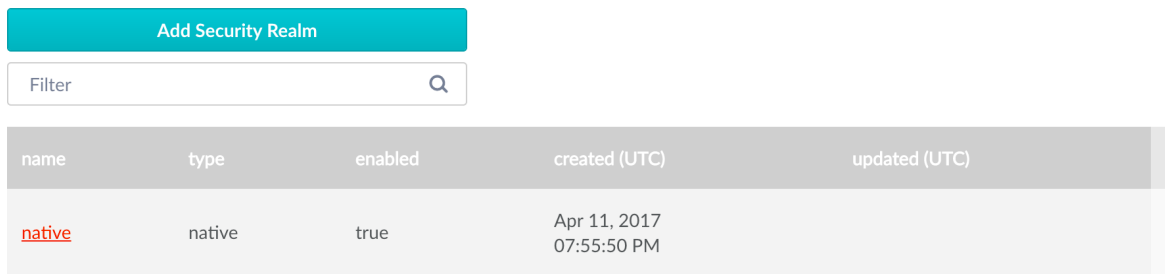
Note: Fusion *doesn't* use permissions from the LDAP for authorization of UI access or API requests. It only obtains group names (optionally), which are mapped to role names. If an Active Directory Security Query Trimming Stage is used, then directory-service permissions *are* used for trimming. If a connector supports security trimming, then connector permissions *are* used for trimming.

To configure Fusion to use an external LDAP as an authentication provider, you'll need to get information about the LDAP server(s) running on your system, either from your system or your sysadmin.

[[1-add-an-ldap-security-realm]] == 1. Add an LDAP Security Realm

1. log in to the Fusion UI as the user admin, or as a different user with corresponding permissions.
2. Click **Devops > Access Control > Security Realms > Add Security Realm**.

Security Realms



name	type	enabled	created (UTC)	updated (UTC)
native	native	true	Apr 11, 2017 07:55:50 PM	

3. Specify info for the new realm:

- name – Name of the security realm. It must be unique. It should be descriptive but short.
- type – Choose **ldap** from the pulldown menu.
- enabled checkbox – Whether Fusion allows user logins for this security realm. The default is yes (checked).
- auto-create users checkbox – Whether a user account is created automatically upon initial authentication. The default is yes (checked). If the checkbox is unchecked, then a Fusion user with admin permissions must create Fusion users.

New Security Realm

* name

realm name

* type

enabled

auto-create users

Enables/disables the auto-creation of Fusion user accounts. This process occurs only when users successfully authenticate for the first time. If disabled, user accounts will need to be manually created.

[[2-specify-static-roles-optional]] == 2. Specify Static Roles (Optional)

Specify one or more Fusion roles for the security realm. These roles are always considered. They don't depend on searching for LDAP groups and mapping group names to Fusion role names.

In a security realm, you can specify these static roles, add to the list of roles dynamically through an LDAP search, or both. If you do neither, Fusion uses only the role(s) and permissions defined for the user.

Roles

These roles are dynamically applied to users.

search

Provides read-only query and write-only signal API access to the Fusion "default" collection.

developer

Developer access with read/write permission required for building/running applications.

admin

Full access to every service. This is the super-admin role.

[[3-specify-ldap-connection-details]] == 3. Specify LDAP Connection Details

Specify the hostname and port of the LDAP server. Check the checkbox if the server is running over SSL.

LDAP Realm

Connection Details

*** host**

The LDAP server host

*** port**

The LDAP server port. Common defaults are 389 and 636 (SSL).

SSL?

[[4-specify-the-authentication-method]] == 4. Specify the Authentication Method

Specify the authentication method:

- Bind - LDAP authentication is carried out via a single "Bind" operation.
- Search - LDAP authentication is carried out indirectly via a Search operation followed by a Bind operation.
- Kerberos - Kerberos authenticates Fusion and an LDAP Search operation is carried out to find group-level authorizations.

Bind

Use the Bind authentication method when the Fusion login username matches a part of the LDAP distinguished name (DN). Specify the remainder of the LDAP DN in the "DN Template" configuration entry, which uses a single pair of curly braces ({}) as a placeholder for the value of the Fusion username.

Authentication Method: bind

- Bind
- Search
- Kerberos

*** Direct Bind Login: DN Template**

A DN template to use for direct bind user logins

Example: uid={},ou=users,ou=system

Search

Use the Search authentication method when the username used for Fusion login *doesn't* match a part of the LDAP DN. The search request returns a valid user DN, which is then used together with the user password for authentication via a Bind request.

1. Construct a search request

The Search authentication method is generally required when working with Microsoft Active Directory servers. In this case, you need to know the username and password of *some* user who has sufficient privileges to query the LDAP server for user and group memberships; this user doesn't have to be the superuser.

In addition to a privileged user DN and password, the Search authentication method requires constructing a search request. There are two parts to the request. The first part is the base DN of the LDAP directory tree that contains user account objects. The second part of the request is a Search Filter object that restricts the results to a matching subset of the information.

Authentication Method: search

- Bind
- Search
- Kerberos

* Search Based Login: Base DN

The base DN to use for search-based user logins
Example: ou=users,ou=system

* Search Based Login: Filter Template

The base DN to use for search-based user logins. A template placeholder of {} will be replaced with the users login-value.
Example: (&(uid=){})(objectClass=inetOrgPerson)

2. Provide the administrator bind DN:

Super User

Bind DN

The admin bind DN - used for search-based user logins and group queries
Example: uid=admin,ou=system

Bind Password

The admin bind password - used for search-based user logins and group queries

Kerberos

Use the Kerberos authentication method when Kerberos is the authentication provider.

Authentication Method: kerberos

- Bind
- Search
- Kerberos

Service Principal

The name of the HTTP service principal

Keytab File

A local file path to a keytab file. This keytab must contain the Service Principal's key. As a fallback, if the KRB5CCNAME environment variable is set, Fusion will use this as a Ticket Cache.

[[5-search-for-ldap-groups-optional]] == 5. Search for LDAP Groups (Optional)

A Fusion role is a bundle of permissions tailored to the access needs of different kinds of users. Access to services and data for LDAP-managed users is controlled by mappings from LDAP users and groups to Fusion roles.

Roles can be assigned globally or restricted to specific LDAP groups. The security realm configuration panel contains a list of all Fusion roles with a checkbox for each, used to assign that role to all users in that realm. LDAP group names can be mapped directly to specific Fusion roles and LDAP group search and filter queries can also be used to map kinds of LDAP users to specific Fusion roles.

Group / Role Mapping

Group Search: Base DN

A base DN for finding groups
Example: ou=groups,ou=system

* Group Search: Name Attribute

The attribute to use for group names

* Group Search: Filter Template

A filter query to use for finding groups. A template placeholder of {} will be replaced with the users' DN.
Example: (&((member=uid={})(uniqueMember=uid={}))((objectClass=groupOfNames)(objectClass=groupOfUniqueNames)))

[[6-map-ldap-groups-to-fusion-roles-optional]] == 6. Map LDAP Groups to Fusion Roles (Optional)

If LDAP group names returned by the search for groups match Fusion role names, you don't need to map the group names to role names. You must map any LDAP group names that don't match to Fusion role names (if you don't, they

won't be used).

Group Mapping

[add new mapping](#)

Engineering	developer
-------------	-----------

[remove](#)

key value pairs of LDAP group names to Fusion role names, one per line

[[7-save-and-test-the-security-realm-configuration]] == 7. Save and Test the Security Realm Configuration

The last part of the form allows you to test the LDAP realm configuration using a valid LDAP username and password:

When the "Update and test settings" button is clicked, the username from the form is turned into a DN according to the DN template, and a Bind operation request is sent to the configured LDAP server.

Test LDAP Connections Settings

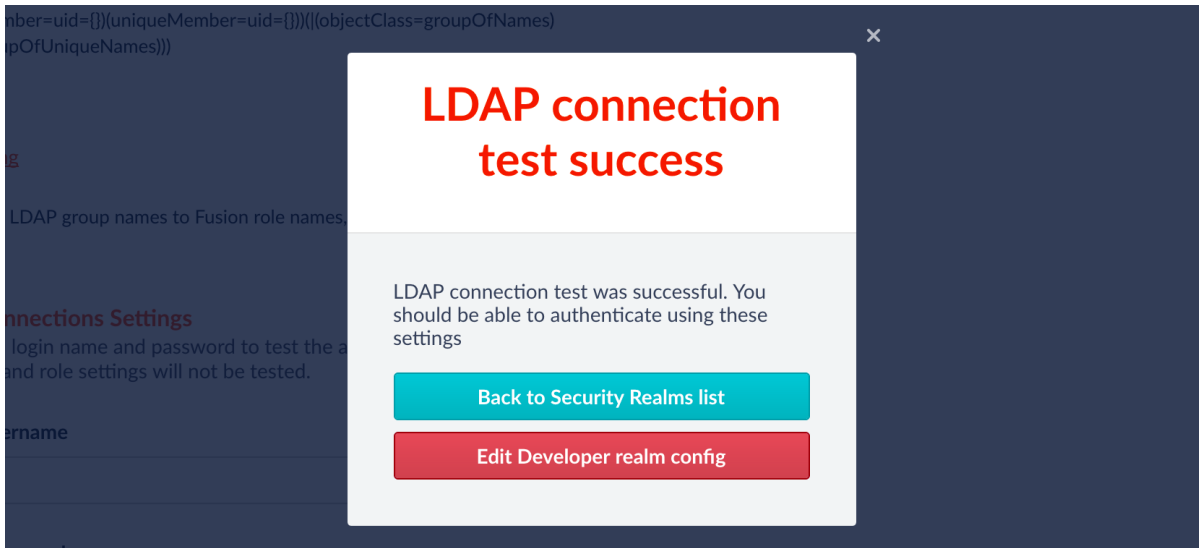
Please provide a login name and password to test the above LDAP connection settings. Group and role settings will not be tested.

Test account username

Test account password

Test on save

Fusion reports whether or not authentication was successful:



Basic LDAP Concepts and Terminology

The [LDAP](#) protocol is used to share information about users, systems, networks, and services between servers on the internet. LDAP servers are used as a central store for usernames, passwords, and user and group permissions. Applications and services use the LDAP protocol to send user login and password information to the LDAP server. The server performs name lookup and password validation. LDAP servers also store Access Control Lists (ACLs) for file and directory objects which specify the users and groups and kinds of access allowed for those objects.

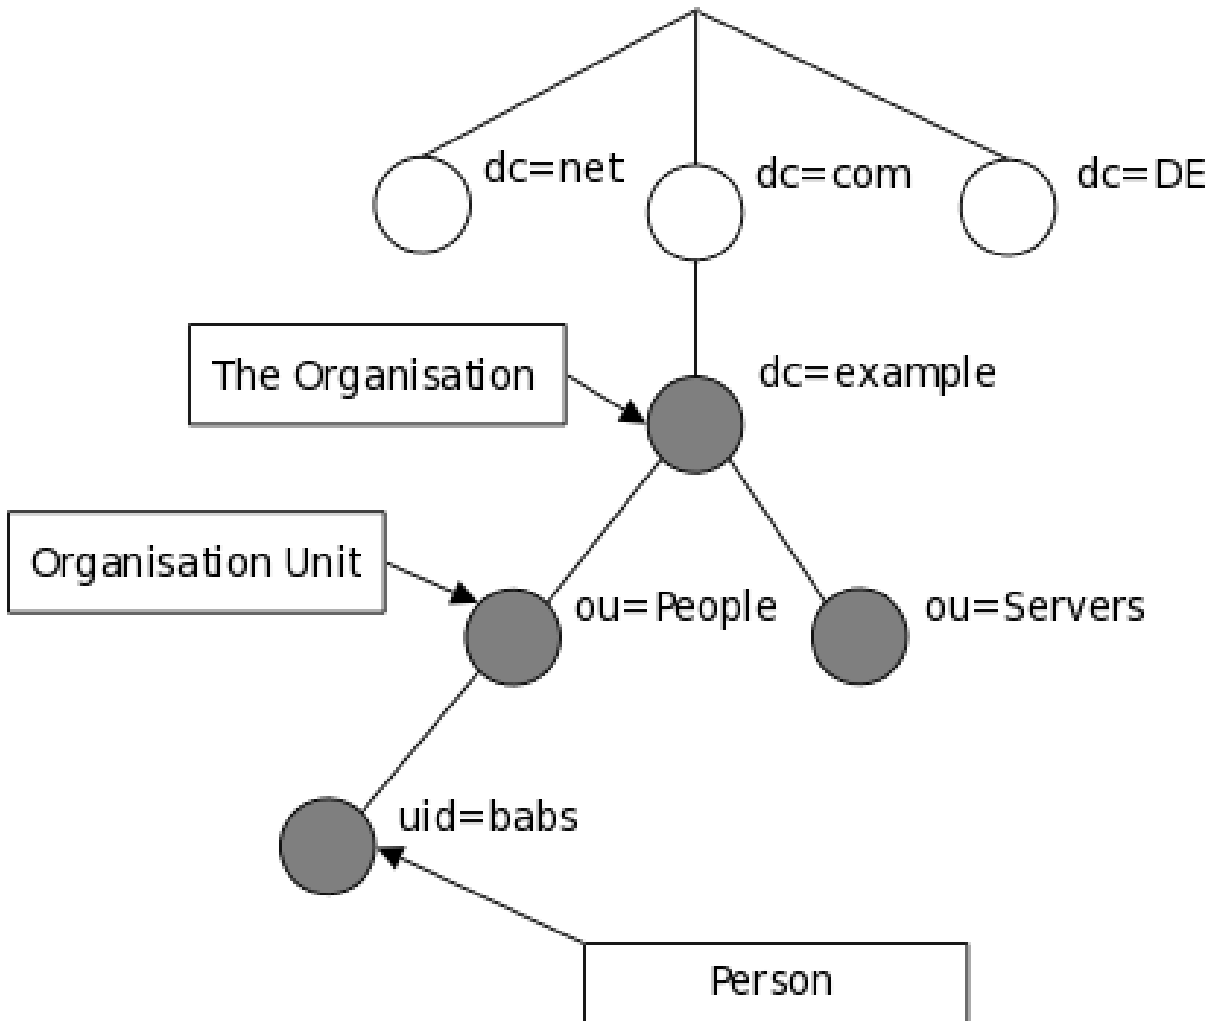
LDAP is an open standard protocol and there are many commercial and open-source LDAP servers available. Microsoft environments generally use Active Directory. Unix servers use AD or other LDAP systems such as OpenLDAP, although many Unix systems don't use LDAP at all. To configure Fusion for LDAP, you'll need to get information about the LDAP server(s) running on your system either from your sysadmin or via system utilities.

Directories and Distinguished Names

An LDAP information store is a Directory Information Tree (DIT). The tree is composed of entry nodes; each node has a single parent and zero or more child nodes. Every node must have at least one attribute which uniquely distinguishes it from its siblings which is used as the node's Relative Distinguished Name (RDN). A node's Distinguished Name (DN) is a globally unique identifier.

The string representation of a DN is specified in [RFC 4514](#). It consists of the node's RDN followed by a comma, followed by the parent node's DN. The string representation of the RDN is the attribute-value pair name, connected by an equals ("=") sign. This recursive definition means that the DN of a node is composed by working from the node back through its parent and ancestor nodes up to the root node.

Here is a small example of a DIT:




The person entry in this tree has the DN: "uid=babs, ou=people, dc=example, dc=com".

Attribute names include many short strings based on English words and abbreviations, e.g.:

Name	Description
cn	commonName
dc	domainComponent
mail	email address
ou	organizationalUnitName
sn	surname
uid	userId

LDAP entry attributes can refer to other LDAP entries by using the DN of the entry as value of that attribute. The following example of a directory which contains user and groups information shows how this works:



This tree contains two organizational units: "ou=people" and "ou=groups". The children of the "group" organizational unit are specific named groups, just as the child nodes of organization unit "people" are specific users. There are three user entries with RDNs "uid=bob", "uid=alice", "uid=bill" and two groups with RDNs "cn=user" and "cn=admin". The dotted lines and group labels around the person nodes indicates group membership. This relationship is declared on the groups nodes by adding an attributes named "member" whose value is a users DN. In the [LDAP data interchange format \(LDIF\)](#), this is written:

```

cn=user,ou=groups,dc=acme,dc=org
  member: uid=bob,ou=people,dc=acme,dc=org
  member: uid=alice,ou=people,dc=acme,dc=org
cn=admin,ou=groups,dc=acme,dc=org
  member: uid=bill,ou=people,dc=acme,dc=org
  
```

See the [Wikipedia's LDAP entry](#) for details.

LDAP Protocol Operations

For authentication purposes, Fusion sends Bind operation requests to the LDAP server. The Bind operation authenticates clients (and the users or applications behind them) to the directory server, establishes authorization identity used for subsequent operations on that connection, and specifies the LDAP protocol version that the client will use.

Depending on the way that the host system uses LDAP to store login information about users and groups, it may be necessary to send Search operation requests to the LDAP server as well. The Search operation retrieves partial or complete copies of entries matching a given set of criteria.

[LDAP filters](#) specify which entries should be returned. These are specified using prefix notation. Boolean operators are "&" for logical AND, "|" for logical OR, e.g., "A AND B" is written "(&(A)(B))". To tune and test search filters for a Unix-based LDAP system, see the [ldapsearch command line utility](#) documentation. For Active Directory systems, see [AD](#)

4.8.10. Configuring Fusion for Kerberos

To configure the Fusion UI service to use Kerberos for user authentication, you must create a Kerberos security realm.

Kerberos is a system that provides authenticated access for users and services on a network. Instead of sending passwords in plaintext over the network, encrypted passwords are used to generate time-sensitive tickets that are used for authentication. SPNEGO provides a mechanism for extending Kerberos to Web applications through the standard HTTP protocol.

Kerberos uses symmetric-key cryptography and a trusted third party called a Key Distribution Center (KDC) to authenticate users to a suite of network services. (By users we mean both end users and client programs). The computers managed by that KDC and any secondary KDCs constitute a realm. When a user authenticates to the KDC, the KDC sends a set of credentials (a ticket) specific to that session back to the user's machine. Kerberos-aware services use the ticket on the user's machine for authentication instead of requiring sign-on with a password. Because tickets are used rather than passwords, this provides the convenience of Single Sign-On (SSO) in addition to security.

A Kerberized process is one that has been configured so it can get tickets from a KDC and negotiate with Kerberos-aware services. When a user sends an HTTP request, Fusion tries to authenticate using the Kerberos/SPNEGO protocol. If the request was sent from a browser, Fusion doesn't display the initial sign-on panel; instead on login, the user sees the main Fusion collections panel.

To Kerberize Fusion, you must:

- Configure the Kerberos client on the server that the Fusion UI service will be running on so that it can talk to the KDC (section Configuring the Kerberos client below).
- Configure the security realm of the Fusion UI (section Configuring Fusion Authentication for Kerberos Realm below).
- Depending on the encryption used by the KDC, you may also need to install additional Java security libraries into the Fusion distribution. These are freely available from Oracle, see download and installation instructions below.

To do this, you need following information, which you can get from your sysadmin:

- Kerberos realm name - in most cases, this is your domain name in upper case.
- KDC name - usually "kerberos." + your domain name.
- Kerberos principal name and password. A principal is a unique identity to which Kerberos can assign tickets. When the entity is a client program, this is called the Service Principal name.
- A keytab file which holds the encrypted credentials.

The usual scenario in an enterprise organization is to have a Kerberos admin create a service principal with a random key password. Then, the admin generates a keytab, which is then used for Fusion service principal authentication.

The Kerberos commands needed for configuration and testing are:

- kinit - obtain and cache a ticket from the KDC (i.e., domain login)
- kdestroy - destroys credentials (i.e., domain logout)
- klist - lists cached credentials
- ktutil - create or add credentials to a keytab file

If your browser is not already configured to use the Kerberos/SPNEGO, you will need to do so in order to test the Fusion

configuration.

Configuring the Kerberos client

Step 1: Edit the Kerberos configuration file.

To configure your local Kerberos client so that it can talk to the Kerberized server, you must edit Kerberos configuration file named `krb5.conf`. On most Unix systems, this file is located at `/etc/krb5.conf`.

This file contains Kerberos configuration information, including the locations of KDCs and admin servers for the Kerberos realms of interest, defaults for the current realm and for Kerberos applications, and mappings of hostnames onto Kerberos realms.

If your organization realm name is "MYORG.ORG", and your KDC server is named "kerberos.myorg.org", then you edit two entries. The first entry is `libdefaults`. Set MYORG.ORG as the default realm:

```
[libdefaults]
  default_realm = MYORG.ORG
```

The second entry is `realms`. Add MYORG.ORG as a realm:

```
[realms]
  MYORG.ORG = {
    kdc = kerberos.myrealm.com
  }
```

For example, for realm LUCIDWORKS.IO, the `krb5.conf` file is just like the above example, except that instead of "myorg.org" we specify "lucidworks.io".

Step 2: Authenticate to Kerberos

The command `kinit` is the Kerberos authentication command. To get started, you authenticate to Kerberos using the Kerberos principal name and password (which you may need to obtain from your sysadmin). For this example, the principal name is "prince".

```
> kinit prince
```

The `kinit` command prompts for a password. Successful authentication is silent. Unsuccessful authentication results in an error message.

The command `klist` shows all cached Kerberos credentials. To check that you've successfully authenticated, run this command:

```
> klist
```

Output should be in this form, but with your data:

```
Credentials cache: API:C980F9F5-415C-4A3E-9C67-883C7D5FFFBE
Principal: prince@MYORG.ORG
```

Issued	Expires	Principal
May 6 15:14:55 2015	May 7 01:12:47 2015	krbtgt/MYORG.ORG@MYORG.ORG

The Service Principal Keytab file

The usual scenario in an enterprise organization is to have a Kerberos admin create a service principal with a random key password. Then, the admin generates a keytab, which is then used for Fusion service principal authentication. If you are your own Kerberos admin, then you will need to create this file for yourself.

Step 3. Create a Keytab file

The command `ktutil` creates the service principal keytab file which holds the encrypted credentials that the Fusion UI Proxy will use for Kerberos authentication. In order to generate the keytab file, you must have a set of cached credentials, therefore, first run the `kinit` command (step 2).

From the command line, run the command `ktutil`. You must enter your password twice.

```
> ktutil -k http-myrealm.org.keytab add -p HTTP/myrealm.org@MYORG.ORG -e aes256-cts-hmac-sha1-96 -V 0
```

The `-k` argument specifies the name of the keytab file which will be created or updated. The command "add" takes the following argument flags:

- `-p` : service principal, format `<service>/<fully.qualified.domain>@REALM`
- `-e` : [encryption type](#). Depending on the encryption type, you may need to download additional Java security libraries for strong encryption.
- `-V` : key version number (kvno). Key version numbers are used in the Kerberos V5 protocol to distinguish between different keys in the same domain.

If successful, this command creates a keytab file called "http-myrealm.org.keytab". Note the directory you're in - you'll need this full path when creating the Proxy realm-config later.

Step 4. Test the Keytab file

The location of this keytab file will be used to configure UI Proxy configuration. Before configuring the Fusion UI Proxy, you should check that the keytab file is valid. Testing the keytab requires the following sequence of steps:

- Clear any existing credentials via command `kdestroy`.
- Log in using the keytab as an argument to the command `kinit -kt <keytab file> <principal>`, where `<principal>` is the name of a principal within the keytab file.
- Examine your credentials via command `klist`.
- Clear credentials via command `kdestroy`, which removes any existing credentials, effectively logging you out of Kerberos.

To remove cached credentials, use the `kdestroy` command. This command succeeds silently. To check that credentials have been removed, re-run the `klist` command:

```
> kdestroy
> klist
klist: krb5_cc_get_principal: No credentials cache file found
```

Use the keytab to login as the service principal, without being prompted for a password:

```
> kinit -t http-myrealm.org.keytab HTTP/myrealm.org
```

Examine your credentials via the command `klist`. The output should be similar to this:

```
Credentials cache: API:51D488FF-5CD9-4E16-98FA-B47743F5B4ED
Principal: HTTP/myrealm.org@MYORG.ORG

Issued                Expires                Principal
Apr  1 09:15:02 2015  Apr  1 19:13:42 2015  krbtgt/MYORG.ORG@MYORG.ORG
```

Logout again with `kdestroy`.

Configuring Fusion Authentication for Kerberos Realm

Once you have tested both the user and service principal logins, you must create the service principal realm-config in the Fusion Authentication Proxy. This allows the Proxy to authenticate to Kerberos as the service principal, without a password.

Step 5. Configure the Fusion Realm

Fusion security realms can be configured either via the Fusion UI Admin tool or the Fusion REST API. The advantage of using the Fusion UI Admin Tool is that a single panel's worth of configuration requires a series of calls to the REST API. It is important to understand the set of configuration properties collected by the Fusion UI and how they are used by the REST API.

A security realm for Kerberos has the following properties:

- `name` : unique string identifier
- `realmType` : "kerberos"
- `enabled` : whether or not the realm is available for users to use with system authentication
- `config` : this property is required for the realm type "kerberos". It takes two key-value pairs:
 - `principal` : the principal service name
 - `keytab` : this must be the full path to the keytab file.

To configure Fusion via the Fusion UI Admin tool, you should be logged in to Fusion as "admin" or as a user who has super-admin privileges.

From the Fusion Admin tool, choose the "Access" control from the left hand side nav bar, and go to the Access tab "Security Realms" (URL: <server>:<port>/admin/security-realms).

Click on the "Add Security Realm" button. On the New Security Realm form, choose type "kerberos" from the pulldown menu so that there are input boxes for the Kerberos realm properties "Service Principal" and "Keytab path". Choose a

realm name, check the "enabled" checkbox, and enter the service principal and the full path to the keytab file.

The New Security Realm form also controls the default roles assigned to a user the first time that they access Fusion via this realm, using the Kerberos/SPNEGO protocol. For example, once you have defined a Kerberos realm "my-kerberos-realm" for domain "MYORG.ORG", when user "any.user" in domain "MYORG.ORG" authenticates to Fusion for the first time via this realm, they will be added to the set Fusion users as username "[any.user@MYORG.ORG](#)" and they will have all default roles.

It is prudent to allow the minimum set of default roles, as all users will have these permissions. Some users will require admin privileges and a few users will require super-admin privileges. There should always be a user with super-admin privileges that can authenticate to Fusion using the native security realm and can then grant permissions to individual users as needed.

Kerberos/SPNEGO HTTP Authentication

[SPNEGO](#) provides a mechanism for extending Kerberos to applications that use the HTTP protocol including web browsers and the `curl` command-line utility.

Step 6. Configure the HTTP client

When a user sends a request to the Kerberized Fusion UI, a SPNEGO request (`http[s]`) is made. If the user is not already authenticated, the Fusion authentication proxy will yield a 401 status code and a Negotiate header. This status/header response triggers compatible clients to fetch a local ticket from their Kerberos "ticket tray". This ticket is then encoded and sent back to the Fusion. The Fusion authentication proxy will then decode the ticket, and perform a SPN.doAs(user) authentication request to the KDC/Authentication Service. Depending on the results, the proxy then successfully executes the original request (along with a session cookie) or a 401 (without the Negotiate). Clients can either choose to use the session cookie or continue authenticating on every request.

Configuring Web Browsers and `curl` for SPNEGO

The `--negotiate` option enables SPNEGO in `curl`.

IE and Safari require no additional configuration to use SPNEGO.

To configure Firefox, access the low level configuration page by loading the `about:config` page. Then go to the `network.negotiate-auth.trusted-uris` preference and add the hostname or the domain of the web server that is HTTP Kerberos SPNEGO protected (if using multiple domains and hostname use comma to separate them).

The Chrome browser must be launched from the command line with several added parameters.

To run Chrome on linux:

```
> google-chrome --enable-plugins --args\  
  --auth-server-whitelist="*KERBEROS_DOMAIN\  
  --auth-negotiate-delegate-whitelist="*KERBEROS_DOMAIN\  
  --auth-schemes="basic,digest,ntlm,negotiate"
```

To run Chrome on a Mac:

```
> open 'Google Chrome.app' --args\  
--auth-server-whitelist="*ROGUECLOUD.COM\  
--auth-negotiate-delegate-whitelist="*KERBEROS_DOMAIN\  
--auth-schemes="basic,digest,ntlm,negotiate"
```

To run Chrome on Windows:

```
chrome.exe --auth-server-whitelist="*KERBEROS_DOMAIN\  
--auth-negotiate-delegate-whitelist="*ROGUECLOUD.COM\  
--auth-schemes="basic,digest,ntlm,negotiate"
```

For more information, see [Using a Web Browser to Access an URL Protected by Kerberos HTTP SPNEGO](#) and <http://www.roguelynn.com/words/apache-kerberos-for-django/>.

Session cookies

A successful Kerberos/SPNEGO login will yield a session cookie, this cookie is identical to the cookie yielded by the Fusion authentication proxy's current POST-login-mechanism.

The expiration policy on the cookie is currently fixed at 8 hours. But has a 1 hour "idle" max, which means if you don't make a request for 1 hour, the cookie is invalidated. Otherwise the lifetime is pushed ahead until the 8 hour max is met.

The name of the cookie is "id" and the value is a UUID. This UUID is a key that maps to an in-memory value containing the real user ID.

Testing and Troubleshooting

Once you have configured the Kerberos security realm, you can test it by logging out of Fusion and shutting down the browser.

Check that you have a valid Kerberos authentication ticket via the klist command.

Now open a new browser session and access the Fusion installation, <domain>:<UI port>.

This should take you directly to the main Fusion panel, bypassing the the Fusion "Welcome" login panel. To view your login profile, click on the profile icon on the top nav bar. Your user name should be your login name + "@" + your domain name.

If instead, you see the Fusion login panel, your browser is not configured for SPNEGO.

If the Fusion display consists only of an empty top nav bar, this indicates an authentication failure. Check that the path to your keytab file is correct. Then check the Fusion logs.

If your KDC uses "AES256 CTS mode with HMAC SHA1-96" for key encryption, the proxy will log this error when attempting to authenticate:

```
GSSEException: Failure unspecified at GSS-API level (Mechanism level: Encryption type AES256 CTS mode with HMAC  
SHA1-96 is not supported/enabled)
```

To get around this, the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy File will need to be downloaded and installed. It can be downloaded from:

- <http://www.oracle.com/technetwork/java/javase/downloads/jce-7-download-432124.html>

Place the jars in your JAVA_HOME/jre/lib/security/ directory, then restart Fusion.

Clicking on the "logout" icon on the top nav bar (rightmost icon or a padlock) takes you back to the main Fusion panel. If you destroy your Kerberos credentials cache via the "kdestroy" command, the next time you logout of Fusion, you will be logged out and the browser will display the Fusion login panel.

References and Tutorials

https://en.wikipedia.org/wiki/Kerberos_%28protocol%29

<http://www.roguelynn.com/words/explain-like-im-5-kerberos/>

<http://thekspace.com/home/component/content/article/54-kerberos-and-spnego.html>

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Managing_Smart_Cards/Using_Kerberos.html#about-kerberos

<http://www.oracle.com/technetwork/articles/idm/weblogic-ss0-kerberos-1619890.html> - section "Install Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files"

<http://www.cisco.com/c/en/us/support/docs/security-vpn/kerberos/16087-1.html>

4.8.11. Configuring Fusion for SAML

[SAML 2.0](#) is a standard for exchanging authentication and authorization data between security domains. The SAML protocol allows web-browser single sign-on (SSO) through a sequence of messages sent to and from the browser, which is the intermediary between Fusion and the SAML authority acting as the Identity Provider (IdP).

Fusion has been tested with the following IdPs (but any IdP should work):

- ADFS
- Okta
- OneLogin
- OpenAM
- Shibboleth

To configure Fusion to use SAML 2.0 for user authentication and authorization you must create a SAML security realm. In addition to configuring the Fusion security realm, you must configure the SAML identity provider to recognize the Fusion application.

Once Fusion is configured for a SAML realm, this realm is added to the list of available realms on the initial Fusion sign-on panel. When the SAML realm is chosen from the list of available realms, the browser then redirects to the IdP which handles user authentication. Upon successful authentication, the IdP sends a response back to the browser which contains authentication and authorization information as well as the URL of the Fusion application. The browser redirects back to the Fusion URL, passing along the SAML message with the user authentication and authorization information. Fusion then issues a session cookie which is used for subsequent user access.

Fusion Configuration for a SAML Realm

You must get the following information about the SAML Identity Provider either from your sys admin or from the IdP directly:

- Identity Provider URL - the URL used by the SAML authority for single sign-on. Usually a URL which ends in "saml/sso", e.g. "https://www.my-idp.com/<my-app-path>/sso/saml"
- Issuer - SAML Issuer Id. A unique ID for that authority, e.g. "http://www.my-idp.com/exk686w2xi5KTuSXz0h7".
- Certificate Fingerprint - the contents of the SAML authority certificate, *without the certificate header and footer*. You must get this certificate from the SAML Identity Provider. The certificate is a text file which has a pair of header and footer lines which say "BEGIN CERTIFICATE" and "END CERTIFICATE", respectively. The fingerprint consists of the lines between the header and the footer. You can cut and paste this information into the text box on the Fusion UI.
- User ID Attribute - an optional attribute. The Identity Provider contains the user database. By default, the Fusion username is the same as the login name known to the Identity Provider. When another field or attribute in the user record stored by the IdP should be used as the Fusion username, that attribute name is the value of the User ID Attribute.

All Fusion security realms require the following information:

- name - must be unique, should be descriptive yet short.
- type - value is "SAML" (one of the choices on the Fusion UI Security Realms config panel).
- "enabled" - default value is true. The "enabled" setting controls whether or not Fusion allows user logins for this security realm.

- "auto-create users" - default is true. This controls whether or not a user account is created automatically upon initial authentication. If false, new user accounts can only be created by a Fusion user with admin privileges.

SAML Authority Identify Provider Configuration for Fusion

The Fusion application must be registered with the SAML Identity Provider. The amount of information varies depending on the SAML authority.

All systems will require the Fusion URL to redirect to upon successful login; this is the protocol, server, and port for the Fusion application, and path "api/saml", e.g. "https://www.my-fusion-app.com:8764/api/saml". If the Fusion application is running behind a load-balancer, then this URL is the load-balancer URL plus path "api/saml". Note that the load-balancer should be session-sticky in order for the sequence of messages that comprise the SAML protocol to run to completion successfully.

Some authorities may require additional information. In particular the SAML 2.0 "AudienceRestriction" tag may be part of the SAML message. This tag specifies the domain for which the SAML trust conditions are valid, which is usually the domain in which the Fusion app is running, e.g. "https://www.my-fusion-app".

Example SAML Realm Configuration

The Fusion endpoint "api/realms-config" returns a JSON list of all the configuration objects for all realms. After configuring a SAML realm named "saml-test" using the okta.com developer preview tool, the configuration object for this realm is:

```
{
  "name": "saml-test",
  "realmType": "saml",
  "enabled": true,
  "config": {
    "autoCreateUsers": true,
    "idpUrl": "https://dev-417804.oktapreview.com/app/dev417804_1/exk686w2xi5KTuSXz0h7/sso/saml",
    "issuer": "http://www.okta.com/exk686w2xi5KTuSXz0h7",

    "certificateFingerprint": "MIIDpDCCAOygAwIBAgIGAVQr4A4NMA0GCSqGSIb3DQEBBQUAMIGSMQswCQYDVQGEwJVUzETMBEG\nA1UECAwKQ2FsaWZvcml5TEWMBQGA1UEBwwNU2FuIEZyYW5jaXNjbzENMAsGA1UECgwET2t0YTEU\nnMBIGA1UECwwLU1NPUHJvdmlkZXIxExARBgNVBA\nMMCMRldi00MTc4MDQxHDAaBgkqhkiG9w0BCQEW\nnDWluZm9Ab2t0YS5jb20wHhcNMTYwNDE5MDAx\nNTI0WWhcNMjYwNDE5MDAxNjI0WjCBKjELMA\nkGA1UE\nnBhMCVVMxExARBgNVBAgMCkNhbG1mb3JuaWE\nXfjAUBgNVBAcMDVNhbiBGcmFuY2lzY28xDTAL\nBgNV\nnBAoMBE9rdGExFDASBgNVBAsM\nC1NTT1Byb3ZpZGVyMRMwEQYDVQDDApkZX\nYtNDE3ODAMRwwGgYJ\nnKoZiHvcNAQkBFg1pbmZvQG9rdGEuY29t\nMIIBIjANBgkqhkiG9w0BAQEF\nAAOCAQ8AMIIBCGKAQEAnuNz0trRFPw2d4x\nXoUvX2oWeZo1TveftaTnB9SWUyjk4og0Wd\nT7rNdbG10eTvB2ezBwXCf24UGGui\nnr1kjkZjiHDqDxKtzQYWPguzLCjh/4P\nxKFGDaiUNKcE1Ig5myiEBTvMv99XtrcI75\nQdUGDhbMiBr\nn2PR5FukW0YepzLBzqY0JSDzX9NYJBKPkz\n+syK4mj0I6dqtYOU+bcTv\njF9sR7jiHtQ+d0Z18rz1Ca\nnyuE3mNUtFJ1IjRy/RARhH1AB6mXbV/de\n1CXmGhKQbqQAbx9S\niKtki9n84gKEwuWdV0jIqcBLGxUQ\n\nngjbsaIVqed2oX+7F2fh6t0g/I8NPn0\nWXOTvA+wIDAQABMA0GCSqGSIb3DQE\nBBQUAA4IBAQCpt+DR\n\nlIiSHO/iVnmLFPgfqrCO/gMv++xnq2\nwOB19YX7HhT1GIY2YzoVphrQpXH3000/\n8AZJ8ApCmq0E9x\n\nnxUTQwQPBanVqlyLtu1Hr0c6dbAqcd5\nPtaEe6Ci33nayWPYdh0mitvIyb/WtWt\nbel9HcdUkoGvka1\n\nng305jnxkhwGLJm4jkzYe+eaYhd6o63/\nJcHqKdGYGf7i2Z\n\nny0D7vxTeBQ+8PbZfsUKg0PLKyTo\ncb\n\nnydSmDPISsA1x0H5zlw+h\nzFIdKgD5vW7QLvpIc\n\nNc2hqki/nWL/CHut0TnUuP/V3boRE\nmDu395n\n\n/u72pgNANfP7+2DTBb+CBTjGUs\nXpKRF",
    "userIdAttribute": ""
  },
  "roleNames": ["search"],
  "id": "52e1c0d2-5e00-4c76-a3d4-57f1381bdb4f",
  "createdAt": "2016-04-19T19:49:04Z",
  "updatedAt": "2016-04-19T20:06:56Z"
}
```


References

<http://docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf>

<http://docs.oasis-open.org/security/saml/v2.0/saml-glossary-2.0-os.pdf>

4.8.12. User Access Request Params

Fusion requests must come from a known user, i.e., a user with a unique user id (UUID). Fusion's ZooKeeper registry tracks all users across all realms. Usernames must be unique within a realm. Fusion creates a globally unique user ID for all users based on the combination of username and realm.

All requests to the Fusion REST API require either a username, password, and security realm name, or the session cookie which contains the unique user ID.

Per-Request Authentication

To pass authentication information with each request, the realmName is specified as a query parameter on the request itself:

```
curl -u joe.smith:password123 "http://www.acme.com:8764/api/apollo/collections?realmName=acmeLDAP"
```

The default realmName parameter is "native", so for native authentication, this parameter can be omitted.

Session Cookies

The Fusion UI service endpoint "api/session" can be used to generate a session cookie which contains the unique user id via a POST request whose body consists of a JSON object which contains the username, password information. For users belonging to a realm other than the native realm, the request parameter "realmName" must be specified. The command to generate a session cookie for the admin user with password "password123" is:

```
curl \
-c cookie -i -X POST -H "Content-type:application/json" -d @- \
http://localhost:8764/api/session?realmName=native \
<<EOF
{ "username" : "admin" , "password" : "password123" }
EOF
```

The curl command takes any number of specialized arguments, followed by the URL of the request endpoint. The arguments used here are:

- **-c** : filename of cookies file. If it exists, cookies are added to it. You can use **-c -** which writes to the terminal window (std out).
- **-i** : include the HTTP-header in the output. Used here to see the cookie returned with the response.
- **-X** : request method, in this case **POST**
- **-H** : request header. The **api/session** endpoint requires **Content-type:application/json**.
- **-d** : Pass POST body as part of the command-line request. To get ready the body from a file, use the syntax **-d @<filename>**. The argument **-d @-** reads the data from stdin.

The header output shows the cookie information:

```
HTTP/1.1 201 Created
Set-Cookie: id=996e4adf-bd04-4058-a926-8ea8ca08c05a;Secure;HttpOnly;Path=/api
Content-Length: 0
Server: Jetty(9.2.11.v20150529)
```

Once the session cookie file has been created, it can be sent along in all subsequent requests to the REST API. For the curl command-line client, the **-b** flag is used to send the contents of the cookie file to the server along with the request.

The following command sends a GET request to the Fusion REST API Collections service to check the status of the "system_metrics" collection. The **-b** flag sends in a freshly generated session cookie.

```
> curl -b cookie -i http://localhost:8764/api/apollo/collections/system_metrics

HTTP/1.1 200 OK
Content-Type: application/json; charset=utf-8
Content-Encoding: gzip
Vary: Accept-Encoding, User-Agent
Content-Length: 278
Server: Jetty(9.2.11.v20150529)

{
  "id" : "system_metrics",
  "createdAt" : "2016-03-04T23:29:47.779Z",
  "searchClusterId" : "default",
  "commitWithin" : 10000,
  "solrParams" : {
    "name" : "system_metrics",
    "numShards" : 1,
    "replicationFactor" : 1
  },
  "type" : "METRICS",
  "metadata" : { }
}
```

If the session cookie has expired, the system returns a 401 Unauthorized code:

```
> curl -b cookie -i http://localhost:8764/api/apollo/collections/system_metrics

HTTP/1.1 401 Unauthorized
Content-Type: application/json; charset=utf-8
Content-Length: 31
Server: Jetty(9.2.11.v20150529)

{"code":"session-idle-timeout"}
```

4.8.13. Configuring Fusion for SSL

Fusion uses the [Java Secure Socket Extension \(JSSE\)](#) framework to enable SSL configuration for secure communication between the Fusion UI and any HTTP client.

To configure Fusion for SSL you must install an SSL certificate and enable SSL in the Fusion UI.

Installing an SSL certificate

The server has a locally-protected private key that is accessible via a [JSEE keystore](#). The keystore maintains both the server certificate and the private key.

Important	In a production environment, SSL certificates must be signed by a trusted Certificate Authority (CA).
-----------	---

To store certificates, you can use the Java [keytool](#), which is part of the JDK. When you have a signed certificate, then you create a JSSE keystore by using the keytool "import" command.

In the following example, we have a signed certificate, in pfx format, in a PKCS#12 file called `fusion.keystore.p12`. The following command creates a new JSSE keystore (JKS):

```
keytool -importkeystore -srckeystore /opt/lucidworks/fusion/apps/jetty/ui/etc/fusion.keystore.p12 \
-srcstoretype pkcs12 -destkeystore /opt/lucidworks/fusion/jetty/ui/etc/keystore -deststoretype JKS
```

If you have the certificate and private key as separate files, then you need to use [openssl](#) to create a PKCS#12 file. For example:

```
openssl pkcs12 -export -out /home/admin/keys/keystore.pkcs12 -in /home/admin/keys/fullchain.pem -inkey
/home/admin/keys/privkey.pem
```

Note	When prompted for a password, do not enter a blank password.
------	---

Now use keytool to import the PKCS#12 file into Java keystore format and optionally delete the PKCS#12 file:

```
keytool -importkeystore -srckeystore /home/admin/keys/keystore.pkcs12 -srcstoretype PKCS12 -destkeystore
{fusion_path}/apps/jetty/ui/etc/keystore
```

Note	If your server certificate is signed by an intermediate CA rather than a root CA, you must add the intermediate certificate to the keystore before you add the server certificate.
------	--

Enabling SSL in the Fusion UI

1. Run the following:

```
cd {fusion_path}/apps/jetty/ui
java -jar {fusion_path}/apps/jetty/home/start.jar --add-to-start=https
```

Note

This step requires an Internet connection. If no connection is available on the Fusion host, run the commands above on a separate host, then copy the resulting `fusion/3.1.x/apps/jetty/ui/etc/keystore` file to the Fusion host.

2. Get the hashed version of your keystore password:

```
cd {fusion_path}
java -cp ./apps/libs/jetty-util-9.3.8.v20160314.jar org.eclipse.jetty.util.security.Password <secret>
```

Where `<secret>` is the password you used for the keystore or the `-storepass` value if you generated the self-signed certificate as per the guide provided here.

3. Edit `fusion/3.1.x/apps/jetty/ui/start.ini` to uncomment and add the hashed password from the previous step to these properties:
 - `jetty.sslContext.keyStorePassword`
 - `jetty.sslContext.keyManagerPassword`
 - `jetty.sslContext.trustStorePassword`

For example:

```
## Keystore password
jetty.sslContext.keyStorePassword=OBF:2uha1vgt1jg01a4b1a4j1jda1vg11ugg

## Keystore type and provider
# jetty.sslContext.keyStoreType=JKS
# jetty.sslContext.keyStoreProvider=

## KeyManager password
jetty.sslContext.keyManagerPassword=OBF:2uha1vgt1jg01a4b1a4j1jda1vg11ugg

## Truststore password
jetty.sslContext.trustStorePassword=OBF:2uha1vgt1jg01a4b1a4j1jda1vg11ugg
```

4. Set the local SSL port by editing the `jetty.ssl.port` property.

For example:

```
## Connector port to listen on
jetty.ssl.port=8864
```

5. Configure ZooKeeper to use HTTPS:

- a. Start ZooKeeper if required.

- b. Run `fusion/3.1.x/apps/solr-dist/server/scripts/cloud-scripts/zkcli.sh` `-zkhost`
`<zkHost1>:<port1>,<zkHost2>:<port2>,...` `-cmd put /clusterprops.json '{"urlScheme":"https"}'`

6. Restart Fusion with `fusion/3.1.x/bin/fusion restart`.

HTTPS should now be enabled on port 8864.

Optional Steps to set up your firewall / load balancer:

- Disallow all requests from 8764 from the outside world. Only localhost should be able to talk to Fusion on the non-SSL port 8764. Block all others.
- If you are using a load balancer or web server in front of Fusion, use it to redirect all HTTP requests to use HTTPS instead.

Disabling HTTP

Note	Only do this if blocking access to HTTP using the firewall is not feasible.
------	---

To entirely disable HTTP, remove the HTTP connector from the Jetty startup configuration:

Note	You can only use this option if your SSL certificate covers a hostname that can be accessed from the local host. For example, if your SSL certificate only covers https://myfusion.com then your local machine must be able to access Fusion using this URL.
------	---

1. Edit `fusion/3.1.x/apps/jetty/ui/start.d/http.ini` to change `--module=http` to `#--module=http`.
2. Edit `fusion/3.1.x/conf/fusion.properties`:
 - a. Make sure the Agent JVM uses the UI's keystore by adding the following to the end of the file:

```
agent.jvmOptions=-Djavax.net.ssl.trustStore={fusion_path}/apps/jetty/ui/etc/keystore
-Djavax.net.ssl.trustStorePassword=<password>
-Djavax.net.ssl.keyStore={fusion_path}/apps/jetty/ui/etc/keystore
-Djavax.net.ssl.keyStorePassword=<password>
```

- b. Uncomment `default.address` and change it to the hostname of the server that is validated by your SSL certificate.

Note	If the hostname saved in <code>default.address</code> is not validated by your ssl certificate, UI will not start because the Agent's liveness detector will not be able to access HTTPS port to see if Fusion is running.
------	--

Note	If you self-signed the certificate, the <code>default.address</code> must match the hostname you specified while signing the certificate. For example if my SSL certificate's validated hostname is <code>localhost</code> , change <code>#default.address = 127.0.0.1</code> to <code>default.address = localhost</code>
------	---

- c. Change `ui.port` to the SSL port you chose for `jetty.ssl.port` in UI's `start.ini` file earlier.
- d. Set `ui.ssl` to true by changing `# ui.ssl=false` to `ui.ssl=true`

3. Restart Fusion.

Configuring Fusion for SSL Solr/SolrCloud

To configure the Fusion HTTP client for a SSL-ed Solr or SolrCloud server, you must specify the `javax.net.ssl` system properties.

See the Solr wiki [instructions for enabling SSL for SolrCloud](#).

Step 1. Edit the configuration file

In the `fusion/3.1.x/conf/fusion.properties` file, add the following properties to the options `api.jvmOptions`, `connectors.jvmOptions`, and `ui.jvmOptions`:

```
-Djavax.net.ssl.keyStore=/path/to/solr-ssl.keystore.jks \  
-Djavax.net.ssl.keyStorePassword=secret \  
-Djavax.net.ssl.trustStore=/path/to/solr-ssl.keystore.jks \  
-Djavax.net.ssl.trustStorePassword=secret
```

Step 2. Register SolrCloud as a search cluster in Fusion

Send a request to the REST API 'searchCluster' endpoint:

```
curl -u admin:pass -H 'Content-type: application/json' -X POST \  
'http://localhost:8764/api/apollo/searchCluster' -d '  
{  
  "id" : "ssl",  
  "connectString" : "localhost:2181",  
  "cloud" : true  
}
```

Step 3. Test: create collection, index data, query collection

Create collection in SolrCloud with configured SSL:

```
curl -u admin:pass -H 'Content-type: application/json' -X POST 'http://localhost:8764/api/apollo/collections' \  
-d '  
{  
  "id" : "mycollection",  
  "searchClusterId" : "ssl"  
}'
```

Index data using an existing pipeline:

```
curl -u admin:pass 'http://localhost:8764/api/apollo/index-pipelines/mycollection- \  
default/collections/mycollection/index' -XPOST -H "Content-type: application/json" -d '{  
  "id": "1",  
  "foo_s": "bar",  
  "spam_s": 42  
}'
```

Query the collection using the default query pipeline:

```
curl -u admin:pass 'http://localhost:8764/api/apollo/query-pipelines/mycollection-  
default/collections/mycollection/select?q=*:*'
```

Generating a self-signed certificate

If don't have signed certificates from an external CA, then you can generate a set of self-signed certificates using the Java [keytool](#) utility to generate a public/private key pair and generate a self-signed certificate.

The keytool option "-genkeypair" generates a public/private key pair. It wraps the public key into an X.509 v3 self-signed certificate, which is stored as a single-element certificate chain. This certificate chain and the private key are stored in a new keystore entry identified by alias. The full set of arguments to this command are:

```
-genkeypair  
  {-alias alias}  
  {-keyalg keyalg}  
  {-keysize keysize}  
  {-sigalg sigalg}  
  [-dname dname]  
  [-keypass keypass]  
  {-startdate value}  
  {-ext ext}*  
  {-validity valDays}  
  {-storetype storetype}  
  {-keystore keystore}  
  [-storepass storepass]  
  {-providerClass provider_class_name {-providerArg provider_arg}}  
  {-v}  
  {-protected}  
  {-Jjavaoption}
```

Arguments of interest are:

- `keyalg` specifies the algorithm to be used to generate the key pair. This must be RSA in order to talk to browser clients IE and Navigator.
- `keysize` specifies the size of each key to be generated. Depends on `keyalg`. For RSA, this should be 2048.
- `dname` specifies the X.500 Distinguished Name to be associated with the alias, and is used as the issuer and subject fields in the self-signed certificate. If no distinguished name is provided at the command line, the user will be prompted for one. An X.500 Distinguished name is a set of named fields, of these, the field named "CN", ("Common Name"), is the internet-facing fully qualified domain name of the server.
- `keypass` is a password used to protect the private key of the generated key pair. If no password is provided, the user is prompted for it. If you press RETURN at the prompt, the key password is set to the same password as that used for the keystore. `keypass` must be at least 6 characters long.
- `ext` is used to embed extensions into the certificate generated. For Fusion, the server certificate should include the "SAN" or [SubjectAlternativeName](#) extension which allows alternative URIs and IP addresses to be associated with this certificate.

Keystore files created by the keystore tool are in the JKS format, which is a proprietary file format capable of storing multiple key-pairs, certificates, and symmetric encryption keys, and with all entries indexed by the keypair alias.

To generate a self-signed certificate for the Fusion UI running as "localhost", we use the following arguments:


```
keytool -genkeypair \  
-alias localhost -keyalg RSA -keysize 2048 \  
-keypass secret -storepass secret \  
-validity 365 -keystore my.keystore.jks \  
-ext SAN=DNS:localhost,IP:127.0.0.1 \  
-dname "CN=localhost, OU=org unit, O=org, L=loc, ST=st, C=country"
```

Note

The value for **CN** must match the hostname you will be using to access Fusion. For example, if will use the URL <https://myfusion.com:8864> to access the UI then the CN should have the value **myfusion.com**.

This keystore file can now be imported to the Fusion UI keystore. To check the generated keystore we use the openssl tool to pretty-print the signed certificate in the file **my.keystore.jks**. This is not strictly necessary; this should always be done before sending the certificate to a CA to get it signed properly in order to verify that the certificate information is complete and correct.

To get from the keystore JKS format to a human-readable printout, it must be converted to the text-based "PEM" format, which is ASCII (Base64) armored data prefixed with a **— BEGIN …** line. This requires three steps.

First, we use the keytool to convert the proprietary JKS format to the **PKCS #12** format:

```
keytool -importkeystore \  
-srckeystore my.keystore.jks -destkeystore my.keystore.p12 \  
-srcstoretype jks -deststoretype pkcs12
```

This command prompts for passwords - as before, the password is "secret".

Next, we use openssl to convert the PKCS format to PEM format:

```
openssl pkcs12 -in my.keystore.p12 -out my.keystore.pem
```

Finally, to pretty-print the certificate, we use the following openssl command:

```
openssl x509 -in my.keystore.pem -text -noout
```

This converts the PEM format to text format, and writes the output to the terminal. The output is:

Certificate:

Data:

Version: 3 (0x2)
Serial Number: 1442779707 (0x55ff123b)
Signature Algorithm: sha256WithRSAEncryption
Issuer: C=country, ST=st, L=loc, O=org, OU=org unit, CN=localhost
Validity
Not Before: Apr 8 07:39:35 2015 GMT
Not After : Apr 7 07:39:35 2016 GMT
Subject: C=country, ST=st, L=loc, O=org, OU=org unit, CN=localhost
Subject Public Key Info:
Public Key Algorithm: rsaEncryption
RSA Public Key: (2048 bit)

Modulus (2048 bit):

00:96:04:46:6c:be:7f:ec:ea:18:fa:28:11:a9:fb:
3d:07:c5:3c:49:39:57:11:24:1d:75:47:5d:76:26:
4b:73:c0:ea:44:7a:a5:59:3a:a7:4b:16:eb:1f:be:
05:f1:2a:be:62:72:2c:67:ec:d3:8b:ad:76:af:dc:
6d:14:ca:c9:75:5a:76:24:80:c6:f8:55:b0:27:6a:
fa:a8:1b:4b:5c:55:93:49:ff:f8:84:67:29:56:80:
ca:c1:d8:c4:8c:b8:57:a4:78:6a:e8:e0:2c:51:74:
fb:fb:52:3d:d3:e9:58:e5:11:79:5f:5a:70:ec:c3:
de:d7:56:36:67:fd:52:dd:73:60:f7:93:9f:00:c0:
36:49:65:7d:77:45:76:34:0e:81:38:96:2b:19:b0:
30:8b:ac:2f:ae:dd:92:41:78:da:47:32:02:f7:0d:
04:f5:51:85:dc:06:58:08:9b:d2:a0:69:52:ac:b2:
7d:c7:bd:16:1d:9e:af:e2:2b:6a:61:8e:cb:a9:ec:
fc:01:fe:6b:34:49:1c:d8:75:8b:ca:ec:ea:fd:93:
0a:8b:34:6b:77:98:ec:83:6f:d2:bc:81:ec:f5:18:
48:41:db:92:da:ef:19:19:27:5b:05:5f:8c:6e:1e:
9d:e5:90:42:6f:36:8f:11:49:05:aa:dd:a5:c9:0a:
fe:81

Exponent: 65537 (0x10001)

X509v3 extensions:

X509v3 Subject Alternative Name:

DNS:localhost, IP Address:127.0.0.1

X509v3 Subject Key Identifier:

E1:D1:66:F6:1F:5A:18:1A:82:33:A7:32:0E:0B:EA:8E:C0:D3:ED:05

Signature Algorithm: sha256WithRSAEncryption

56:a8:31:c0:94:8f:5a:80:27:17:36:ba:e0:ff:e6:59:13:9e:
67:6b:1e:7d:67:04:5c:e1:88:9e:d9:89:11:ca:88:a0:21:4c:
90:a8:8b:9d:b5:ba:0f:92:65:da:c5:a9:91:82:17:46:55:43:
82:78:92:b1:f9:f5:2b:65:5e:b0:93:03:3b:94:83:bb:fe:9b:
09:a9:f3:82:d2:4d:b5:72:e9:ee:75:15:31:3d:18:a7:c1:e8:
45:44:1d:40:d2:eb:96:b7:01:41:dd:9d:1c:31:e6:45:4a:c2:
3d:ec:22:1a:35:ec:38:62:e8:3d:b3:30:10:d3:88:09:5a:87:
54:87:fb:92:d6:0e:74:52:3c:f2:c3:f5:70:61:ea:72:3f:cd:
65:72:34:6f:51:94:13:e0:7a:73:bb:57:c8:ad:98:f7:3f:43:
4d:75:96:db:cf:2e:e6:82:1c:c2:97:38:2d:37:06:c4:27:db:
87:82:6b:c6:01:71:f7:e9:1f:69:62:0d:cc:54:e9:f4:25:86:
6a:e2:38:72:c3:9c:53:b4:6e:4f:ae:8a:09:36:14:f0:10:57:
9c:c9:a8:a3:a5:e6:db:d1:d4:39:95:f3:54:95:4f:2f:db:59:
b6:bd:77:00:77:c2:9d:4f:d9:04:d5:af:33:bb:2e:0f:65:a9:
74:ff:66:f2

References and tutorials

[Transport Layer Security](#) (wikipedia.org)

[Public Key Certificate](#)(wikipedia.org)

[OpenSSL Cookbook](#) (free ebook)

[OpenSSL Command Line Utilities](#) (openssl.org wiki)

[Java Tutorials: Generating and Verifying Certificates](#)

[IBM developerWorks: What is the JSSE all about?](#)

4.8.14. Web Authentication Cookie FAQs

Does the application use web authentication cookies?

Yes, we use a session cookie for maintaining authenticated user identification.

Is the cookie used for non-authentication purposes?

No, the cookie contains exactly one value, the session ID.

Is the cookie set to the narrowest/lowest path or domain needed in order to prevent inadvertent or unauthorized sharing of cookies by other web applications?

Yes, the path is set to /api only (the narrowest path). As recommended by [OWASP](#), we do not directly set the domain attribute, so the default ends up being the origin server.

Are the cookies non-persistent?

The cookies are session based, and not persisted beyond logout or via timeout.

Is the value of the cookie not predictable and does it provide 64-bit entropy?

The cookie value is a Java UUID, which uses the SecureRandom class. A little research leads me to believe it's a 128 bit value, with 122 bit randomness.

Are default values not used for the name of the cookie?

As recommended by OWASP, the cookie name is vague/meaningless. It is simply, "id".

Is the cookie set via SSL channel and are the 'secure' and 'HTTPOnly' attributes set?

If the web server is running under SSL, then the cookie is set to secure and HttpOnly is set to true.

Can the cookie be manually deleted through a logout button that sets the cookie value to null or the cookie value is rendered invalid on the server after a period of inactivity?

Yes, the cookie can be manually deleted. There is also a timeout mechanism:

- 8 hour absolute lifetime - it never lasts longer than 8 hours
- 1 hour soft lifetime - if the last request time was > than 1 hour, the session is destroyed. Otherwise, the lifetime is bumped up 1 more hour, until the maximum 8 hour limit is met.

Is the cookie cleared during the authentication of a user?

Yes.

Chapter 5. Spark and Machine Learning

[Apache Spark](#) is an open-source cluster-computing framework. Spark improves on previous previous MapReduce implementations by use of Resilient Distributed Datasets (RDDs), a distributed memory abstraction that lets programmers perform in-memory computations on large clusters in a fault-tolerant manner. To schedule and run jobs on the nodes in the cluster, Spark uses [Akka](#) which is a toolkit and runtime for building highly concurrent, distributed, and resilient message-driven applications on the JVM.

Fusion manages a Spark cluster which is used for all signal aggregation processes. As of Fusion 2.4, this Spark cluster can also be used to train and compile machine learning models as well as to run experiment-management tasks via the Spark Jobs API.

The topics covered in the section are:

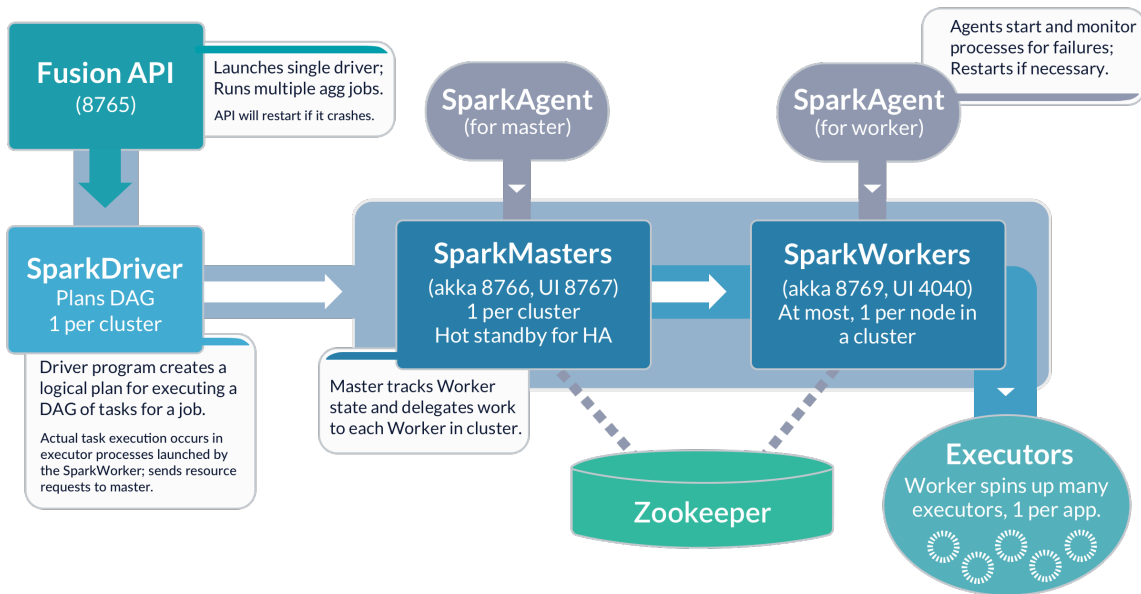
- Spark Concepts and Terminology - Spark integration in Fusion - basic terminology and schematics.
- Spark Getting Started - Starting Spark processes and working with the shell and the Spark UI.
- Spark Drivers - Running different kinds of Spark applications.
- Spark Configuration - Configuration settings for Spark applications.
- Scaling Spark Aggregations - Working with very large data sets.
- Troubleshooting - Common mistakes and how to fix them.
- Machine Learning Models in Fusion - Example of using a Spark cluster to train a sentiment classifier for tweets.

5.1. Further Reading

- [Machine Learning in Lucidworks Fusion](#)
- [Apache Spark Key Terms, Explained](#)
- [Apache Spark on Wikipedia](#)
- [Akka \(toolkit\) on Wikipedia](#)
- [Reactive App presentation](#)

5.2. Spark Concepts and Terminology

The following schematic shows the Spark components available from Fusion:



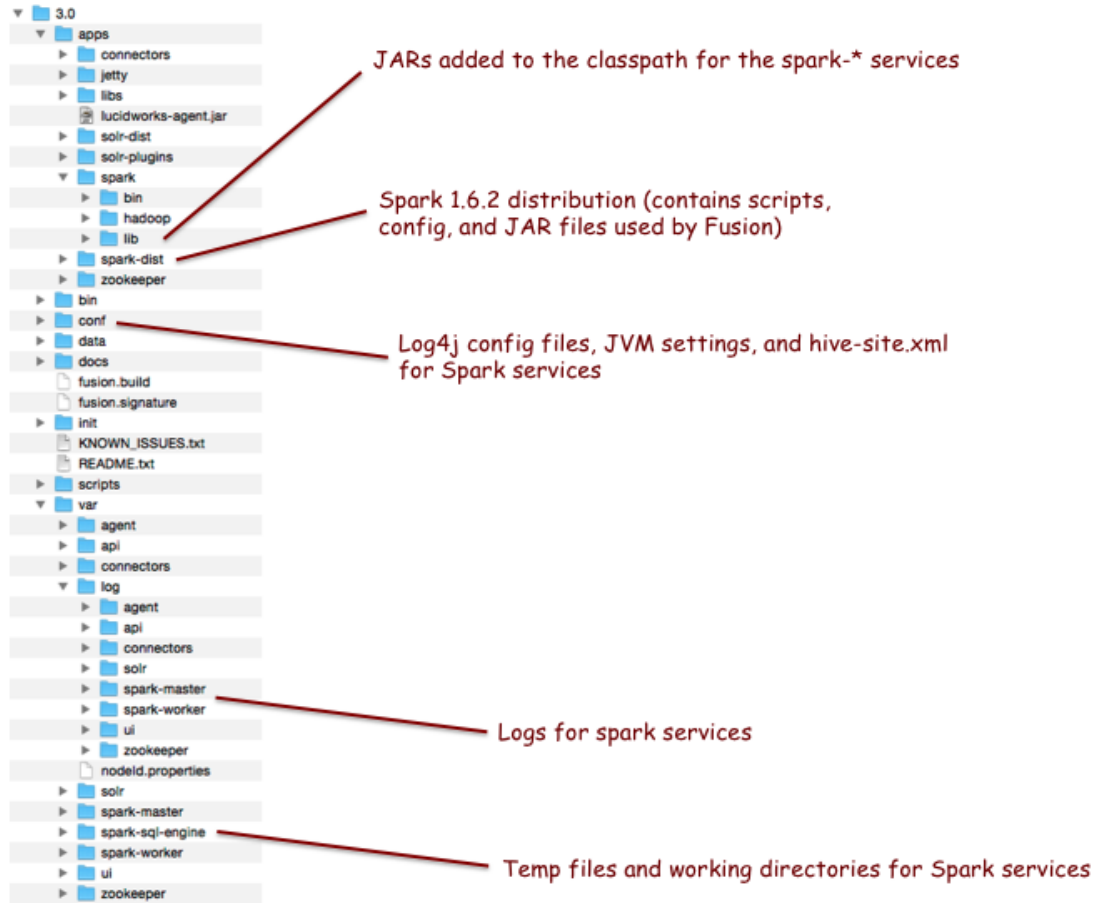
5.2.1. Spark Fusion Terminology

- **application**: an active SparkContext in Spark Master UI, which consists of a classpath and configuration. Jobs submitted to the cluster always run as classes in a specific application, ie. using the application's classpath and configuration.
- **SparkDriver**: JVM process launched by the Fusion API service to execute Fusion jobs in Spark. SparkDriver creates and manages SparkContext for the Fusion application, and stops SparkContext when it's no longer needed.
- **spark-master**: Agent-managed Fusion service that coordinates worker processes and applications in a Spark cluster. You should run at least 2 spark-master processes per cluster to achieve high-availability. ZooKeeper determines which spark-master process is the leader and handles fail-over.
- **spark-worker**: Agent-managed Fusion service that launches executors for Spark applications. Spark-workers communicate with the master to launch executors for an application.
- **spark-sql-engine**: Agent-managed Fusion service that runs Spark's thrift-based SQL engine. Provides JDBC access to a Spark cluster.
- **spark-shell**: Wrapper script provided with Fusion to launch the Spark Scala REPL shell with the correct master URL (pulled from Fusion's API) and shaded Fusion JAR added.
- **CoarseGrainedExecutorBackend**: Executor process(es) launched by a spark-worker to execute the tasks for a specific application, such as the spark-shell.
- **dynamic allocation**: Spark configuration setting that allows the Spark master to reclaim allocated resources (cpu & memory) from an "application" if it is not using the resources.
- **shaded JAR**: The Fusion API service creates an uber-jar containing all of the dependencies needed to use spark-solr and Fusion classes within a spark job. The classes that conflict with classes on spark's classpath are shaded to ensure that Fusion classes use the correct version.
- **akka**: Akka is a toolkit and runtime for building highly concurrent, distributed, and resilient message-driven applications on the JVM. Akka uses the Actor model to hide all the thread-related code and provides simple

interfaces which allow you to more easily implement a scalable and fault-tolerant system. Spark is built on top of Akka.

5.2.2. Spark Fusion Directories

In Fusion 3.0, the following directories are used for Spark logs and components:



5.3. Spark Getting Started

The public GitHub repo <https://github.com/lucidworks/fusion-spark-bootcamp> contains examples and labs for learning how to use Fusion's Spark features.

In this section, you'll walk through some basic concepts of using Spark in Fusion. For further exposure, you should work through the labs in the Fusion Spark Bootcamp.

5.3.1. Starting Spark Master and Worker Processes

In Fusion 3.0, the Fusion run script `fusion/3.1.x/bin/fusion` with the argument `start` doesn't start the spark-master and spark-worker processes. This reduces the number of Java processes needed to run Fusion and therefore reduces memory and CPU consumption.

Jobs that depend on Spark, e.g. aggregations, will still execute in what Spark calls "local" mode. When in local mode, Spark executes tasks in-process in the driver application JVM. Local mode is intended for jobs that consume/produce small datasets. One caveat of using local mode is that a persistent Spark UI is not available, but you can access the driver/job application UI at port `:4040` while the local SparkContext is running.

Tip	To allow the spark-master and spark-worker processes to start and stop via <code>bin/fusion start</code> and <code>bin/fusion stop</code> , we recommend adding them to the <code>group.default</code> definition in <code>conf/fusion.properties</code> : <pre>group.default = zookeeper, solr, api, connectors, ui, spark-master, spark-worker</pre>
-----	---

To scale-out Spark in Fusion to support larger data sets and to speed up processing, you should start the spark-master and spark-worker processes:

```
$ cd /path/to/fusion/3.1.x
$ bin/spark-master start
$ bin/spark-worker start
```

After starting the master and worker services, direct your browser to <http://localhost:8767> to view the Spark UI. The display should be similar to:



Spark 1.6.2 **Spark Master at spark://192.168.1.9:8766**

URL: spark://192.168.1.9:8766
REST URL: spark://192.168.1.9:6066 (cluster mode)
Alive Workers: 1
Cores in use: 8 Total, 1 Used
Memory in use: 2.0 GB Total, 1024.0 MB Used
Applications: 2 Running, 17 Completed
Drivers: 0 Running, 0 Completed
Status: ALIVE

Workers

Worker Id	Address	State	Cores	Memory
worker-20161005115307-192.168.1.9-8769	192.168.1.9:8769	ALIVE	8 (1 Used)	2.0 GB (1024.0 MB Used)

Check the following logs if you do not see the master UI or at least one worker in the ALIVE state.


```
fusion/3.1.x/var/log/spark-master/spark-master.log
fusion/3.1.x/var/log/spark-worker/spark-worker.log
```

Use the Fusion API to get the status of the Spark master via the following request:

```
curl http://localhost:8765/api/v1/spark/master/status
```

This should return a response of the form:

```
[ {
  "url" : "spark://192.168.1.9:8766",
  "status" : "ALIVE",
  "workers" : [ {
    "id" : "worker-20161005175058-192.168.1.9-8769",
    "host" : "192.168.1.9",
    "port" : 8769,
    "webuiaddress" : "http://192.168.1.9:8082",
    "cores" : 8,
    "coresused" : 0,
    "coresfree" : 8,
    "memoryused" : 0,
    "memoryfree" : 2048,
    "state" : "ALIVE",
    "lastheartbeat" : 1475711489460
  } ], ...
```

If you have multiple Spark masters running in a Fusion cluster, each will be shown in the status but only one will be ALIVE; the other masters will be in STANDBY mode.

Tip

If you are operating a multi-node Spark cluster, we recommend running multiple Spark master processes to achieve high-availability if the active one fails, the standby will take over.

5.3.2. Running a Job in the Spark Shell

Once you have started the Spark master and worker, launch the Fusion Spark shell wrapper:

```
$ bin/spark-shell
```

It may take a few minutes to load the first time as the script needs to download the shaded Fusion JAR file from the API service. When the shell is initialized, you'll see the prompt:

```
scala>
```

Type `:paste` to activate paste mode in the shell and paste in the following Scala code:

```

val readFromSolrOpts = Map(
  "collection" -> "logs",
  "fields" -> "host_s,port_s,level_s,message_t,thread_s,timestamp_tdt"
)
val logsDF = spark.read.format("solr").options(readFromSolrOpts).load
logsDF.registerTempTable("fusion_logs")
var sqlDF = spark.sql("""
| SELECT COUNT(*) as num_values, level_s as level
| FROM fusion_logs
| GROUP BY level_s
| ORDER BY num_values desc
| LIMIT 10""").stripMargin)
sqlDF.show(10,false)

```

Note

In Fusion 2.4.x, you'll need to specify the Solr zkhost in the readFromSolrOpts Map, but in 3.x we set it automatically using Fusion configuration API.

Make sure the zkhost val is correct for your environment!

Type Ctrl+D to execute the script. Your results should look similar to the following:

```

scala> :paste
// Entering paste mode (ctrl-D to finish)

val readFromSolrOpts = Map(
  "collection" -> "logs",
  "fields" -> "host_s,port_s,level_s,message_t,thread_s,timestamp_tdt"
)
val logsDF = spark.read.format("solr").options(readFromSolrOpts).load
logsDF.registerTempTable("fusion_logs")
var sqlDF = spark.sql("""
| SELECT COUNT(*) as num_values, level_s as level
| FROM fusion_logs
| GROUP BY level_s
| ORDER BY num_values desc
| LIMIT 10""").stripMargin)
sqlDF.show(10,false)

// Exiting paste mode, now interpreting.

warning: there was one deprecation warning; re-run with -deprecation for details
+-----+-----+
|num_values|level|
+-----+-----+
|2260      |INFO |
|604       |WARN |
+-----+-----+

readFromSolrOpts: scala.collection.immutable.Map[String,String] = Map(collection -> logs, fields -> host_s,port_s,level_s,message_t,thread_s,timestamp_tdt)
logsDF: org.apache.spark.sql.DataFrame = [host_s: string, port_s: string ... 4 more fields]
sqlDF: org.apache.spark.sql.DataFrame = [num_values: bigint, level: string]

```

Congratulations, you just ran your first Fusion Spark job that reads data from Solr and performs a simple aggregation!

5.3.3. The Spark Shell UI

The Spark Master UI allows us to dig into the details of the Spark job. This handy guide helps you understand the Spark UI: <https://jaceklaskowski.gitbooks.io/mastering-apache-spark/content/spark-webui.html>.

In your browser (<http://localhost:8767>), there should be a job named "Spark shell" under running applications (the application ID will be different than the following screenshot):

Running Applications

Application ID	Name	Cores	Memory per Node	Submitted Time	User	State	Duration
app-20161005155449-0000	(kill) Spark shell	4	500.0 MB	2016/10/05 15:54:49	timpotter	RUNNING	4.0 min

Click on the application ID and then on the **Application Detail UI** link:

Spark Jobs (?)

Total Uptime: 8.3 min
 Scheduling Mode: FIFO
 Completed Jobs: 1

▶ Event Timeline

Completed Jobs (1)

Job Id	Description	Submitted	Duration	Stages: Succeeded/Total	Tasks (for all stages): Succeeded/Total
0	show at <console>-44	2016/10/05 15:58:33	7 s	2/2	202/202

Notice the tabs at the top that allow you to dig into details about the running application. Take a moment to explore the UI. Can answer the following questions about your application:

- How many tasks were needed to execute this job?
- Which JARs were added to the classpath for this job? Hint: look under the Environment tab.
- How many executor processes were used to run this job? Why? Hint: take a look at all the Spark configuration properties under the Environment tab.
- How many rows were read from Solr for this job? Hint: look under the SQL tab

For the above run, the answers are:

- 202 tasks were needed to execute this job.
- The Environment tab shows that one of the JAR files is named "fusion.jar" and was "Added By User". The fusion.jar file is a copy of the shaded JAR created by the API service:

Classpath Entries

Resource	Source
/Users/timpotter/dev/lw/sstk-local/fusion-dev/3.0-SNAPSHOT/apps/spark-dist/conf/	System Classpath
/Users/timpotter/dev/lw/sstk-local/fusion-dev/3.0-SNAPSHOT/apps/spark-dist/lib/datanucleus-api-jdo-3.2.6.jar	System Classpath
/Users/timpotter/dev/lw/sstk-local/fusion-dev/3.0-SNAPSHOT/apps/spark-dist/lib/datanucleus-core-3.2.10.jar	System Classpath
/Users/timpotter/dev/lw/sstk-local/fusion-dev/3.0-SNAPSHOT/apps/spark-dist/lib/datanucleus-rdbms-3.2.9.jar	System Classpath
/Users/timpotter/dev/lw/sstk-local/fusion-dev/3.0-SNAPSHOT/apps/spark-dist/lib/spark-assembly-1.6.2-hadoop2.6.0.jar	System Classpath
http://192.168.1.9:64835/jars/fusion.jar	Added By User

- It took 2 executor processes to run this job. Each executor has 2 CPUs allocated to it and the `bin/spark-shell` script asked for 4 total CPUs for the shell application.
- This particular job read about 21K rows from Solr, but this number will differ based on how long Fusion has been running.

The key take-away is that you can see how Spark interacts with Solr using the UI.

5.3.4. Spark Job Tuning


Returning to the first question, why were 202 tasks were needed to execute this job?

Details for Query 2

Submitted Time: 2016/10/05 16:17:41

Duration: 0.2 s

Succeeded Jobs: 2



The reason is that SparkSQL defaults to using 200 partitions when performing distributed group by operations, see property: `spark.sql.shuffle.partitions`. Let's adjust that so Spark only uses 4 tasks since our data set is so small. In the Spark shell, execute the following Scala:

```
sqlContext.setConf("spark.sql.shuffle.partitions", "4")
```


You just need to re-execute the final query and show commands:

```
val readFromSolrOpts = Map(
  "collection" -> "logs",
  "fields" -> "host_s,port_s,level_s,message_t,thread_s,timestamp_tdt"
)
val logsDF = spark.read.format("solr").options(readFromSolrOpts).load
logsDF.registerTempTable("fusion_logs")
var sqlDF = spark.sql("""
| SELECT COUNT(*) as num_values, level_s as level
| FROM fusion_logs
| GROUP BY level_s
| ORDER BY num_values desc
| LIMIT 10""").stripMargin
sqlDF.show(10,false)
```

Now if you look at the Job UI, you'll see a new job that executed with only 6 executors! You've just had your first experience with tuning Spark jobs.

5.4. Spark Driver Processes

A Spark "driver" is an application that creates a SparkContext for executing one or more jobs in the Spark cluster. The following diagram depicts the driver's role in a Spark cluster:



In the diagram above, the spark-master service in Fusion is the Cluster Manager.

If your Spark job performs any collect operations, then the result of the collect (or collectAsMap) is sent back to the driver from all the executors. Consequently, if the result of the collect is too big to fit into memory, you will encounter OOM issues (or other memory related problems) when running your job.

All Fusion jobs run on Spark using a driver process started by the API service. There are three types of drivers in Fusion:

- Default driver: Executes built-in Fusion jobs, such as a signal aggregation job or a metrics rollup job.
- Scripted job driver: Executes custom script jobs; a separate driver is needed to isolate the classpath for custom Scala scripts.
- Spark-shell driver: Launch using `fusion/3.1.x/bin/spark-shell`.

5.4.1. Default Driver

Navigate to the Fusion UI and select the system metrics collection in the UI. Select one of the built-in Aggregation jobs, such as hourlyMetricsRollup-counters. In the diagram above, the spark-master service in Fusion is the Cluster Manager.

You must delete any existing driver applications before launching the job. Even if you haven't started any jobs by hand, Fusion's API service may have started one automatically, because Fusion ships with built-in jobs that run in the background which perform rollups of metrics in the system_metrics collection. Therefore, before you try to launch a job, you should run the following command:

```
> curl -XDELETE http://localhost:8765/api/v1/spark/driver
```

Wait a few seconds and use the Spark UI to verify that no Fusion-spark application (e.g., "Fusion-20161005224611") is

running.

In a terminal window or windows, set up a `tail -f` job on the api and spark-driver-default logs

```
$ cd /path/to/fusion/3.1.x
$ tail -f var/log/api/api.log var/log/api/spark-driver-default.log
```

Now, start any aggregation job from the UI. It doesn't matter whether or not this job performs any work; the goal of this exercise is to show what happens in Fusion and Spark when you run an aggregation job. You should see activity in both logs related to starting the driver application and running the selected job. The Spark UI will now show a Fusion-spark app:

Running Applications

Application ID	Name	Cores	Memory per Node	Submitted Time	User	State	Duration
app-20161005164611-0003	(kill) Fusion-20161005224611	8	2.0 GB	2016/10/05 16:46:11	timpotter	RUNNING	32 s

Use the `ps` command to get more details on this process:

```
$ ps waux | grep SparkDriver
```

The output should show that the Fusion SparkDriver is a JVM process started by the API service; it is not managed by the Fusion agent. Within a few minutes, the Spark UI will update itself:

```
URL: spark://192.168.1.9:8766
REST URL: spark://192.168.1.9:6066 (cluster mode)
Alive Workers: 1
Cores in use: 8 Total, 0 Used
Memory in use: 2.0 GB Total, 0.0 B Used
Applications: 1 Running, 4 Completed
Drivers: 0 Running, 0 Completed
Status: ALIVE
```

Workers

Worker Id	Address	State	Cores
worker-20161005115307-192.168.1.9-8769	192.168.1.9:8769	ALIVE	8 (0 Used)

Running Applications

Application ID	Name	Cores	Memory per Node	Submitted Time
app-20161005165054-0004	(kill) Fusion-20161005225054	0	2.0 GB	2016/10/05 16:50:54

Notice that the application no longer has any cores allocated and that all of the memory available is not being used (0.0B Used of 2.0 GB Total). This is because we launch our driver applications with `spark.dynamicAllocation.enabled=true`. This setting allows the Spark master to reclaim CPU & memory from an application if it is not actively using the resources allocated to it.

Both driver processes (default and scripted) manage a SparkContext. For the default driver, the SparkContext will be shut down after waiting a configurable (`fusion.spark.idleTime`: default 5 mins) idle time. The scripted driver shuts down the SparkContext after every scripted job is run to avoid classpath pollution between jobs.

5.4.2. Scripted Driver

Fusion supports custom script jobs, see: https://doc.lucidworks.com/fusion/2.4/REST_API_Reference/Spark-Jobs-API.html#job-types

Script jobs require a separate driver to isolate the classpath for custom Scala scripts, as well as to isolate the classpath between the jobs, so that classes compiled from scripts don't pollute the classpath for subsequent scripted jobs. For this reason the SparkContext that each scripted job uses is immediately shut down after the job is finished and recreated for new jobs. This adds some startup overhead for scripted jobs. Refer to the [apachelogs lab](#) in the [Fusion Spark Bootcamp project](#) for a complete example of a custom script job.

To troubleshoot problems with a script job, start by looking for errors in the script driver log `fusion/3.1.x/var/log/api/spark-driver-scripted.log`.

5.4.3. Spark Drivers in a Multi-node Cluster

To find out which node is running the Spark driver which node is running the driver when running a multi-node Fusion deployment which has several nodes running Fusion's API services, you can query the driver status via the following call to the Spark Jobs API endpoint:

```
$ curl http://localhost:8765/api/v1/spark/driver/status
```

This returns a status report:

```
{
  "/spark-drivers/15797426d56T537184c2" : {
    "id" : "15797426d56T537184c2",
    "hostname" : "192.168.1.9",
    "port" : 8601,
    "scripted" : false
  }
}
```

Note

This endpoint only exists in 3.x and 2.4.3 and higher.

5.5. Spark Configuration

Spark has a number of configuration properties. In this section, we'll cover some of the key settings you'll need to use Fusion's Spark integration.

For the full set of Fusion's spark-related configuration properties, see the Spark Jobs API.

5.5.1. Spark Master / Worker Resource Allocation

Note	If you co-locate Spark workers and Solr nodes on the same server, then be sure to reserve some CPU for Solr to avoid a compute intensive Spark job from starving Solr of CPU resources.
------	---

Number of Cores Allocated

When a worker process joins the Spark cluster, it looks in the Fusion configuration (stored in Fusion's ZooKeeper) for the property `spark.worker.cores` (Fusion 2.4) or `fusion.spark.worker.cores` (Fusion 3.x). If this setting is left unspecified, then the worker will use all available cores.

For example, in the screenshot below, we have a 3-node cluster, where each worker uses all available cores (8) for a total of 24 cores:

The screenshot shows the Spark Master interface for a 3-node cluster. The status bar indicates 3 workers, 24 total cores (all used), and 6.0 GB of memory. A table below lists the three workers, each using 8 cores and 2.0 GB of memory. Annotations with red arrows point to the worker count and the 'Cores' column, explaining that each worker uses all 8 cores because the `spark.worker.cores` property is not configured.

Spark Master at spark://10.44.174.130:8766

URL: spark://10.44.174.130:8766
REST URL: spark://10.44.174.130:8066 (cluster mode)
Alive Workers: 3
Cores in use: 24 Total, 0 Used
Memory in use: 6.0 GB Total, 0.0 B Used
Applications: 0 Running, 28 Completed
Drivers: 0 Running, 0 Completed
Status: ALIVE

Workers

Worker Id	Address	State	Cores	Memory
worker-20161006162823-10.44.174.130-8769	10.44.174.130:8769	ALIVE	8 (0 Used)	2.0 GB (0.0 B Used)
worker-20161006162824-10.44.174.128-8769	10.44.174.128:8769	ALIVE	8 (0 Used)	2.0 GB (0.0 B Used)
worker-20161006162907-10.44.174.129-8769	10.44.174.129:8769	ALIVE	8 (0 Used)	2.0 GB (0.0 B Used)

To change the CPU usage per worker, you need to use the Fusion configuration API to update this setting, as in the following examples.

Fusion 2.4

```
$ curl -u <name>:<password> -H 'Content-type:application/json' -X PUT -d '6' \  
http://localhost:8764/api/apollo/configurations/spark.worker.cores
```

Fusion 3.0

```
$ curl -u <name>:<password> -H 'Content-type:application/json' -X PUT -d '6' \  
http://localhost:8764/api/apollo/configurations/fusion.spark.worker.cores
```

You can also over-allocate cores to a spark-worker, which usually is recommended for hyper-threaded cores by setting the property `spark-worker.envVars` to `SPARK_WORKER_CORES=<number of cores>` in the `fusion.properties` file on all nodes

hosting a spark-worker. For example, a r4.2xlarge instance in EC2 has 8 CPU cores, but the following configuration will improve utilization and performance:

```
spark-worker.envVars=SPARK_WORKER_CORES=16
```

After making this change to your spark worker nodes, you must restart the spark-worker process on each node:

```
$ cd /path/to/fusion/3.1.x
$ bin/spark-worker restart
```

Memory Allocation

The amount of memory allocated to each worker process is controlled by Fusion property `fusion.spark.worker.memory` which specifies the total amount of memory available for all executors spun up by that Spark Worker process. This is the quantity seen in the memory column against a worker entry in the Workers table.

Note	The JVM memory setting (-Xmx) for the spark-worker process configured in the <code>fusion.properties</code> file in Fusion 3.x and in the <code>spark-worker</code> script in Fusion 2.x controls how much memory the spark-worker needs to manage executors and not how much memory should be made available to your job(s). Typically, 512m to 1g is sufficient for the spark-worker JVM process.
------	---

Fusion 3.0:

```
$ curl -u <name>:<password> -H 'Content-type:application/json' -X PUT -d '8g' \
http://localhost:8764/api/apollo/configurations/fusion.spark.worker.memory
```

Fusion 2.4:

```
$ curl -u <name>:<password> -H 'Content-type:application/json' -X PUT -d '8g' \
http://localhost:8764/api/apollo/configurations/spark.worker.memory
```

The Spark worker process manages executors for multiple jobs running concurrently. For certain types of aggregation jobs you can also configure the per executor memory, but this may impact how many jobs you can run concurrently in your cluster. Unless explicitly specified using the parameter `spark.executor.memory`, Fusion calculates the amount of memory that can be allocated to the executor

In Fusion 3.x, the aggregation Spark jobs always get half the memory of what is assigned to the workers. This is controlled by `fusion.spark.executor.memory.fraction` property which is set to default value `0.5`. E.g., Spark workers in 3.x have 4 Gb of memory by default and the executors for aggregator Spark jobs are assigned 2 Gb for each executor. To let fusion aggregation jobs use more memory of the workers, increase `fusion.spark.executor.memory.fraction` property to `1`. This property should be used instead of the spark executor memory property.

Fusion 3.0:

```
$ curl -u <name>:<password> -H 'Content-type:application/json' -X PUT -d '1' \
http://localhost:8764/api/apollo/configurations/fusion.spark.executor.memory.fraction
```

In Fusion 2.4, you must set the per-executor memory directly via property `spark.executor.memory`:

```
$ curl -u <name>:<password> -H 'Content-type:application/json' -X PUT -d '8g' \
http://localhost:8764/api/apollo/configurations/spark.executor.memory
```

After making these changes and restarting the workers, when we run a Fusion job, we get the following (screenshots taken from running instance of Fusion 2.4):

```
curl -u ... -H 'Content-type:application/json' -XPUT -d '8g' "http://localhost:8764/api/apollo/configurations/spark.executor.memory"
curl -u ... -H 'Content-type:application/json' -XPUT -d '8g' "http://localhost:8764/api/apollo/configurations/spark.worker.memory"
curl -u ... -H 'Content-type:application/json' -XPUT -d '6' "http://localhost:8764/api/apollo/configurations/spark.worker.cores"
```

The screenshot shows the Spark Master web interface. At the top, it says "Spark Master at spark://10.44.174.130:8766". Below this, it lists various metrics: URL, REST URL, Alive Workers (3), Cores in use (18 Total, 18 Used), Memory in use (24.0 GB Total, 24.0 GB Used), Applications (1 Running, 31 Completed), Drivers (0 Running, 0 Completed), and Status (ALIVE). There are two red arrows pointing from text annotations to the "Cores in use" and "Memory in use" lines.

Now only 18 cores available (6 x 3)

8gb per worker

Worker Id	Address	State	Cores	Memory
worker-20161007154244-10.44.174.130-8769	10.44.174.130:8769	ALIVE	6 (6 Used)	8.0 GB (8.0 GB Used)
worker-20161007154515-10.44.174.128-8769	10.44.174.128:8769	ALIVE	6 (6 Used)	8.0 GB (8.0 GB Used)
worker-20161007154543-10.44.174.129-8769	10.44.174.129:8769	ALIVE	6 (6 Used)	8.0 GB (8.0 GB Used)

Running Applications

Application ID	Name	Cores	Memory per Node	Submitted Time	User	State	Duration
app-20161007154752-0031	(kill) Fusion-20161007154751	18	8.0 GB	2016/10/07 15:47:52	ec2-user	RUNNING	0.7 s

Fusion driver app will use all cores available by default ... and all memory per executor

Cores per Driver Allocation

The configuration property `fusion.spark.cores.fraction` allows you to limit the number of cores used by the Fusion driver applications (default and scripted). For instance, in the screenshot above, we see 18 total CPUs available.

We set the cores fraction property to 0.5 via the following command:

```
$ curl -u <name>:<password> -H 'Content-type:application/json' -X PUT -d '0.5' \
http://localhost:8764/api/apollo/configurations/fusion.spark.cores.fraction
```

This cuts the number of available cores in half, as shown in the following screenshot:

Running Applications

Application ID	Name	Cores	Memory per Node	Submitted Time	User	State	Duration
app-20161007160422-0035	(kill) Fusion-20161007160421	9	8.0 GB	2016/10/07 16:04:22	ec2-user	RUNNING	1 s

5.5.2. Ports used by Spark in Fusion

The following list shows the default port number used by Spark processes. If that a port number is not available, Spark will use the next available port by adding a '+1' to the assigned port. E.g., if 4040 is not available, Spark will use 4041 (if available, or 4042 ...etc).

Process	Port number
Spark master web UI	8767
Spark worker web UI	8082
SparkContext web UI	4040
Spark master listening port	8766
Spark worker listening port	8769
Spark driver listening port	random (spark.driver.port)
Spark executor listening port	random (spark.executor.port)
Spark blockmanager port	random (spark.blockManager.port)
Spark file server port	random (spark.fileserver.port)
Spark REPL class server port	random, can be overridden through spark.replClassServer.port

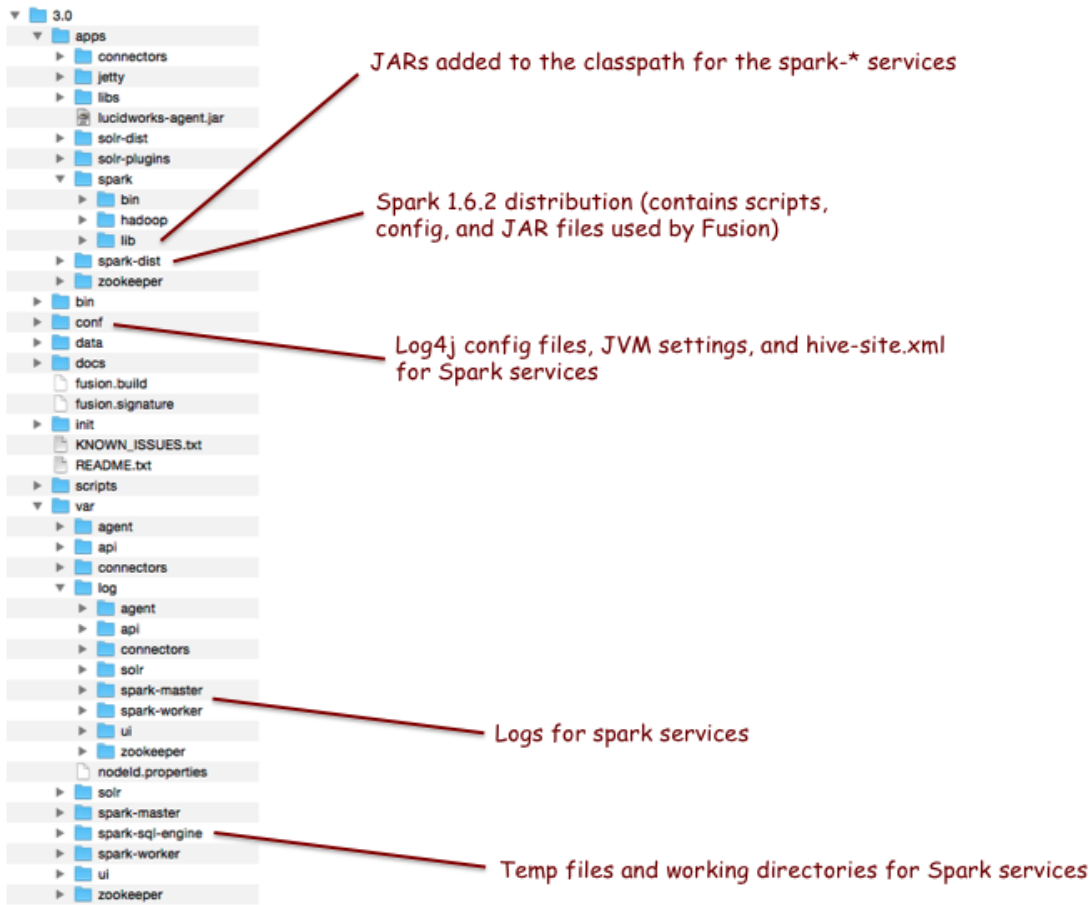
5.5.3. Directories and Temporary Files

Shaded jar file

The shaded jar file is downloaded to the `var/api/work` folder. If one of the jars in the (api) has changed, then a new shaded jar will be created with an updated name.

Temporary work directories

Temporary work dirs (spark-workDir-*) are created in `var/` when an application is running. They are removed after the driver is shut down or closed.



5.5.4. Connection Configurations for an SSL-enabled Solr cluster

You'll need to set these Java system properties used by SolrJ:

- `javax.net.ssl.trustStore`
- `javax.net.ssl.trustStorePassword`
- `javax.net.ssl.trustStoreType`

For the following Spark configuration properties:

- `spark.executor.extraJavaOptions`
- `fusion.spark.driver.jvmArgs`
- `spark.driver.extraJavaOptions`

```
$ curl -H 'Content-type:application/json' -X PUT \  
-d '-Djavax.net.ssl.trustStore=/opt/app/jobs/ssl/solrtrust.jks -Djavax.net.ssl.trustStorePassword=changeit  
-Djavax.net.ssl.trustStoreType=jks' \  
"http://localhost:8765/api/v1/configurations/spark.executor.extraJavaOptions"  
  
$ curl -H 'Content-type:application/json' -X PUT \  
-d '-Djavax.net.ssl.trustStore=/opt/app/jobs/ssl/solrtrust.jks -Djavax.net.ssl.trustStorePassword=changeit  
-Djavax.net.ssl.trustStoreType=jks' \  
"http://localhost:8765/api/v1/configurations/fusion.spark.driver.jvmArgs"  
  
$ curl -H 'Content-type:application/json' -X PUT \  
-d '-Djavax.net.ssl.trustStore=/opt/app/jobs/ssl/solrtrust.jks -Djavax.net.ssl.trustStorePassword=changeit  
-Djavax.net.ssl.trustStoreType=jks' \  
"http://localhost:8765/api/v1/configurations/spark.driver.extraJavaOptions"
```

5.6. Scaling Spark Aggregations

For this section, we'll walk through the process of running a simple aggregation on 130M signals (created from synthetic data).

One of the most common issues when running an aggregation job over a large signals data set is task timeout issues in Stage 2 (foreachPartition). This is typically due to slowness indexing aggregated jobs back into Solr or JavaScript functions. The solution is to increase the number of partitions of the aggregated RDD (the input to Stage 2). You need to increase the default parallelism by setting the following configuration property:

```
> curl -u ... -H 'Content-type:application/json' -X PUT -d '72'
"$FUSION_API/configurations/spark.default.parallelism"
```

After making the change, the foreachPartition stage of the job will use 72 partitions:

Details for Job 0

Status: RUNNING

Job Group: aggregation_157aaefd2c6Tab6f880f

Active Stages: 1

Pending Stages: 1

▶ Event Timeline

▶ DAG Visualization

Active Stages (1)

Stage Id	Pool Name	Description	Submitted	Duration	Tasks: Succeeded/Total
0	default	perfJob mapToPair at EventAggregator.java:329 +details (kill)	2016/10/09 19:34:23	7.7 min	56/144

Pending Stages (1)

Stage Id	Pool Name	Description	Submitted	Duration	Tasks: Succeeded/Total	In
1		foreachPartition at AggregatorSparkJob.java:202 +details	Unknown	Unknown	0/72	

You can increase the number of rows read per page, default is 10000, by passing the rows parameter when starting your aggregation job, such as:

```
> curl -u ... -XPOST "$FUSION_API/aggregator/jobs/perf_signals/perfJob?rows=20000&sync=false"
```

For example, we were able to read 130M signals from Solr in 18 minutes at ~120K rows / sec using rows=20000 vs. 21 minutes using the default 10000.

5.7. Spark Troubleshooting

Troubleshooting Tips and Techniques

5.7.1. Job Hung in Waiting Status

Check the logs for a message that looks like:

```
2016-10-07T11:51:44,800 - WARN [Timer-0:Logging$class@70] - Initial job has not accepted any resources; check your cluster UI to ensure that workers are registered and have sufficient resources
```

If you see this, then it means your job has requested more CPU or memory than is available. For instance, if you ask for 4g but there is only 2g available, then the job will just hang in WAITING status.

5.7.2. Lost Executor Due to Heartbeat Timeout

If you see errors like the following:

```
2016-10-09T19:56:51,174 - WARN [dispatcher-event-loop-5:Logging$class@70] - Removing executor 1 with no recent heartbeats: 160532 ms exceeds timeout 120000 ms

2016-10-09T19:56:51,175 - ERROR [dispatcher-event-loop-5:Logging$class@74] - Lost executor 1 on ip-10-44-188-82.ec2.internal: Executor heartbeat timed out after 160532 ms

2016-10-09T19:56:51,178 - WARN [dispatcher-event-loop-5:Logging$class@70] - Lost task 22.0 in stage 1.0 (TID 166, ip-10-44-188-82.ec2.internal): ExecutorLostFailure (executor 1 exited caused by one of the running tasks) Reason: Executor heartbeat timed out after 160532 ms
```

This is most likely due to an OOM in the executor JVM (preventing it from maintaining the heartbeat with the application driver). However, we've seen cases where tasks fail, but the job still completes, so you'll need to wait it out to see if the job recovers.

Another situation when this may occur is when a shuffle size (incoming data for a particular task) exceeds 2GB. This is hard to predict in advance because it depends on job parallelism and the number of records produced by earlier stages. The solution is to re-submit the job with increased job parallelism.

5.7.3. Spark Master won't start on EC2

<http://deploymentzone.com/2014/01/06/aws-instances-and-java-net-unknownHostException/>

5.8. Machine Learning Models in Fusion

Fusion provides the following tools required for the model training process:

- Solr can easily store all your training data.
- Spark jobs perform the iterative machine learning training tasks.
- Fusion’s blob store facility makes the final model available for processing new data.

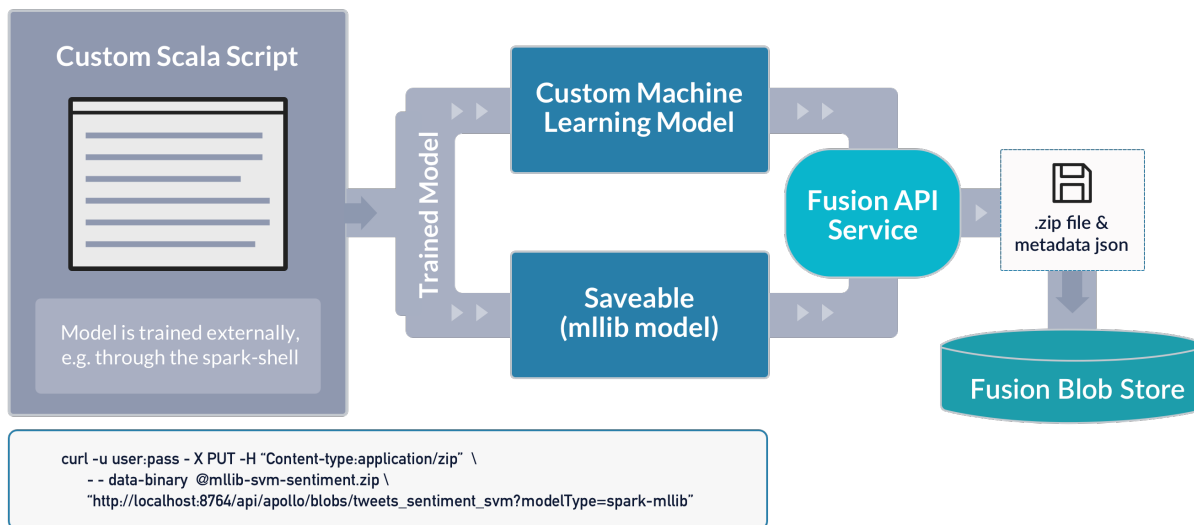
5.8.1. Training Models

Note: The approach for training models explained in this section still works in Fusion 3.1. A new approach introduced in Fusion 3.1 lets you create model-training jobs in the Fusion UI. See [Machine Learning in Lucidworks Fusion](#) for more information.

An example Scala script to train an SVM-based sentiment classifier for tweets is provided in the Lucidworks GitHub repo for the spark-solr project, see file [SVMExample.scala](#). The SVMExample code illustrates how to read training data from Solr into Spark using the SparkSQL DataFrame API, how to run the iterations necessary to train an SVM model, and how to store the final model.

The following diagram depicts this process:

Supervised Machine Learning Model Training Workflow

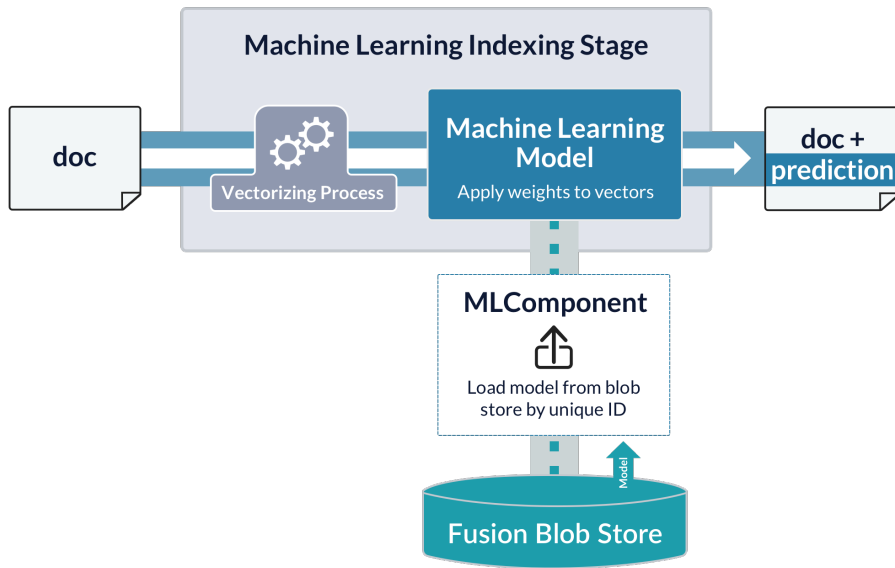


The class `com.lucidworks.spark.fusion.FusionMLModelSupport` is used to save the final model into Fusion’s blob store. It uploads the Spark MLib model directory together with a file named `spark-mllib.json`.

5.8.2. Model Prediction

Fusion’s blob store requires all stored objects have a unique ID. Once the model is stored in the Fusion blobstore, it is available to Fusion’s index and query Machine Learning pipeline stages, which use the model to make predictions for new data in pipeline documents and queries. The following diagram shows how this process works:

Supervised Machine Learning Model Serving Workflow



5.8.3. Model Checking

To test the goodness of your model in Fusion, first create either a document index pipeline or a query processing pipeline which contains a Machine Learning stage that uses your model to make predictions on your data, and then send a document or query through that pipeline which contains data for which you know what the predicted value should be. For example, given a trained sentiment classifier and an index stage configured to use it, the following document should be classified as a highly positive tweet, with a value of (close to) 1.0 in the "sentiment_d" field:

```
{ "id": "tweets-2",  
  "fields": [  
    { "name": "tweet_txt",  
      "value": "I am super excited that spring is finally here, yay! #happy" }  
  ]  
}
```

5.8.4. Metadata file "spark-mlib.json"

The file "spark-mlib.json" contains metadata about the model implementation, in particular, how it derives feature vectors from a document or query. The JSON object has the following attributes:

- "id" - a string label that used a unique ID for the Fusion blobstore, e.g., "tweets_sentiment_svm".
- "modelName" - the name of the spark-mlib class or the custom Java class that implements the com.lucidworks.spark.ml.MLModel interface.
- "featureFields" - a list of one or more field names.
- "vectorizer" - specified the processing required to derive a vector of features from the contents of the document fields listed in the "featureFields" entry.

The following example shows the "spark-mlib.json" file for the model with id "tweets_sentiment_svm":

```

{
  "id":"tweets_sentiment_svm",
  "modelClassName":"org.apache.spark.mllib.classification.SVMModel",
  "featureFields":["tweet_txt"],
  "vectorizer":[
    {
      "lucene-analyzer": {
        "analyzers":[{"
          "name":"std_tok_lower",
          "tokenizer":{"type":"standard"},
          "filters":[{"type":"lowercase"}]},
          "fields":[{"regex":".+","analyzer":"std_tok_lower"}]}
        }, {
          "hashingTF":{"numFeatures":"1000000"}
        }
      ]
    }
  }
}

```

The "vectorizer" consists of two steps: a lucene-analyzer step followed by a hashingTF step. The lucene-analyzer step can use any Lucene analyzer to perform text analysis. For more information about using the lucene-analyzer, see: <https://lucidworks.com/blog/2016/04/13/spark-solr-lucenetextanalyzer/>.

Other available vectorizer operations include: the MLlib normalizer, standard scaler, and ChiSq selector. See [SVMExample.scala](#) for an example of how to use the standard scaler.

Reference Manual

Chapter 6. Connectors Configuration Reference

Fusion comes with a standard set of built-in connectors:

- Local Filesystem connector
- File Upload connector
- JDBC connector
- Web connector

Built-in connectors are in [fusion/3.1.x/apps/connectors/bootstrap-plugins/](#).

Additional connectors are available for download at <http://lucidworks.com/connectors/>. You can look in [fusion/3.1.x/apps/connectors/plugins/](#) to see which additional connectors are currently installed.

6.1. List of connectors

Database connectors

Couchbase

The Couchbase connector uses the Cross-Datacenter Replication (XDCR) feature of Couchbase to retrieve data stored in Couchbase continuously in real-time.

[Download](#)

JDBC

The JDBC connector fetches documents from a relational database via SQL queries. Under the hood, this connector implements the Solr [DataImportHandler \(DIH\)](#) plugin.

Built-in

MongoDb

Retrieve data from a MongoDB instance.

[Download](#)

Filesystem connectors

Box.com

The Box connector retrieves data from a Box.com cloud-based data repository. To fetch content from multiple Box users, you must create a Box app that uses OAuth 2.0 with JWT server authentication. For limited testing using a single user account, you can create a Box app that uses Standard OAuth 2.0 authentication.

[Download](#)

Dropbox

The Dropbox connector retrieves data from a Dropbox cloud-based data repository.

[Download](#)

File Upload

The File Upload connector provides a convenient way to quickly ingest data from your local filesystem. It's available in the Quickstart interface in addition to the Index Workbench and the Datasources page.

Built-in

FTP

Retrieve documents using the File Transfer Protocol (FTP).

[Download](#)

Google Drive

The Google Drive connector is used to index the documents in a Google Drive account.

[Download](#)

HDFS

[Hadoop Distributed File System \(HDFS\)](#). It traverses the Hadoop file system as it would a regular Unix filesystem.

[Download](#)

Local Filesystem

A filesystem-based data store is a network of nodes to be traversed, where each node (such as a Unix file directory) provides information about its child nodes (such as the files in that directory) or references other nodes (such as links in an HTML document).

Built-in

S3

The S3 connector can access AWS S3 buckets in native format.

[Download](#)

SolrXML

The SolrXML connector indexes XML files formatted according to Solr's XML structure. It is not a generic XML file crawler; it can only index SolrXML-formatted documents.

[Download](#)

Windows Share

The Windows Share connector can access content in a Windows Share or Server Message Block (SMB)/Common Internet File System (CIFS) filesystem.

[Download](#)

Hadoop cluster connectors

Hadoop

The Hadoop Connector is a MapReduce-enabled crawler which leverages the scaling qualities of [Apache Hadoop](#).

[Download](#)

Push content connectors

Solr Push Endpoint

The Solr Push Endpoint accepts documents and pushes them to Solr using the Fusion index pipelines.

[Download](#)

Repository connectors

Alfresco

The Alfresco Connector is a crawler for the Alfresco server, which adheres to the [Content Management Interoperability Services \(CMIS\)](#) standard.

[Download](#)

Azure

The Azure connector is used to crawl an Azure instance. Its connector type is "lucid.azure" and its plugin type is "azure".

[Download](#)

Drupal

This connector uses Drupal's [Services 7.x3.11Module REST API](#).

[Download](#)

GitHub

The GitHub connector retrieves data from GitHub repositories using the [GitHub REST API](#).

[Download](#)

JIRA

The JIRA connector retrieves data from a instance of [Atlassian's JIRA](#) issue tracking system.

[Download](#)

Salesforce

[Salesforce REST API](#) to extract data from a Salesforce repository via a Salesforce Connected App.

[Download](#)

ServiceNow

The ServiceNow Datasource retrieves data from ServiceNow repository via the [ServiceNow REST API](#). ServiceNow records are stored in named tables.

[Download](#)

SharePoint

The SharePoint connector retrieves content and metadata from an on-premises SharePoint repository.

[Download](#)

SharePoint Online

The SharePoint Online connector retrieves data from cloud-based SharePoint repositories. Authentication requires a Sharepoint user who has permissions to access Sharepoint via the SOAP API. This user must be registered with the Sharepoint Online authentication server; it is not necessarily the same as the user in Active Directory or LDAP.

[Download](#)

Solr Index

A Solr connector pulls documents from an external standalone Solr instance or SolrCloud cluster using Solr's [javabin](#) response type and streaming response parser.

[Download](#)

Subversion

This connector requires a Subversion client that is compatible with javahl.

[Download](#)

Zendesk

The Zendesk connector uses the [Zendesk REST API](#) to retrieve tickets and their associated comments and attachments from a Zendesk repository.

[Download](#)

Script connectors

Javascript

The Javascript connector allows users to write ad-hoc document retrieval routines to fetch content from filesystems and websites.

[Download](#)

Social media connectors

Jive

Retrieve content from a [Jive](#) instance.

[Download](#)

Web connectors

Web

The Web connector is used to retrieve data from a Web site using HTTP and starting from a specified URL.

Built-in

6.2. Installing a connector

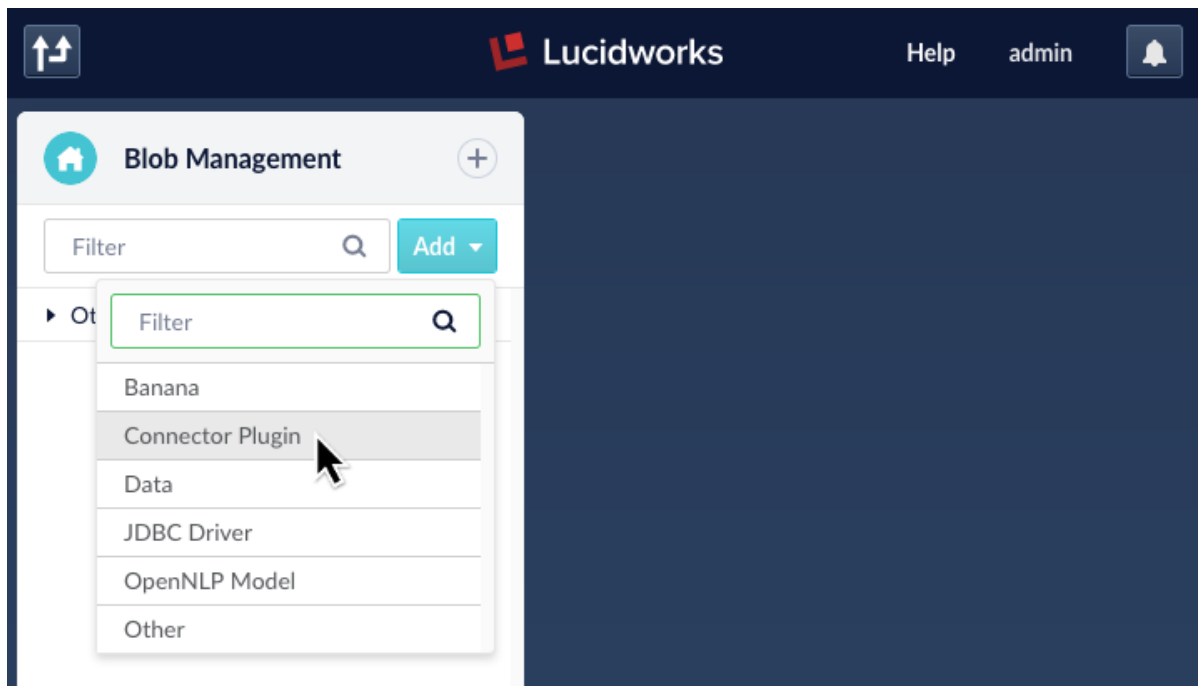
Connectors are installed by uploading them to the blob store. You can do this using the Fusion UI or the Blob Store API. Both methods are explained below.

6.2.1. Installing a connector using the Fusion UI

1. Download the connector zip file from <http://lucidworks.com/connectors/>.

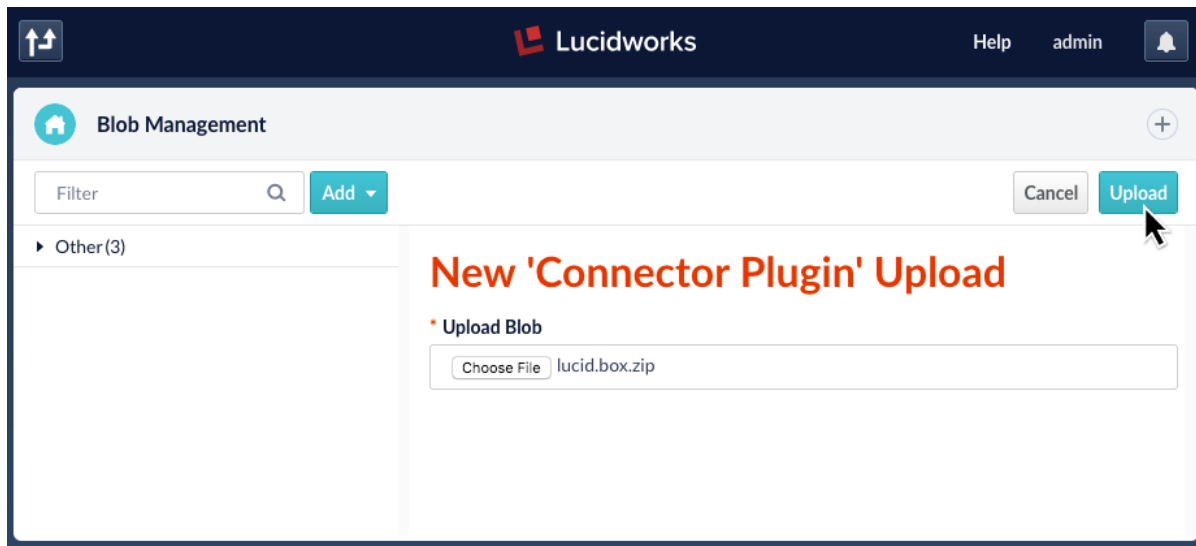
Do not expand the archive; Fusion consumes it as-is.

2. In the Fusion UI, navigate to **DevOps > Blobs**.
3. Click **Add**.
4. Select **Connector Plugin**.



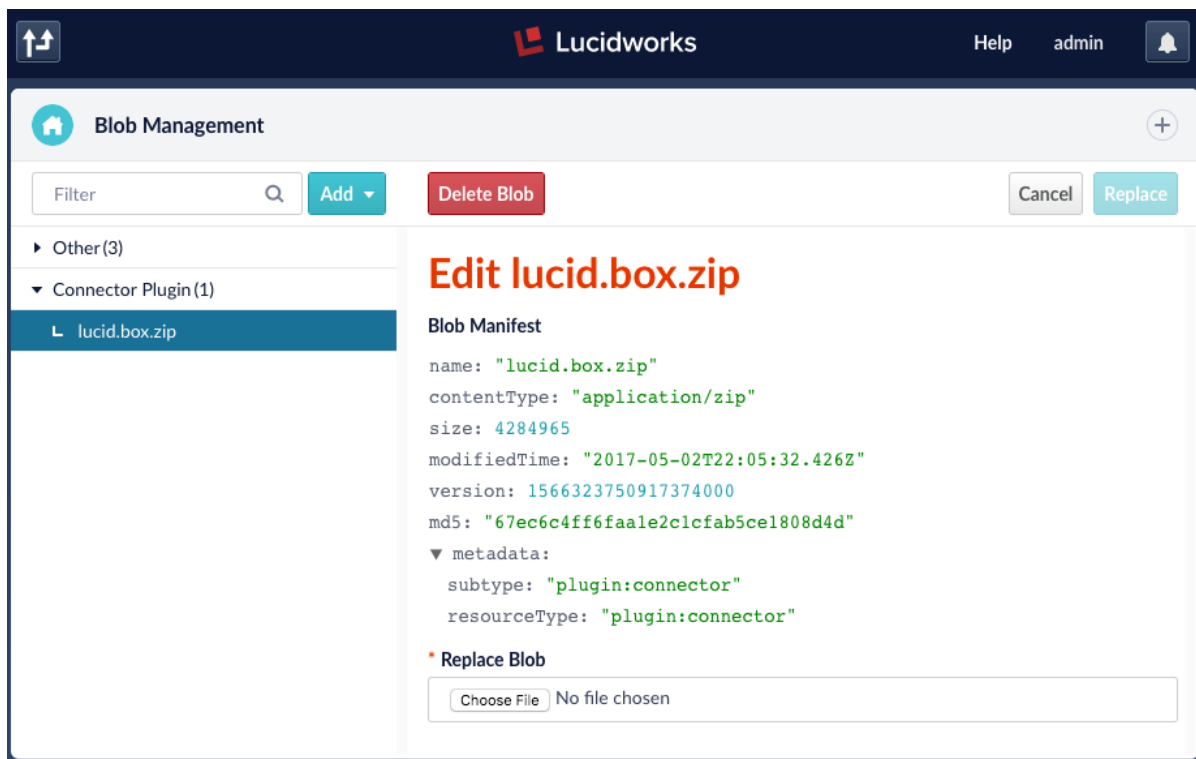
The "New 'Connector Plugin' Upload" panel appears.

5. Click **Choose File** and select the downloaded zip file from your file system.



6. Click **Upload**.

The new connector's blob manifest appears.



From this screen you can also delete or replace the connector.

6.2.2. Installing a connector using the API

1. Download the connector zip file from <http://lucidworks.com/connectors/>.

Do not expand the archive; Fusion consumes it as-is.

2. Upload the connector zip file to Fusion's blob store.

Specify an arbitrary blob ID, and a `resourceType` value of `plugin:connector`, as in this example:

```
curl -H 'content-type:application/zip' -X PUT
'localhost:8765/api/v1/blobs/myplugin?resourceType=plugin:connector' --data-binary @myplugin.zip
```

Fusion automatically publishes the event to the cluster, and the listeners perform the connector installation process on each node.

Tip	If the blob ID is identical to an existing one, the old connector will be uninstalled and the new connector will be installed in its place. To get the list of existing blob IDs, run: <code>curl -u user:pass localhost:8764/api/apollo/blobs</code>
-----	---

3. Look in `fusion/3.1.x/apps/connectors/plugins/` to verify that the new connector is installed.

6.3. Updating a connector

On Unix, you can update a connector by simply uploading the new one. Fusion overwrites the old one, and no restart is needed.

On Windows, a different procedure is needed:

How to update a Fusion connector on Windows

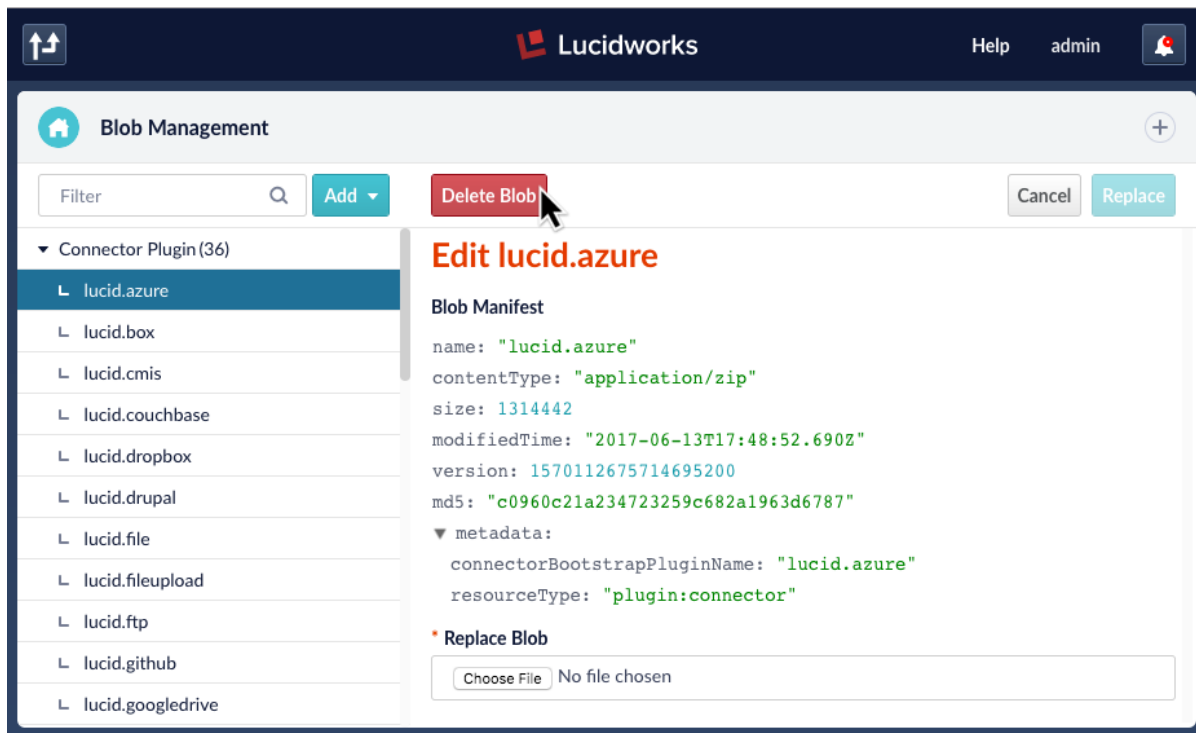
1. Delete the old connector, as explained below.
2. Restart Fusion.
3. Upload the new connector.

6.4. Deleting a connector

You can delete a connector using the Fusion UI or the Blob Store API.

6.4.1. Deleting a connector using the Fusion UI

1. In the Fusion UI, navigate to **DevOps > Blobs**.
2. Under **Connector Plugin**, select the connector to delete.
3. Click **Delete Blob**.



Fusion prompts you to confirm that you want to delete the blob.

4. Click **Yes, Delete**.

The connector disappears from the blob list.

6.4.2. Deleting a connector using the REST API

1. Get the list of blobs of the connector plugin type:

```
curl -u user:pass http://localhost:8764/api/apollo/blobs?resouType=plugin:connector
```

2. Locate the connector you want to delete, and copy its ID.

For example, the Jive connector ID is `lucid.jive`:

```
{
  "name" : "lucid.jive",
  "contentType" : "application/zip",
  "size" : 125302,
  "modifiedTime" : "2017-06-13T17:49:20.171Z",
  "version" : 1570112704530612224,
  "md5" : "7032bf2c038bb2d1e27aee82c056c0fb",
  "metadata" : {
    "connectorBootstrapPluginName" : "lucid.jive",
    "resourceType" : "plugin:connector"
  }
}
```

1. Delete the connector as follows:

```
curl -u user:pass -X DELETE http://localhost:8764/api/apollo/blobs/<id>
```

For example

```
curl -u user:pass -X DELETE http://localhost:8764/api/apollo/blobs/lucid.jive
```

A null response indicates success. You can verify that the connector is deleted like this:

```
curl -u user:pass http://localhost:8764/api/apollo/blobs | grep lucid.jive
```

6.5. Alfresco Connector and Datasource Configuration



width=600%

The Alfresco Connector is a crawler for the Alfresco server, which adheres to the [Content Management Interoperability Services](#) (CMIS) standard.

This connector was developed for the [Alfresco Community 5.0d](#) content repositories and hasn't been tested on any other versions of Alfresco or any other CMIS server.

6.5.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.cmis_password Password	type: <code>string</code>
f.cmis_repo_url CMIS Connection URL	The AtomPub connection URL, e.g., <code>'http://[host]:[port]/alfresco/api/-default-public/cmisis/versions/1.1/atom'</code> . type: <code>string</code>
f.cmis_session_cache_size Session Cache Size	The maximum number of session connections to the repository to cache. Default is 5. type: <code>integer</code> default value: <code>'5'</code>
f.cmis_username Username	type: <code>string</code>
f.maxSizeBytes Max file size (bytes)	The maximum size, in bytes, of a document to crawl. type: <code>integer</code> default value: <code>'4194304'</code>

Security Trimming

Property	Description
enable_security_trimming Enable Security Trimming	type: object

Link Discovery

Property	Description
restrictToTree Restrict to sub-directories and child pages	If true, only documents found in a tree below the start links will be fetched. type: boolean default value: 'true'
restrictToTreeAllowSubdomains Allow sub-domains in restrictToTree	If true, any sub-domain will be allowed, even if the crawl is restricted to the tree of items found below the start links. type: boolean default value: 'false'
restrictToTreeUseHostAndPath Use paths in restrictToTree	If true, the path in start links will be used to restrict items fetched. For example, if the start link is 'http://host.com/US', this option will limit all followed URLs to this path. type: boolean default value: 'false'
restrictToTreeIgnoredHostPrefixes Ignored host prefixes	List of host prefixes to ignore when checking links for restrictToTree link-legality checks. For example, 'www.' can be ignored so links with the same domain are allowed. type: array of string

Limit Documents

Property	Description
f.maxSizeBytes Max file size (bytes)	The maximum size, in bytes, of a document to crawl. type: integer default value: '4194304'
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1'
includeExtensions Included file extensions	File extensions to be fetched. This will limit this datasource to only these file extensions. type: array of string
includeRegexes Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: array of string
excludeExtensions Excluded file extensions	File extensions that should not to be fetched. This will limit this datasource to all extensions except this list. type: array of string
excludeRegexes Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string default value: ['.IMAP\s+Attachments.', '.IMAP\s+Home.', '.Data\s+Dictionary.']

Property	Description
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: boolean default value: 'true'
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: integer default value: '-1'

Crawl Performance

Property	Description
f.cmis_session_cache_size Session Cache Size	The maximum number of session connections to the repository to cache. Default is 5. type: integer default value: '5'
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: integer default value: '1'
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: integer default value: '5'

Property	Description
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: integer default value: '0'
fetchDelayMSPerHost Fetch delay per host	If true, the 'Fetch delay (ms)' property will be applied for each host. type: boolean default value: 'false'
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: integer default value: '5'
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: boolean default value: 'true'
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: boolean default value: 'true'

Property	Description
parserRetryCount Max Parser Retries	The maximum number of times the configured parser will try getting content before giving up type: <code>integer</code> default value: <code>'0'</code> exclusiveMaximum: true exclusiveMinimum: false maximum: 5 minimum: 0

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: <code>boolean</code> default value: <code>'false'</code>
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a <code>'dedupeSignature_s'</code> field. Note this may cause errors about 'immense terms' in that field. type: <code>boolean</code> default value: <code>'false'</code>
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: <code>string</code>

Property	Description
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content)\{' function, but can use any combination of document fields. The function must return a string. type: <code>string</code>

Recrawl Rules

Property	Description
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: <code>boolean</code> default value: 'true'
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: <code>boolean</code> default value: 'false'
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: <code>boolean</code> default value: 'false'
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: <code>integer</code> default value: '-1'
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: <code>array of string</code>

Property	Description
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: <code>array of string</code>
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: <code>string</code>
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: <code>boolean</code> default value: 'false'

Crawl History

Property	Description
retainOutlinks Retain links in the crawl db	Set to true for links found during fetching to be stored in the crawl db. This increases precision in certain recrawl scenarios, but requires more memory and disk space. type: <code>boolean</code> default value: 'false'
aliasExpiration Alias expiration	The number of crawls after which an alias will expire. The default is 1 crawl. type: <code>integer</code> default value: '1'
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: <code>string</code> default value: 'in-memory' enum: \{ in-memory on-disk }

Property	Description
<p>indexCrawlDBToSolr</p> <p>Index crawl database to Solr</p>	<p>EXPERIMENTAL: Set to true to index the crawl-database into a 'crawldb_' collection in Solr.</p> <p>type: <code>boolean</code></p> <p>default value: 'false'</p>
<p>reevaluateCrawlDBOnStart</p> <p>Reevaluate crawl db on start?</p>	<p>Reevaluate existing crawl db entries for legality on startup?</p> <p>type: <code>boolean</code></p> <p>default value: 'false'</p>

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>{</code></p> <p>operation : <code>{</code></p> <p>display name: Operation</p> <p>type: string</p> <p>default value: 'copy'</p> <p>description : The type of mapping to perform: move, copy, delete, add, set, or keep.</p> <p>enum: <code>{ copy move delete set add keep }</code></p> <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.6. Azure Connector and Datasource Configuration



width=600%

The Azure connector is used to crawl an Azure instance. Its connector type is "lucid.azure" and its plugin type is "azure".

Note	This connector is included with Fusion 2.1.3 and above. Older versions of Fusion had separate Azure Table and Azure Blob connectors.
Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.

6.6.1. Configuration

Property	Description
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: <code>boolean</code> default value: <code>'true'</code>
max_bytes Max Bytes	The maximum size, in bytes, of a document to crawl. type: <code>integer</code> default value: <code>'10485760'</code>
max_connections Maximum number of connections	Maximum number of simultaneous connections to the repository. This value usually does not need to be changed. type: <code>integer</code> default value: <code>'5000'</code>

Property	Description
max_threads Maximum number of threads	The maximum number of threads to use for fetching data. Each thread will create a new connection to the repository, which may make overall throughput faster, but will also require more system resources including CPU and memory. type: <code>integer</code> default value: '5'
service_type Azure service type <i>required</i>	The Azure service type to crawl, either Blobs or Tables. type: <code>string</code> enum: \{ Azure Table Azure Blob }
storage_account Storage Account <i>required</i>	The Azure storage account name. type: <code>string</code> minLength: 1
storage_container Storage Container	The name of an Azure Blob container. type: <code>string</code> minLength: 1
table_filter_statement Table Filter Statement	A filter to apply to the crawl. Use Azure's syntax for filtering table queries. type: <code>string</code>
tables Table name	The Azure table to index. type: <code>string</code> minLength: 1
token_secret Token Secret <i>required</i>	A valid Azure Access Key for authentication. type: <code>string</code> minLength: 1

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } + `source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + `target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + }</p>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
<p>unmapped</p> <p>Unmapped Fields</p>	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

ConnectorDb Configuration

Property	Description
<p>aliases</p> <p>Process Aliases?</p>	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
<p>inlinks</p> <p>Process Inlinks?</p>	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

description

Description

Optional description for datasource instance.

id

Datasource ID

Unique name for datasource instance.

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.7. Box.com Connector and Datasource Configuration



width=600%

The Box connector retrieves data from a Box.com cloud-based data repository. To fetch content from multiple Box users, you must create a Box app that uses OAuth 2.0 with JWT server authentication. For limited testing using a single user account, you can create a Box app that uses Standard OAuth 2.0 authentication.

Note In Fusion 3.1, the Box connector uses an improved approach for crawling large Box data repositories. If you use the Box connector, we recommend that you upgrade to Fusion 3.1 or later.

6.7.1. How the Box Connector Works

When you crawl a Fusion datasource that uses the Box connector, the Box connector performs a two-step process to crawl a Box data repository:

1. **Build a prefetch index** – The Box connector crawls file metadata *user-by-user*. It creates a distributed *prefetch index* that describes the structure of files in the repository. The prefetch index contains basic file metadata—file IDs and the directory relationships. Fusion stores the prefetch index in Solr as a system collection.
2. **Build the file index** The Box connector crawls files *file-by-file*. It uses the prefetch index to fetch the contents of files and metadata. It indexes the documents through Fusion’s index pipeline.

The prefetch index lets the Box connector crawl files randomly, file-by-file; instead of user-by-user. This gets around Box rate limits.

The initial crawl of a Box data repository can take a long time (hours or days). After the initial crawl, both the prefetch and main parts of the crawl are *incremental*, and they proceed much more quickly.

6.7.2. Overview of Steps

Note: These steps are for a multi-user Box.com data repository. For limited testing using a single user account, you can create a Box app that uses Standard OAuth 2.0 authentication.

Following is an overview of the steps required to set up Box and Fusion, and to crawl a Box data repository.

Set Up Box:

1. Sign up for a Box developer account.
2. Enable 2-step verification.
3. Create a Box app that Fusion can use to crawl the Box files.
4. Configure your app to use a Box service account.

Set Up Fusion:

1. Install Fusion’s Box Connector.
2. Create datasources in Fusion that use the Box connector.

Crawl a Box Data Repository:

1. Crawl the Fusion datasources.

6.7.3. Set Up Box

Set up Box so that Fusion can crawl Box data repositories.

Step 1: Sign Up for a Box Developer Account

If you already have an account, proceed to Step 2.

1. Open the [Box Developers web portal](#).
2. In the upper right corner, click **Sign Up**.
3. Enter the requested information and click **Submit**.
4. Open the confirmation email and click **Verify Email**.
5. Log in to your account.

Step 2: Enable 2-Step Verification

1. Log in to your Box developer account as the Admin.
 - a. Open the [Box Developers web portal](#).
 - b. In the upper right corner, click **Log In**.
2. Create the Box account that you want to use for crawling.
 - a. Open the [Box Admin Console](#).
 - b. Click **Users and Groups > Create account**.
 - c. Enter the Name and Email for the user, and then click **Add user**.
 - d. Click the user you just created to enter its user settings.
 - e. Make this user a [Co-Admin](#) by selecting Co-Admin checkbox. Once clicked, a pane titled "User is granted the following administrative privileges" appears. It is very important that you select all of the following:
 - Manage users
 - Manage groups
 - View users' content
 - Log in to users' accounts
 - Run new reports and access existing reports
 - f. Click **Save**.
 - g. Close the Admin Console browser tab.
3. Enable 2-step verification for unrecognized logins:
 - a. Open the [Account Settings](#) page. (You can reach this page from the drop-down menu under your initials.)
 - b. Under **Authentication**, select **Require 2-step verification for unrecognized logins**.
 - c. Choose your **Country** and enter a **Mobile Phone Number**, and then click **Continue**.
 - d. Enter the verification code you receive, and then click **Continue**.
 - e. If you are using a new mobile device, Box will send you a second code. Enter it, and then click **Submit**.

- f. Click **Save Changes**.

Step 3: Create a Box App that Fusion Can Use to Crawl the Box files

Create a Box app that uses OAuth 2.0 with JWT server authentication.

If you already have an app, configure it.

1. Open the [Box Developer's Console](#).
2. Click **Create New App**.
3. Click **Custom App**, and then click **Next**.
4. Click **OAuth 2.0 with JWT (Server Authentication)**, and then click **Next**.
5. Name your app, and then click **Create App**. The name must be globally unique across all apps created by all Box users.
6. Click **View Your App**.

Step 4: Configure Your App to Use a Box Service Account

1. Use OpenSSL to create a private/public key pair:
 - a. Install OpenSSL if you need to. Windows instructions are [here](#).
 - b. Open a Command Prompt window and run these commands to generate a private/public key pair:

```
openssl genrsa -aes128 -out private_key.pem 2048
```

Enter a password for the private key when prompted.

```
openssl rsa -pubout -in private_key.pem -out public_key.pem
```

In the current directory of the Command Prompt, you will now have private and public key files, [private_key.pem](#) and [public_key.pem](#) respectively.

2. Open the [Box Dev Console](#), log in as Admin if you are asked to log in, and then click your app.
3. Click **Configuration**.
4. Configure scopes and advanced features:
 - a. Under **Application Scopes**, deselect **Manage groups**.
 - b. Under **Advanced Features**, select **Generate User Access Tokens** and **Perform Actions as Users**.
 - c. Under **Application Access**, select **Enterprise**.
 - d. Click **Save Changes**.
5. In the **Add and Manage Public Keys** area, click **Add a public key** and paste the contents of the [public_key.pem](#) file (generated from the Box key creation) into the text box.
 - a. Make a note of the new **Public Key ID** that you just created.
6. Under **OAuth 2.0 Credentials**, click **COPY** for the Client ID.
7. Authorize your app:

- a. Open the [Box Admin Console](#).
- b. Click **Settings > Enterprise Settings** (or **Business Settings**) > **Apps**.
- c. Under **Custom Applications**, click **Authorize New App**.
- d. In the **API Key** box, paste the Client ID credential you copied in step 6, and then click **Next**.
- e. Read the App Authorization dialog and click **Authorize**.
- f. Close the Admin Console browser tab.

Note	If you change your app's configuration later, you must repeat this step to re-authorize your app.
------	---

8. Close the Dev Console browser tab.

6.7.4. Set Up Fusion

Set up Fusion to crawl Box data repositories.


Step 5: Install Fusion's Box Connector

1. Navigate to the [Fusion connectors download page](#).
2. On the Connectors page, provide the requested contact information, select **Box**, and then click **Submit**.
3. You will receive an email. Open it and click **Get My Connectors**.
4. On the Connector Downloads page, click the Box connector to download it. Don't expand the archive.
5. Open the Fusion UI and click **Devops > Blobs**.
6. Click **Add**.
7. Select **Connector Plugin**.
8. Click **Choose File**, select the file, and then click **Open**.
9. Click **Upload**.

Step 6: Create Datasources

Create datasources that use the Box connector to access the Box data repository.

For each datasource:

1. From the Fusion launcher, click **Search > Home  > Datasources**.
2. Click **Add**.
3. Choose **Box.com**.
4. Fill in the form. Note the following regarding configuration settings to use:

Setting	Notes
Start Links	<p>Each start link defined for the datasource must consist of a numeric Box file ID or directory ID, preceded by a / (slash). The root directory of any Box account has ID 0 (zero). To crawl your entire Box repository, enter '/0'. These images indicate with underlines where you can get a folder ID or file ID. Select a folder or file at Box.com.</p> <p>Folder ID:</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;"> 🔒 Secure https://app.box.com/folder/34192617287 </div> <p>You would enter the start link /34192617287.</p> <p>File ID:</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;"> 🔒 Secure https://app.box.com/file/204871656422 </div> <p>You would enter the start link /204871656422:</p>
API Key	In the Box Developer Console , select the app. On the Configuration tab under OAuth 2.0 Credentials , use the Client ID .
API Secret	In the Box Developer Console , select the app. On the Configuration tab under OAuth 2.0 Credentials , use the Client Secret .
JWT App User ID	Email address that you use to sign in to your Box co-admin account. Use the Co-admin account you created earlier for this.
JWT Public Key Id	In the Box Developer Console , select the app. On the Configuration tab, under Add and Manage Public Keys , use the ID for a public key.
JWT Private Key File	Full path to the private-key file you created that matches your JWT Public Key Id
JWT Private Key Password	Passphrase for the private key (from the private-key file you created during Box key creation)
Distributed crawl collection name	Collection that will contain the index that results from the crawl
Box.com children responses per page	Use the default value of 1000.
Nested folder depth limit	Generally, you want a number that will crawl all documents, so keep the default value. For testing, you could reduce the number substantially to speed up the crawl.

Setting	Notes
Number of partition buckets	Divide the number of files by 5000. Use that number or 10000, whichever is smaller.
Number distributed crawl datasources	Use 1 to 27.
Number of pre-fetch index creator threads	A number between 2 and 5. Use 2 for small datasources and 5 for huge datasources (over 10 million files).

5. Click **Save**.

6.7.5. Crawl a Box Data Repository


Crawl a Box data repository.

Step 7: Crawl the Fusion Datasources


Crawl the datasources, which use Fusion’s Box connector to access the Box data repository. Fusion’s Box connector uses the prefetch index to fetch the contents of each file from Box.com, get metadata from both the distributed index and Box.com, and index the documents through Fusion’s index pipeline.

You can:

- **Run the crawl now:**

1. From the Fusion launcher, click **Search > Home**  **> Datasources**.
2. Click the datasource.
3. Click **Start Crawl**.

- **Schedule the crawl:**

1. From the Fusion launcher, click **Devops > Home**  **> Scheduler**.
2. Click the the row for the job that corresponds to the datasource.
3. Specify schedule information, and then click **Save**.

See a tutorial about the complete configuration process here (full-screen recommended):

6.7.6. Box Authorization, Access, and Refresh Tokens

Fusion supports two methods of authentication with the Box API:

- JSON Web Token (JWT)
- OAuth2

Authentication Using JWT

Box.com has released a [Box Developer Edition](#). The Box Developer Edition offers a new functionality where app users will no longer have to create their own Box accounts to use an application.

App Auth uses the [JSON Web Token \(JWT\)](#) authentication architecture to establish a trusted connection with Box, allowing an application to provision and manage a new type of Box account that removes the friction of multiple logins for users or the difficulty of managing services.

For this option, Fusion needs the inputs below to crawl your Box data.

Required options are highlighted.

UI Label, API Name	Description
JWT App User ID <code>f.fs.appUserId</code>	Email address that you use to sign in to your Box co-admin account. Use the Co-admin account you created earlier for this.
JWT Public Key ID <code>f.fs.publicKeyId</code>	The public key prefix registered in Box Auth that you want to use to authenticate with
JWT Private Key File Path <code>f.fs.privateKeyFile</code>	Full path to the private-key file that Box downloaded (for the private key that corresponds to the public key you chose for JWT Public Key Id)
JWT Private Key File Password <code>f.fs.privateKeyPassword</code>	Passphrase for the private key (from the private-key file)
Tip	The biggest advantage to using the JWT App Auth Users approach is that you don't have to generate new refresh tokens. The public/private key file combination remain valid indefinitely.

6.7.7. Authentication Using OAuth 2.0

For limited testing using a single user account, you can create a Box app that uses Standard OAuth 2.0 authentication.

1. Log in to your Box developer account as the Admin.
 - a. Open the [Box Developers web portal](#).
 - b. In the upper right corner, click **Log In**.
2. Open the page for creating a new app and click **Create New App**.
3. Click **Custom App**, and then click **Next**.
4. Click **Standard OAuth 2.0 (User Authentication)**, and then click **Next**.
5. Name your app, and then click **Create App**. The name must be globally unique across all apps created by all Box users.
6. Click **View Your App**.
7. On the Configuration page:
 - a. Click the Authentication Method **Standard OAuth 2.0 (User Authentication)**.
 - b. Set the **Redirect URI** to <http://localhost> or <http://0.0.0.0>. This address is not used by Fusion, but cannot be left blank.
 - c. Click **Save Changes**.

6.7.8. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.addFileMetadata Add file metadata	Set to true to add information about documents found in the filesystem to the document, such as document owner, group, or ACL permissions. type: <code>boolean</code> default value: 'true'
f.fs.apiKey API Key	The Box API Key. type: <code>string</code>
f.fs.apiSecret API Secret	The Box API Secret. type: <code>string</code>
f.fs.appUserId JWT App User ID	(JWT only) The JWT App User ID with access to crawl. type: <code>string</code>
f.fs.childrenPageSize Box.com children responses per page	The number of results to get from Box.com API's children() methods. Default is the max of 1000, Range can be 1-1000. type: <code>integer</code> default value: '1000'
f.fs.distributedCrawlCollectionName Distributed crawl collection name	The collection name of the Distributed Crawl Collection. If you do not specify one, it will use 'system_box_distributed_crawl'. type: <code>string</code>

Property	Description
f.fs.distributedCrawlDatasourceIndex Distributed crawl datasource index	Distributed Job index. Zero-based index of what distributed job index this data source represents. Must be in range [0, numDatasources]. For example, if you have 3 jobs in an distributed crawl, the index can be either 0, 1 or 2. Each data source must have a unique distributedJobIndex. Once the pre-fetch index is created, this index is used to signify the chunk of the file IDs that this node is responsible for indexing from the Distributed Crawl Collection. type: <i>integer</i> default value: '0'
f.fs.excludedExtensions Exclude extensions	Comma separated list of extensions. No Box Files or Folders that have a filename that ends with any of these extensions will be crawled. Case will be ignored. E.g. .txt,.xls,.DS_Store type: <i>string</i>
f.fs.generatedSharedLinksAccess Generated Shared Link Access	Only applicable when Generate Shared Links when Absent is selected... Sets the shared link access setting. Can be left blank (the default) or set to open, company or collaborators type: <i>string</i>
f.fs.generatedSharedLinksExpireDays Generated Shared Link Expires After Days	Only applicable when Generate Shared Links when Absent is selected... this will control how many days the shared links stay valid for. 0 for unlimited. type: <i>integer</i> default value: '0'
f.fs.isGenerateSharedLinkPermissionCanDownload Generated Shared Link Has Download Permission	Only applicable when Generate Shared Links when Absent is selected... On the box shared link, is the "can download" permission granted? type: <i>boolean</i>

Property	Description
f.fs.isGenerateSharedLinkPermissionCanPreview Generated Shared Link Has Preview Permission	Only applicable when Generate Shared Links when Absent is selected... On the box shared link, is the "can preview" permission granted? type: <code>boolean</code>
f.fs.isGenerateSharedLinkWhenAbsent Generate Shared Link When Absent	If this is selected, the crawler will automatically create a shared link for any non-shared documents it finds while crawling. Note: This will change all documents to 'Shared' in your Box view. Use with caution. type: <code>boolean</code>
f.fs.max_request_attempts Box Max Request Retries	If Box API throws an error when trying to get a file, how many times do we retry before giving up? type: <code>integer</code> default value: '20'
f.fs.nestedFolderDepth Nested folder depth limit	Maximum depth of nested folders that will be crawled. Range: [1, int-max]. Default is int-max. type: <code>integer</code> default value: '2147483647'
f.fs.numDistributedDatasources Number distributed crawl datasources	Number of separate datasource jobs that will be running in this distributed crawl. In other words, how many datasources are part of this crawl? This value is needed in order to distribute work evenly amongst multiple jobs. type: <code>integer</code> default value: '1'
f.fs.numPreFetchIndexCreationThreads Number of pre-fetch index creator threads	The number of concurrent threads that will create the Distributed Pre-fetch Index. Default: 5 type: <code>integer</code> default value: '5'

Property	Description
f.fs.partitionBucketCount Number of partition buckets	Number of partition buckets to be used during the full crawl. Default is 5000. type: <i>integer</i> default value: '5000'
f.fs.privateKeyFile JWT Private Key File	(JWT only) Path to the private key file. type: <i>string</i>
f.fs.privateKeyPassword JWT Private Key Password	(JWT only) The password you entered for the private key file. type: <i>string</i>
f.fs.proxyHost Proxy host	The address to use when connecting through the proxy. type: <i>string</i>
f.fs.proxyPort Proxy port	The port to use when connecting through the proxy. type: <i>integer</i>
f.fs.proxyType Proxy type	Type of proxy to use, if any. Allowed values are 'HTTP' and 'SOCKS'. Leave empty for no proxy. type: <i>string</i>
f.fs.publicKeyId JWT Public Key Id	(JWT only) The public key prefix from the box.com public keys. type: <i>string</i>
f.fs.refreshToken OAuth Refresh Token	OAuth Refresh token (Not needed for JWT). type: <i>string</i>
f.fs.refreshTokenFile OAuth Refresh Token File	File that stores the refresh token for the next session. type: <i>string</i> default value: 'refresh_token.txt'

Property	Description
f.fs.user_excludes User exclude regexes	In addition to the user filter, you can here optionally specify regexes matching user names that should not be crawled. type: <code>array of string</code>
f.fs.user_filter_term User Filter Term	If you specify a user filter term, then a users files will only be crawled if their login starts with the user filter term. Can be comma separated list of multiple filter terms. Example: a,b,c,v would be all box users that have a login starting with a,b,c, or v. This value can be empty to return all results. type: <code>string</code>
f.maxSizeBytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: <code>integer</code> default value: '4194304'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of documents to fetch. type: <code>integer</code> default value: '0'

Limit Documents

Property	Description
f.maxSizeBytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: <code>integer</code> default value: '4194304'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of documents to fetch. type: <code>integer</code> default value: '0'

Property	Description
f.addFileMetadata Add file metadata	Set to true to add information about documents found in the filesystem to the document, such as document owner, group, or ACL permissions. type: boolean default value: 'true'
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1'
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: boolean default value: 'true'
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: integer default value: '-1'
addFileMetadata	type: ``
maxSizeBytes	type: ``
minSizeBytes	type: ``

Security Trimming

Property	Description
enable_security_trimming Enable Security Trimming	type: object object attributes: \{ }

Crawl Performance

Property	Description
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: integer default value: '1'
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: integer default value: '5'
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: integer default value: '0'
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: integer default value: '5'
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: boolean default value: 'true'

Property	Description
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: boolean default value: 'true'
f.fs.max_request_attempts Box Max Request Retries	If Box API throws an error when trying to get a file, how many times do we retry before giving up? type: integer default value: '20'

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: boolean default value: 'false'
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: boolean default value: 'false'
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: string

Property	Description
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content){}' function, but can use any combination of document fields. The function must return a string. type: <code>string</code>

Recrawl Rules

Property	Description
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: <code>boolean</code> default value: 'false'
batch_incremental_crawling Batch Incremental crawling	When enabled, the re-crawl processes will retrieve just the new, modified and deleted files from Box file system. This feature only works if the user is an enterprise admin user. type: <code>boolean</code> default value: 'true'
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: <code>boolean</code> default value: 'false'
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: <code>boolean</code> default value: 'false'
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: <code>integer</code> default value: '-1'

Property	Description
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: array of string
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: array of string
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: string
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: boolean default value: 'false'

Crawl History

Property	Description
retainOutlinks Retain links in the crawl db	Set to true for links found during fetching to be stored in the crawl db. This increases precision in certain recrawl scenarios, but requires more memory and disk space. type: boolean default value: 'false'
aliasExpiration Alias expiration	The number of crawls after which an alias will expire. The default is 1 crawl. type: integer default value: '1'

Property	Description
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: <code>string</code> default value: 'in-memory' enum: \{ in-memory on-disk }
indexCrawlDBToSolr Index crawl database to Solr	EXPERIMENTAL: Set to true to index the crawl-database into a 'crawldb_' collection in Solr. type: <code>boolean</code> default value: 'false'

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>{</code></p> <p>operation : <code>{</code></p> <p>display name: Operation</p> <p>type: string</p> <p>default value: 'copy'</p> <p>description : The type of mapping to perform: move, copy, delete, add, set, or keep.</p> <p>enum: <code>{ copy move delete set add keep }</code></p> <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.8. Couchbase Connector and Datasource Configuration



width=600%

The Couchbase connector uses the Cross-Datacenter Replication (XDCR) feature of Couchbase to retrieve data stored in Couchbase continuously in real-time.

This connector has been tested for compatibility with Couchbase Server 2.5.1 Enterprise Edition.

6.8.1. Indexing and Commits

Because this connector retrieves data continuously, two properties are available to control the frequency of commits to Solr, which makes the documents available for user queries. The properties define the maximum number of documents to queue for a commit (set to 50,000 by default) and the maximum amount of time to wait between commits (set to 120 seconds, or 2 minutes). Documents will be committed when one of those thresholds is reached first, meaning that if 2 minutes have passed and there are only 20,000 documents, a commit will occur. Similarly, if only 1 minute has passed and there are 50,000 documents in the queue, a commit will occur. These properties can be adjusted for your own requirements if needed.

Note	This connector retrieves data continuously. You can limit the number of documents it fetches during testing by setting the maximum number of documents retrieved, or you can manually stop the connector with the Fusion UI or Connector Datasources API.
------	---

6.8.2. Splitting Couchbase Documents

Because Couchbase has a flexible data model, documents may have a nested JSON structure. It is possible to split nested documents with a `splitpath` property, which uses an XPath-style path to the element to split on. These paths do not accept wildcards.

For example, if you have a document that looks like this:

```
{
  "first": "John",
  "last": "Doe",
  "grade": 8,
  "exams": [
    {
      "subject": "Maths",
      "test" : "term1",
      "marks":90},
    {
      "subject": "Biology",
      "test" : "term1",
      "marks":86}
  ]
}
```

If we want to split this document on the 'exams' element and create two documents each with a different subject, we

would define `"splitpath":"/exams"` in our datasource definition (if using the Fusion UI to configure the datasource, you would enter the path without quotes).

The output from retrieving the document will look like this:

```
{
  "first": "John",
  "last": "Doe",
  "grade": 8,
  "exams": [
    {
      "subject": "Maths",
      "test" : "term1",
      "marks":90
    }
  ]
},
{
  "first": "John",
  "last": "Doe",
  "grade": 8,
  "exams": [
    {
      "subject": "Biology",
      "test" : "term1",
      "marks":86
    }
  ]
}
```

6.8.3. Field Mapping with Couchbase

The Couchbase connector has built-in field mapping allows mapping Couchbase fields to fields in your schema. The mapping configuration defines a field from your schema and an XPath-style path to the field in the Couchbase JSON document.

The field mapping can accept wildcards and double-wildcards to map fields automatically. Wildcards can be used, but only at the end of the path definition.

- `field_name=""` and `field_path=/docs/*` - maps all the fields under docs to the same name as given in JSON.
- `field_name=""` and `field_path=/docs/**` - maps all the fields under docs and their children fields to the same name as given in JSON.
- `field_name=searchField` and `field_path=/docs/*` - maps all the fields under /docs to a single field named 'searchField'.
- `field_name=searchField` and `field_path=/docs/**` - maps all the fields under /docs and their children fields to a single field named 'searchField'.

If mapping is not defined, a default mapping will be assigned, in the format of the second example above, i.e., `field_name=""` and `field_path=/docs/**`.

Example

This example some simple field mapping, using a single document such as this:

```

{
  "first": "John",
  "last": "Doe",
  "grade": 8,
  "exams": [
    {
      "subject": "Maths",
      "test"   : "term1",
      "marks": 90 },
    {
      "subject": "Biology",
      "test"   : "term1",
      "marks": 86 }
  ]
}

```

When we configure the datasource, we can define our field mapping as follows:

```

"field_mapping": [
  {
    "field_name": "points_i",
    "field_path": "/exams/marks"
  },
  {
    "field_name": "",
    "field_path": "/*"
  }
]

```

Two mappings are defined. The first will map the '/exams/marks' field from Couchbase to the 'points_i' field in Solr. The second maps all top-level and child fields from Couchbase to either the same field name in Solr or to a dynamic field rule.

After retrieving the document, it will look like this:

```

{
  "first_s": "John",
  "last_s": "Doe",
  "grade_i": 8,
  "exams": [
    {
      "subject_s": "Maths",
      "test_s"   : "term1",
      "points_i": 90},
    {
      "subject_s": "Biology",
      "test_s"   : "term1",
      "points_i": 86}
  ]
}

```

The 'marks' field from the original document has been mapped to the 'points_i' field; most of the other fields have been mapped to appropriate dynamic field rules.

Note that the representation of the document above is after it has been retrieved from Couchbase, but before it has been processed by the index pipelines. Since the index pipelines contain several stage types that can further transform the document, such as the Apache Tika Parser stage and the Field Mapping stage, the document that ends up indexed to Solr may be different from the document representation above. Some small iterations of crawling are recommended to be sure the documents are indexed as required.

6.8.4. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
client_host Couchbase client host	Host where Couchbase replica is running, which must be reachable from Couchbase server. type: <code>string</code> default value: <code>'127.0.0.1'</code> minLength: 1
client_port Couchbase client port	Port number where Couchbase client will be started. type: <code>integer</code> default value: <code>'9876'</code>
cluster_name Cluster name <i>required</i>	Connector's cluster name in Couchbase server. type: <code>string</code> minLength: 1
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: <code>boolean</code> default value: <code>'true'</code>

Property	Description
max_docs Max documents	The maximum number of documents to process before stopping. The default, '-1', means no maximum and this connector will continue processing content indefinitely. type: <i>integer</i> default value: '-1'
num_vbuckets Number of Couchbase VBuckets	Number of VBuckets used by Couchbase. Values should be 64 if Couchbase server is installed on Mac OS or 1024 for other platforms. Note that the number of VBuckets must be consistent for both Couchbase server and Couchbase Connector. Connector will try to obtain correct value from Couchbase server and the value specified will be used in case of failure. type: <i>integer</i> enum: \{ 64 1024 }
password Password <i>required</i>	Couchbase server's valid password. type: <i>string</i>
server_host Couchbase server host	Host where Couchbase server is running. type: <i>string</i> default value: '127.0.0.1' minLength: 1
server_port Couchbase server port	Port number where Couchbase server is running. type: <i>integer</i> default value: '8091'

Property	Description
source_buckets Source buckets <i>required</i>	Couchbase buckets to synchronize with. type: array of object minimum number of items (minItems): 1 object attributes: \{ fieldmapping : \{ display name: Field mappings type: array of object } name (required) : \{ display name: Bucket name type: string minLength : 1 } splitpath : \{ display name: Splitpath type: string default value: '/' } \}
username Username <i>required</i>	Couchbase server's valid username. type: string
verify_access Validate access	Set to true to require successful connection to the filesystem before saving this datasource. type: boolean default value: 'true'

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
unmapped Unmapped Fields	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }</p> <pre> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
inlinks Process Inlinks?	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
description Description	Optional description for datasource instance.
id Datasource ID	Unique name for datasource instance.
pipeline Pipeline ID	The index pipeline used to process documents.

6.9. Dropbox Connector and Datasource Configuration



width=600%

The Dropbox connector retrieves data from a Dropbox cloud-based data repository.

To fetch content from Dropbox, you must have an OAuth token accessible to the crawler so it can properly authenticate.

If you want to crawl all of the files owned by the account, you can simply add '/' as the startLink.

6.9.1. Dropbox Authentication

In order to authenticate to Dropbox, you must have a access token. This is obtained by first creating an app in the Dropbox platform as follows:

- Go to <https://www.dropbox.com/developers/apps> and sign in with your Dropbox account.
- Navigate to the "Your apps" page (from "App Console" in the left-hand navigation bar)
- Click "Create app". You will be taken through a short wizard-like selection process. You will be presented with the following questions:
 - What type of app do you want to create? Choose "Dropbox API app"
 - What type of data does your app need to store on Dropbox? Choose "Files and datastores"
 - Can your app be limited to its own folder? Choose "No"
 - What type of files dose your app need access to? Choose "All file types"
 - Provide a name for the app and click "Create App".
- After saving your app, you will see the settings screen for the new app. Scroll to the OAuth 2 section, then click the Generate button under "Generated access token". A long string will appear on the screen. Copy or save this token so that you can use when configuring the datasource.

Paste this access token into the "Dropbox OAuth Access Token" field in the Fusion UI. If using the REST API, this is the value of the property 'f.fs.oauthAccessToken'.

6.9.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.addFileMetadata Add file metadata	Set to true to add information about documents found in the filesystem to the document, such as document owner, group, or ACL permissions. type: boolean default value: 'true'
f.fs.connect_timeout Connect timeout	The connection timeout is the timeout in making the initial connection; i.e. completing the TCP connection handshake. type: integer default value: '60000'
f.fs.oAuthAccessToken Dropbox OAuth Access Token	Dropbox OAuth Token for the registered application. type: string
f.fs.read_timeout Read timeout	The read timeout is the timeout on waiting to read data. Specifically, if the server fails to send a byte seconds after the last byte, a read timeout error will be raised. type: integer default value: '60000'
f.maxSizeBytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: integer default value: '4194304'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of documents to fetch. type: integer default value: '0'

Limit Documents

Property	Description
f.maxSizeBytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: integer default value: '4194304'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of documents to fetch. type: integer default value: '0'
f.addFileMetadata Add file metadata	Set to true to add information about documents found in the filesystem to the document, such as document owner, group, or ACL permissions. type: boolean default value: 'true'
restrictToTree Restrict to sub-directories and child pages	If true, only documents found in a tree below the start links will be fetched. type: boolean default value: 'true'
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1'
includeExtensions Included file extensions	File extensions to be fetched. This will limit this datasource to only these file extensions. type: array of string

Property	Description
includeRegexes Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: array of string
excludeExtensions Excluded file extensions	File extensions that should not to be fetched. This will limit this datasource to all extensions except this list. type: array of string
excludeRegexes Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: boolean default value: 'true'
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: integer default value: '-1'

Crawl Performance

Property	Description
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: integer default value: '1'

Property	Description
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: integer default value: '5'
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: integer default value: '0'
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: integer default value: '5'
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: boolean default value: 'true'
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: boolean default value: 'true'

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: boolean default value: 'false'
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: boolean default value: 'false'
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: string
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content){}' function, but can use any combination of document fields. The function must return a string. type: string

Recrawl Rules

Property	Description
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: boolean default value: 'true'

Property	Description
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: boolean default value: 'false'
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: boolean default value: 'false'
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: integer default value: '-1'
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: array of string
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: array of string
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: string
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: boolean default value: 'false'

Crawl History

Property	Description
retainOutlinks Retain links in the crawldb	Set to true for links found during fetching to be stored in the crawldb. This increases precision in certain recrawl scenarios, but requires more memory and disk space. type: boolean default value: 'false'
aliasExpiration Alias expiration	The number of crawls after which an alias will expire. The default is 1 crawl. type: integer default value: '1'
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: string default value: 'in-memory' enum: \{ in-memory on-disk }
indexCrawlDBToSolr Index crawl database to Solr	EXPERIMENTAL: Set to true to index the crawl-database into a 'crawldb_' collection in Solr. type: boolean default value: 'false'
reevaluateCrawlDBOnStart Reevaluate crawldb on start?	Reevaluate existing crawldb entries for legality on startup? type: boolean default value: 'false'

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>{</code></p> <p>operation : <code>{</code></p> <p>display name: Operation</p> <p>type: string</p> <p>default value: 'copy'</p> <p>description : The type of mapping to perform: move, copy, delete, add, set, or keep.</p> <p>enum: <code>{ copy move delete set add keep }</code></p> <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.10. Drupal 7.x Connector and Datasource Configuration



width=600%

This connector uses Drupal's [Services 7.x3.11Module REST API](#).

6.10.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.cacheSize Cache size (number of entries)	The number of entries to cache when making REST requests. type: <code>integer</code> default value: '2000'
f.comment Drupal comment service	Name of the Comment resource to be able to index comment data. If you did not create an alias for the 'comment' object, keep the default. type: <code>string</code> default value: 'comment'
f.drupal_password Drupal password	Password to access the REST service, if required. type: <code>string</code>
f.drupal_username Drupal username	Optional username, only required if the REST service requires authentication. type: <code>string</code>
f.endpoint Drupal Endpoint	Name of the REST endpoint defined when you added the REST service to Drupal. type: <code>string</code> default value: 'rest'

Property	Description
f.file Drupal file service	Name of the File resource to be able to index file data. If you did not create an alias for the 'file' object, keep the default. type: <code>string</code> default value: 'file'
f.node Drupal node service	Name of the Node resource to be able to index node data. If you did not create an alias for the 'node' object, keep the default. type: <code>string</code> default value: 'node'
f.pageSize Node Page Size (number of nodes per page)	The number of items that will be returned. The Drupal default without this value is 20, this allows you to request more items and reduce the overall number of Node requests to fetch all content. type: <code>integer</code> default value: '100'
f.taxonomy_term Drupal Taxonomy Term service	Name of the Taxonomy Term resource to be able to index taxonomy term data. If you did not create an alias for the 'taxonomy_term' object, keep the default. type: <code>string</code> default value: 'taxonomy_term'
f.taxonomy_vocabulary Drupal Taxonomy Vocabulary service	Name of the Taxonomy Vocabulary resource to be able to index taxonomy data. If you did not create an alias for the 'taxonomy_vocabulary' object, keep the default. type: <code>string</code> default value: 'taxonomy_vocabulary'
f.timeoutMS Connection timeout (ms)	Time in ms to wait for a server response. type: <code>integer</code> default value: '10000'

Property	Description
f.user Drupal user service	Name of the User resource to be able to login, if authenticating to the REST service is required. If you did not create an alias for the 'user' object, keep the default. type: <code>string</code> default value: 'user'

Drupal URL Aliases

Property	Description
f.endpoint Drupal Endpoint	Name of the REST endpoint defined when you added the REST service to Drupal. type: <code>string</code> default value: 'rest'
f.node Drupal node service	Name of the Node resource to be able to index node data. If you did not create an alias for the 'node' object, keep the default. type: <code>string</code> default value: 'node'
f.user Drupal user service	Name of the User resource to be able to login, if authenticating to the REST service is required. If you did not create an alias for the 'user' object, keep the default. type: <code>string</code> default value: 'user'
f.file Drupal file service	Name of the File resource to be able to index file data. If you did not create an alias for the 'file' object, keep the default. type: <code>string</code> default value: 'file'

Property	Description
f.comment Drupal comment service	Name of the Comment resource to be able to index comment data. If you did not create an alias for the 'comment' object, keep the default. type: <code>string</code> default value: 'comment'
f.taxonomy_vocabulary Drupal Taxonomy Vocabulary service	Name of the Taxonomy Vocabulary resource to be able to index taxonomy data. If you did not create an alias for the 'taxonomy_vocabulary' object, keep the default. type: <code>string</code> default value: 'taxonomy_vocabulary'
f.taxonomy_term Drupal Taxonomy Term service	Name of the Taxonomy Term resource to be able to index taxonomy term data. If you did not create an alias for the 'taxonomy_term' object, keep the default. type: <code>string</code> default value: 'taxonomy_term'

Limit Documents

Property	Description
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: <code>integer</code> default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: <code>integer</code> default value: '-1'
includeExtensions Included file extensions	File extensions to be fetched. This will limit this datasource to only these file extensions. type: <code>array of string</code>

Property	Description
includeRegexes Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: array of string
excludeExtensions Excluded file extensions	File extensions that should not to be fetched. This will limit this datasource to all extensions except this list. type: array of string
excludeRegexes Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: boolean default value: 'true'
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: integer default value: '-1'

Crawl Performance

Property	Description
f.pageSize Node Page Size (number of nodes per page)	The number of items that will be returned. The Drupal default without this value is 20, this allows you to request more items and reduce the overall number of Node requests to fetch all content. type: integer default value: '100'

Property	Description
f.cacheSize Cache size (number of entries)	The number of entries to cache when making REST requests. type: <code>integer</code> default value: '2000'
f.timeoutMS Connection timeout (ms)	Time in ms to wait for a server response. type: <code>integer</code> default value: '10000'
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: <code>integer</code> default value: '1'
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: <code>integer</code> default value: '5'
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: <code>integer</code> default value: '0'
fetchDelayMSPerHost Fetch delay per host	If true, the 'Fetch delay (ms)' property will be applied for each host. type: <code>boolean</code> default value: 'false'

Property	Description
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: <code>integer</code> default value: <code>'5'</code>
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: <code>boolean</code> default value: <code>'true'</code>
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: <code>boolean</code> default value: <code>'true'</code>
parserRetryCount Max Parser Retries	The maximum number of times the configured parser will try getting content before giving up type: <code>integer</code> default value: <code>'0'</code> exclusiveMaximum: true exclusiveMinimum: false maximum: 5 minimum: 0

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: boolean default value: 'false'
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: boolean default value: 'false'
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: string
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content){}' function, but can use any combination of document fields. The function must return a string. type: string

Recrawl Rules

Property	Description
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: boolean default value: 'true'

Property	Description
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: boolean default value: 'false'
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: boolean default value: 'false'
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: integer default value: '-1'
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: array of string
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: array of string
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: string
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: boolean default value: 'false'

Crawl History

Property	Description
retainOutlinks Retain links in the crawldb	Set to true for links found during fetching to be stored in the crawldb. This increases precision in certain recrawl scenarios, but requires more memory and disk space. type: boolean default value: 'false'
aliasExpiration Alias expiration	The number of crawls after which an alias will expire. The default is 1 crawl. type: integer default value: '1'
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: string default value: 'in-memory' enum: \{ in-memory on-disk }
indexCrawlDBToSolr Index crawl database to Solr	EXPERIMENTAL: Set to true to index the crawl-database into a 'crawldb_' collection in Solr. type: boolean default value: 'false'
reevaluateCrawlDBOnStart Reevaluate crawldb on start?	Reevaluate existing crawldb entries for legality on startup? type: boolean default value: 'false'

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>{</code></p> <p>operation : <code>{</code></p> <p>display name: Operation</p> <p>type: string</p> <p>default value: 'copy'</p> <p>description : The type of mapping to perform: move, copy, delete, add, set, or keep.</p> <p>enum: <code>{ copy move delete set add keep }</code></p> <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.11. FTP Connector and Datasource Configuration



width=600%

Retrieve documents using the File Transfer Protocol (FTP).

The configuration property "url" specifies the protocol (`ftp`), the host address, and the path to crawl. By default, all files linked to from this URL will be processed. There are several configuration properties available to limit the crawl.

6.11.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
add_failed_docs Add failed documents	Set to true to add documents even if they partially fail processing. Failed documents will be added with as much metadata as available, but may not include all expected fields. type: <code>boolean</code> default value: <code>'false'</code>
bounds Crawl bounds	Limits the crawl to a specific directory sub-tree, hostname or domain. type: <code>string</code> default value: <code>'tree'</code> enum: <code>\{ tree host domain none }</code>
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>

Property	Description
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: boolean default value: 'true'
crawl_depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1' exclusiveMinimum: false minimum: -1
crawl_item_timeout Fetch timeout	Time in milliseconds to fetch any individual document. type: integer default value: '600000' exclusiveMinimum: true minimum: 0
exclude_paths Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string
include_extensions Included file extensions	List the file extensions to be fetched. Note: Files with possible matching MIME types but non-matching file extensions will be skipped. Extensions should be listed without periods, using whitespace to separate items (e.g., 'pdf zip'). type: array of string

Property	Description
include_paths Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: array of string
index_directories Index directories	Set to true to add directories to the index as documents. If set to false, directories will not be added to the index, but they will still be traversed for documents. type: boolean default value: 'false'
max_bytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: integer default value: '10485760' exclusiveMinimum: false minimum: -1
max_docs Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1' exclusiveMinimum: false minimum: -1
max_threads Fetch threads	The maximum number of threads to use for fetching data. Note: Each thread will create a new connection to the repository, which may make overall throughput faster, but this also requires more system resources, including CPU and memory. type: integer default value: '1'

Property	Description
maximum_connections Maximum fetch connections	Maximum number of concurrent connections to the filesystem. A large number of documents could cause a large number of simultaneous connections to the repository and lead to errors or degraded performance. In some cases, reducing this number may help performance issues. type: integer default value: '1000'
password Password	Password for the user. type: string
url Start link <i>required</i>	A starting URI for this datasource. The URI must be fully-qualified, and include the protocol, host, port and path, as appropriate. type: string minLength: 1 pattern: ..
username Username	Username with permissions to access the repository, if necessary. type: string
verify_access Validate access	Set to true to require successful connection to the filesystem before saving this datasource. type: boolean default value: 'true'

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{"operation" ⇒ "move", "source" ⇒ "fetch_time", "target" ⇒ "fetch_time_dt"}{"operation" ⇒ "move", "source" ⇒ "ds:description", "target" ⇒ "description"}</code></p> <p>object attributes: <code>{</code></p> <ul style="list-style-type: none"> <code>operation</code> : <code>{</code> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: <code>{ copy move delete set add keep }</code> <pre> } + 'source' _ (required)_ : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
unmapped Unmapped Fields	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } }</p>

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
inlinks Process Inlinks?	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

description

Description

Optional description for datasource instance.

id

Datasource ID

Unique name for datasource instance.

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.12. GitHub Connector and Datasource Configuration



width=600%

The GitHub connector retrieves data from GitHub repositories using the [GitHub REST API](#).

GitHub repository access may require authentication, in which case GitHub credentials or OAuth token Authentication will be used to perform GitHub requests. See the [GitHub Authentication](#) page for more information.

6.12.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.blobs Content (blobs)	type: <code>boolean</code> default value: <code>'false'</code>
f.branches Branches	type: <code>boolean</code> default value: <code>'false'</code>
f.collaborators Collaborators	type: <code>boolean</code> default value: <code>'false'</code>
f.commit_comments Commit comments	type: <code>boolean</code> default value: <code>'false'</code>
f.commit_diffs Commit diffs	type: <code>boolean</code> default value: <code>'false'</code>
f.commits Commits	type: <code>boolean</code> default value: <code>'true'</code>

Property	Description
f.github_accept_self_signed_ssl_certificate Accept self signed SSL certificate	type: <code>boolean</code> default value: 'true'
f.github_base_url GitHub Base URL	type: <code>string</code> default value: 'https://github.com'
f.github_enterprise GitHub Enterprise	type: <code>boolean</code> default value: 'false'
f.github_oauth_token GitHub OAuth Token	type: <code>string</code>
f.github_password GitHub password	type: <code>string</code>
f.github_username GitHub username	type: <code>string</code>
f.issue_comments Issue comments	type: <code>boolean</code> default value: 'false'
f.issues Issues	type: <code>boolean</code> default value: 'true'
f.milestones Milestones	type: <code>boolean</code> default value: 'false'
f.pull_request_comments Pull request comments	type: <code>boolean</code> default value: 'false'
f.pull_requests Pull requests	type: <code>boolean</code> default value: 'true'

Property	Description
f.releases Releases	type: <code>boolean</code> default value: 'false'
f.tags Tags	type: <code>boolean</code> default value: 'false'

Limit Documents

Property	Description
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: <code>integer</code> default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: <code>integer</code> default value: '-1'
includeExtensions Included file extensions	File extensions to be fetched. This will limit this datasource to only these file extensions. type: <code>array of string</code>
includeRegexes Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: <code>array of string</code>
excludeExtensions Excluded file extensions	File extensions that should not to be fetched. This will limit this datasource to all extensions except this list. type: <code>array of string</code>

Property	Description
excludeRegexes Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: boolean default value: 'true'
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: integer default value: '-1'

Crawl Performance

Property	Description
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: integer default value: '1'
fetchThreads Fetch threads	The number of threads to use during fetching. The value is 1 when 'enterprise:false', if 'enterprise:true' the threads can be incremented type: integer default value: '1'

Property	Description
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: integer default value: '0'
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: integer default value: '5'
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: boolean default value: 'true'
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: boolean default value: 'true'

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: boolean default value: 'false'

Property	Description
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: <code>boolean</code> default value: 'false'
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: <code>string</code>
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content)\{' function, but can use any combination of document fields. The function must return a string. type: <code>string</code>

Recrawl Rules

Property	Description
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: <code>boolean</code> default value: 'true'
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: <code>boolean</code> default value: 'false'
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: <code>boolean</code> default value: 'false'

Property	Description
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: integer default value: '-1'
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: array of string
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: array of string
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: string
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: boolean default value: 'false'

Crawl History

Property	Description
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: string default value: 'in-memory' enum: \{ in-memory on-disk }
indexCrawlDBToSolr Index crawl database to Solr	EXPERIMENTAL: Set to true to index the crawl-database into a 'crawl-db_' collection in Solr. type: boolean default value: 'false'

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>{</code></p> <p>operation : <code>{</code></p> <p>display name: Operation</p> <p>type: string</p> <p>default value: 'copy'</p> <p>description : The type of mapping to perform: move, copy, delete, add, set, or keep.</p> <p>enum: <code>{ copy move delete set add keep }</code></p> <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
pipeline Pipeline ID	The index pipeline used to process documents.

6.13. Google Drive Connector and Datasource Configuration



width=600%

The Google Drive connector is used to index the documents in a Google Drive account.

6.13.1. Google Drive authentication

There are two methods of Google Drive authentication for Fusion:

- authentication for access to site-wide documents

Use a Google account with admin-level access rights to configure access to all shared documents owned by the users in your organization.

- authentication for access to per-user documents

This type of authentication gives you access to your own documents in Google Drive. Admin-level access rights are not required.

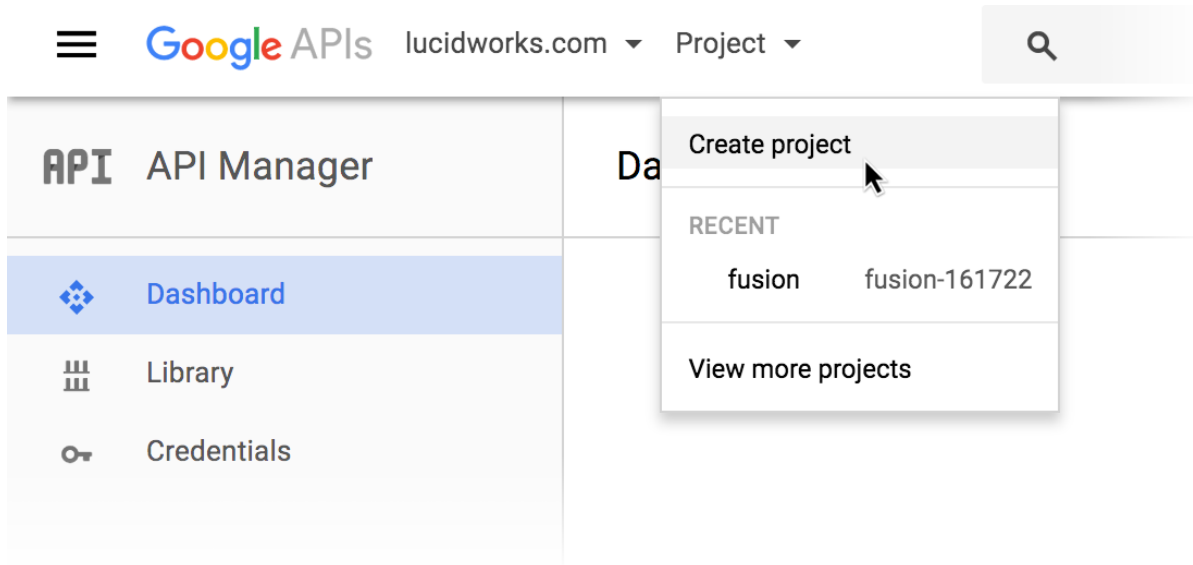
In both cases, you'll get a client ID, client secret, and refresh token from Google. These become part of your datasource configuration in Fusion.

Authentication for access to site-wide documents

In order to access all the shared documents by users in your organization, you must configure the Google Drive API and the Admin SDK. See the instructions in the [knowledge base](#).

How to configure authentication for access to site-side documents

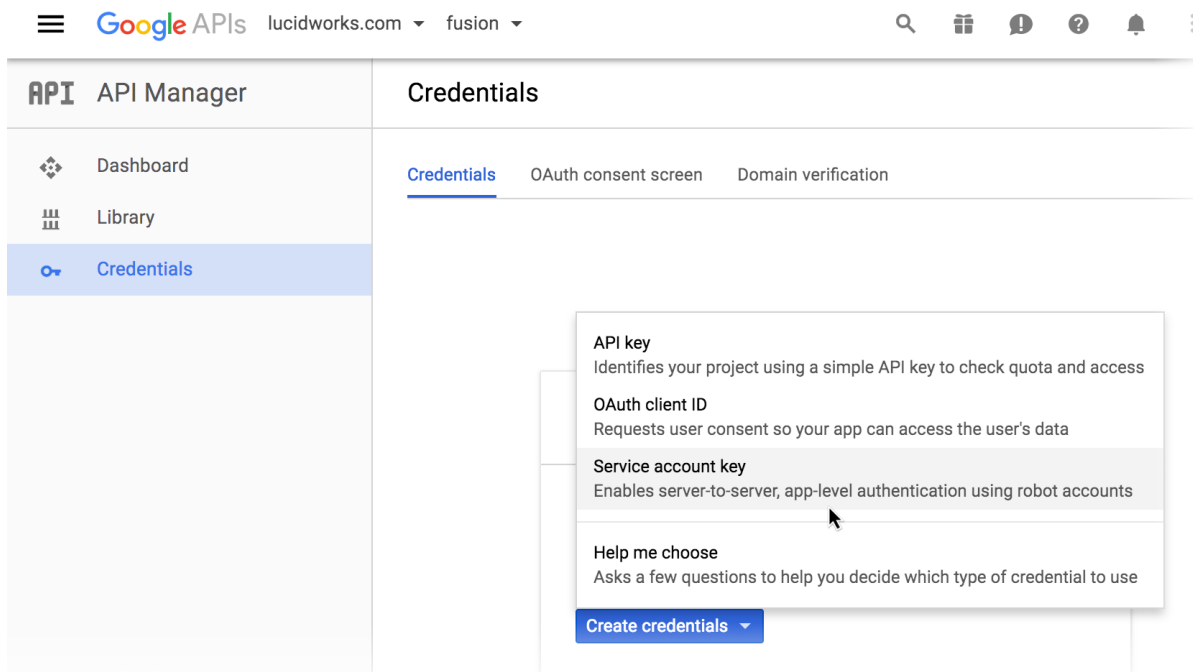
1. Log in to Google as a user with *admin-level access rights*.
2. Go to <https://console.developers.google.com/>.
3. Create a Google project for Fusion:
4. In the upper left, open the **Project** menu and select **Create Project**:



5. Enter a new project name, such as "fusion".
6. Click **Create**.
7. In the new project, click **Enable API**.
8. Under "Google Apps APIs", click **Drive API**.
9. Click **Enable**.

Google may prompt you to create credentials. Do not create credentials here; we'll do that a few steps later.


10. Click **Library**, then **Admin SDK**.
11. Click **Enable**.
12. Create a service account key:
13. Navigate to **Credentials** > **Create Credentials** > **Service account key**:



14. From the **Service account** list, select **Compute Engine default service account**.

15. Under "Key type", select **P12**.

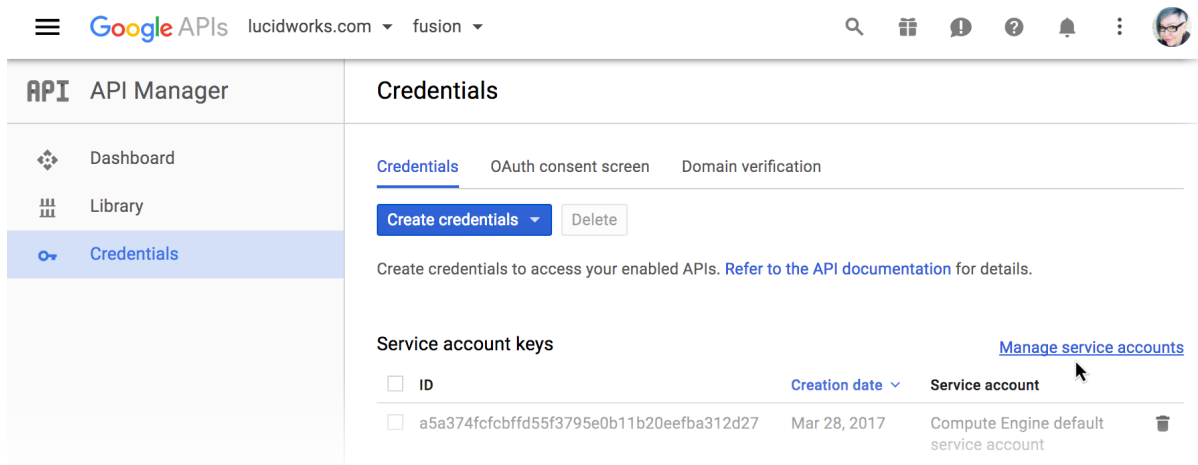
16. Click **Create**.



A new private key downloads automatically to your local drive. Google prompts you to store it securely and save the displayed password. The key and password will not be provided to you again.

17. Create a service account:


18. Click **Manage service accounts**:



19. Click **Create Service Account**.

20. Enter a service account name, such as "fusion-service-account".

21. From the **Role** list, select **Project > Service account actor**:



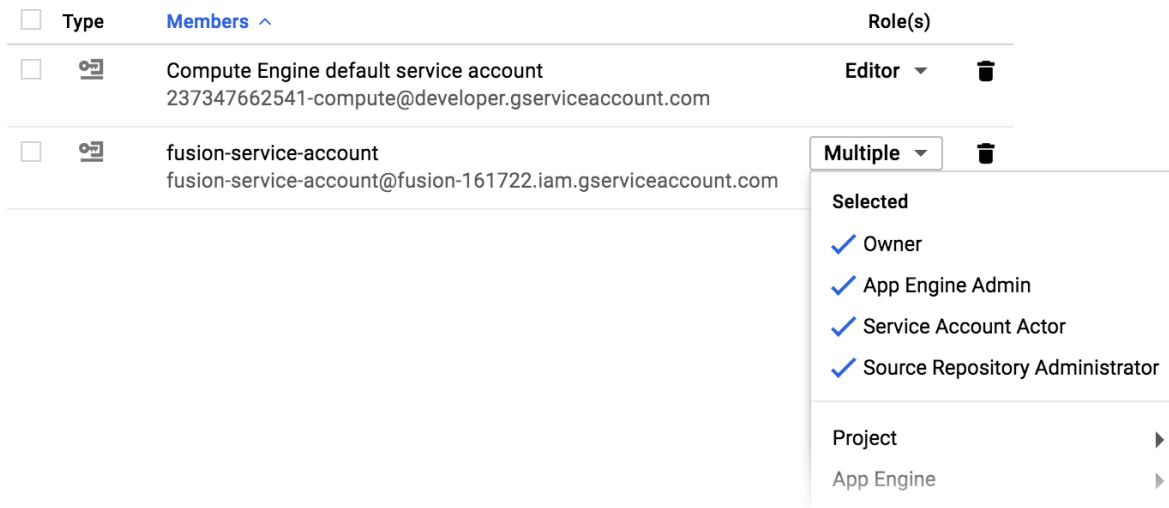
22. Select **Enable G Suite Domain-wide Delegation**.
23. Enter an arbitrary name under **Product name for the consent screen**.
24. Click **Create**.

Google displays the list of service accounts.

25. Next to the "fusion-service-account", click **View Client ID**.

You may need to scroll to the right in order to see this link.

26. Copy the client ID and service account name. Save them in a separate location.
27. Click the menu in the upper left and select **IAM & Admin**.
28. Next to the "fusion-service-account" project, select its permissions as shown below:



29. Go to <https://admin.google.com>.

30. Navigate to **Admin Console > Security**.

Remember, you must be logged in as a user with admin-level access rights.

31. Navigate to **Show more > Advanced settings > Manage API client access**.

32. Create a new API client:

- In the **Client Name** field, enter the client ID from your service account (above).
- In the **One or More API Scopes** field, enter the following:

<https://www.googleapis.com/auth/admin.directory.group>,<https://www.googleapis.com/auth/admin.directory.group.readonly>,<https://www.googleapis.com/auth/admin.directory.user>,<https://www.googleapis.com/auth/admin.directory.user.alias.readonly>,<https://www.googleapis.com/auth/drive>,<https://www.googleapis.com/auth/drive.readonly>

c. Click **Authorize**.

The new API client authorization appears in the list:

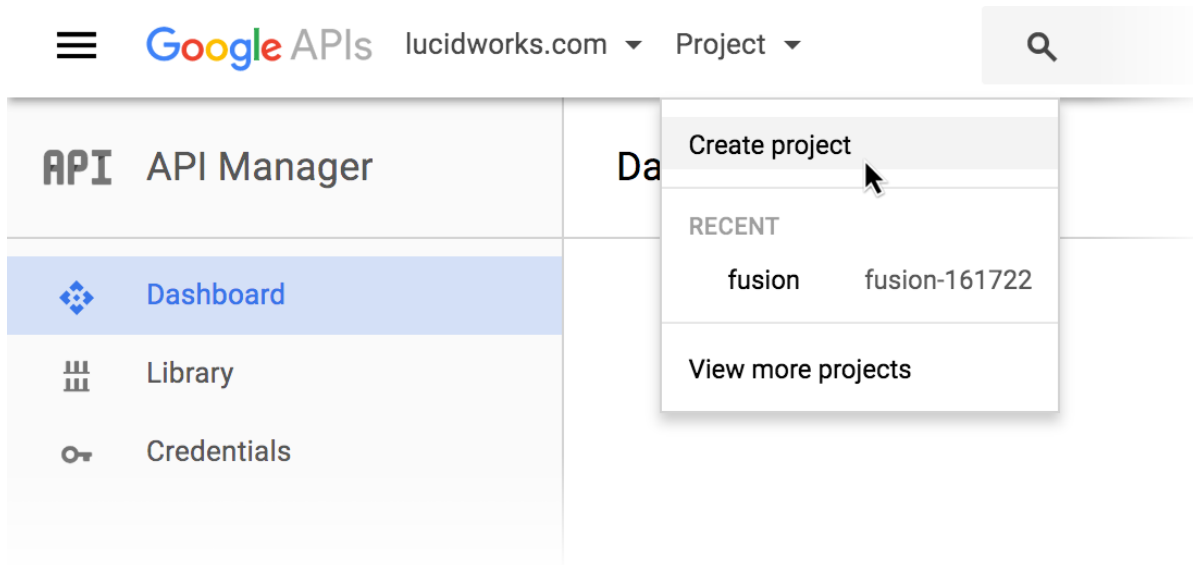
The screenshot shows the Google Admin console interface. At the top, there's a blue header with the Google Admin logo and a search bar. Below the header, the 'Security' section is active. A yellow notification bar says 'Your settings have been saved.' The main content area is titled 'Manage API client access' and includes a brief description. Below this, there's a section for 'Authorized API clients' with a table of registered clients. Each row in the table shows a Client Name, One or More API Scopes, and a 'Remove' button. The first client has ID 112497234733170822870 and scopes for admin.directory.group, admin.directory.group.readonly, admin.directory.user, and admin.directory.user.alias.readonly. The second client has ID 114856940503494649515 and scopes for admin.directory.group, admin.directory.group.readonly, admin.directory.user, and admin.directory.user.alias.readonly. The third client has ID 115023438108627703333 and scopes for admin.directory.group, admin.directory.group.readonly, admin.directory.user, and admin.directory.user.alias.readonly.

Authentication for access to per-user documents

These instructions allow you to configure Google to allow Fusion to crawl a specific user's Google Drive, including documents that other users have shared with them.

How to configure authentication for access to per-user documents

- Log in to Google as a user with *admin-level access rights*.
- Go to <https://console.developers.google.com/>.
- Create a Google project for Fusion:
- In the upper left, open the **Project** menu and select **Create Project**:



5. Enter a new project name, such as "fusion".
6. Click **Create**.
7. Create the client ID and client secret:
8. In the new project, click **Enable API**.
9. Under "Google Apps APIs", click **Drive API**.
10. Click **Enable**.

Google may prompt you to create credentials, if this is the first time you've enabled this API.

11. Click **Credentials**, then **Create Credentials > OAuth client ID**.
12. Select **Web application**.
13. Enter a name for this Web application, such as "Fusion search".
14. In the **Authorized Javascript origins** field, enter "https://developers.google.com".
15. In the **Authorized redirect URIs** field, enter "https://developers.google.com/oauthplayground" and press Return on your keyboard.
16. In the **Authorized redirect URIs** field, enter "http://<fusion-host>:8764/oauth-redirect", specifying the hostname of your Fusion instance.
17. Click **Create**.

Google displays the new client ID and client secret.

18. Copy the client ID and client secret. Save them in a separate location.
19. Click **OK**.
20. Go to <https://developers.google.com/oauthplayground/>.
21. In the upper right, click the gear icon.

The OAuth 2.0 configuration window opens.

22. Select **Use your own OAuth credentials**.

? ↔ ⚙️

OAuth 2.0 configuration

OAuth flow: Server-side

OAuth endpoints: Google

Authorization endpoint:

Token endpoint:

Access token location: Authorization header w/ Bearer prefix

Access type: Offline

Force prompt: Consent Screen

Use your own OAuth credentials

You will need to list the URL <https://developers.google.com/oauthplayground> as a valid redirect URI in your [Google APIs Console](#)'s project. Then enter the client ID and secret assigned to a web application on your project below:

OAuth Client ID:

OAuth Client secret:


Note: Your credentials will be sent to our server as we need to proxy the request. Your credentials will not be logged.

[Close](#)

23. Enter your client ID and client secret.
24. Click **Close**.
25. Add the credentials to the datasource configuration in the Fusion UI:
26. In the Google Drive datasource configuration panel, enter a string for the **Datasource ID**.
27. Enter the Google client ID and client secret.
28. Click **Get Refresh Token**.

A new browser window opens, and Google prompts you for permission to access the documents:

▼ Fusion search would like to:



By clicking **Allow**, you allow this app and Google to use your information in accordance with their respective terms of service and privacy policies. You can change this and other [Account Permissions](#) at any time.



29. Click **Allow**.

Fusion automatically populates the **Google Drive OAuth Refresh Token** field.

30. In the **Startlinks** field, enter a starting URL to which this user has access.

See below for details about the format for this value.

31. Click **Save**.

6.13.2. Configuration

The **StartLinks** values for this connector must be one of the following:

- **root** to crawl the whole organization or per-user drive.
- A folder ID.

For example, if your folder's URL is <https://drive.google.com/drive/folders/0B1u0p7N096R6MWgma3gwUj4j> then the start link in the connector configuration should be `0B1u0p7N096R6MWgma3gwUj4j`.

- A document ID.

For example, if your document's URL is https://docs.google.com/document/d/10HRr5gD00etzgEL9fsyxryf_AfiKqDQ8cn12YXQ/edit then the start link in the connector configuration should be `10HRr5gD00etzgEL9fsyxryf_AfiKqDQ8cn12YXQ`.

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.addFileMetadata Add file metadata	Set to true to add information about documents found in the filesystem to the document, such as document owner, group, or ACL permissions. type: <code>boolean</code> default value: 'true'
f.fs.RecrawlCollectionName Incremental crawling collection name	The collection name for incremental crawling type: <code>string</code> default value: 'system_google_drive_recrawl'
f.fs.additional_item_filters Additional Item Filters	In https://developers.google.com/drive/v3/web/search-parameters#fn4 there are additional search parameters you can add to filter the files returned by google to be indexed. Example: <code>modifiedTime > '2012-06-04T12:00:00'</code> type: <code>string</code>
f.fs.batchPageSize Incremental crawling batch page size	Incremental crawling batch page size type: <code>integer</code> default value: '100' exclusiveMaximum: false exclusiveMinimum: false maximum: 100 minimum: 1
f.fs.clientID Google Drive OAuth Client ID	Google OAuth Client ID for a registered application with access to the Drive API. type: <code>string</code>
f.fs.clientSecret Google Drive OAuth Client Secret	Google OAuth Client Secret for the registered application. type: <code>string</code>

Property	Description
f.fs.connectTimeout Google Drive Connect timeout (ms)	Determines how long, in milliseconds, a request to the Google Drive API is allowed to take to connect prior to timing out type: <code>integer</code> default value: '20000'
f.fs.indexTrash Index trash	Set to true to index files in users Trash folders type: <code>boolean</code> default value: 'false'
f.fs.mime_type_excludes Mime Type Excludes	A comma-separated list of the Mime types to exclude from this crawl. NOTE: This is only used if the "Mime Type Includes" field is empty. type: <code>string</code>
f.fs.mime_type_includes Mime Type Includes	A comma-separated list of the Mime types to include in this crawl. Includes supercede excludes. type: <code>string</code>
f.fs.readTimeout Google Drive read timeout (ms)	Determines how long, in milliseconds, a request to the Google Drive API is allowed to attempt to read content prior to timing out type: <code>integer</code> default value: '20000'
f.fs.refreshToken Google Drive OAuth Refresh Token	OAuth Refresh Token to allow re-authorization of the connector to the Drive API. type: <code>string</code>
f.fs.serviceAccountEmail Service Account Email	For Service Account configuration only - NOTE: Required only when 'Apply Group Security Filtering' is checked. This service account email must be assigned to your Google project as a Service Actor in the console. It must have ability to list groups for a user, list users, and read google drive content. type: <code>string</code>

Property	Description
f.fs.serviceAccountId Service Account ID	For Service Account configuration - specifies the Service Account ID to use to connect Fusion to Google Drive. type: <i>string</i>
f.fs.serviceAccountPrivateKeyFile Service Account P12 Private Key File	For Service Account configuration - specifies the private key file in P12 private key format. type: <i>string</i>
f.fs.serviceAccountPrivateKeyFilePassword Service Account P12 Private Key Password.	For Service Account configuration - specifies the private key file in P12 private key format. type: <i>string</i>
f.maxSizeBytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: <i>integer</i> default value: '4194304'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of documents to fetch. type: <i>integer</i> default value: '0'

Limit Documents

Property	Description
f.maxSizeBytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: <i>integer</i> default value: '4194304'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of documents to fetch. type: <i>integer</i> default value: '0'

Property	Description
f.addFileMetadata Add file metadata	Set to true to add information about documents found in the filesystem to the document, such as document owner, group, or ACL permissions. type: boolean default value: 'true'
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1'
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: boolean default value: 'true'
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: integer default value: '-1'
indexTrash	type: ``

Security Trimming

Property	Description
enable_security_trimming Enable Security Trimming <i>required</i>	type: object object attributes: \{ }

Crawl Performance

Property	Description
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: integer default value: '1'
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: integer default value: '5'
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: integer default value: '0'
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: integer default value: '5'
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: boolean default value: 'true'

Property	Description
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: boolean default value: 'true'
connectTimeout	type: ``
readTimeout	type: ``
parserRetryCount Max Parser Retries	The maximum number of times the configured parser will try getting content before giving up type: integer default value: '0' exclusiveMaximum: true exclusiveMinimum: false maximum: 5 minimum: 0

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: boolean default value: 'false'
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: boolean default value: 'false'

Property	Description
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: <code>string</code>
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content)\{\}' function, but can use any combination of document fields. The function must return a string. type: <code>string</code>

Recrawl Rules

Property	Description
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawl db. type: <code>boolean</code> default value: 'false'
batch_incremental_crawling Batch Incremental crawling	Batch Incremental crawling type: <code>boolean</code> default value: 'true'
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: <code>boolean</code> default value: 'false'
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: <code>boolean</code> default value: 'false'

Property	Description
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: <code>integer</code> default value: <code>'-1'</code>
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: <code>array of string</code>
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: <code>array of string</code>
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: <code>string</code>
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: <code>boolean</code> default value: <code>'false'</code>

Crawl History

Property	Description
retainOutlinks Retain links in the crawldb	Set to true for links found during fetching to be stored in the crawldb. This increases precision in certain recrawl scenarios, but requires more memory and disk space. type: <code>boolean</code> default value: <code>'false'</code>

Property	Description
aliasExpiration Alias expiration	The number of crawls after which an alias will expire. The default is 1 crawl. type: integer default value: '1'
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: string default value: 'in-memory' enum: \{ in-memory on-disk }
indexCrawlDBToSolr Index crawl database to Solr	EXPERIMENTAL: Set to true to index the crawl-database into a 'crawl_db_' collection in Solr. type: boolean default value: 'false'

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" } { "operation" ⇒ "move", "source" ⇒ "driveMimeType", "target" ⇒ "driveMimeType_s" }</code></p> <p>object attributes: <code>\{</code></p> <ul style="list-style-type: none"> <code>operation</code> : <code>\{</code> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: <code>\{ copy move delete set add keep }</code> <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.14. HDFS Connector and Datasource Configuration



width=600%

[Hadoop Distributed File System \(HDFS\)](#). It traverses the Hadoop file system as it would a regular Unix filesystem.

This connector can only be used with the default Hadoop shipped with Solr.

See also the Hadoop connector, a connector for HDFS filesystem which uses [MapReduce](#) to distribute the crawl processes. When there is a lot of content to process, the MapReduce-enabled connector will be significantly faster.

6.14.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
add_failed_docs Add failed documents	Set to true to add documents even if they partially fail processing. Failed documents will be added with as much metadata as available, but may not include all expected fields. type: <code>boolean</code> default value: <code>'false'</code>
bounds Crawl bounds	Limits the crawl to a specific directory sub-tree, hostname or domain. type: <code>string</code> default value: <code>'tree'</code> enum: <code>\{ tree host domain none }</code>
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>

Property	Description
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: boolean default value: 'true'
converter Converter	Fully-qualified classname for a custom converter to produce valid SolrInputDocuments extracted from Hadoop Sequence or MapReduce files. type: string
crawl_depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1' exclusiveMinimum: false minimum: -1
crawl_item_timeout Fetch timeout	Time in milliseconds to fetch any individual document. type: integer default value: '600000' exclusiveMinimum: true minimum: 0
exclude_paths Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string

Property	Description
include_extensions Included file extensions	List the file extensions to be fetched. Note: Files with possible matching MIME types but non-matching file extensions will be skipped. Extensions should be listed without periods, using whitespace to separate items (e.g., 'pdf zip'). type: array of string
include_paths Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: array of string
index_directories Index directories	Set to true to add directories to the index as documents. If set to false, directories will not be added to the index, but they will still be traversed for documents. type: boolean default value: 'false'
kerberos_keytab Kerberos keytab	Full path to the Kerberos keytab file. type: string
kerberos_user Kerberos principal	Kerberos principal name, i.e., ' username@YOUR-REALM.COM '. type: string
max_bytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: integer default value: '10485760' exclusiveMinimum: false minimum: -1

Property	Description
max_docs Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1' exclusiveMinimum: false minimum: -1
max_threads Fetch threads	The maximum number of threads to use for fetching data. Note: Each thread will create a new connection to the repository, which may make overall throughput faster, but this also requires more system resources, including CPU and memory. type: integer default value: '1'
maximum_connections Maximum fetch connections	Maximum number of concurrent connections to the filesystem. A large number of documents could cause a large number of simultaneous connections to the repository and lead to errors or degraded performance. In some cases, reducing this number may help performance issues. type: integer default value: '1000'
url Start link <i>required</i>	A starting URI for this datasource. The URI must be fully-qualified, and include the protocol, host, port and path, as appropriate. type: string minLength: 1 pattern: ..
verify_access Validate access	Set to true to require successful connection to the filesystem before saving this datasource. type: boolean default value: 'true'

Initial Mappings	Description
unmapped Unmapped Fields	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }</p> <pre> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
inlinks Process Inlinks?	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

description

Description

Optional description for datasource instance.

id

Datasource ID

Unique name for datasource instance.

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.15. Hadoop Connector and Datasource Configuration



width=600%

The Hadoop Connector is a MapReduce-enabled crawler which leverages the scaling qualities of [Apache Hadoop](#).

A Hadoop connector creates a series of MapReduce-enabled jobs which index raw content into Fusion. The Hadoop connector can be run using on of the following Apache Hadoop distributions:

- [Apache Hadoop](#) v2.x
- [Cloudera CDH](#) v4.x, v5.x
- [Hortonworks Data Platform](#) v2.x
- [MapR](#) v4.x, v5.x
- [Pivotal HD](#) v1.1

The Hadoop connector name is `lucid.hadoop` and for each Hadoop distribution has its own plugin type. All plugin types take the same set of configuration properties.

There is a non-MapReduce enabled connector for HDFS filesystem; see [page HDFS Connector and Datasource Configuration](#) for details.

The Hadoop crawlers take full advantage of the scaling abilities of the MapReduce architecture and will use all of the nodes available in the cluster just like any other MapReduce job. This has significant ramifications for performance since it is designed to move a lot of content, in parallel, as fast as possible (depending on the system's capabilities), from its raw state to the Fusion index. The Hadoop crawlers work in three stage:

1. Create one or more SequenceFiles from the raw content. This can be done in one of two ways:
2. If the source files are available in a shared Hadoop filesystem, prepare a list of source files and their locations as a SequenceFile. The raw contents of each file are not processed until step 2.
3. If the source files are not available, prepare a list of source files and the raw content, stored as a BehemothDocument. This process is currently done sequentially and can take a significant amount of time if there is a large number of documents and/or if they are very large.
4. Run a MapReduce job to extract text and metadata from the raw content using Apache Tika. This is similar to the Fusion approach of extracting content from crawled documents, except it is done with MapReduce.
5. Run a MapReduce job to send the extracted content from HDFS to the index pipeline for further processing.

Note:

The first step of the crawl process converts the input content into a SequenceFile. In order to do this, the entire contents of that file must be read into memory so that it can be written out in the SequenceFile. Thus, you should be careful to ensure that the system does not load into memory a file that is larger than the Java heap size of the process. In certain cases, Behemoth can work with existing files such as SequenceFiles to convert them to Behemoth SequenceFiles. Contact Lucidworks for possible alternative approaches.

The processing approach is currently "all or nothing" when it comes to ingesting the raw content and all 3 steps must be completed each time, regardless of whether the raw content hasn't changed. Future versions may allow the crawler to

restart from the SequenceFile conversion process. In the meantime, incremental crawling is not supported for this connector.

6.15.1. Hadoop Installation and Configuration

The Connector services must be able to access the Hadoop client in file `$HADOOP_HOME/bin/hadoop`, so it must either be installed on one of the nodes of the Hadoop cluster (such as the `nameNode`), or a client supported by your specific distribution must be installed on the same server as the Connectors. The Hadoop client must be configured properly to access the Hadoop cluster so the crawler is able to access the Hadoop cluster for content processing.

Please note, instructions for setting up any of the supported Hadoop distributions is beyond the scope of this document. We recommend reading one of the many tutorials found online or one of the books on Hadoop.

This connector writes to the `hadoop.tmp.dir` and the `/tmp` directory in HDFS, so Fusion should be started by a user who has read/write permissions for both.

Permission Issues

Using any flavor of Hadoop, you will need to be aware of the way Hadoop and systems based on Hadoop (such as CDH, MapR, etc.) handle permissions for services that communicate with other nodes.

Hadoop services execute under specific user credentials: a quadruplet consisting of user name, group name, numeric user id, numeric group id. Installations that follow the manual usually use user 'mapr' and group 'mapr' (or similar), with numeric ids assigned by the operating system (e.g., `uid=1000, gid=20`). When the system is configured to enforce user permissions (which is the default in some systems), any client that connects to Hadoop services has to use a quadruplet that exists on the server. This means that ALL values in this quadruplet must be equal between the client and the server, i.e., an account with the same user, group, uid, and gid must exist on both client and server machines.

When a client attempts to access a resource on Hadoop filesystems (or the JobTracker, which also uses this authentication method) it sends its credentials, which are looked up on the server, and if an exactly matching record is found then those local permissions will be used to determine read/write access. If no such account is found then the user is treated as "other" in the sense of the permission model.

This means that the crawlers for the HDFS data source should be able to crawl Hadoop or MapR filesystems without any authentication, as long as there is a read (and execute for directories) access for "other" users granted on the target resources. Authenticated users will be able to access resources owned by their equivalent account.

However, the Hadoop crawling described on this page require write access to a `/tmp` directory to use as a working directory. In many cases, this directory does not exist, or if it does, it doesn't have write access to "other" (not authenticated) users. Therefore users of these data sources should make sure that there is a `/tmp` directory on the target filesystem that is writable using their local user credentials, be it a recognized user, group, or "other". If a local user is recognized by the server then it's enough to create a `/tmp` directory that is owned by that user. If there is no such user, then the `/tmp` directory must be modified to have write permissions for "other" users. The working directory can be modified to be another directory that can be used for temporary working storage that has the correct permissions.

Configuration for a Kerberos Hadoop Cluster

Kerberos is a system that provides authenticated access for users and services on a network. Instead of sending passwords in plaintext over the network, encrypted passwords are used to generate time-sensitive tickets which are used for authentication. Kerberos uses symmetric-key cryptography and a trusted third party called a Key Distribution Center (KDC) to authenticate users to a suite of network services. When a user authenticates to the KDC, the KDC sends a set of credentials (a ticket) specific to that session back to the user's machine.

To work with a Kerberized Hadoop cluster you must have a set of credentials. These are generated by running the "kinit" program. The datasource can be configured to run this program, in which case, the following information must be specified: the full path to the program, the Kerberos principal name, the location of a keytab file and the name of the file in which to store the ticket.

For more information on Kerberized Hadoop Clusters, see http://www.cloudera.com/content/cloudera/en/documentation/archives/cloudera-manager-4/v4-5-2/Configuring-Hadoop-Security-with-Cloudera-Manager/cmeechs_topic_4_5.html.

6.15.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>
fusion_batchsize Batch Size	Fusion Client Batch Size type: <code>integer</code> default value: <code>'500'</code> exclusiveMinimum: true minimum: 1
fusion_buffer_timeoutms Timeout (ms)	Fusion Client Timeout (ms). type: <code>integer</code> default value: <code>'1000'</code> exclusiveMinimum: true minimum: 1
fusion_endpoints List of Fusion Endpoints <i>required</i>	type: <code>array of string</code> minimum number of items (<code>minItems</code>): 1 default value: <code>[http://localhost:8764,]</code>

Property	Description
fusion_fail_on_error Fail on Error	Fusion Client Fail on Error type: <code>boolean</code> default value: 'false'
fusion_login_app_name Config App Name	Login Config App Name FusionClient by default. type: <code>string</code> default value: 'FusionClient'
fusion_login_config Login Config	The file path of Login Configuration for Fusion kerberized, it must be placed in every mapper/reduce node. type: <code>string</code>
fusion_password Password	Fusion client User's password, leave empty if kerberos is use. type: <code>string</code>
fusion_realm Fusion client's Authentication	Fusion's Realm, If 'native' is selected the password is mandatory. If 'kerberos' is selected the Login Configuration is mandatory. type: <code>string</code> default value: 'NATIVE' enum: \{ NATIVE KERBEROS }
fusion_user User/Principal <i>required</i>	Fusion client's User or Principal if Kerberos is chosen. type: <code>string</code>
hadoop_home Hadoop home <i>required</i>	Path to the Hadoop home directory where \$HADOOP_HOME/bin/hadoop can be found. The connector requires access to either a full Hadoop installation, or a Hadoop client provided by your Hadoop distribution that has been configured to access the Hadoop installation. type: <code>string</code> minLength: 1

Property	Description
hadoop_input Input source <i>required</i>	Hadoop input source file/directory type: <i>string</i> minLength: 1
hadoop_mapper Mapper <i>required</i>	Hadoop Ingest Mapper type: <i>string</i> default value: 'CSV' enum: \{ CSV DIRECTORY GROK REGEX SEQUENCE_FILE SOLR_XML WARC ZIP \}
job_jar Job Jar <i>required</i>	Path and name of the Hadoop job jar. Unless you are using a custom job jar, the default provided by Fusion is preferred. type: <i>string</i> default value: 'lucidworks-hadoop-job-2.2.7.jar' minLength: 1
kinit_cache 'kerberos' cache	Full path of 'kerberos' cache. If this path does not exist, it will be created. type: <i>string</i>
kinit_cmd 'kinit' command	Full path to the 'kinit' binary. type: <i>string</i> default value: 'kinit'
kinit_keytab 'kerberos' keytab	Full path to the Kerberos keytab file. type: <i>string</i>
kinit_principal 'kerberos' principal	Kerberos principal name, i.e., username@YOUR-REALM.COM type: <i>string</i>

Property	Description
<p>mapper_args</p> <p>Job Jar arguments</p>	<p>Parameters for the Hadoop job.</p> <p>type: array of object</p> <p>object attributes: \{ arg_name : \{ display name: name type: string enum: \{ csvFieldMapping csvDelimiter csvFirstLineComment csvStrategy idField add.subdirectories grok.uri grok.config.path grok.additional.patterns com.lucidworks.hadoop.ingest.RegexIngestMapper.regex com.lucidworks.hadoop.ingest.RegexIngestMapper.groups _to_fields com.lucidworks.hadoop.ingest.RegexIngestMapper.match }</p> <pre> } + 'arg_value' : \{ + display name: value + type: 'string' + } + } </pre>
<p>reducers</p> <p>Number of Reducers</p>	<p>(Expert) Depending on the OutputFormat and your system resources, you may wish to have Hadoop do a reduce step first so as to not open too many connections to the output resource</p> <p>type: integer</p> <p>default value: '0'</p> <p>exclusiveMinimum: false</p> <p>minimum: 0</p>
<p>run_kinit</p> <p>Run 'kinit'</p>	<p>If your Hadoop installation requires job requests to authenticate with Kerberos, this option will allow Fusion to run 'kinit' to get a valid ticket.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } } 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } }</p>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
<p>unmapped</p> <p>Unmapped Fields</p>	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

ConnectorDb Configuration

Property	Description
<p>aliases</p> <p>Process Aliases?</p>	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
<p>inlinks</p> <p>Process Inlinks?</p>	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
description Description	Optional description for datasource instance.
id Datasource ID	Unique name for datasource instance.
pipeline Pipeline ID	The index pipeline used to process documents.

6.16. JDBC Connector and Datasource Configuration

The JDBC connector fetches documents from a relational database via SQL queries. Under the hood, this connector implements the Solr [DataImportHandler \(DIH\)](#) plugin.

6.16.1. SQL queries for document retrieval

The JDBC connector is configured with the SQL statement used to retrieve a resultset from the database. Each row in the result set is treated as a Solr document. This statement is specified as the value of the required property `sql_select_statement`.

The column names will be used as Solr document field names; use the SQL "AS" keyword to rename column names as needed. Column and table names should be treated as if they are case-insensitive, even though some databases allow use of mixed case names. All Solr documents must have a unique key, which is the Fusion "id" field. Therefore, the results set must contain a column "id" which can be used as a unique key for the resulting document.

```
SELECT customer_id AS id, * from customers
```

Delta queries

Delta queries provide incremental updates to the contents of a collection by indexing only those records in the database which have been changed since the database was last indexed by this connector. The SQL statement is specified as the value of the property `delta_sql_query`.

Delta queries select only primary key values, therefore, the query must use the primary key and it must also have a "WHERE" clause which specifies a "last_modified" condition as follows:

```
SELECT customer_id AS id from customers WHERE last_modified > $
```

The dollar-sign character '\$' is required; it holds the last successful import time from the database.

Nested queries

Nested queries are used to index information which is stored in the database across a series of tables where there is a one-to-many or many-to-many relationship between them. This statement is specified as the value of the property `nested_queries`.

A nested query is used in conjunction with the SQL query specified by the `sql_select_statement` statement. The dollar-sign character '\$' specifies the primary key in the resultset retrieved by the `sql_select_statement` statement.

The following example shows the pair of query, nested query statements used to index list of tags assigned to documents:


```
SELECT id FROM document
SELECT tag FROM tag INNER JOIN document_tag ON document_tag.tag_id=tag.id WHERE document_tag.doc_id=$
```

6.16.2. Uploading a JDBC driver

Fusion stores JDBC drivers in the blob store. You can upload a driver using the Fusion UI or the Blob Store API.

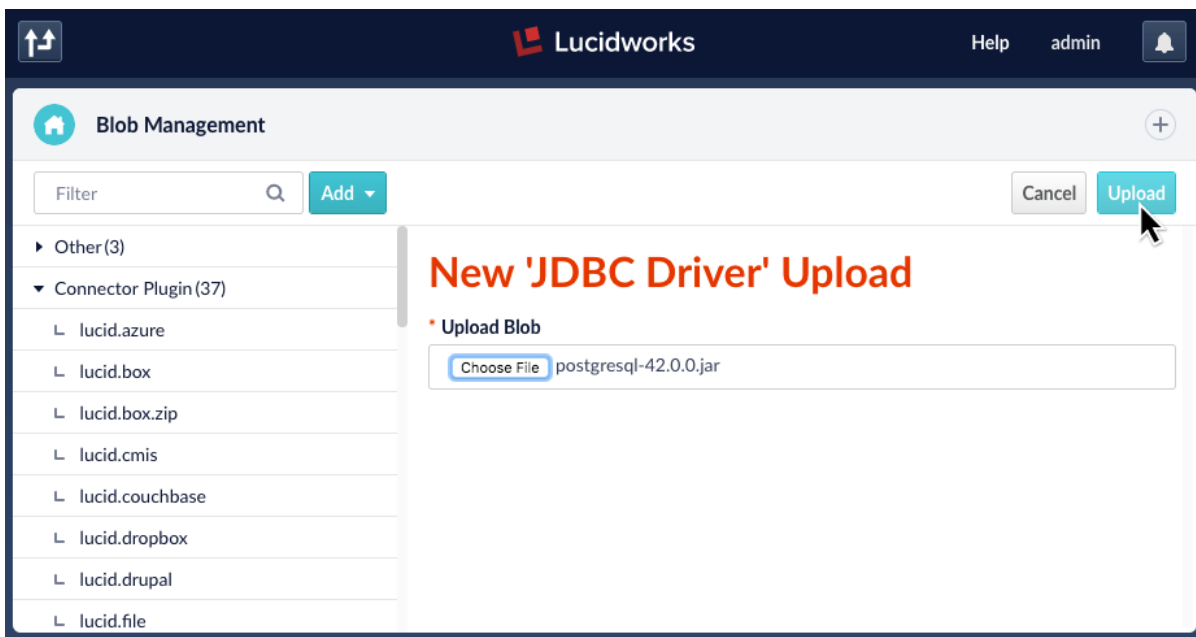
To upload a JDBC driver using the Fusion UI

1. In the Fusion UI, navigate to **DevOps > Blob Management**.
2. Click **Add**.
3. Select **JDBC Driver**.




The "New 'JDBC Driver' Upload" panel appears.

4. Click **Choose File** and select the .jar file from your file system.



5. Click **Upload**.

The new driver's blob manifest appears.



From this screen you can also delete or replace the driver.

To install a JDBC driver using the API

1. Upload the .jar file to Fusion's blob store using the `/blobs/{id}` endpoint.

Specify an arbitrary blob ID, and a `resourceType` value of `plugin:connector`, as in this example:

```
curl -u user:pass -H "content-type:application/java-archive" -H "content-length:707261" -X PUT --data-binary @postgresql-42.0.0.jar http://localhost:8764/api/apollo/blobs/mydriver?resourceType=driver:jdbc
```

Success response:

```
{
  "name" : "mydriver",
  "contentType" : "application/java-archive",
  "size" : 707261,
  "modifiedTime" : "2017-06-09T19:00:48.919Z",
  "version" : 0,
  "md5" : "c67163ca764bfe632f28229c142131b5",
  "metadata" : {
    "subtype" : "driver:jdbc",
    "drivers" : "org.postgresql.Driver",
    "resourceType" : "driver:jdbc"
  }
}
```

Fusion automatically publishes the event to the cluster, and the listeners perform the driver installation process on each node.

Tip	If the blob ID is identical to an existing one, the old driver will be uninstalled and the new driver will be installed in its place. To get the list of existing blob IDs, run: <code>curl -u user:_password_ localhost:8764/api/apollo/blobs</code>
-----	---

2. To verify the uploaded driver, run:

```
curl -u user:pass http://localhost:8764/api/apollo/blobs/<id>/manifest
```

Where the <id> is the name specified during upload, such as "mydriver" above. A success response looks like this:

```
{
  "name" : "mydriver",
  "contentType" : "application/java-archive",
  "size" : 707261,
  "modifiedTime" : "2017-06-09T19:05:17.897Z",
  "version" : 1569755095787110400,
  "md5" : "c67163ca764bfe632f28229c142131b5",
  "metadata" : {
    "subtype" : "driver:jdbc",
    "drivers" : "org.postgresql.Driver",
    "resourceType" : "driver:jdbc"
  }
}
```

6.16.3. Indexing binary data

The JDBC connector in Fusion does not automatically discover and index binary data you may have stored in your database (such as PDF files). However, you can configure Fusion to recognize and extract binary data correctly by modifying the datasource configuration file. This file is created when the datasource is first run, and then it is created in `fusion/3.1.x/data/connectors/lucid.jdbc/datasources/ datasourceID/conf`. The name of the file will include the name of the datasource, as in `dataconfig_datasourceName.xml`. If you are familiar with Solr's DIH, you will recognize this as a standard `dataconfig.xml` file.

Follow these steps to modify the configuration file:

1. Add a `name` attribute for the database containing your binary data to the `dataSource` entry.
2. Set the `convertType` attribute for the `dataSource` to `false`. This prevents Fusion from treating binary data as strings.
3. Add a `FieldStreamDataSource` to stream the binary data to the Tika entity processor.
4. Specify the `dataSource` name in the `root` entity.
5. Add an entity for your `FieldStreamDataSource` using the `TikaEntityProcessor` to take the binary data from the `FieldStreamDataSource`, parse it, and specify a field for storing the processed data.
6. Reload the Solr core to apply your configuration changes.

6.16.4. Troubleshooting

When using the JDBC connector, it is recommended that you work closely with your database administrator to formulate efficient and robust queries.

One source of possible problems is the driver being used. In some cases, indexing may fail due to problems with the driver, in particular older versions of Oracle's JDBC driver. If you have checked that your connection information is correct and your database is allowing the connection, you may want to research if there are any known bugs with the driver you are using.

With Oracle databases, note that column names not enclosed in double-quotes are converted to upper-case, but Solr field names are case sensitive. If your column-to-field mapping is not happening properly, check your SQL statement for any lower-case names not enclosed in double-quotes.

Dates can also be problematic. Solr has a different date format than many relational databases. If you want date and time fields to be indexed properly, you may need to convert database dates into the proper format using date/string convert functions. In Oracle this is the TO_CHAR function; in Microsoft SQL, this is the DATEPART function.

In MySQL databases, dates are allowed to be 0-strings, such as 0000-00-00, which is not acceptable to JDBC. If you have legacy date data you may need to add the query parameter "zeroDateTimeBehavior=convertToNull" to your JDBC request string, as in `jdbc:mysql://localhost/myDatabase?zeroDateTimeBehavior=convertToNull`. This will convert the zero-string dates to null values that can be added to the index.

Finally, database timeouts are another problematic area. There are several possible solutions to this, from increasing the timeout in the JDBC request (with "netTimeoutForStreamingResults"), altering the SQL statement to page the results, or dumping the records to CSV and indexing them with another connector.

6.16.5. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
clean_in_full_import_mode Clean in full import mode	Clean old records when doing a full import. This will remove all records from the index before re-indexing. type: <code>boolean</code> default value: <code>'true'</code>
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>

Property	Description
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: boolean default value: 'true'
delta_import_query Delta Import SQL Query	A SQL statement to retrieve the delta records. If not defined, the original SQL Statement will be used to retrieve delta records, which may be error prone in the case of complex statements. If that occurs, you can define a secondary SQL statement that will only be used for delta imports. type: string
delta_sql_query Delta SQL Query	A SQL statement to select the delta (records added or changed) since the last run of the datasource. For example, 'SELECT id FROM customers WHERE last_modified > \$'. The \$ indicates the date and time of the last import. type: string
driver JDBC Driver <i>required</i>	The class name of the JDBC driver to use to connect to the database. Only JDBC4 drivers will appear in this list. If a JDBC3 driver has been uploaded but does not appear in the list, it may be possible to manually enter the class name. type: string minLength: 1
fetch_size JDBC fetch size	The number of documents to retrieve per batch. type: integer
max_docs Max Documents	The maximum number of documents to crawl. Use -1 to index all documents found. type: integer default value: '-1'

Property	Description
nested_queries Nested queries	A nested query to join data from multiple tables. The nested query will be used with the SQL Statement and must include the primary key with the \$ character. For example, 'SELECT tag FROM tag WHERE document_tag.doc_id=\$'. type: array of string
password Password <i>required</i>	The password of the account used for authentication and data access. type: string
primary_key Primary Key	The column name of the primary key for the table. type: string
sql_select_statement SQL Statement <i>required</i>	A SQL SELECT statement to choose the records to be retrieved. type: string
url URL <i>required</i>	A URL to the database, starting with jdbc and vendor name, e.g., 'jdbc:mysql://localhost/test' type: string minLength: 1
username Username <i>required</i>	The username of a database account used for authentication and data access. type: string
verify_access Validate access	Set to true to require successful connection to the filesystem before saving this datasource. type: boolean default value: 'true'

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
<p>unmapped</p> <p>Unmapped Fields</p>	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + }</pre>

ConnectorDb Configuration

Property	Description
<p>aliases</p> <p>Process Aliases?</p>	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
<p>inlinks</p> <p>Process Inlinks?</p>	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
description Description	Optional description for datasource instance.
id Datasource ID	Unique name for datasource instance.
pipeline Pipeline ID	The index pipeline used to process documents.

6.17. JIRA Connector and Datasource Configuration



width=600%

The JIRA connector retrieves data from a instance of [Atlassian's JIRA](#) issue tracking system.

It uses the JIRA REST API to retrieve the following JIRA elements:

- Projects
- Issues
- Comments
- Worklogs
- Attachments

JIRA access to projects and issues may be restricted to certain users. Access of restricted information requires a JIRA username and password.

The JIRA connector first requests a list of all projects, and for each project, if finds all issues. For each issue, the information on summary, priority, assignee, etc. is retrieved. Worklogs, comments, and attachments are treated as new entries (issue links) and thus are indexed as a new document, not as a component of the issue itself.

6.17.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.cacheSize Cache size (number of entries)	Number of entries to cache when making REST requests. type: <code>integer</code> default value: '2000'
f.index_attachments Index Attachments	Select to index attachments on issues. type: <code>boolean</code> default value: 'true'

Property	Description
f.index_comments Index Comments	Select to index comments on issues. type: boolean default value: 'true'
f.index_issues Index Documents	Select to index issues. type: boolean default value: 'true'
f.index_projects Index Projects	Select to index project-level information. type: boolean
f.index_worklogs Index Worklogs	Select to index worklogs on issues. type: boolean default value: 'false'
f.issueFieldsToIndex Issue fields to index	A list of fields to be indexed for JIRA issues type: array of string default value: 'assigneeissuetypepriorityprojectreporterstatussummaryupdated'
f.jira_password JIRA password	Password for the administrator user, if required. type: string
f.jira_username JIRA username	Username for accessing the JIRA REST API. This user must be a project-level administrator to index single projects, or a global or system-level administrator to index all projects. Leave empty if your JIRA system does not require authentication to see projects, issues, comments and attachments. type: string

Property	Description
f.maxResults Max results per page	The number of items to retrieve with each request. type: <code>integer</code> default value: '100'
f.maxSizeBytes Maximum file size (bytes)	Maximum size, in bytes, of a project description, issue description, worklog, issue comment, or attachment. type: <code>integer</code> default value: '32000'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of a project description, issue description, worklog, issue comment, or attachment. type: <code>integer</code> default value: '0'
f.retryDelay Retry delay(in ms)	The milliseconds to wait until repeat a failed request. type: <code>integer</code> default value: '100' exclusiveMinimum: false minimum: 100
f.stopRetry Stop retry after a given delay	The minutes to wait until stop retry. type: <code>integer</code> default value: '5' exclusiveMinimum: false minimum: 1
f.timeoutMS Connection timeout (ms)	Time in milliseconds to wait for server response. type: <code>integer</code> default value: '10000'

Security Trimming

Property	Description
enable_security_trimming Enable Security Trimming	type: object

Limit Documents

Property	Description
f.maxSizeBytes Maximum file size (bytes)	Maximum size, in bytes, of a project description, issue description, worklog, issue comment, or attachment. type: integer default value: '32000'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of a project description, issue description, worklog, issue comment, or attachment. type: integer default value: '0'
f.maxResults Max results per page	The number of items to retrieve with each request. type: integer default value: '100'
f.index_projects Index Projects	Select to index project-level information. type: boolean
f.index_issues Index Documents	Select to index issues. type: boolean default value: 'true'
f.index_comments Index Comments	Select to index comments on issues. type: boolean default value: 'true'

Property	Description
f.index_attachments Index Attachments	Select to index attachments on issues. type: boolean default value: 'true'
f.index_worklogs Index Worklogs	Select to index worklogs on issues. type: boolean default value: 'false'
f.["assignee", "issuetype", "priority", "project", "reporter", "status", "summary", "updated"]	type: ``
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1'
includeExtensions Included file extensions	File extensions to be fetched. This will limit this datasource to only these file extensions. type: array of string
includeRegexes Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: array of string
excludeExtensions Excluded file extensions	File extensions that should not to be fetched. This will limit this datasource to all extensions except this list. type: array of string

Property	Description
excludeRegexes Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: boolean default value: 'true'
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: integer default value: '-1'

Crawl Performance

Property	Description
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: integer default value: '1'
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: integer default value: '5'
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: integer default value: '0'

Property	Description
fetchDelayMSPerHost Fetch delay per host	If true, the 'Fetch delay (ms)' property will be applied for each host. type: boolean default value: 'false'
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: integer default value: '5'
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: boolean default value: 'true'
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: boolean default value: 'true'

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: boolean default value: 'false'

Property	Description
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: <code>boolean</code> default value: 'false'
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: <code>string</code>
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content)\{' function, but can use any combination of document fields. The function must return a string. type: <code>string</code>

Recrawl Rules

Property	Description
f.timeoutMS Connection timeout (ms)	Time in milliseconds to wait for server response. type: <code>integer</code> default value: '10000'
f.cacheSize Cache size (number of entries)	Number of entries to cache when making REST requests. type: <code>integer</code> default value: '2000'
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: <code>boolean</code> default value: 'true'

Property	Description
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: boolean default value: 'false'
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: boolean default value: 'false'
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: integer default value: '-1'
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: array of string
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: array of string
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: string
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: boolean default value: 'false'

Crawl History

Property	Description
retainOutlinks Retain links in the crawldb	Set to true for links found during fetching to be stored in the crawldb. This increases precision in certain recrawl scenarios, but requires more memory and disk space. type: boolean default value: 'false'
aliasExpiration Alias expiration	The number of crawls after which an alias will expire. The default is 1 crawl. type: integer default value: '1'
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: string default value: 'in-memory' enum: \{ in-memory on-disk }
indexCrawlDBToSolr Index crawl database to Solr	EXPERIMENTAL: Set to true to index the crawl-database into a 'crawldb_' collection in Solr. type: boolean default value: 'false'
reevaluateCrawlDBOnStart Reevaluate crawldb on start?	Reevaluate existing crawldb entries for legality on startup? type: boolean default value: 'false'

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>\{</code> operation : <code>\{</code> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: <code>\{ copy move delete set add keep }</code> <code>}</code> <code>\{</code> source` _ (required)_ : <code>\{</code> + display name: Source Field + type: `string` + description : The name of the field to be mapped. + <code>}</code> + target` : <code>\{</code> + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + <code>}</code> + <code>}</code></p>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
      mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
      mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.18. Javascript Connector and Datasource Configuration



width=600%

The Javascript connector allows users to write ad-hoc document retrieval routines to fetch content from filesystems and websites.

It provides a property `f.script` which is a JavaScript program that is compiled by the JDK. This program returns a content item which is handed off to the fetcher.

The script engine works exactly the same as the JavaScript Index and JavaScript Query Pipeline stages. The JavaScript program must be standard [ECMAScript](#).

You can use any Java class available to the connectors JDK classloader to manipulate that object within a function. As in Java, to access Java classes by their simple names instead of their fully specified class names, e.g. to be able to write `String` instead of `java.lang.String`, these classes must be imported. The `java.lang` package is not imported by default, because its classes would conflict with `Object`, `Boolean`, `Math`, and other built-in JavaScript objects. To import a Java class, use the `JavaImporter` object and the `with` statement, which limits the scope of the imported Java packages and classes.

```
var imports = new JavaImporter(java.lang.String);
...
with (imports) {
    var name = new String("foo"); ...
}
```

For global variables, you can reference these objects using the `Java.type` API extension. See this tutorial for details: <http://winterbe.com/posts/2014/04/05/java8-nashorn-tutorial/>

6.18.1. The JavaScript Program

The Javascript context provides the following variables:

- `id`, type `java.lang.String` - the id of the object to fetch. This is almost always the URI of the datasource to connect to and fetch content.
- `lastModified`, type `long` - the time since the epoch from which the item was last touched.
- `signature`, type `java.lang.String` - an optional string meant to be used to compare versions of the ID being fetched, e.g. an ETag in a web-crawl.
- `content`, type `crawler.common.MutableObject` - a Content object that can be modified and returned, for fine grained control over the return. See the section on return types below.
- `_fetcher`, type `Fetcher` - the current `Fetcher` instance (usually type `JavascriptFetcher`), used to interact with the `Fetcher`, including getting a `WebFetcher` instance using `_fetcher.getWebFetcher()`
- `_context`, type `java.util.Map` - a map used to store data to persist across calls to `fetch()`, e.g. an instance of `WebFetcher` obtained using `_fetcher.getWebFetcher()`.

The program must return one of the following kinds of objects:

- String—A string object. This will be converted to UTF-8 bytes and added as the raw content on a `common.crawler.Content` object and returned from the `fetch()` method
- `byte []`—A byte array. This array will be set on a `common.crawler.Content` object and returned from the `fetch()` method
- `common.crawler.MutableContent`: If you wish to have complete control over the return from `fetch()`, make changes to the content object provided in the Context and return it. **DO NOT CREATE A NEW OBJECT.**
- An array of Objects. These will be converted to Embedded Content (the Fetcher will return a parent Content object that has a "Container" `discardMessage`. The Embedded Content on that container will consist of calling `toString()` on the objects in the array. Thus, it is best if the array is simply
- A JavaScript Map. The map will be converted to fields on the Content item returned

If the JavaScript script is implemented as a function, the return statement must return one of the above types. If the script is not function-based, than the last line in the script must evaluate to one of these object types.

6.18.2. Examples

Return content as a `java.lang.String`

```
var str = new java.lang.String("Java");
str;
```

Return content as a byte array

```
var bytes = new java.lang.String("Java");
bytes.getBytes('UTF-8');
```

Return content as a JavaScript array

```
var strings = ["hi", "bye"];
strings;
```

Return content as a JavaScript map

```
var map = {"hi": "bye", "bye": "hi", "number":1};
map;
```

Leverage the Fetcher

```

var webFetcher = _context.get("webFetcher");
if (null == webFetcher) {
    webFetcher = _fetcher.getWebFetcher();
    // it's possible to pass config options to getWebFetcher() as a map as well, e.g.:
    // _fetcher.getWebFetcher({"f.discardLinkURLQueries" : false });
    _context.put("webFetcher", webFetcher);
}
var webContent = webFetcher.fetch(id, lastModified, signature);
var jsoupDoc = webContent.getDocument();
if (null != jsoupDoc) {
    // modify the Jsoup document or web-content as-needed here, adding new links, removing sections etc.
    // ...
    // ...
    webContent.setRawContent(jsoupDoc.toString().getBytes("UTF-8"));
}
webContent;

```

6.18.3. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.script Script	JavaScript program to fetch documents. type: <code>string</code>

Link Discovery

Property	Description
restrictToTree Restrict to sub-directories and child pages	If true, only documents found in a tree below the start links will be fetched. type: <code>boolean</code> default value: 'false'
restrictToTreeAllowSubdomains Allow sub-domains in restrictToTree	If true, any sub-domain will be allowed, even if the crawl is restricted to the tree of items found below the start links. type: <code>boolean</code> default value: 'false'

Property	Description
restrictToTreeUseHostAndPath Use paths in restrictToTree	If true, the path in start links will be used to restrict items fetched. For example, if the start link is 'http://host.com/US', this option will limit all followed URLs to this path. type: boolean default value: 'false'
restrictToTreeIgnoredHostPrefixes Ignored host prefixes	List of host prefixes to ignore when checking links for restrictToTree link-legality checks. For example, 'www.' can be ignored so links with the same domain are allowed. type: array of string default value: [` www. `]

Limit Documents

Property	Description
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1'
includeExtensions Included file extensions	File extensions to be fetched. This will limit this datasource to only these file extensions. type: array of string
includeRegexes Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: array of string

Property	Description
excludeExtensions Excluded file extensions	File extensions that should not to be fetched. This will limit this datasource to all extensions except this list. type: <code>array of string</code>
excludeRegexes Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: <code>array of string</code>

Crawl Performance

Property	Description
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: <code>integer</code> default value: '1'
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: <code>integer</code> default value: '5'
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: <code>integer</code> default value: '0'
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: <code>integer</code> default value: '5'

Property	Description
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: <code>boolean</code> default value: 'true'
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: <code>boolean</code> default value: 'true'

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: <code>boolean</code> default value: 'false'
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: <code>boolean</code> default value: 'false'
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: <code>string</code>

Property	Description
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content)\{' function, but can use any combination of document fields. The function must return a string. type: <code>string</code>

Recrawl Rules

Property	Description
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: <code>boolean</code> default value: 'true'
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: <code>boolean</code> default value: 'false'
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: <code>boolean</code> default value: 'false'
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: <code>integer</code> default value: '-1'
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: <code>array of string</code>

Property	Description
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: <code>array of string</code>
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: <code>string</code>
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: <code>boolean</code> default value: 'false'
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: <code>boolean</code> default value: 'true'
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: <code>integer</code> default value: '-1'

Crawl History

Property	Description
retainOutlinks Retain links in the crawl db	Set to true for links found during fetching to be stored in the crawl db. This increases precision in certain recrawl scenarios, but requires more memory and disk space. type: <code>boolean</code> default value: 'false'

Property	Description
aliasExpiration Alias expiration	The number of crawls after which an alias will expire. The default is 1 crawl. type: integer default value: '1'
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: string default value: 'in-memory' enum: \{ in-memory on-disk }
indexCrawlDBToSolr Index crawl database to Solr	EXPERIMENTAL: Set to true to index the crawl-database into a 'crawl_db_' collection in Solr. type: boolean default value: 'false'

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>{</code></p> <p>operation : <code>{</code></p> <p>display name: Operation</p> <p>type: string</p> <p>default value: 'copy'</p> <p>description : The type of mapping to perform: move, copy, delete, add, set, or keep.</p> <p>enum: <code>{ copy move delete set add keep }</code></p> <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
pipeline Pipeline ID	The index pipeline used to process documents.

6.19. Jive Connector and Datasource Configuration



width=600%

Retrieve content from a [Jive](#) instance.

Important	If you are upgrading from Fusion 2.0 or earlier, you must remove the old connector, add the new one, and then install the Lucidworks add-on for Jive. See the instructions below.
-----------	---

6.19.1. Updating the Jive connector

Perform all of the steps below as the user that owns the Fusion installation files and processes.

Removing the Old Jive Connector

1. Remove all existing Jive datasources, using either the UI or the REST API. Any datasource that is not a Jive type of datasource can be retained.
 - a. In the UI, go to the Datasources screen for each collection that includes a Jive datasource, and select the X icon on the right side of the row. Confirm the delete action to remove the datasource.
 - b. Using the Connector Datasources API, you can perform a call similar to this one:

```
curl -u user:password -X DELETE http://localhost:8764/api/apollo/connectors/datasources/<id>
```

Be sure to use the correct username and password, and also replace the hostname, port and datasource ID as appropriate.

2. Stop Fusion.
3. On the filesystem, find the file `fusion/3.1.x/apps/connectors/plugins/lucid.anda/connectors.json`. From this file, remove the following line:

```
"jive": "com.lucidworks.connectors.anda.type.JiveAndaType",
```

Be sure to save your changes.

4. In the directory `fusion/3.1.x/apps/connectors/plugins/lucid.anda/lib/`, remove the file `lucid.jive.jar` by deleting it or saving it outside of your Fusion installation directory tree (i.e., to a `/tmp` directory or similar).

Adding the New Jive Connector

1. In the directory `fusion/3.1.x/apps/connectors/plugins/`, make a new directory named `lucid.jive`.
2. Copy the files `connectors.json` and `lucid.jive.jar` to the directory created in the previous step, `fusion/3.1.x/apps/connectors/plugins/lucid.jive`.
3. In order for the UI to display the correct name for the new connector, we must update a datasource mapping file.

In `fusion/3.1.x/apps/jetty/ui/webapps/root/WEB-INF/classes/public/data/datasource-name-map.json`, remove the following line:

```
"lucid.anda.jive" : "Jive",
```

And add the following line (anywhere, as long as it is a new line):

```
"lucid.jive.jive" : "Jive",
```

When adding the new line, be sure to add a comma to the previous line in order to preserve correct JSON. If this file cannot be parsed, you will see odd-looking datasource names in the UI.

4. Start Fusion.

Installing the Jive Add-on

The add-on allows using OAuth for authenticating the crawler to the Jive API. This is a recommended way for Jive users to use the API, as documented by Jive at <https://community.jivesoftware.com/docs/DOC-157031>. The add-on is used to get a client ID and secret that are supplied to the connector.

If the client ID and secret cannot be retrieved or used, then Basic authentication is used as a fallback.

How to install the Jive Add-on

Following the instructions at <https://community.jivesoftware.com/docs/DOC-141123>, use the Jive UI to upload the add-on located in your Fusion distribution: `fusion/3.1.x/apps/connectors/resources/lucid.jive/lucidworks-jive-addon.zip`

6.19.2. Security Trimming of Results

If "Enable Security Trimming" is enabled, the Jive connector will use the `visibility` field while indexing permission metadata on content. This data is stored in the `acl_ss` field for each document.

The value of the `visibility` field impacts the permissions assigned to a document. The following list describes how the types of permissions found in the `visibility` field of a document are used.

All	The value stored in the <code>acl_ss</code> field will be "all".
People	The document will include a list of users who are authorized to view the content. This list will be stored in the <code>acl_ss</code> field as user email addresses.
Place	A request will be made to determine the type of group. The group type will determine the permissions stored.
Open or Members Only	The value stored in the <code>acl_ss</code> field will be "all".

Private or Secret	The value stored in the <code>acl_ss</code> field will be the name of the group.
-------------------	--

6.19.3. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Property	Description
batch_size Batch Size	Number of items to fetch in each batch request. The Jive API limit is 100. If you are having problems with Jive timing out on requests, decreasing the batch size might help the crawl continue, but will extend the overall time the crawl takes. type: <code>integer</code> default value: '95'
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>
contents_to_crawl Contents to crawl <i>required</i>	List of content types to crawl, separated by commas and no spaces. The connector can only retrieve objects from Jive's Content entity, as described in https://developers.jivesoftware.com/api/v3/cloud/rest/ContentEntity.html . type: <code>string</code> default value: 'All' minLength: 1
diagnostic_mode Diagnostic Mode	Enable to print more detailed information to the logs about each request. type: <code>boolean</code> default value: 'false'

Property	Description
<p data-bbox="126 163 454 197">enable_security_trimming</p> <p data-bbox="126 231 435 264">Enable Security Trimming</p>	<p data-bbox="813 163 1490 310">Security trimming restricts query results to records the user is allowed to access by indexing user access information with other document metadata. Enable this option to index user access metadata.</p> <p data-bbox="813 344 967 378">type: boolean</p> <p data-bbox="813 411 1045 445">default value: 'false'</p>

Property	Description
<p>excluded_content_types</p> <p>Excluded content types</p>	<p>Content MIME types to exclude from fetching. Files with a content type contained in this list will not be downloaded but available metadata will be added to the index. By default, this list includes many content types Apache Tika is not able to process.</p> <p>type: <code>array of string</code></p> <p>default value: [<code>video/x-ms-wm, text/csv, application/postscript, video/x-ms-wmv, image/vnd.dwg, image/tiff, audio/x-aiff, audio/x-ms-wma, application/mac-binhex40, application/winhlp, application/x-font-ttf, video/x-ms-asf, application/illustrator, chemical/x-pdb, application/x-iso9660-image, application/x-msdownload, audio/x-pn-realaudio, application/x-ms-installer, image/gif, video/x-msvideo, application/vnd.visio, audio/mpeg, application/x-emf, image/vnd.dxf, application/x-dosexec, application/x-shockwave-flash, audio/midi, image/vnd.adobe.photoshop, application/octet-stream, application/x-java-jnlib, image/x-ms-bmp, image/jpeg, application/x-msaccess, video/mp4, text/plain, video/quicktime, application/x-dtbrsource+xml, text/x-expect, application/vnd.ms-htmlhelp, audio/basic, video/mpeg, application/x-quattro-pro, application/vnd.rn-realmedia, audio/x-wav, application/vnd.ms-cab-compressed, audio/ogg, video/ogg, application/ogg, audio/vorbis, audio/x-oggflac, audio/x-oggpcm, audio/opus, audio/speex, video/daala, video/theora, video/x-dirac, video/x-ogm, video/x-ogguvs, video/x-oggyuv, video/x-oggrgb, audio/x-flac, application/x-msmetafile, application/x-roxio-toast, application/mp4, audio/mp4, audio/mp4a-latm, video/mp4v-es, video/x-m4v, audio/mpeg4-generic, audio/vnd.sealedmedia.softseal.mpeg, audio/x-mpegurl, video/bmpeg, video/mpeg4-generic, application/x-msdownload;format=pe, application/x-dosexec, application/x-msdownload;format=pe32, application/x-msdownload;format=pe64, application/x-msdownload;format=pe-itanium, application/x-msdownload;format=pe-armLE, application/x-msdownload;format=pe-arm7, application/x-tika-msoffice,]</code></p>
<p>fetch_announcements</p> <p>Fetch Announcements</p>	<p>When enabled system and place announcements are fetched</p> <p>type: <code>boolean</code></p> <p>default value: <code>'true'</code></p>

Property	Description
fetch_deactivated_users Fetch deactivated Jive users	If enabled, deactivated users from Jive will be fetched type: <code>boolean</code> default value: 'false'
fetch_users Fetch Jive users	If enabled, users from Jive will be fetched type: <code>boolean</code> default value: 'false'
jive_client_id Jive Add-on Client Id	Jive Client ID provided by the Jive Add-on. If the Add-on has been installed in your Jive instance, OAuth will be used to authenticate to the Jive instance. If this is not provided, the username and password will be used instead. type: <code>string</code>
jive_client_secret Jive Add-on Client Secret	Jive Secret Key for the provided Client ID. type: <code>string</code>
jive_instance_url Jive instance URL <i>required</i>	Address of your Jive instance. type: <code>string</code> minLength: 1
jive_password Jive password <i>required</i>	Password for the provided user. type: <code>string</code> minLength: 1
jive_username Jive username <i>required</i>	Username to access the Jive instance. If this is provided and the Client ID is also provided, the Client ID will be preferred for OAuth-based authentication to Jive. The username and password are defined as a backup in case OAuth is not available. type: <code>string</code> minLength: 1

Property	Description
max_file_size Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. Files with size greater than this value will not be downloaded but metadata will be added. type: <code>integer</code> default value: '10485760'
max_retries Maximum number of retries.	The maximum number of retries for a failed request. type: <code>integer</code> default value: '5'
places_filter Places Filter	Filter to control the indexation of contents under specific places, this can not be set with Places to crawl at the same time type: <code>array of object</code> object attributes: \{ <ul style="list-style-type: none"> <code>object_id</code>: \{ <ul style="list-style-type: none"> display name: Object Id type: <code>string</code> <code>object_type_id</code>: \{ <ul style="list-style-type: none"> display name: Object Type Id type: <code>string</code>
places_to_crawl Places to crawl	List of places to crawl, separated by commas and no spaces type: <code>string</code> default value: 'All'
proxy_address HTTP proxy address	Address of the HTTP proxy, if required. This should be entered in the format <code>://</code> : type: <code>string</code>

Property	Description
request_delay Request delay (ms)	The amount of time, in milliseconds, to wait before perform each request. type: <code>integer</code> default value: <code>'0'</code>
requests_timeout Connection timeout (ms)	Time in milliseconds to wait for server response. type: <code>integer</code> default value: <code>'60000'</code>
retry_delay Retry delay (ms)	The number of milliseconds to wait before retrying a failed request. type: <code>integer</code> default value: <code>'15000'</code>

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
<p>unmapped</p> <p>Unmapped Fields</p>	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

ConnectorDb Configuration

Property	Description
<p>aliases</p> <p>Process Aliases?</p>	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
<p>inlinks</p> <p>Process Inlinks?</p>	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

description

Description

Optional description for datasource instance.

id

Datasource ID

Unique name for datasource instance.

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.20. Local Filesystem Connector and Datasource Configuration

A filesystem-based data store is a network of nodes to be traversed, where each node (such as a Unix file directory) provides information about its child nodes (such as the files in that directory) or references other nodes (such as links in an HTML document).

The crawler captures information about the node, e.g., filename, permissions, date of creation, last modification, and last access, as well as the contents of the nodes. The extent of the network of nodes to be traversed is discovered during the crawl.

The connector provides rules to limit the crawl and re-crawling. These rules use datasource configuration properties to limit the extent of the network (depth of nodes to explore) as well as limiting processing to a subset of files based on file names and file size. An overall limit can be set on number of files retrieved during a crawl.

6.20.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.addFileMetadata Add file metadata	Set to true to add information about documents found in the filesystem to the document, such as document owner, group, or ACL permissions. type: <code>boolean</code> default value: <code>'true'</code>
f.maxSizeBytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: <code>integer</code> default value: <code>'4194304'</code>
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of documents to fetch. type: <code>integer</code> default value: <code>'0'</code>

Limit Documents

Property	Description
f.maxSizeBytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: integer default value: '4194304'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of documents to fetch. type: integer default value: '0'
f.addFileMetadata Add file metadata	Set to true to add information about documents found in the filesystem to the document, such as document owner, group, or ACL permissions. type: boolean default value: 'true'
restrictToTree Restrict to sub-directories and child pages	If true, only documents found in a tree below the start links will be fetched. type: boolean default value: 'true'
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1'
includeExtensions Included file extensions	File extensions to be fetched. This will limit this datasource to only these file extensions. type: array of string

Property	Description
includeRegexes Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: <code>array of string</code>
excludeExtensions Excluded file extensions	File extensions that should not to be fetched. This will limit this datasource to all extensions except this list. type: <code>array of string</code>
excludeRegexes Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: <code>array of string</code>

Crawl Performance

Property	Description
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: <code>integer</code> default value: '1'
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: <code>integer</code> default value: '5'
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: <code>integer</code> default value: '0'

Property	Description
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: integer default value: '5'
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: boolean default value: 'true'
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: boolean default value: 'true'

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: boolean default value: 'false'
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: boolean default value: 'false'

Property	Description
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: <code>string</code>
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content)\{\}' function, but can use any combination of document fields. The function must return a string. type: <code>string</code>

Recrawl Rules

Property	Description
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: <code>boolean</code> default value: 'true'
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: <code>integer</code> default value: '-1'
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: <code>boolean</code> default value: 'true'
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: <code>boolean</code> default value: 'false'

Property	Description
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: boolean default value: 'false'
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: integer default value: '-1'
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: array of string
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: array of string
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: string
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: boolean default value: 'false'

Crawl History

Property	Description
retainOutlinks Retain links in the crawl db	Set to true for links found during fetching to be stored in the crawl db. This increases precision in certain recrawl scenarios, but requires more memory and disk space. type: boolean default value: 'false'
aliasExpiration Alias expiration	The number of crawls after which an alias will expire. The default is 1 crawl. type: integer default value: '1'
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: string default value: 'in-memory' enum: \{ in-memory on-disk }
indexCrawlDBToSolr Index crawl database to Solr	EXPERIMENTAL: Set to true to index the crawl-database into a 'crawl db_' collection in Solr. type: boolean default value: 'false'

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>{</code></p> <p>operation : <code>{</code></p> <p>display name: Operation</p> <p>type: string</p> <p>default value: 'copy'</p> <p>description : The type of mapping to perform: move, copy, delete, add, set, or keep.</p> <p>enum: <code>{ copy move delete set add keep }</code></p> <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.21. MongoDB Datasource and Connector Configuration

```
<a href="http://lucidworks.com/connectors/">" width="45.0" tmp="false">] <a href="http://lucidworks.com/connectors/">Connector download</a>
```

Retrieve data from a MongoDB instance.

This is an add-on connector that you can [download](#) and install.

For the initial data ingest from a MongoDB database, select the option to do a full synchronization of the content in MongoDB. Subsequent ingest runs can use the `oplog` in MongoDB to discover new content and updates to existing content (updated or removed documents). A full synchronization can be done by selecting the full synchronization option and re-running the data ingest job.

6.21.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
batch_size_solr_commit Batch size Solr commit	The number of documents every time <code>solr_commit</code> will be made. type: <code>integer</code> default value: <code>'1000'</code>
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>
collections MongoDB Collections to index	The MongoDB collections to index, in the format <code>'databaseName.collection'</code> . Multiple collections can be separated by commas. The default <code>'.'</code> option crawls all databases (limited by user access) and their related collections. type: <code>string</code> default value: <code>'.'</code> minLength: 1

Property	Description
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: boolean default value: 'true'
list_credentials Credentials	Credentials for Mongo databases type: array of object object attributes: \{ <pre> database : \{ display name: Database type: string } id : \{ display name: Auth Config id type: string } password : \{ display name: Password type: string } username : \{ display name: Username type: string } } </pre>
list_hosts Hosts <i>required</i>	Host and ports of Mongo nodes type: array of object default value: [{"host" => "localhost", "port" => 27017},] object attributes: \{ <pre> host : \{ display name: Host type: string } port : \{ display name: Port type: integer } } </pre>

Property	Description
process_oplog Process OPLog <i>required</i>	Process updates from the oplog. Disable this option to perform a full synchronization of content in MongoDB collections with the index. type: boolean default value: 'true'
read_preferences Read Preference Modes	Read preference describes how MongoDB clients route read operations to the members of a replica set. type: string default value: 'primary' enum: \{ primary primary preferred secondary secondary preferred nearest \}
tag_set_list Read Preference Tag Sets	A list of Tag Sets used for non-primary read modes type: array of object default value: [] object attributes: \{ tag_set : \{ display name: Tag Set type: array of object default value: '' } \}
verify_access Validate access	Set to true to require successful connection to the filesystem before saving this datasource. type: boolean default value: 'true'

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
unmapped Unmapped Fields	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } }</p>

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
inlinks Process Inlinks?	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
description Description	Optional description for datasource instance.
id Datasource ID	Unique name for datasource instance.
pipeline Pipeline ID	The index pipeline used to process documents.

6.22. Solr Push Endpoint Datasource and Configuration



width=600%

The Solr Push Endpoint accepts documents and pushes them to Solr using the Fusion index pipelines.

In Fusion 3.0 and earlier, this is called the Push connector.

It uses the embedded JettySolrRunner to push the documents. This requires defining a port for the JettySolrRunner that is not already in use by any other process. The documents can then be sent to Fusion at that port, and they will be consumed by Fusion.

6.22.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: <code>boolean</code> default value: <code>'true'</code>
port Port <i>required</i>	The port that will be used to push content to Solr. type: <code>integer</code>
url URL	The endpoint to which the external application will send the documents. type: <code>string</code>

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } } 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } }</p>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
unmapped Unmapped Fields	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } }</p>

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
inlinks Process Inlinks?	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
description Description	Optional description for datasource instance.
id Datasource ID	Unique name for datasource instance.
pipeline Pipeline ID	The index pipeline used to process documents.

6.23. S3 Connector and Datasource Configuration

The S3 connector can access AWS S3 buckets in native format.

6.23.1. Bucket Permissions

The connector uses the S3 API to request data from S3. It calls the `listBucket` service, which lists all buckets owned by the user account supplied to the connector.

When creating an S3 datasource using the UI, Fusion automatically verifies that the user information supplied has access to the bucket defined in the URL property. If the bucket is not in the list returned by S3, datasource creation may fail. At crawl time, if the bucket is not in the list returned by S3, the crawl will fail.

Permission errors when trying to create or crawl the datasource may be caused by incorrect username or password, or they may be due to user account permissions. The user account must have List Bucket permissions for the account which owns the bucket that the crawler is trying to access.

6.23.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
add_failed_docs Add failed documents	Set to true to add documents even if they partially fail processing. Failed documents will be added with as much metadata as available, but may not include all expected fields. type: <code>boolean</code> default value: <code>'false'</code>
aws_region AWS Region Bucket	AWS Region bucket on GET requests with AWS Signature Version 4 enabled type: <code>string</code> default value: <code>'us-west-2'</code> enum: <code>{ us-gov-west-1 us-east-1 us-west-1 us-west-2 eu-west-1 eu-central-1 ap-southeast-1 ap-southeast-2 ap-northeast-1 ap-northeast-2 sa-east-1 cn-north-1 }</code>

Property	Description
bounds Crawl bounds	Limits the crawl to a specific directory sub-tree, hostname or domain. type: string default value: 'tree' enum: \{ tree host domain none }
collection Collection	Collection documents will be indexed to. type: string pattern: ^[a-zA-Z0-9_-]+\$
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: boolean default value: 'true'
crawl_depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1' exclusiveMinimum: false minimum: -1
crawl_item_timeout Fetch timeout	Time in milliseconds to fetch any individual document. type: integer default value: '600000' exclusiveMinimum: true minimum: 0

Property	Description
exclude_paths Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string
include_extensions Included file extensions	List the file extensions to be fetched. Note: Files with possible matching MIME types but non-matching file extensions will be skipped. Extensions should be listed without periods, using whitespace to separate items (e.g., 'pdf zip'). type: array of string
include_paths Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: array of string
index_directories Index directories	Set to true to add directories to the index as documents. If set to false, directories will not be added to the index, but they will still be traversed for documents. type: boolean default value: 'false'
max_bytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: integer default value: '10485760' exclusiveMinimum: false minimum: -1

Property	Description
max_docs Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: <i>integer</i> default value: '-1' exclusiveMinimum: false minimum: -1
max_threads Fetch threads	The maximum number of threads to use for fetching data. Note: Each thread will create a new connection to the repository, which may make overall throughput faster, but this also requires more system resources, including CPU and memory. type: <i>integer</i> default value: '1'
maximum_connections Maximum fetch connections	Maximum number of concurrent connections to the filesystem. A large number of documents could cause a large number of simultaneous connections to the repository and lead to errors or degraded performance. In some cases, reducing this number may help performance issues. type: <i>integer</i> default value: '1000'
password AWS Secret Key <i>required</i>	The AWS Secret Key associated with the Access Key. type: <i>string</i>
retryDelay Retry Delay	The retry time delay, in milliseconds. type: <i>integer</i> default value: '1000' exclusiveMinimum: false minimum: 1000

Property	Description
stopRetry Max Retry Time	The maximum time to retry failed requests, in minutes. type: integer default value: '5' exclusiveMinimum: false minimum: 1
url S3 URL <i>required</i>	A fully-qualified S3 URL, including bucket and sub-bucket paths, as required, e.g., 's3://{bucketName}/{path}'. type: string minLength: 1 pattern: ..
use_sigv4 Use AWS Signature Version 4	Additional HTTP header on GET requests to retrieve encrypted objects by SSE-KMS type: boolean default value: 'false'
username AWS Key <i>required</i>	An AWS Access Key ID that can access the content. type: string
verify_access Validate access	Set to true to require successful connection to the filesystem before saving this datasource. type: boolean default value: 'true'

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{"operation" ⇒ "move", "source" ⇒ "fetch_time", "target" ⇒ "fetch_time_dt"}{"operation" ⇒ "move", "source" ⇒ "ds:description", "target" ⇒ "description"}</code></p> <p>object attributes: <code>{</code></p> <ul style="list-style-type: none"> <code>operation</code> : <code>{</code> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: <code>{ copy move delete set add keep }</code> <pre> } + 'source' _ (required)_ : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
<p>unmapped</p> <p>Unmapped Fields</p>	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + }</pre>

ConnectorDb Configuration

Property	Description
<p>aliases</p> <p>Process Aliases?</p>	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
<p>inlinks</p> <p>Process Inlinks?</p>	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

description

Description

Optional description for datasource instance.

id

Datasource ID

Unique name for datasource instance.

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.24. Salesforce Connector and Datasource Configuration



width=600%

Use the [Salesforce REST API](#) to extract data from a Salesforce repository via a Salesforce Connected App.

This service has an [authentication procedure](#) and [API usage limits](#).

Note	You need to have the Consumer Key and Consumer Secret of the Salesforce Connected App, which must be created by a Salesforce administrator for the account. Instructions on how to do this follow the datasource configuration details.
Note	Salesforce authentication requires a security token as well as a user password. A security token is an automatically-generated key from Salesforce. For more information on security tokens see Reset Your Security Token in the online help.

6.24.1. Security Trimming

When you enable security trimming for a Salesforce connector, the system uses the Fusion username and tries finding the identical ID in Salesforce. If it finds this ID, then it uses the permissions given to that ID in Salesforce and applies a filter to the search query. If the ID is not found then it applies a filter to block all documents from being shown.

The security trimming feature assumes that the Fusion username is the same as Salesforce alias. When retrieving data from Salesforce, the connector retrieves the user Id based on the Fusion username (Salesforce alias) via the query:

```
Select Id from User WHERE alias={username}
```

If it isn't found, then the connector will create a security filter to deny access to all documents.

If you are logged into the Fusion UI or send a request to the REST API with username "admin", the Salesforce connector uses that as the value of property "salesforce_username". Therefore, in Fusion, you must have accounts which reflect the Salesforce accounts.

6.24.2. Creating a Salesforce Connected App

Before creating the datasource, a Salesforce administrator for the account must create a Salesforce Connected App in Salesforce:

- Go to Setup > Create > Apps and in the Connected Apps section and click the new button
- Add values to the following properties:
 - Connected App Name
 - API Name
 - Contact Email

- Check the property Enable OAuth Settings
- More properties should be displayed:
 - Callback URL: You can use a dumb HTTPS URL e.g. <https://mydomain.com/auth>
 - Selected OAuth Scopes: Select Full Access(full) scope
- Save the new Connected App
- After this step a message will be displayed, click the continue button.
- In the section API (Enable OAuth Settings) the consumer key and consumer secret should be displayed, which will be used on the datasource

For the Salesforce credentials the user must be a System Administrator or one with no restrictions (able to access all the objects and fields)

6.24.3. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Property	Description
collection Collection type: <code>string</code> pattern: <code>^[a-zA-Z0-9_]+\$</code>	Collection documents will be indexed to.
connected_app_client_id Salesforce Client Id <i>required</i> type: <code>string</code> minLength: 1	Salesforce connected app Consumer Key
connected_app_client_secret Salesforce Client secret <i>required</i> type: <code>string</code> minLength: 1	Salesforce connected app Consumer Secret
diagnostic_mode Diagnostic Mode type: <code>boolean</code> default value: <code>'false'</code>	

Property	Description
enable_security_trimming Enable Security Trimming	Security trimming restricts query results to records the user is allowed to access by indexing user access information with other document metadata. Enable this option to index user access metadata. type: <code>boolean</code> default value: 'false'
is_production_environment Is production environment	Set to true for production instances and false for sandboxes. type: <code>boolean</code> default value: 'true'
objects_to_crawl Objects to crawl	List of objects to crawl, separated by commas and no spaces. type: <code>array of string</code> default value: [<code>Case</code> , <code>CaseComment</code> , <code>CaseHistory</code> , <code>CaseFeed</code> , <code>FeedComment</code> , <code>Asset</code> , <code>Account</code> , <code>Contact</code> , <code>Opportunity</code> , <code>Product2</code> ,]
salesforce_password Salesforce password <i>required</i>	type: <code>string</code> minLength: 1
salesforce_username Salesforce username <i>required</i>	type: <code>string</code> minLength: 1
soql_query Custom SOQL query	Optional SOQL query to limit or join records fetched for indexing. type: <code>string</code> maxLength: 15000

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px;"> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
unmapped Unmapped Fields	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } }</p>

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
inlinks Process Inlinks?	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
description Description	Optional description for datasource instance.
id Datasource ID	Unique name for datasource instance.
pipeline Pipeline ID	The index pipeline used to process documents.

6.25. ServiceNow Connector and Datasource Configuration



width=600%

The ServiceNow Datasource retrieves data from ServiceNow repository via the [ServiceNow REST API](#). ServiceNow records are stored in named tables.

The ServiceNow connector fetches records from these tables and transforms them into Fusion documents according to the datasource configuration. Access to the ServiceNow requires both a ServiceNow username and password, as well as an [OAuth](#) client password and token.

6.25.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
batch_size Batch size	Number of records to fetch in batch requests. The default is 100 to avoid the ServiceNow REST request timeout of 60 seconds. Only increase the batch size if you are sure a higher batch size will not trigger the timeout. type: <code>integer</code> default value: '100' exclusiveMaximum: false exclusiveMinimum: false maximum: 10000 minimum: 1
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>

Property	Description
diagnostic_mode Diagnostic Mode	Enable to print more detailed information to the logs about each request. type: boolean default value: 'false'
enable_security_trimming Enable Security Trimming	Security trimming restricts query results to records the user is allowed to access by indexing user access information with other document metadata. Enable this option to index user access metadata. type: boolean default value: 'false'
exclude_field_name_value Exclude records by field values	Records can be excluded based on values of specific fields. Enter exclusions in the format fieldName=fieldValue, e.g., workflow_state=review. type: array of string
oauth_application_client_id OAuth application Client Id <i>required</i>	OAuth application Client ID. This ID is created after registering the OAuth endpoint for the ServiceNow instance. type: string minLength: 1
oauth_application_client_secret OAuth application Client Secret <i>required</i>	OAuth application Client Secret. This key is created after registering the OAuth endpoint for the ServiceNow instance. type: string minLength: 1
servicenow_instance_url ServiceNow Instance URL <i>required</i>	The ServiceNow instance address. type: string minLength: 1

Property	Description
servicenow_password ServiceNow password <i>required</i>	Password to access the ServiceNow instance. type: <i>string</i> minLength: 1
servicenow_username ServiceNow username <i>required</i>	A user with access to all of the tables configured below. This user should also have access to related tables, as appropriate, in order to retrieve content referenced in other tables. type: <i>string</i> minLength: 1
tables_to_crawl ServiceNow tables to crawl	ServiceNow tables to fetch content from. At least one table name should be entered type: <i>array of string</i> minimum number of items (<i>minItems</i>): 1

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
unmapped Unmapped Fields	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } + }</p>

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
inlinks Process Inlinks?	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
description Description	Optional description for datasource instance.
id Datasource ID	Unique name for datasource instance.
pipeline Pipeline ID	The index pipeline used to process documents.

6.26. SharePoint Connector and Datasource Configuration



width=600%

The SharePoint connector retrieves content and metadata from an on-premises SharePoint repository.

To retrieve content from cloud-based SharePoint repositories, see the SharePoint Online connector.

This connector can access a SharePoint repository running on the following platforms:

- Microsoft SharePoint 2010
- Microsoft SharePoint 2013
- Microsoft SharePoint 2016

See this tutorial about configuring a SharePoint datasource and enabling security trimming:

When crawling, the connector discovers SharePoint contents in the following order: sites, then sub-sites (children). A site may contain:

- Sub-sites
- Generic Lists
 - List Items
 - Attachments
- Document Libraries
 - Folders
 - Documents

When the connector re-crawls a SharePoint repository, each previously crawled URL is accessed before any newly discovered objects, but no order is guaranteed. The connector uses a cache to store retrieved parent objects to avoid unnecessary requests. The last modified date of each object is retrieved to determine if it has changed since the last crawl. If it has changed, a new request is made to retrieve the changes. If it has not changed, the object is skipped and no additional request is made.

The connector uses SOAP to connect to and retrieve documents, lists, and other objects for indexing. It does not access a SharePoint site in the same way that a regular user does, and it needs additional privileges to use the SOAP interface to SharePoint.

The connector can be configured to work with Active Directory (AD) or LDAP to retrieve the ACLs for each object, which can then be used for security trimming at query time. In order to use security trimming to restrict user access to SharePoint objects, the authenticated user must have sufficient privileges to read every document in the system and determine which users can access them. The permissions requirements are explained below.

6.26.1. SharePoint Permissions

SharePoint security trimming restricts access to documents based on user permissions. There are two types of [permissions in SharePoint](#):

- Site permissions, which are:
 - managed by SharePoint
 - customizable for each site or subsite
 - inherited by subsites as the default permissions
 - grantable to users and groups
- User permissions, which are:
 - assigned by group membership, when groups have been configured and provided permissions
 - assigned directly to the user

These permissions are stored as ACLs. When the SharePoint server is configured with security trimming set to "true", then documents retrieved from SharePoint have the set of all ACLs stored in a `acl_ss` field on each document.

At search time, the ACLs are used to verify if a user has access to a document. This is configured in a query pipeline with a Security Trimming Query Stage.

To crawl all the sites and subsites, the authenticated user must belong the site administrators group. If not, Fusion can still crawl and complete the job, but the crawled data will be limited by the user's privileges. In addition, a WARNING message will appear in the `connector.log` indicating that the user is not site administrator and therefore unable to get sites from site collections. The message starts with `Authorization Error (401)`.

Required Permissions

The SharePoint datasource must be configured with the name of a user who has sufficient permissions to crawl the entire site. These permissions require use of a custom Permission Policy. The required permissions correspond to the concept of Site Collection Auditor, a permission type which is not the same as a Site Administrator, but requires almost all of the Site Administrator privileges.

You will need to work with your SharePoint administrator to ensure that the account used by Fusion has all of the permissions listed in the following table:

Permission Type	Permission	Description
Site Collection Auditor		Full Read access for the entire site collection, including reading permissions and configuration data.
List	View Items	View items in lists and documents in document libraries.
List	Open Items	View the source of documents with server-side file handlers.
List	View Versions	View past versions of a list item or document.
Site	Browse Directories	Enumerate files and folders in a Web site using SharePoint Designer and WebDAV interfaces.
Site	View Pages	View pages in a Web site.

Permission Type	Permission	Description
Site	Enumerate Permissions	Enumerate permissions on the Web site, list, folder, document, or list item.
Site	Browse User Information	View information about users of the Web site.
Site	Use Remote Interfaces	Use SOAP, WebDAV, Client Object Model, or SharePoint Designer interfaces to access the Web site.
Site	Open	Open a Web site, list or folder in order to access items inside that container.

Troubleshooting Permission Issues

When the connector is configured using a SharePoint username without sufficient privileges, the Fusion connectors log file `fusion/3.1.x/var/log/connectors/connectors.log` contains an error like the following:

```
crawler.common.sharepoint.exception.SharePointException: Server was unable to process request. ---> Attempted
to perform an unauthorized operation. at
crawler.common.sharepoint.service.BaseService.analyzeResponse(BaseService.java:194) ~[classes/:?] at
crawler.common.sharepoint.service.SiteDataService.getContentBySiteOrList(SiteDataService.java:169)
~[classes/:?] at com.lucidworks.permissions.Main.test1(Main.java:50) [classes/:?] at
com.lucidworks.permissions.Main.main(Main.java:32) [classes/:?]
```

This user's permissions may be sufficient to connect via SOAP and read the documents, but not sufficient to get the ACLs and other associated metadata. This may result in complete lack of access to documents, or access to unauthorized documents. Confirm that the configured SharePoint user has the required privileges.

6.26.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.avoid_ssl_hostname_verification Avoid SSL hostname verification	Enable this in cases when the CN on the SSL certificate does not match the host name of the server. type: <code>boolean</code> default value: 'true'
f.domain Sharepoint domain	Authentication domain for the Sharepoint user. type: <code>string</code>

Property	Description
f.enable_http_headers_debugging Enable HTTP headers debugging	Prints DEBUG level information to the logs. type: boolean default value: 'false'
f.log_sharepoint_xml Log Sharepoint Soap XML	When analyzing Sharepoint crawls, it can be helpful to log the Soap XML messages between Fusion and Sharepoint. By selecting this, Sharepoint will log the xml of all messages to the connectors log. type: boolean default value: 'false'
f.maxSizeBytes Maximum file size (bytes)	Maximum size, in bytes, of a document to crawl. type: integer default value: '4194304'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of a document to crawl. type: integer default value: '0'
f.password Sharepoint password	Password for the Sharepoint user. type: string
f.remove_prepend_ids Remove preprended IDs	If fields have been defined to include PrependIds, this option will remove those IDs before indexing. type: boolean default value: 'true'
f.replace_invalid_xml_entities Replace invalid xml entities	Replace invalid xml entities to avoid the error: \"Unmarshalling Error: Illegal character entity: expansion character ...\" type: boolean default value: 'false'

Property	Description
f.sharepoint_online Sharepoint Online	type: boolean default value: 'false'
f.sharepoint_services_timeout Sharepoint services timeout	Time in milliseconds to wait for a server response. type: integer default value: '600000'
f.username Sharepoint username	Name of a Sharepoint user who has the required permissions to access Sharepoint via the SOAP API. type: string

Authentication

Property	Description
f.username Sharepoint username <i>required</i>	Name of a Sharepoint user who has the required permissions to access Sharepoint via the SOAP API. type: string
f.password Sharepoint password <i>required</i>	Password for the Sharepoint user. type: string
f.domain Sharepoint domain	Authentication domain for the Sharepoint user. type: string
f.avoid_ssl_hostname_verification Avoid SSL hostname verification	Enable this in cases when the CN on the SSL certificate does not match the host name of the server. type: boolean default value: 'true'

Property	Description
f.sharepoint_services_timeout Sharepoint services timeout	Time in milliseconds to wait for a server response. type: integer default value: '600000'
f.enable_http_headers_debugging Enable HTTP headers debugging	Prints DEBUG level information to the logs. type: boolean default value: 'false'
f.remove_prepend_ids Remove prepended IDs	If fields have been defined to include PrependIds, this option will remove those IDs before indexing. type: boolean default value: 'true'

Security Trimming

Property	Description
enable_security_trimming Enable Security Trimming <i>required</i>	type: object object attributes: \{ }

Limit Documents

Property	Description
includeRegexes Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: array of string
excludeRegexes Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string

Property	Description
includeExtensions Included file extensions	File extensions to be fetched. This will limit this datasource to only these file extensions. type: array of string
excludeExtensions Excluded file extensions	File extensions that should not to be fetched. This will limit this datasource to all extensions except this list. type: array of string
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1'
f.maxSizeBytes Maximum file size (bytes)	Maximum size, in bytes, of a document to crawl. type: integer default value: '4194304'
f.minSizeBytes Minimum file size (bytes)	Minimum size, in bytes, of a document to crawl. type: integer default value: '0'
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: boolean default value: 'true'

Property	Description
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: <code>integer</code> default value: <code>'-1'</code>

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: <code>boolean</code> default value: <code>'false'</code>
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: <code>boolean</code> default value: <code>'false'</code>
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: <code>string</code>
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content)\{' function, but can use any combination of document fields. The function must return a string. type: <code>string</code>

Recrawl Rules

Property	Description
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: <code>boolean</code> default value: <code>'true'</code>
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: <code>boolean</code> default value: <code>'false'</code>
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: <code>boolean</code> default value: <code>'false'</code>
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: <code>integer</code> default value: <code>'-1'</code>
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: <code>array of string</code>
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: <code>array of string</code>
refreshScript Recrawl script	A JavaScript function ('shouldRefresh0') to customize the items recrawled. type: <code>string</code>

Property	Description
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: <code>boolean</code> default value: 'false'

Crawl Performance

Property	Description
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: <code>integer</code> default value: '1'
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: <code>integer</code> default value: '5'
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: <code>integer</code> default value: '0'
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: <code>integer</code> default value: '5'
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: <code>boolean</code> default value: 'true'

Property	Description
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: <code>boolean</code> default value: 'true'
parserRetryCount Max Parser Retries	The maximum number of times the configured parser will try getting content before giving up type: <code>integer</code> default value: '0' exclusiveMaximum: true exclusiveMinimum: false maximum: 5 minimum: 0

Crawl History

Property	Description
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: <code>string</code> default value: 'in-memory' enum: \{ in-memory on-disk }
indexCrawlDBToSolr Index crawl database to Solr	EXPERIMENTAL: Set to true to index the crawl-database into a 'crawlddb_' collection in Solr. type: <code>boolean</code> default value: 'false'

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>{</code></p> <p>operation : <code>{</code></p> <p>display name: Operation</p> <p>type: string</p> <p>default value: 'copy'</p> <p>description : The type of mapping to perform: move, copy, delete, add, set, or keep.</p> <p>enum: <code>{ copy move delete set add keep }</code></p> <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.27. Solr Connector and Datasource Configuration



width=600%

A Solr connector pulls documents from an external standalone Solr instance or SolrCloud cluster using Solr's [javabin](#) response type and streaming response parser.

For Solr v4.7 and greater, cursorMark deep-paging is used. For earlier versions of Solr, standard paging (start+rows) is used.

The following Solr components and parameters can be configured:

- collection/core (also allows default/empty core)
- query (: by default)
- filter queries
- query parser
- request handler (defaults to /select)
- stored fields to retrieve

Also, since cursorMark deep paging should be used when possible:

- sort spec (default: id asc)

This connector can be configured to store information about datasources and the data ingested in a ConnectorDB crawlDb instance.

6.27.1. Limitations

1. Cannot do incremental crawls. (May be possible to do so in the future using source Solr docs' version field.)
2. Cannot do manual filtered deep paging.
3. Doesn't check that both sort spec and field list contain uniqueKey field.
4. Cannot handle encrypted connection to Solr

6.27.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Property	Description
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: <code>boolean</code> default value: <code>'true'</code>
solr_base_url Standalone Solr server base URL	If using a single Solr instance, enter the base URL, e.g., http://solrhost.example.com:8983/solr/ . type: <code>string</code>
solr_field_list Stored fields to retrieve	Fields to fetch from the source Solr instance/cluster, which must be stored fields. Multiple field names should be separated with commas. type: <code>string</code> default value: <code>'*'</code> minLength: 1
solr_filter_queries Filter queries to execute	Filter queries to select documents from the source Solr instance/cluster. Multiple filter queries should be separated with commas. type: <code>string</code>
solr_page_size Page size	Number of rows per request to Solr. type: <code>integer</code> default value: <code>'100'</code>

Property	Description
solr_query Query to execute	Query to select documents from the source Solr instance/cluster. If not defined, the default query : will be used. type: <code>string</code> default value: <code>‘:’</code> minLength: 1
solr_query_parser Query parser	The query parser to use for the request. type: <code>string</code>
solr_request_handler Solr request handler	The request handler to use for the request to the Solr instance/cluster. type: <code>string</code> default value: <code>‘/select’</code> minLength: 1
solr_sort_spec Sort specification	Sort order for the request. The uniqueKey field must be included as one of the sorted fields. type: <code>string</code> default value: <code>‘id asc’</code> minLength: 1
source_collection Source collection	Collection or Core name in the source Solr instance/cluster. If not defined, the default core will be used. type: <code>string</code>
verify_access Validate access	Set to true to require successful connection to the filesystem before saving this datasource. type: <code>boolean</code> default value: <code>‘true’</code>

Property	Description
zk_host_string SolrCloud ZooKeeper host string	If using a SolrCloud instance, enter the ZooKeeper connect string, e.g., zkServerA:2181,zkServerB:2181,zkServerC:2181/solr. type: <code>string</code>

Field Mapping

Initial Mappings	Description
mappings Field Mappings	List of mapping rules type: <code>array of object</code> default value: <code>{ "operation" => "move", "source" => "version", "target" => "external_version_s" }</code> object attributes: <code>\{</code> <code>operation</code> : <code>\{</code> display name: Operation type: <code>string</code> default value: <code>'copy'</code> description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: <code>\{ copy move delete set add keep }</code> <code>source</code> : <code>\{</code> display name: Source Field type: <code>'string'</code> description : The name of the field to be mapped. <code>target</code> : <code>\{</code> display name: Target Field type: <code>'string'</code> description : The name of the field to be mapped to. <code>}</code> <code>}</code> <code>}</code>
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: <code>boolean</code> default value: <code>'false'</code>

Initial Mappings	Description
unmapped Unmapped Fields	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } }</p>

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
inlinks Process Inlinks?	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
description Description	Optional description for datasource instance.
id Datasource ID	Unique name for datasource instance.
pipeline Pipeline ID	The index pipeline used to process documents.

6.28. SolrXML Connector and Datasource Configuration



width=600%

The SolrXML connector indexes XML files formatted according to Solr's XML structure. It is not a generic XML file crawler; it can only index SolrXML-formatted documents.

Per the Solr standard, all XML files must include the `<add>` tag in order for the documents to be added to the Fusion index.

6.28.1. The SolrXML Format

As described in the Solr Reference Guide [section on using Solr's updateHandlers](#), an XML document formatted for Solr must conform to a very specific structure. There are three general elements that are used:

- `<add>` introduces one or more documents to be added to the index.
- `<doc>` introduces the fields that make up a single document.
- `<field>` defines the content for each field of the document.

For example, this is very simple XML including only one document:

```
<add>
  <doc>
    <field name="id">doc1</field>
    <field name="title">My Solr Document</field>
    <field name="body">This is the body of my document.</field>
  </doc>
</add>
```

The fields can be any field that is defined in your schema, or you can use dynamic field rules to create fields during indexing.

The elements can take some attributes to define document overwrites, commit rules and field or document boosts. See the Solr Reference Guide [section on XML-formatted updates](#) for more details.

6.28.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Property	Description
collection Collection	Collection documents will be indexed to. type: string pattern: <code>^[a-zA-Z0-9_-]+\$</code>
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: boolean default value: 'true'
exclude_paths Exclude paths	An array of regular expression patterns that indicate documents to be excluded from the index. Multiple expressions can be separated by commas. type: array of string default value: []
generate_unique_key Generate unique key?	If true, a unique identifier will be added to each document. In most cases, this is an 'id' field, unless it was changed in your implementation. If your documents already include an ID field, you can set this to false. type: boolean default value: 'true'
include_datasource_metadata Include datasource metadata?	Set True to add '_lw_data_source_s' and '_lw_data_source_type_s' fields to each document in addition to fields found in the file. These fields will ensure these documents are associated with this datasource for faceting, information shown in the UI, or later document removal. type: boolean default value: 'true'

Property	Description
include_paths Include paths	An array of regular expression patterns that indicate documents to be included in the index. Multiple expressions can be separated by commas. type: array of string default value: [.*\. xml,]
max_docs Max documents	The maximum number of documents to crawl. Use -1 to index all documents found. type: integer default value: '-1'
path Path <i>required</i>	Name of the file to read, or directory containing files to read. type: string minLength: 1
url URL	Read-only value that shows the absolute path. type: string minLength: 1
verify_access Validate access	Set to true to require successful connection to the filesystem before saving this datasource. type: boolean default value: 'true'

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
unmapped Unmapped Fields	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } + }</p>

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
inlinks Process Inlinks?	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
description Description	Optional description for datasource instance.
id Datasource ID	Unique name for datasource instance.
pipeline Pipeline ID	The index pipeline used to process documents.

6.29. Subversion Connector and Datasource Configuration



width=600%

This connector requires a Subversion client that is compatible with JavaHL.

Fusion ships with binaries for Windows and Linux which can be used if a JavaHL-compatible client is not already installed. These are found in subdirectories of the `fusion/3.1.x/apps/connectors/resources/lucid.anda/subversion` directory.

6.29.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Connector-specific Properties

Property	Description
f.subversion_batch_size Subversion number of revisions in batch	Number of records to fetch in each batch request. type: <code>integer</code> default value: '100'
f.subversion_configuration_directory Subversion configuration directory	Path to a configuration directory that includes a file named 'servers' to support self-signed certificates. If the Subversion repository does not use SSL or does not use self-signed certificates, this can be left empty. type: <code>string</code>
f.subversion_native_library Subversion JavaHL library path	Path to the <code>libsvnjavahl-1.so</code> (Linux) or <code>libsvnjavahl-1.dll</code> (Windows) library. If using the libraries included with Fusion, you do not need to specify the path. type: <code>string</code> default value: 'fusion-svn-binaries'
f.subversion_password Subversion Password	Password for user, if required. type: <code>string</code>

Property	Description
f.subversion_username Subversion Username	Username with permissions to access the repository, if the Subversion repository requires authentication. type: <code>string</code>

Limit Documents

Property	Description
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: <code>integer</code> default value: <code>'-1'</code>
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: <code>integer</code> default value: <code>'-1'</code>
includeExtensions Included file extensions	File extensions to be fetched. This will limit this datasource to only these file extensions. type: <code>array of string</code>
includeRegexes Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: <code>array of string</code>
excludeExtensions Excluded file extensions	File extensions that should not to be fetched. This will limit this datasource to all extensions except this list. type: <code>array of string</code>
excludeRegexes Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: <code>array of string</code>

Property	Description
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: <code>boolean</code> default value: <code>'true'</code>
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: <code>integer</code> default value: <code>'-1'</code>

Crawl Performance

Property	Description
f.subversion_batch_size Subversion number of revisions in batch	Number of records to fetch in each batch request. type: <code>integer</code> default value: <code>'100'</code>
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: <code>integer</code> default value: <code>'1'</code>
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: <code>integer</code> default value: <code>'5'</code>
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: <code>integer</code> default value: <code>'0'</code>

Property	Description
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: <code>integer</code> default value: '5'
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: <code>boolean</code> default value: 'true'
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: <code>boolean</code> default value: 'true'

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: <code>boolean</code> default value: 'false'
dedupeSaveSignature Save dedupe signature	If true, the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: <code>boolean</code> default value: 'false'

Property	Description
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: <code>string</code>
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content)\{' function, but can use any combination of document fields. The function must return a string. type: <code>string</code>

Recrawl Rules

Property	Description
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: <code>boolean</code> default value: 'true'
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: <code>boolean</code> default value: 'false'
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: <code>boolean</code> default value: 'false'
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: <code>integer</code> default value: '-1'

Property	Description
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: <code>array of string</code>
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: <code>array of string</code>
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: <code>string</code>
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: <code>boolean</code> default value: 'false'

Crawl History

Property	Description
retainOutlinks Retain links in the crawl db	Set to true for links found during fetching to be stored in the crawl db. This increases precision in certain recrawl scenarios, but requires more memory and disk space. type: <code>boolean</code> default value: 'false'
aliasExpiration Alias expiration	The number of crawls after which an alias will expire. The default is 1 crawl. type: <code>integer</code> default value: '1'

Property	Description
crawlDBType Crawl database type	The type of crawl database to use, in-memory or on-disk. type: <code>string</code> default value: 'in-memory' enum: \{ in-memory on-disk }
indexCrawlDBToSolr Index crawl database to Solr	EXPERIMENTAL: Set to true to index the crawl-database into a 'crawldb_' collection in Solr. type: <code>boolean</code> default value: 'false'
reevaluateCrawlDBOnStart Reevaluate crawlDB on start?	Reevaluate existing crawlDB entries for legality on startup? type: <code>boolean</code> default value: 'false'

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>{</code></p> <p>operation : <code>{</code></p> <p>display name: Operation</p> <p>type: string</p> <p>default value: 'copy'</p> <p>description : The type of mapping to perform: move, copy, delete, add, set, or keep.</p> <p>enum: <code>{ copy move delete set add keep }</code></p> <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
pipeline Pipeline ID	The index pipeline used to process documents.

6.30. Twitter Search Connector and Datasource Configuration

The Twitter Search connector Twitter Search, uses [Twitter's search API](#) to query Twitter for tweets that match specific parameters. It allows querying for any keyword, location or other query terms.

6.30.1. Registering for Twitter Credentials

In order to successfully configure either Twitter connector, you must first register your application with Twitter and accept their terms of service. The registration process will provide you with the required OAuth tokens you need to access either API. To get the tokens, follow these steps:

1. Make sure you have a Twitter account, and go to <https://dev.twitter.com/> and sign in.
2. After signing in, choose 'My Applications' from the pull down menu at the upper right that shows a thumbnail of your Twitter profile picture (if you have one). Then choose "Create New App" and fill out the required details. The callback field can be skipped, but you must accept the Terms of Service. To save your information, choose "Create Your Twitter Application" to register your application.
3. The next page will contain the Consumer Key and Consumer Secret, which you will need to configure the data source in Fusion.
4. At the bottom of the same page, choose "Create My Access Token".
5. The next page will contain the Access Token and Token Secret, which you will also need to configure the data source in Fusion.

While you need a Twitter account to register an application, you do not use your Twitter username and password to configure this data source. The APIs will only use the Consumer Key, Consumer Secret, Access Token, and Token Secret information as authentication, so store it where you can access it while configuring the data source.

6.30.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
access_token Access Token <i>required</i>	The OAuth Access Token is provided by Twitter when registering the application. type: <code>string</code> minLength: 1
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>

Property	Description
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: boolean default value: 'true'
consumer_key Consumer Key <i>required</i>	The OAuth Consumer Key is provided by Twitter when registering the application. type: string minLength: 1
consumer_secret Consumer Secret <i>required</i>	The OAuth Consumer Secret is provided by Twitter when registering the application. type: string minLength: 1
lang Language	Restrict Tweets to the given language type: string default value: 'en' enum: \{ aa ab ae af ak am an ar as av ay az ba be bg bh bi bm bn bo br bs ca ce ch co cr cs cu cv cy da de dv dz ee el en eo es et eu fa ff fi fj fo fr fy ga gd gl gn gu gv ha he hi ho hr ht hu hy hz ia id ie ig ii ik in io is it iu iw ja ji jv ka kg ki kj kk kl km kn ko kr ks ku kv kw ky la lb lg li ln lo lt lu lv mg mh mi mk ml mn mo mr ms mt my na nb nd ne ng nl nn no nr nv ny oc oj om or os pa pi pl ps pt qu rm rn ro ru rw sa sc sd se sg si sk sl sm sn so sq sr ss st su sv sw ta te tg th ti tk tl tn to tr ts tt tw ty ug uk ur uz ve vi vo wa wo xh yi yo za zh zu }
max_docs Max Documents	The maximum number of documents to pull down, as a long. -1 for no limit type: integer default value: '-1'

Property	Description
queries Queries	A list of queries to perform type: array of string
sleep Sleep	The amount of time, in milliseconds, to sleep when listening so as to not get throttled type: integer default value: '10000'
token_secret Token Secret <i>required</i>	The OAuth Token Secret is provided by Twitter when registering the application. type: string minLength: 1

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
<p>unmapped</p> <p>Unmapped Fields</p>	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

ConnectorDb Configuration

Property	Description
<p>aliases</p> <p>Process Aliases?</p>	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
<p>inlinks</p> <p>Process Inlinks?</p>	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

description

Description

Optional description for datasource instance.

id

Datasource ID

Unique name for datasource instance.

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.31. Twitter Stream Connector and Datasource Configuration

The Twitter Stream connector uses [Twitter's streaming API](#) to continually index Twitter. The datasource can be configured to limit tweets or it can be run indefinitely, until Twitter cuts off your access or you stop the datasource. This connector will only retrieve tweets created after the datasource has been started.

This connector type is "lucid.twitter.search" and the plugin type is "twitter_search".

6.31.1. Registering for Twitter Credentials

In order to successfully configure either Twitter connector, you must first register your application with Twitter and accept their terms of service. The registration process will provide you with the required OAuth tokens you need to access either API. To get the tokens, follow these steps:

1. Make sure you have a Twitter account, and go to <https://dev.twitter.com/> and sign in.
2. After signing in, choose 'My Applications' from the pull down menu at the upper right that shows a thumbnail of your Twitter profile picture (if you have one). Then choose "Create New App" and fill out the required details. The callback field can be skipped, but you must accept the Terms of Service. To save your information, choose "Create Your Twitter Application" to register your application.
3. The next page will contain the Consumer Key and Consumer Secret, which you will need to configure the data source in Fusion.
4. At the bottom of the same page, choose "Create My Access Token".
5. The next page will contain the Access Token and Token Secret, which you will also need to configure the data source in Fusion.

While you need a Twitter account to register an application, you do not use your Twitter username and password to configure this data source. The APIs will only use the Consumer Key, Consumer Secret, Access Token, and Token Secret information as authentication, so store it where you can access it while configuring the data source.

6.31.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
access_token Access Token <i>required</i>	The OAuth Access Token is provided by Twitter when registering the application. type: <code>string</code> minLength: 1

Property	Description
collection Collection	Collection documents will be indexed to. type: <i>string</i> pattern: <code>^[a-zA-Z0-9_-]+\$</code>
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: <i>boolean</i> default value: 'true'
consumer_key Consumer Key <i>required</i>	The OAuth Consumer Key is provided by Twitter when registering the application. type: <i>string</i> minLength: 1
consumer_secret Consumer Secret <i>required</i>	The OAuth Consumer Secret is provided by Twitter when registering the application. type: <i>string</i> minLength: 1
filter_follow Filter Follow	Set of users (user ids) to track type: <i>array of string</i>
filter_locations Filter Locations	Set of bounding boxes (e.g. 'left,bottom,right,top' lat/long coordinates) type: <i>array of string</i>
filter_track Filter Track	Keywords or phrases to track type: <i>array of string</i>

Property	Description
max_docs Max Documents	The maximum number of documents to pull down, as a long. -1 for no limit type: integer default value: '-1'
sleep Sleep	The amount of time, in milliseconds, to sleep when listening so as to not get throttled type: integer default value: '10000'
token_secret Token Secret <i>required</i>	The OAuth Token Secret is provided by Twitter when registering the application. type: string minLength: 1
url URL	The URL used by the Twitter Stream type: string default value: 'https://stream.twitter.com'

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
<p>unmapped</p> <p>Unmapped Fields</p>	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

ConnectorDb Configuration

Property	Description
<p>aliases</p> <p>Process Aliases?</p>	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
<p>inlinks</p> <p>Process Inlinks?</p>	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

description

Description

Optional description for datasource instance.

id

Datasource ID

Unique name for datasource instance.

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.32. Web Connector and Datasource Configuration

The Web connector is used to retrieve data from a Web site using HTTP and starting from a specified URL.

6.32.1. Crawling JavaScript Web sites

As of Fusion 3.1, the Web connector includes a `f.crawlJS`/"Evaluate Javascript" option. When this option is enabled, the connector crawls links rendered from JavaScript evaluation.

Note	This feature requires Oracle JDK with JavaFX , or OpenJDK with OpenJFX .
------	--

6.32.2. Limiting Crawl Scope

The connector works by going to the seed page (the "startURIs" specified in the configuration form), collecting the content for indexing, and extracting any links to other pages. It then follows those links to collect content on other pages, extracting links to those pages, etc.

When creating a Web data source, pay attention to the "Max depth" and "Restrict To Tree" parameters (also known as "c.depth" and "c.restrictToTree" in the REST API). These properties will help limit the scope of your crawl to prevent an "unbounded" crawl that could continue for a long time, particularly if you are crawling a site with links to many pages outside the main site. An unbounded crawl may also cause memory errors in your system.

The connector keeps track of URIs it has seen, and many of the properties relate to managing the resulting database of entries. If the connector finds a standard redirect, it will track that the redirected URI has an alias, and will not re-evaluate the URI on its next runs until the alias expiration has passed. Documents that were found to be duplicates, if de-duplication is enabled, are also added to the alias list and are not re-evaluated until the alias expiration has passed.

Regular expressions can be used to restrict the crawl either by defining URI patterns that should be followed or URI patterns that should not be followed.

Additionally, specific patterns of the URI can also be defined to define URIs that should not be followed.

6.32.3. Extracting Content from Pages

The connector supports several approaches to extracting and filtering content from pages. When analyzing the HTML of a page, the connector can specifically include or exclude elements based on the HTML tag, the tag ID, or the tag class (such as a 'div' tag, or the '#content' tag ID).

Specific tags can be selected to become fields of the document if needed. For example, all content from `<h1>` tags can be pulled into a 'h1' field, and with field mapping be transformed into document titles.

For even more advanced capabilities, you can use jsoup selectors to find elements in the content to include or exclude from the content.

While field mapping is generally a function of the index pipeline, you can define some initial mapping to occur during the crawl. The 'initial mappings' property for each web datasource is pre-defined with three mappings, to move 'fetchDates' to a 'fetchDates_dts' field, to move 'lastModified' to a 'lastModified_dt' field and to move 'length' to a 'length_l' field.

Finally, the crawler is able to do de-duplication of crawled content. You can define a specific field to use for this de-

duplication (such as title, or another field), or you can use the full raw content as the default.

6.32.4. Sitemap Processing

As of Fusion 1.1.2, crawling sitemaps is supported. Simply add the URL(s) of the sitemap to the `f.sitemapURLs` property ("Sitemap URLs" in the UI) and all of the URLs found in a sitemap will be added to the list of URLs to crawl. If your site has a sitemap index, i.e., a sitemap that points to other sitemaps, that is also supported and the URLs found through each sitemap will be added to the list of URLs to crawl.

If you want to configure your datasource to only crawl the sitemap file, you must add the sitemap URL to both the `startLinks` property (because that is a required property for a datasource) and also to the `f.sitemapsURL` property so it is properly treated as a sitemap by the connector when it starts.

6.32.5. Website Authentication

The Web connector supports Basic, Digest, Form and NTLM authentication to websites.

The credentials for a crawl are stored in a credentials file that should be placed in ``fusion/3.1.x/data/connectors/container/lucid.anda/datasourceName`` where the "datasourceName" corresponds to the name given to the datasource. After creating a datasource, this directory should be created for you. The file should be a JSON formatted file, ending with the '.json' file extension. When defining the datasource, you would pass the name of the file with the 'Authentication file' property in the UI (or 'f.credentialsFile' property if using the REST API).

All types of authentication require the credentials file to include a property called "type" which defines the type of authentication to use. After that, the required properties will vary depending on the type of authentication chosen.

Form-based Authentication

To use basic form-based authentication, use "form" for the type. The other properties are:

- `ttl` - The "time to live" for the session that will be created after authentication. This will have the crawler log in again after the specified time so the crawl activity doesn't fail due to an expired session. This value is defined in seconds.
- `action` - The action to take to log in, i.e., the URL for the login form.
- `params` - The parameters for the form, likely the username and password, but any other required properties. In the example below, we are passing two parameters, the 'os_username' and the 'os_password' that are the properties expected by the system we would like to crawl.

Here is an example using form-based authentication:

```
[ {
  "credential" : {
    "type" : "form",
    "ttl" : 300000,
    "action" :
"http://some.server.com/login.action?os_destination=%2Fpages%2Fviewpage.action%3Ftitle%3DAcme%2B5%2BDocumentat
ion%26spaceKey%3DAcme5",
    "params" : {
      "os_username" : "username",
      "os_password" : "password"
    }
  }
} ]
```

Complex Form-based Authentication

Some websites do not manage their own authentication, but rather trust a third-party authority to authenticate the user. An example of this would be websites that use SAML to log in a user via a central single-signon authority. In order to configure fusion to log in to a website like this, use "smartForm" for the type. The other properties are:

- ttl - the "time to live" for the session that will be created after authentication. This will have the crawler re-login after the specified time so the crawl activity doesn't fail due to an expired session. This value is defined in seconds.
- loginUrl - the URL on which the first page that initializes the login chain is located
- params - a list of parameters to use for the form logins, likely the username and password, but could be other required properties. In the example below, we are passing two parameters, the 'os_username' and the 'os_password' that are the properties expected by the system we would like to crawl. Additionally we expect that once that login has happened, that a new form will be presented to the user which then posts back to where we came from. No data need to be entered in this form, which is why we include an empty { } in the params list.

Here is an example using form-based authentication:

```
[ {
  "credential" : {
    "type" : "smartForm",
    "ttl" : 300000,
    "loginUrl" : "http://some.example.com/login",
    "params" : [{
      "os_username" : "username",
      "os_password" : "password"
    }, {
    } ]
  }
} ]
```

In order to figure out what params you need to specify, turn off JavaScript in your browser and walk through the login chain. Though you normally only see a single login form on your screen, you may be surprised to find many more forms you need to submit before you get logged in when JavaScript is not available to perform those form submissions automatically. Each form in that chain needs to be represented in list of `params`. If no user input is required, simply include an empty { }.

Basic and Digest Authentication

Basic and Digest authentication are simple HTTP authentication methods still in use in some places. To use either of these types use "basic" or "digest" in the credentials file for the 'type' property. Other properties are:

- host - the host of the site.
- port - the port, if any.
- userName - the username to use for authentication.
- password - the password for the userName.
- realm - the realm for the site, if any.

Example basic auth configuration:

```
[ {
  "credential" : {
    "type" : "basic",
    "ttl" : 300000,
    "userName" : "usr",
    "password" : "pswd",
    "host": "hostname.examplerdomain.com"
    "port": 443
  }
}
```

NTLM Authentication

To use NTLM authentication, use "ntlm" in the credentials file for the 'type' property. The other properties available are:

- host - the host of the site.
- port - the port, if any.
- userName - the username to use for authentication.
- password - the password for the userName.
- realm - the realm for the site, if any.
- domain - the domain.
- workstation - the workstation, as needed.

Example NTLM credential configuration:

```
[ {"credential" :
  { "type" : "ntlm",
    "ttl" : 300000,
    "port" : 80,
    "host" : "someHost",
    "domain" : "someDomain",
    "userName" : "someUser",
    "password" : "XXXXXXXX"
  }
}
```

6.32.6. Configuration

Tip

When entering configuration values in the UI, use *unescaped* characters, such as `\t` for the tab character. When entering configuration values in the API, use *escaped* characters, such as `\\t` for the tab character.

Connector-specific Properties

Property	Description
f.allowAllCertificates Allow all HTTPS certificates	If false, security checks will be performed on all SSL/TLS certificate signers and origins. This means self-signed certificates would not be supported. type: boolean default value: 'false'
f.appendTrailingSlashToLinks Add trailing slash to link URLs	If true, a trailing '/' will be added to link URLs when the URL does not end in a dot ('.'). type: boolean default value: 'false'

Property	Description
<p>f.basicAuth</p> <p>Basic Authentication</p>	<p>Settings for Basic authentication</p> <p>type: array of object</p> <p>object attributes: \{</p> <p>host (<i>required</i>): \{</p> <p>display name: Host</p> <p>type: string</p> <p>description : The host of the site</p> <p>}</p> <p>id: \{</p> <p>display name: Auth Config id</p> <p>type: string</p> <p>description : Auth Config id</p> <p>}</p> <p>password: \{</p> <p>display name: Password</p> <p>type: string</p> <p>description : The password for the user</p> <p>}</p> <p>port (<i>required</i>): \{</p> <p>display name: Port</p> <p>type: integer</p> <p>description : The port, if any</p> <p>}</p> <p>realm: \{</p> <p>display name: Realm</p> <p>type: string</p> <p>description : The realm for the site, if any</p> <p>}</p> <p>userName: \{</p> <p>display name: User</p> <p>type: string</p> <p>description : The username to use for authentication</p> <p>}</p> <p>}</p>
<p>f.cookieSpec</p> <p>Cookie spec</p>	<p>type: string</p> <p>default value: 'browser-compatibility'</p> <p>enum: \{ browser-compatibility rfc-2965 best-match ignore-all }</p>

Property	Description
f.crawlJS Evaluate Javascript	Crawl links rendered from javascript evaluation. Requires oracle or openjdk with openjfx type: boolean default value: 'false'
f.credentialsFile Authentication credentials filename	Name of the file that contains the credentials for sites that require authentication to access. This file must be located in '\$FUSION_HOME/data/connectors/container/lucid.web/'. type: string
f.defaultCharSet Default character set	Default character set to use when one is not declared in the HTTP headers. type: string default value: 'UTF-8'
f.defaultMIMEType Default MIME type	Default MIME type to use when one is not declared in the HTTP headers. type: string default value: 'application/octet-stream'

Property	Description
<p>f.digestAuth</p> <p>Digest Authentication</p>	<p>Settings for Digest authentication</p> <p>type: array of object</p> <p>object attributes: \{</p> <p>host (<i>required</i>): \{</p> <p>display name: Host</p> <p>type: string</p> <p>description : The host of the site</p> <p>}</p> <p>id: \{</p> <p>display name: Auth Config id</p> <p>type: string</p> <p>description : Auth Config id</p> <p>}</p> <p>password: \{</p> <p>display name: Password</p> <p>type: string</p> <p>description : The password for the user</p> <p>}</p> <p>port (<i>required</i>): \{</p> <p>display name: Port</p> <p>type: integer</p> <p>description : The port, if any</p> <p>}</p> <p>realm: \{</p> <p>display name: Realm</p> <p>type: string</p> <p>description : The realm for the site, if any</p> <p>}</p> <p>userName: \{</p> <p>display name: User</p> <p>type: string</p> <p>description : The username to use for authentication</p> <p>}</p> <p>}</p>
<p>f.discardLinkURLQueries</p> <p>Discard queries in link URLs</p>	<p>If true, query parameters found in URLs will be removed before being added to the discovery queue.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
f.excludeSelectors Jsoup exclusive selectors	Jsoup-formatted selectors for elements to exclude from the crawled content. Syntax for jsoup selectors is available at http://jsoup.org/apidocs/org/jsoup/select/Selector.html . type: <code>array of string</code>
f.excludeTagClasses Excluded tag classes	HTML tag classes of elements to exclude from the crawled content. type: <code>array of string</code>
f.excludeTagIDs Excluded tag IDs	HTML tag IDs of elements to exclude from the crawled content. type: <code>array of string</code>
f.excludeTags Excluded tags	HTML tag names of elements to exclude from the crawled content. type: <code>array of string</code>
f.filteringRootTags Root elements to filter	Root HTML elements whose child elements will be used to extract content. By default 'body' and 'head' elements are already included. type: <code>array of string</code> default value: 'bodyhead'

Property	Description
<p>f.formAuth</p> <p>Form Authentication</p>	<p>Settings for Form based authentication</p> <p>type: array of object</p> <p>object attributes: \{</p> <p>action (<i>required</i>): \{</p> <p>display name: URL</p> <p>type: string</p> <p>description : The URL of the authentication endpoint</p> <p>\}</p> <p>id: \{</p> <p>display name: Auth Config id</p> <p>type: string</p> <p>description : Auth Config id</p> <p>\}</p> <p>params: \{</p> <p>display name: Parameters</p> <p>type: object</p> <p>\ :</p> <p>\}</p> <p>password: \{</p> <p>display name: Password</p> <p>type: string</p> <p>description : The password to use for the authentication request. This will be copied into the "Parameters" using the "Password Parameter" name as the key</p> <p>\}</p> <p>passwordParamName: \{</p> <p>display name: Password Parameter</p> <p>type: string</p> <p>description : Name of the parameter containing the password</p> <p>\}</p> <p>ttl: \{</p> <p>display name: TTL</p> <p>type: number</p> <p>description : The "time to live" for the session that will be created after authentication.</p> <p>\}</p> <p>\}</p>
<p>f.includeSelectors</p> <p>Jsoup inclusive selectors</p>	<p>Jsoup-formatted selectors for elements to include in the crawled content.</p> <p>type: array of string</p>

Property	Description
f.includeTagClasses Included tag classes	HTML tag classes of elements to include in the crawled content. type: <code>array of string</code>
f.includeTagIDs Included tag IDs	HTML tag IDs of elements to include in the crawled content. type: <code>array of string</code>
f.includeTags Included tags	HTML tag names of elements to include in the crawled content. type: <code>array of string</code>
f.jsAjaxTimeout AJAX Timeout	The time in milliseconds after which an AJAX request will be ignored when considering whether all AJAX requests have completed. Maximum: 180,000ms i.e. 3 minutes type: <code>integer</code> default value: '10000' exclusiveMaximum: false exclusiveMinimum: false maximum: 180000 minimum: -1
f.jsPageLoadTimeout Timeout	The time to wait in milliseconds for a page load to complete. If the timeout is -1, page loads can be indefinite. Maximum: 180,000ms i.e. 3 minutes type: <code>integer</code> default value: '10000' exclusiveMaximum: false exclusiveMinimum: false maximum: 180000 minimum: -1

Property	Description
f.jsScriptTimeout Script Timeout	The time to wait in milliseconds wait for an asynchronous script to finish execution. If the timeout is -1, then the script will be allowed to run indefinitely. Maximum: 30,000ms type: integer default value: '10000' exclusiveMaximum: false exclusiveMinimum: false maximum: 180000 minimum: -1
f.maxSizeBytes Max file size (bytes)	Maximum size, in bytes, of a document to fetch. type: integer default value: '4194304'

Property	Description
<p>f.ntlmAuth</p> <p>NTLM Authentication</p>	<p>Settings for NTLM authentication</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> domain : \{ <ul style="list-style-type: none"> display name: Domain type: string description : The NTLM Domain host (required) : \{ <ul style="list-style-type: none"> display name: Host type: string description : The host of the site id : \{ <ul style="list-style-type: none"> display name: Auth Config id type: string description : Auth Config id password : \{ <ul style="list-style-type: none"> display name: Password type: string description : The password for the user port (required) : \{ <ul style="list-style-type: none"> display name: Port type: integer description : The port, if any realm : \{ <ul style="list-style-type: none"> display name: Realm type: string description : The realm for the site, if any userName : \{ <ul style="list-style-type: none"> display name: User type: string description : The username to use for authentication workstation : \{ <ul style="list-style-type: none"> display name: Workstation type: string description : The NTLM Workstation name <p>\}</p>

Property	Description
f.obeyRobots Obey robots.txt	If true, Allow, Disallow and other rules found in a robots.txt file will be obeyed. type: <code>boolean</code> default value: 'true'
f.obeyRobotsDelay Obey robots.txt Crawl-Delay	If true, Crawl-Delay rules in robots.txt will be obeyed. Disabling this option will speed up crawling, but is considered negative behavior for sites you do not control. type: <code>boolean</code> default value: 'true'
f.proxy HTTP proxy address	Address of the HTTP proxy, if required. This should be entered in the format host:port. type: <code>string</code>
f.respectMetaEquivRedirects Respect refresh redirects	If true, the connector will follow metatags with refresh redirects such as . type: <code>boolean</code> default value: 'false'

Property	Description
<p>f.samlAuth</p> <p>SAML Authentication</p>	<p>Settings for SAML based authentication</p> <p>type: array of object</p> <p>object attributes: \{</p> <p>action (<i>required</i>): \{</p> <p>display name: URL</p> <p>type: string</p> <p>description : The URL of the authentication endpoint</p> <p>}</p> <p>id: \{</p> <p>display name: Auth Config id</p> <p>type: string</p> <p>description : Auth Config id</p> <p>}</p> <p>params : \{</p> <p>display name: Parameters</p> <p>type: object</p> <p>\ :</p> <p>}</p> <p>password : \{</p> <p>display name: Password</p> <p>type: string</p> <p>description : The password to use for the authentication request. This will be copied into the "Parameters" using the "Password Parameter" name as the key</p> <p>}</p> <p>passwordParamName : \{</p> <p>display name: Password Parameter</p> <p>type: string</p> <p>description : Name of the parameter containing the password</p> <p>}</p> <p>ttl : \{</p> <p>display name: TTL</p> <p>type: number</p> <p>description : The "time to live" for the session that will be created after authentication.</p> <p>}</p> <p>}</p>

Property	Description
f.scrapeLinksBeforeFiltering Scrape links before filtering	If true, links will be extracted from documents before any other document processing has occurred. By default, links are extracted after all other document processing. type: <code>boolean</code> default value: <code>'false'</code>
f.selectorFields Jsoup selector fields	List of Jsoup selectors for elements to put into their separate field in the index. The field will have the same name as the element. Syntax for jsoup selectors is available at http://jsoup.org/apidocs/org/jsoup/select/Selector.html . type: <code>array of string</code>
f.sitemapURLs Sitemap URLs	URLs for sitemaps, to be used a basis for link discovery. Rules found in sitemaps will not be processed. type: <code>array of string</code>
f.tagClassFields Tag-class fields	HTML tag classes of elements to put into their own field in the index. The field will have the same name as the tag class. type: <code>array of string</code>
f.tagFields Tag fields	HTML tags of elements to put into their own field in the index. The field will have the same name as the tag. type: <code>array of string</code>
f.tagIDFields Tag-ID fields	HTML tag IDs of elements to put into their own field in the index. The field will have the same name as the tag ID. type: <code>array of string</code>
f.timeoutMS Connection timeout (ms)	Time in milliseconds to wait for server response. type: <code>integer</code> default value: <code>'10000'</code>
f.userAgentEmail HTTP user-agent email address	Email address to use as part of connector identification. type: <code>string</code>

Property	Description
f.userAgentName HTTP user-agent name	Name the connector should use when identifying itself to a website in order to crawl it. type: <code>string</code> default value: 'Lucidworks-Anda/2.0'
f.userAgentWebAddr HTTP user-agent web address	Web address to use as part of connector identification. type: <code>string</code>

Crawl Authorization

Property	Description
f.proxy HTTP proxy address	Address of the HTTP proxy, if required. This should be entered in the format host:port. type: <code>string</code>
f.allowAllCertificates Allow all HTTPS certificates	If false, security checks will be performed on all SSL/TLS certificate signers and origins. This means self-signed certificates would not be supported. type: <code>boolean</code> default value: 'false'
f.obeyRobots Obey robots.txt	If true, Allow, Disallow and other rules found in a robots.txt file will be obeyed. type: <code>boolean</code> default value: 'true'

Property	Description
<p>f.basicAuth</p> <p>Basic Authentication</p>	<p>Settings for Basic authentication</p> <p>type: array of object</p> <p>object attributes: \{</p> <p>host (required): \{</p> <p>display name: Host</p> <p>type: string</p> <p>\}</p> <p>id: \{</p> <p>display name: Auth Config id</p> <p>type: string</p> <p>\}</p> <p>password: \{</p> <p>display name: Password</p> <p>type: string</p> <p>\}</p> <p>port (required): \{</p> <p>display name: Port</p> <p>type: integer</p> <p>\}</p> <p>realm: \{</p> <p>display name: Realm</p> <p>type: string</p> <p>\}</p> <p>userName: \{</p> <p>display name: User</p> <p>type: string</p> <p>\}</p> <p>\}</p>

Property	Description
<p>f.digestAuth</p> <p>Digest Authentication</p>	<p>Settings for Digest authentication</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> host (required): \{ <ul style="list-style-type: none"> display name: Host type: string id: \{ <ul style="list-style-type: none"> display name: Auth Config id type: string password: \{ <ul style="list-style-type: none"> display name: Password type: string port (required): \{ <ul style="list-style-type: none"> display name: Port type: integer realm: \{ <ul style="list-style-type: none"> display name: Realm type: string userName: \{ <ul style="list-style-type: none"> display name: User type: string <p>\}</p>

Property	Description
<p>f.ntlmAuth</p> <p>NTLM Authentication</p>	<p>Settings for NTLM authentication</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> domain : \{ <ul style="list-style-type: none"> display name: Domain type: string host (required) : \{ <ul style="list-style-type: none"> display name: Host type: string id : \{ <ul style="list-style-type: none"> display name: Auth Config id type: string password : \{ <ul style="list-style-type: none"> display name: Password type: string port (required) : \{ <ul style="list-style-type: none"> display name: Port type: integer realm : \{ <ul style="list-style-type: none"> display name: Realm type: string userName : \{ <ul style="list-style-type: none"> display name: User type: string workstation : \{ <ul style="list-style-type: none"> display name: Workstation type: string <p>\}</p>

Property	Description
<p>f.formAuth</p> <p>Form Authentication</p>	<p>Settings for Form based authentication</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> action (<i>required</i>): \{ <ul style="list-style-type: none"> display name: URL type: string id: \{ <ul style="list-style-type: none"> display name: Auth Config id type: string params: \{ <ul style="list-style-type: none"> display name: Parameters type: object password: \{ <ul style="list-style-type: none"> display name: Password type: string passwordParamName: \{ <ul style="list-style-type: none"> display name: Password Parameter type: string ttl: \{ <ul style="list-style-type: none"> display name: TTL type: number <p>\}</p>

Property	Description
f.samlAuth SAML Authentication	Settings for SAML based authentication type: array of object object attributes: \{ action (<i>required</i>): \{ display name: URL type: string } id : \{ display name: Auth Config id type: string } params : \{ display name: Parameters type: object } password : \{ display name: Password type: string } passwordParamName : \{ display name: Password Parameter type: string } ttl : \{ display name: TTL type: number } }
f.credentialsFile Authentication credentials filename	Name of the file that contains the credentials for sites that require authentication to access. This file must be located in '\$FUSION_HOME/data/connectors/container/lucid.web/'. type: string

Link Discovery

Property	Description
f.sitemapURLs Sitemap URLs	URLs for sitemaps, to be used a basis for link discovery. Rules found in sitemaps will not be processed. type: array of string

Property	Description
restrictToTreeAllowSubdomains Allow sub-domains in restrictToTree	If true, any sub-domain will be allowed, even if the crawl is restricted to the tree of items found below the start links. type: boolean default value: 'false'
restrictToTreeUseHostAndPath Use paths in restrictToTree	If true, the path in start links will be used to restrict items fetched. For example, if the start link is 'http://host.com/US', this option will limit all followed URLs to this path. type: boolean default value: 'false'
restrictToTreeIgnoredHostPrefixes Ignored host prefixes	List of host prefixes to ignore when checking links for restrictToTree link-legality checks. For example, 'www.' can be ignored so links with the same domain are allowed. type: array of string default value: [` www. `]
f.respectMetaEquivRedirects Respect refresh redirects	If true, the connector will follow metatags with refresh redirects such as . type: boolean default value: 'false'

Limit Documents

Property	Description
restrictToTree Restrict to sub-directories and child pages	If true, only URLs that match the startLinks URL domain will be followed type: boolean default value: 'true'

Property	Description
depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1'
maxItems Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: integer default value: '-1'
includeExtensions Included file extensions	File extensions to be fetched. This will limit this datasource to only these file extensions. type: array of string
includeRegexes Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: array of string
excludeExtensions Excluded file extensions	File extensions that should not to be fetched. This will limit this datasource to all extensions except this list. type: array of string
excludeRegexes Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string
f.maxSizeBytes Max file size (bytes)	Maximum size, in bytes, of a document to fetch. type: integer default value: '4194304'

Crawler ID

Property	Description
f.userAgentName HTTP user-agent name	Name the connector should use when identifying itself to a website in order to crawl it. type: <i>string</i> default value: 'Lucidworks-Anda/2.0'
f.userAgentEmail HTTP user-agent email address	Email address to use as part of connector identification. type: <i>string</i>
f.userAgentWebAddr HTTP user-agent web address	Web address to use as part of connector identification. type: <i>string</i>

Document Parsing

Property	Description
f.appendTrailingSlashToLinks Add trailing slash to link URLs	If true, a trailing '/' will be added to link URLs when the URL does not end in a dot ('.'). type: <i>boolean</i> default value: 'false'
f.defaultCharSet Default character set	Default character set to use when one is not declared in the HTTP headers. type: <i>string</i> default value: 'UTF-8'
f.defaultMIMEType Default MIME type	Default MIME type to use when one is not declared in the HTTP headers. type: <i>string</i> default value: 'application/octet-stream'
f.filteringRootTags Root elements to filter	Root HTML elements whose child elements will be used to extract content. By default 'body' and 'head' elements are already included. type: <i>array of string</i> default value: [` body ` ` head `]

Property	Description
f.scrapeLinksBeforeFiltering Scrape links before filtering	If true, links will be extracted from documents before any other document processing has occurred. By default, links are extracted after all other document processing. type: <code>boolean</code> default value: 'false'
f.includeTags Included tags	HTML tag names of elements to include in the crawled content. type: <code>array of string</code>
f.includeTagClasses Included tag classes	HTML tag classes of elements to include in the crawled content. type: <code>array of string</code>
f.includeTagIDs Included tag IDs	HTML tag IDs of elements to include in the crawled content. type: <code>array of string</code>
f.includeSelectors Jsoup inclusive selectors	Jsoup-formatted selectors for elements to include in the crawled content. type: <code>array of string</code>
f.excludeTags Excluded tags	HTML tag names of elements to exclude from the crawled content. type: <code>array of string</code>
f.excludeTagClasses Excluded tag classes	HTML tag classes of elements to exclude from the crawled content. type: <code>array of string</code>
f.excludeTagIDs Excluded tag IDs	HTML tag IDs of elements to exclude from the crawled content. type: <code>array of string</code>

Property	Description
f.excludeSelectors Jsoup exclusive selectors	Jsoup-formatted selectors for elements to exclude from the crawled content. Syntax for jsoup selectors is available at http://jsoup.org/apidocs/org/jsoup/select/Selector.html . type: array of string
f.tagFields Tag fields	HTML tags of elements to put into their own field in the index. The field will have the same name as the tag. type: array of string
f.tagIDFields Tag-ID fields	HTML tag IDs of elements to put into their own field in the index. The field will have the same name as the tag ID. type: array of string
f.tagClassFields Tag-class fields	HTML tag classes of elements to put into their own field in the index. The field will have the same name as the tag class. type: array of string
f.selectorFields Jsoup selector fields	List of Jsoup selectors for elements to put into their separate field in the index. The field will have the same name as the element. Syntax for jsoup selectors is available at http://jsoup.org/apidocs/org/jsoup/select/Selector.html . type: array of string

Crawl Performance

Property	Description
chunkSize Fetch batch size	The number of items to batch for each round of fetching. A higher value can make crawling faster, but memory usage is also increased. The default is 1. type: integer default value: '1'

Property	Description
fetchThreads Fetch threads	The number of threads to use during fetching. The default is 5. type: integer default value: '5'
fetchDelayMS Fetch delay	Number of milliseconds to wait between fetch requests. The default is 0. This property can be used to throttle a crawl if necessary. type: integer default value: '0'
fetchDelayMSPerHost Fetch delay per host	If true, the 'Fetch delay (ms)' property will be applied for each host. type: boolean default value: 'true'
emitThreads Emit threads	The number of threads used to send documents from the connector to the index pipeline. The default is 5. type: integer default value: '5'
failFastOnStartLinkFailure Fail crawl if start links are invalid	If true, when Fusion cannot connect to any of the provided start links, the crawl is stopped and an exception logged. type: boolean default value: 'true'
retryEmit Retry emits	Set to true for emit batch failures to be retried on a document-by-document basis. type: boolean default value: 'true'

Property	Description
parserRetryCount Max Parser Retries	The maximum number of times the configured parser will try getting content before giving up type: integer default value: '0' exclusiveMaximum: true exclusiveMinimum: false maximum: 5 minimum: 0
f.timeoutMS Connection timeout (ms)	Time in milliseconds to wait for server response. type: integer default value: '10000'
f.obeyRobotsDelay Obey robots.txt Crawl-Delay	If true, Crawl-Delay rules in robots.txt will be obeyed. Disabling this option will speed up crawling, but is considered negative behavior for sites you do not control. type: boolean default value: 'true'

Dedupe

Property	Description
dedupe Dedupe documents	If true, documents will be deduplicated. Deduplication can be done based on an analysis of the content, on the content of a specific field, or by a JavaScript function. If neither a field nor a script are defined, content analysis will be used. type: boolean default value: 'false'

Property	Description
dedupeSaveSignature Save dedupe signature	If true,the signature used for dedupe will be stored in a 'dedupeSignature_s' field. Note this may cause errors about 'immense terms' in that field. type: boolean default value: 'false'
dedupeField Dedupe field	Field to be used for dedupe. Define either a field or a dedupe script, otherwise the full raw content of each document will be used. type: string
dedupeScript Dedupe script	Custom javascript to dedupe documents. The script must define a 'genSignature(content)\{' function, but can use any combination of document fields. The function must return a string. type: string

Javascript Evaluation

Property	Description
f.crawlJS Evaluate Javascript	Crawl links rendered from javascript evaluation. Requires oracle or openjdk with openjfx type: boolean default value: 'false'
f.jsPageLoadTimeout Timeout	The time to wait in milliseconds for a page load to complete. If the timeout is -1, page loads can be indefinite. Maximum: 180,000ms i.e. 3 minutes type: integer default value: '10000' exclusiveMaximum: false exclusiveMinimum: false maximum: 180000 minimum: -1

Property	Description
f.jsScriptTimeout Script Timeout	The time to wait in milliseconds wait for an asynchronous script to finish execution. If the timeout is -1, then the script will be allowed to run indefinitely. Maximum: 30,000ms type: integer default value: '10000' exclusiveMaximum: false exclusiveMinimum: false maximum: 180000 minimum: -1
f.jsAjaxTimeout AJAX Timeout	The time in milliseconds after which an AJAX request will be ignored when considering whether all AJAX requests have completed. Maximum: 180,000ms i.e. 3 minutes type: integer default value: '10000' exclusiveMaximum: false exclusiveMinimum: false maximum: 180000 minimum: -1

Recrawl Rules

Property	Description
refreshAll Recrawl all items	Set to true to always recrawl all items found in the crawldb. type: boolean default value: 'true'

Property	Description
refreshStartLinks Recrawl start links	Set to true to recrawl items specified in the list of start links. type: boolean default value: 'false'
refreshErrors Recrawl errors	Set to true to recrawl items that failed during the last crawl. type: boolean default value: 'false'
refreshOlderThan Recrawl age	Number of seconds to recrawl items whose last fetched date is longer ago than this value. type: integer default value: '-1'
refreshIDPrefixes Recrawl ID prefixes	A prefix to recrawl all items whose IDs begin with this value. type: array of string
refreshIDRegexes Recrawl ID regexes	A regular expression to recrawl all items whose IDs match this pattern. type: array of string
refreshScript Recrawl script	A JavaScript function ('shouldRefresh()') to customize the items recrawled. type: string
forceRefresh Force recrawl	Set to true to recrawl all items even if they have not changed since the last crawl. type: boolean default value: 'false'

Property	Description
delete Delete dead URIs	Set to true to remove documents from the index when they can no longer be accessed as unique documents. type: <code>boolean</code> default value: <code>'true'</code>
deleteErrorsAfter Fetch failure allowance	Number of fetch failures to tolerate before removing a document from the index. The default of -1 means no fetch failures will be removed. type: <code>integer</code> default value: <code>'-1'</code>

Crawl History

Property	Description
f.discardLinkURLQueries Discard queries in link URLs	If true, query parameters found in URLs will be removed before being added to the discovery queue. type: <code>boolean</code> default value: <code>'false'</code>
retainOutlinks Retain links in the crawl db	Set to true for links found during fetching to be stored in the crawl db. This increases precision in certain recrawl scenarios, but requires more memory and disk space. type: <code>boolean</code> default value: <code>'false'</code>
aliasExpiration Alias expiration	The number of crawls after which an alias will expire. The default is 1 crawl. type: <code>integer</code> default value: <code>'1'</code>

Property	Description
<p>crawlDBType</p> <p>Crawl database type</p>	<p>The type of crawl database to use, in-memory or on-disk.</p> <p>type: string</p> <p>default value: 'in-memory'</p> <p>enum: \{ in-memory on-disk }</p>
<p>indexCrawlDBToSolr</p> <p>Index crawl database to Solr</p>	<p>EXPERIMENTAL: Set to true to index the crawl-database into a 'crawldb_' collection in Solr.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Field Mapping

Property	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "charSet", "target" ⇒ "charSet_s" } { "operation" ⇒ "move", "source" ⇒ "fetchDate", "target" ⇒ "fetchDate_dt" } { "operation" ⇒ "move", "source" ⇒ "lastModified", "target" ⇒ "lastModified_dt" } { "operation" ⇒ "move", "source" ⇒ "signature", "target" ⇒ "dedupeSignature_s" } { "operation" ⇒ "move", "source" ⇒ "contentSignature", "target" ⇒ "signature_s" } { "operation" ⇒ "move", "source" ⇒ "length", "target" ⇒ "length_l" } { "operation" ⇒ "move", "source" ⇒ "mimeType", "target" ⇒ "mimeType_s" } { "operation" ⇒ "move", "source" ⇒ "parent", "target" ⇒ "parent_s" } { "operation" ⇒ "move", "source" ⇒ "owner", "target" ⇒ "owner_s" } { "operation" ⇒ "move", "source" ⇒ "group", "target" ⇒ "group_s" }</code></p> <p>object attributes: <code>\{</code> operation : <code>\{</code> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: <code>\{ copy move delete set add keep }</code> <code>}</code> <code>\{</code> source` _ (required)_ : <code>\{</code> + display name: Source Field + type: `string` + description : The name of the field to be mapped. + <code>}</code> + target` : <code>\{</code> + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + <code>}</code> + <code>}</code></p>

Property	Description
reservedFieldsMappingAllowed Allow System Fields Mapping?	type: boolean default value: 'false'
unmapped Unmapped Fields	If fields do not match any of the field mapping rules, these rules will apply. type: object object attributes: \{ <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }

```

    } +
    'source' : \{ +
      display name: Source Field +
      type: 'string' +
      description : The name of the field to be
      mapped. +
    } +
    'target' : \{ +
      display name: Target Field +
      type: 'string' +
      description : The name of the field to be
      mapped to. +
    } +
  }

```

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB. type: boolean default value: 'false'

Property	Description
inlinks Process Inlinks?	Keep track of incoming links. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.33. Websphere Connector and Datasource Configuration

WebSphere Commerce Suite, a.k.a IBM WebSphere is an IBM product that offers full stack software for companies that do eCommerce business. User interactions with the IBM storefront eCommerce website can be captured as Fusion signals and can be used for product recommendations on the storefront.

WebSphere search is powered by Solr and the product is distributed with Solr jar files and deployment scripts for different development and production configurations. In order to use the WebSphere-deployed Solr with Fusion index and query pipelines, is it necessary to install and configure an additional set of Solr plugins. The plugins, installation and configuration information are publically available from:

- <https://github.com/LucidWorks/fusion-solr-plugins>

6.34. Windows Share Connector and Datasource Configuration



width=600%

The Windows Share connector can access content in a Windows Share or Server Message Block (SMB)/Common Internet File System (CIFS) filesystem.

See this tutorial about configuring a Windows Share datasource and enabling security trimming:

6.34.1. Access Control Lists (ACLs)

The connector is able to retrieve and store ACL details when crawling with the 'smb' type. There are several properties available to define how the datasource should read the user and group information found in Active Directory, and when security trimming is enabled, document results will take user authorizations into consideration.

For each document, the `acl` field is populated with data that can be used at search time to filter the results so that only people that have been granted access at the user level or through group membership can see them. Two kinds of tokens are stored: Allow and Deny. The format used is as follows:

Allow: ``WINA<SID>``

Deny: ``WIND<SID>``

Where `SID` is the security identifier commonly used in Microsoft Windows systems. There are some well known SIDs that can be used in the `acl` field to make documents that are crawled through some other mechanism than by using SMB data source behave, from the `acl pow`, the same way as the crawled SMB content:

SID	Description
S-1-1-0	Everyone.
S-1-5-domain-500	A user account for the system administrator. By default, it is the only user account that is given full control over the system.
S-1-5-domain-512	Domain Admins: a global group whose members are authorized to administer the domain. By default, the Domain Admins group is a member of the Administrators group on all computers that have joined a domain.
S-1-5-domain-513	Domain Users.

Note that some of the listed SIDs contain a `domain` token. This means that the actual SIDs differ from system to system. To find out the SIDs for particular user in particular system you can use the information provided by the Windows command line tool `whoami` by executing command `whoami /all`.

You can populate the `acl` field in your documents with these Windows SIDs to make them searchable in Fusion. For example, if you wanted to make some documents available to "Everyone" you would populate the `acl` field with the `WINAS-1-1-0` token. If you wanted to make all docs from one data source available to everybody you can use the literal definitions in the data source configuration.

6.34.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
Property	Description
add_failed_docs Add failed documents	Set to true to add documents even if they partially fail processing. Failed documents will be added with as much metadata as available, but may not include all expected fields. type: <code>boolean</code> default value: 'false'
bounds Crawl bounds	Limits the crawl to a specific directory sub-tree, hostname or domain. type: <code>string</code> default value: 'tree' enum: <code>\{ tree host domain none \}</code>
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>
commit_on_finish Solr commit on finish	Set to true for a request to be sent to Solr after the last batch has been fetched to commit the documents to the index. type: <code>boolean</code> default value: 'true'

Property	Description
crawl_depth Max crawl depth	Number of levels in a directory or site tree to descend for documents. type: integer default value: '-1' exclusiveMinimum: false minimum: -1
crawl_item_timeout Fetch timeout	Time in milliseconds to fetch any individual document. type: integer default value: '600000' exclusiveMinimum: true minimum: 0
enable_security_trimming Enable security trimming	Set to true to fetch and index access control information from files. type: object object attributes: \{ } }
exclude_paths Exclusive regexes	Regular expressions for URI patterns to exclude. This will limit this datasource to only URIs that do not match the regular expression. type: array of string
include_extensions Included file extensions	List the file extensions to be fetched. Note: Files with possible matching MIME types but non-matching file extensions will be skipped. Extensions should be listed without periods, using whitespace to separate items (e.g., 'pdf zip'). type: array of string

Property	Description
include_paths Inclusive regexes	Regular expressions for URI patterns to include. This will limit this datasource to only URIs that match the regular expression. type: <code>array of string</code>
index_directories Index directories	Set to true to add directories to the index as documents. If set to false, directories will not be added to the index, but they will still be traversed for documents. type: <code>boolean</code> default value: 'false'
kerberos_keytab Kerberos keytab	Full path to the Kerberos keytab file. type: <code>string</code>
kerberos_user Kerberos principal	Kerberos principal name, i.e., ' username@YOUR-REALM.COM '. type: <code>string</code>
max_bytes Maximum file size (bytes)	Maximum size (in bytes) of documents to fetch or -1 for unlimited file size. type: <code>integer</code> default value: '10485760' exclusiveMinimum: false minimum: -1
max_docs Max items	Maximum number of documents to fetch. The default (-1) means no limit. type: <code>integer</code> default value: '-1' exclusiveMinimum: false minimum: -1

Property	Description
max_threads Fetch threads	The maximum number of threads to use for fetching data. Note: Each thread will create a new connection to the repository, which may make overall throughput faster, but this also requires more system resources, including CPU and memory. type: <i>integer</i> default value: '1'
maximum_connections Maximum fetch connections	Maximum number of concurrent connections to the filesystem. A large number of documents could cause a large number of simultaneous connections to the repository and lead to errors or degraded performance. In some cases, reducing this number may help performance issues. type: <i>integer</i> default value: '10000'
password Password	Password for the user. Do not set a password if Kerberos authentication is used. type: <i>string</i>
url Start link <i>required</i>	A starting URI for this datasource. The URI must be fully-qualified, and include the protocol, host, port and path, as appropriate. type: <i>string</i> minLength: 1 pattern: ..
username Username	A user in the Windows domain with READ permissions for the Windows Share. Do not set username if Kerberos authentication is used. type: <i>string</i>

Property	Description
verify_access Validate access	Set to true to require successful connection to the filesystem before saving this datasource. type: <code>boolean</code> default value: 'true'
windows_domain Windows Domain	Authentication domain for the user. Not required if using Kerberos authentication. type: <code>string</code>
with_kerberos Use Kerberos Authentication	Select to use Kerberos for authentication instead of username and password. Kerberos authentication requires information for the advanced properties 'Kerberos principal' and 'Kerberos keytab'. If 'false', Username and Password must be entered. type: <code>boolean</code> default value: 'false'

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>default value: <code>{ "operation" ⇒ "move", "source" ⇒ "fetch_time", "target" ⇒ "fetch_time_dt" }</code> <code>{ "operation" ⇒ "move", "source" ⇒ "ds:description", "target" ⇒ "description" }</code></p> <p>object attributes: <code>{</code></p> <ul style="list-style-type: none"> operation : <code>{</code> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: <code>{ copy move delete set add keep }</code> <pre> } + 'source' _ (required)_ : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
unmapped Unmapped Fields	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep }</p> <pre> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

ConnectorDb Configuration

Property	Description
aliases Process Aliases?	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
inlinks Process Inlinks?	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property

Description

description

Description

Optional description for datasource instance.

id

Datasource ID

Unique name for datasource instance.

parserId

Parser

The parser used to process raw content.

pipeline

Pipeline ID

The index pipeline used to process documents.

6.35. Zendesk Connector and Datasource Configuration



width=600%

The Zendesk connector uses the [Zendesk REST API](#) to retrieve tickets and their associated comments and attachments from a Zendesk repository.

It retrieves all tickets with all fields (e.g., customer, assignee, priority, status) as well as access restrictions for users and groups. ACLs can be used for security trimming at query time. The types of items retrieved are:

- tickets and their associated metrics: e.g., time elapsed until first response, time to close.
- ticket comment counts, comment ids
- ticket comment attachment URLs

Items retrieved are returned as individual Solr documents, therefore, for a given Zendesk ticket, there will be multiple documents:

- the Zendesk ticket itself
- one document per comment
- one document per comment attachment

Documents have fields for Zendesk type and reference field to parent documents, e.g., a comment document will have field "ticket_id" pointing back to the Zendesk ticket.

Incremental recrawls allow updates to the Fusion collection to add new tickets and record further changes to existing tickets without having to retrieve the entire Zendesk repository contents.

6.35.1. Authorization

The Zendesk user must have administrator privileges in order to retrieve all tickets and associated information. All communication with the Zendesk API is encrypted with SSL.

6.35.2. Required Configuration Properties

A Zendesk datasource must be configured with the following properties:

- Authentication Key - username or email
- Authentication Value - the password or API token
- Token Auth - a flag to indicate whether the auth key/value should be treated as username/password or email/token
- Base URL - the API url to an instance of the Zendesk API
- Organization ID - set to restrict indexing to only tickets that belong to the Organization

6.35.3. Configuration Properties

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Property	Description
auth_key Authentication Key <i>required</i>	An email address of a valid Zendesk user which must have an admin role. type: <code>string</code> minLength: 1
auth_value Authentication Value <i>required</i>	This value can be a user password, or an API token. If an API token is used, the Token Auth must be set to true. type: <code>string</code>
base_url Base API URL <i>required</i>	A properly formatted v2 Zendesk API URL. Example: https://yourcompany.zendesk.com/api/v2 type: <code>string</code>
collection Collection	Collection documents will be indexed to. type: <code>string</code> pattern: <code>^[a-zA-Z0-9_-]+\$</code>
organization_id Organization ID	The ID of a Zendesk Organization. This optional setting will allow the Connector to skip records that don't have the same Organization ID. type: <code>string</code>
use_token_auth Token Auth <i>required</i>	Used to determine how the authentication credentials will be formatted. If set to true, the <code>auth_key</code> value will be formatted according to the Zendesk token auth specifications (/). type: <code>boolean</code> default value: <code>'true'</code>

Field Mapping

Initial Mappings	Description
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{ operation : \{ display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } } } 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } }</p>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Initial Mappings	Description
<p>unmapped</p> <p>Unmapped Fields</p>	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

ConnectorDb Configuration

Property	Description
<p>aliases</p> <p>Process Aliases?</p>	<p>Keep track of original URI-s that resolved to the current URI. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>
<p>inlinks</p> <p>Process Inlinks?</p>	<p>Keep track of incoming links. This negatively impacts performance and size of DB.</p> <p>type: boolean</p> <p>default value: 'false'</p>

Property	Description
inv_aliases Process Inverted Aliases?	Keep track of target URI-s that the current URI resolves to. This negatively impacts performance and size of DB. type: <code>boolean</code> default value: 'false'
type Implementation Class Name	Fully qualified class name of ConnectorDb implementation. type: <code>string</code> default value: 'com.lucidworks.connectors.db.impl.MapDbConnectorDb' minLength: 1

General Configuration

Property	Description
description Description	Optional description for datasource instance.
id Datasource ID	Unique name for datasource instance.
pipeline Pipeline ID	The index pipeline used to process documents.

Chapter 7. Parsers

Parsers were introduced in Fusion 3.0 to provide more fine-grained configuration for inbound data. Parsers are configured in stages, much like index pipelines and query pipelines. They can include conditional parsing and nested parsing, and can be configured via the Fusion UI or the Parsers API.

Connectors receive the inbound data, convert it into a byte stream, and send the byte stream through the configured parsing stages. The stream moves through the parser stage by stage until it has been successfully parsed, then proceeds to the index pipeline.

Each parsing stage evaluates whether the inbound stream matches the stage's default media type or filename extension. The first stage that finds a match can output one or both of the following:

- Zero or more pipeline documents for consumption by the index pipeline
- Zero or more new input streams for re-parsing

This recursive approach is useful for containers (zip or tar files, for example). The output of the container parsing may be another container or a stream of uncompressed content which requires its own parsing.

There are a few static fields that impact the overall configuration and are accessible whenever you have selected the parser in the Index Workbench:

- Document ID Source Field
- Enable Automatic Media Type Detection
- Maximum Recursion Depth

7.1. Built-in parsing stages

These stages are available for configuration:

- HTML parser stage
- XML parser stage
- CSV parser stage
- JSON parser stage
- Text parser stage
- Archive parser stage
- Apache Tika parser stage
- Fallback parser stage

Datasources which use connectors that retrieve fixed-structure content (like Twitter or Jira) have hard-coded parsers and do not expose any configurable parser options.

7.1.1. HTML parser stage

This parser stage processes the following HTML elements:

- `<title>`
- `<body>` (with tags removed)
- `<meta>`
- `<a>` and `<link>`

Additionally, you can configure [JSoup selectors](#) to extract specific HTML and CSS elements from a document and map them to PipelineDocument fields. For example, you could use this to process navigational `DIV` elements one way, then process content-ful `DIV` elements another way.

See HTML parser stage for configuration details.

Note	The HTML Transformation index pipeline stage is deprecated in favor of this parser stage.
------	---

7.1.2. XML parser stage

The XML parser stage parses whole XML documents by default, but it can also be configured to parse only specific nodes without loading the entire document into memory. It can also split an XML document into multiple documents. XPATH-like expressions are used to select specific nodes to parse, such as `/posts/row` or `/posts/record`. Nested XML elements are flattened.

7.1.3. CSV parser stage

This parser breaks down incoming CSV files into the most efficient components for Fusion to index. It produces one new document per row from the CSV input, excluding comment rows and header rows.

See CSV parser stage for configuration details.

7.1.4. JSON parser stage

JSON parsing converts JSON content from a single document field into one or more new documents. This parser uses Solr's [JsonRecordReader](#) to split JSON into sub-documents.

See JSON parser stage for configuration details.

7.1.5. Text parser stage

The Plain Text parser can split a text file by lines or consume it into a single document.

Options for treatment of this filetype include:

- Plain Text Parser Fields
- Number of header rows to skip
- Split on line end or not
- Comment character
- Skip empty lines
- Charset

See Text parser stage for configuration details.

7.1.6. Archive parser stage

The Archive parser stage can parse the majority of common archive and compressed file formats. They are parsed into their constituent documents, which can then be parsed further or sent straight to the index pipeline. The following archive formats are supported:

- tar
- zip
- jar
- 7z
- ar
- arj
- Unix dump
- cpio

See Archive parser stage for configuration details.

7.1.7. Apache Tika parser stage

Apache Tika is a versatile parser that supports many types of unstructured document formats, such as HTML, PDF, Microsoft Office documents, OpenOffice, RTF, audio, video, images, and more. A complete list of supported formats is available at <http://tika.apache.org/>.

See Apache Tika parser stage for configuration details.

7.1.8. Fallback parser stage

The Fallback parser stage is useful for processing data that Fusion does not have a specified parsing process for. Fallback does not technically parse data, since it does not know what to do with it, it simply copies the raw bytes into a Solr document. If your Fusion parser stage configuration encounters data it does not know how to parse, such as someone's proprietary data file format, it will copy it as-is, whereas if it encounters recognizable data in more common file types, such as PDFs, Fusion will parse the text and metadata using Tika.

The Fallback parser acts as the final stage that attempts to parse any documents that haven't been parsed already. When the correct parsing stage lands on the data, it executes accordingly.

See Fallback parser stage for configuration details.

7.2. Configuring parsers

When you configure a datasource, you can use the Index Workbench or the Parsers API to create a parser. A parser consists of an ordered list of parser stages, some global parser parameters, and the stage-specific parameters. You can re-order the stages list by dragging them up or down in the Index Workbench.

Any parser stage can be added to the same parser multiple times if different configuration options are needed for different stages. Datasources with fixed-structure data will also be parsed by Fusion, but with default settings that do not need to be customized.

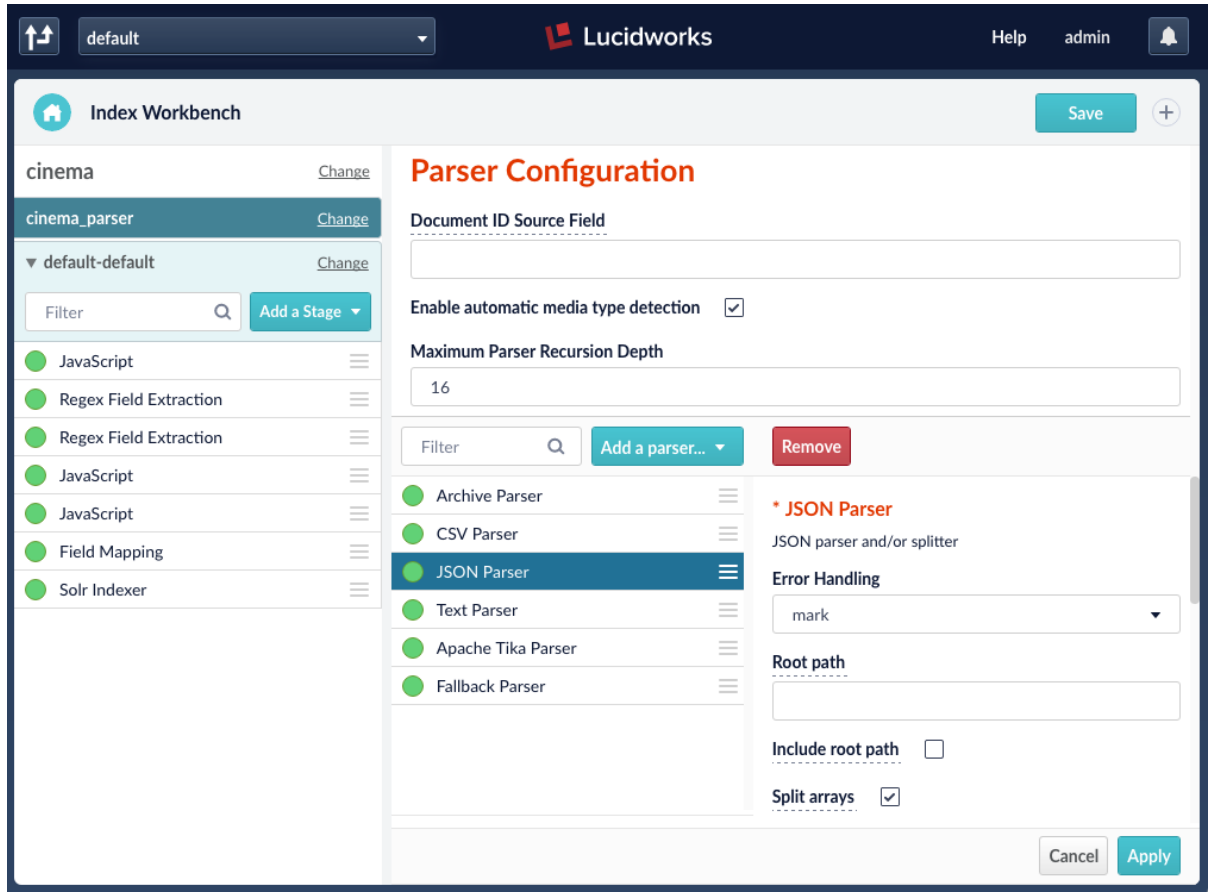
There is no limit to the number of stages that can be included in a parser. The order in which they run is also completely

flexible and can be linear or recursive. When the end of the parsing sequence is reached, a default parser stage automatically attempts to parse anything that has not yet been matched.

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

7.2.1. Parser configuration in the Fusion UI

In the Search context, select a collection, then navigate to **Home > Index Workbench** and click the parser, usually called "*datasource-name*_parser". Clicking a specific stage opens its configuration panel.



7.2.2. Parser configuration in the REST API

The Parsers API provides a programmatic interface for viewing, creating, and modifying parsers, as well as sending documents directly to a parser.

- To get all currently-defined parsers: <http://localhost:8765/api/v1/parsers/>
- To get the parser schema: http://localhost:8765/api/v1/parsers/_schema

Here's a very simple parser example, for parsing JSON input:

```
{ "id": "simple-json",
  "type": "json",
    arbitrary parser-specific options here.
  "prettify": false
}
```

The example below shows a parser that can parse JSON input, as well as JSON that is inside zip, tar, or gzip containers, or any combination (such as .tar.gz). The order of the stages begins with the outermost containers and ends with the innermost content.

```
{ "id": "default-json",
  "type": "composite",
  "parsers": [
    { "id" : "zip-parser",
      "type" : "zip" },
    { "type" : "gz" },
    { "type" : "tar" },
    { "id": "json-parser",
      "type": "json",
      "prettify": false
    }
  ]
}
```

ID is optional, just as in pipeline stages. Many parser stages require no configuration other than `type`.

7.3. Parser index pipeline stage

The parsers themselves only parse whole documents. Parsing of content embedded in fields is performed separately by the Parser Index Pipeline Stage. This stage identifies the field or context that requires parsing, the appropriate parser to use, and what to do with the parsed content.

7.4. Apache Tika Parser Stage

Apache Tika is a versatile parser that supports many types of unstructured document formats, such as HTML, PDF, Microsoft Office documents, OpenOffice, RTF, audio, video, images, and more. A complete list of supported formats is available at <http://tika.apache.org/>.

7.4.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Global configuration

These configuration options apply to the parser as a whole.

Property	Description
<code>idField</code>	A document field to use as the document ID.
<code>enableMediaTypeDetection</code>	Automatically detect the Content-Type of each document; disable this to use <code>application/octet-stream</code> .
<code>maxParserDepth</code>	Maximum number of times a stage may recurse over any document before proceeding to the next stage.

Apache Tika parser stage configuration

Property	Description
<code>errorHandling</code>	One of the following: <code>ignore</code> - Ignore errors, drop the current record, and continue parsing the next record or document. <code>log</code> - Log errors, drop the current record, and continue parsing the next record or document. <code>fail</code> - Generate an exception and stop parsing. <code>mark</code> (default) - Create a marker document that is emitted instead of the bad record. The error document contains common metadata gathered so far, plus error message and error class. The parser may also add more details about the error condition.
<code>includeImages</code>	Include images; default false.
<code>flattenCompound</code>	"True" to flatten compound documents; default false.
<code>addFailedDocs</code>	Add failed documents; default false.
<code>addOriginalContent</code>	Add original document content (raw bytes); default true.

Property	Description
<code>contentEncoding</code>	Content transport encoding (per RFC 1341), one of <code>binary</code> or <code>base64</code> ; default <code>binary</code> .
<code>returnXml</code>	"True" to return parsed content as XML (instead of HTML); default false.
<code>keepOriginalStructure</code>	"True" to return original XML and HTML instead of Tika XML output.
<code>extractHtmlLinks</code>	Collect links explicitly declared in document structure (for example, using HTML tags, bookmarks, and so on); default true.
<code>extractOtherLinks</code>	Use regex-based heuristic extractor to collect likely links from plain text content in all fields; default false.
<code>excludeContentTypes</code>	An array of content types to exclude from parsing.
<code>mediaTypes</code>	An array of types for this parser, which must match the pattern: <code>^\[\]\\$</code>
<code>pathPatterns</code>	Specify a file name or pattern that must be matched for this parser to run. Forward slashes (/) are used to join names of files inside archives with the archive name. <code>syntax</code> - One of "glob" or "regex". <code>pattern</code> - The filename or pattern to match. Glob examples: <code>z.txt</code> or <code>.md</code> or <code>/a//b/f.txt</code> Regex examples: <code>z.txt\$</code> or <code>.\.txt\$</code> or <code>/a/[\/]b/f.txt\$</code>
<code>inheritMediaTypes</code>	"True" to inherit acceptable types from the parser.

7.5. Archive Parser Stage

The Archive parser stage can parse the majority of common archive and compressed file formats. They are parsed into their constituent documents, which can then be parsed further or sent straight to the index pipeline. The following archive formats are supported:

- tar
- zip
- jar
- 7z
- ar
- arj
- Unix dump
- cpio

7.5.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Global configuration

These configuration options apply to the parser as a whole.

Property	Description
<code>idField</code>	A document field to use as the document ID.
<code>enableMediaTypeDetection</code>	Automatically detect the Content-Type of each document; disable this to use <code>application/octet-stream</code> .
<code>maxParserDepth</code>	Maximum number of times a stage may recurse over any document before proceeding to the next stage.

Archive parser stage configuration

Property	Description
<code>errorHandling</code>	<p>One of the following:</p> <p><code>ignore</code> - Ignore errors, drop the current record, and continue parsing the next record or document.</p> <p><code>log</code> - Log errors, drop the current record, and continue parsing the next record or document.</p> <p><code>fail</code> - Generate an exception and stop parsing.</p> <p><code>mark</code> (default) - Create a marker document that is emitted instead of the bad record. The error document contains common metadata gathered so far, plus error message and error class. The parser may also add more details about the error condition.</p>
<code>alwaysDetect</code>	When true, always force content type detection, otherwise only if not declared.
<code>mediaTypes</code>	An array of types for this parser, which must match the pattern: <code>^\[\]\\$</code>
<code>pathPatterns</code>	<p>Specify a file name or pattern that must be matched for this parser to run. Forward slashes (/) are used to join names of files inside archives with the archive name.</p> <p><code>syntax</code> - One of "glob" or "regex".</p> <p><code>pattern</code> - The filename or pattern to match.</p> <p>Glob examples: <code>z.txt</code> or <code>.md</code> or <code>/a//b/f.txt</code></p> <p>Regex examples: <code>z.txt\$</code> or <code>.\.txt\$</code> or <code>/a/[\/]b/f.txt\$</code></p>
<code>inheritMediaTypes</code>	"True" to inherit acceptable types from the parser.

7.6. CSV Parser Stage

This parser breaks down incoming CSV files into the most efficient components for Fusion to index. It produces one new document per row from the CSV input, excluding comment rows and header rows.

7.6.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Global configuration

These configuration options apply to the parser as a whole.

Property	Description
<code>idField</code>	A document field to use as the document ID.
<code>enableMediaTypeDetection</code>	Automatically detect the Content-Type of each document; disable this to use <code>application/octet-stream</code> .
<code>maxParserDepth</code>	Maximum number of times a stage may recurse over any document before proceeding to the next stage.

CSV parser stage configuration

Property	Description
<code>errorHandling</code>	One of the following: <code>ignore</code> - Ignore errors, drop the current record, and continue parsing the next record or document. <code>log</code> - Log errors, drop the current record, and continue parsing the next record or document. <code>fail</code> - Generate an exception and stop parsing. <code>mark</code> (default) - Create a marker document that is emitted instead of the bad record. The error document contains common metadata gathered so far, plus error message and error class. The parser may also add more details about the error condition.
<code>charset</code> <i>required</i>	The default is <code>detect</code> , to auto-detect the character set.

Property	Description
<code>ignoreBOM</code> <i>required</i>	Ignore Byte-Order Mark (BOM) if present and always use the configured character set. When set to false, a valid BOM character set overrides the configured default character set.
<code>delimiter</code>	Delimiter character between fields. The following characters are valid: <code>,</code> (comma) <code>\\t</code> (tab) <code>` `</code> (space) <code>`</code>
<code>`</code> (pipe) <code>^</code> (carat) Default is comma if auto-detection is disabled.	<code>quote</code>
Quote character; default is a double quote (") if auto-detection is disabled.	<code>quoteEscape</code>
Quote escape character, default is a double quote (") if auto-detection is disabled.	<code>autoDetect</code>
Attempt to guess the delimiter, quote, quote escape, and comment characters.	<code>trimWhitespace</code>
Trim off leading and trailing whitespace from columns; default "true".	<code>hasHeaders</code>
Treat the first row as column headers; default "false".	<code>headers</code>
List of column headers, overrides file headers if present.	<code>skipEmptyLines</code>
Skip any empty lines encountered; default "true".	<code>lineSeparator</code>
Line separator character.	<code>nullValue</code>
A string value to replace nulls with; no default.	<code>emptyValue</code>
A string value to replace empty strings with; no default.	<code>includeRowNumber</code>
Include the row number (line number) in the emitted documents; default "true".	<code>comment</code>
Character at start of row to indicate a comment, default is hash (#) if auto-detection is disabled.	<code>commentHandling</code>

Property	Description
<p>How to handle comments, one of the following:</p> <p><code>ignore</code> - Ignore all comments (default).</p> <p><code>as_field</code> - Add each comment as a field in the document.</p> <p><code>as_document</code> - Add each comment to a separate document.</p>	<code>maxLength</code>
Maximum number of characters to allow for a single read line; default 10485760 (10MB).	<code>maxNumColumns</code>
Maximum number of columns to allow for a single row; default 1000.	<code>maxColumnChars</code>
Maximum number of characters a single column value can have; default 10485760 (10MB).	<code>columnHandling</code>
<p>What to do when a row has too many or too few columns, do one of the following:</p> <p><code>error</code> - Throw an error.</p> <p><code>align</code> - Align the column.</p> <p><code>default</code> - Do nothing special.</p>	<code>fillValue</code>
A string value to use when aligning the columns (when <code>columnHandling</code> is "align").	<code>mediaTypes</code>
An array of types for this parser, which must match the pattern: <code>^\[\]\\$</code>	<code>pathPatterns</code>
<p>Specify a file name or pattern that must be matched for this parser to run. Forward slashes (/) are used to join names of files inside archives with the archive name.</p> <p><code>syntax</code> - One of "glob" or "regex".</p> <p><code>pattern</code> - The filename or pattern to match.</p> <p>Glob examples: <code>z.txt</code> or <code>.md</code> or <code>/a//b/f.txt</code></p> <p>Regex examples: <code>z.txt\$</code> or <code>.\.txt\$</code> or <code>/a/[\/]b/f.txt\$</code></p>	<code>inheritMediaTypes</code>

7.7. Fallback Parser Stage

The Fallback parser stage is useful for processing data that Fusion does not have a specified parsing process for. Fallback does not technically parse data, since it does not know what to do with it, it simply copies the raw bytes into a Solr document. If your Fusion parser stage configuration encounters data it does not know how to parse, such as someone's proprietary data file format, it will copy it as-is, whereas if it encounters recognizable data in more common file types, such as PDFs, Fusion will parse the text and metadata using Tika.

The Fallback parser acts as the final stage that attempts to parse any documents that haven't been parsed already. When the correct parsing stage lands on the data, it executes accordingly.

7.7.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Global configuration

These configuration options apply to the parser as a whole.

Property	Description
<code>idField</code>	A document field to use as the document ID.
<code>enableMediaTypeDetection</code>	Automatically detect the Content-Type of each document; disable this to use <code>application/octet-stream</code> .
<code>maxParserDepth</code>	Maximum number of times a stage may recurse over any document before proceeding to the next stage.

Fallback parser stage configuration

Property	Description
<code>errorHandling</code>	One of the following: <code>ignore</code> - Ignore errors, drop the current record, and continue parsing the next record or document. <code>log</code> - Log errors, drop the current record, and continue parsing the next record or document. <code>fail</code> - Generate an exception and stop parsing. <code>mark</code> (default) - Create a marker document that is emitted instead of the bad record. The error document contains common metadata gathered so far, plus error message and error class. The parser may also add more details about the error condition.

Property	Description
<code>includeImages</code>	Include images; default false.
<code>flattenCompound</code>	"True" to flatten compound documents; default false.
<code>addFailedDocs</code>	Add failed documents; default false.
<code>addOriginalContent</code>	Add original document content (raw bytes); default true.
<code>charset</code> <i>required</i>	The default is <code>detect</code> , to auto-detect the character set.
<code>returnXml</code>	"True" to return parsed content as XML (instead of HTML); default false.
<code>keepOriginalStructure</code>	"True" to return original XML and HTML instead of Tika XML output.
<code>extractHtmlLinks</code>	Collect links explicitly declared in document structure (for example, using HTML tags, bookmarks, and so on); default true.
<code>extractOtherLinks</code>	Use regex-based heuristic extractor to collect likely links from plain text content in all fields; default false.
<code>excludeContentTypes</code>	An array of content types to exclude from parsing.
<code>mediaTypes</code>	An array of types for this parser, which must match the pattern: <code>^\[\]\\$</code>
<code>pathPatterns</code>	Specify a file name or pattern that must be matched for this parser to run. Forward slashes (/) are used to join names of files inside archives with the archive name. <code>syntax</code> - One of "glob" or "regex". <code>pattern</code> - The filename or pattern to match. Glob examples: <code>z.txt</code> or <code>.md</code> or <code>/a//b/f.txt</code> Regex examples: <code>z.txt\$</code> or <code>.\.txt\$</code> or <code>/a/[\/]b/f.txt\$</code>
<code>inheritMediaTypes</code>	"True" to inherit acceptable types from the parser.

7.8. HTML Parser Stage

This parser stage processes the following HTML elements:

- `<title>`
- `<body>` (with tags removed)
- `<meta>`
- `<a>` and `<link>`

Additionally, you can configure [JSoup selectors](#) to extract specific HTML and CSS elements from a document and map them to PipelineDocument fields. For example, you could use this to process navigational `DIV` elements one way, then process content-ful `DIV` elements another way.

HTML and CSS elements can be selected for extraction into new documents or fields:

- To create new documents from selected elements, configure `recordSelector`.
- To create new fields from selected elements, configure `mappings`.

Title, body, metadata, and links are only populated in the parent document. Both of these parameters support [JSoup selectors](#), which provides a rich syntax for selecting HTML and CSS elements.

Note	The HTML Transformation index pipeline stage is deprecated in favor of this parser stage.
------	---

7.8.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Global configuration

These configuration options apply to the parser as a whole.

Property	Description
<code>idField</code>	A document field to use as the document ID.
<code>enableMediaTypeDetection</code>	Automatically detect the Content-Type of each document; disable this to use <code>application/octet-stream</code> .
<code>maxParserDepth</code>	Maximum number of times a stage may recurse over any document before proceeding to the next stage.

HTML parser stage configuration

Property	Description
<code>mediaTypes</code>	An array of types for this parser, which must match the pattern: <code>^\[\]\\$</code>

Property	Description
<code>pathPatterns</code>	<p>Specify a file name or pattern that must be matched for this parser to run. Forward slashes (/) are used to join names of files inside archives with the archive name.</p> <p><code>syntax</code> - One of "glob" or "regex".</p> <p><code>pattern</code> - The filename or pattern to match.</p> <p>Glob examples: <code>z.txt</code> or <code>.md</code> or <code>/a//b/f.txt</code></p> <p>Regex examples: <code>z.txt\$</code> or <code>.\.txt\$</code> or <code>/a[/\]/b/f.txt\$</code></p>
<code>inheritMediaTypes</code>	"True" to inherit acceptable types from the parser.
<code>errorHandling</code>	<p>One of the following:</p> <p><code>ignore</code> - Ignore errors, drop the current record, and continue parsing the next record or document.</p> <p><code>log</code> - Log errors, drop the current record, and continue parsing the next record or document.</p> <p><code>fail</code> - Generate an exception and stop parsing.</p> <p><code>mark</code> (default) - Create a marker document that is emitted instead of the bad record. The error document contains common metadata gathered so far, plus error message and error class. The parser may also add more details about the error condition.</p>
<code>charset</code> <code>required</code>	The default is <code>detect</code> , to auto-detect the character set.
<code>recordSelector</code>	Create child documents from each HTML file, using JSoup selectors . Only one selector may be configured at once, but a selector may be complex.
<code>mappings</code>	<p>Extract parts of the document into new fields, using JSoup selectors.</p> <p>If <code>mappings</code> is configured and <code>recordSelector</code> is not, then additional metadata (if configured; see below) is populated only in the parent document.</p>
<code>keepParent</code>	"True" to keep the parent document after child documents are created. This property no effect if <code>recordSelector</code> is not specified.

Property	Description
<code>extractHtmlLinks</code>	Collect links explicitly declared in document structure (for example, using HTML tags, bookmarks, and so on); default true.
<code>extractBodyText</code>	Extract the content of all elements in the <code><body></code> as a single text field.

7.9. JSON Parser Stage

JSON parsing converts JSON content from a single document field into one or more new documents. This parser uses Solr's [JsonRecordReader](#) to split JSON into sub-documents.

7.9.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Global configuration

These configuration options apply to the parser as a whole.

Property	Description
<code>idField</code>	A document field to use as the document ID.
<code>enableMediaTypeDetection</code>	Automatically detect the Content-Type of each document; disable this to use <code>application/octet-stream</code> .
<code>maxParserDepth</code>	Maximum number of times a stage may recurse over any document before proceeding to the next stage.

JSON parser stage configuration

Property	Description
<code>errorHandling</code>	One of the following: <code>ignore</code> - Ignore errors, drop the current record, and continue parsing the next record or document. <code>log</code> - Log errors, drop the current record, and continue parsing the next record or document. <code>fail</code> - Generate an exception and stop parsing. <code>mark</code> (default) - Create a marker document that is emitted instead of the bad record. The error document contains common metadata gathered so far, plus error message and error class. The parser may also add more details about the error condition.
<code>rootPath</code>	Use only children of this JSON pointer. The default is <code>"/</code> .
<code>splitArrays</code>	First split top-level arrays into multiple documents, and then apply other rules.

Property	Description
<code>mappings</code>	<p>Extract parts of the document into specified fields.</p> <p><code>path</code> - JSONPath expression.</p> <p><code>target</code> - The target field.</p>
<code>mediaTypes</code>	<p>An array of types for this parser, which must match the pattern: <code>^\[\]\\$</code></p>
<code>pathPatterns</code>	<p>Specify a file name or pattern that must be matched for this parser to run. Forward slashes (/) are used to join names of files inside archives with the archive name.</p> <p><code>syntax</code> - One of "glob" or "regex".</p> <p><code>pattern</code> - The filename or pattern to match.</p> <p>Glob examples: <code>z.txt</code> or <code>.md</code> or <code>/a//b/f.txt</code></p> <p>Regex examples: <code>z.txt\$</code> or <code>.\.txt\$</code> or <code>/a/[\/]/b/f.txt\$</code></p>
<code>inheritMediaTypes</code>	<p>"True" to inherit acceptable types from the parser.</p>

7.10. Text Parser Stage

The Plain Text parser can split a text file by lines or consume it into a single document.

Options for treatment of this filetype include:

- Plain Text Parser Fields
- Number of header rows to skip
- Split on line end or not
- Comment character
- Skip empty lines
- Charset

7.10.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Global configuration

These configuration options apply to the parser as a whole.

Property	Description
<code>idField</code>	A document field to use as the document ID.
<code>enableMediaTypeDetection</code>	Automatically detect the Content-Type of each document; disable this to use <code>application/octet-stream</code> .
<code>maxParserDepth</code>	Maximum number of times a stage may recurse over any document before proceeding to the next stage.

Text parser stage configuration

Property	Description
<code>errorHandling</code>	<p>One of the following:</p> <p><code>ignore</code> - Ignore errors, drop the current record, and continue parsing the next record or document.</p> <p><code>log</code> - Log errors, drop the current record, and continue parsing the next record or document.</p> <p><code>fail</code> - Generate an exception and stop parsing.</p> <p><code>mark</code> (default) - Create a marker document that is emitted instead of the bad record. The error document contains common metadata gathered so far, plus error message and error class. The parser may also add more details about the error condition.</p>
<code>charset</code> <i>required</i>	The default is <code>detect</code> , to auto-detect the character set.
<code>ignoreBOM</code> <i>required</i>	Ignore Byte-Order Mark (BOM) if present and always use the configured character set. When set to false, a valid BOM character set overrides the configured default character set.
<code>splitLines</code>	Split text into lines to create multiple records; default false.
<code>skipHeaderLines</code>	Skip a number of header lines; default 0.
<code>trimWhitespace</code>	Trim off leading and trailing whitespace from lines; default false.
<code>skipEmptyLines</code>	Skip any empty lines encountered; default false.
<code>outputField</code>	Name of the output field where text is stored; default <code>body</code> .
<code>maxLength</code>	Maximum number of characters to allow for the body, <code>-1</code> for unlimited; default 1MB.
<code>maxLineLength</code>	Maximum number of characters to allow for any single line, default 1MB.
<code>commentField</code>	Name of the output field where comment is stored; default <code>comment</code> .
<code>comment</code>	Characters at start of line to indicate a comment; default <code>#</code> (hash).

Property	Description
<code>commentHandling</code>	How to handle comments, one of the following: <code>ignore</code> - Ignore comments and remove them from the text. <code>include</code> - Include comments as-is (default). <code>as_field</code> - Add comments as a field.
<code>mediaTypes</code>	An array of types for this parser, which must match the pattern: <code>^\[\]\\$</code>
<code>pathPatterns</code>	Specify a file name or pattern that must be matched for this parser to run. Forward slashes (/) are used to join names of files inside archives with the archive name. <code>syntax</code> - One of "glob" or "regex". <code>pattern</code> - The filename or pattern to match. Glob examples: <code>z.txt</code> or <code>.md</code> or <code>/a//b/f.txt</code> Regex examples: <code>z.txt\$</code> or <code>.\.txt\$</code> or <code>/a/[\/]b/f.txt\$</code>
<code>inheritMediaTypes</code>	"True" to inherit acceptable types from the parser.

7.11. XML Parser Stage

The XML parser stage parses whole XML documents by default, but it can also be configured to parse only specific nodes without loading the entire document into memory. It can also split an XML document into multiple documents. XPATH-like expressions are used to select specific nodes to parse, such as `/posts/row` or `/posts/record`. Nested XML elements are flattened.

To create new documents from selected elements, configure `rootPaths`.

Note	The XML Transformation index pipeline stage is deprecated in favor of this parser stage.
------	--

7.11.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

Global configuration

These configuration options apply to the parser as a whole.

Property	Description
<code>idField</code>	A document field to use as the document ID.
<code>enableMediaTypeDetection</code>	Automatically detect the Content-Type of each document; disable this to use <code>application/octet-stream</code> .
<code>maxParserDepth</code>	Maximum number of times a stage may recurse over any document before proceeding to the next stage.

XML parser stage configuration

Property	Description
<code>mediaTypes</code>	An array of types for this parser, which must match the pattern: <code>^\[\]\\$</code>
<code>pathPatterns</code>	Specify a file name or pattern that must be matched for this parser to run. Forward slashes (/) are used to join names of files inside archives with the archive name. <code>syntax</code> - One of "glob" or "regex". <code>pattern</code> - The filename or pattern to match. Glob examples: <code>z.txt</code> or <code>.md</code> or <code>/a//b/f.txt</code> Regex examples: <code>z.txt\$</code> or <code>.\.txt\$</code> or <code>/a[\/]/b/f.txt\$</code>
<code>inheritMediaTypes</code>	"True" to inherit acceptable types from the parser.
<code>errorHandling</code>	One of the following: <code>ignore</code> - Ignore errors, drop the current record, and continue parsing the next record or document. <code>log</code> - Log errors, drop the current record, and continue parsing the next record or document. <code>fail</code> - Generate an exception and stop parsing. <code>mark</code> (default) - Create a marker document that is emitted instead of the bad record. The error document contains common metadata gathered so far, plus error message and error class. The parser may also add more details about the error condition.
<code>rootPaths</code>	Read XML elements that can be found on specified XML paths and parse them into separate documents.
<code>maxSize</code>	The maximum size of a document.

Property	Description
<code>listHandling</code>	One of the following: <ul style="list-style-type: none">• <code>multivalued</code> + Create a single multivalued field containing all items. * <code>index_numbered</code> + Create a separate index-numbered field per list item.

Chapter 8. Pipeline Stages Reference

Index Pipeline stages are used to create and modify PipelineDocument objects to control how incoming data is indexed in Fusion's Solr core.

Query Pipeline stages are used to modify Request objects and Response objects to control how query results are returned to the end user.

8.1. Pipeline Stage Properties

Fusion associates a pipeline stage definition with a unique ID and stores the definition so that stages can be reused across pipelines and applications. In addition to an ID, all stages have the following properties:

- **type** (required): an enumeration, one of the defined Fusion pipeline stage types, e.g., "index-logging". If the Fusion UI is used to define the stage, this property is filled in automatically.
- **label** (optional) : a string field with a maximum length of 255 characters. The label is displayed on the Fusion UI.
- **skip** (optional): a boolean value, If true, pipeline processing bypasses this stage altogether. The default is **false**.
- **condition** (optional): a JavaScript expression that evaluates to true (1) or false (0). If this condition evaluates to false, this stage is skipped. The default is **true**.

8.2. Conditional Processing

Each Fusion pipeline stage contains a "Condition" field at the beginning of the stage configuration found in the Fusion UI. A conditional script can be used to dynamically turn a stage on/off. For example, in the index pipeline you can create a field mapping stage and create a condition to only execute the stage if a document contains the field 'sample_field', and to skip the stage otherwise.

```
doc.hasField("sample_field");
```

For the example above, if the pipeline document has the field 'sample_field', then the conditional script returns 'true' and the pipeline stage is executed. If the pipeline document does not have the 'sample_field', the condition returns 'false' and the pipeline stage is skipped.

The JavaScript expression specified in the condition property of a pipeline stage has access to the pipeline objects. For example, query processing can be conditional on information in the query Request object:

```
request.hasParam("fusion-user-name") && request.getFirstParam("fusion-user-name").equals("SuperUser");
```

The above condition first checks that the property "fusion-user-name" is specified, then checks for a particular value.

8.3. Index Pipeline Stages

Index Pipeline stages are used to create and modify PipelineDocument objects. Use the Index Workbench to configure stages in a pipeline and preview the results.

See these reference topics for details about each index pipeline stage:

8.3.1. Document transformation

- Apache Tika Parser
- CSV Parsing
- HTML Transformation
- JSON Parsing
- XML Transformation

8.3.2. Document filtering and enrichment

- Detect Language
- Exclude Documents
- Format Signals
- Include Documents
- JDBC Lookup
- REST Query

8.3.3. Field transformation

- Date Parsing
- Field Mapping
- Filter Short Fields
- Find and Replace
- Regex Field Extraction
- Regex Field Filter
- Regex Field Replacement
- Resolve Multivalued Fields
- Solr Dynamic Field Name Mapping

8.3.4. Natural language processing

- Detect Sentences
- Gazetteer Lookup Extraction
- OpenNLP NER Extraction
- Tag Part-of-Speech

8.3.5. Indexing

- Solr Indexer
- Solr Partial Update Indexer

8.3.6. Troubleshooting

- Logging
- Send PagerDuty Message
- Send SMTP Email
- Send Slack Message
- Write Log Message

8.3.7. Advanced

- Call Pipeline
- Exclusion Filter
- Javascript
- Machine Learning
- Set Property
- Update Experiment

8.4. Apache Tika Parser Index Stage

The Apache Tika Parser index stage type includes rules for parsing documents with [Apache Tika](#). Fusion uses Tika v1.13. (Note that components of the Solr distribution included with Fusion contain their own Tika jar files; these are not used by Fusion.)

To parse a CSV document, you should use a CSV Parsing Index Stage instead of an Apache Tika Parser stage.

8.4.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `tika-parser`.

Configuration Properties

Property	Description, Type
addFailedDocs Add Failed Documents	type: <code>boolean</code> default value: 'false'
addOriginalContent Add Original Document Content (Raw Bytes)	type: <code>boolean</code> default value: 'true'
contentEncoding Content Transport Encoding of the Content (per RFC1341)	type: <code>string</code> default value: 'binary' enum: <code>{ binary base64 }</code>
contentField Field Name Where Content is Expected	type: <code>string</code> default value: ' <code>raw_content</code> '
excludeContentTypes Content Types to Exclude	List of content types to exclude from parsing type: <code>array of string</code>
extractHtmlLinks Extract XHTML Links	Collect links explicitly declared in document structure (e.g. using HTML tags, bookmarks, etc) type: <code>boolean</code> default value: 'true'

Property	Description, Type
extractOtherLinks Extract Other Links	Use regex-based heuristic extractor to collect likely links from plain text content in all fields. type: <code>boolean</code> default value: 'false'
flattenCompound Flatten Compound Documents	type: <code>boolean</code> default value: 'false'
includeContentTypes Content Types to Include	List of content types to parse type: <code>array of string</code>
includeImages Include Images	type: <code>boolean</code> default value: 'false'
keepOriginalStructure Return Original XML and HTML Instead of Tika XML Output	type: <code>boolean</code> default value: 'false'
returnXml Return Parsed Content as XML or HTML	type: <code>boolean</code> default value: 'false'
zipBombCompressionRatio Maximum input-to-output byte ratio	Maximum number of output bytes fusion will generate per input byte. If you are indexing highly compressed files, you may increase this value to avoid triggering 'Zip Bomb' detection type: <code>integer</code> default value: '200'
zipBombMaxDepth Maximum nesting depth	Returns the maximum XML element nesting level. If you are indexing highly nested files, you may increase this value to avoid triggering 'Zip Bomb' detection type: <code>integer</code> default value: '200'

Property	Description, Type
zipBombMaxPackageEntryDepth Maximum package entry depth	Sets the maximum package entry nesting level. If you are indexing highly nested files, you may increase this value to avoid triggering 'Zip Bomb' detection type: <code>integer</code> default value: '20'

8.5. Call Pipeline Index Stage

The Call Pipeline index stage (called the Fusion Pipeline stage in versions earlier than 3.0) provides a means of composing index pipelines.

8.5.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `apollo`.

Configuration Properties

Property	Description, Type
collection Collection <i>required</i>	type: <code>string</code>
pipeline Pipeline ID <i>required</i>	type: <code>string</code>

8.6. Date Parsing Index Stage

The Date Parsing Stage (previously called the Date Parser stage) is an index pipeline stage that performs parsing and normalization of date/time data in document fields which uses the Fusion DateUtils library. The resulting date/time information is available both as a timestamp in UTC time zone as well as a local date/time in the original local time zone.

The time zone name, offset and the epoch time are stored in separate fields, too. Additionally the formatted dates can be split into their components, and each component added to separate document fields.

Note that this stage works only with data that consists solely of the date/time information, i.e. it will not work correctly if dates are a part of a larger piece of text.

8.6.1. Timestamp splitting options

Splitting options help in processing timestamp information without resorting to scripting - e.g. in order to index day of week information it's more convenient and faster to split the timestamp in this stage, and then just discard other components that are not needed (using a field mapping stage), rather than using a JavaScript stage to parse and split the timestamp manually.

Please note that time zone name and time zone offset, as well as epoch time, are always added as separate fields regardless of the splitting options. E.g. for a field named `test` these values will be added as fields `tz.test`, `tz_offset.test`, and `epoch.test`.

The option `splitLocal` splits the timestamp in its original timezone, while the option `splitUTC` first converts the timestamp to UTC and then splits it. The resulting date and time components are stored in fields that follow patterns `<part>.local.<fieldName>` and `<part>.utc.<fieldName>` respectively.

The following parts are extracted and added to the document:

- **year** - year component
- **month** - month in year, from 1 to 12
- **day** - day in month, from 1 to 31
- **yday** - day in year, from 1 to 356
- **weekday** - day of week, 1 being Monday and 7 being Sunday
- **week** - week in year, from 1 to 52. Note: in the standard ISO8601 week algorithm, the first week of the year is that in which at least 4 days are in the year. As a result of this definition, day 1 of the first week may be in the previous year, which will be indicated by `weekyear`. The opposite is also true - last day of the last week may be in the next year, and `weekyear` will show the next year.
- **weekyear** - year corresponding to the week value. This can be either the current year or previous one, or the next one.
- **hour** - hour in day, from 0 to 23
- **min** - minute in hour, from 0 to 59
- **sec** - second in minute, from 0 to 59
- **ms** - millisecond in second, from 0 to 999

Example: given this normalized timestamp in the original timezone `2015-01-01 00:00:00.000 Europe/Warsaw` in a field

`test`, the corresponding normalized UTC timestamp will be `2014-12-31T23:00:00.00Z`.

Example: `splitLocal` parsing

The following table shows the additional fields added to a document as the result of applying `splitLocal` parsing to the contents a field named `test` which contains the value `2015-01-01 00:00:00.000 Europe/Warsaw`:

Field name	value
<code>tz.test</code>	Europe/Warsaw
<code>tz_offset.test</code>	+01:00
<code>epoch.test</code>	1420066800000

Example: `splitUTC` parsing

The following table shows the additional fields added to a document as the result of applying `splitUTC` parsing to the contents a field named `test` which contains the value `2015-01-01 00:00:00.000 Europe/Warsaw`:

Field name	value
<code>tz.test</code>	Europe/Warsaw
<code>tz_offset.test</code>	+01:00
<code>epoch.test</code>	1420066800000
<code>year.utc.test</code>	2014
<code>year.local.test</code>	2015
<code>month.utc.test</code>	12
<code>month.local.test</code>	1
<code>day.utc.test</code>	31
<code>day.local.test</code>	1
<code>yday.utc.test</code>	365
<code>yday.local.test</code>	1
<code>weekday.utc.test</code>	3
<code>weekday.local.test</code>	4
<code>week.utc.test</code>	1
<code>week.local.test</code>	1
<code>weekyear.utc.test</code>	2015
<code>weekyear.local.test</code>	2015
<code>hour.utc.test</code>	23
<code>hour.local.test</code>	0
<code>min.utc.test</code>	0

Field name	value
min.local.test	0
sec.utc.test	0
sec.local.test	0
ms.utc.test	0
ms.local.test	0

Note: The following:

- weekday is different - UTC day of week is Wednesday, and local day of week is already Thursday.
- yday in UTC points to the last day of the year, while it's the first day of the year in local time zone, similarly with day.
- week and weekyear are the same in both cases - because according to the ISO 8601 definition all days of this week belong to year 2015 so it doesn't matter whether it's Wednesday or Thursday.

8.6.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `date-parsing`.

Configuration Properties

Property	Description, Type
dateFormats Date Formats	Custom date formats, or empty for default formats type: <code>array of string</code>
defaultLocale Default Locale	Locale to assume if different from <code>Locale.ENGLISH</code> . This uses IETF BCP 47 codes. type: <code>string</code>
defaultTimezone Default Timezone	Timezone to assume if one is not present in the incoming date type: <code>string</code>

Property	Description, Type
ignoreInvalid Ignore Invalid	When false invalid date strings will cause the whole document to be rejected. When true invalid values are silently discarded type: boolean default value: 'false'
requireTimezone Require Timezone	Accept only formats that explicitly specify the timezone type: boolean default value: 'false'
sourceFields Source Fields	type: array of string
splitLocal Split Local Date Into Parts	Split local date (in the local timezone) into parts and store in .local. fields type: boolean default value: 'false'
splitUTC Split UTC Date Into Parts	Split UTC date (in the UTC timezone) into parts and store in .utc. fields type: boolean default value: 'false'

8.7. Detect Sentences Index Stage

The Detect Sentences index stage (called the Sentence Detection stage in versions earlier than 3.0) operates over one of more fields in the Pipeline Document and annotates each field with sentence boundary information. These annotations can be used by downstream indexing stages. A Detect Sentences stage can be used in tandem with a Tag Part-of-Speech Index Stage to provide part-of-speech annotations for the individual tokens in the field.

This stage uses [Apache OpenNLP](#) project's [Sentence Detection](#) tool. The OpenNLP documentation states:

The OpenNLP Sentence Detector can detect that a punctuation character marks the end of a sentence or not. In this sense a sentence is defined as the longest white space trimmed character sequence between two punctuation marks. The first and last sentence make an exception to this rule. The first non whitespace character is assumed to be the begin of a sentence, and the last non whitespace character is assumed to be a sentence end. _

Fusion comes with a set of OpenNLP language models for english. These data files are found in the directory: [fusion/3.1.x/data/nlp/models](#). The included sentence model is `en-sent.bin`. More models are available from the OpenNLP [models SourceForge repository](#).

Model files must be uploaded to Fusion using the Fusion Blob Store service via the REST API (see examples below).

8.7.1. Sentence Detection in a NLP Pipeline

The following video shows how to use a Sentence Detection index stage as part of an NLP pipeline:

8.7.2. Stage Setup

This is an example of how to upload a sentence model file to the Fusion blob:

INPUT

```
curl -u user:pass -X PUT --data-binary @en-pos-maxent.bin -H 'Content-type: text/plain'  
http://localhost:8764/api/apollo/blobs/en-pos-maxent.bin
```

OUTPUT

```
{  
  "name" : "en-sent.bin",  
  "contentType" : "text/plain",  
  "size" : 5696197,  
  "modifiedTime" : "2015-07-15T06:57:48.636Z",  
  "version" : 0,  
  "md5" : "db2cd70395b9e2e4c6b9957015a10607"  
}
```

This is an example setup of this stage using the previously loaded .bin file:

INPUT

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id":"DetectSentences1", "type": "detect-sentences","sentenceModel":"en-sent.bin","source": ["A test sentence"]}'  
http://localhost:8764/api/apollo/index-stages/instances
```

OUTPUT

```
{  
  "type" : "detect-sentences",  
  "id" : "DetectSentences1",  
  "sentenceModel" : "en-sent.bin",  
  "source" : [ "A test sentence" ],  
  "skip" : false,  
  "label" : "detect-sentences",  
  "type" : "detect-sentences"  
}
```

8.7.3. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `part-of-speech`.

Configuration Properties

Property	Description, Type
posModel Part of Speech Model <i>required</i>	type: <code>string</code> reference: model:open-nlp
source Source Fields <i>required</i>	type: <code>array of string</code> minimum number of items (<code>minItems</code>): 1
tokenizerModel Tokenizer Model <i>required</i>	type: <code>string</code> reference: model:open-nlp

8.8. Exclude Documents Index Stage

The Exclude Documents stage drops all documents that match all of the specified rules (Boolean AND). If some field has multiple values then at least one value must match against specified pattern. No further processing is done on any matching documents, thus they will not be indexed into a Fusion collection. All non-matching documents are passed to the next stage in the pipeline. Rules are defined using regular expression field matching.

8.8.1. Examples

Give the "simple-exclude" pipeline a stage that excludes certain document types:

```
curl -u user:pass -X POST -H "Content-type: application/json" 'http://localhost:8764/api/apollo/index-pipelines' -d '{
  "id" : "simple-exclude",
  "stages" : [ {
    "type" : "exclude-doc",
    "matchRules" : [ {
      "field" : "document_type",
      "pattern" : "(xls|xlsx|xlst|doc|docx)"
    }
  ]
}]
}'
```

Send a text document through the "simple-exclude" pipeline:

```
curl -u user:pass 'http://localhost:8764/api/apollo/index-pipelines/simple-exclude/collections/logs/index?simulate=true&echo=true' -H 'Content-type: application/json' -d '{
  "document_type": "txt"
}'
```

The response is document metadata, indicating the document passed the stage:

```
[ {
  "id" : "93da43ff-4218-4f24-a690-23b530926104",
  "fields" : [ {
    "name" : "document_type",
    "value" : "txt",
    "metadata" : { },
    "annotations" : [ ]
  } ]
} ]
```

Send an XLS document through the "simple-exclude" pipeline:

```
curl -u user:pass 'http://localhost:8764/api/apollo/index-pipelines/simple-exclude/collections/logs/index?simulate=true&echo=true' -H 'Content-type: application/json' -d '{
  "document_type": "xls"
}'
```

The empty response indicates the document was dropped:

```
[ ]
```

8.8.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `exclude-doc`.

Configuration Properties

Property	Description, Type
matchRules Fields and Patterns <i>required</i>	type: array of object minimum number of items (minItems): 1 object attributes: \{ field (required) : \{ display name: Field type: string description : The name of the field to match } pattern (required) : \{ display name: Regex Pattern type: string description : Pattern to match the field value against. The value may be a regex pattern. format : regex } }

8.9. Exclusion Filter Index Stage

The Exclusion Filter index stage is used to remove fields or documents that match items in a pre-defined exclusion list.

There are two ways to supply an exclusion list:

- Upload a file containing a newline-separated list, using the Blob Store. When configuring the index stage, reference the list by its blob name in the `location` property (**Exclusion List URI** in the Fusion UI).
- When configuring the index stage, enter an array of values for exclusion in the `excludeValues` property (**Exclusion List** in the Fusion UI).

The Exclusion Filter stage can be configured using one or both of these methods; Fusion combines them into one list. If `regexPattern` is configured, the pattern is applied to the field before the result is compared to the combined list.

By default, any matching *field* is excluded from indexing. To exclude the whole document, set `skipDocument` to "true" (**Skip Document** in the Fusion UI).

8.9.1. Uploading an exclusion list

Before you can configure the `location` property, you must upload one or more exclusion lists to Fusion using the Blob Store API.

Fusion comes with an example exclusion list at `fusion/3.1.x/data/nlp/excludes/excludes.txt`. Here is an example of how to upload this file using `curl`, where `admin:pass` are the credentials for an admin-level user:

```
curl -u admin:pass -X PUT --data-binary @data/nlp/excludes/excludes.txt -H 'Content-type: text/plain'
http://localhost:8764/api/apollo/blobs/excludes.txt
```

8.9.2. Example

Use an exclusion list for entities found in the author field:

```
{
  "type" : "exclusion-filter",
  "id" : "iw",
  "filters" : [ {
    "sourceField" : "author_s",
    "location" : "excludes.txt",
    "caseSensitive" : false
  } ],
  "skip" : false
} ]
```

8.9.3. Configuration

Tip

When entering configuration values in the UI, use *unescaped* characters, such as `\t` for the tab character. When entering configuration values in the API, use *escaped* characters, such as `\\t` for the tab character.

When using Fusion's REST-API, the ID of this stage is: **exclusion-filter**.

Configuration Properties

Property	Description, Type
filters	type: array of object
Filtering Rules	object attributes: <pre>\{ caseSensitive (required): \{ display name: Case Sensitive type: boolean default value: 'false' } excludeValues: \{ display name: Values to Exclude type: array of string } location: \{ display name: Exclusion List URI (Blob name) type: string reference : other } regexPattern: \{ display name: Regex Expression type: object group : \{"default" ⇒ 0, "title" ⇒ "Regex Capture Group", "type" ⇒ "integer"} pattern (required): \{"format" ⇒ "regex", "title" ⇒ "Regex expression", "type" ⇒ "string"} } skipDocument: \{ display name: Skip Document type: boolean default value: 'false' } sourceField (required): \{ display name: Source Field type: string } }</pre>

8.10. Field Mapping Index Stage

The Field Mapping stage was renamed for Fusion 3.0; it was called the Field Mapper stage in previous versions.

A Field Mapping stage is used to do customized mapping of the fields in an Index Pipeline document to fields in a the Solr schema.

8.10.1. Field Mapping Stage Properties

A Field Mapping stage specification consists of three things:

- a unique ID
- a set of mapping rules that specify operations applied to named fields as a triple: { `source`, `target`, `operation` }.
- a set of rules called "unmapped" rules which specify operations applied to fields whose name doesn't match any of the mapping rules, also a triple { `source`, `target`, `operation` }.

Mapping Rules and Unmapped Rules

Each rule has the following properties:

Property	Description
source	The name of the source field. This will be the name of the field in the Pipeline document that should be mapped to another field. Java regular expressions can be used in the source field by surrounding the regular expression with forward slashes ('/'). For example, <code>/(.)text(.\\)/</code> is a valid expression that will find field names in the incoming document that contain the string 'text' between any number of preceding or succeeding characters. If a regular expression is not used, the value supplied for the source will be treated as a literal field name and will be matched ignoring the case (for example, "text" will match "tExt" or "Text", etc.).
target	The name of the target field. If the value for the <code>source</code> was a regular expression, then this can also be a regular expression. It can also contain substitutions using references to capture groups (using Java's Matcher.replaceAll()). Otherwise, the source field name will be simply substituted by the value of target according to the operation rules described below.

Property	Description
operation	<p>What to do with the field during mapping. Several options are available:</p> <ul style="list-style-type: none"> • copy: Content contained in fields matching <code>source</code> will be copied to <code>target</code>. • move: Content contained in fields matching <code>source</code> will be moved to <code>target</code> (it may also help to think of this as the field name being replaced by the value of <code>target</code>). • delete: Content contained in fields matching <code>source</code> will be dropped from the document and not indexed. In this case, the <code>target</code> can be null or not defined at all. • add: The literal value of <code>target</code> will be added to the <code>source</code> if <code>source</code> is a regular expression. If <code>source</code> is not a regular expression, <code>target</code> will be added as a new field. • set: The literal value of <code>target</code> will be set as the new value of <code>source</code> if <code>source</code> is a regular expression. If <code>source</code> is not a regular expression, <code>target</code> will be set as a new field. • keep: Content contained in fields matching <code>source</code> will be retained and unchanged, and the fields will be added to a list of known fields and they will not be affected by however the <code>renameUnknown</code> rule has been set.

Note that the mapping rules are applied in the order in which they are defined, which may have an impact on the final effects of the mapping process.

8.10.2. Field Mapping Behavior

The field mapping rules are applied in a specific order.

1. A copy of the Pipeline document is prepared. All further operations are applied to this copy.
2. The rules are traversed only once, in the order of their declaration in the `mapping` property. This means it is possible to do multiple operations on a field. However, note that if fields are moved (renamed), further operations should reference the new field name.
3. Before each rule is evaluated, the current list of field names is prepared and sorted in alphabetic ascending order.
4. The current rule is applied to field values for each matching name from the list of names prepared in step 3. New field names resulting from the current rule do not effect the snapshot list of field names; in order for a rule to be applied to a new field name, it will be included in a later round of the evaluation cycle.
5. The process is repeated for each rule, and a list of matching source fields is noted.
6. If the document contains any fields that were not affected by any mapping rule, the `renameUnknown` option is applied

if it has been set to true.

7. Finally, the resulting transformed document is returned to the next stage of the index pipeline.

8.10.3. Examples

Map several fields:

```
{
  "id": "mapping-text",
  "type": "field-mapping",
  "mappings": [
    {
      "operation": "move",
      "source": "plaintextcontent",
      "target": "body"
    },
    {
      "operation": "add",
      "source": "content-length",
      "target": "fileSize"
    },
    {
      "operation": "move",
      "source": "/file(.*)/",
      "target": "lastModified"
    },
    {
      "operation": "delete",
      "source": "last-printed"
    },
    {
      "operation": "copy",
      "source": "mimetype",
      "target": "content_type"
    }
  ],
  "unmapped": {
    "source": "/(.*)/",
    "target": "$1_ss",
    "operation": "move"
  },
  "skip" : false
}
```

Set the `urlX` field based on the value of the `employee_id` field:

```
{
  "id": "set-field",
  "type": "field-mapping",
  "mappings": [
    {
      "operation": "set",
      "source": "urlX",
      "target": "https://mydomain.com/<employee_id>"
    }
  ],
  "skip" : false
}
```

8.10.4. Configuration

Tip

When entering configuration values in the UI, use *unescaped* characters, such as `\t` for the tab character. When entering configuration values in the API, use *escaped* characters, such as `\\t` for the tab character.

When using Fusion's REST-API, the ID of this stage is: `field-mapping`.

Configuration Properties

Property	Description, Type
<p>mappings</p> <p>Field Mappings</p>	<p>List of mapping rules</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre> } + 'source` _ (required)_ : \{ + display name: Source Field + type: `string` + description : The name of the field to be mapped. + } + 'target` : \{ + display name: Target Field + type: `string` + description : The name of the field to be mapped to. + } + } </pre>
<p>reservedFieldsMappingAllowed</p> <p>Allow System Fields Mapping?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

Property	Description, Type
<p>unmapped</p> <p>Unmapped Fields</p>	<p>If fields do not match any of the field mapping rules, these rules will apply.</p> <p>type: object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> operation : \{ <ul style="list-style-type: none"> display name: Operation type: string default value: 'copy' description : The type of mapping to perform: move, copy, delete, add, set, or keep. enum: \{ copy move delete set add keep } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

8.11. Field Parser Index Stage

Parse the contents of a field using a pre-defined parser.

8.11.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `field-parser`.

Configuration Properties

Property	Description, Type
includeMetadata Include Metadata	Include parser metadata as document fields. type: <code>boolean</code> default value: 'true'
inheritFields Inherit Fields	Should generated documents include the original document's fields. type: <code>boolean</code> default value: 'true'
parserId Parser ID <i>required</i>	The ID of the parser to use, as defined by the Parsers API type: <code>string</code> default value: 'default' reference: parser
sourceContentType Source Content Type	The Content-Type of the data stored in the source field. type: <code>string</code>
sourceEncoding Source Field Encoding	The content transport encoding of the data stored in the source field (per RFC1341). type: <code>string</code> default value: 'binary' enum: <code>{ binary base64 }</code>

Property	Description, Type
sourceField Source Field	The document field where the content should be read from. type: <code>string</code> default value: <code>'raw_content'</code>

8.12. Filter Short Fields Index Stage

The Filter Short Fields Index Stage was renamed for Fusion 3.0; it was called the Short Field Filter stage in previous versions.

A Filter Short Fields index stage removes short field values from a pipeline document according to a set of filters where each filter specifies a field name and a minimum length. All field values less than the specified length will be removed from the document.

For the REST API, this stage type is "filter-short-fields".

8.12.1. Example Stage Specification

Remove entities in the `author_s` field that are less than 3 characters:

```
{
  "type" : "filter-short-fields",
  "id" : "short-stage",
  "filters" : [ {
    "sourceField" : "author_s",
    "length" : 3
  } ],
  "skip" : false
}
```

8.12.2. Configuration

Tip

When entering configuration values in the UI, use *unescaped* characters, such as `\t` for the tab character. When entering configuration values in the API, use *escaped* characters, such as `\\t` for the tab character.

When using Fusion's REST-API, the ID of this stage is: `short-filter`.

Configuration Properties

Property	Description, Type
<p>filters</p> <p>Filters</p>	<p>type: array of object</p> <p>object attributes: \{</p> <p> length (<i>required</i>): \{</p> <p> display name: Minimum Length</p> <p> type: integer</p> <p> default value: '3'</p> <p> description : Minimum length of permissible entities.</p> <p> Fields with lesser values will be removed.</p> <p> exclusiveMinimum : false</p> <p> minimum : 1</p> <p> }</p> <p> sourceField (<i>required</i>): \{</p> <p> display name: Source Field</p> <p> type: string</p> <p> }</p> <p>}</p>

8.13. Find and Replace Index Stage

The Find and Replace stage is used to standardize or remove a set of specific passages, phrases, or words from a document field using exact string matching. This stage is configured with one or more rules. Each rule specifies one or more texts to match against and a *single* text string that is the replacement value. When the replacement string is the empty string, this will delete the matched text from the document field.

There are two different ways of setting up a Find and Replace stage:

- Using Fusion's blob store to store the list of texts to be standardized. This is known as "Find List Replace" For each standardization, there is a file which contains the set of texts to be standardized, one per line, with file suffix `.lst`. The `.lst` file is uploaded into Fusion's `../REST_API_Reference/Blob-Store-API.html[Blob Store]`. The required properties for each rule are: the document field to scan; the name of the list file uploaded to the blob store, inclusive of suffix `.lst`; and the replace text value.
- Specifying a set of find/replace pairs directly using the Fusion UI or REST API. This is known as "Find Replace". The requires properties for each rule are: the document field to scan; a list of find/replace pairs.

8.13.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `find-replace`.

Configuration Properties

Property	Description, Type
<p>findListReplaceRules</p> <p>Find List Replace rules</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> caseSensitive : \{ <ul style="list-style-type: none"> display name: Case Sensitive type: boolean default value: 'true' listLocation (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Blob Name of the List type: string reference : other replacementValue : \{ <ul style="list-style-type: none"> display name: Replacement Value type: string sourceField (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Source Field type: string <p>\}</p>
<p>findReplaceRules</p> <p>Find replace rules</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> caseSensitive : \{ <ul style="list-style-type: none"> display name: Case sensitive type: boolean default value: 'true' keyValues (<i>required</i>): \{ <ul style="list-style-type: none"> display name: find and replace strings list type: array of object sourceField (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Source Field type: string <p>\}</p>

8.14. Format Signals Index Stage

The Format Signals stage (called the Signal Formatter stage in versions earlier than 3.0) normalizes both the fields and field contents of a PipelineDocument to ensure that certain pre-defined fields for signals are populated.

8.14.1. Date/time parsing and formatting

Timestamp data can be obtained from the following fields, in this order of precedence:

- timestamp
- timestamp_tdt
- timestamp_dt
- epoch - value in this field is treated as a number of milliseconds since epoch, and UTC zone is assumed.

It's now possible to specify the locale to be used for parsing timestamps by setting the "timestampLocale" property in the stage configuration. If this property is null then the default platform locale will be used.

Output document will carry the following two fields:

- "timestamp" - containing the ISO8601 timestamp
- "tz_timestamp_txt" - containing the "zoned format" of the timestamp with normalized components.

8.14.2. Example Stage Specification

The Format Signals stage defined as part of the default 'signals_ingest' pipeline included with Fusion.

Note that this stage does not define a list of allowed types.

```
{
  "type" : "format-signals",
  "id" : "ingest-signals",
  "flatten" : true,
  "undefinedType" : "general",
  "skip" : false,
  "label" : "format-signals",
  "type" : "format-signals"
}
```

8.14.3. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `signal-formatter`.

Configuration Properties

Property	Description, Type
allowedTypes Allowed Types	List of allowed signal types. type: array of string
flatten Flatten	Flatten nested values. type: boolean default value: 'true'
timestampLocale Timestamp Locale	Use this locale when parsing timestamp information. Null uses platform default locale. type: string
undefinedType Undefined Type	Signal type when undefined. Null discards events with undefined type. type: string

8.15. Gazetteer Lookup Extraction Index Stage

The Gazetteer Lookup Extraction index stage (called the Gazetteer Lookup Extractor stage in versions earlier than 3.0) uses predefined lists of words and phrases to process specified text fields in a document. A gazetteer is a set of lookup lists over names of people, places, or things. These lookup lists are used to find occurrences of these names in text. The matched items are saved into separate fields on the document for downstream processing.

8.15.1. Gazetteers and OpenNLP Tools

The following video shows how to configure a Gazetteer Lookup Extraction stage in combination with OpenNLP:

8.15.2. Uploading Lookup Lists to Fusion Blob Store

Fusion includes a number of lookup lists in the directory `fusion/3.1.x/data/nlp/gazetteer`. To use the supplied lists or a list of your own data, each must list be uploaded to Fusion using the Blob Store API in order to make the list contents available to the Gazetteer Lookup Extraction stage.

For example, to identify color names, you would first compile a list of color terms, one entry per line in a text file with suffix `.lst` and then upload that file using the Fusion REST API endpoint `api/apollo/blobs/<listfilename>`, as per the following example which uses the `\curl` command-line utility, where 'admin' is the name of a user with admin privileges, and 'pass' is that user's password:

```
curl -u admin:pass -X PUT --data-binary @data/nlp/gazetteer/colours.lst -H 'Content-type: text/plain'
http://localhost:8764/api/apollo/blobs/colours.lst
```

8.15.3. Name Lookup Example

Define a lookup-extractor to identify mentions of certain celebrities in text field `description_t`:

```
{
  "type" : "lookup-extractor",
  "id" : "peopleLookup",
  "rules" : [ {
    "source" : [ "description_t" ],
    "target" : "celebrities_ss",
    "entityTypes" : [ {
      "name" : "person_female",
      "definitions" : [ "person_female.lst" ]
    } ],
    "additionalEntities" : [ {
      "name" : "players",
      "definitions" : [ "sharapova", "murray" ]
    }, {
      "name" : "actors",
      "definitions" : [ "pitt", "jolie" ]
    } ],
    "caseSensitive" : false
  } ],
  "skip" : false
}
```

8.15.4. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion’s REST-API, the ID of this stage is: `lookup-extractor`.

Configuration Properties

Property	Description, Type
<p>rules</p> <p>Extraction Rules</p>	<p>type: <code>array of object</code></p> <p>object attributes: <code>{</code></p> <ul style="list-style-type: none"> <code>additionalEntities</code> : <code>{</code> <li style="padding-left: 20px;">display name: Additional Entities <li style="padding-left: 20px;">type: <code>array of object</code> <li style="padding-left: 20px;"><code>}</code> <code>caseSensitive</code> : <code>{</code> <li style="padding-left: 20px;">display name: Case Sensitive <li style="padding-left: 20px;">type: <code>boolean</code> <li style="padding-left: 20px;">default value: ‘false’ <li style="padding-left: 20px;"><code>}</code> <code>entityTypes (required)</code> : <code>{</code> <li style="padding-left: 20px;">display name: Entity Types <li style="padding-left: 20px;">type: <code>array of object</code> <li style="padding-left: 20px;"><code>}</code> <code>source (required)</code> : <code>{</code> <li style="padding-left: 20px;">display name: Source Fields <li style="padding-left: 20px;">type: <code>array of string</code> <li style="padding-left: 20px;">minItems : 1 <li style="padding-left: 20px;"><code>}</code> <code>target (required)</code> : <code>{</code> <li style="padding-left: 20px;">display name: Target Field <li style="padding-left: 20px;">type: <code>string</code> <li style="padding-left: 20px;"><code>}</code> <code>writeMode</code> : <code>{</code> <li style="padding-left: 20px;">display name: Write Mode <li style="padding-left: 20px;">type: <code>string</code> <li style="padding-left: 20px;">default value: ‘append’ <li style="padding-left: 20px;">description : What to do if document has target field already <li style="padding-left: 20px;">enum: <code>{ overwrite append }</code> <div style="border: 1px solid #ccc; border-radius: 5px; padding: 5px; margin-top: 10px;"> <pre> } + } </pre> </div>

8.16. Include Documents Index Stage

This stage passes documents to the next stage in the pipeline if they match one or more of the specified rules (Boolean OR). If some field has multiple values then at least one value must match against specified pattern. All non-matching documents are dropped. Rules are defined using regular expression field matching.

8.16.1. Examples

Give the "simple-include" pipeline a stage that includes only certain document types:

```
curl -u user:pass -X POST -H "Content-type: application/json" 'http://localhost:8764/api/apollo/index-pipelines' -d '{
  "id": "simple-include",
  "stages": [ {
    "type": "include-doc",
    "matchRules": [ {
      "field": "document_type",
      "pattern": "(xls|xlsx|xlst|doc|docx)"
    } ]
  } ]
}'
```

Response:

```
{
  "id": "simple-include",
  "stages": [ {
    "type": "include-doc",
    "id": "f701f96b-780e-4355-9dd3-6e53a89afe3e",
    "matchRules": [ {
      "field": "document_type",
      "pattern": "(xls|xlsx|xlst|doc|docx)"
    } ],
    "type": "include-doc",
    "skip": false,
    "label": "include-doc"
  } ],
  "properties": { }
}
```

Send a text document through the "simple-include" pipeline:

```
curl -u user:pass 'http://localhost:8764/api/apollo/index-pipelines/simple-include/collections/logs/index?simulate=true&echo=true' -H 'Content-type: application/json' -d '{
  "document_type": "txt"
}'
```

The empty response indicates the document was dropped:


```
[ ]
```

Send an XLS document through the pipeline:

```
curl -u user:pass 'http://localhost:8764/api/apollo/index-pipelines/simple-include/collections/logs/index?simulate=true&echo=true' -H 'Content-type: application/json' -d '{
  "document_type": "xls"
}'
```

The response is document metadata, indicating the document passed the stage:

```
{
  "id" : "9e7d1c2e-343a-49de-bc6a-1d1fc25fa93f",
  "fields" : [ {
    "name" : "document_type",
    "value" : "xls",
    "metadata" : { },
    "annotations" : [ ]
  } ]
}
```

8.16.2. Configuration

Tip

When entering configuration values in the UI, use *unescaped* characters, such as `\t` for the tab character. When entering configuration values in the API, use *escaped* characters, such as `\\t` for the tab character.

When using Fusion's REST-API, the ID of this stage is: `include-doc`.

Configuration Properties

Property	Description, Type
<p>matchRules</p> <p>Fields and Patterns</p> <p><i>required</i></p>	<p>type: array of object</p> <p>minimum number of items (minItems): 1</p> <p>object attributes: \{</p> <p> field (<i>required</i>): \{</p> <p> display name: Field</p> <p> type: string</p> <p> description : The name of the field to match</p> <p> }</p> <p> pattern (<i>required</i>): \{</p> <p> display name: Regex Pattern</p> <p> type: string</p> <p> description : Pattern to match the field value against.</p> <p> The value may be a regex pattern.</p> <p> format : regex</p> <p> }</p> <p>}</p>

8.17. JavaScript Index Stage

For a complete description of the JavaScript Index stage and additional examples, see: [Custom JavaScript Stages For Index Pipelines](#).

8.17.1. Examples

Drop a document by ID

```
function(doc) {
  var id = doc.getId();
  if (id !== null) {
    var pattern = "https://www.mydomain.com/links/contact/?";

    // 0 means the pattern was found so drop the doc
    return (id.indexOf(pattern) == 0) ? null : doc;
  }

  return doc;
}
```

Format Date to Solr Date

```
// For example:
// From: 26/Mar/2015:14:38:48 -0700
// To: 2015-03-26T14:38:48Z (Solr format)
function(doc) {
  if (doc.getId() !== null) {
    var inboundPattern = "dd/MMM/yyyy':HH:mm:ss Z"; // modify this to match the format of the inbound date
    var solrDatePattern = "yyyy-MM-dd'T'HH:mm:ss'Z"; // leave this alone
    var dateFieldName = "apachelogtime"; // change this to your date field name

    var solrFormatter = new java.text.SimpleDateFormat(solrDatePattern);
    var apacheParser = new java.text.SimpleDateFormat(inboundPattern);

    var dateString = doc.getFirstFieldValue(dateFieldName);
    logger.info("**** dateString: " + dateString);
    var inboundDate = apacheParser.parse(dateString);
    logger.info("**** inboundDate: " + inboundDate.toString());
    var solrDate = solrFormatter.format(inboundDate);
    logger.info("**** solrDate: " + solrDate.toString());

    doc.setField(dateFieldName, solrDate.toString());
  }

  return doc;
}
```

Replace whitespace and newlines

```

function(doc) {
  if (doc.getId() !== null) {
    var fields = ["col1", "col2", "col3"];

    for (i = 0; i < fields.length; i++ ) {
      var field = fields[i];
      var value = doc.getFirstFieldValue(field);
      logger.info("BEFORE: Field " + field + ": *" + value + "*");

      if (value !== null) {
        value = value.replace(/^\s+/, ""); // remove leading whitespace
        logger.info("AFTER: Field " + field + ": *" + value + "*");
        value = value.replace(/\s+$/, ""); // remove trailing whitespace
        logger.info("AFTER: Field " + field + ": *" + value + "*");
        value = value.replace(/\s+/g, " "); // multiple whitespace to one space
        logger.info("AFTER: Field " + field + ": *" + value + "*");

        doc.setField(field, value);
      }
    }
  }

  return doc;
}

```

Split the values in a field

```

//Split On a delimiter. In this case, a newline
function(doc){
  if (doc.getId() !== null) {
    var fromField = "company2_ss";
    var toField = "company2_ss";
    var delimiter = "\n";

    var oldList = doc.getFieldValues(fromField);

    var values = [];

    // parse the entries one at a time
    doc.removeFields(toField); // clear out the target field
    for (i = 0; i < oldList.size(); i++) {
      values[i] = oldList.get(i);

      // get the list of strings split by the delimiter
      newList = values[i].split(delimiter);
      for(j = 0; j < newList.length; j++){
        doc.addField(toField, newList[j]);
      }
    }
  }
  return doc;
}

```

8.17.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `javascript-index`.

Configuration Properties

Property	Description, Type
script Script Body <i>required</i>	type: <code>string</code>

8.18. JDBC Index Stage

The JDBC Index Stage is used to connect to a database, lookup one or more values, and then inject them into the context. The properties for setting stage are identical to the JDBC Query Stage, except for the "rows" property, which defaults to **-1** for the index stage (which returns all rows in the database), and defaults to **10** returned rows for the query stage.

Note	You must first upload the JDBC driver to Fusion, see the Connector JDBC API page.
------	---

8.18.1. Example

An example of a JDBC Index Stage setup

Upload stage config via POST to Fusion REST API endpoint ``api/apollo/index-stages/instances``

```
curl -u user:password -X POST -H "Content-Type: application/json" -d '{"id": "jdbc-index-test", "type": "jdbc-index-lookup", "driver": "postgresql-9.3-1101-jdbc4.jar", "connection": "jdbc:postgresql:database", "username": "user", "password": "password1", "preparedStatement": "select ID as id from DATABASE;"}' http://localhost:8764/api/apollo/index-stages/instances
```

Response

```
{
  "type" : "jdbc-index-lookup",
  "id" : "jdbc-index-test",
  "driver" : "postgresql-9.3-1101-jdbc4.jar",
  "connection" : "jdbc:postgresql:database",
  "username" : "user",
  "password" : "password1",
  "preparedStatement" : "select ID as id from DATABASE;",
  "fetchSize" : -1,
  "join" : true,
  "rows" : -1,
  "skip" : false,
  "label" : "jdbc-index-lookup",
  "type" : "jdbc-index-lookup"
}
```

8.18.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `jdbc-index-lookup`.

Configuration Properties

Property	Description, Type
connection Connection URI <i>required</i>	The connection string for the database. type: <i>string</i>
driver JDBC Driver <i>required</i>	The fully qualified class name of the JDBC Driver to use. type: <i>string</i>
fetchSize Fetch Size	The JDBC Fetch Size to use. If -1, use the driver default. type: <i>integer</i> default value: '-1'
join Join With Document	If true, the results will be added on to the document using the prefix key and the row id, else the results will be put in the pipeline context using type: <i>boolean</i> default value: 'true'
password Password <i>required</i>	The password to connect to the database. type: <i>string</i>
prefix Result Prefix Key	The string to use as a prefix for all values. type: <i>string</i>
preparedStatement SQL Prepared Statement <i>required</i>	The SQL Prepared Statement to execute when bound with values. type: <i>string</i>

Property	Description, Type
<p>preparedStatementKeys</p> <p>Prepared Statement Keys</p>	<p>The keys in the Pipeline Context to use to map request attributes into the prepared statement. These must map to the '?'s in your prepared statement. They must also be able to be resolved as the first parameter of that name in a request.</p> <p>type: array of string</p>
<p>rows</p> <p>Rows</p>	<p>The number of rows to return. -1 for all rows (be wary of memory usage for this).</p> <p>type: integer</p> <p>default value: '-1'</p>
<p>username</p> <p>Username</p> <p><i>required</i></p>	<p>The username to connect to the database.</p> <p>type: string</p>

8.19. Logging Index Stage

The Logging Index stage prints messages to the Connectors log file, default location `fusion/3.1.x/var/log/connectors/connectors.log`.

The verbosity of this message is controlled by the property `detailed`. When detailed logging is true the current PipelineDocument object is pretty-printed to the Connectors log file.

In a production environment logging stages should be configured with property `skip` set to `true`, if possible. Use of detailed logging may impact performance.

8.19.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `index-logging`.

Configuration Properties

Property	Description, Type
detailed	type: <code>boolean</code>
Detailed Logging	default value: 'false'

8.20. Machine Learning Index Stage

The Fusion machine learning indexing stage uses a compiled machine learning model to analyze a field or fields of a PipelineDocument and stores the results of analysis in a new field of either the PipelineDocument or PipelineContext object. This stage was introduced in Fusion version 2.4.

You must use [Spark's MLlib API](#) to create a supervised machine learning model and upload this model into Fusion's blob store collection. Complete details are available in section: Machine Learning Models in Fusion.

Successful use of this stage requires a proper understanding of both the model and your data. The machine learning model is described by its spark-mllib.json file, which contains the model specification as a JSON object. This object contains attribute "featureFields" which takes as its value a list of one or more field names. The contents of these fields are processed into the vector of features which the model operates on. If these fields aren't present in the document being analyzed, then the result is either an empty prediction or a configurable default value. If the contents of these fields differ greatly from the data used to compile the model, the predictions made by the model will be unreliable.

8.20.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `machine-learning`.

Configuration Properties

Property	Description, Type
defaultValue Default Value	Value to provide if a prediction cannot be made for a document. type: <code>string</code>
docFeatureFieldName Document Feature Field	Name of the field to extract document features from (model input). type: <code>string</code> default value: <code>'body_t'</code> minLength: 1
failOnError Fail on Error	Flag to indicate if this stage should throw an exception if an error occurs while generating a prediction for a document. type: <code>boolean</code> default value: <code>'false'</code>

Property	Description, Type
<p>includeRawPredictions</p> <p>Include Raw Predictions and Scores</p>	<p>Flag to indicate that raw predictions and scores, in addition to the best prediction, should be set on the document.</p> <p>type: boolean</p> <p>default value: 'false'</p>
<p>modelId</p> <p>Machine Learning Model ID</p> <p><i>required</i></p>	<p>The ID of the ML model stored in the Fusion blob store.</p> <p>type: string</p>
<p>predictionFieldName</p> <p>Prediction Field Name</p> <p><i>required</i></p>	<p>Name of the field to store the prediction (model output) in the document.</p> <p>type: string</p>
<p>storeInContext</p> <p>Store the Prediction in the Context</p>	<p>Flag to indicate that the prediction should be set as a context property instead of setting a field on the document.</p> <p>type: boolean</p> <p>default value: 'false'</p>

8.21. OpenNLP NER Extraction Index Stage

Named Entity Recognition (NER) is the task of finding the names of persons, organizations, locations, and/or things in a passage of free text. The OpenNLP NER Extraction index stage (previously called the OpenNLP NER Extractor stage) uses a set of rules to find named entities in a field in the Pipeline Document (the "source") and populates a new fields (the "target") with these entities.

This stage uses [Apache OpenNLP project's Named Entity Recognition tool](#) (the Name Finder tool). The OpenNLP documentation states

The Name Finder tool can detect named entities and numbers in text. To be able to detect entities the Name Finder needs a model. The model is dependent on the language and entity type it was trained for. The OpenNLP projects offers a number of pre-trained name finder models which are trained on various freely available corpora. They can be downloaded at our model download page. To find names in raw text the text must be segmented into tokens and sentences.

See this video tutorial for a demonstration of how to configure this stage:

Models are available from the OpenNLP [models SourceForge repository](#).

The Fusion directory `fusion/3.1.x/data/nlp` contains a set of NER models for English, as well as sentence, token, and part-of-speech models.

Before they can be used, model files must be uploaded to Fusion using the Fusion Blob Store service via the REST API. Here is an example of how to upload the sentence model file from the `fusion/3.1.x` using the curl command-line utility, where `admin` is the name of a user with admin privileges, and `pass` is the password:

```
curl -u admin:pass -X PUT --data-binary @data/nlp/models/en-sent.bin -H 'Content-type: application/octet-stream' http://localhost:8764/api/apollo/blobs/en-sent.bin
```

8.21.1. Example Specification

Specification of a stage which extracts names of people and places from field named 'body':

```

{
  "type" : "nlp-extraction",
  "id" : "nd",
  "rules" : [ {
    "source" : [ "body_s" ],
    "target" : "content",
    "sentenceModel" : "en-sent.bin",
    "tokenizerModel" : "en-token-1.bin",
    "entityTypes" : [ {
      "name" : "organization",
      "definition" : "en-ner-organization-1.bin"
    }, {
      "name" : "person",
      "definition" : "en-ner-person-1.bin"
    } ]
  } ],
  "skip" : false
} ]
}

```

8.21.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `nlp-extractor`.

Configuration Properties

Property	Description, Type
<p>rules</p> <p>Extractor Rules</p> <p><i>required</i></p>	<p>type: array of object minimum number of items (minItems): 1</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> entityTypes : \{ <ul style="list-style-type: none"> display name: Entity Types type: array of object sentenceModelLocation (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Sentence Model type: string reference : model:open-nlp source (<i>required</i>) : \{ <ul style="list-style-type: none"> display name: Source Fields type: array of string minItems : 1 target (<i>required</i>) : \{ <ul style="list-style-type: none"> display name: Target Field type: string tokenizerModelLocation (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Tokenizer Model type: string reference : model:open-nlp writeMode : \{ <ul style="list-style-type: none"> display name: Write Mode type: string default value: 'append' description : What to do if document has target field already enum: \{ overwrite append } <p style="text-align: right;">\} + \}</p>

8.22. Regex Field Extraction Index Stage

The Regex Field Extraction stage (called the Regular Expression Extractor stage in versions earlier than 3.0) is used to extract entities from documents based on matching regular expressions. The resulting regex matches over the contents of the source field are copied to the target field. The regular expression, source, and target fields are defined properties of this stage.

If using the REST API, this stage type is named "regex-extractor".

8.22.1. Example Stage Specification

Define a *regex-field-extraction* stage to apply a regular expression that looks for storage capabilities of products when it appears in the product 'name' field, and store it in a special field:

```
{
  "type" : "regex-field-extraction",
  "id" : "storagesize-regex-extraction",
  "rules" : [ {
    "source" : [ "name" ],
    "target" : "storage_size_ss",
    "pattern" : "(\\d{1,20}\\s{0,3})(GB|MB|TB|KB|mb|gb|tb|kb)",
    "annotateAs" : "storage_size"
  } ],
  "skip" : false
}
```

8.22.2. Configuration

Tip

When entering configuration values in the UI, use *unescaped* characters, such as `\t` for the tab character. When entering configuration values in the API, use *escaped* characters, such as `\\t` for the tab character.

When using Fusion's REST-API, the ID of this stage is: `regex-extractor`.

Configuration Properties

Property	Description, Type
<p>rules</p> <p>Regex Rules</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> annotateAs : \{ <ul style="list-style-type: none"> display name: Annotation Name type: string description : the name of the annotation to be added to the source field to mark the boundaries of the match group : \{ <ul style="list-style-type: none"> display name: Regex Capture Group type: integer default value: '0' exclusiveMinimum : false minimum : 0 noMatchValue : \{ <ul style="list-style-type: none"> display name: No Match Literal Value type: string pattern (<i>required</i>) : \{ <ul style="list-style-type: none"> display name: Regex Pattern type: string format : regex returnIfNoMatch : \{ <ul style="list-style-type: none"> display name: Return if no Match type: string default value: 'null' enum: \{ null input_string value } <p> } + 'source' _ (required)_ : \{ + display name: Source Fields + type: 'array' of 'string' + minItems : 1 + } + 'target' _ (required)_ : \{ + display name: Target Field + type: 'string' + } + 'writeMode' : \{ + display name: Write Mode + type: 'string' + default value: ''append'' + description : What to do if document has target field already + enum: \{ overwrite append } </p>

8.23. Regex Field Filter Index Stage

The Regex Field Filter Index Stage (called the Regular Expression Filter stage in versions earlier than 3.0) removes a field or fields from a PipelineDocument according to a set of filters where each filter specifies a field name and a regular expression. If a field value matches the regular expression, the field is deleted from the document. The regex patterns follow [Java regular expression pattern rules](#).

8.23.1. Example Stage Specification

Create a regex-filter to find Social Security Numbers and drop them from documents:

```
{
  "type" : "regex-field-replacement",
  "id" : "ssnFilter",
  "skip" : false,
  "filters" : [ {
    "sourceField" : "notes_t",
    "pattern" : "^\\d{3}-\\d{2}-\\d{4}$"
  } ]
}
```

8.23.2. Configuration

Tip

When entering configuration values in the UI, use *unescaped* characters, such as `\t` for the tab character. When entering configuration values in the API, use *escaped* characters, such as `\\t` for the tab character.

When using Fusion's REST-API, the ID of this stage is: `regex-filter`.

Configuration Properties

Property	Description, Type
<p>filters</p> <p>Filters</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> exclusive : \{ <ul style="list-style-type: none"> display name: Filter is Exclusive? type: boolean default value: 'true' description : Exclude only matching values when true, otherwise include only matching values. pattern (<i>required</i>) : \{ <ul style="list-style-type: none"> display name: Filter Pattern type: string format : regex sourceField (<i>required</i>) : \{ <ul style="list-style-type: none"> display name: Source Field type: string <p>\}</p>
<p>reservedFieldsFilteringAllowed</p> <p>Allow System Fields Filtering?</p>	<p>type: boolean</p> <p>default value: 'false'</p>

8.24. Regex Field Replacement Index Stage

The Regex Field Replacement index stage, introduced in Fusion 3.0, lets you match against a source field and append or overwrite that field or another one with a replacement string. It uses the [java.util.regex](#) library.

One notable feature is support for flags in the prefix of the pattern, such as `(?iu)\w+` where the `i` flag means case-insensitive and `u` means Unicode-aware case folding. You can also configure the desired behavior when there is no match. Additionally, you can also use capture groups in your replacement configuration using `$1` to denote the first capture group.

8.24.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `regex-replacement`.

Configuration Properties

Property	Description, Type
<p>rules</p> <p>Regex Rules</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> noMatchValue : \{ <ul style="list-style-type: none"> display name: No Match Literal Value type: string pattern (required) : \{ <ul style="list-style-type: none"> display name: Regex Pattern type: string format : regex replaceWhich : \{ <ul style="list-style-type: none"> display name: Replace Which type: string default value: 'all' description : Replace first or all matches enum: \{ all first } <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'replacement' _ (required)_ : \{ + display name: Regex Replacement + type: 'string' + } + 'returnIfNoMatch' : \{ + display name: Return if no Match + type: 'string' + default value: 'input_string' + enum: \{ null input_string value } </pre> <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + 'source' _ (required)_ : \{ + display name: Source Fields + type: 'array' of 'string' + minItems : 1 + } + 'target' : \{ + display name: Target Field + type: 'string' + } + 'writeMode' : \{ + display name: Write Mode + type: 'string' + default value: 'overwrite' + description : What to do if document has target field already + enum: \{ overwrite append } </pre> <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> } + } </pre>

8.25. Detect Sentences Index Stage

The Detect Sentences index stage (called the Sentence Detection stage in versions earlier than 3.0) operates over one of more fields in the Pipeline Document and annotates each field with sentence boundary information. These annotations can be used by downstream indexing stages. A Detect Sentences stage can be used in tandem with a Tag Part-of-Speech Index Stage to provide part-of-speech annotations for the individual tokens in the field.

This stage uses [Apache OpenNLP](#) project's [Sentence Detection](#) tool. The OpenNLP documentation states:

The OpenNLP Sentence Detector can detect that a punctuation character marks the end of a sentence or not. In this sense a sentence is defined as the longest white space trimmed character sequence between two punctuation marks. The first and last sentence make an exception to this rule. The first non whitespace character is assumed to be the begin of a sentence, and the last non whitespace character is assumed to be a sentence end. _

Fusion comes with a set of OpenNLP language models for english. These data files are found in the directory: [fusion/3.1.x/data/nlp/models](#). The included sentence model is `en-sent.bin`. More models are available from the OpenNLP [models SourceForge repository](#).

Model files must be uploaded to Fusion using the Fusion Blob Store service via the REST API (see examples below).

8.25.1. Sentence Detection in a NLP Pipeline

The following video shows how to use a Sentence Detection index stage as part of an NLP pipeline:

8.25.2. Stage Setup

This is an example of how to upload a sentence model file to the Fusion blob:

INPUT

```
curl -u user:pass -X PUT --data-binary @en-pos-maxent.bin -H 'Content-type: text/plain'  
http://localhost:8764/api/apollo/blobs/en-pos-maxent.bin
```

OUTPUT

```
{  
  "name" : "en-sent.bin",  
  "contentType" : "text/plain",  
  "size" : 5696197,  
  "modifiedTime" : "2015-07-15T06:57:48.636Z",  
  "version" : 0,  
  "md5" : "db2cd70395b9e2e4c6b9957015a10607"  
}
```

This is an example setup of this stage using the previously loaded .bin file:

INPUT

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id":"DetectSentences1", "type": "detect-sentences","sentenceModel":"en-sent.bin","source": ["A test sentence"]}'  
http://localhost:8764/api/apollo/index-stages/instances
```

OUTPUT

```
{  
  "type" : "detect-sentences",  
  "id" : "DetectSentences1",  
  "sentenceModel" : "en-sent.bin",  
  "source" : [ "A test sentence" ],  
  "skip" : false,  
  "label" : "detect-sentences",  
  "type" : "detect-sentences"  
}
```

8.25.3. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `part-of-speech`.

Configuration Properties

Property	Description, Type
posModel Part of Speech Model <i>required</i>	type: <code>string</code> reference: model:open-nlp
source Source Fields <i>required</i>	type: <code>array of string</code> minimum number of items (<code>minItems</code>): 1
tokenizerModel Tokenizer Model <i>required</i>	type: <code>string</code> reference: model:open-nlp

8.26. Resolve Multivalued Fields Index Stage

The Resolve Multivalued Fields stage (previously called the Multi-Value Resolver stage) allows you to choose one value from a set of one or more field values using a set of pre-defined rules, based on either field name or field type. For each field name or field type rule, a strategy is defined. There are 6 available strategies:

- **DEFAULT**: leave all values intact. This may cause problems if a field or dynamic field rule is not defined as multi-valued in the schema.
- **PICK_FIRST**: choose the first value and discard all others.
- **PICK_LAST**: choose the last value and discard all others.
- **PICK_BY_CREATOR**: based on the name of the "creator" metadata, choose the name of 'creator' as defined in the creatorStrategy property for the field name. Only the last matching value will be retained.
- **PICK_MAX**: choose the maximum value and discard all others. If the values are non-numeric, then all non-numeric values will be discarded.
- **PICK_MIN**: choose the minimum value and discard all others. If the values are non-numeric, then all non-numeric values will be discarded.

8.26.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `multivalue-resolver`.

Configuration Properties

Property	Description, Type
creatorStrategy Creator Strategies	Mapping of fields to a 'creator' name type: <code>array of object</code> object attributes: <code>{</code> <code>creatorName (required) : {</code> display name: Creator Name type: <code>string</code> <code>}</code> <code>fieldName (required) : {</code> display name: Field Name type: <code>string</code> <code>}</code> <code>}</code>

Property	Description, Type
<p>fieldStrategy</p> <p>Field Strategies</p>	<p>Mapping of fields to a strategy</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> fieldName (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Field Name type: string resolverStrategy (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Resolver Strategy type: string default value: 'pick_first' enum: \{ pick_first pick_last pick_by_creator concatenate_unique pick_min pick_max default } <p style="text-align: right;">} +</p> <p style="text-align: right;">}</p>
<p>typeStrategy</p> <p>Type Strategies</p>	<p>Mapping of types to a strategy. These supersede field strategies.</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> fieldName (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Field Name type: string resolverStrategy (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Resolver Strategy type: string default value: 'pick_first' enum: \{ pick_first pick_last pick_by_creator concatenate_unique pick_min pick_max default } <p style="text-align: right;">} +</p> <p style="text-align: right;">}</p>

8.27. REST Query Index Stage

The REST Query index stage (called the Query REST (RPC) Client stage in versions earlier than 3.0) performs Remote Procedure Calls to external services. It merges the response from that service with a document being processed by the pipeline. Calls to the external system are made for each document.

This stage can call external systems in any of the following ways:

- any HTTP or HTTPS request: `<http-protocol>://<path>`
- a Solr request: `solr://<collection>/...`
- a Fusion service request: `service://<serviceName>/<path>`

The call syntax allows the use of document field variables. For example, variable `${docField}` will be replaced with the value of the field named `docField` for the document being processed.

Several types of responses are also supported:

- POJOs
- JSON
- XML

8.27.1. Call Parameters

In the Fusion UI, select **include** to display the fields for the call parameters (`params` in the REST API). Parameters are specified as key-value pairs:

Endpoint URI / <code>uri</code>	A fully-qualified service URI. This can be an HTTP call, a Solr request, or a Fusion service call <ul style="list-style-type: none">• any HTTP or HTTPS request: <code><http-protocol>://<path></code>• a Solr request: <code>solr://<collection>/...</code>• a Fusion service request: <code>service://<serviceName>/<path></code>
Call Method / <code>method</code>	The method to use for the RPC call. Supported methods are <code>GET</code> , <code>POST</code> , <code>PUT</code> , and <code>DELETE</code> .
Query parameters / <code>queryParams</code>	Query parameters to be passed with the request.
Request protocol headers / <code>headers</code>	Request protocol headers, such as <code>"Content-Type" : "application/json"</code> .
Request entity (as string) / <code>entity</code>	The request body. This parameter can only be configured via the REST API; see the example below.

The `params` can also take variable substitution expressions. Variables are expressed as `${fieldName}`, where `fieldName` is a name of the current document's field (or `id` for the document's id). Only the first value of a multi-valued field is used for substitution, and this value is treated as a string.

For example, a `queryParams` could be constructed as follows:

```
"params" : {
  "uri" : "solr://collection1/select",
  "method" : "GET",
  "queryParams" : {
    "q" : "{!field f=a_txt v=${prodName}}",
    "prodName" : "${prodName}"
  }
}
```

In this example, the variable `${prodName}` will be replaced with the string value of the 'prodName' field in the current document being processed. If the current document contains "iPhone 6" as the value of 'prodName', the resulting query in this example will be a fielded search `q=a_txt:iPhone 6`.

8.27.2. Mapping Rules

The `mappingRules` property takes the following, specified as key-value pairs:

- `path` - an [XPath expression](#). It's assumed that this expression always returns a node list, and each returned node is processed separately and its converted value is added to a multi-valued field.
- `target` - the name of the target field in the current document where the value(s) will be stored.
- `append` - if true, values extracted from the RPC response will be appended to the target multi-valued field. When this is false, the default, existing values in the target field will be discarded and replaced with the values from the RPC response.
- `xml` - if true, the extracted DOM nodes will be separately serialized to XML and the resulting XML-formatted text will be added to the target field. When this is false, the default, the extracted DOM nodes will be flattened and converted to a list of fields. Field names in this case will correspond to XML element names, dot-separated, and element attributes will be represented as fields with `@attributeName` suffix.

8.27.3. Examples

Create a REST Query stage to merge results from another Solr system:

```

{
  "type" : "indexing-rpc",
  "id" : "BasicSolrCall",
  "mappingRules" : [ {
    "path" : "//result/doc[1]/str[@name='foo_s']/text()",
    "target" : "foo_s",
    "append" : true,
    "xml" : false
  }, {
    "path" : "//result/doc/arr[@name='a_txt']/*",
    "target" : "doc_txt",
    "append" : true,
    "xml" : false
  } ],
  "debug" : true,
  "params" : {
    "uri" : "solr://collection1/select",
    "method" : "GET",
    "queryParams" : {
      "q" : "a_txt:${doc_value}"
    },
    "headers" : { }
  },
  "skip" : false,
  "label" : "indexing-rpc"
}

```

Upload a stopwords list:

```

{
  "type" : "indexing-rpc",
  "id" : "demo",
  "debug" : true,
  "params" : {
    "uri" : "http://localhost:8765/api/v1/stopwords/movies",
    "method" : "PUT",
    "headers" : { "Content-Type" : "application/json" },
    "entity" : "New stopwords list"
  },
  "skip" : false,
  "label" : "Indexing RPC Demo",
  "type" : "indexing-rpc"
}

```

8.27.4. Configuration

Tip

When entering configuration values in the UI, use *unescaped* characters, such as `\t` for the tab character. When entering configuration values in the API, use *escaped* characters, such as `\\t` for the tab character.

When using Fusion's REST-API, the ID of this stage is: `indexing-rpc`.

Configuration Properties

Property	Description, Type
debug Add Debugging Information	Setting to true will add a number of properties to either the context (in the Query case) or the document (in the indexing case) type: boolean default value: 'false'
hasNoSideEffects Run in simulation mode	To run this stage in simulation mode, set to 'true'. type: boolean default value: 'false'
mappingRules Mapping of Returned Values (as XPath Expressions) to Document Fields	type: array of object object attributes: \{ <ul style="list-style-type: none"> append : \{ <ul style="list-style-type: none"> display name: Append to Existing Values in Target Field type: boolean default value: 'false' } path (required) : \{ <ul style="list-style-type: none"> display name: XPath expression type: string } target (required) : \{ <ul style="list-style-type: none"> display name: Target field type: string } xml : \{ <ul style="list-style-type: none"> display name: Add as an XML Fragment type: boolean default value: 'false' } }

Property	Description, Type
<p>params</p> <p>Call Parameters</p> <p><i>required</i></p>	<p>type: object</p> <p>object attributes: \{</p> <p> entity : \{</p> <p> display name: Request entity (as string)</p> <p> type: string</p> <p> }</p> <p> headers : \{</p> <p> display name: Request protocol headers</p> <p> type: object</p> <p> }</p> <p> method : \{</p> <p> display name: Call method</p> <p> type: string</p> <p> description : One of GET, POST, PUT, or DELETE</p> <p> enum: \{ get put post delete }</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <pre> } + 'queryParams' : \{ + display name: Query parameters + type: 'object' + } + 'uri' : \{ + display name: Endpoint URI + type: 'string' + } + } </pre> </div>

8.28. Send PagerDuty Message Index Stage

This stage sends a [PagerDuty](#) Message from Fusion, for alerting, monitoring, and more, using Fusion's Messaging Services.

Read more about PagerDuty integration in Fusion on our [blog](#).

8.28.1. Enabling PagerDuty Messaging

Before you can use the PagerDuty pipeline stage, you must enable PagerDuty messaging in Fusion:

1. Click **Applications > System > Messaging Services**.
2. Select **PagerDuty Message Service** from the drop-down menu.
3. Enter the following information:
 - [PagerDuty service key](#)
 - PagerDuty Service API URL; this should be https://events.pagerduty.com/generic/2010-04-15/create_event.json.
4. Click **Save message service**.

8.28.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `pagerduty-message`.

Configuration Properties

Property	Description, Type
client Client	The name of the monitoring client that is triggering this event. type: <code>string</code> default value: 'Fusion' minLength: 1
clientURL Client URL	The URL of the monitoring client that is triggering this event. type: <code>string</code> default value: 'fusion-monitoring.yourdomain.com' minLength: 1

Property	Description, Type
<p>errorKey</p> <p>Message Response Failure Key</p>	<p>The name of the key to store a boolean if sending a message failed. If set, you can check the MessageResponse errorCode and other attributes for the reason.</p> <p>type: string</p> <p>default value: 'messageResponseFailure'</p>
<p>eventType</p> <p>Event Type</p> <p><i>required</i></p>	<p>Select the Pager Duty Event Type.</p> <p>type: string</p> <p>default value: 'trigger'</p> <p>enum: \{ trigger acknowledge resolve }</p>
<p>incidentContextImages</p> <p>Incident Context Images</p>	<p>type: array of object</p> <p>object attributes: \{</p> <p>alt : \{</p> <p>display name: Alternate Text</p> <p>type: string</p> <p>description : HTML 'alt' tag for the image.</p> <p>} </p> <p>href (required) : \{</p> <p>display name: Target Link</p> <p>type: string</p> <p>description : URL to open when clicked on the image.</p> <p>} </p> <p>src (required) : \{</p> <p>display name: Source</p> <p>type: string</p> <p>description : HTML 'src' tag for the image. Should be always secure connection (https://).</p> <p>} </p> <p>} </p>

Property	Description, Type
<p>incidentContextLinks</p> <p>Incident Context Links</p>	<p>type: array of object</p> <p>object attributes: \{</p> <p>href (<i>required</i>): \{</p> <p>display name: Target Link</p> <p>type: string</p> <p>description : URL to open when clicked on the link.</p> <p>\}</p> <p>text : \{</p> <p>display name: Text</p> <p>type: string</p> <p>description : Arbitrary text explaining the URL.</p> <p>\}</p> <p>\}</p>
<p>incidentDescription</p> <p>Description</p> <p><i>required</i></p>	<p>A short description of the problem that led to this trigger. This field (or a truncated version) will be used when generating phone calls, SMS messages and alert emails. It will also appear on the incidents tables in the PagerDuty UI. The maximum length is 1024 characters.</p> <p>type: string</p> <p>default value: 'Sample Description'</p> <p>maxLength: 1024</p> <p>minLength: 1</p>
<p>incidentDetails</p> <p>Incident Details</p>	<p>type: array of object</p> <p>object attributes: \{</p> <p>name (<i>required</i>): \{</p> <p>display name: Name</p> <p>type: string</p> <p>\}</p> <p>value (<i>required</i>): \{</p> <p>display name: Value</p> <p>type: string</p> <p>\}</p> <p>\}</p>

Property	Description, Type
incidentKey Incident Key	Identifies the incident to which this trigger event should be applied. If there's no open (i.e. unresolved) incident with this key, a new one will be created. If there's already an open incident with a matching key, this event will be appended to that incident's log. type: <code>string</code> default value: <code>'Incident'</code> minLength: 1
responseKey Message Response Context Key	The name of the key to store the MessageResponse under in the Pipeline Context. type: <code>string</code> default value: <code>'messageResponse'</code>
storeInContext Add to Pipeline Context	Put the generated Message later in the pipeline. type: <code>boolean</code> default value: <code>'false'</code>

8.29. Send Slack Message Index Stage

This stage sends a [Slack](#) message from Fusion, for alerting, reporting, and more, using Fusion’s Messaging Services.

8.29.1. Enabling Slack Messaging

Before you can use the Slack pipeline stage, you must enable Slack messaging in Fusion:

1. Click **Applications > System > Messaging Services**.
2. Select **Slack Message Service** from the drop-down menu.
3. Enter the following information:
 - [Slack auth token](#)
 - Message template

The default is `<subject> : <body>`, which are configured with `messageSubjectTemplate` and `messageBodyTemplate` below. See Messaging Services Templates for details on the template language.

- Optionally, you can configure a proxy or the error reporting channel name.
4. Click **Save message service**.

8.29.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion’s REST-API, the ID of this stage is: `slack-message`.

Configuration Properties

Property	Description, Type
errorKey Message Response Failure Key	The name of the key to store a boolean if sending a message failed. If set, you can check the MessageResponse errorCode and other attributes for the reason. type: <code>string</code> default value: <code>'messageResponseFailure'</code>
from From <i>required</i>	Who the message is from. May be a string template similar to the body and subject. type: <code>string</code>

Property	Description, Type
messageBodyTemplate Message Body Template <i>required</i>	A Message Template that is used to create the message body to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: <code>string</code> default value: <code>'`Processing Document`'</code>
messageSubjectTemplate Message Subject Template <i>required</i>	A Message Template that is used to create the message subject to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: <code>string</code> default value: <code>'Hello'</code>
responseKey Message Response Context Key	The name of the key to store the MessageResponse under in the Pipeline Context. type: <code>string</code> default value: <code>'messageResponse'</code>
storeInContext Add to Pipeline Context	Put the generated Message later in the pipeline. type: <code>boolean</code> default value: <code>'false'</code>
to To <i>required</i>	Who to send the message to. May be a string template similar to the body and subject. type: <code>array of string</code>

8.30. Send SMTP Email Index Stage

The Send SMTP Email index stage was renamed for Fusion 3.0; it was called the Send email (via SMTP) stage in previous versions.

This stage sends an SMTP message from Fusion, for alerting, reporting, and more, using Fusion’s Messaging Services.

8.30.1. Enabling Email Messaging

Before you can use the Email pipeline stage, you must enable Email messaging in Fusion:

1. Click **Applications > System > Messaging Services**.
2. Select **SMTP Message Service** from the drop-down menu.
3. Verify that the default settings are sufficient.
4. Click **Save message service**.

8.30.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion’s REST-API, the ID of this stage is: `smtp-message`.

Configuration Properties

Property	Description, Type
errorKey Message Response Failure Key	The name of the key to store a boolean if sending a message failed. If set, you can check the MessageResponse errorCode and other attributes for the reason. type: <code>string</code> default value: <code>'messageResponseFailure'</code>
from From <i>required</i>	Who the message is from. May be a string template similar to the body and subject. type: <code>string</code>

Property	Description, Type
messageBodyTemplate Message Body Template <i>required</i>	A Message Template that is used to create the message body to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: <code>string</code> default value: ``Processing Document`
messageSubjectTemplate Message Subject Template <i>required</i>	A Message Template that is used to create the message subject to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: <code>string</code> default value: `Hello`
responseKey Message Response Context Key	The name of the key to store the MessageResponse under in the Pipeline Context. type: <code>string</code> default value: `messageResponse`
smtpPassword SMTP Password <i>required</i>	The SMTP password for the user credentials. type: <code>string</code>
smtpUser SMTP Username <i>required</i>	The SMTP user to send the message from. type: <code>string</code>
storeInContext Add to Pipeline Context	Put the generated Message later in the pipeline. type: <code>boolean</code> default value: `false`

Property	Description, Type
to To <i>required</i>	Who to send the message to. May be a string template similar to the body and subject. type: <code>array of string</code>

8.31. Set Property Index Stage

The Set Property Index Stage is used to set a value on a document, or into the Pipeline Context for downstream consumption by specifying a series of simple matching conditions. These conditions match against whether a field exists, and simple substring matches on the field contents. For more complex logic, use a JavaScript Index Stage.

8.31.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `index-set-property`.

Configuration Properties

Property	Description, Type
falseOutputValue False Output Value	The value to set when nothing matches. If null/empty, then nothing will be output (same behavior as in Fusion 2.2). May be a String Template. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md type: <code>string</code>
ifEquals If Value Equals	Check whether a value equals another value. Valid entries for this are: another field/context key, the word 'null' or 'not null' (without the quotes) or a literal string (e.g. "match") in quotes (single or double). The value may be a String Template. type: <code>string</code>
outputKey Output Key	The name of the key to insert into the output location (document or context). May be a String Template. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md type: <code>string</code> default value: 'setPropertyKey'

Property	Description, Type
outputType Output Type <i>required</i>	Select whether the flag should be set on the document or in the Pipeline Context. type: <i>string</i> default value: 'context' enum: \{ document context }
outputValue Output Value	The value to set. May be a String Template. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md type: <i>string</i> default value: 'true'
regularExpression Regular Expression Match	Apply the regular expression to the source fields and or context keys. If any of them match, than set the property. The implementation returns true if the regular expression matches anywhere in the string. type: <i>string</i>
source Source	The source fields and context keys to check conditions against. If only the source fields are set and nothing else, then only set the property if all fields are present and non-null on the document. May be a String Template. type: <i>array of string</i>
whatMatchedKey What Matched Key	The name of the context key to use to store a space separated list of what conditionals matched. type: <i>string</i> default value: 'whatMatched'

8.32. Solr Dynamic Field Name Mapping

The Solr Dynamic Field Name Mapping index stage was introduced in Fusion 3.0. This field can be used to re-name existing fields to match Solr dynamic field names, so that they will be typed correctly by Solr. It also can be used to specify fields which should be exempt from name changes as well as to set the dynamic type suffixes to a specified field. Note that the field that is processed by a Field Mapping stage prior to this stage will not be renamed, even if the stage config matches the field name.

8.32.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `solr-dynamic-field-name-mapping`.

Configuration Properties

Property	Description, Type
advancedTextFieldsIndexing Text Fields Advanced Indexing	This option enables indexing of a text data (not exceeding certain max length) into both tokenized and non-tokenized fields. For example a field 'name' containing "John Smith", if not hinted or finalized in previous stages, without this option enabled will be indexed into a "text" type field by default (name_t). But if the option is enabled, it will be indexed into both 'name_t' and 'name_s' fields allowing relevant search using 'name_t' field (by matching to a 'Smith' query) and also proper faceting and sorting using 'name_s' field (using value 'John Smith' for sorting or faceting). type: <code>boolean</code> default value: 'true'
duplicateSingleValuedFields Duplicate Single-Valued Fields as Multi-Valued Fields	This option enables indexing of a field data into both single-valued and multi-valued Solr fields. For example, a field 'phone' hinted as a "string" type and containing only one value will be indexed into a 'phone_s without this option enabled. However, if the option is enabled, the 'phone' field will be indexed into both 'phone_s' (single-valued) and 'phone_ss' (multi-valued) fields. type: <code>boolean</code> default value: 'false'

Property	Description, Type
fieldsToExclude Field Not To Map	The list of Pipeline Document fields to exclude from processing. type: <code>array of string</code>
maxTextLengthForAdvancedIndexing Max Length for Advanced Indexing of Text Fields	type: <code>integer</code> default value: '100'

8.33. Solr Indexer Stage

A Solr Indexer stage transforms a Fusion PipelineDocument object into a Solr document and sends it to Solr for indexing into a collection.

A PipelineDocument object contains fields which take as their values either a string or list of strings. Solr fields have a rich variety of types. The Solr Indexer stage transforms PipelineDocument field values into Solr document fields. The Solr Indexer stage can be configured so that it will try to ensure that all document fields are valid Solr fields. This feature is convenient, but offers very little control over how fields and field values are transformed, especially with respect to dates. A Date Parsing stage offers greater control over date values.

8.33.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `solr-index`.

Configuration Properties

Property	Description, Type
bufferDocsForSolr Buffer Documents and Send Them to Solr in Batches	type: <code>boolean</code> default value: <code>'false'</code>
dateFormats Additional Date Formats	type: <code>array of string</code>
enforceSchema Map to Solr Schema	type: <code>boolean</code> default value: <code>'true'</code>
params Additional Update Request Parameters	type: <code>array of object</code> object attributes: <code>{</code> <code>key (required) : {</code> display name: Parameter Name type: <code>string</code> } <code>value : {</code> display name: Parameter Value type: <code>string</code> } }

Property	Description, Type
<p>unmapped</p> <p>Unmapped Fields Mapping</p>	<p>type: object</p> <p>object attributes: \{</p> <p>operation : \{</p> <p>display name: Operation</p> <p>type: string</p> <p>default value: 'copy'</p> <p>description : The type of mapping to perform: move, copy, delete, add, set, or keep.</p> <p>enum: \{ copy move delete set add keep }</p> <pre> } + 'source' : \{ + display name: Source Field + type: 'string' + description : The name of the field to be mapped. + } + 'target' : \{ + display name: Target Field + type: 'string' + description : The name of the field to be mapped to. + } + } </pre>

8.34. Solr Partial Update Indexer Stage

The Solr Partial Update Indexer Stage updates one or more fields of an existing Solr document in a collection managed by Fusion. It provides an alternative to the Solr Indexer stage. The Solr Partial Update Indexer stage was introduced in [Fusion 2.1](#).

When a data feed consists of an ongoing flow of messages about known documents in a collection, such as item price, inventory counts, or weather conditions at a location, this stage provides fast indexing throughput and can be configured to enforce data atomicity to guarantee that the index always reflects the most recent update.

This stage is configured with a set of update directives based on Solr's [atomic updates](#). At run time, it creates a Solr update by applying these directives to the data from a Fusion PipelineDocument object and then submits this update to Solr's update handler.

Note	Solr's atomic update functionality requires that the schema for a collection is configured so that all fields have the attribute <code>stored="true"</code> , excepting fields which are <code><copyField/></code> destinations which must be configured as <code>stored="false"</code> .
------	---

8.34.1. Example Stage Specification

Configuration for a Partial Updater Stage in JSON:

```
{ "type" : "solr-partial-update-index",
  "enforceSchema" : false,
  "solrDocIdFieldName" : "id",
  "solrDocIdFieldValue" : "<doc.id>",
  "updatedFields" : [
    { "updateType" : "set", "fieldName" : "statusValue", "values" : "<doc.statusValue>" },
    { "updateType" : "set", "fieldName" : "lastCommunicationTime", "values" : "<doc.lastCommunicationTime>" }
  ],
  "concurrencyControlEnabled" : true,
  "skip" : false,
  "label" : "solr-partial-update-index",
}
```

The expression `<doc.X>` will evaluate to the contents of the current PipelineDocument's field named "X".

8.34.2. Types of Update Operations

The set of update operations are based on operations supported by Solr. They are:

- 'add' - add a new value or values to an existing Solr document field, or add a new field and value(s).
- 'set' - change the value or values in an existing Solr document field.
- 'remove' - remove all occurrences of the value or values from an existing Solr document field.
- 'removeregex' - remove all occurrences of the values which match the regex or list of regexes from an existing Solr document field.
- 'increment' - increment the the numeric value of existing Solr document field by a specific amount.
- 'decrement' - decrement the the numeric value of existing Solr document field by a specific amount.

In addition, this stage introduces experimental "Positional" operations which can be used to add, set or remove exactly one element of a field which takes a list of values (i.e, a multi-valued field).

- 'positionalUpdates' - used to add or set the value at specific position.
- 'positionalRemoves' - used to delete an element at a specific position.

When a collection contains two or more multi-value fields which are maintained in parallel so that taken together, they act like a table stored column by column, a positional update operation updates several data cells across one row of the table. To maintain this kind of column-oriented table, the positional delete directive must specify all the fields in the document which logically comprise the table.

8.34.3. Document Identifier Field

A Fusion collection is a Solr collection managed by Fusion. Underlyingly, a Solr document is a list of named, typed fields. The Solr [unique key field](#) stores a string which is the unique identifier for that document. There is at most one UniqueKey field per document, which is defined in the Solr schema. The UniqueKey field value is required. For collections created via Fusion, the UniqueKey field is named "id". Other document fields may also store string values which can be used as a unique identifier.

Solr uses the UniqueKey field to find the document to be updated. If the data feed information contains a document identifier which is different than the identifier value stored in the UniqueKey field, then this stage must do a Solr lookup to find the UniqueKey value.

8.34.4. Optimistic Concurrency

Solr's [Optimistic Concurrency](#) is a mechanism which checks whether or not a document has changed between the point at which an update request was submitted and the point at which the request is processed. Solr documents have an internal field named "*version*" which is updated whenever there is any change made to any of the other fields in that document. When optimistic concurrency control is on, update requests will be discarded if the current version of the document has changed since that request was made. This guarantees that the document will always reflect the most recent update. However, this requires an additional Solr lookup to get the current document version number, which is submitted as part of the update request.

8.34.5. Performance Considerations

In order to send a single update request to Solr, without preliminary lookup requests:

- The document identifier field should match the Solr collection's UniqueKey identifier field.
- Optimistic Concurrency should be turned off.
- Positional updates are experimental and potentially expensive, since all the values for all fields being updated must be fetched into memory in order to perform positional operations.

8.34.6. Solr Date Formats

```
"yyyy-MM-dd'T'HH:mm:ss'Z'", // Solr format without milliseconds
"yyyy-MM-dd'T'HH:mm:ss.SSS'Z'", // standard Solr format, with literal "Z" at the end
"yyyy-MM-dd'T'HH:mm:ss.SS'Z'", // standard Solr format, with literal "Z" at the end
"yyyy-MM-dd'T'HH:mm:ss.S'Z'" // standard Solr format, with literal "Z" at the end
```

See <https://cwiki.apache.org/confluence/display/solr/Working+with+Dates>

8.34.7. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion’s REST-API, the ID of this stage is: `solr-partial-update-index`.

Configuration Properties

Property	Description, Type
concurrencyControlEnabled Enable Concurrency Control	Select to enable Optimistic Concurrency Control in Solr, guaranteeing that the document update will not be overridden by another Partial Update to the same document. If disabled, in the case of an edit collision, the last committed update to the document will win. type: <code>boolean</code> default value: ‘true’
customRouteFieldName Custom Route Field Name	This option is used when custom shard routing is configured in Solr so the document route is defined by value of Solr document’s field (defined as ‘router.field’ when created the collection). If set here, the field with this name will be transferred to the partial update Solr document from the pipeline document. type: <code>string</code>
dateFormats Additional Date Formats	type: <code>array of string</code>
deletedFields Deletions	Fields to Delete from Solr Document. type: <code>array of object</code> object attributes: <code>\{</code> <code>fields (required) : \{</code> display name: Field type: <code>string</code> <code>}</code> <code>}</code>

Property	Description, Type
<p>enforceSchema</p> <p>Map to Solr Schema</p>	<p>type: boolean</p> <p>default value: 'true'</p>
<p>params</p> <p>Additional Update Request Parameters</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> key (required): \{ <ul style="list-style-type: none"> display name: Parameter Name type: string value: \{ <ul style="list-style-type: none"> display name: Parameter Value type: string <p>\}</p>
<p>positionalRemovals</p> <p>Positional Removals</p>	<p>Update Field or Group of Fields to remove value at a specific position. See documentation for additional information.</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> fields (required): \{ <ul style="list-style-type: none"> display name: Fields List type: string description : The field [field ...] list of fields (Solr field names) where removal of a value at specified position should happen. position (required): \{ <ul style="list-style-type: none"> display name: Position type: string description : The position at which the field value will be removed. Could be 'first', 'last' or numeric value (position index). <p>\}</p>

Property	Description, Type
<p>positionalUpdates</p> <p>Positional Updates</p>	<p>Update Field or Group of Fields to update (add or set) value at a specific position. See documentation for additional information.</p> <p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> fieldsAndValues (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Fields and Values type: string description : The field:value [,field:value ...] list. The values will be changed for specified fields at specified position. The list separator is comma (,) if the comma present in the field value, escape it with a backslash (\). Quotation marks("") can be used to enclose Field value to preserve the white spaces, if needed. \} position (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Position type: string description : The position at which the new value(s) will be changed. Could be 'first', 'last' or numeric value (position index). \} positionalUpdateType : \{ <ul style="list-style-type: none"> display name: Update Type type: string default value: 'set' description : The Update Type enum: \{ set add \} <div style="border: 1px solid #ccc; border-radius: 5px; padding: 5px; margin-top: 10px;"> <pre> } + }</pre> </div>
<p>rejectUpdatesIfDocNotPresent</p> <p>Reject Update if Solr Document is not Present</p>	<p>Whether to reject the update attempt if the document with given id is not present in Solr. This is not typical situation since the updates usually are performed on existing documents, however you may disable this to attempt update even if the document is not present. If the concurrency control is disabled, enabling this flag will force set the <i>version</i> field to 1, or to 0 otherwise.</p> <p>type: boolean</p> <p>default value: 'true'</p>

Property	Description, Type
<p>solrDocIdFieldName</p> <p>Solr Document ID Field Name</p> <p><i>required</i></p>	<p>type: <code>string</code></p> <p>default value: <code>'id'</code></p>
<p>solrDocIdFieldValue</p> <p>Solr Document ID Field Value</p> <p><i>required</i></p>	<p>type: <code>string</code></p> <p>default value: <code>''</code></p>
<p>updateAllDocFields</p> <p>Process All Pipeline Doc Fields</p>	<p>If this option is set, the Partial Update Stage will process pipeline document fields even if they are not set by Updates and Deletions instructions here. In this case those fields will be included into the partial update document and will be processed by Solr according to atomic update rules, i.e. non-map field value(s) will be treated as a 'set' update for the field, and Map field values will be processed as an atomic update defined in the Map. The Map structure should comply to Solr atomic update rules. Note that the Partial Update stage does NOT validate consistency of fields that are not Updates or Deletions configured here, it just sends them to Solr 'as is'.</p> <p>type: <code>boolean</code></p> <p>default value: <code>'false'</code></p>

Property	Description, Type
<p>updatedFields</p> <p>Updates</p>	<p>Fields to update (set, add or remove field values) in the Solr Document.</p> <p>type: array of object</p> <p>object attributes: \{</p> <p>fieldName (<i>required</i>): \{</p> <p>display name: Field Name</p> <p>type: string</p> <p>description : The Solr Document Field to update.</p> <p>\}</p> <p>updateType : \{</p> <p>display name: Update Type</p> <p>type: string</p> <p>default value: 'set'</p> <p>description : The Update Type</p> <p>enum: \{ set add remove remove_regex increment decrement \}</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <pre> } + 'values' _ (required)_ : \{ + display name: Value + type: 'string' + description : For increment operation only one value (positive or negative int) is allowed. For add, set, remove or remove_regex a single value or list of values can be specified. The list separator is a comma (,) if the comma should be present in the field value, escape it with a backslash (\). Quotation marks("") can be used to enclose Field value to preserve the white spaces, if needed. + } + } </pre> </div>

8.35. Tag Part-of-Speech Index Stage

The Tag Part-of-Speech Index stage (previously called the Part of Speech stage) operates over one or more fields in the Pipeline Document. It marks sentences with part of speech information as annotations which can be used by downstream indexing stages. Therefore this stage requires a Detect Sentences stage defined over these fields earlier in the pipeline.

This stage uses [Apache OpenNLP](#) project's [Part of Speech Tagger](#) to mark tokens with their corresponding word type based on the token itself and the context of the token. The OpenNLP documentation states:

"A token might have multiple pos tags depending on the token and the context. The OpenNLP POS Tagger uses a probability model to predict the correct pos tag out of the tag set. To limit the possible tags for a token a tag dictionary can be used which increases the tagging and runtime performance of the tagger."

Fusion comes with a set of OpenNLP language models for english. These data files are found in the directory: [fusion/3.1.x/data/nlp/models](#).

More models are available from the OpenNLP [models SourceForge repository](#). Model files must be uploaded to Fusion using the Fusion Blob Store service via the REST API.

8.35.1. Part-of-speech Tagging in a NLP Pipeline

The following video shows how to use a Part-of-speech indexing stage as part of an NLP pipeline:

8.35.2. Stage Setup

Here is an example of how to upload a part-of-speech model file to the Fusion blob store:

INPUT

```
curl -u user:pass -X PUT --data-binary @en-pos-maxent.bin -H 'Content-type: text/plain'  
http://localhost:8764/api/apollo/blobs/en-pos-maxent.bin
```

OUTPUT

```
{  
  "name" : "en-pos-maxent.bin",  
  "contentType" : "text/plain",  
  "size" : 5696197,  
  "modifiedTime" : "2015-07-15T06:57:48.636Z",  
  "version" : 0,  
  "md5" : "db2cd70395b9e2e4c6b9957015a10607"  
}
```

This is an example setup of this stage using the previously loaded .bin file:

INPUT

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id":"TagPartofSpeech1", "type": "tag-part-of-speech", "tokenizerModel":"en-pos-maxent.bin", "posModel":"en-pos-perceptron.bin", "source": ["sample", "text", "for", "NLP"]}' http://localhost:8764/api/apollo/index-stages/instances
```

OUTPUT

```
{
  "type" : "tag-part-of-speech",
  "id" : "TagPartofSpeech1",
  "posModel" : "en-pos-perceptron.bin",
  "tokenizerModel" : "en-sent.bin",
  "source" : [ "sample", "text", "for", "NLP" ],
  "skip" : false,
  "label" : "tag-part-of-speech",
  "type" : "tag-part-of-speech"
}
```

8.35.3. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `part-of-speech`.

Configuration Properties

Property	Description, Type
posModel Part of Speech Model <i>required</i>	type: <code>string</code> reference: model:open-nlp
source Source Fields <i>required</i>	type: <code>array of string</code> minimum number of items (<code>minItems</code>): 1
tokenizerModel Tokenizer Model <i>required</i>	type: <code>string</code> reference: model:open-nlp

8.36. Update Experiment Stage

The Update Experiment index stage was introduced in Fusion 2.4. This stage is part of Fusion’s Machine Learning framework. It is highly experimental and subject to change.

The Machine Learning framework provides tools to evaluate alternative methods of computing and displaying search results. For example, an experiment may consist of a system which has two or more different query pipelines running, and users are randomly served search results using some variant of the system via instrumented pages that capture user response. In such a system, an Update Experiment stage would be used to feed these responses directly back into the running experiment.

The Experiment Update stage checks each document for an experiment ID, a variant ID and a value and updates the experiment accordingly. A default experiment ID can be provided if none is found in the document. By default these documents are silently discarded after processing, but the stage can be configured to forward them down the pipeline to the next stages.

8.36.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion’s REST-API, the ID of this stage is: `experiment-update`.

Configuration Properties

Property	Description, Type
defaultExperimentId Default Experiment ID	type: <code>string</code>
experimentIdField Field With Experiment ID	type: <code>string</code>
experimentVariantField Field With Variant ID	type: <code>string</code>
forward Forward Messages to the Next Stage	type: <code>boolean</code> default value: <code>'false'</code>

8.37. Write Log Message Index Stage

The Write Log Message Index stage (called the Log a Message stage in versions earlier than 3.0) is an extension of Fusion's Logging Index Stage, which logs any message sent to the configured logging system using the Messaging Services.

8.37.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `logging-message`.

Configuration Properties

Property	Description, Type
errorKey Message Response Failure Key	The name of the key to store a boolean if sending a message failed. If set, you can check the MessageResponse errorCode and other attributes for the reason. type: <code>string</code> default value: 'messageResponseFailure'
from From <i>required</i>	Who the message is from. May be a string template similar to the body and subject. type: <code>string</code>
logLevel Log Level	The Log Level. May be: debug, info, warn, error type: <code>string</code> enum: <code>{ debug info warn error }</code>
messageBodyTemplate Message Body Template <i>required</i>	A Message Template that is used to create the message body to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: <code>string</code> default value: ' `Processing Document` '

Property	Description, Type
messageSubjectTemplate Message Subject Template <i>required</i>	A Message Template that is used to create the message subject to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: string default value: 'Hello'
responseKey Message Response Context Key	The name of the key to store the MessageResponse under in the Pipeline Context. type: string default value: 'messageResponse'
storeInContext Add to Pipeline Context	Put the generated Message later in the pipeline. type: boolean default value: 'false'
to To <i>required</i>	Who to send the message to. May be a string template similar to the body and subject. type: array of string

8.38. XML Transformation Index Stage

The XML Transformation stage (previously called the XML Transform Stage) allows you to process an XML document into one or more Solr documents and to specify mappings between elements and document fields. A common use case for an XML Transformation stage in a pipeline is when the XML document is a container-like document which contains a set of inner elements, each of which should be treated as a separate document. A parent ID field can be used to relate these multiple documents back to the containing document.

8.38.1. Pipeline Configuration

The default XML processing provided by the Apache Tika Parser index stage extracts all text from an XML into a single document field called `content`. This not only flattens the document contents, it loses all information about the containing elements in the document. To process XML documents using an XML Transformation stage, the index pipeline must have as its initial processing stage an Apache Tika Parser index stage which is configured to pass the document through to the XML Transformation stage *as raw XML*, via the following configuration:

- UI checkbox "Add original document content" **unchecked** / REST API property "addOriginalContent" set to **false**
- UI checkbox "Return parsed content as XML or HTML" **checked** / REST API property "keepOriginalStructure" set to **true**
- UI checkbox "Return original XML and HTML instead of Tika XML output" **checked** / REST API property "returnXml" set to **true**

With this configuration, the Tika parser stage decodes the raw input stream of bytes into a string containing the entire XML document which is returned in the PipelineDocument field `body`.

The pipeline must have a Field Mapping stage after the XML Transformation stage, before the Solr Indexer stage. The Field Mapping stage is used to remove the following fields from the document:

- `raw-content`
- `Content-Type`
- `Content-Length`
- `parsing`
- `parsing_time`

8.38.2. XML Transforms

The XML Transformation stage uses a Solr `XPathRecordReader` which is a streaming XML parser that supports *only a limited subset of XPath selectors*. It provides exact matching on element attributes and it can only extract the element text, not attribute values.

Examples of allowed XPath specifications where "a", "b", "c" are any element tags, likewise "attrName" is any attribute name:

```
/a/b/c
/a/b/c[@attrName='someValue']
/a/b/c[@attrName=]/d
/a/b/c/@attrName
//b//...
```

Note	When specifying the list of mappings , for each mapping, the specification for the xpath attribute must include the full path, i.e., the xpath attribute will include the rootXPath . See the example configuration below.
------	--

8.38.3. Example Stage Specification

Definition of an XML-Transformation stage that extracts elements from a MEDLINE/Pubmed article abstract:

```

{ "type" : "xml-transform",
  "id" : "n0j2a9k9",
  "rootXPath" : "/MedlineCitationSet/MedlineCitation",
  "bodyField" : "body",
  "mappings" : [ {
    "xpath" : "/MedlineCitationSet/MedlineCitation/Article/ArticleTitle",
    "field" : "article-title_txt",
    "multivalued" : false
  }, {
    "xpath" : "/MedlineCitationSet/MedlineCitation/Article/Abstract/AbstractText",
    "field" : "article-abstract_txt",
    "multivalued" : true
  }, {
    "xpath" : "/MedlineCitationSet/MedlineCitation/MeshHeadingList/MeshHeading/DescriptorName",
    "field" : "mesh-heading_txt",
    "multivalued" : true
  }, {
    "xpath" : "/MedlineCitationSet/MedlineCitation/PMID",
    "field" : "pmid_txt",
    "multivalued" : false
  } ],
  "keepParent" : false,
  "skip" : false,
  "label" : "medline_xml_transform",
}

```

Template for a minimal index pipeline that includes an XML-Transformation stage. Replace the XPath and field names in the XML-Transformation stage according to your data.

```

{
  "id" : "xml-pipeline-default",
  "stages" : [ {
    "type" : "tika-parser",
    "includeImages" : false,
    "flattenCompound" : false,
    "addFailedDocs" : false,
    "addOriginalContent" : false,
    "contentField" : "_raw_content_",
    "returnXml" : true,
    "keepOriginalStructure" : true,
    "extractHtmlLinks" : false,
    "extractOtherLinks" : false,
    "csvParsing" : false,
    "skip" : false,
    "label" : "tika",
    "sourceField" : "_raw_content_"
  }, {
    "type" : "xml-transformation",
    "rootXPath" : "/ROOTS/ROOT",
    "bodyField" : "body",
    "mappings" : [ {
      "xpath" : "/ROOTS/ROOT/element",
      "field" : "element-field_t",
      "multivalue" : false
    } ],
    "keepParent" : false,
    "skip" : false,
    "label" : "xml"
  }, {
    "type" : "field-mapping",
    "mappings" : [ {
      "source" : "parsing",
      "operation" : "delete"
    }, {
      "source" : "parsing_time",
      "operation" : "delete"
    }, {
      "source" : "Content-Type",
      "operation" : "delete"
    }, {
      "source" : "Content-Length",
      "operation" : "delete"
    } ],
    "skip" : false,
    "label" : "field mapping"
  }, {
    "type" : "solr-index",
    "enforceSchema" : true,
    "bufferDocsForSolr" : false,
    "skip" : false,
    "label" : "solr-index"
  } ]
}

```

8.38.4. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `xml-transform`.

Configuration Properties

Property	Description, Type
bodyField Body Field Name	type: <code>string</code> default value: 'body'
keepParent Keep Parent Document	type: <code>boolean</code> default value: 'false'
mappings XPath Mappings	type: <code>array of object</code> object attributes: <code>{</code> <code>field (required): {</code> display name: Field type: <code>string</code> <code>}</code> <code>multivalue: {</code> display name: Multi Value type: <code>boolean</code> default value: 'false' <code>}</code> <code>xpath (required): {</code> display name: XPath Expression type: <code>string</code> <code>}</code> <code>}</code>

Property	Description, Type
metadata Additional Metadata	type: array of object object attributes: \{ field (<i>required</i>): \{ display name: Field type: string } value (<i>required</i>): \{ display name: Value type: string } \}
parentIdField Parent ID Field Name	type: string
rootXPath Root XPath <i>required</i>	type: string

Chapter 9. Query Pipeline Stages

A query pipeline is made up of a series of query stages that process incoming search queries.

A pipeline stage definition associates a unique ID with a set of properties. These definitions are registered with the Fusion API service and stored in ZooKeeper for re-use across pipelines and search applications.

Fusion includes a number of specialized query stages as well as a JavaScript stage that allows advanced processing via a JavaScript program.

Use the The Query Workbench to configure stages in a query pipeline.

See these reference topics for details about each query pipeline stage:

9.1. Setup

- Active Directory Security Trimming
- Field Facet
- More Like This
- Query Fields
- Security Trimming

9.2. Results relevancy

- Block Documents
- Boost Documents
- Landing Pages
- Parameterized Boosting
- Recommend More Like This stage
- Boost with Signals stage
- Recommend Items for User stage
- Recommend Items for Item stage

9.3. Fetch data

- JDBC Lookup
- REST Query
- Solr Query
- Solr Subquery

9.4. Troubleshooting

- Logging

- Send PagerDuty Message
- Send Slack Message
- Send SMTP Email
- Write Log Message

9.5. Advanced

- Additional Query Parameters
- Javascript
- Retrieve Stored Parameters

9.6. Other

- Analytics Catalog Query
- Call Pipeline
- Experiment Query
- Machine Learning
- Parameterized Faceting
- Return Query Parameters
- Rollup Aggregation

9.7. Active Directory Security Trimming Stage

An Active Directory Security Trimming query-pipeline stage retrieves an Active Directory user's security identifiers to build a security filter. This restricts the documents in the query result to only those documents for which a user has access permissions. Security trimming is commonly used in business to authenticate between administrative users and normal users, or to limit the site access of website users according to a login/password.

9.7.1. Example Stage Setup

Active Directory Security Trimming query stage setup:

Input

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id":"ADSecurity", "type": "active-directory-security-trimming", "server":"ldap://hostname:port", "bindName":"ADuser@example.com", "bindPassword":"login1"}' http://localhost:8764/api/apollo/query-stages/instances
```

Output

```
{
  "type" : "active-directory-security-trimming",
  "id" : "ADSecurity",
  "server" : "ldap://hostname:port",
  "bindName" : "ADuser@example.com",
  "bindPassword" : "login1",
  "enableCache" : true,
  "cacheSize" : 1000,
  "expirationTime" : 3600,
  "skip" : false,
  "label" : "active-directory-security-trimming",
  "type" : "active-directory-security-trimming"
}
```

9.7.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `active-directory-security-trimming`.

Configuration Properties

Property	Description, Type
bindName Bind Name <i>required</i>	username in userPrincipalName format (e.g. user@example.com) type: string format: @.
bindPassword Bind Password <i>required</i>	type: string
cacheSize Cache Size	type: integer default value: '1000'
datasources Restrict filter to Datasource(s)	A list of Fusion datasources to which security-trimming should be restricted, allowing content from other datasources to pass through un-filtered; if empty, all matching content is subject to filtering. type: array of string
enableCache Enable Cache	type: boolean default value: 'true'
expirationTime Cache Expiration Time	(in seconds) type: integer default value: '3600'
filterAttribute Filter Attribute	Active Directory attribute to use as the security-trimming filter criterion type: string enum: \{ objectSid sAMAccountName userPrincipalName }

Property	Description, Type
<p>overrideUserIdentityHandling</p> <p>Override Default User Identity Handling?</p>	<p>Default handling first attempts to take the user identity from a 'fusion-user-id' http-header, which is the logged-in user ID from the Fusion proxy service. If that value is empty, a 'username' query parameter is tried instead. When this DataSource property is enabled, the specified source and key properties are used explicitly, without any fallback behavior.</p> <p>type: boolean</p> <p>default value: 'false'</p>
<p>server</p> <p>Active Directory Url</p> <p><i>required</i></p>	<p>E.g. ldap://hostname:port</p> <p>type: string</p> <p>format: ldap://(:\d)?</p>
<p>userIdentityKey</p> <p>User ID key</p>	<p>e.g. username, userID, etc.</p> <p>type: string</p> <p>default value: 'username'</p>
<p>userIdentitySource</p> <p>User ID source</p>	<p>Specify whether the value comes from an http header or query parameter.</p> <p>type: string</p> <p>default value: 'query_param'</p> <p>enum: \{ query_param header }</p>

9.8. Additional Query Parameters Stage

This stage was renamed to Additional Query Parameters in Fusion 3.0. It was called the Set Query Params stage in previous versions.

The Additional Query Parameters query-pipeline stage is used to set, append, and remove [Solr query parameters](#). This stage takes a list of query parameters. Each parameter is specified as a triple consisting of parameter name, parameter value, and update policy.

Available update policies are: **replace**, **append**, **remove** and **default**. The policy 'default' means that this parameter will be added only if it has not yet been set via the request or by a default specification in the Solr config file solrconfig.xml.

9.8.1. Example Stage Specification

Set the response format to JSON:

```
{ "type":"set-params",
  "id":"default-params",
  "params": [
    {"key":"wt", "value":"json", "policy":"default"}
  ],
  "skip":false
}
```

Return a field named "id", and 10 rows of results:

```
{ "type":"set-params",
  "id":"default-params",
  "params": [
    {"key":"fl", "value":"id", "policy":"append"},
    {"key":"rows", "value":"10", "policy":"replace"}
  ],
  "skip":false
}
```

9.8.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: **set-params**.

Configuration Properties

Property	Description, Type
<p>params</p> <p>Parameters and Values</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> key (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Parameter Name type: string policy: \{ <ul style="list-style-type: none"> display name: Update Policy type: string default value: 'append' enum: \{ replace append remove default } <pre> } + 'value' : \{ + display name: Parameter Value + type: 'string' + } + } </pre>

9.9. Analytics Catalog Stage

The Fusion Analytics Catalog query-pipeline stage lets you define views of data stored in Fusion collections. You can query the objects in the Analytics Catalog using SQL, [Solr Streaming Expressions](#), or regular Solr queries. See also the Catalog API.

This query stage was introduced in Fusion version 3.0.

9.9.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `analytics-catalog-query`.

Configuration Properties

Property	Description, Type
catalogProjectId Catalog Project ID <i>required</i>	Analytics catalog project ID. type: <code>string</code> default value: 'fusion'
contextKey Context Key	The key used to bind query results into the pipeline context; if not set, the query results are returned to the client. type: <code>string</code>
params Additional Params to be Included in the Solr Query Request. Only Applies if the Query Type is 'solr'.	type: <code>array of object</code> object attributes: <code>{</code> <code>key (required) : {</code> display name: Parameter Name type: <code>string</code> } <code>value : {</code> display name: Parameter Value type: <code>string</code> } }

Property	Description, Type
<p>queryText</p> <p>Query</p> <p><i>required</i></p>	<p>Query to execute; accepts any valid Solr query, SQL, or Solr streaming expression. Replaceable parameters will be applied before query execution using values pulled from request parameters and the pipeline context.</p> <p>type: <i>string</i></p>
<p>queryType</p> <p>Query Type</p>	<p>Query type: solr, sql, or streaming_expression</p> <p>type: <i>string</i></p> <p>default value: 'sql'</p> <p>enum: \{ solr sql streaming_expression }</p>
<p>requestHandler</p> <p>Solr Request Handler.</p>	<p>The Solr request handler to send a query to; only applies if query type is 'solr' or 'streaming_expression'. Defaults to /select for solr queries and /stream for streaming expressions.</p> <p>type: <i>string</i></p>

9.10. Block Documents Stage

The Block Documents query-pipeline stage removes documents from the result based on a Block Documents rule which consists of the following elements:

- **field** - the document field to filter on.
- **keywords** - the words, phrases, or regex used as the filter.
- **mode** - filtering logic applied to keywords, one of:
 - **exact** - exact matching on any item in the keywords list.
 - **phrase** - phrase matching on the items in the keywords list.
 - **regex** - treat items in the keywords list as a regex.
 - **match** - requires a match for every item in the keyword list, but doesn't require phrase matching.
- **blocks** - a list of document IDs for documents which are always removed from the query result.

The block documents rule is used to craft a Solr query. The **keywords** are added to the **q** Solr query parameter, by default. The configuration property **queryParam** can be used to specify a different Solr query parameter to use as the keywords filter. The rest of the rule is processed into the **fq** Solr query parameter.

9.10.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: **blocks**.

Configuration Properties

Property	Description, Type
queryParam	type: string
Query Parameters for Matching	default value: 'q'

Property	Description, Type
<p>rules</p> <p>Block Rules</p>	<p>type: array of object</p> <p>object attributes: \{</p> <p> blocks : \{</p> <p> display name: Blocks</p> <p> type: array of string</p> <p> description : The documents to block.</p> <p> }</p> <p> field (required) : \{</p> <p> display name: Field</p> <p> type: string</p> <p> default value: 'id'</p> <p> description : The name of the Solr Field to use for blocks.</p> <p> }</p> <p> keyword (required) : \{</p> <p> display name: Keyword</p> <p> type: string</p> <p> description : Search keywords to match on.</p> <p> }</p> <p> mode (required) : \{</p> <p> display name: Match Strategy</p> <p> type: string</p> <p> default value: 'exact'</p> <p> description : How to match the keywords.</p> <p> enum: \{ exact phrase regex match }</p> <div style="border: 1px solid #ccc; border-radius: 5px; padding: 5px; margin-top: 10px;"> <p> } +</p> <p>}</p> </div>

9.11. Boost Documents Stage

The Boost Documents query-pipeline stage adds boosting parameters to matched documents based on specific search terms. Boosts are defined with a term value to boost and the boost factor to add. The boosting parameters are added to the `bq` Solr query parameter.

The Boost Documents rule consists of the following elements:

- `field` - the document field to filter on
- `keywords` - the words, phrases, or regex used as the filter
- `mode` - filtering logic applied to keywords, one of:
 - `exact` - the keyword and the query must match exactly. This is case-sensitive.
 - `phrase` phrase matching on the items in the keywords list
 - `regex` treat items in the keywords list as a regex
 - `match` must match every item in the keyword list, but doesn't require phrase matching
- `boosts` - consists of a list of pairs
 - `value` - one or more terms used for boosting
 - `boost` - the numeric boost value. Default `100`.

9.11.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `boosts`.

Configuration Properties

Property	Description, Type
<code>queryParam</code>	type: <code>string</code>
Query Parameter for Matching	default value: <code>'q'</code>

Property	Description, Type
<p>rules</p> <p>Boost Rules</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> boosts : \{ <ul style="list-style-type: none"> display name: Boosts type: array of object field (<i>required</i>) : \{ <ul style="list-style-type: none"> display name: Field type: string default value: 'id' description : The name of the Solr Field to use for boosting. keyword (<i>required</i>) : \{ <ul style="list-style-type: none"> display name: Keyword type: string description : Search keywords to match on. mode (<i>required</i>) : \{ <ul style="list-style-type: none"> display name: Match Strategy type: string default value: 'exact' description : How to match the keywords. enum: \{ exact phrase regex match } <p style="text-align: right;">} +</p> <p style="text-align: right;">}</p>

9.12. Boost with Signals Stage

This stage was renamed to Boost with Signals for Fusion 3.1; it was called the Recommendation Boosting stage in previous Fusion versions.

The Boost with Signals query-pipeline stage uses aggregated signals to selectively boost items in the set of search results. Boost with Signals implements one type of *collaborative filtering*. Signals contain relevance feedback, that is, information about what users do. For example, they might contain information about clicks or page views. Boosting occurs when the same *query* comes in again.

See Collaborative Filtering for more information.

Tip	This stage accesses the <code>signals_aggr</code> collection. Before using it, verify that the following permission is set: <code>GET:/solr/<collection-name>_signals_aggr/select</code>
-----	---

9.12.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `recommendation`.

Configuration Properties

Property	Description, Type
aggrType Aggregation Type	type: <code>string</code> default value: <code>*</code>
boostId Solr Field to Boost On	Which Solr field to use when applying recommendation boosts. type: <code>string</code> default value: <code>'id'</code>
boostingMethod Boost Method <i>required</i>	The boost method to use. <code>query-parser</code> should be chosen if <code>defType!=edismax</code> for main query. type: <code>string</code> default value: <code>'query-param'</code> enum: <code>\{ query-param query-parser }</code>

Property	Description, Type
boostingParam Boost Param <i>required</i>	'Boost' multiplies scores by the boost values whereas 'bq' adds optional clauses to main query. type: string default value: 'bq' enum: \{ boost bq }
numRecommendations Number of Recommendations	type: integer default value: '10'
numSignals Number of Signals	Number of signals to process when getting recommended items. type: integer default value: '100'
queryParams Solr Query parameters	Parameters for querying Signal aggregation collection type: array of object default value: <code>{ "key" => "qf", "value" => "query_t" }{ "key" => "pf", "value" => "query_t^3" }{ "key" => "pf", "value" => "query_t~2^7" }{ "key" => "pf", "value" => "query_t~0^1" }{ "key" => "mm", "value" => "50%" }{ "key" => "defType", "value" => "edismax" }{ "key" => "sort", "value" => "score desc, weight_d desc" }</code> object attributes: \{ key (required) : \{ display name: Parameter Name type: string } value : \{ display name: Parameter Value type: string } }

Property	Description, Type
rollupField Rollup Field	Field to use for rolling up documents that have same doc id's type: <code>string</code>
rollupWeightField Rollup weight field	Field to use for signal weights type: <code>string</code>
scaleRange Scale Boosts	Scale the boost values to a [min,max] range type: <code>object</code> object attributes: \{ <ul style="list-style-type: none"> <code>scaleMax</code> : \{ <ul style="list-style-type: none"> display name: Maximum value of the scale range type: <code>number</code> <code>scaleMin</code> : \{ <ul style="list-style-type: none"> display name: Minimum value of the scale range type: <code>number</code>

9.13. Run Query Pipeline Stage

The Run Query query-pipeline stage was introduced in Fusion 3.0.

9.13.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `run-query-pipeline`.

Configuration Properties

Property	Description, Type
collection Collection to Query	type: <code>string</code>
pipelineToRun Query Pipeline to Call <i>required</i>	type: <code>string</code> default value: 'default'
solrHandler Solr Handler to Use	type: <code>string</code>

9.14. Experiment Query Parameters Stage

This stage was introduced in Fusion 2.4.

The Experiment Query Parameters query-pipeline stage is part of Fusion’s Machine Learning framework. It is highly experimental and subject to change.

The Machine Learning framework provides tools to evaluate alternative methods of computing and displaying search results. For example, an experiment may consist of a system which has two or more different query pipelines running, and users are randomly served search results using some variant of the system via instrumented pages that capture user response. In such a system, the Experiment query stage selects a variant from a a running experiment and inject its properties into the current PipelineContext.

9.14.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion’s REST-API, the ID of this stage is: `experiment-query`.

Configuration Properties

Property	Description, Type
<code>defaultExperimentId</code>	type: <code>string</code>
Default Experiment ID	

9.15. Facets Stage

The Facets query-pipeline stage (previously called the Facet stage) is used to add a [Solr Field Facet query](#) to the search query pipeline. A Field facet query computes the top values for a field and returns the list of those values along with a count of the subset of documents in the search results which match that term. Field faceting works best over fields which contain a single label or set of labels from a finite, controlled lexicon such as product category. Facet field parameters can be tuned for performance, see: [Facet Field Configuration](#).

It's possible to specify more than one field facets. For each field facet you must specify the field name plus the following additional parameters:

- Limit – the maximum number of terms to be returned. Default 100.
- Offset – the number of top facet values to skip in the response. Default 0.
- Sort - the order in which to list facet values: `count` ordering is by documents per term, descending, and `index` ordering is sorted on term values themselves.
- Missing - the number of documents in the results set which have no value for the facet field.
- Choice of facet method (advanced) - specify Solr algorithm used to calculate facet counts. (See [Facet Method Configuration](#) for details). One of:
 - `enum` - small number of distinct categories
 - `fc` ("field cache") - many different values in the field, each document has low number of values, multi-valued field
 - `fcs` ("single value string fields") - good for rapidly changing indexes

For further details see: [Solr Wiki Faceting Overview](#).

9.15.1. Configuration

When using Fusion's REST-API, the ID of this stage is: `facet`.

Configuration Properties

Property	Description, Type
<p>fieldFacets</p> <p>Facet Fields</p>	<p>type: array of object</p> <p>object attributes: \{</p> <p>enumCacheMinDf : \{</p> <p>display name: Enum Cache Minimum DF</p> <p>type: integer</p> <p>exclusiveMinimum : false</p> <p>minimum : 0</p> <p>}</p> <p>field (<i>required</i>) : \{</p> <p>display name: Field</p> <p>type: string</p> <p>description : The field whose values you want to facet on.</p> <p>}</p> <p>limit : \{</p> <p>display name: Limit</p> <p>type: integer</p> <p>description : Maximum number of facets to return.</p> <p>}</p> <p>method : \{</p> <p>display name: Method</p> <p>type: string</p> <p>enum: \{ enum fc fcs }</p>

```

} +
`minCount` : \{ +
  display name: Minimum Count +
  type: `integer` +
  default value: `1` +
  description : Lower threshold of term counts to
be included. +
} +
`missing` : \{ +
  display name: Count Missing +
  type: `boolean` +
  default value: `false` +
  description : Optionally include a 'missing'
facet bucket for documents without the selected
field. +
} +
`offset` : \{ +
  display name: Offset +
  type: `integer` +
  description : Offset into list of resulting
facets. +
} +
`prefix` : \{ +
  display name: Prefix +
  type: `string` +
  description : Prefix of terms to facet on. +
} +
`sort` : \{ +
  display name: Sort +

```

9.16. JavaScript Stage

For a complete description of the JavaScript query-pipeline stage, see [Custom JavaScript Stages for Query Pipelines](#).

```

type: 'string' +
+
enum: \{ count index   }

```

9.16.1. Configuration

<p>Tip</p>	<p>When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.</p>				
<p>When using Fusion's REST-API, the ID of this stage is: <code>javascript-query</code>.</p> <p>Configuration Properties</p>	<pre> type: 'integer' + exclusiveMinimum : false + minimum : 0 + } + } </pre>				
<table border="1"> <thead> <tr> <th data-bbox="126 646 808 697">Property</th> <th data-bbox="815 646 1490 697">Description, Type</th> </tr> </thead> <tbody> <tr> <td data-bbox="126 705 808 911"> <p>script</p> <p>Script Body</p> <p><i>required</i></p> </td> <td data-bbox="815 705 1490 911"> <p>One context variable 'request' is visible to the script.</p> <p>type: <code>string</code></p> </td> </tr> </tbody> </table>	Property	Description, Type	<p>script</p> <p>Script Body</p> <p><i>required</i></p>	<p>One context variable 'request' is visible to the script.</p> <p>type: <code>string</code></p>	
Property	Description, Type				
<p>script</p> <p>Script Body</p> <p><i>required</i></p>	<p>One context variable 'request' is visible to the script.</p> <p>type: <code>string</code></p>				

9.17. JDBC Lookup Stage

This stage was renamed to JDBC Lookup in Fusion 3.0; it was called the JDBC Query Lookup stage in previous Fusion versions.

The JDBC Lookup query-pipeline stage is used to call out to a database as part of a pipeline stage, to inject results into either the context/request or the pipeline document. For example, if you needed to look up a user from a DB and add their profile information onto a request for downstream use in a pipeline, the JDBC Lookup Stage would facilitate this.

Note	You must first upload the JDBC driver to Fusion, see the Connector JDBC API page.
------	---

9.17.1. Example

An example of setup for a JDBC Lookup query-pipeline stage

Upload stage config via POST to Fusion REST API endpoint `api/apollo/query-stages/instances``

```
curl -u user:pass -X POST -H "Content-Type: application/json" -d '{"id": "jdbc-quer", "type": "jdbc-query-lookup", "driver": "postgresql-9.3-1101-jdbc4.jar", "connection": "jdbc:postgresql:database", "username": "user", "password": "password1", "preparedStatement": "select ID as id from DATABASE;"}' http://localhost:8764/api/apollo/query-stages/instances`
```

Response

```
{
  "type" : "jdbc-query-lookup",
  "id" : "jdbc-quer",
  "driver" : "postgresql-9.3-1101-jdbc4.jar",
  "connection" : "jdbc:postgresql:database",
  "username" : "user",
  "password" : "password1",
  "preparedStatement" : "select ID as id from DATABASE;",
  "fetchSize" : -1,
  "join" : true,
  "rows" : 10,
  "skip" : false,
  "label" : "jdbc-query-lookup",
  "type" : "jdbc-query-lookup"
}
```

9.17.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `jdbc-query-lookup`.

Configuration Properties

Property	Description, Type
connection Connection URI <i>required</i>	The connection string for the database. type: <i>string</i>
driver JDBC Driver <i>required</i>	The fully qualified class name of the JDBC Driver to use. type: <i>string</i>
fetchSize Fetch Size	The JDBC Fetch Size to use. If -1, use the driver default. type: <i>integer</i> default value: '-1'
join Join With Request	If true, the results will be added on to the request using the prefix key and the row id, else the results will be put in the pipeline context using type: <i>boolean</i> default value: 'true'
password Password <i>required</i>	The password to connect to the database. type: <i>string</i>
prefix Result Prefix Key	The string to use as a prefix for all values. type: <i>string</i>
preparedStatement SQL Prepared Statement <i>required</i>	The SQL Prepared Statement to execute when bound with values. type: <i>string</i>

Property	Description, Type
<p>preparedStatementKeys</p> <p>Prepared Statement Keys</p>	<p>The keys in the Request/Header/Context to use to map request attributes into the prepared statement. These must map to the '?'s in your prepared statement. They must also be able to be resolved as the first parameter of that name in a request.</p> <p>type: array of string</p>
<p>rows</p> <p>Rows</p>	<p>The number of rows to return.</p> <p>type: integer</p> <p>default value: '10'</p>
<p>username</p> <p>Username</p> <p><i>required</i></p>	<p>The username to connect to the database.</p> <p>type: string</p>

9.18. Landing Pages Stage

The Landing Pages query-pipeline stage provides the mechanism by which specific search queries will be pinned to certain URLs. This stage returns one or more URLs which can be used for redirection. It doesn't perform a redirection; this must be done by the calling application. The redirection URLs are returned in a separate section of the Fusion response object, with attribute name `fusion`.

This stage is configured using Landing Page Rules, which consist of the following:

- `keyword` - words, phrases, or a regex
- `mode` - filtering logic applied to the query, one of:
 - `exact` - the keyword and the query must match exactly. This is case-sensitive.
 - `phrase` phrase matching on the items in the keywords list
 - `regex` treat items in the keywords list as a regex
 - `match` must match every item in the keyword list, but doesn't require phrase matching
- `url` - a list of URLs

9.18.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `landing-pages`.

Configuration Properties

Property	Description, Type
matchCount Maximum Matches	type: <code>integer</code> default value: '1' exclusiveMinimum: false minimum: 1
queryParam Query Parameter for Matching	type: <code>string</code> default value: 'q'

Property	Description, Type
<p>rules</p> <p>Landing Page Rules</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> keyword (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Keyword type: string description : Search keywords to match on. mode (<i>required</i>): \{ <ul style="list-style-type: none"> display name: Match Strategy type: string default value: 'exact' description : How to match the keywords. enum: \{ exact phrase regex match } <p style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <pre> } + 'url' _ (required)_ : \{ + display name: Redirect URL + type: 'string' + description : The URL to return for the redirect. For 'regex' match types, the URL can contain %s characters which will get populated by the regex capture groups. + } + </pre> </p> <p>}</p>

9.19. Logging Stage

The Logging query-pipeline stage prints messages to the API log file; the default location is `fusion/3.1.x/var/log/api/api.log`.

The verbosity of this message is controlled by the property `detailed`. If true, then the current Request object will be pretty-printed to the log file. If false, only the basic information about this stage will be logged.

In a production environment logging stages should be configured with property `skip` set to `true`, if possible. Use of detailed logging may impact performance.

9.19.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `query-logging`.

Configuration Properties

Property	Description, Type
<code>detailed</code>	type: <code>boolean</code>
Detailed Logging	default value: 'false'

9.20. Machine Learning Stage

The Machine Learning query-pipeline stage uses a compiled machine learning model to analyze a field or fields of a Query Request object and stores the results of analysis in a new field added to either the Request or the PipelineContext object. You must use [Spark's MLlib API](#) to create a supervised machine learning model and upload this model into Fusion's blob store collection. Complete details are available in section: Machine Learning Models in Fusion

Successful use of this stage requires a proper understanding of both the model and your data. The machine learning model is described by its spark-mllib.json file, which contains the model specification as a JSON object. This object contains attribute "featureFields" which takes as its value a list of one or more field names. The contents of these fields are processed into the vector of features which the model operates on. If these fields aren't present in the request, then the result is either an empty prediction or a configurable default value. If the contents of these fields differ greatly from the data used to compile the model, the predictions made by the model will be unreliable.

9.20.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `machine-learning-query`.

Configuration Properties

Property	Description, Type
defaultValue Default Value	Value to provide if a prediction cannot be made for a document. type: <code>string</code>
failOnError Fail on Error	Flag to indicate if this stage should throw an exception if an error occurs while generating a prediction for a document. type: <code>boolean</code> default value: 'false'
modelId Machine Learning Model ID <i>required</i>	The ID of the ML model stored in the Fusion blob store. type: <code>string</code> minLength: 1

Property	Description, Type
<p>predictionFieldName</p> <p>Prediction Field Name</p> <p><i>required</i></p>	<p>Name of the field to store the prediction (model output) in the document.</p> <p>type: <i>string</i></p> <p>minLength: 1</p>
<p>queryFeatureFieldName</p> <p>Query Feature Field</p>	<p>Name of the field to extract query features from (model input) in the document.</p> <p>type: <i>string</i></p> <p>default value: 'q'</p> <p>minLength: 1</p>
<p>storeInContext</p> <p>Store the Prediction in the Context</p>	<p>Flag to indicate that the prediction should be set as a context property instead of setting a field on the document.</p> <p>type: <i>boolean</i></p> <p>default value: 'false'</p>

9.21. Parameterized Boosting Stage

This stage was renamed to Parameterized Boosting in Fusion 3.0; it was called the Advanced Boosting stage in previous Fusion versions.

The Parameterized Boosting query-pipeline stage reads the `boostValues` (in `List<DocumentResult>` format) from the context variable (added by a Rollup Aggregation stage or a JavaScript stage), and adds boosts to the main query using 'bq' or 'boost' based on the stage configuration. The weights for the boost values can also be scaled.

9.21.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `adv-boost`.

Configuration Properties

Property	Description, Type
boostFieldName Boost Field <i>required</i>	The field name to boost the values on. type: <code>string</code>
boostingMethod Boost Method <i>required</i>	The boost method to use. <code>query-parser</code> should be chosen if <code>defType!=edismax</code> for main query. type: <code>string</code> default value: 'query-param' enum: <code>\{ query-param query-parser }</code>
boostingParam Boost Param <i>required</i>	'Boost' multiplies scores by the boost values whereas 'bq' adds optional clauses to main query. type: <code>string</code> default value: 'boost' enum: <code>\{ boost bq }</code>

Property	Description, Type
key Context Key <i>required</i>	The key name to read from context for boost id and values. type: string
scaleRange Scale Boosts	Scale the boost values to a [min,max] range type: object object attributes: \{ <ul style="list-style-type: none"> scaleMax : \{ <ul style="list-style-type: none"> display name: Maximum value of the scale range type: number scaleMin : \{ <ul style="list-style-type: none"> display name: Minimum value of the scale range type: number

9.22. Parameterized Faceting Stage

This stage was introduced in Fusion 2.4 as the Auto Facet stage. In Fusion 3.0, it was renamed to the Parameterized Faceting stage.

The Parameterized Faceting query-pipeline stage facets documents in the query pipeline. The stage queries Solr to determine the most frequent value for a given field, and then uses that value to look up stored parameters to automatically select the most appropriate facet for a query.

9.22.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `auto-facet`.

Configuration Properties

Property	Description, Type
<p>facetField</p> <p>Facet Field</p> <p><i>required</i></p>	<p>type: object</p> <p>object attributes: \{</p> <p>enumCacheMinDf : \{</p> <p>display name: Enum Cache Minimum DF</p> <p>type: integer</p> <p>exclusiveMinimum : false</p> <p>minimum : 0</p> <p>\}</p> <p>field : \{</p> <p>display name: Field</p> <p>type: string</p> <p>description : The field whose values you want to facet on.</p> <p>\}</p> <p>limit : \{</p> <p>display name: Limit</p> <p>type: integer</p> <p>description : Maximum number of facets to return.</p> <p>\}</p> <p>method : \{</p> <p>display name: Method</p> <p>type: string</p> <p>enum: \{ enum fc fcs \}</p>

```

} +
`minCount` : \{ +
  display name: Minimum Count +
  type: `integer` +
  default value: `1` +
  description : Lower threshold of term counts to
be included. +
} +
`missing` : \{ +
  display name: Count Missing +
  type: `boolean` +
  default value: `false` +
  description : Optionally include a 'missing'
facet bucket for documents without the selected
field. +
} +
`offset` : \{ +
  display name: Offset +
  type: `integer` +
  description : Offset into list of resulting
facets. +
} +
`prefix` : \{ +
  display name: Prefix +
  type: `string` +
  description : Prefix of terms to facet on. +
} +
`sort` : \{ +
  display name: Sort +

```

Property	Description, Type
failQueryOnError Fail Query on Error	Fail the query request if an error occurs when processing this stage. type: <code>boolean</code> default value: 'false'
fallbackValue Fallback Value	Optional fallback value to use for the facet field if one cannot be determined by querying Solr. type: <code>string</code>
handler Request Handler <i>required</i>	Solr query request handler to send the auto-facet query request to. type: <code>string</code> default value: 'select'
headers Headers	type: <code>array of object</code> object attributes: \{ <code>key (required)</code> : \{ display name: Parameter Name type: <code>string</code> } <code>value</code> : \{ display name: Parameter Value type: <code>string</code> } \}
inherit Inherit Parameters From Parent?	Select box if a child category should inherit its parents' stored parameters. type: <code>boolean</code> default value: 'false'
method HTTP Method <i>required</i>	HTTP method used to send the auto-facet query request. type: <code>string</code> default value: 'GET' enum: \{ GET POST \}

Property	Description, Type
<p>params</p> <p>Additional Query Parameters</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none">key (<i>required</i>): \{<ul style="list-style-type: none">display name: Parameter Nametype: string } value: \{ <ul style="list-style-type: none">display name: Parameter Valuetype: string }

9.23. Query Fields Stage

This stage was renamed to Query Fields in Fusion 3.0. It was called the Search Fields stage in previous Fusion versions.

The Query Fields query-pipeline stage defines common Solr query parameters for the [edismax query parser](#). An alternative to this stage is the Additional Query Parameters stage.

9.23.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `search-fields`.

Configuration Properties

Property	Description, Type
minimumMatch Minimum Should Match	type: <code>string</code>
queryFields Search Fields	type: <code>array of object</code> object attributes: <code>{</code> <code>boost</code> : <code>{</code> display name: Field Boost type: <code>number</code> <code>}</code> <code>field</code> : <code>{</code> display name: Field Name type: <code>string</code> <code>}</code> <code>}</code>
returnFields Return Fields	type: <code>array of string</code> default value: <code>'*score'</code>
rows Number of Results	type: <code>integer</code> default value: <code>'10'</code>
sort Results Sort Order	type: <code>string</code>

Property	Description, Type
start Result Offset	type: <code>integer</code> default value: '0'

9.24. Recommend Items for Item Stage

The Recommend Items for Item query-pipeline stage uses signals about users' item choices to recommend similar items based on a specific item. For example, a Recommend Items for Item stage can recommend a list of similar books on a web page that displays information about a book.

Similarity can be based on different things, for example, click patterns, people who bought this also bought that, percentage match of document tags, etc.

This pipeline stage uses items-for-item recommendations that have been precomputed by an ALS (Alternating Least Squares) Spark recommender.

See Collaborative Filtering for more information.

9.24.1. Prerequisites

Enable recommendations:

Before creating a Recommend Items for Item stage, enable recommendations.

- **In the Fusion UI** – With Query Workbench open, click **Settings > Enable Recommendations**.
- **Using the REST API** – Use this command to enable recommendations:

```
`curl -u admin:<password> -X PUT http://%3Chostname%3E:%3Cport%3E/api/v1/collections/%3Ccollection-name%3E/features/recommendations -H 'content-type: application/json' -d '{"enabled":true}'`
```

Note: When you enable recommendations, Fusion creates a query pipeline that already contains this stage, and that is configured for boosting. The query pipeline is `<collection>_items_for_item_recommendations`.

9.24.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `item-recommendation`.

Configuration Properties

Property	Description, Type
boostFieldName	The field name to boost the values on.
Boost Field	type: <code>string</code> default value: <code>'id'</code>

Property	Description, Type
boostingMethod Boost Method <i>required</i>	The boost method to use. query-parser should be chosen if defType!=edismax for main query. type: <i>string</i> default value: 'query-param' enum: \{ query-param query-parser }
boostingParam Boost Param <i>required</i>	'Boost' multiplies scores by the boost values whereas 'bq' adds optional clauses to main query. type: <i>string</i> default value: 'boost' enum: \{ boost bq }
collection Recommendation Collection	If left blank, the default recommendation collection for the collection being queried will be used. type: <i>string</i>
foldInUpdates Estimate Recent Results	Update recommendations based on user activity that has happened since the last recommendation job run type: <i>boolean</i>
itemIdField Item ID Field	the name of the field in the recommendation collection where user ID is stored type: <i>string</i> default value: 'itemId'
itemIdParam Item ID Request Parameter	The name of the request parameter containing the item ID type: <i>string</i> default value: 'item_id'
modelID Model ID	type: <i>string</i> default value: '*'

Property	Description, Type
modelIdField Model ID Field	the name of the field in the recommendation collection where model ID is stored type: string default value: 'modelId'
numRecommendations Number of Recommendations	type: integer default value: '10'
otherItemIdField Recommended item ID Field	the name of the field in the recommendation collection where similar item recommendation is stored type: string default value: 'otherItemId'
resultsLocation Results Location	If As Response is chosen, then the result of the RPC call will be the one and only response. In all other cases, the stage will put the response from the REST/RPC call into the target location using the resultsKey. type: string default value: 'As Boosts' enum: \{ As Boosts As Response \}
scaleRange Scale Boosts	Scale the boost values to a [min,max] range type: object object attributes: \{ scaleMax : \{ display name: Maximum value of the scale range type: number } scaleMin : \{ display name: Minimum value of the scale range type: number } \}

Property	Description, Type
similarityScoreField Similarity Score Field	the name of the field in the recommendation collection where item similarity is stored type: <i>string</i> default value: 'sim'

9.25. Recommend Items for User Stage

The Recommend Items for User query-pipeline stage uses signals about item choices to recommend other similar items for a specific user. Personalization for the user can be based on the user's search history, browsing history, or purchase history, etc.

This pipeline stage uses items-for-user recommendations that have been precomputed by an ALS (Alternating Least Squares) Spark recommender.

See Collaborative Filtering for more information.

9.25.1. Prerequisites

Enable recommendations:

Before creating a Recommend Items for Item stage, enable recommendations.

- **In the Fusion UI** – With Query Workbench open, click **Settings** > **Enable Recommendations**.
- **Using the REST API** – Use this command to enable recommendations:

```
`curl -u admin:<password> -X PUT http://%3Chostname%3E:%3Cport%3E/api/v1/collections/%3Ccollection-name%3E/features/recommendations -H 'content-type: application/json' -d '{"enabled":true}'`
```

Note: When you enable recommendations, Fusion creates a query pipeline that already contains this stage, and that is configured for boosting. The query pipeline is `<collection>_items_for_user_recommendations`.

9.25.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `user-recommendation`.

Configuration Properties

Property	Description, Type
boostFieldName	The field name to boost the values on.
Boost Field	type: <code>string</code> default value: <code>'id'</code>

Property	Description, Type
boostingMethod Boost Method <i>required</i>	The boost method to use. query-parser should be chosen if defType!=edismax for main query. type: <i>string</i> default value: 'query-param' enum: \{ query-param query-parser }
boostingParam Boost Param <i>required</i>	'Boost' multiplies scores by the boost values whereas 'bq' adds optional clauses to main query. type: <i>string</i> default value: 'boost' enum: \{ boost bq }
collection Recommendation Collection	If left blank, the default recommendation collection for the collection being queried will be used. type: <i>string</i>
foldInUpdates Estimate Recent Results	Update recommendations based on user activity that has happened since the last recommendation job run type: <i>boolean</i>
itemIdField Item ID Field	the name of the field in the recommendation collection where item ID is stored type: <i>string</i> default value: 'itemId'
modelID Model ID	type: <i>string</i> default value: '*'
modelIdField Model ID Field	the name of the field in the recommendation collection where model ID is stored type: <i>string</i> default value: 'modelId'

Property	Description, Type
numRecommendations Number of Recommendations	type: <code>integer</code> default value: '10'
rawSignalsCollection Signals Collection	The collection to use to fetch recent user interactions, if 'Estimate Recent Results' is true. type: <code>string</code>
resultsLocation Results Location	If As Response is chosen, then the result of the RPC call will be the one and only response. In all other cases, the stage will put the response from the REST/RPC call into the target location using the resultsKey. type: <code>string</code> default value: 'As Boosts' enum: \{ As Boosts As Response \}
scaleRange Scale Boosts	Scale the boost values to a [min,max] range type: <code>object</code> object attributes: \{ <code>scaleMax</code> : \{ display name: Maximum value of the scale range type: <code>number</code> \} <code>scaleMin</code> : \{ display name: Minimum value of the scale range type: <code>number</code> \} \}
userIdField User ID Field	the name of the field in the recommendation collection where user ID is stored type: <code>string</code> default value: 'userId'
userIdParam User ID Request Parameter	The name of the request parameter containing the user ID type: <code>string</code> default value: 'user_id'

Property	Description, Type
weightField Weight Field	the name of the field in the recommendation collection where weight of the recommendation is stored type: <i>string</i> default value: 'weight'

9.26. Recommend More Like This Stage

This stage was renamed to Recommend More Like This in Fusion 3.1. It was introduced in Fusion 3.0, where it was called the Recommend Similar Items stage (and the More Like This stage).

The recommendations from Recommend More Like This are not based on collaborative filtering. An alternative query-pipeline stage, the Recommend Items for Item stage, uses collaborative filtering.

See Recommendations for more information.

9.26.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `more-like-this`.

Configuration Properties

Property	Description, Type
boost boost	Specifies if the query will be boosted by the interesting term relevant type: <code>boolean</code>
count Count	Specifies the number of similar documents to be returned for each result. type: <code>integer</code>
docId DocId Field name	Specifies the name of the id field we are finding more like this terms on type: <code>string</code> default value: 'id'
interestingTerms interesting terms	Controls how the More Like This component presents the interesting terms. Supports 3 settings, list lists the terms, none lists no terms and details lists the terms with the boosts type: <code>string</code> enum: <code>{ list none details }</code>

Property	Description, Type
matchInclude match include	Specifies whether the response should include the matched doc type: boolean
matchOffset match offset	Specifies an offset to the main query to find the doc on which the MoreLikeThis query should operate. By default it is 0 type: integer
maxdf maxdf	Specify the frequency at which words will be ignored which occur in more than this many docs type: integer default value: '10000'
maxntp maxntp	Sets the max number of tokens to parse in each example doc that is not stored with TV support type: integer
maxqt maxqt	Sets the max number of query terms that will be included in any generate query type: integer
maxwl maxwl	Sets the maximum word length type: integer
mindf mindf	Specify the frequency at which words will be ignored which occur in at least this many docs type: integer default value: '2'
mintf mintf	Specify the frequency below which terms will be ignored in the source doc type: integer

Property	Description, Type
minwl minwl	Sets the minimum word length for words to be recognized by the MoreLikeThis type: integer default value: '3'
moreLikeThisFields More Like This Fields	Specifies the name of the field you want to run the mlt on. NOTE: If you don't supply any fields we will default to using the body field. type: array of string
useQueryParser Use Query Parser <i>required</i>	Specifies whether to use the MLT Query Parser. Note, if you choose to use this you MUST specify a document id to run the MLT Query on and a Field to run the MLT with. type: boolean default value: 'true'

9.27. Query RPC Stage

The Query RPC query-pipeline stage was introduced in Fusion version 3.0.

9.27.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `query-rpc`.

Configuration Properties

Property	Description, Type
debug Add Debugging Information	Setting to true will add a number of properties to either the context (in the Query case) or the document (in the indexing case) type: <code>boolean</code> default value: 'false'
hasNoSideEffects Run in simulation mode	To run this stage in simulation mode, set to 'true'. type: <code>boolean</code> default value: 'false'

Property	Description, Type
<p>mappingRules</p> <p>Mapping of Returned Values (as XPath Expressions) to Context, Request or Response Values</p>	<p>type: array of object</p> <p>object attributes: \{</p> <ul style="list-style-type: none"> append : \{ <ul style="list-style-type: none"> display name: Append to Existing Values in Target Location. Only Valid when the Target Location is 'Request' type: boolean default value: 'false' path (<i>required</i>) : \{ <ul style="list-style-type: none"> display name: XPath Expression type: string target (<i>required</i>) : \{ <ul style="list-style-type: none"> display name: Target Key type: string targetLocation : \{ <ul style="list-style-type: none"> display name: Target Location type: string default value: 'Request' enum: \{ Request Response Context } <p style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <pre> } + 'xml' : \{ + display name: Add as an XML Fragment + type: 'boolean' + default value: 'false' + } + } </pre> </p>

Property	Description, Type
<p>params</p> <p>Call Parameters</p> <p><i>required</i></p>	<p>type: object</p> <p>object attributes: \{</p> <p>entity : \{</p> <p>display name: Request entity (as string)</p> <p>type: string</p> <p>\}</p> <p>headers : \{</p> <p>display name: Request protocol headers</p> <p>type: object</p> <p>\}</p> <p>method : \{</p> <p>display name: Call method</p> <p>type: string</p> <p>description : One of GET, POST, PUT, or DELETE</p> <p>enum: \{ get put post delete \}</p> <pre> } + 'queryParams' : \{ + display name: Query parameters + type: 'object' + } + 'uri' : \{ + display name: Endpoint URI + type: 'string' + } + } </pre>
<p>resultsKey</p> <p>Results Key</p> <p><i>required</i></p>	<p>The name of the key to store the results object under. See the documentation for the type of objects stored.</p> <p>type: string</p> <p>default value: 'queryRPC'</p>
<p>resultsLocation</p> <p>Results Location</p>	<p>If As Response is chosen, then the result of the RPC call will be the one and only response. In all other cases, the stage will put the response from the REST/RPC call into the target location using the resultsKey.</p> <p>type: string</p> <p>default value: 'Request'</p> <p>enum: \{ Request Response Context As Response \}</p>

Property	Description, Type
useIncomingRequestEntity Use the Incoming Request Entity	If an input entity is POSTed or PUT from the client, we can pass that along to the RPC target. type: <code>boolean</code> default value: 'false'

9.28. Retrieve Stored Parameters Stage

This stage was renamed to Retrieve Stored Parameters in Fusion 3.0. It was called the Stored Parameters stage in prior Fusion versions.

The Retrieve Stored Parameters stage is used to add parameters to the downstream Solr query stage in the pipeline. The primary use case for this query stage is in eCommerce applications to organize and enhance search query results, where the parameters added to the Solr query are used for faceting. This stage is used in query pipelines for applications where the documents in the collection are to be classified according to a known category hierarchy that is stored in Fusion as an auxiliary "stored parameters" collection.

The information in a taxonomy is meta-information about the categories used to classify a set of things. For an eCommerce site, the set of things are items in the product catalog, which are stored in a Fusion collection, referred to as the primary collection. The taxonomy itself is specified as a set of categories which are stored in an auxiliary collection and naming conventions relate the primary and auxiliary collections so that for a primary collection named "COLL", the auxiliary taxonomy is stored in a collection named "COLL_stored_parameters".

9.28.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `stored-parameters`.

Configuration Properties

Property	Description, Type
inherit Inherit Parameters From Parent?	Select box if a child category should inherit its parents' stored parameters. type: <code>boolean</code> default value: 'false'
key Key	The parameter name from the context/query whose value will be used to query the stored parameters collection. type: <code>string</code>

9.29. Return Query Parameters Stage

This stage was renamed to Return Query Parameters in Fusion 3.0. In prior Fusion versions, it was called the Return Query Params stage.

The Return Query Parameters query-pipeline stage is used to return query parameters as a pipeline response. The stage takes no additional properties beyond id, type, label, skip, and condition.

9.29.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `return-queryparams`.

Configuration Properties

Property	Description, Type
----------	-------------------

9.30. Rollup Aggregation Stage

The Rollup Aggregation query-pipeline stage (previously called the Rollup Aggregator stage) is used to rollup Solr results from a context variable. If using the REST API, this stage type is named "rollup-rec-aggr".

This stage reads the Solr results (SolrResponse.class) from the context and rolls up over a single field product a list of unique IDs and also aggregates the weights (any numeric filed in Solr) for those IDs using any of the statistical aggregation functions available. The result from aggregation is saved back in the context and can later be used in the Parameterized Boosting stage.

9.30.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `rollup-rec-aggr`.

Configuration Properties

Property	Description, Type
excludeResultsKey Results to Exclude	The key containing a set of results to exclude from this rollup. type: <code>string</code>
key Key <i>required</i>	The key name to use from context to read Solr results. type: <code>string</code>
maxRows Max rows	The maximum number of results to return type: <code>integer</code> default value: '10'
resultKey Result Key <i>required</i>	The key name to which the results should be saved. type: <code>string</code>

Property	Description, Type
rollupField Rollup Field <i>required</i>	The field to rollup on. type: <code>string</code>
sort Sort results	If enabled, the output is sorted based on weight field if it is not null type: <code>boolean</code> default value: 'true'
weightField Weight Field	The numerical field to consider as weight. type: <code>string</code>
weightFunction Weight Arithmetic Function	The arithmetic function to use for weight fields on documents with same rollup field. type: <code>string</code> default value: 'sum' enum: \{ sum mean max min stddev variance geoMean sumOfSquares sumOfLogs \}

9.31. Security Trimming Stage

The Security Trimming query-pipeline stage restricts query results according to the user ID. While indexing the content, the Fusion connectors service stores security ACL metadata associated with the crawled items and indexes them as fields. The Security Trimming stage matches this information against the ID of the user running the search query.

9.31.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `security-trimming`.

Configuration Properties

Property	Description, Type
datasources Restrict filter to Datasource(s)	A list of Fusion datasources to which security-trimming should be restricted, allowing content from other datasources to pass through un-filtered; if empty, all matching content is subject to filtering. type: <code>array of string</code>
overrideUserIdentityHandling Override Default User Identity Handling?	Default handling first attempts to take the user identity from a 'fusion-user-id' http-header, which is the logged-in user ID from the Fusion proxy service. If that value is empty, a 'username' query parameter is tried instead. When this DataSource property is enabled, the specified source and key properties are used explicitly, without any fallback behavior. type: <code>boolean</code> default value: 'false'
userIdentityKey User ID key	e.g. username, userID, etc. type: <code>string</code> default value: 'username'

Property	Description, Type
userIdentitySource User ID source	Specify whether the value comes from an http header or query parameter. type: <i>string</i> default value: 'query_param' enum: \{ query_param header }

9.32. Send PagerDuty Message Stage

This stage sends a [PagerDuty](#) Message from Fusion, for alerting, monitoring, and more, using Fusion's Messaging Services.

Read more about PagerDuty integration in Fusion on our [blog](#).

9.32.1. Enabling PagerDuty Messaging

Before you can use the PagerDuty pipeline stage, you must enable PagerDuty messaging in Fusion:

1. Click **Applications > System > Messaging Services**.
2. Select **PagerDuty Message Service** from the drop-down menu.
3. Enter the following information:
 - [PagerDuty service key](#)
 - PagerDuty Service API URL; this should be https://events.pagerduty.com/generic/2010-04-15/create_event.json.
4. Click **Save message service**.

9.32.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `pagerduty-message`.

Configuration Properties

Property	Description, Type
client Client	The name of the monitoring client that is triggering this event. type: <code>string</code> default value: 'Fusion' minLength: 1
clientURL Client URL	The URL of the monitoring client that is triggering this event. type: <code>string</code> default value: 'fusion-monitoring.yourdomain.com' minLength: 1

Property	Description, Type
<p>errorKey</p> <p>Message Response Failure Key</p>	<p>The name of the key to store a boolean if sending a message failed. If set, you can check the MessageResponse errorCode and other attributes for the reason.</p> <p>type: string</p> <p>default value: 'messageResponseFailure'</p>
<p>eventType</p> <p>Event Type</p> <p><i>required</i></p>	<p>Select the Pager Duty Event Type.</p> <p>type: string</p> <p>default value: 'trigger'</p> <p>enum: \{ trigger acknowledge resolve }</p>
<p>incidentContextImages</p> <p>Incident Context Images</p>	<p>type: array of object</p> <p>object attributes: \{</p> <p>alt : \{</p> <p>display name: Alternate Text</p> <p>type: string</p> <p>description : HTML 'alt' tag for the image.</p> <p> }</p> <p>href (required) : \{</p> <p>display name: Target Link</p> <p>type: string</p> <p>description : URL to open when clicked on the image.</p> <p> }</p> <p>src (required) : \{</p> <p>display name: Source</p> <p>type: string</p> <p>description : HTML 'src' tag for the image. Should be always secure connection (https://).</p> <p> }</p> <p> }</p>

Property	Description, Type
<p>incidentContextLinks</p> <p>Incident Context Links</p>	<p>type: array of object</p> <p>object attributes: \{</p> <p>href (<i>required</i>): \{</p> <p>display name: Target Link</p> <p>type: string</p> <p>description : URL to open when clicked on the link.</p> <p>\}</p> <p>text : \{</p> <p>display name: Text</p> <p>type: string</p> <p>description : Arbitrary text explaining the URL.</p> <p>\}</p> <p>\}</p>
<p>incidentDescription</p> <p>Description</p> <p><i>required</i></p>	<p>A short description of the problem that led to this trigger. This field (or a truncated version) will be used when generating phone calls, SMS messages and alert emails. It will also appear on the incidents tables in the PagerDuty UI. The maximum length is 1024 characters.</p> <p>type: string</p> <p>default value: 'Sample Description'</p> <p>maxLength: 1024</p> <p>minLength: 1</p>
<p>incidentDetails</p> <p>Incident Details</p>	<p>type: array of object</p> <p>object attributes: \{</p> <p>name (<i>required</i>): \{</p> <p>display name: Name</p> <p>type: string</p> <p>\}</p> <p>value (<i>required</i>): \{</p> <p>display name: Value</p> <p>type: string</p> <p>\}</p> <p>\}</p>

Property	Description, Type
<p>incidentKey</p> <p>Incident Key</p>	<p>Identifies the incident to which this trigger event should be applied. If there's no open (i.e. unresolved) incident with this key, a new one will be created. If there's already an open incident with a matching key, this event will be appended to that incident's log.</p> <p>type: <code>string</code></p> <p>default value: <code>'Incident'</code></p> <p>minLength: 1</p>
<p>responseKey</p> <p>Message Response Context Key</p>	<p>The name of the key to store the MessageResponse under in the Pipeline Context.</p> <p>type: <code>string</code></p> <p>default value: <code>'messageResponse'</code></p>
<p>storeInContext</p> <p>Add to Pipeline Context</p>	<p>Put the generated Message later in the pipeline.</p> <p>type: <code>boolean</code></p> <p>default value: <code>'false'</code></p>

9.33. Send SMTP Email Stage

This stage was renamed to Send SMTP Email in Fusion 3.0; it was called the Send email (via SMTP) stage in previous Fusion versions.

This stage sends an SMTP message from Fusion, for alerting, reporting, and more, using Fusion’s Messaging Services.

9.33.1. Enabling Email Messaging

Before you can use the Email pipeline stage, you must enable Email messaging in Fusion:

1. Click **Applications > System > Messaging Services**.
2. Select **SMTP Message Service** from the drop-down menu.
3. Verify that the default settings are sufficient.
4. Click **Save message service**.

9.33.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion’s REST-API, the ID of this stage is: `smtp-message`.

Configuration Properties

Property	Description, Type
errorKey Message Response Failure Key	The name of the key to store a boolean if sending a message failed. If set, you can check the MessageResponse errorCode and other attributes for the reason. type: <code>string</code> default value: <code>'messageResponseFailure'</code>
from From <i>required</i>	Who the message is from. May be a string template similar to the body and subject. type: <code>string</code>

Property	Description, Type
messageBodyTemplate Message Body Template <i>required</i>	A Message Template that is used to create the message body to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: <code>string</code> default value: <code>'`Processing Document`'</code>
messageSubjectTemplate Message Subject Template <i>required</i>	A Message Template that is used to create the message subject to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: <code>string</code> default value: <code>'Hello'</code>
responseKey Message Response Context Key	The name of the key to store the MessageResponse under in the Pipeline Context. type: <code>string</code> default value: <code>'messageResponse'</code>
smtpPassword SMTP Password <i>required</i>	The SMTP password for the user credentials. type: <code>string</code>
smtpUser SMTP Username <i>required</i>	The SMTP user to send the message from. type: <code>string</code>
storeInContext Add to Pipeline Context	Put the generated Message later in the pipeline. type: <code>boolean</code> default value: <code>'false'</code>

Property	Description, Type
to To <i>required</i>	Who to send the message to. May be a string template similar to the body and subject. type: <code>array of string</code>

9.34. Send Slack Message Stage

This stage sends a [Slack](#) message from Fusion, for alerting, reporting, and more, using Fusion’s Messaging Services.

9.34.1. Enabling Slack Messaging

Before you can use the Slack pipeline stage, you must enable Slack messaging in Fusion:

1. Click **Applications > System > Messaging Services**.
2. Select **Slack Message Service** from the drop-down menu.
3. Enter the following information:
 - [Slack auth token](#)
 - Message template

The default is `<subject> : <body>`, which are configured with `messageSubjectTemplate` and `messageBodyTemplate` below. See Messaging Services Templates for details on the template language.

- Optionally, you can configure a proxy or the error reporting channel name.
4. Click **Save message service**.

9.34.2. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion’s REST-API, the ID of this stage is: `slack-message`.

Configuration Properties

Property	Description, Type
errorKey Message Response Failure Key	The name of the key to store a boolean if sending a message failed. If set, you can check the MessageResponse errorCode and other attributes for the reason. type: <code>string</code> default value: <code>'messageResponseFailure'</code>
from From <i>required</i>	Who the message is from. May be a string template similar to the body and subject. type: <code>string</code>

Property	Description, Type
messageBodyTemplate Message Body Template <i>required</i>	A Message Template that is used to create the message body to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: <code>string</code> default value: <code>'`Processing Document`'</code>
messageSubjectTemplate Message Subject Template <i>required</i>	A Message Template that is used to create the message subject to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: <code>string</code> default value: <code>'Hello'</code>
responseKey Message Response Context Key	The name of the key to store the MessageResponse under in the Pipeline Context. type: <code>string</code> default value: <code>'messageResponse'</code>
storeInContext Add to Pipeline Context	Put the generated Message later in the pipeline. type: <code>boolean</code> default value: <code>'false'</code>
to To <i>required</i>	Who to send the message to. May be a string template similar to the body and subject. type: <code>array of string</code>

9.35. Solr Query Stage

This stage was renamed to Solr Query in Fusion 3.1. It was called the Query Solr stage in prior Fusion versions.

The Solr Query query-pipeline stage transforms the Fusion query pipeline Request object into a Solr query and sends it to Solr.

9.35.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `solr-query`.

Configuration Properties

Property	Description, Type
allowFederatedSearch Allow Federated Search	Allow Solr 'collection' and 'shards' parameters. type: <code>boolean</code> default value: 'false'
allowedRequestHandlers Configure Request Handlers Allowed for Queries	type: <code>array</code> of <code>string</code> minimum number of items (<code>minItems</code>): 0
httpMethod HTTP Method	HTTP method for querying Solr. type: <code>string</code> default value: 'POST' enum: <code>\{ POST GET \}</code>

9.36. Write Log Message Stage

This stage was renamed to Write Log Message in Fusion 3.0. In prior Fusion version, it was called the Log a Message stage.

The Write Log Message query-pipeline stage is an extension of Fusion's Logging Query Stage, which logs any message sent to the configured logging system using the Messaging Services.

9.36.1. Configuration

Tip	When entering configuration values in the UI, use <i>unescaped</i> characters, such as <code>\t</code> for the tab character. When entering configuration values in the API, use <i>escaped</i> characters, such as <code>\\t</code> for the tab character.
-----	---

When using Fusion's REST-API, the ID of this stage is: `logging-message`.

Configuration Properties

Property	Description, Type
errorKey Message Response Failure Key	The name of the key to store a boolean if sending a message failed. If set, you can check the MessageResponse errorCode and other attributes for the reason. type: <code>string</code> default value: 'messageResponseFailure'
from From <i>required</i>	Who the message is from. May be a string template similar to the body and subject. type: <code>string</code>
logLevel Log Level	The Log Level. May be: debug, info, warn, error type: <code>string</code> enum: <code>\{ debug info warn error }</code>
messageBodyTemplate Message Body Template <i>required</i>	A Message Template that is used to create the message body to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: <code>string</code> default value: <code>'`Processing Document`'</code>

Property	Description, Type
messageSubjectTemplate Message Subject Template <i>required</i>	A Message Template that is used to create the message subject to send. See https://github.com/antlr/stringtemplate4/blob/master/doc/index.md for details on the template language. type: string default value: 'Hello'
responseKey Message Response Context Key	The name of the key to store the MessageResponse under in the Pipeline Context. type: string default value: 'messageResponse'
storeInContext Add to Pipeline Context	Put the generated Message later in the pipeline. type: boolean default value: 'false'
to To <i>required</i>	Who to send the message to. May be a string template similar to the body and subject. type: array of string

Chapter 10. REST API Reference

See Index of API Services Endpoints for a list of all endpoints in Fusion’s API service component.

Fusion API service are designed to be accessed via Fusion’s authentication proxy module which is part of the Fusion UI service (default port 8764). All applications should use this method to access the API service. However, for development purposes it is possible to access the API service directly:

Development	Production
<pre><code>http://&lt;fusion-host&gt;;8765/api/v1/&lt;endpoint&gt;</code></pre> <p>Direct access to the API service, bypassing the authentication proxy. <i>Access to this port should be restricted.</i></p>	<pre><code>http://&lt;fusion-host&gt;;8764/api/apollo/&lt;endpoint&gt;</code></pre> <p>Access via the authentication proxy, using either a username/password pair or a session cookie with the user ID.</p>
Important	In a production environment, do not expose port 8765 to users. Using your firewall software or the Jetty configuration of the API server, make it accessible only to the auth proxy service and the connectors service.

10.1. Listing all Fusion component services

The Fusion `introspect` endpoint lists basic information about endpoints and parameters for all Fusion endpoints, including the Connectors services endpoints:

```
curl -u user:pass http://localhost:8764/api/apollo/introspect
```

10.2. Fusion Components REST API Reference Pages

- Blob Store API
- Catalog API
- Collection Features API
- Collections API
- Configurations API
- Connectors APIs
- History API
- Index Pipelines API
- Index Profiles API
- Index Stages API

- Messaging API
- Nodes API
- Objects API
- Parsers API
- Query Pipelines API
- Query Profiles API
- Query Stages API
- Realms API
- Recommendations API
- Reporting API
- Roles API
- Scheduler API
- Search Cluster API
- Sessions API
- Signals API
- Signals Aggregator API
- Solr API
- SolrAdmin API
- Synonyms Editor API
- System API
- Usage API
- User API
- ZooKeeper Import/Export API

10.3. Authentication and Authorization APIs

The Fusion access control component handles user authentication and authorization. It runs in the same process as the Fusion UI and provides the following sets of endpoints:

- User API: create, update, delete and list user records in Fusion.
- Roles API: create, update, delete and list roles which grant users different levels of permissions in the system.
- Realms API: create, update, delete and list realms.
- Sessions API: create an authenticated session and session cookie.

10.3.1. Realms API

Realms are used to authenticate users across several different user access control systems.

Fusion supports these types of security realms.

Create, Update, Delete or List Realms

The path for this request is:

`/api/realm-configs/<id>`

where `<id>` is the ID of a realm. The ID is optional for a GET request and omitted from a POST request.

A GET request returns the configured realms. If ID is omitted, all realms will be returned.

A POST request creates a new realm. If the request is successful, a new ID will be generated.

A PUT request updates a realm.

A DELETE request removes the realm.

Input

Parameter	Description
name <i>Required</i>	The name of the realm. This name will appear on the login screen of the UI, and will appear in user records to identify the realm they belong to.
enabled <i>Required</i>	If true , the realm is available for users to use with system authentication.
realmType <i>Required</i>	String value for realm type. Supported realm types are <code>native</code> , <code>ldap</code> , <code>kerberos</code> , <code>saml</code> , and <code>trusted-http</code> .

Native realms have users whose usernames and passwords are stored in the Fusion database. Authenticating users with an LDAP system creates a user record in Fusion, which includes a property for the realm the user belongs to. This Fusion user record is used by administrators to grant users access permissions for the UI or REST API services. LDAP realms connect to an LDAP server to verify the user's ID and password.

Configuration for an LDAP security realm requires the following additional properties:

Parameter	Description
host	The hostname of the LDAP server.
port	The port to use when connecting to the LDAP server.
ssl	If true , SSL will be used when connecting to the LDAP server.
bindDN	A string consisting of the LDAP server DN (Distinguished Name) and a single pair of curly braces ({}), which is a placeholder for the username.

Output

When creating a new realm, the output will include the properties for the realm just created, or an error to indicate a problem with the entry.

For a GET request, the output will include all defined properties of the realm.

For a DELETE or a PUT request, no output will be returned.

Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Get details of the default 'native' realm:

REQUEST

```
curl -u user:pass http://localhost:8764/api/realm-configs/86df9b5b-4a1c-4b0b-bc10-25aee55fef63
```

RESPONSE

```
{
  "enabled": true,
  "id": "86df9b5b-4a1c-4b0b-bc10-25aee55fef63",
  "name": "native",
  "realmType": "native"
}
```

Create a realm to support LDAP authentication:

REQUEST

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"realmType":"ldap", "name":"dev-ldap", "enabled":true, "config":{"host":"localhost", "port":10636, "ssl":true, "bindDn":"uid={},ou=users,dc=security,dc=example,dc=com"} }' http://localhost:8764/api/realm-configs
```

RESPONSE

```
{
  "realmType":"ldap",
  "name":"dev-ldap",
  "enabled":true,
  "config":{
    "bindDn":"uid={},ou=users,dc=security,dc=example,dc=com",
    "ssl":true,
    "port":10636,
    "host":"localhost"
  }
}
```

10.3.2. Roles API

Roles are groups of permissions that allow access to the UI and the REST APIs. See Roles for details.

Security Realms can be configured to use LDAP group membership to assign Roles to users. See the LDAP configuration instructions for details.

Create, Update or Delete Roles

The endpoint for this request can take the role ID as a request parameter:

`/api/roles/<id>`

The role ID string is generated by Fusion when the role is created.

A GET request returns the configured roles for a specific ID. If the ID is omitted from the path, all roles will be returned.

A POST request creates a new role. When creating a new role, the request path is `/api/roles`. If the role is created, the request returns the role ID.

A PUT request updates an existing role.

A DELETE request will remove the role configuration.

Role Specification

To create or update a Role via a POST or PUT request, the request body is a JSON object with the following attributes:

Property	Description
name <i>Required</i>	A string containing the role name.
desc <i>Optional</i>	A string containing a brief text description, for display on the Access Control "ROLES" panel.
permissions <i>Optional</i>	A list of permissions, specified in JSON notation. See section Permissions for details.
uiPermissions <i>Optional</i>	A list of names of UI components.

The following example describes a role with permissions to access Fusion Dashboards for collection "mdb1":

```
{
  "name": "view-dashboard-mdb1",
  "desc": "can access/use analytics dashboard \"mdb1\" but not allowed to change dashboard controls.",
  "permissions": [
    {"methods": ["GET"], "path": "/solr/system_banana/*"},
    {"methods": ["GET"], "path": "/solr/{id}/*", "params": {"id": ["mdb1"]}},
    {"methods": ["GET"], "path": "/solr/{id}/admin/luke", "params": {"id": ["mdb1"]}},
    {"methods": ["GET"], "path": "/collections/system_banana"}
  ],
  "uiPermissions": [
    "dashboards",
    "fields"
  ]
}
```

Examples

Get the details for the role with id '3416c03a-31df-4103-b446-358f6790af3e':

REQUEST

```
curl -u user:pass http://localhost:8764/api/roles/3416c03a-31df-4103-b446-358f6790af3e
```

RESPONSE

```
{
  "id": "3416c03a-31df-4103-b446-358f6790af3e",
  "name": "search",
  "createdAt": "2016-03-09T20:01:48Z",
  "permissions": [
    {"methods": ["GET"], "path": "/query-pipelines/*/collections/*/select"},
    {"methods": ["GET"], "path": "/query-pipelines"},
    {"methods": ["GET"], "path": "/solr/*/schema"},
    {"methods": ["GET"], "path": "/prefs/apps/search/*"},
    {"methods": ["GET"], "path": "/collections/**"},
    {"methods": ["GET"], "path": "/solr/*/admin/luke"}
  ],
  "uiPermissions": [
    "search",
    "collections"
  ],
  "desc": "Provides read-only/required permissions for the Fusion Search UI."
}
```

10.3.3. Sessions API

The session API is used to create sessions using defined realms, such as LDAP.

A session can be saved into a cookies file that can be re-used for subsequent requests. Sessions time out after 10 minutes of no activity, or after 8 hours.

Create a Session

The path for this request is:

```
/api/session?realmName=<realmName>
```

where the query parameter *realmName* takes as its value the name of a realm to authenticate against.

Input

Parameter	Description
username <i>Required</i>	The username to use in authentication.
password <i>Required</i>	The password to use in authentication.

Output

The output will include a cookie ID in the HTTP response header. This can be saved to a file and re-used with subsequent REST API requests.

Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Create a session against an LDAP server and store it in a file named 'cookies':

REQUEST

```
curl -c cookies -i -H "content-type:application/json" -X POST -d '{"username":"myUser",  
"password":"myPassword"}' http://localhost:8764/api/session?realmName=myLDAP
```

RESPONSE

```
HTTP/1.1 201 Created  
Set-Cookie: id=840a33d4-b650-49f2-87a4-85412e99b37c;HttpOnly;Path=/api  
Content-Length: 0  
Server: Jetty(9.1.4.v20140401)
```

Note	In this case, we got a response because we set <code>curl</code> to include the HTTP in the output. Otherwise, we would not know for sure the session was created.
------	--

Use the cookie in another cURL request to see all collections:

```
curl -b cookies http://localhost:8764/api/apollo/collections
```


10.3.4. User API

The User API allows you to create, update, and remove user accounts. This API should only be called to manage users in the native security realm. Users from other security realms are managed directly by Fusion’s auth proxy.

Create, Update, Delete or List Users

The path for this request is:

`/api/users/<id>`

where `<id>` is the user ID.

A GET request lists information about the user. The ID can be omitted in a GET request to get all users.

A POST request creates a new user, while a PUT updates a user record.

DELETE will remove the user.

Input

Parameter	Description
username <i>Required</i>	The username. This is distinct from their ID, which is assigned by the system as a unique identifier.
password <i>Required</i>	The user’s password. Required when creating a new user. The user’s password is not returned in the output of any request.
passwordConfirm <i>Required</i>	When creating a user or updating a user’s password, you must confirm the defined password.
realmName <i>Required</i>	The realm the user belongs to, which defines how they authenticate against the system.
permissions <i>Optional</i>	The permissions that have been defined for this user that are not inherited from their assigned role.
inheritedPermissions <i>Optional</i>	The user’s specific permissions that are inherited from their role assignment.
roleNames <i>Optional</i>	The list of user’s roles, which define some or all of the permissions they have.

Output

When creating a user with a POST request or listing users with GET, the user properties will be returned.

When updating or removing a user with a PUT or DELETE, no output will be returned.

Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Get all the configured users of the system:

REQUEST

```
curl -u user:pass http://localhost:8764/api/users
```

RESPONSE

```
[
  {
    "realmName": "native",
    "username": "admin",
    "id": "2856ba33-80bd-400d-99dc-3d181bc68d9a",
    "roleNames": ["admin"],
    "permissions": [],
    "createdAt": "2015-07-01T03:18:06Z",
    {
      "realmName": "native",
      "username": "collection-admin",
      "id": "9780a33c-c49d-48e3-a869-bd65951aea8f",
      "roleNames": ["ui-user", "collection-admin"],
      "permissions": [],
      "createdAt": "2015-07-01T03:18:06Z"
    }
  }
]
```

Add a new user named 'guest':

REQUEST

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"username": "guest",
"password": "password456", "passwordConfirm": "password456", "realmName": "native"}'
http://localhost:8764/api/users
```

RESPONSE

```
{
  "realmName": "native",
  "username": "guest",
  "id": "2f5b52a7-550d-407d-b592-32ab42afe3ca",
  "roleNames": [],
  "permissions": [],
  "createdAt": "2015-08-06T11:42:15Z"
}
```

Update a user to include the role named "admin":

REQUEST

```
curl -u user:pass -X PUT -H 'Content-type: application/json' -d '{"name": "joe.smith", "realmName": "myLDAP",
"roleNames": ["admin"]}' http://localhost:8764/api/users/aefa7ffc-23f1-45ac-b326-f7bb007d3b9d
```

RESPONSE

None.

10.4. Blob Store API

The Blob Store REST API allows storing binary objects in Solr. The primary use case for this is to store entity extraction models, lookup lists or exclusion lists for use in index pipelines. This may include the entity extraction models and lookup lists included with Fusion in the `fusion/3.1.x/data/nlp` directory, or files that you have created on your own.

Blobs uploaded to Solr with this REST API are stored in the 'system_blobs' collection.

10.4.1. Blob Types

A `resourceType` query parameter can be used to specify the a blob type. For example, specify `plugin:connector` when uploading a connector, like this:

```
curl -H 'content-type:application/zip' -X PUT
'localhost:8765/api/v1/blobs/myplugin?resourceType=plugin:connector' --data-binary @myplugin.zip
```

The complete list of valid values for `resourceType` is below:

Type	Description
<code>catalog</code>	An analytics catalog
<code>driver:jdbc</code>	A JDBC driver
<code>plugin:connector</code>	A connector plugin
<code>model:ml-model</code>	A machine learning model
<code>model:open-nlp</code>	An OpenNLP model
<code>file-upload</code>	Any uploaded file, such as from the Quickstart or the Index Workbench.
<code>banana</code>	A Banana dashboard
<code>other</code>	A blob of unknown type If no <code>resourceType</code> is specified on upload, "other" is assigned by default.

10.4.2. Examples

Upload a file to the blob store:

REQUEST

```
curl -u user:pass -X PUT --data-binary @airports.lst -H 'Content-type: text/plain'
http://localhost:8764/api/apollo/blobs/airports.lst
```

RESPONSE

```
{
  "name" : "airports.lst",
  "contentType" : "text/plain",
  "size" : 66,
  "modifiedTime" : "2014-12-03T22:26:16.436Z",
  "version" : 0,
  "md5" : "fbc581898cb426f6bdcabc3f52253594"
}
```

Upload an OpenNLP sentence model binary file to the blob store:

REQUEST

```
curl -u user:pass -X PUT --data-binary @data/nlp/models/en-sent.bin -H 'Content-type: application/octet-stream' http://localhost:8764/api/apollo/blobs/sentenceModel.bin
```

Note

In this example that we have changed the name of the blob during upload by giving it a different ID. The file is named 'en-sent.bin' but we have defined the ID of this to 'sentenceModel.bin'. When we use this blob in an index pipeline, we would refer to it by the ID we have given it.

Get the manifest for a sentence OpenNLP model we've previously saved in the blob store:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/blobs/sentenceModel.bin/manifest
```

RESPONSE

```
{
  "name" : "sentenceModel.bin",
  "contentType" : "application/octet-stream",
  "size" : 98533,
  "modifiedTime" : "2014-09-08T18:50:07.559Z",
  "version" : 1478704189996531712,
  "md5" : "3822c5f82cb4ba139284631d2f6b7fde"
}
```

Upload a JDBC driver, using slashes in the blob name:

REQUEST

```
curl -u user:pass -X PUT --data-binary @mysql-connector-java-5.1.42-bin.jar -H 'Content-length: 996444' -H 'Content-Type: application/zip' http://localhost:8764/api/apollo/blobs/good/to/go/mysql-connector-java-5.1.42-bin.jar?resourceType=driver:jdbc
```

RESPONSE

```
{
  "name" : "good/to/go/mysql-connector-java-5.1.42-bin.jar",
  "contentType" : "application/zip",
  "size" : 996444,
  "modifiedTime" : "2017-04-04T15:58:32.856Z",
  "version" : 0,
  "md5" : "b1946ac92492d2347c6235b4d2611184",
  "metadata" : {
    "subtype" : "driver:jdbc",
    "resourceType" : "driver:jdbc"
  }
}
```

Get the JDBC driver that was uploaded:

REQUEST

```
curl -u user:pass -H "Accept: application/zip" http://localhost:8764/api/apollo/blobs/jtlds-1.3.1-
src.zip?resourceType=driver:jdbc -o jtlds-1.3.1-src.zip
```

RESPONSE

```
[ {
  "name" : "good/to/go/sentenceModel.bin",
  "contentType" : "application/octet-stream",
  "size" : 6,
  "modifiedTime" : "2017-04-04T06:21:53.465Z",
  "version" : 1563727666574524416,
  "md5" : "b1946ac92492d2347c6235b4d2611184",
  "metadata" : {
    "subtype" : "driver:jdbc",
    "resourceType" : "driver:jdbc"
  }
} ]
```

10.5. Catalog API

The Fusion Catalog is a collection of one or more analytics projects, and each project is a collection of data assets, such as tables or relations. Fusion comes with a built-in project called "fusion".

The Fusion Catalog API provides access to assets by data analysis applications that can perform SQL or Solr queries. It includes endpoints for finding, retrieving, and manipulating projects and assets using basic keyword and metadata-driven search.

By default, non-admin Fusion users do not have access to Catalog objects. However, the Catalog API itself does not enforce any permissions, so a user who bypasses the auth proxy has full access to all projects and assets. An admin can grant permissions to Catalog endpoints for users; see Access Control.

10.5.1. Intra-shard splits

If your Spark cluster has more available executor slots than the number of shards, then you can increase parallelism when reading from Solr by splitting each shard into sub-ranges using a split field. The sub range splitting enables faster fetching from Solr by increasing the number of tasks in Solr. This should only be used if there are enough computing resources in the Spark cluster.

Shard splitting is enabled by default, with two sub-ranges per shard. See Configuration options below for shard splitting parameters.

10.5.2. Body attributes

For PUT and POST requests, these are valid JSON body attributes:

Name	Type	Description
<code>projectId</code>	String	The project name
<code>name</code>	String	The asset name
<code>assetType</code>	DataAssetType	One of: + * project * table * relation * field * udf * metric
<code>description</code>	String	A string describing this asset
<code>sourceUri</code>	String	A URI to the data source
<code>owner</code>	String	The user that owns the asset
<code>ownerEmail</code>	String	The owner's email address
<code>tags</code>	Set<String>	A set of arbitrary category strings
<code>format</code>	String	The format of the underlying data source
<code>options</code>	List<String>	A list of options for the underlying data source. See Configuration options below for valid options.
<code>filters</code>	List<String>	A set of Solr query parameters to filter the request
<code>sql</code>	String	A SQL statement to execute

Name	Type	Description
cacheOnLoad	boolean	'True' to cache the dataset in Spark on catalog project initialization
dependsOn	List<String>	A list of other assets to load before initializing this data asset
createdOn	Date	The asset's creation date, in ISO-8601 format; otherwise the current timestamp is used

10.5.3. Configuration options

Name	Description	Default
collection	The Solr collection name.	None
zkhost	A ZooKeeper connect string is the list of all servers and ports for the current ZooKeeper cluster. For example, if running a single-node Fusion developer deployment with embedded ZooKeeper, the connect string is <code>localhost:9983/lwfusion/3.1.0/solr</code> . If you have an external 3-node ZooKeeper cluster running on servers "zk1.acme.com", "zk2.acme.com", "zk3.acme.com", all listening on port 2181, then the connect string is <code>zk1.acme.com:2181,zk2.acme.com:2181,zk3.acme.com:2181</code>	The <code>connectString</code> of the default search cluster
query	A Solr query that limits the rows to load into Spark. For example, to only load documents that mention "solr": <pre>options("query","body_t:solr")</pre>	:

Name	Description	Default
<p><code>fields</code></p>	<p>A subset of fields to retrieve for each document in the results, such as:</p> <pre data-bbox="591 264 1029 352">options("fields", "id,author_s,fa vorited_b,...")</pre> <p>You can also specify an alias for a field using Solr's field alias syntax, such as <code>author:author_s</code>. If you want to invoke a function query, such as <code>rord()</code>, then you'll need to provide an alias, such as <code>ord_user:ord(user_id)</code>. If the return type of the function query is something other than <code>int</code> or <code>long</code>, then you'll need to specify the return type after the function query, such as:</p> <pre data-bbox="591 814 1029 898">foo:div(sum(x,100),max(y,1)):dou ble</pre>	<pre>===== ===== ===== ===== ===== =====</pre>
<p>Tip</p>	<p>If you request Solr function queries, then the library must use the <code>/select</code> Solr handler to make the request as exporting function queries through <code>/export</code> is not supported by Solr.</p>	<pre>===== ===== ===== ===== =====</pre>

Name	Description	Default
By default, all stored fields for each document are pulled back from Solr.	<code>rows</code>	<p>The number of rows to retrieve from Solr per request; do not confuse this with <code>max_rows</code> (see below). This is not the maximum number of rows to read from Solr. All matching rows on the backend are read. The <code>rows</code> parameter is the page size.</p> <p>Behind the scenes, the implementation uses either deep paging cursors or Streaming API and response streaming, so it is usually safe to specify a large number of rows. By default, the implementation uses 1000 rows but if your documents are smaller, you can increase this to 10000. Using too large a value can put pressure on the Solr JVM's garbage collector.</p> <p>Example: <code>options("rows", "10000")</code></p>
<code>1000</code>	<code>max_rows</code>	<p>The maximum number of rows; only applies when using the <code>/select</code> handler. The library will issue the query from a single task and let Solr do the distributed query processing.</p> <p>No paging is performed, that is, the <code>rows</code> param is set to <code>max_rows</code> when querying. Consequently, this option should not be used for large <code>max_rows</code> values, rather you should just retrieve all rows using multiple Spark tasks and then re-sort with Spark if needed.</p> <p>Example: <code>options("max_rows", "100")</code></p>

Name	Description	Default
None	<code>request_handler</code>	<p>Set the Solr request handler for queries. This option can be used to export results from Solr via <code>/export</code> handler which streams data out of Solr. See Exporting Result Sets for more information.</p> <p>The <code>/export</code> handler needs fields to be explicitly specified. Please use the <code>fields</code> option or specify the fields in the query.</p> <p>Example: <code>options("request_handler", "/export")</code></p>
<code>/select</code>	<code>splits</code>	<p>Enable shard splitting on default field <i>version</i>.</p> <p>Example: <code>options("splits", "true")</code></p> <p>The above option is equivalent to <code>options("split_field", "version")</code></p>
False	<code>split_field</code>	<p>The field to split on can be changed using <code>split_field</code> option.</p> <p>Example: <code>options("split_field", "id")</code></p>
<code>version</code>	<code>splits_per_shard</code>	<p>Split the shard into evenly-sized splits using filter queries. You can also split on a string-based keyword field but it should have sufficient variance in the values to allow for creating enough splits to be useful. In other words, if your Spark cluster can handle 10 splits per shard, but there are only 3 unique values in a keyword field, then you will only get 3 splits.</p> <p>Keep in mind that this is only a hint to the split calculator and you may end up with a slightly different number of splits than what was requested.</p> <p>Example: <code>options("splits_per_shard", "30")</code></p>

Name	Description	Default
20	<code>flatten_multivalued</code>	Flatten multi-valued fields from Solr. Example: <code>options("flatten_multivalued", "false")</code>
true	<code>dv</code>	Fetch the docValues that are indexed but not stored by using function queries. Should be used for Solr versions lower than 5.5.0. Example: <code>options("dv", "true")</code>
false	<code>sample_seed</code>	Read a random sample of documents from Solr using the specified seed. This option can be useful if you just need to explore the data before performing operations on the full result set. By default, if this option is provided, a 10% sample size is read from Solr, but you can use the <code>sample_pct</code> option to control the sample size. Example: <code>options("sample_seed", "5150")</code>
None	<code>sample_pct</code>	The size of a random sample of documents from Solr; use a value between 0 and 1. Example: <code>options("sample_pct", "0.05")</code>
0.1	<code>skip_non_dv</code>	Skip all fields that are not docValues. Example: <code>options("skip_non_dv", "true")</code>

10.5.4. Examples

Define a "movielens" project:

```
FUSION=localhost:8764
curl -u user:pass -X POST -H "Content-type:application/json" \
-d '{
  "name": "movielens",
  "assetType": "project",
  "description": "tables and views for the movielens project",
  "tags": ["movies","users"],
  "cacheOnLoad": false
}' "http://{fusion_path}/api/apollo/catalog"
```

Add a "ratings" table to the "movielens" project:

```
curl -u user:pass -X POST -H "Content-type:application/json" -d '{
  "name": "ratings",
  "assetType": "table",
  "projectId": "movielens",
  "description": "movie ratings data",
  "tags": ["movies"],
  "format": "solr",
  "cacheOnLoad": true,
  "options": ["collection -> movielens_ratings", "fields -> user_id,movie_id,rating,rating_timestamp"]
}' "http://{fusion_path}/api/apollo/catalog/movielens/assets"
```

Issue a SQL statement against the "ratings" table:

```
curl -u user:pass -X POST -H "Content-type:application/json" -d '{
  "name": "ratings",
  "assetType": "table",
  "projectId": "movielens",
  "description": "movie ratings data",
  "tags": ["movies"],
  "format": "solr",
  "cacheOnLoad": true,
  "options": ["collection -> movielens_ratings", "fields -> user_id,movie_id,rating,rating_timestamp"]
}' "http://{fusion_path}/api/apollo/catalog/movielens/query"
```

Issue a SQL query against the "movielens" project:

```
curl -u user:pass -X POST -H "Content-Type:application/json" -d '{
  "sql":"SELECT m.title as title, solr.aggCount as aggCount FROM movies m INNER JOIN (SELECT movie_id, COUNT(*)
as aggCount FROM ratings WHERE rating >= 4 GROUP BY movie_id ORDER BY aggCount desc LIMIT 10) as solr ON
solr.movie_id = m.movie_id ORDER BY aggCount DESC"
}' http://localhost:8764/api/apollo/catalog/movielens/query
```

Load a catalog table from a Postgres database:

```
curl -u user:pass -X POST -H "Content-type:application/json" -d '{
  "projectId": "nyc_taxi",
  "assetType": "table",
  "name": "trips",
  "sourceUri": "http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml",
  "owner": "Joe Example",
  "ownerEmail": "examplejoe@gmail.com",
  "description": "The NYC taxi trip data stored in Postgres using tools provided by
https://github.com/toddwschneider/nyc-taxi-data",
  "tags": ["nyc", "taxi", "postgres", "trips"],
  "format": "jdbc",
  "cacheOnLoad": true,
  "options": ["url -> ${nyc_taxi_jdbc_url}", "dbtable -> trips", "partitionColumn -> id", "numPartitions ->
4", "lowerBound -> 0", "upperBound -> $MAX(id)", "fetchSize -> 1000"],
  "filters": ["pickup_latitude >= -90 AND pickup_latitude <= 90 AND pickup_longitude >= -180 AND
pickup_longitude <= 180", "dropoff_latitude >= -90 AND dropoff_latitude <= 90 AND dropoff_longitude >= -180
AND dropoff_longitude <= 180"],
  "sql": "SELECT
id,cab_type_id,vendor_id,pickup_datetime,dropoff_datetime,store_and_fwd_flag,rate_code_id,passenger_count,trip
_distance,fare_amount,extra,mta_tax,tip_amount,tolls_amount,ehail_fee,improvement_surcharge,total_amount,payme
nt_type,trip_type, concat_ws(', ',pickup_latitude,pickup_longitude) as pickup,
concat_ws(', ',dropoff_latitude,dropoff_longitude) as dropoff FROM trips"
}' "http://{fusion_path}/api/apollo/catalog/nyc_taxi/assets"
```

Create a data asset using a streaming expression:

```
curl -u user:pass -X POST -H "Content-type:application/json" -d '{
  "name": "movie_ratings",
  "assetType": "table",
  "projectId": "movielens",
  "description": "movie ratings data",
  "tags": ["movies"],
  "format": "solr",
  "cacheOnLoad": true, "options": ["collection -> movielens_ratings", "expr ->
hashJoin(search(movielens_ratings,q=\"*:*\",fl=\"movie_id,user_id,rating\",sort=\"movie_id
asc\",qt=\"\\/export\",partitionKeys=\"movie_id\"),hashed=search(movielens_movies,q=\"*:*\",fl=\"movie_id,title
\",sort=\"movie_id asc\",qt=\"\\/export\",partitionKeys=\"movie_id\"),on=\"movie_id\")"]
}' "http://{fusion_path}/api/apollo/catalog/movielens/assets"
```

Send a Solr query:

```
curl -u user:pass -X POST -H "Content-Type:application/json" -d '{
  "solr": "*:*",
  "requestHandler": "/select",
  "collection": "movielens_movies",
  "params": {
    "facet": "on",
    "facet.field": "genre",
    "rows": 0
  }
}' http://localhost:8764/api/apollo/catalog/movielens/query
```

Send a Solr query using a streaming expression:

```
curl -u user:pass -X POST -H "Content-Type:application/json" --data-binary @streaming_join.json  
http://localhost:8764/api/apollo/catalog/movielens/query
```

```
{  
  "solr": "hashJoin(search(movielens_ratings, q=*, qt=\"/export\", fl=\"user_id,movie_id,rating\",  
    sort=\"movie_id asc\", partitionKeys=\"movie_id\"), hashed=search(movielens_movies, q=*,  
    fl=\"movie_id,title\", qt=\"/export\", sort=\"movie_id asc\", partitionKeys=\"movie_id\"), on=\"movie_id\"),  
  "collection": "movielens_ratings",  
  "requestHandler": "/stream"  
}
```

10.6. Collections API

The Collections API manages Fusion collections. It provides endpoints for creating, updating, and deleting collection, as well as endpoints for getting a collection's status and usage statistics.

Fusion maintains internal system collections for logs, blobs, and metrics data which operate in conjunction with collections created by users. The Collections API is used to manage all Fusion collections.

10.6.1. Examples

Create a new collection called 'newCollection', with appropriate SolrCloud environment settings:

REQUEST

```
curl -u user:pass -X PUT -H 'Content-type: application/json' -d
'{"solrParams":{"replicationFactor":1,"numShards":1}}'
http://localhost:8764/api/apollo/collections/newCollection
```

RESPONSE

```
{
  "id" : "newCollection",
  "createdAt" : "2014-09-19T18:46:52.954Z",
  "searchClusterId" : "default",
  "solrParams" : {
    "name" : "newCollection",
    "numShards" : 1,
    "replicationFactor" : 1
  },
  "type" : "DATA",
  "metadata" : { }
}
```

Create a collection named 'local-collection1' that refers to 'collection1' in a pre-existing SolrCloud cluster named 'Solr4.10' (see also the section Search Clusters):

REQUEST

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id":"local-collection1",
"searchClusterId":"Solr4.10", "solrParams":{"name":"collection1"}}'
http://localhost:8764/api/apollo/collections
```

RESPONSE


```
{
  "id" : "local-collection1",
  "createdAt" : "2014-09-19T18:48:45.396Z",
  "searchClusterId" : "Solr4.10",
  "solrParams" : {
    "name" : "collection1"
  },
  "type" : "DATA",
  "metadata" : { }
}
```

Delete a collection, but keep the associated signals and searchLogs collections:

REQUEST

```
curl -u user:pass -X DELETE http://localhost:8764/api/apollo/collections/newCollection?purge=false
```

Get the status of the 'demo' collection:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/collections/demo/status
```

RESPONSE

```
{
  "maxShardsPerNode" : 1,
  "replicationFactor" : 1,
  "shards" : {
    "shard1" : {
      "range" : "80000000-7fffffff",
      "state" : "active",
      "replicas" : {
        "core_node1" : {
          "state" : "active",
          "core" : "demo_shard1_replica1",
          "leader" : true,
          "base_url" : "http://localhost:8983/solr",
          "node_name" : "localhost:8983_solr"
        }
      }
    }
  }
}
```

Get stats for the 'demo' collection:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/collections/demo/stats
```

RESPONSE

```
{
  "collectionId" : "demo",
  "documentCount" : 536,
  "requestCount" : 6,
  "qps" : 28.34716542561849,
  "sizeInBytes" : 7646045,
  "lastModified" : "2014-05-19T19:58:33.545Z"
}
```

10.7. Collection Features API

The Collection Features API allows the following settings to be specified for a collection:

Property	Description
dynamicSchema	Modifies the Solr schema to be "managed", which means it's possible for Fusion to use Solr's schema API to manage the schema. It also sets Solr to operate in 'schemaless' mode, which means fields do not need to be pre-defined in the schema for them to be added to Solr's index. Note that this applies to the Solr included with Fusion, and does not modify an existing Solr cluster, if you have one already.
searchLogs	Creates a parallel collection for the storage of log data which is used to generate search query reports.
signals	Creates a parallel collection for the storage of signals data (such as user clicks, or ratings). Signals will need to be indexed and aggregated in order to be used. See the section on Signals for more information.
partitionByTime	Partition the corresponding Solr collection by time; see Time Series Indexing.

10.7.1. Examples

List the status of the features for the "demo" collection:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/collections/demo/features
```

RESPONSE

```
[ {  
  "name" : "dynamicSchema",  
  "collectionId" : "demo",  
  "params" : { },  
  "enabled" : false  
}, {  
  "name" : "searchLogs",  
  "collectionId" : "demo",  
  "params" : { },  
  "enabled" : false  
}, {  
  "name" : "signals",  
  "collectionId" : "demo",  
  "params" : { },  
  "enabled" : false  
} ]
```

Enable signals for a collection named 'demo':

REQUEST

```
curl -u user:pass -X PUT -H Content-type:application/json -d '{"enabled":true}'  
http://localhost:8764/api/apollo/collections/demo/features/signals
```

10.8. Connector APIs

Connectors run as a standalone module. Requests for connector services are routed through the Lucidworks proxy layer, which manages the authorizations for users. This means that all connector requests must contain authentication credentials, and access to the connector REST APIs can be limited to certain users.

There are several REST APIs for Connectors, listed here:

- Connector Datasources API
- Connector History API
- Connector JDBC API
- Connector Jobs API
- Connector Plugins API
- Connectors Crawl Database API
- Connector Status API

10.8.1. Connector Datasources API

The connector datasources API is used to create and configure datasources. See Connectors and Datasources for general information, Connectors and Datasource Reference for details on configuring specific datasources.

List a Datasource or All Datasources

The path for this request is:

```
/api/apollo/connectors/datasources?collection=<collectionName>
```

A GET request will return the definitions for all datasources for that collection. If the collection parameter is omitted, all datasource definitions will be returned.

Input

None.

Output

The output will include all datasources that match the request, with all properties.

Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Get datasources assigned to the "demo" collection:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/datasources?collection=demo
```

RESPONSE

```
[ {
  "id" : "database",
  "created" : "2014-05-04T19:47:22.867Z",
  "modified" : "2014-05-04T19:47:22.867Z",
  "connector" : "lucid.jdbc",
  "type" : "jdbc",
  "description" : null,
  "pipeline" : "conn_solr",
  "properties" : {
    "db" : null,
    "commit_on_finish" : true,
    "verify_access" : true,
    "sql_select_statement" : "select CONTENTID as id from CONTENT;",
    "debug" : false,
    "collection" : "demo",
    "password" : "password",
    "url" : "jdbc:postgresql://localhost:5432/db",
    "nested_queries" : null,
    "clean_in_full_import_mode" : true,
    "username" : "user",
    "delta_sql_query" : null,
    "commit_within" : 900000,
    "primary_key" : null,
    "driver" : "org.postgresql.Driver",
    "max_docs" : -1
  }
} ]
```

Create or Update a Datasource

A POST request will create a new datasource. The path for this request is:

</api/apollo/connectors/datasources>

A PUT request will update an existing datasource. The path for this request is:

</api/apollo/connectors/datasources/<id>>

where *<id>* is the name of the datasource to be updated.

Input

Property	Description
id <i>Required</i>	A unique identifier for this datasource. When creating a datasource, this property is required.
connector <i>Required</i>	The connector that will be used for this datasource.
type <i>Required</i>	The type of datasource, which must be supported by the defined connector.
description <i>Optional</i>	A description of the datasource.

Property	Description
pipeline <i>Optional</i>	The index pipeline that should be used for this datasource.
properties <i>Required</i>	<p>The properties for the datasource. The available properties will depend on the connector and type chosen.</p> <p>Not all available properties are required, but all datasources have at least one property that must be set to create the datasource.</p>

Output

A successful request returns the datasource definition. A failed request returns a JSON-formatted error message.

Examples

Create a datasource to index Solr-formatted XML files:

REQUEST

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id":"SolrXML",
"connector":"lucid.solrxml", "type":"solrxml", "properties":{"path":"/Applications/solr-
4.10.2/example/exampledocs", "generate_unique_key":false, "collection":"MyCollection}}'
http://localhost:8764/api/apollo/connectors/datasources
```

RESPONSE


```
{
  "id" : "SolrXML",
  "created" : "2015-05-18T15:47:51.199Z",
  "modified" : "2015-05-18T15:47:51.199Z",
  "connector" : "lucid.solrxml",
  "type" : "solrxml",
  "properties" : {
    "commit_on_finish" : true,
    "verify_access" : true,
    "generate_unique_key" : false,
    "collection" : "MyCollection",
    "include_datasource_metadata" : true,
    "include_paths" : [ "*.xml" ],
    "initial_mapping" : {
      "id" : "a35c9ff3-dbb6-434b-af40-597722c2986a",
      "skip" : false,
      "label" : "field-mapping",
      "type" : "field-mapping"
    },
  },
  "path" : "/Applications/apache-repos/solr-4.10.2/example/exampledocs",
  "exclude_paths" : [ ],
  "url" : "file:/Applications/apache-repos/solr-4.10.2/example/exampledocs/",
  "max_docs" : -1
}
```

Change the `max_docs` value for the above datasource:

REQUEST

```
curl -u user:pass -X PUT -H 'Content-type: application/json' -d '{"id":"SolrXML", "connector":"lucid.solrxml", "type":"solrxml", "properties":{"path":"/Applications/solr-4.10.2/example/exampledocs", "max_docs":10}}'
http://localhost:8764/api/apollo/connectors/datasources/SolrXML
```

RESPONSE

```
true
```

Delete a Datasource or All Datasources

The path for a DELETE request is:

</api/apollo/connectors/datasources/<id>>

where `<id>` is the name of the datasource to be deleted. If the id is omitted, all datasources will be deleted. The delete request takes one parameter:

Parameter	Description
force	If false , the default, the delete will wait for any running datasource jobs to complete before deleting the datasource. If you want to abort jobs before deleting, you can change this to true .

Input

None.

Output

The request succeeds silently. A failed request returns a JSON-formatted error message.

Examples

Delete the datasource named 'database':

REQUEST

```
curl -u user:pass -X DELETE http://localhost:8764/api/apollo/connectors/datasources/database
```

RESPONSE

If successful, no response.

10.8.2. Connector History API

The connector history REST API provides details of the last 50 runs of a named datasource. It also provides a way to clear the history.

Get or Delete the History for a Datasource

To GET or DELETE the history, the request path is:

```
/api/apollo/connectors/history/<id>
```

where *<id>* is the name of a datasource.

A GET request takes an optional parameter:

Parameter	Description
cumulative	If false, the default, details of each datasource run will be returned. Set to true if you would like to see a cumulative summary of the datasource runs.

Input

None.

Output

When the cumulative query parameter is set to true, the output will include the total number of documents processed as input to the pipeline, skipped, or failed as well as the total number of documents processed as output. Also shown will be the total number of runs of the datasource.

When the cumulative query parameter is set to false, details of each datasource job run will be shown for up to 50 past runs. The details shown will include any messages returned by the job (such as errors), the start and stop times of the job, the number of documents processed as input and output (including skipped or failed), and data for each stage of the index pipeline.

Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Show the cumulative history for a datasource named "Lucid5Docs":

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/history/Lucid5Docs?cumulative=true
```

RESPONSE

```
{
  "total_time" : 160000,
  "input" : 378,
  "skipped" : 5,
  "failed" : 7,
  "total_runs" : 7,
  "output" : 366
}
```

Show the detailed history for a datasource named "Lucid5Docs":

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/history/Lucid5Docs
```

RESPONSE

This output has been truncated to only show one successful run.

```
{
  "message" : null,
  "id" : "Lucid5Docs",
  "crawl_stopped" : "2014-06-05T15:16:08+0000",
  "pipeline" : {
    "stats" : {
      "counters" : {
        "stage.field-mapping::myMapping" : {
          "processed" : 58
        },
        "stage.logging::conn_logging" : {
          "info" : 116,
          "processed" : 116
        },
        "stage.solr-index::solr-default" : {
          "processed" : 58
        },
        "stage.tika-parser::tika" : {
          "input" : 58,
          "processed" : 58
        }
      },
      "gauges" : { },
      "histograms" : { },
      "meters" : { },
      "timers" : { }
    },
    "history" : {
      "events" : [ ]
    }
  },
  "crawl_started" : "2014-06-05T15:15:45+0000",
  "crawl_state" : "FINISHED",
  "counters" : {
    "input" : 60,
    "skipped" : 1,
    "failed" : 1,
    "output" : 58
  },
  "job_id" : "Lucid5Docs"
}
```

10.8.3. Connector JDBC API

Note	This API is deprecated and will be removed in a future release. To install and manage JDBC drivers, use the Blob Store API or the Blob Manager. See the JDBC Connector for instructions.
------	--

This REST API allows you to upload JDBC drivers for use with a JDBC datasource. It can also be used to find the driver class name needed for JDBC datasource configuration.

Drivers are uploaded to a directory named `fusion/3.1.x/data/connectors/lucid.jdbc/jdbc/<collection>`, where `<collection>` is the collection name used during the upload.

Note	When choosing a driver, be sure to use the latest version available that is compatible with your database and the Java version you are using with Fusion.
------	---

Upload a Driver

The path to upload a .jar file containing a JDBC driver via a POST request is:

```
/api/apollo/connectors/plugins/lucid.jdbc/resources/jdbc?collection=<collectionName>
```

where `<collectionName>` is the collection for which the driver will be used.

Input

The .jar file is a binary file. If using curl, use the `--form` option, or if using another client, use the supported method for calling binary files.

Output Content

No output will be returned.

Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Upload a Postgres driver to be used by a collection named 'docs':

REQUEST

```
curl -u user:pass -X POST --form file=@/path/postgresql-9.3-1101.jdbc4.jar  
http://localhost:8764/api/apollo/connectors/plugins/lucid.jdbc/resources/jdbc?collection=docs
```

RESPONSE

None.

List Available Drivers

A GET request will return all of the loaded drivers for a collection. The request path is:

```
/api/apollo/connectors/plugins/lucid.jdbc/resources/jdbc?collection=<collectionName>
```

where <collectionName> is the name of the collection.

Input

None.

Output

The response will include the name of the .jar file loaded and the available driver class names. The driver class name is the value that should be used when configuring a Database datasource.

Examples

Return drivers for a collection named 'docs':

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/plugins/lucid.jdbc/resources?collection=docs
```

RESPONSE

```
[ {  
  "name" : "postgresql-9.3-1101.jdbc4.jar",  
  "properties" : {  
    "type" : "jar"  
  }  
}, {  
  "name" : "org.postgresql.Driver",  
  "properties" : {  
    "type" : "class"  
  }  
} ]
```

Delete a Driver

The path to remove the .jar file of a JDBC driver from a specific collection via a DELETE request is:

```
/connectors/v1/connectors/plugins/lucid.jdbc/resources/<driver>?collection=<collection>
```

where <driver> is the name of the driver being removed, and <collection> is the name of the collection to which it is attached.

Input

None.

Output

No output will be returned. The removal of the JDBC driver can be confirmed by sending the above request to list available drivers and verifying its absence.

Example

Delete the Postgres driver from the collection named 'docs':

REQUEST:

```
curl -u admin:password -X DELETE http://localhost:8984/connectors/v1/connectors/plugins/lucid.jdbc/resources/postgresql-9.3-1101-jdbc4.jar?collection=docs
```

RESPONSE:

None.

Note that the DELETE request is sent to the connector port (default: 8984), rather than the Fusion UI port (default: 8764).

10.8.4. Connector Jobs API

The Connector Jobs API provides methods to start, monitor, and stop a datasource. A datasource is specific connector instance that connects to a defined repository, collects content, and sends it through an index pipeline.

The Connector Jobs API can also provide detailed information about the indexing job and each stage of the index pipeline.

To define and launch a job, simply use the datasource id with a POST request to start crawling.

Start, Stop or Check Status of a Job

The request path is:

`/api/apollo/connectors/jobs/<id>`

where `<id>` is the name of the datasource. A GET request will return the status of a job. A POST will create and start a job. PUT will start an existing job. DELETE will stop the job.

DELETE requests have the following parameters which control how running jobs are stopped:

Parameter	Description
abort	When false , the default, the job will be allowed to finish processing before stopping. Use true , if you need the job to stop immediately.
wait_time	The wait_time allows you to configure the number of milliseconds to wait while stopping a job before the job is aborted.

If the job id is not specified, the DELETE request has the following parameters:

Parameter	Description
connector	The name of a connector. This would allow you to stop all jobs using the 'lucid.dih' connector, for example, to stop all database crawls.
collection	The name of a collection, to stop all jobs for a single named collection.

Input

None.

Output

The output will include the state of the job (RUNNING, FINISHED, etc.), the start time, and documents retrieved so far (in `counters`).

Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Start crawling a datasource named "MyDocs":

REQUEST

```
curl -u user:pass -X POST http://localhost:8764/api/apollo/connectors/jobs/MyDocs
```

RESPONSE

```
{
  "id" : "MyDocs",
  "dataSourceId" : "MyDocs",
  "state" : "RUNNING",
  "message" : null,
  "startTime" : 1397840639000,
  "endTime" : -1,
  "finished" : false,
  "counters" : { },
  "exception" : null,
  "running" : true
}
```

Get Detailed Job Statistics

To GET detailed job statistics, the request path is:

```
`/api/apollo/connectors/jobs/<id>/pipeline`
```

where *<id>* is the name of the datasource.

Input

None.

Output

The output will show the results of each stage of the index pipeline, including the number of documents processed at each stage.

For example, the field-mapping stage would just show the number of documents that passed through the stage. The tika-parser stage, however, will show the number of documents input from the connector and the number output to the next stage. In a solr-index stage, the number of documents processed would indicate the number of documents that were added to the index.

If a stage encountered errors for any or all documents, the error would be shown for each document. If all documents failed due to a badly formed index pipeline, the output of this REST API may be quite lengthy.

Examples

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/jobs/TwitterSearch/pipeline
```

RESPONSE

```
{
  "stats" : {
    "counters" : {
      "stage.field-mapping::twitter-mapping" : {
        "processed" : 101
      },
      "stage.logging::conn_logging" : {
        "info" : 202,
        "processed" : 202
      },
      "stage.solr-index::solr-default" : {
        "command.ok.commit" : 1,
        "processed" : 201
      },
      "stage.tika-parser::tika" : {
        "input" : 101,
        "processed" : 202
      }
    },
    "gauges" : { },
    "histograms" : { },
    "meters" : { },
    "timers" : { }
  },
  "history" : {
    "events" : [ ]
  }
}
```

10.8.5. Connector Plugins API

The Connector Plugins API allows you to identify the available connectors and plugins as well as listing the datasource configuration properties for a specific connector plugin type. See [Connectors and Datasources](#) for discussion of the relationship between the two.

List Connectors

A GET request lists all available connectors. The path for this request is one of:

[/api/apollo/connectors](#)

[/api/apollo/connectors/plugins](#)

Both endpoints have the same behavior, so either can be used.

Input

None.

Output

A JSON list of connector names.

Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Get all available connectors:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/plugins
```

RESPONSE

```
[ "lucid.twitter.search", "lucid.push", "lucid.hadoop.mapr", "lucid.azure.table", "lucid.anda",  
  "lucid.hadoop.cloudera", "lucid.azure.blob", "lucid.hortonworks", "lucid.mongodb", "lucid.hadoop",  
  "lucid.hadoop.apache2", "lucid.fs", "lucid.hadoop.apache1", "lucid.jdbc", "lucid.hadoop.intel",  
  "lucid.twitter.stream", "lucid.solrxml" ]
```

List Connector Types

A GET request lists all supported plugin types for a specific connector. The path for this request is one of:

[/api/apollo/connectors/plugins/<name>](#)

[/api/apollo/connectors/plugins/<name>/types](#)

where *<name>* is a specific connector. These endpoints act the same, so either can be used.

To see the complete datasource specification for this connector plugin type, the GET request path is:

```
/api/apollo/connectors/plugins/<name>/types/<type>
```

where *<name>* is a specific connector and *<type>* is a supported plugin type for this connector.

Input

None.

Output

The output is a JSON schema that details each available property of each supported type for each connector requested.

The schema will include default values for any property, the data type for the property, if the property is editable, if it has a list of allowed values, and if it is required when creating a datasource for that type.

Examples

List the property specifications for the "file" type of the "lucid.fs" connector (the output has been truncated for space):

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/plugins/lucid.fs/types/file
```

RESPONSE

```
{
  "props": [ {
    "description": "general",
    "name": "---"
  }, {
    "read_only": false,
    "default_value": null,
    "description": "datasource.collection",
    "hints": [ "summary" ],
    "allowed_values": null,
    "name": "collection",
    "type": "string",
    "required": false
  }, {
    "read_only": false,
    "default_value": false,
    "description": "datasource.debug",
    "hints": [ "advanced" ],
    "allowed_values": null,
    "name": "debug",
    "type": "boolean",
    "required": false
  }, {
    "read_only": false,
    "default_value": null,
    "description": "datasource.initial_mapping",
    "hints": [ "advanced" ],
    "allowed_values": null,
    "name": "initial_mapping",
```

```

    "type" : "map",
    "required" : false
  }, {
    "read_only" : false,
    "default_value" : null,
    "description" : "datasource.db",
    "hints" : [ "advanced" ],
    "allowed_values" : null,
    "name" : "db",
    "type" : "map",
    "required" : false
  }, {
    "description" : "root path",
    "name" : "---"
  }, {
    "read_only" : false,
    "default_value" : null,
    "description" : "datasource.path",
    "hints" : [ "summary" ],
    "allowed_values" : null,
    "name" : "path",
    "type" : "string",
    "required" : true
  }, {
    "description" : "Filesystem URL",
    "name" : "---"
  }, {
    "read_only" : true,
    "default_value" : null,
    "description" : "datasource.url",
    "hints" : [ "summary" ],
    "allowed_values" : null,
    "name" : "url",
    "type" : "string",
    "required" : false
  },
  ...
] }

```

Drop ConnectorDB State

A DELETE request will drop the tables in the connector database for the named plugin. The request path is:

DELETE /api/apollo/connectors/plugins/<name>/state?collection=<collectionId>

where <name> is the name of the connector plugin for which to drop the state and <collectionId> is the name of a collection associated with this plugin. The collection parameter is optional.

Input

None.

Output Content

None.

Examples

Delete the crawl database for the SolrXML plugin:

REQUEST

```
curl -u user:pass -X DELETE http://localhost:8764/api/apollo/plugins/lucid.solrxml/state
```

RESPONSE

None.

10.8.6. Connector Status API

The connector registry loads all of the plugin classes and their dependencies. The Connector Status API reports the status of this registry. It can also reload or shut down the registry.

Reload, Quit or Get Status of Connector Registry

The path for the request is:

</api/apollo/connectors/status>

A GET request will return the connector registry status. A DELETE request will shut down the registry. A PUT will reload the registry. The PUT request takes an optional parameter:

Parameter	Description
initRegistry	If true , reinitialize the connectors registry.

Input

None.

Output Content

The output of a GET request will include the status of the registry and any messages reported.

The output of a PUT or a DELETE request will be empty.

Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Get the status of the connectors registry:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/status
```

RESPONSE

```
{
  "status" : "OK",
  "node" : "localhost",
  "messages" : [ {
    "message" : "2 data sources loaded, 2 data sources verified.",
    "time" : "2014-06-10T17:44:10.023Z"
  } ],
  "type" : "Local ConnectorManager"
}
```


10.8.7. Connectors Crawl Database API

Some of the connectors use a crawl database to track documents that have been seen by prior crawls and are able to use this information to understand which documents are new or have been updated or removed and take appropriate action in the index. The connectors that support this are currently lucid.fs and lucid.solrxml.

This API allows looking into the crawl database and dropping tables or clearing the database.

Note:

This API can be used only with connectors which support it, which at the current time are the 'lucid.fs' and 'lucid.solrxml' connectors.

The 'lucid.anda' connector also uses a crawl database, but it is not the same database, and does not have a REST API or other interface to access it.

Get Statistics for a Datasource Database or Drop Database

The path for this request is:

```
/api/apollo/connectors/datasources/<id>/db
```

where <id> is the name of the datasource.

A GET request will return statistics from the crawl database associated with a specific datasource. DELETE will drop the tables, meaning that the history of any crawl will be removed and all documents found on the next crawl will be treated as brand new documents and will be submitted to the index pipeline.

Input

None.

Output

The output from a GET request will include several sections detailing the database structure:

- counters: the counters section reports the document counts of database activities, such as table inserts.
- ops: the ops section reports on database operations that have occurred, such as initiating tables, retrieving items, processing items and table drops.
- tables: the tables section lists the tables in the database with a count of the number of items in each table. Inspecting the items is described in the next section.

The output from a DELETE request will be empty. When dropping the database, note that no documents will be removed from the index. However, the crawl database will be empty, so on the next datasource run, all documents will be treated as though they were never seen by the connectors.

Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Get the crawl database statistics for the datasource named "SolrXML":

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/datasources/SolrXML/db
```

RESPONSE

```
{
  "counters" : {
    "new" : 14,
    "processed.insert" : 14
  },
  "ops" : {
    "initTable" : 4,
    "dropTable" : 7,
    "flush" : 1,
    "getItem" : 28,
    "renameTable" : 2,
    "commitUpdates" : 1,
    "listTables" : 2,
    "finishProcessing" : 14,
    "beginUpdates" : 1,
    "insertItem" : 14
  },
  "tables" : {
    "deleted" : {
      "count" : 0
    },
    "discarded" : {
      "count" : 0
    },
    "errors" : {
      "count" : 0
    },
    "items" : {
      "count" : 14
    }
  }
}
```

Get Table Statistics or Drop the Table

The path for this request is:

```
/api/apollo/connectors/datasources/<id>/db/<table>
```

where *<id>* is the name of the datasource and *<table>* is the name of a database table.

A GET request will return the table statistics. A DELETE request will drop the table and clear its data.

Input

None.

Output

The output from a GET request will be the statistics for the named table. This is usually the item count.

The output from a DELETE request will be empty.

When dropping tables, be aware that the 'items' table does not delete documents from the index, but instead changes the database so database considers them new documents. When dropping other tables, such as the 'errors' table, it will merely clear out old error messages.

Examples

Get the statistics for the 'items' table in the 'SolrXML' datasource's connector database:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/datasources/SolrXML/db/items
```

RESPONSE

```
{
  "count" : 14
}
```

[Back to Top](#)

Get or Delete Table Items

The path for this request is:

[/api/apollo/connectors/datasources/<id>/db/items/<item>](#)

where *<id>* is the name of the datasource and *<item>* is the name of a specific item in the table. If no item name is specified, the request will get all items.

A GET request retrieves information about an item or items.

A DELETE request removes the information *from the Crawl Database only*. Note that this doesn't affect the Solr Index.

A request takes two optional parameters:

Parameter	Description
start	The starting key, which is the document ID. If empty, response will start at the first row of the table. Used with a GET request only.
rows	The number of rows to return. The default is to return all records. Used with a GET request only.

Input

None.

Output

The output of a GET request will include information on when the document was fetched, if it contained any links to other documents, and the size of the document.

The output of a DELETE request will be empty. Note that this does not delete a document from the index, it only changes the database so if or when the document is crawled again, the database considers it a new document.

Examples

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/connectors/datasources/SolrXML/db/items/items?rows=5
```

RESPONSE

```
{
  "/Applications/solr-4.8.0/example/exampledocs/gb18030-example.xml" : {
    "timestamp" : "1398117855000",
    "fetchUri" : null,
    "fetchTime" : 1402503143632,
    "docsCount" : 1,
    "outlinks" : null,
    "discarded" : false,
    "discardMessage" : null,
    "byteSize" : 1331,
    "exception" : null,
    "fetched" : true
  },
  "/Applications/solr-4.8.0/example/exampledocs/hd.xml" : {
    "timestamp" : "1398117855000",
    "fetchUri" : null,
    "fetchTime" : 1402503143691,
    "docsCount" : 1,
    "outlinks" : null,
    "discarded" : false,
    "discardMessage" : null,
    "byteSize" : 2241,
    "exception" : null,
    "fetched" : true
  }
}
```

10.9. Experiments API (experimental)

Use the Experiments API to compare different configuration variants and determine which ones are most successful. For example, configure different variants that use different query pipelines or recommendations, then analyze and compare search activity to see which variant best meets your goals.

Experiments allow you to evaluate multiple *variants* (named sets of configuration parameters) that yield measurable effects. An *Overall Evaluation Criterion* (OEC) is used to measure the *profit* value for each variant so they can be compared quantitatively.

Experiment types include A/B testing, multi-armed bandits optimization, evaluating efficiency of ML models, and search relevance. Experiments can be based on any metrics that are well-defined and computable for a given variant.

Note	This was introduced in Fusion 3.0 and is subject to change.
------	---

10.9.1. Experiment life cycle

Created	<p>The experiment and its variants are defined.</p> <p>There is no in-memory or saved data for any of the variants yet, so the OEC is the same for all variants. The experiment is in idle state.</p>
Edited	<p>The experiments configuration has been modified.</p> <p>Editing an experiments configuration make sense only before the experiment is started, or after all of its previous data is cleared. Otherwise there's a risk of mixing data from old configuration with the new data, which would make results meaningless.</p>
Started	<p>The experiment is ready for updates and new iterations (draws), consisting of "select variant" and "update" operations.</p> <p>In the current framework, experiments may be started/stopped multiple times (so the start/stop operations are equivalent to suspend/resume). If the service is restarted, the current state of experiment is initialized from any previously-saved state, if available.</p>
Read	<p>Current in-memory state is reported as well as the last saved results (if available).</p>

Updated	<p>One or more variants have been updated.</p> <p>When the experiment is running, users can perform “select variant” and “update variant” operations repeatedly, not necessarily in order (though lack of updates may skew the automatic variant selection for bandit experiments).</p>
Recomputed	<p>The experiment results are recomputed.</p> <p>The experiment must be stopped before recomputing. When this operation is completed, the in-memory state will be updated to reflect the computed results for each variant (including the OEC), and this state will be persisted (current implementation uses Fusions blob store).</p>
Stopped	<p>The experiment is in an idle state. The experiment will reject any new updates until it’s started again.</p>
Reset	<p>The experiment’s in-memory state is reset to the “zero” state, that is, a state that it would assume on “start” when no previously-saved data is available.</p> <p>Unlike “cleared” (below), this operation does <i>not</i> affect the previously-saved results.</p>
Cleared	<p>Both the experiment’s in-memory state <i>and</i> previously-saved results are discarded and the experiment is put in its “zero” state.</p> <p>Implementations may also discard any auxiliary data related to this experiment, such as raw events, logs of interactions, models, and so on.</p>

10.9.2. Experiment types

A/B testing and multi-armed bandits

Multi-armed bandit algorithms are designed to balance the exploration of each variant with the exploitation of the best variants in order to maximize the profit as measured by the OEC. An A/B test is the simplest type of multi-armed bandit experiment, comparing just two variants: A and B.

The algorithm determines which variant to select next for the trial and update. The update after each draw drives the algorithm, so if there are no updates it’s likely that the same variant will be proposed repeatedly. The OEC is calculated in real time and can be reported immediately (the “recompute” step is a no-op), which makes this class of experiment very lightweight.

The simplest setup is to create just two variants (A and B). Variant A is customarily considered the “ground truth”, “gold standard”, or “control” configuration, and variant B is the test configuration or “treatment” to compare with the control

configuration.

It is also common to perform A/A testing, that is, create a separate variant with exactly the same configuration as the control. A/A testing is useful for detecting any systemic errors.

The multi-armed bandits implementation allows you to perform both A/B and A/A testing as a part of the same experiment. To do this, create three variants: two variants with exactly the same “control” configuration and a third variant with the “treatment” configuration.

Implemented bandit algorithms

Epsilon Greedy	<p>Explore with probability <code>epsilon</code> and exploit with probability <code>1 - epsilon</code>.</p> <p>When <code>epsilon == 0.0</code>, the algorithm uses <i>annealing</i> (automatically decreasing <code>epsilon</code> based on the number of draws so far).</p> <p>This algorithm explores by selecting from all of the arms at random. It makes one of these random exploratory decisions with probability <code>epsilon</code>, otherwise always selecting the best arm.</p>
Softmax	<p>This algorithm explores by randomly selecting from all of the available arms with probabilities that are approximately proportional to the estimated value of each of the arms.</p> <p>If other arms are noticeably worse than the best arm, they're chosen with very low probability and the algorithm converges quickly to exploit the best variant. If the arms all have similar values, they're each chosen nearly equally often and the algorithm may never converge.</p> <p>When its parameter (called <i>temperature</i>) is close to 0.0, there is little randomness; the algorithm almost always selects the best arm (no exploration, only exploitation). As the temperature increases to <code>Inf</code>, it picks arms more randomly, thus increasing exploration at the cost of exploitation.</p> <p>Typically a value of 0.1 yields good results when one of the arms is clearly better than others. A value of 0.0 causes the algorithm to use annealing, that is, gradually decreasing temperature over time.</p>

UCB1	<p>Upper Confidence Bounds type 1 (UCB1) algorithm.</p> <p>This algorithm is deterministic and uses no parameters, which makes it much easier to use when the potential outcomes of experiment variants are difficult to predict. However, its accuracy and performance is somewhat lower than the best-tuned Epsilon Greedy or Softmax.</p> <p>UCB1 first ensures that it has played each arm at least once, avoiding the problem of cold start (though it means that you must update it at least as many times as there are arms). Then it selects arms based on their accumulated value <i>and</i> a bonus for the relative lack of knowledge about the arm (the inverse of visit counts for that arm). This results in occasional selection of lesser-known arms with lower values. However, over time a strong preference for the best arm(s) develops.</p>
------	--

OEC update strategy

Bandit algorithms use the numeric `value` property in the update payload and combine it with the accumulated value of the variant.

By default, the result is calculated as an arithmetic average of all values seen so far. However, as the number of draws increases to infinity, the impact of the recent updates becomes negligible.

Alternative implementation of update strategy is provided (`AlphaUpdateStrategy`), which proportionally decays the older accumulated value using a parameter `alpha` (ranging between 0.0 - 1.0). When `alpha` is close to 0.0, new values have a small impact on the accumulated value. When `alpha` is close to 1.0, new values are practically equal to the accumulated value. Common practice is to use `alpha` somewhere between 0.1 - 0.25.

IR quality metrics based on signals

This experiment type uses the `QualityAggregator` and signal data from arbitrary Spark DataFrame sources. Each set of signals represents data collected for a variant. These could be collected by external applications and sent to Fusion’s `/signals` or `/index-pipelines` endpoints.

Aggregation jobs are executed on collections of signals produced by different variants. The “recompute” operation uses aggregation jobs and related aggregation functions to compute the summary statistics and based on this it determines the OEC of each variant.

Signals are expected to represent multiple named lists of ranked results, each individual signal containing the following properties:

name	The list’s name. For example, for click log this is the <code>query_s</code> field.
item	The list’s item. For example, for click log this is the <code>doc_id_s</code> field that represents the clicked document ID.

rank	The rank of the item in the list. For example, for click log this is the <code>params.position_s</code> field that represents (1-based) rank of the clicked document on the list of results for a query.
tag	Optional property to separate signals belonging to different variants.

The source of signals for each variant is provided in the experiment variant’s configuration (as `input` property), and it represents a Spark DataFrame source (with `format` and `options` properties).

In case of this experiment type the decision about a variant selection is usually made elsewhere - in fact, it could be driven by a multi-armed bandits experiment running in parallel. By default each variant in turn is returned from “select variant” in a round-robin fashion. Updates will be usually provided somewhere else, too, eg. using the `/signals` or `/index-pipelines` endpoints, although the experiment API provides a pass-through to the `/signals` endpoint.

Arbitrary aggregations

This experiment type uses regular aggregation configurations.

Each experiment variant is configured with an existing aggregation ID. When the “recompute” action is invoked, the respective aggregation job is executed, and its summary statistics are used for the OEC calculation.

Machine learning pipelines

In this experiment type, each variant corresponds to a different configuration of a machine learning pipeline. Variant selection could be external or driven by multi-armed bandits.

Models build and updated using each variant’s configuration will be stored in Fusion’s blob store and only a reference to their location will be stored in the variant’s configuration (although for the purpose of updating they may need to be all loaded in memory).

10.9.3. Integration with query pipelines and index pipelines

Experiment Query Stage

This query stage lets you select a variant from a running experiment and inject its properties into the current `PipelineContext`. These properties can then be used by other query pipeline stages.

The Solr Index Stage uses these properties as overrides for request parameters. The properties that this stage sets (including the experiment ID and variant ID) are also returned in response, and can be further propagated to external applications, which can then include these ID-s in feedback to the experiment.

Experiment Update Stage

This indexing stage provides a way to apply updates sent to `/index-pipelines` or `/signals` directly to the experiment - in case of e.g. multi-armed bandits this method offers a much quicker turnaround than using aggregations.

Input documents are checked for the presence of experiment ID / variant ID and the value to be used for updating the experiment. By default these documents are silently discarded after processing, but the stage can be configured to forward them down the pipeline to the next stages.

10.10. Groups API

The Groups API allows you to view and define groups of Fusion objects.

10.11. Index Pipelines API

The Index Pipelines API provides methods for managing a set of named index pipelines. Every pipeline is made up of one or more stages. Stages can be defined during the creation of a pipeline, or stages can be defined separately and included into one or more pipelines. For details of the REST API for index stages, see Index Stages API.

Document processing proceeds stage by stage in a linear fashion. The order of the stages in a pipeline is the order in which they were defined. At installation, Fusion includes several pre-configured pipelines. See Index Pipelines for details on these default pipelines.

For more information about structuring documents for indexing, see Pushing Documents to a Pipeline.

10.11.1. Examples

List the 'default' pipeline: REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/index-pipelines/default
```

RESPONSE

```
{
  "id" : "default",
  "stages" : [ {
    "type" : "solr-index",
    "id" : "solr-default",
    "skip" : false
  } ]
}
```

Create an index pipeline named 'my-index-pipeline' with three stages, one of which does not yet exist:

REQUEST

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id":"my-index-pipeline","stages":[{"id":"tika","type":"tika-parser","includeImages":true},{"id":"conn_mapping","type":"ref"}, {"id":"solr-default","type":"ref"}]}'
http://localhost:8764/api/apollo/index-pipelines
```

RESPONSE

```
{
  "id" : "my-index-pipeline",
  "stages" : [ {
    "type" : "tika-parser",
    "id" : "tika",
    "includeImages" : true,
    "flattenCompound" : false,
    "addFailedDocs" : false,
    "addOriginalContent" : true,
    "contentField" : "_raw_content_",
    "skip" : false,
    "label" : "tika-parser"
  }, {
    "type" : "ref",
    "id" : "conn_mapping",
    "skip" : false,
    "label" : "ref"
  }, {
    "type" : "ref",
    "id" : "solr-default",
    "skip" : false,
    "label" : "ref"
  } ]
}
```

Reload the 'my-index-pipeline' pipeline:

INPUT

```
curl -u user:pass -X PUT http://localhost:8764/api/apollo/index-pipelines/my-index-pipeline/refresh
```

Index two JSON documents through a pipeline named 'conn_solr' and a collection named 'my-docs':

INPUT

```
curl -u user:pass -X POST -H "Content-Type: application/vnd.lucidworks-document" -d '[{"id": "myDoc1", "fields": [{"name": "title", "value": "My first document"}, {"name": "body", "value": "This is a simple document."}]}, {"id": "myDoc2", "fields": [{"name": "title", "value": "My second document"}, {"name": "body", "value": "This is another simple document."}]]' http://localhost:8764/api/apollo/index-pipelines/conn_solr/collections/my-docs/index
```

OUTPUT

```
[ {
  "id" : "myDoc1",
  "fields" : [ {
    "name" : "content",
    "value" : "This is a simple document.",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "title",
    "value" : "My first document",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "parsing_s",
    "value" : "no_raw_data",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "parsing_time_l",
    "value" : [ "java.lang.Long", 7 ],
    "metadata" : { },
    "annotations" : [ ]
  } ],
  "metadata" : { },
  "commands" : [ ]
}, {
  "id" : "myDoc2",
  "fields" : [ {
    "name" : "content",
    "value" : "This is another simple document.",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "title",
    "value" : "My second document",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "parsing_s",
    "value" : "no_raw_data",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "parsing_time_l",
    "value" : [ "java.lang.Long", 0 ],
    "metadata" : { },
    "annotations" : [ ]
  } ],
  "metadata" : { },
  "commands" : [ ]
} ]
```

Index a PDF document with the 'conn_solr' pipeline:

INPUT

```
curl -u user:pass -X POST -H "Content-Type: application/pdf" --data-binary @/solr/core/src/test-files/mailling_lists.pdf http://localhost:8764/api/apollo/index-pipelines/conn_solr/collections/my-docs/index
```

OUTPUT

```
[ {
  "id" : "d6c7757e-33d9-4fbb-aa38-eef84d679ca9",
  "fields" : [ {
    "name" : "fileSize_l",
    "value" : "8582",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "parsing_s",
    "value" : "no_raw_data",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "pageCount_i",
    "value" : "2",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "parsing_time_l",
    "value" : [ "java.lang.Long", 1171 ],
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "parsing_time_l",
    "value" : [ "java.lang.Long", 0 ],
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "attr_pdf:encrypted_",
    "value" : "false",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "attr_X-Parsed-By_",
    "value" : "org.apache.tika.parser.pdf.PDFParser",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "attr_pdf:PDFVersion_",
    "value" : "1.3",
    "metadata" : {
```

```

    "creator" : "tika-parser"
  },
  "annotations" : [ ]
}, {
  "name" : "attr_producer_",
  "value" : "FOP 0.20.5",
  "metadata" : {
    "creator" : "tika-parser"
  },
  "annotations" : [ ]
}, {
  "name" : "content",
  "value" : "\nSolr Mailing Lists\n\nTable of contents\n1 ",
  "metadata" : { },
  "annotations" : [ ]
}, {
  "name" : "attr_dc:format_",
  "value" : "application/pdf; version=1.3",
  "metadata" : {
    "creator" : "tika-parser"
  },
  "annotations" : [ ]
}, {
  "name" : "mimeType_s",
  "value" : "application/pdf",
  "metadata" : {
    "creator" : "tika-parser"
  },
  "annotations" : [ ]
} ],
"metadata" : { },
"commands" : [ ]
} ]

```

Index a JSON document through the 'conn_solr' pipeline into a collection called 'docs', using the "command" option:

INPUT

```

curl -u user:pass -X POST -H "Content-Type: application/vnd.lucidworks-document" -d '{"id":
"myDoc2","commands": [{"name":"delete","value": "myDoc2"}]},{"id": "myDoc1","commands":
[{"name":"delete","value": "myDoc1"}, {"name":"commit","value": "true"}]}'
http://localhost:8764/api/apollo/index-pipelines/conn_solr/collections/docs/index

```

OUTPUT

```
[ {
  "id" : "myDoc2",
  "fields" : [ ],
  "commands" : [ {
    "name" : "delete",
    "params" : { }
  } ]
}, {
  "id" : "myDoc1",
  "fields" : [ ],
  "commands" : [ {
    "name" : "delete",
    "params" : { }
  }, {
    "name" : "commit",
    "params" : { }
  } ]
} ]
```

Index two simple documents through a pipeline named 'conn_solr' and a collection named 'my-docs' and get a detailed output of the pipeline process:

INPUT

```
curl -u user:pass -X POST -H "Content-Type: application/vnd.lucidworks-document" -d '[{"id": "myDoc1", "fields": [{"name": "title", "value": "My first document"}, {"name": "body", "value": "This is a simple document."}], {"id": "myDoc2", "fields": [{"name": "title", "value": "My second document"}, {"name": "body", "value": "This is another simple document."}]}]' http://localhost:8764/api/apollo/index-pipelines/conn_solr/collections/my-docs/debug
```

OUTPUT

The output will include how each document passed through each stage. (In the example output below, we have truncated the 'field-mapping' stage for space.)

```
{
  "stages" : [ {
    "type" : "tika-parser",
    "id" : "conn_tika",
    "includeImages" : true,
    "flattenCompound" : false,
    "addFailedDocs" : true,
    "addOriginalContent" : true,
    "contentField" : "_raw_content_",
    "skip" : false
  }, {
    "type" : "field-mapping",
    "id" : "conn_mapping",
    "mappings" : [
      ...
    ],
    "unmapped" : {
      "source" : "/(.*)/",
      "target" : "attr_$1_"
    }
  }
]
```



```

    "operation" : "move"
  },
  "skip" : false
}, {
  "type" : "multivalued-resolver",
  "id" : "conn_multivalued_resolver",
  "typeStrategy" : [ {
    "fieldName" : "string",
    "resolverStrategy" : "pick_last"
  } ],
  "skip" : false
}, {
  "type" : "solr-index",
  "id" : "conn_solr",
  "enforceSchema" : true,
  "skip" : false
} ],
"output" : [ {
  "stageType" : "tika-parser",
  "stageId" : "conn_tika",
  "context" : {
    "simulate" : false,
    "stageIndex" : 0,
    "collection" : "docs",
    "async" : false
  },
  "docs" : [ {
    "id" : "6b5c10f1-d941-41a6-957f-f677f5ad0fd5",
    "fields" : [ {
      "name" : "attr_id_",
      "value" : "myDoc1",
      "metadata" : { },
      "annotations" : [ ]
    }, {
      "name" : "parsing_time_l",
      "value" : [ "java.lang.Long", 0 ],
      "metadata" : { },
      "annotations" : [ ]
    }, {
      "name" : "parsing_s",
      "value" : "no_raw_data",
      "metadata" : {
        "creator" : "tika-parser"
      },
      "annotations" : [ ]
    }, {
      "name" : "attr_fields_",
      "value" : [ "java.util.ArrayList", [ {
        "name" : "title",
        "value" : "My first document"
      } ], {
        "name" : "body",
        "value" : "This is a simple document."
      } ] ],
      "metadata" : { },
      "annotations" : [ ]
    } ],
    "metadata" : { },
    "annotations" : [ ]
  } ],
  "metadata" : { },

```

```

    "commands" : [ ]
  }, {
    "id" : "4dac3c4e-d7f5-4cbd-96dc-e2eae69711e3",
    "fields" : [ {
      "name" : "attr_id_",
      "value" : "myDoc2",
      "metadata" : { },
      "annotations" : [ ]
    }, {
      "name" : "parsing_time_l",
      "value" : [ "java.lang.Long", 0 ],
      "metadata" : { },
      "annotations" : [ ]
    }, {
      "name" : "parsing_s",
      "value" : "no_raw_data",
      "metadata" : {
        "creator" : "tika-parser"
      },
      "annotations" : [ ]
    }, {
      "name" : "attr_fields_",
      "value" : [ "java.util.ArrayList", [ {
        "name" : "title",
        "value" : "My second document"
      }, {
        "name" : "body",
        "value" : "This is another simple document."
      } ] ],
      "metadata" : { },
      "annotations" : [ ]
    } ],
    "metadata" : { },
    "commands" : [ ]
  } ]
}, {
  "stageType" : "field-mapping",
  "stageId" : "conn_mapping",
  "context" : {
    "simulate" : false,
    "stageIndex" : 1,
    "collection" : "docs",
    "async" : false
  },
  "docs" : [ {
    "id" : "6b5c10f1-d941-41a6-957f-f677f5ad0fd5",
    "fields" : [ {
      "name" : "attr_id_",
      "value" : "myDoc1",
      "metadata" : { },
      "annotations" : [ ]
    }, {
      "name" : "parsing_s",
      "value" : "no_raw_data",
      "metadata" : {
        "creator" : "tika-parser"
      },
      "annotations" : [ ]
    } ]
  } ]
} ]

```

```

}, {
  "name" : "parsing_time_l",
  "value" : [ "java.lang.Long", 0 ],
  "metadata" : { },
  "annotations" : [ ]
}, {
  "name" : "attr_fields_",
  "value" : [ "java.util.ArrayList", [ {
    "name" : "title",
    "value" : "My first document"
  } ], {
    "name" : "body",
    "value" : "This is a simple document."
  } ] ],
  "metadata" : { },
  "annotations" : [ ]
} ],
"metadata" : { },
"commands" : [ ]
}, {
  "id" : "4dac3c4e-d7f5-4cbd-96dc-e2eae69711e3",
  "fields" : [ {
    "name" : "attr_id_",
    "value" : "myDoc2",
    "metadata" : { },
    "annotations" : [ ]
  } ], {
    "name" : "parsing_s",
    "value" : "no_raw_data",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  } ], {
    "name" : "parsing_time_l",
    "value" : [ "java.lang.Long", 0 ],
    "metadata" : { },
    "annotations" : [ ]
  } ], {
    "name" : "attr_fields_",
    "value" : [ "java.util.ArrayList", [ {
      "name" : "title",
      "value" : "My second document"
    } ], {
      "name" : "body",
      "value" : "This is another simple document."
    } ] ],
    "metadata" : { },
    "annotations" : [ ]
  } ],
  "metadata" : { },
  "commands" : [ ]
} ]
}, {
  "stageType" : "multivalued-resolver",
  "stageId" : "conn_multivalued_resolver",
  "context" : {
    "simulate" : false,

```

```

"stageIndex" : 2,
"collection" : "docs",
"async" : false
},
"docs" : [ {
  "id" : "6b5c10f1-d941-41a6-957f-f677f5ad0fd5",
  "fields" : [ {
    "name" : "attr_id_",
    "value" : "myDoc1",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "parsing_s",
    "value" : "no_raw_data",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "parsing_time_l",
    "value" : [ "java.lang.Long", 0 ],
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "attr_fields_",
    "value" : [ "java.util.ArrayList", [ {
      "name" : "title",
      "value" : "My first document"
    } ], {
      "name" : "body",
      "value" : "This is a simple document."
    } ] ],
    "metadata" : { },
    "annotations" : [ ]
  } ],
  "metadata" : { },
  "commands" : [ ]
}, {
  "id" : "4dac3c4e-d7f5-4cbd-96dc-e2eae69711e3",
  "fields" : [ {
    "name" : "attr_id_",
    "value" : "myDoc2",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "parsing_s",
    "value" : "no_raw_data",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "parsing_time_l",
    "value" : [ "java.lang.Long", 0 ],
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "attr_fields_",

```

```

    "value" : [ "java.util.ArrayList", [ {
      "name" : "title",
      "value" : "My second document"
    }, {
      "name" : "body",
      "value" : "This is another simple document."
    } ] ],
    "metadata" : { },
    "annotations" : [ ]
  } ],
  "metadata" : { },
  "commands" : [ ]
} ]
}, {
  "stageType" : "solr-index",
  "stageId" : "conn_solr",
  "context" : {
    "simulate" : false,
    "stageIndex" : 3,
    "collection" : "docs",
    "async" : false
  },
  "docs" : [ {
    "id" : "6b5c10f1-d941-41a6-957f-f677f5ad0fd5",
    "fields" : [ {
      "name" : "attr_id_",
      "value" : "myDoc1",
      "metadata" : { },
      "annotations" : [ ]
    }, {
      "name" : "parsing_s",
      "value" : "no_raw_data",
      "metadata" : {
        "creator" : "tika-parser"
      },
      "annotations" : [ ]
    }, {
      "name" : "parsing_time_l",
      "value" : [ "java.lang.Long", 0 ],
      "metadata" : { },
      "annotations" : [ ]
    }, {
      "name" : "attr_fields_",
      "value" : [ "java.util.ArrayList", [ {
        "name" : "title",
        "value" : "My first document"
      }, {
        "name" : "body",
        "value" : "This is a simple document."
      } ] ],
      "metadata" : { },
      "annotations" : [ ]
    } ],
    "metadata" : { },
    "commands" : [ ]
  }, {
    "id" : "4dac3c4e-d7f5-4cbd-96dc-e2eae69711e3",
    "fields" : [ {

```

```

    "name" : "attr_id_",
    "value" : "myDoc2",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "parsing_s",
    "value" : "no_raw_data",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "parsing_time_l",
    "value" : [ "java.lang.Long", 0 ],
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "attr_fields_",
    "value" : [ "java.util.ArrayList", [ {
      "name" : "title",
      "value" : "My second document"
    } ], {
      "name" : "body",
      "value" : "This is another simple document."
    } ] ],
    "metadata" : { },
    "annotations" : [ ]
  } ],
  "metadata" : { },
  "commands" : [ ]
} ]
} ]
}

```

10.12. Index Profiles API

Index profiles are used to send documents to consistent endpoints, and to share pipelines among different collections. For information on using this service through the UI, see [Index Profiles](#).

Examples

Create a profile named 'testProfile' for the 'docs' collection and associate it with the pipeline named 'docs-default':

REQUEST

```
curl -u user:pass -X PUT -H 'Content-type: application/json' -d '"docs-default"'
http://localhost:8764/api/apollo/collections/docs/index-profiles/testProfile
```

RESPONSE

None.

Index two simple documents with the profile named 'testProfile':

REQUEST

```
curl -u user:pass -X POST -H "Content-Type: application/vnd.lucidworks-document" -d '[{"id":
"myDoc1","fields": [{"name":"title", "value":"My first document"}, {"name":"body", "value": "This is a simple
document."}]}, {"id": "myDoc2","fields": [{"name":"title","value": "My second document"}, {"name":"body",
"value": "This is another simple document."}]]' http://localhost:8764/api/apollo/collections/docs/index-
profiles/testProfile/index
```

RESPONSE

```

[ {
  "id" : "myDoc1",
  "fields" : [ {
    "name" : "content",
    "value" : "This is a simple document.",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "title",
    "value" : "My first document",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "parsing_s",
    "value" : "no_raw_data",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "parsing_time_l",
    "value" : [ "java.lang.Long", 6 ],
    "metadata" : { },
    "annotations" : [ ]
  } ],
  "metadata" : { },
  "commands" : [ ]
}, {
  "id" : "myDoc2",
  "fields" : [ {
    "name" : "content",
    "value" : "This is another simple document.",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "title",
    "value" : "My second document",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "parsing_s",
    "value" : "no_raw_data",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "parsing_time_l",
    "value" : [ "java.lang.Long", 0 ],
    "metadata" : { },
    "annotations" : [ ]
  } ],
  "metadata" : { },
  "commands" : [ ]
} ]

```


10.13. Index Stages API

The Index Stages API provides endpoints to: * manage query stage instances. * list query stage configuration properties * test processing on a set of queries

An index pipeline is comprised of index stages. Each index stage has a name and a type. The name identifies the stage instance, and the type identifies its class. Every stage type has a number of properties, which can be configured for a particular index stage instance. See the section Index Pipeline Stages for a taxonomy of index stage types.

10.13.1. Examples

View the configuration properties for index stage type "regex-extractor":

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/index-stages/schema/regex-extractor
```

RESPONSE

```

{
  "type" : "object",
  "title" : "Regex Field Extraction",
  "description" : "This stage allows you to extract entities using regular expressions",
  "properties" : {
    "rules" : {
      "type" : "array",
      "title" : "Regex Rules",
      "items" : {
        "type" : "object",
        "required" : [ "pattern" ],
        "properties" : {
          "source" : {
            "type" : "array",
            "title" : "Source Fields",
            "items" : {
              "type" : "string"
            }
          },
          "target" : {
            "type" : "string",
            "title" : "Target Field"
          },
          "pattern" : {
            "type" : "string",
            "title" : "Regex Pattern",
            "format" : "regex"
          },
          "annotateAs" : {
            "type" : "string",
            "title" : "Annotation Name"
          }
        }
      }
    }
  }
}

```

See all defined index pipeline stages, regardless of type:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/index-stages/instances
```

RESPONSE

```
[{
  "type" : "tika-parser",
  "id" : "conn_tika",
  "includeImages" : true,
  "flattenCompound" : false,
  "addFailedDocs" : true,
  "addOriginalContent" : true,
  "skip" : false
},
{
  "type" : "index-logging",
  "id" : "detailed-logging",
  "detailed" : true,
  "skip" : false,
  "label" : "detailed-index-logging",
}]
```

See details of an index-stage named 'conn_tika':

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/index-stages/instances/conn_tika
```

RESPONSE

```
{
  "type" : "tika-parser",
  "id" : "conn_tika",
  "includeImages" : true,
  "flattenCompound" : false,
  "addFailedDocs" : true,
  "addOriginalContent" : true,
  "skip" : false
}
```

Create a an index stage:

REQUEST

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id": "storagesize-regex-extractor",
"type":"regex-extractor", "rules": [{"source":["name"], "target":"storage_size_ss",
"pattern":"(\\d{1,20}\\s{0,3}(GB|MB|TB|KB|mb|gb|tb|kb))", "annotateAs":"storage_size"}]}'
http://localhost:8764/api/apollo/index-stages/instances
```

RESPONSE

```
{
  "type" : "regex-extractor",
  "id" : "storagesize-regex-extractor",
  "rules" : [ {
    "source" : [ "name" ],
    "target" : "storage_size_ss",
    "pattern" : "(\\d{1,20}\\s{0,3}(GB|MB|TB|KB|mb|gb|tb|kb))",
    "annotateAs" : "storage_size"
  } ],
  "skip" : false
}
```

Delete an index stage:

REQUEST

```
curl -u user:pass -X DELETE http://localhost:8764/api/apollo/index-stages/instances/storagesize-regex-extractor
```

No response is returned. To check that the stage is no longer defined, list all index stage instances.

Send a document through the index stage named 'conn_tika':

REQUEST

```
curl -u user:pass -X POST -H "Content-Type: application/json" -d '[{"id": "myDoc4", "fields": [{"name": "title", "value": "Another little document document"}, {"name": "body", "value": "This is a simple document."}]]'
http://localhost:8764/api/apollo/index-stages/instances/conn_tika/docs/test
```

RESPONSE

```
[ {
  "id" : "7b8a1d5b-9e42-40eb-8059-5804c4b4fc6b",
  "fields" : [ {
    "name" : "id",
    "value" : "myDoc4",
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "parsing_time",
    "value" : [ "java.lang.Long", 0 ],
    "metadata" : { },
    "annotations" : [ ]
  }, {
    "name" : "parsing",
    "value" : "no_raw_data",
    "metadata" : {
      "creator" : "tika-parser"
    },
    "annotations" : [ ]
  }, {
    "name" : "fields",
    "value" : [ "java.util.ArrayList", [ {
      "name" : "title",
      "value" : "Another little document document"
    }, {
      "name" : "body",
      "value" : "This is a simple document."
    } ] ],
    "metadata" : { },
    "annotations" : [ ]
  } ],
  "metadata" : { },
  "commands" : [ ]
} ]
```

10.14. Jobs API

The Jobs API allows you to define schedules for objects that perform work. You can schedule datasources, tasks, and Spark jobs.

This topic explains how to format requests to the Jobs API. To learn how to work with jobs and schedules in the Fusion UI, see Jobs and Schedules.

10.14.1. Job runtime configuration

Job configuration requests have these parameters:

- `<type>:<id>` path parameter (required)

The resource identifier takes the form `<type>:<id>`, such as `datasource:movie-db`, `spark:dailyMetricsRollup-counters`, or `task:delete-old-system-logs`.

See Job types below.

- `enabled` body parameter

"True" or "false" to enable or disable the job.

- `triggers` body parameter

This parameter defines one or more conditions that trigger the job. See Job triggers below.

Below is an example request that configures a job:

```
PUT /jobs/datasource:movie-db/schedule
{
  "enabled": true,
  "triggers": [
    {"type": "cron", "enabled": true, "value": "* /10 * * * *"}
  ]
}
```

Job types

The job type must be specified as part of the resource identifier in `PUT`, `POST`, and `DELETE` calls. These are the possible job types:

<code>datasource</code>	<p>A job to ingest data according to the specified datasource configuration, such as <code>datasource:movie-db</code>. Datasources are created using the Connector Datasources API or the Fusion UI.</p> <p>See Datasource Jobs.</p>
-------------------------	--

<code>spark</code>	<p>A Spark job to process data, such as <code>spark:dailyMetricsRollup-counters</code>. Spark jobs are created using the Spark Jobs API or the Fusion UI.</p> <p>See Spark Jobs.</p>
<code>task</code>	<p>A job to perform an HTTP call or log cleanup, such as <code>task:delete-old-system-logs</code>. Tasks are created using the Tasks API or the Fusion UI.</p> <p>See Tasks.</p>

Job triggers

Each job can have multiple triggers, and each trigger configuration requires `enabled` and `type` parameters.

Additional parameters are required by each trigger type, described below:

Trigger type	Parameters
<p><code>interval</code></p> <p>Run the job at a regular interval over some time range.</p>	<ul style="list-style-type: none"> <code>startTime</code> <p>The date and time at which the job will first run, as an ISO-8601 date-time string.</p> <code>endTime</code> <p>The date and time at which the job will run last, as an ISO-8601 date-time string.</p> <code>interval</code> <p>The interval at which the job will run, as an integer whose unit is specified by <code>repeatUnit</code> below.</p> <code>repeatUnit</code> <p>One of: <code>minute</code>, <code>hour</code>, or <code>day</code>.</p>
<p><code>cron</code></p> <p>Run the job on a cron-style schedule.</p>	<p><code>value</code></p> <p>A cron string.</p>

Trigger type	Parameters
<p>job</p> <p>Run the job after another job succeeds or fails.</p>	<ul style="list-style-type: none"> • jobID The ID of the job that triggers this one. • runOnSuccess "True" to run when the specified job succeeds. • runOnFailure "True" to run when the specified job fails. • runOnAbort "True" to run when the specified job is aborted.

10.14.2. Examples

Get all datasource jobs

Input

```
curl -u user:pass http://localhost:8764/api/apollo/jobs/?type=datasource
```

Output

```
[ {
  "resource" : "datasource:ratings-db",
  "enabled" : true,
  "startedBy" : "acb38917-611a-46e1-a3cf-4ed46a30fd24"
}, {
  "resource" : "datasource:movies-db",
  "enabled" : true,
  "startedBy" : "acb38917-611a-46e1-a3cf-4ed46a30fd24",
  "status" : "success",
  "lastStartTime" : "2017-05-17T20:58:10.503Z",
  "lastEndTime" : "2017-05-17T20:58:42.020Z"
} ]
```

Start a datasource job now

Input

```
curl -u user:pass -X POST -H "Content-Type: application/json"
http://localhost:8764/api/apollo/jobs/datasource:movies-db/actions -d '{"action": "start"}'
```

Output


```
{
  "resource" : "datasource:movies-db",
  "action" : "start",
  "accepted" : true,
  "message" : "Started job datasource:movies-db"
}
```

Get the job history for a datasource

Input

```
curl -u user:pass http://localhost:8764/api/apollo/jobs/datasource:movies-db/history
```

Output

```
[ {
  "resource" : "datasource:movies-db",
  "startTime" : "2017-05-17T21:01:46.743Z",
  "endTime" : "2017-05-17T21:02:05.720Z",
  "status" : "success"
}, {
  "resource" : "datasource:movies-db",
  "startTime" : "2017-05-17T20:58:10.503Z",
  "endTime" : "2017-05-17T20:58:42.020Z",
  "status" : "success"
} ]
```

Define a trigger for a datasource job

Input

```
curl -u user:pass -X PUT -H "Content-Type: application/json"
http://localhost:8764/api/apollo/jobs/datasource:ratings-db/schedule -d '{"triggers": [{"type": "cron",
"enabled": "true", "value": "*/10 * * * *"}]}'
```

Output

```
{
  "resource" : "datasource:ratings-db",
  "triggers" : [ {
    "type" : "cron",
    "enabled" : true,
    "expression" : "0 15 10 ? * MON-FRI",
    "type" : "cron"
  } ]
}
```

10.15. Links API

The Links API manages the links that represent the relationships between Fusion objects. You can see these links in Object Explorer by pressing Ctrl+K (on a PC) or Command +K (on a Mac) from any screen in the Fusion UI. You can also export and import links using the Objects API.

Links are structured as a tuple of (subject, object, linkType). An example would be (foo, bar, dependsOn), which would read as "foo dependsOn bar".

10.15.1. Link types and link reversibility

The possible values of linkType are as follows:

<p><code>dependsOn</code></p>	<p>The subject depends on the object.</p> <p>In this example, a datasource called "web_pokeapi_co-default" depends on a parser by the same name:</p> <pre data-bbox="820 747 1487 957">{ "subject" : "datasource:web_pokeapi_co-default", "object" : "parser:web_pokeapi_co-default", "linkType" : "dependsOn" }</pre>
<p><code>supports</code></p>	<p>The subject is depended on by the object (the reverse of <code>dependsOn</code>).</p> <p>In this example, a parser called "web_pokeapi_co-default" supports a datasource by the same name:</p> <pre data-bbox="820 1220 1487 1430">{ "subject" : "parser:web_pokeapi_co-default", "object" : "datasource:web_pokeapi_co-default", "linkType" : "supports" }</pre>

<p>isPartOf</p>	<p>The subject is part of the object, where the object is a group.</p> <p>In this example, an aggregation called "click-signals-default" is part of a group called "signals-object-group-default":</p> <pre>{ "subject" : "aggregation:click-signals-default", "object" : "group:signals-object-group-default", "linkType" : "isPartOf" }</pre>
<p>hasPart</p>	<p>The subject is a group, of which the object is a member (the reverse of isPartOf).</p> <p>In this example, an aggregation called "click-signals-default" is part of a group called "signals-object-group-default":</p> <pre>{ "subject" : "group:signals-object-group-default", "object" : "aggregation:click-signals-default", "linkType" : "hasPart" }</pre>
<p>relatesTo</p>	<p>The subject relates to the object. This is a uni-directional relationship that associates two objects that have none of the relationships described above.</p> <p>For example, when recommendations are enabled for the default collection, Fusion creates a group of objects which relates to the default collection:</p> <pre>{ "subject" : "group:recommendations-default", "object" : "collection:default", "linkType" : "relatesTo" }</pre>

With the exception of the **relatesTo** link type, links are reversible. **dependsOn** is the reverse of **supports** (and vice versa), while **isPartOf** is the reverse of **hasPart** (and vice versa).

For example, here is the group of links for the **signals-default** collection:

```
[
  {
    "subject": "group:4e6ee9a6-d4ab-413a-ac70-a27c1445bb60",
    "object": "collection:default_signals_aggr",
    "linkType": "hasPart"
  },
  {
    "subject": "group:4e6ee9a6-d4ab-413a-ac70-a27c1445bb60",
    "object": "spark:default_click_signals_aggregation",
    "linkType": "hasPart"
  },
  {
    "subject": "group:4e6ee9a6-d4ab-413a-ac70-a27c1445bb60",
    "object": "collection:default",
    "linkType": "relatesTo"
  },
  {
    "subject": "group:4e6ee9a6-d4ab-413a-ac70-a27c1445bb60",
    "object": "collection:default_signals",
    "linkType": "hasPart"
  }
]
```

When we request the links for one of the `hasPart` objects shown above, we get the reverse representation:

```
GET http://localhost:8765/api/v1/links?subject=collection:default_signals_aggr
```

```
[
  {
    "subject": "collection:default_signals_aggr",
    "object": "group:4e6ee9a6-d4ab-413a-ac70-a27c1445bb60",
    "linkType": "isPartOf"
  }
]
```

10.15.2. Examples

Get links to the default collection:

```
curl http://localhost:8765/api/v1/links?subject=collection:default
```

Output:

```
[ { "subject": "collection:default", "object": "index-pipeline:default", "linkType": "supports" } ]
```

Create a new link:

```
curl -X PUT http://localhost:8765/api/v1/links -H 'content-type: application/json' -d '{"linkType": "supports", "object": "index-pipeline:default", "subject": "collection:default"}'
```

Output:

```
{ "subject" : "collection:default", "object" : "index-pipeline:default", "linkType" : "supports" }
```

Delete a link:

```
curl -X DELETE http://localhost:8765/api/v1/links?subject=collection:default
```

10.16. Messaging API

The Messaging Service provides implementations to support sending messages and alerts under a coherent REST API. Service instances can be added and removed via REST API. Messages can be scheduled or sent immediately.

10.16.1. Attributes

The Messaging API supports the generic concept of a message, and each implementation interprets attributes in context. For example, "to:" in the context of a Slack message references the intended channel, whereas "to:" in the context of an email(SMTP) references the intended email address of its recipient.

Message Attribute	Description
<code>id</code>	An application specific id for tracking the message. Must be unique.
<code>type</code>	The type of message to send, one of the following: <code>* log *</code> <code>* smtp *</code> <code>* slack *</code> <code>* pagerduty</code>
<code>to</code>	One or more destinations for the message, as a list
<code>from</code>	Who/what the message is from
<code>subject</code>	The subject of the message
<code>body</code>	The main body of the message
<code>schedule</code>	Used to pass messages at a later time, or at a recurring basis.
<code>messageServiceParams</code>	Passes a map of any service-specific parameters, such as the user name and password

10.16.2. Examples

Display current messaging services:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/messaging/service
```

RESPONSE

```
[ {
  "type": "logging-message-service-config",
  "defaultLogLevel" : "info",
  "template": "<to>:<from>:<subject>:<body>",
  "serviceType": "log"
} ]
```

Send a message using Slack

REQUEST

```
curl -u user:pass -X POST -H 'Content-Type: application/json' -d '{"id": "myID", "type": "slack", "subject": "test", "body": "@recipient: this is a test", "to": ["recipient"], "from": "sender"}'
http://localhost:8764/api/apollo/messaging/send
```

RESPONSE

```
[ {
  "id": "myID",
  "type": "slack",
  "subject": "test",
  "body": "@recipient: this is a test",
  "to": "recipient",
  "from": "sender"
} ]
```

Get all scheduled messages

```
curl -u user:pass http://localhost:8764/api/apollo/messaging/scheduled
```

RESPONSE

```
[
  {
    "id": "myID",
    "type": "slack",
    "subject": "Updates",
    "body": "@recipient:here is the latest version",
    "to": ["updatechannel"],
    "from": "sender",
    "schedule":{"id":"slack", "creatorType":"human", "creatorId":"admin", "repeatUnit":"DAY", "interval":1,
    "startTime":"2015-05-21T06:44:00.000Z", "active":true}
  }
]
```

Show the current status of messaging services:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/messaging/status
```

RESPONSE

```
[ {  
  "status" : "ok",  
  "node" : "http://xxx.xx.xx.x:8764/api/apollo/messaging",  
  "messages" : [ ]  
} ]
```


10.17. Objects API

The Objects API lets you migrate objects between Fusion instances by exporting and importing. Fusion objects include all your searchable data, plus pipelines, aggregations, and other configurations on which your collections depend.

You can select all objects, or limit the operation to specific object types or IDs. In addition to export/import endpoints, a validation endpoint is provided for troubleshooting.

This service was introduced in Fusion 3.0.

Note	By default, system-created collections are not exported. See below for details.
------	---

10.17.1. Object export and import

Collections and encrypted values are treated specially; details are provided below. During import, conflicts are resolved according to the specified import policy.

For objects other than collections, no implicit filtering is performed; all objects are included by default. However, on export you can filter by type and ID.

Supported objects

Fusion lets you export and import these types of objects:

- `collection`
- `index-pipeline`
- `query-pipeline`
- `search-cluster`
- `datasource`
- `banana`
- `parser`
- `group`
- `link`
- `task`
- `job`
- `spark`

Exporting and importing collections

Collections are processed with these dependent objects:

- `features`
- `index profiles`
- `query profiles`

Datasources, parser configurations, and pipeline configurations are not included when collections are exported or imported. These must be exported and imported explicitly.

Only user-created collections are included by default. Certain types of collections are excluded:

- the "default" collection
- collections whose type is not DATA
- collections whose names start with "system_"
- "Secondary" collections, that is, collections created by features

Instead, create the same features on the target system; this automatically creates the corresponding secondary collections.

You can override these exclusions by specifying a collection, like this:

```
http://localhost:8764/api/apollo/objects/export?collection.ids=default
```

Encrypted passwords

Some objects, such as datasources and pipelines, include encrypted passwords for accessing remote data.

- On export, these encrypted values are replaced with `${secret.n.nameOfProperty}`.
- On import, the original, plaintext passwords must be provided in a JSON map:

```
{"secret.1.bindPassword" : "abc", "secret.2.bindPassword" : "def"}
```

The file must be supplied as multipart form data.

Note	Variables that do not start with <code>secret.</code> are ignored.
------	--

Import policies

On import, the `importPolicy` parameter is required. It specifies what to do if any object in the import list already exists on the target system:

<code>abort</code>	If there are conflicts, then import nothing.
<code>merge</code>	If there are conflicts, then skip the conflicting objects.
<code>overwrite</code>	If there are conflicts, then overwrite or delete/create the conflicting objects on the target system.

Filtering on export

On export, there are two ways to specify the objects to include:

- by type

You can specify a list of object types to export all objects of those types. Valid values:

- `collection`

- `index-pipeline`
- `query-pipeline`
- `search-cluster`
- `datasource`
- `banana`
- `parser`
- `group`
- `link`
- `task`
- `job`
- `spark`
- by type and ID

The `type.ids` parameter lets you list the IDs to match for the specified object type.

The `type` and `type.ids` parameters can be combined as needed.

Exporting linked objects

Related Fusion objects are linked. You can view linked objects using the Links API or the Object Explorer.

When exporting a specific Fusion object, you can also export its linked objects without specifying each one individually. To export all objects linked to the specified object, include the `deep="true"` query parameter in your request. See the example below. When `deep` is "true", Fusion follows these link types:

- `DependsOn`
- `HasPart`
- `RelatesTo`

10.17.2. Validation

Objects are validated before import. If any objects fail validation, the whole import request is rejected. A separate endpoint is available for validating objects without importing them.

Validation includes checking whether an object already exists on the target system and whether the user is authorized to create or modify the object.

For collection objects, the following special validation is performed:

- We check the `searchClusterId` of each collection and verify that a cluster with this ID exists on the target system or in the import file (error).
- We check that features, index profiles, and query profiles belong only to the collections specified in the import file (error).
- We check that a feature exists on the target system for each feature in the import file (error).
- We check for index profiles or query profiles that do not exist on the target system or in the import file (warning).

For `job` objects, which contain schedule configurations, Fusion only imports them if their associated `task`, `datasource`, or `spark` objects are also present, either on the target host or in the import file.

10.17.3. Status messages

Validation completed with no errors	The validation method was called and no errors found, though there may be warnings.
Validation found errors	The validation was called and errors found. Validation does not stop on the first error, so the complete list of errors is reported.
Validation was not completed because of system error	The validation was interrupted by system error.
Import was not performed because validation errors exist	The import method was called, but import didn't start because of validation errors.
Import was not performed because of input data error	The import method was called, but import didn't start, because Fusion could not find a substitution for one of the secret values in objects in import.
Import was not completed because of system error	The validation found no errors and import started, but it was interrupted by system error.
Import was completed	Validation found no errors and import finished successfully.

10.17.4. Examples

Export all objects

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export
```

Export all datasources

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export?type=datasource
```

Export a specific datasource and its linked objects

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export?export?datasource.ids=movies_csv-ml-movies&deep=true
```

Export all datasources and pipelines, plus two specific parsing configurations

```
curl -u user:pass http://localhost:8764/api/apollo/objects/export?type=datasource,index-pipeline,query-pipeline&parser.ids=cinema_parser,metafilis_parser
```

Import objects from a file and stop if there are conflicts

```
curl -u user:pass -H "Content-Type:multipart/form-data" -X POST -F  
'importData=@/Users/admin/Fusion/export.json'  
http://localhost:8764/api/apollo/objects/import?importPolicy=abort
```

Import objects, substitute the password variables, and merge any conflicts

```
curl -u user:pass -H "Content-Type:multipart/form-data" -X POST -F  
'importData=@/Users/admin/Fusion/export.json' -F 'variableValues=@password_file.json'  
http://localhost:8764/api/apollo/objects/import?importPolicy=merge
```

Note

`password_file.json` must contain plaintext passwords.

10.18. Parsers API

The Parser API works much like the Index Pipelines API and the Query Pipelines API, providing CRUD operations over parser configurations.

This service was introduced in Fusion 3.0.

To configure an index pipeline to use a specific parser, see the Index Pipelines API.

10.19. Query Pipelines API

The Query Pipelines API allows you to define parameters that should be applied to all queries processed by the system. A query pipeline consists of one or more stages. Query pipeline stages can be defined with the Query Stages API. See the Query Language Cheat Sheet for help constructing queries to send through a query pipeline.

10.19.1. Query Pipeline Definition Properties

Property	Description
id <i>Required</i>	Required only when creating a new pipeline.
stages <i>Required</i>	<p>A JSON map of the stages to include with the pipeline. The stages can exist already or they can be defined while defining the pipeline.</p> <p>The stage must include the stage definitions:</p> <ul style="list-style-type: none">• id: the ID of the stage to include.• stageParams: Any parameters for the stage. These are only required for a stage that has been pre-defined and you wish to modify the properties.• type: The stage type. You can define a stage directly on a pipeline, or you can use a pre-existing stage. If you use a pre-existing stage, you must use "ref", as a reference to an existing stage.• skip: If the stage should be skipped during pipeline processing. By default, this is 'false', and does not need to be defined unless you would like to skip a stage in the future. <p>If the stage supports setting other properties, as defined in the section Query Pipeline Stages, then those can be defined while defining the stage within the pipeline.</p>

10.19.2. Examples

Get the definition for the default query pipeline:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/query-pipelines/default
```

RESPONSE

```

{
  "id" : "default",
  "stages" : [ {
    "type" : "recommendation",
    "id" : "recommendation",
    "numRecommendations" : 10,
    "useFq" : true,
    "numSignals" : 100,
    "aggrType" : "click@doc_id_s-query_s-filters_s",
    "skip" : false
  }, {
    "type" : "solr-query",
    "id" : "solr",
    "skip" : false
  } ]
}

```

Create a new query pipeline, specifying a facet stage that is new and a solr-query stage that was previously defined:

REQUEST

```

curl -u user:pass -X POST -H 'Content-Type: application/json' -d '{"id":"my-pipeline",
"stages":[{"type":"facet", "id":"myFacets", "fieldFacets":[{"field":"data_source_s"}, {"field":"author_s",
"minCount":5, "missing":true}, {"field":"isRetweet_b"}], "skip":false}, {"id":"mySolr", "type":"ref",
"skip":false}]}' http://localhost:8764/api/apollo/query-pipelines

```

RESPONSE


```

{
  "id" : "my-pipeline",
  "stages" : [ {
    "type" : "facet",
    "id" : null,
    "fieldFacets" : [ {
      "field" : "data_source_s",
      "prefix" : null,
      "sort" : null,
      "limit" : null,
      "offset" : null,
      "missing" : null,
      "method" : null,
      "threads" : null,
      "minCount" : null,
      "enumCacheMinDf" : null
    }, {
      "field" : "author_s",
      "prefix" : null,
      "sort" : null,
      "limit" : null,
      "offset" : null,
      "missing" : true,
      "method" : null,
      "threads" : null,
      "minCount" : 5,
      "enumCacheMinDf" : null
    }, {
      "field" : "isRetweet_b",
      "prefix" : null,
      "sort" : null,
      "limit" : null,
      "offset" : null,
      "missing" : null,
      "method" : null,
      "threads" : null,
      "minCount" : null,
      "enumCacheMinDf" : null
    } ],
    "skip" : false
  }, {
    "type" : "solr-query",
    "id" : null,
    "skip" : false
  } ]
}

```

Query a pipeline named 'default' and a collection named "docs" for the term "solr". Also limit the response to only document 'titles' and return it in JSON format:

REQUEST

```

curl -u user:pass http://localhost:8764/api/apollo/query-
pipelines/default/collections/docs/select?q=solr&fl=title&wt=json

```

RESPONSE

```
{
  "response": {
    "numFound": 106,
    "start": 0,
    "docs": [
      {
        "title": [
          "Solr and SolrAdmin APIs - Fusion Documentation - Lucidworks"
        ]
      },
      {
        "title": [
          "Solr and SolrAdmin APIs - Fusion Documentation - Lucidworks"
        ]
      },
      {
        "title": [
          "Search Clusters - Fusion Documentation - Lucidworks"
        ]
      },
      {
        "title": [
          "Search Clusters - Fusion Documentation - Lucidworks"
        ]
      },
      {
        "title": [
          "Searching - Fusion Documentation - Lucidworks"
        ]
      }
    ]
  },
  "responseHeader": {
    "status": 0,
    "QTime": 1,
    "params": {
      "fl": "title",
      "lw.pipelineId": "default",
      "q": "solr",
      "wt": "json",
      "rows": "5",
      "defType": "edismax"
    }
  }
}
```

10.20. Query Profiles API

The Query Profiles API allows users to manage query profiles. Query profiles allow users to change query pipelines during search time while pointing the search toward a static endpoint. This gives flexibility to test different stage combinations without reconfiguration. For information, see Query Profiles.

10.20.1. Examples

Create a profile named 'testProfile' for the 'docs' collection and associate it with the pipeline named 'docs-default':

REQUEST

```
curl -u user:pass -X PUT -H 'Content-type: application/json' -d '{"docs-default"'
http://localhost:8764/api/apollo/collections/docs/query-profiles/testProfile
```

RESPONSE

None.

Query for the term 'title:fusion' in the 'docs' collection, using the profile named 'newProfile':

INPUT

```
curl -u user:pass http://localhost:8764/api/apollo/collections/docs/query-
profiles/newProfile/select?q=title:fusion&wt=json&fl=title
```

Note that we have also added a few other query parameters, such as 'wt=json' to get the results in JSON format, and 'fl=title' to only retrieve document titles.

OUTPUT

```
{
  "response": {
    "numFound": 85,
    "start": 0,
    "docs": [
      {
        "title": [
          "Fusion Services - Fusion Documentation - Lucidworks"
        ]
      },
      {
        "title": [
          "Lucidworks Fusion Documentation - Fusion Documentation - Lucidworks"
        ]
      },
      {
        "title": [
          "Installation - Fusion Documentation - Lucidworks"
        ]
      },
      {
```

```

        "title": [
            "Searching - Fusion Documentation - Lucidworks"
        ]
    },
    {
        "title": [
            "Collections - Fusion Documentation - Lucidworks"
        ]
    },
    {
        "title": [
            "Recommendations - Fusion Documentation - Lucidworks"
        ]
    },
    {
        "title": [
            "Connectors - Fusion Documentation - Lucidworks"
        ]
    },
    {
        "title": [
            "Reporting - Fusion Documentation - Lucidworks"
        ]
    },
    {
        "title": [
            "Schedules - Fusion Documentation - Lucidworks"
        ]
    },
    {
        "title": [
            "Profiles - Fusion Documentation - Lucidworks"
        ]
    }
]
},
"responseHeader": {
    "status": 0,
    "QTime": 1,
    "params": {
        "facet": "true",
        "fl": "title",
        "lw.pipelineId": "docs-default",
        "q": "title:fusion",
        "wt": "json",
        "defType": "edismax",
        "rows": "10"
    }
}
}
}

```

10.21. Query Stages API

The Query Stages API provides endpoints to: * manage query stage instances. * list query stage configuration properties * test processing on a set of queries

10.21.1. Examples

Get the properties for the "apply-defaults" type:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/query-stages/schema/set-params
```

RESPONSE

```
{
  "type" : "object",
  "title" : "Set Query Params",
  "description" : "This stage allows you to set, append, and remove query params",
  "properties" : {
    "params" : {
      "type" : "array",
      "title" : "Parameters and Values",
      "items" : {
        "type" : "object",
        "required" : [ "key" ],
        "properties" : {
          "key" : {
            "type" : "string",
            "title" : "Parameter Name"
          },
          "value" : {
            "type" : "string",
            "title" : "Parameter Value"
          },
          "policy" : {
            "type" : "string",
            "title" : "Update Policy",
            "enum" : [ "replace", "append", "remove" ],
            "default" : "append"
          }
        }
      }
    }
  }
}
```

See all defined query pipeline stages, regardless of type:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/query-stages/instances
```

RESPONSE

```
[{
  "type" : "query-logging",
  "id" : "detailed-logging",
  "detailed" : true,
  "skip" : false,
  "label" : "detailed-query-logging",
}]
```

Add a new query stage:

REQUEST

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d'{ "type" : "query-logging", "id" : "detailed-logging", "detailed" : true }' http://localhost:8764/api/apollo/query-stages/instances
```

RESPONSE

```
{
  "type" : "query-logging",
  "id" : "detailed-logging",
  "detailed" : true,
  "skip" : false,
  "label" : "query-logging"
}
```

Update a query stage:

REQUEST

```
curl -u user:pass -X PUT -H 'Content-type: application/json' -d'{ "type" : "query-logging", "id" : "detailed-logging", "detailed" : true, "label" : "detailed-query-logging" }' http://localhost:8764/api/apollo/query-stages/instances/detailed-logging
```

RESPONSE

```
{
  "type" : "query-logging",
  "id" : "detailed-logging",
  "detailed" : true,
  "skip" : false,
  "label" : "detailed-query-logging",
}
```

Note that all required elements must be included in the update.

Delete a query stage:

REQUEST

```
curl -u user:pass -X DELETE http://localhost:8764/api/apollo/query-stages/instances/detailed-logging
```

No response is returned. To check that the stage is no longer defined, list all query stage instances.

Test that a set-params stage defines properties correctly:

REQUEST

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"type":"set-params",  
"params":[{"key":"rows", "value":"2", "policy":"append"}]}' http://localhost:8764/api/apollo/query-  
stages/solr/test?q=*:*
```

RESPONSE

```
{  
  "request" : {  
    "headers" : {  
      "User-Agent" : [ "curl/7.30.0" ],  
      "Content-Type" : [ "application/json" ],  
      "Accept" : [ "*/*" ],  
      "Host" : [ "localhost:8764" ],  
      "Content-Length" : [ "80" ]  
    },  
    "params" : {  
      "q" : [ "*:*" ],  
      "rows" : [ "2" ]  
    }  
  },  
  "response" : null,  
  "totalTime" : 0  
}
```

10.22. Recommendations API

The Recommendations REST API uses signals and aggregated signals to return a list of items that can be used for recommendations.

To use the REST API Recommendations service to get recommendations for items in some collection, that collection must have associated signals and aggregated-signals system collections. How good the recommendations are depends on how well the information in the signals and aggregated signals collections, which is derived from observed user behavior, matches user behavior going forward.

In addition to these endpoints, is also possible to get recommendations as part of a query request.

Fusion 3.1 introduces new UI functionality for enabling recommendations and two new recommenders, which you can use in query pipelines: user-item recommenders and item-similarity recommenders.

See Recommendations for a discussion of when to use the API and when to use recommender query stages.

10.22.1. Recommendation types

The API includes separate endpoints for retrieving different types of recommendations:

<code>itemsForQuery</code>	Get items for a defined query. The response is a list of <i>document IDs</i> and their weights.
<code>queriesForItem</code>	Get queries for a defined item (a document ID). This finds the top queries that led users to the defined item. The response is a list of <i>query terms</i> and their weights.

10.22.2. Output

The output includes the following sections:

<code>header</code>	The query parameters (in a section named <code>queryParams</code>) and the total time it took to process the query.
<code>items</code>	Depending on the recommendation type: <ul style="list-style-type: none"><code>itemsForQuery</code> The document IDs and the weights of aggregated events that match the query. This type supports a <code>debug</code> option that adds a <code>debug</code> section to the output.<code>queriesForItem</code> The queries and the weights of aggregated events that match the document ID.

10.22.3. Examples

Below are examples for each recommendation type.

itemsForQuery

Get the top items for the query 'ipod':

INPUT

```
curl -u user:pass
http://localhost:8764/api/apollo/recommend/lucidworks102/itemsForQuery?q=ipod&fq=count_d:4&debug=true
```

OUTPUT

```
{
  "header" : {
    "queryParams" : {
      "aggrType" : "*",
      "rows" : 10,
      "collection" : "lucidworks102",
      "aggrRows" : 100,
      "debug" : true,
      "q" : "ipod",
      "fq" : [ "count_d:4" ]
    },
    "totalTime" : 5
  },
  "items" : [ {
    "weight" : 1.0726584E-11,
    "docId" : "8771929"
  }, {
    "weight" : 3.865899E-12,
    "docId" : "9225439"
  }, {
    "weight" : 9.230597E-12,
    "docId" : "3109302"
  } ],
  "debugInfo" : {
    "aggrTime" : 1,
    "queryTime" : 4,
    "solrParams" : {
      "mm" : [ "50%" ],
      "pf" : [ "query_t^3", "query_t~2^7", "query_t~0^1" ],
      "fl" : [ "id", "doc_id_s", "weight_d" ],
      "sort" : [ "score desc,weight_d desc" ],
      "q" : [ "ipod" ],
      "qf" : [ "query_t" ],
      "collection" : [ "lucidworks102_signals_aggr" ],
      "fq" : [ "aggr_type_s:*", "count_d:4" ],
      "rows" : [ "100" ],
      "defType" : [ "edismax" ]
    }
  }
}
```

queriesForItem

INPUT

```
curl -u user:pass http://localhost:8764/api/apollo/recommend/lucidworks102/queriesForItem?docId=9225439
```

OUTPUT

```
{
  "header" : {
    "queryParams" : {
      "aggrType" : "*",
      "rows" : 10,
      "collection" : "lucidworks102",
      "docId" : "9225439"
    },
    "totalTime" : 8
  },
  "items" : [ {
    "query" : "ipod",
    "weight" : 3.865899E-12
  }, {
    "query" : "columbusday ipod mp3 20111009",
    "weight" : 3.5141304E-12
  }, {
    "query" : "apple itouch",
    "weight" : 2.3619889E-12
  }, {
    "query" : "ipod 4th generation",
    "weight" : 1.6436526E-12
  }, {
    "query" : "ipod touch 4th generation",
    "weight" : 9.674966E-13
  }, {
    "query" : "onlinemidnightsale ipod mp3players",
    "weight" : 9.568035E-13
  }, {
    "query" : "ipod touch",
    "weight" : 7.774231E-13
  }, {
    "query" : "itouch",
    "weight" : 7.707221E-13
  } ]
}
```

10.23. Reporting API

The reporting API provides a set of reports for key metrics about user searches over a primary Fusion collection by running reports against the auxiliary searchLog and signals collections. Reports include:

- 'topClicked' - documents which have received the most clicks.
- 'topQueries' - queries ordered by frequency.
- 'lessThanN' - queries which return less than N documents.
- 'topN' - query terms ordered by frequency.
- 'histo' - a histogram over all queries binned by length of query execution time.
- 'dateHisto' - shows relative rate of queries over time.

The reports available for a particular collection depends on whether or not the auxiliary searchLogs and signals collection have been created. Collections created with the Fusion UI will have both searchLogs and signals by default. The 'topClicked' report requires both searchLogs and signals auxiliary collections. All other reports require searchLogs auxiliary collections.

10.23.1. Report Configuration Information

POST requests should always send a JSON object containing the Report configuration. If no special configuration is required, this is the empty object "{}".

The report configuration object contains the following attributes, depending on the report:

- "n" value is a positive integer, used for both "topN" and "lessThanN" reports. E.g.: "n" : 1
- "num" value is a positive integer, used for both "topClicked" report. E.g.: "num" : 5
- "field" : required for report "topN", specifies the field where the search terms are stored. Default field is "q_txt". E.g.: "field" : "q_txt" See Search Query Reporting for details.
- "rangeStart", "rangeEnd", "interval" : attributes used to restrict histogram report range and set bin size accordingly. E.g.: "rangeStart": 0, "rangeEnd": 1000, "interval" : 1000
- "dateRangeStart", "dateRangeEnd", "timeInterval" : attributes required for date range histogram report. E.g.: "dateRangeStart": "NOW/DAY-1DAY", "dateRangeEnd": "NOW/DAY+1DAY", "timeInterval": "+1DAY"

10.23.2. Examples

See reports available for collection "bb_catalog" which has searchLogs and signals enabled:

```
> curl -u user:pass http://localhost:8764/api/apollo/reports/bb_catalog  
[ "histo", "topClicked", "topQueries", "lessThanN", "dateHisto", "topN" ]
```

See reports available for system collection "system_blobs" which doesn't have any auxiliary collections:

```
> curl -u user:pass http://localhost:8764/api/apollo/reports/system_blobs  
[ ]
```

Identify queries over collection "bb_catalog" for which no matching documents are found, i.e., queries which return less than 1 result:

```
> curl -u user:pass -X POST -H 'Content-type: application/json' -d @- \
> http://localhost:8764/api/apollo/reports/bb_catalog/lessThanN \
> <<EOF
> {"n":1}
> EOF

[ {
  "key" : "ipad",
  "count" : 3,
  "percentage" : 0.375,
  "token" : "eyJmaWx0ZXJzIjpbIm51bWRvY3NfbDpbKiBUTyAxXSIsInFfczppcGFkIl19"
}, {
  "key" : "id:2125233",
  "count" : 2,
  "percentage" : 0.25,
  "token" : "eyJmaWx0ZXJzIjpbIm51bWRvY3NfbDpbKiBUTyAxXSIsInFfczppZFxc0jIxmJyUyMzMlXX0="
}, {
  "key" : "ipod",
  "count" : 1,
  "percentage" : 0.125,
  "token" : "eyJmaWx0ZXJzIjpbIm51bWRvY3NfbDpbKiBUTyAxXSIsInFfczppcG9kIl19"
}, {
  "key" : "typewriter",
  "count" : 1,
  "percentage" : 0.125,
  "token" : "eyJmaWx0ZXJzIjpbIm51bWRvY3NfbDpbKiBUTyAxXSIsInFfczpp0eXBld3JpdGVyIl19"
}, {
  "key" : "unicorn",
  "count" : 1,
  "percentage" : 0.125,
  "token" : "eyJmaWx0ZXJzIjpbIm51bWRvY3NfbDpbKiBUTyAxXSIsInFfczpp1bmljb3JuIl19"
} ]
```

Drill down on "lessThanN" report to examine information for "key": "ipad" by token ID:

```
> curl -u user:pass \
>
http://localhost:8764/api/apollo/reports/bb_catalog/lessThanN/eyJmaWx0ZXJzIjpbIm51bWRvY3NfbDpbKiBUTyAxXSIsInFfczppcGFkIl19

{
  "numFound" : 3,
  "start" : 0,
  "maxScore" : 0.0,
  "docs" : [ {
    "id" : "cdcdd42c-66f2-499e-a940-33d980596d36",
    "collection_s" : "bb_catalog",
    "q_txt" : [ "ipad" ],
    "q_s" : "ipad",
    "qtime_l" : 1,
    "totaltime_l" : 2,
    "numdocs_l" : 0,
    "timestamp_dt" : "2015-08-31T20:40:48.096Z",
```

```

"httpmethod_s" : "POST",
"req_q_ss" : [ "ipad" ],
"req_debug_ss" : [ "true" ],
"req_json.nl_ss" : [ "arrarr" ],
"req_echoParams_ss" : [ "all" ],
"req_lw.pipelineId_ss" : [ "bb_catalog-default" ],
"req_fl_ss" : [ "*,score" ],
"req_start_ss" : [ "0" ],
"req_isFusionQuery_ss" : [ "true" ],
"req_sort_ss" : [ "score desc" ],
"req_rows_ss" : [ "10" ],
"req_wt_ss" : [ "json" ],
"_version_" : 1511054270105911296
}, {
  "id" : "d4e22f4e-ae27-4662-82ed-17c68111f0d5",
  "collection_s" : "bb_catalog",
  "q_txt" : [ "ipad" ],
  "q_s" : "ipad",
  "qtime_l" : 3,
  "totaltime_l" : 4,
  "numdocs_l" : 0,
  "timestamp_dt" : "2015-09-01T13:45:28.008Z",
  "httpmethod_s" : "POST",
  "req_debug_ss" : [ "true" ],
  "req_json.nl_ss" : [ "arrarr" ],
  "req_echoParams_ss" : [ "all" ],
  "req_lw.pipelineId_ss" : [ "bb_catalog-default" ],
  "req_fl_ss" : [ "*,score" ],
  "req_start_ss" : [ "0" ],
  "req_isFusionQuery_ss" : [ "true" ],
  "req_rows_ss" : [ "10" ],
  "req_bq_ss" : [ "id:1945531^4.090439397841692", "id:2339322^1.5108471289277077",
  "id:1945595^1.0636971555650234", "id:1945674^0.40656840801239014", "id:2842056^0.33429211378097534",
  "id:2408224^0.43880610167980194", "id:2339386^0.39254774153232574", "id:2319133^0.32736557722091675",
  "id:9924603^0.19560790061950684", "id:1432551^0.18906432390213013" ],
  "req_q_ss" : [ "ipad" ],
  "req_defType_ss" : [ "edismax" ],
  "req_wt_ss" : [ "json" ],
  "req_facet_ss" : [ "true" ],
  "_version_" : 1511118736467165184
}, {
  "id" : "a249d93c-9232-4ea7-a99a-fcf01b6c2c2f",
  "collection_s" : "bb_catalog",
  "q_txt" : [ "ipad" ],
  "q_s" : "ipad",
  "qtime_l" : 0,
  "totaltime_l" : 2,
  "numdocs_l" : 0,
  "timestamp_dt" : "2015-09-01T13:46:41.309Z",
  "httpmethod_s" : "POST",
  "req_q_ss" : [ "ipad" ],
  "req_debug_ss" : [ "true" ],
  "req_json.nl_ss" : [ "arrarr" ],
  "req_echoParams_ss" : [ "all" ],
  "req_lw.pipelineId_ss" : [ "bb_catalog-default" ],
  "req_fl_ss" : [ "*,score" ],
  "req_start_ss" : [ "0" ],
  "req_isFusionQuery_ss" : [ "true" ],

```

```

    "req_rows_ss" : [ "10" ],
    "req_wt_ss" : [ "json" ],
    "req_facet_ss" : [ "true" ],
    "req_bq_ss" : [ "id:1945531^4.090439397841692", "id:2339322^1.5108471289277077",
    "id:1945595^1.0636971555650234", "id:1945674^0.40656840801239014", "id:2842056^0.33429211378097534",
    "id:2408224^0.43880610167980194", "id:2339386^0.39254774153232574", "id:2319133^0.32736557722091675",
    "id:9924603^0.19560790061950684", "id:1432551^0.18906432390213013" ],
    "_version_" : 1511118813327785984
  } ]
}

```

Get all of the top queries, regardless of date, pass in empty date range specification:

```

> curl -u user:pass -X POST -H 'Content-type: application/json' -d '{}' \
> http://localhost:8764/api/apollo/reports/bb_catalog/topQueries

[ {
  "key" : "ipad",
  "count" : 42,
  "percentage" : 0.7118644,
  "token" : "eyJmaWx0ZXJzIjpbInFfczppcGFkI119"
}, {
  "key" : " *.*",
  "count" : 10,
  "percentage" : 0.16949153,
  "token" : "eyJmaWx0ZXJzIjpbInFfczpcXCpcXDpcXCoiXX0="
}, {
  "key" : "id:2125233",
  "count" : 2,
  "percentage" : 0.033898305,
  "token" : "eyJmaWx0ZXJzIjpbInFfczppZFxcOjIxMjUyMzMiXX0="
}, {
  "key" : "typewriter",
  "count" : 2,
  "percentage" : 0.033898305,
  "token" : "eyJmaWx0ZXJzIjpbInFfczpp0eXBld3JpdGVyI119"
}, {
  "key" : "unicorn",
  "count" : 2,
  "percentage" : 0.033898305,
  "token" : "eyJmaWx0ZXJzIjpbInFfczpp1bmljb3JuI119"
}, {
  "key" : "ipod",
  "count" : 1,
  "percentage" : 0.016949153,
  "token" : "eyJmaWx0ZXJzIjpbInFfczppcG9kI119"
} ]

```

_Drill down on topQueries report for item with "key" : "ipod", "token": "eyJmaWx0ZXJzIjpbInFfczppcG9kI119"

```

> curl -u user:pass \
> http://localhost:8764/api/apollo/reports/bb_catalog/topQueries/eyJmaWx0ZXJzIjpbInFfczppcG9kIl19

{
  "numFound" : 1,
  "start" : 0,
  "maxScore" : 0.0,
  "docs" : [ {
    "id" : "4a6f7f5e-3d13-4f20-b59e-6188ce4c5783",
    "collection_s" : "bb_catalog",
    "q_txt" : [ "ipod" ],
    "q_s" : "ipod",
    "qtime_l" : 1,
    "totaltime_l" : 2,
    "numdocs_l" : 0,
    "timestamp_dt" : "2016-04-05T17:51:56.197Z",
    "httpmethod_s" : "POST",
    "req_debug_ss" : [ "true" ],
    "req_json.nl_ss" : [ "arrarr" ],
    "req_echoParams_ss" : [ "all" ],
    "req_lw.pipelineId_ss" : [ "default" ],
    "req_fl_ss" : [ "*,score" ],
    "req_start_ss" : [ "0" ],
    "req_isFusionQuery_ss" : [ "true" ],
    "req_rows_ss" : [ "10" ],
    "req_q_ss" : [ "ipod" ],
    "req_defType_ss" : [ "edismax" ],
    "req_qf_ss" : [ "doc_id_s" ],
    "req_wt_ss" : [ "json" ],
    "req_facet_ss" : [ "true" ],
    "_version_" : 1530793784714985472
  } ]
}

```

Run "topN" report over collection "bb_catalog", return top-ranking query, search field "q_txt":

```

> curl -u user:pass -X POST -H 'Content-type: application/json' -d @- \
> http://localhost:8764/api/apollo/reports/bb_catalog/topN \
> <<EOF
> { "num" : 1, "field" : "q_txt" }
> EOF

[ {
  "key" : "ipad",
  "count" : 42,
  "percentage" : 0.7118644,
  "token" : "eyJmaWx0ZXJzIjpbInFfdHh0OmLwYWQiXX0="
} ]

```

Run "topClicked" report, return 5 most-clicked documents:

```

> curl -u user:pass -X POST -H 'Content-type: application/json' -d @- \
> http://localhost:8764/api/apollo/reports/bb_catalog/topClicked \
> <<EOF
> {"num":5}
> EOF

[ {
  "key" : "2842056",
  "count" : 42636,
  "percentage" : 0.0107869385,
  "token" : "eyJmaWx0ZXJzIjpbInR5cGVfczpjbjG1jayIsImRvY19pZF9z0jI4NDIwNTYiXX0="
}, {
  "key" : "1945531",
  "count" : 23510,
  "percentage" : 0.0059480467,
  "token" : "eyJmaWx0ZXJzIjpbInR5cGVfczpjbjG1jayIsImRvY19pZF9z0jE5NDU1MzEiXX0="
}, {
  "key" : "2842092",
  "count" : 22683,
  "percentage" : 0.0057388153,
  "token" : "eyJmaWx0ZXJzIjpbInR5cGVfczpjbjG1jayIsImRvY19pZF9z0jI4NDIwOTIiXX0="
}, {
  "key" : "9225377",
  "count" : 21603,
  "percentage" : 0.0054655746,
  "token" : "eyJmaWx0ZXJzIjpbInR5cGVfczpjbjG1jayIsImRvY19pZF9z0jkyMjUzNzc iXX0="
}, {
  "key" : "9755322",
  "count" : 20993,
  "percentage" : 0.005311244,
  "token" : "eyJmaWx0ZXJzIjpbInR5cGVfczpjbjG1jayIsImRvY19pZF9z0jk3NTUzMjIiXX0="
} ]

```

Get a histogram of the number of documents returned for queries over range 0 to 2000, interval 500 (4 bins):


```

> curl -u user:pass -X POST -H 'Content-type: application/json' -d @- \
> http://localhost:8764/api/apollo/reports/bb_catalog/histo \
> <<EOF
> {"field": "numdocs_l", "rangeStart": 0, "rangeEnd": 100, "interval": "25"}
> EOF

-X POST -H 'Content-type: application/json' -d @- http://localhost:8764/api/apollo/reports/bb_catalog/histo
<<EOF
> {"field": "numdocs_l", "rangeStart": 0, "rangeEnd": 2000, "interval": 500 }
> EOF
[ {
  "key" : "0",
  "count" : 10,
  "percentage" : 0.16949153,
  "token" : "eyJmaWx0ZXJzIjpbIm51bWRvY3NfbDpbMCBUTyA1MDB9I119"
}, {
  "key" : "500",
  "count" : 0,
  "percentage" : 0.0,
  "token" : "eyJmaWx0ZXJzIjpbIm51bWRvY3NfbDpbNTAwIFRPIDEwMDB9I119"
}, {
  "key" : "1000",
  "count" : 39,
  "percentage" : 0.66101694,
  "token" : "eyJmaWx0ZXJzIjpbIm51bWRvY3NfbDpbMTAwMCBUTyAxNTAwfSjdfQ=="
}, {
  "key" : "1500",
  "count" : 0,
  "percentage" : 0.0,
  "token" : "eyJmaWx0ZXJzIjpbIm51bWRvY3NfbDpbMTUwMCBUTyAyMDAwfSjdfQ=="
} ]

```

Get a date histogram for the last two days, with an interval of 1 day:

```

> curl -u user:pass -X POST -H Content-type:application/json -d @- \
> http://localhost:8764/api/apollo/reports/bb_catalog/dateHisto \
> <<EOF
> {"dateRangeStart": "NOW/DAY-1DAY", "dateRangeEnd": "NOW/DAY+1DAY", "timeInterval": "+1DAY"}
> EOF

[ {
  "key" : "2016-04-04T00:00:00Z",
  "count" : 0,
  "percentage" : 0.0,
  "token" :
  "eyJmaWx0ZXJzIjpbInRpbWVzdGFtcF9kdDpbTk9XL0RBWS0xREFZIFRPIE5PVy9EQVkrMURBWV0iLCJ0aW11c3RhbnBfZHQ6WzIwMTZcXC0wN
  FxcLTA0VDAwXW6MDBCXOWMFogVE8gMjAxNi0wNC0wNFQwMDowMDowMFoMURBWX0iXX0="
}, {
  "key" : "2016-04-05T00:00:00Z",
  "count" : 7,
  "percentage" : 1.0,
  "token" :
  "eyJmaWx0ZXJzIjpbInRpbWVzdGFtcF9kdDpbTk9XL0RBWS0xREFZIFRPIE5PVy9EQVkrMURBWV0iLCJ0aW11c3RhbnBfZHQ6WzIwMTZcXC0wN
  FxcLTA1VDAwXW6MDBCXOWMFogVE8gMjAxNi0wNC0wNFQwMDowMDowMFoMURBWX0iXX0="
} ]

```

10.24. Scheduler API

The scheduler REST API is used to define a schedule for system activities and manage the jobs that result from the schedule.

Note	As of Fusion 3.1, this API is deprecated in favor of the Jobs API.
------	--

All of the Fusion services are available for scheduling, as are any Solr activities and any other HTTP-based URI.

See Schedules for more information about all of the options available for scheduling, and instructions for configuring scheduler jobs using the Fusion UI.

10.24.1. Schedule Definition Properties

Parameter	Description
id <i>Required</i>	The ID of the schedule. This is required when creating a schedule with a POST request.
creatorType <i>Optional</i>	The type of user that created the schedule. If you have resources creating schedules programmatically, you could define a type that helps distinguish those schedules from others created by people.
creatorId <i>Optional</i>	An ID for the user that created the schedule.
startTime <i>Optional</i>	A time when the schedule should be started. If this is not set, it will be set to the date/time when the schedule was created.
endTime <i>Optional</i>	A time when the schedule should be terminated, i.e., when it will not run any more in the future. If not set, the schedule will run until it is disabled (i.e., "active" is set to false) or deleted.
repeatUnit <i>Optional</i>	A unit of time that when combined with the interval property defines how often the schedule will run. Allowed values are: <ul style="list-style-type: none">• "millisecond" or "ms"• "second" or "sec"• "minute" or "min"• "hour" or "hr"• "day"• "week"• "month" These time units are not case sensitive.

Parameter	Description
interval <i>Optional</i>	An integer that when combined with the repeatUnit property defines how often the schedule will run. If this is not set, or set to a number lower than '1' (i.e., '0'), the schedule will only be run once.
active <i>Optional</i>	If true , the schedule will be executed according to the defined interval. If false , the schedule will be disabled and will not run at the defined time.
callParams <i>Required</i>	<p>The callParams define the API call to execute at the defined time intervals. It allows several properties:</p> <ul style="list-style-type: none"> • uri: a fully-qualified service URI. This can be an HTTP call, a Solr request, or a Fusion service call. Supported URI schemas are: <ul style="list-style-type: none"> ◦ <code>http://</code> or <code>https://</code> ◦ <code>solr://{collection}/...</code> <p>A SolrCloud request.</p> ◦ <code>service://{serviceName}/{path}</code> <p>A load-balanced Fusion service request.</p> • method: The method to use, as in GET, POST, PUT, or DELETE. • queryParams: query parameters to be passed with the request. For Solr calls, this could be parameters such as "q", "fq", "commit", etc. • headers: any additional protocol headers, such as "Content-type". • entity: an optional payload to be sent with the request.

10.24.2. Examples

Run a crawl of the datasource "LocalDocs2" every day:

REQUEST

```
curl -u user:pass -X POST -H 'Content-Type: application/json' -d '{"id":"1", "creatorType":"human", "creatorId":"admin1", "repeatUnit":"DAY", "interval":1, "active":true, "callParams":{"uri":"service://connectors/jobs/LocalDocs2", "method":"POST"}}'
http://localhost:8764/api/apollo/scheduler/schedules
```

RESPONSE

```

{
  "id" : "1",
  "trigger" : "1_1400007488512_-9223372036854775808",
  "schedule" : {
    "id" : "1",
    "creatorType" : "human",
    "creatorId" : "admin1",
    "createTime" : "2014-05-13T18:58:08.512Z",
    "startTime" : "2014-05-13T18:58:08.512Z",
    "endTime" : null,
    "repeatUnit" : "DAY",
    "interval" : 1,
    "active" : true,
    "callParams" : {
      "uri" : "service://connectors/jobs/LocalDocs2",
      "method" : "POST",
      "queryParams" : { },
      "headers" : { },
      "entity" : null
    }
  }
}

```

Issue a commit to the SolrCloud collection "demo" every hour:

REQUEST

```

curl-u user:pass -X POST -H 'Content-Type: application/json' -d '{"id":"2", "creatorType":"human",
"creatorId":"admin1", "repeatUnit":"HOUR", "interval":1, "active":true,
"callParams":{"uri":"solr://demo/update", "method":"GET", "queryParams":{"stream.body":"<commit/>"}}}'
http://localhost:8764/api/apollo/scheduler/schedules

```

RESPONSE

```
{
  "id" : "2",
  "trigger" : "2_1400011854443_-9223372036854775807",
  "schedule" : {
    "id" : "2",
    "creatorType" : "human",
    "creatorId" : "admin1",
    "createTime" : "2014-05-13T20:10:54.443Z",
    "startTime" : "2014-05-13T20:10:54.443Z",
    "endTime" : null,
    "repeatUnit" : "HOUR",
    "interval" : 1,
    "active" : true,
    "callParams" : {
      "uri" : "solr://demo/update",
      "method" : "GET",
      "queryParams" : {
        "stream.body" : "<commit/>"
      },
      "headers" : { },
      "entity" : null
    }
  }
}
```

Set schedule "2" to inactive:

REQUEST

```
curl -u user:pass -X PUT -H 'Content-Type: application/json' -d '{"creatorType":"human", "creatorId":"admin1", "repeatUnit":"HOUR", "interval":1, "active":false, "callParams":{"uri":"solr://demo/update", "method":"GET", "queryParams":{"stream.body":"<commit/>"}}}' http://localhost:8764/api/apollo/scheduler/schedules/2
```

There will be no response if the PUT request was successful.

List all scheduled jobs:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/scheduler/jobs
```

RESPONSE

```
[ {
  "id" : "2",
  "trigger" : "2_1400011854443_-9223372036854775807",
  "schedule" : {
    "id" : "2",
    "creatorType" : "human",
    "creatorId" : "admin1",
    "createTime" : "2014-05-13T20:10:54.443Z",
    "startTime" : "2014-05-13T20:10:54.443Z",
    "endTime" : null,
    "repeatUnit" : "HOUR",
    "interval" : 1,
    "active" : true,
    "callParams" : {
      "uri" : "solr://demo/update",
      "method" : "GET",
      "queryParams" : {
        "stream.body" : "<commit/>"
      },
      "headers" : { },
      "entity" : null
    }
  }
}, {
  "id" : "1",
  "trigger" : "1_1400007488512_-9223372036854775808",
  "schedule" : {
    "id" : "1",
    "creatorType" : "human",
    "creatorId" : "admin1",
    "createTime" : "2014-05-13T18:58:08.512Z",
    "startTime" : "2014-05-13T18:58:08.512Z",
    "endTime" : null,
    "repeatUnit" : "DAY",
    "interval" : 1,
    "active" : true,
    "callParams" : {
      "uri" : "service://connectors/jobs/LocalDocs2",
      "method" : "POST",
      "queryParams" : { },
      "headers" : { },
      "entity" : null
    }
  }
} ]
```

List the history of job ID '1':

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/history/scheduler/items/1
```

RESPONSE

```

{
  "events" : [ {
    "start" : "2014-05-16T15:34:49.008Z",
    "end" : "2014-05-16T15:34:49.435Z",
    "source" : "scheduler",
    "type" : "execute",
    "status" : "ok",
    "details" : {
      "status" : 200,
      "entity" : "{\n \"id\" : \"TwitterSearch\", \n \"dataSourceId\" : \"TwitterSearch\", \n \"state\" :
\"RUNNING\", \n \"message\" : null, \n \"startTime\" : 1400254489000, \n \"endTime\" : -1, \n \"finished\" :
false, \n \"counters\" : { }, \n \"exception\" : null, \n \"running\" : true\n}"
    },
    "error" : null
  }, {
    "start" : "2014-05-16T15:38:32.536Z",
    "end" : "2014-05-16T15:38:32.559Z",
    "source" : "scheduler",
    "type" : "execute",
    "status" : "ok",
    "details" : {
      "status" : 200,
      "entity" : "{\n \"id\" : \"TwitterSearch\", \n \"dataSourceId\" : \"TwitterSearch\", \n \"state\" :
\"RUNNING\", \n \"message\" : null, \n \"startTime\" : 1400254712000, \n \"endTime\" : -1, \n \"finished\" :
false, \n \"counters\" : { }, \n \"exception\" : null, \n \"running\" : true\n}"
    }
  ]
}

```

10.25. Search Cluster API

The search cluster API allows users to connect Fusion with any existing Solr instances that are already present.

Once the Solr cluster is registered with Fusion, requests can be proxied through Fusion to it. The possible requests include search requests, but they can also be content indexing requests, such as the content crawled with a connector.

Once the searchCluster has been configured, the user can create Fusion collections that refer to the Solr collections that have been previously defined.

10.25.1. Search Cluster Definition Properties

Property	Description
id <i>Required</i>	The ID of the search cluster. This is only required when creating a new cluster definition with a POST request.
connectString <i>Required</i>	The string to use to connect to the existing Solr cluster or standalone instance. If the existing Solr is running in SolrCloud mode, use the connect string for the ZooKeeper ensemble. If the existing Solr is running as a standalone instance, use the full URL for the Solr instance.
cloud <i>Required</i>	Defines if the "cluster" being defined is a SolrCloud cluster (true) or a standalone Solr instance (false).
bufferCommitWithin <i>Optional</i>	Defines a commitWithin property for the buffer when writing to this cluster. If not defined, the system will default to 10,000 milliseconds.
bufferFlushInterval <i>Optional</i>	Defines how often to flush the buffer when writing to this cluster. If not defined, the system will default to 1000 milliseconds.
bufferSize <i>Optional</i>	Defines the size of the buffer. If not defined, the system will default to 100 items in the buffer.
concurrency <i>Optional</i>	Defines the maximum number of concurrent /parallel requests to Solr servers when Fusion index pipeline Solr Indexer stage has property bufferDocsForSolr set to true.
zkClientTimeout <i>Optional</i>	The maximum amount of time to wait when communicating with the ZooKeeper ensemble for a SolrCloud instance.
zkConnectTimeout <i>Optional</i>	The maximum amount of time to wait when trying to connect to the ZooKeeper ensemble for a SolrCloud instance.

10.25.2. Examples

Create a new search cluster that is an existing SolrCloud cluster:

REQUEST

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id":"mySolrCluster",
"connectString":"10.0.1.6:5001,10.0.1.6:5002,10.0.1.6:5003", "cloud":true}'
http://localhost:8764/api/apollo/searchCluster
```

RESPONSE

```
{
  "id" : "mySolrCluster",
  "connectString" : "10.0.1.6:5001,10.0.1.6:5002,10.0.1.6:5003",
  "cloud" : true,
}
```

Create a 'cluster' that is a standalone Solr instance:

REQUEST

```
curl -u user:pass -X POST -H 'Content-type: application/json' -d '{"id":"myOtherSolrCluster",
"connectString":"http://localhost:8983/solr", "cloud":false}' http://localhost:8764/api/apollo/searchCluster
```

RESPONSE

```
{
  "id" : "myOtherSolrCluster",
  "connectString" : "http://localhost:8983/solr",
  "cloud" : false,
}
```

Show the status of each node of 'mySolrCluster':

REQUEST

```
curl http://localhost:8764/api/apollo/searchCluster/mySolrCluster/nodes
```

RESPONSE

```
[ {
  "name" : "10.0.1.11:7574_solr",
  "baseUrl" : "http://10.0.1.11:7574/solr",
  "state" : "active"
}, {
  "name" : "10.0.1.8:7574_solr",
  "baseUrl" : "http://10.0.1.8:7574/solr",
  "state" : "active"
} ]
```

Show the system information for one named node:

REQUEST

```
curl http://10.0.1.8:8764/api/apollo/searchCluster/mySolrCluster/systemInfo?nodeName=10.0.1.8:7574_solr
```

RESPONSE

```
{
  "10.0.1.8:7574_solr" : {
    "mode" : "solrcloud",
    "lucene" : {
      "solr-spec-version" : "4.8.0",
      "lucene-spec-version" : "4.8.0"
    },
    "jvm" : {
      "version" : "1.8.0_121 25.121-b13",
      "name" : "Oracle Corporation Java HotSpot(TM) 64-Bit Server VM",
      "processors" : 4,
      "memory" : {
        "raw" : {
          "free" : 66736272,
          "total" : 204800000,
          "max" : 204800000,
          "used" : 138063728,
          "used%" : 67.4139296875
        }
      }
    },
    "system" : {
      "name" : "Mac OS X",
      "version" : "10.9.3",
      "arch" : "x86_64",
      "systemLoadAverage" : 2.130859375,
      "committedVirtualMemorySize" : 2963378176,
      "freePhysicalMemorySize" : 9321914368,
      "freeSwapSpaceSize" : 1073741824,
      "processCpuTime" : 313176000000,
      "totalPhysicalMemorySize" : 17179869184,
      "totalSwapSpaceSize" : 1073741824,
      "openFileDescriptorCount" : 208,
      "maxFileDescriptorCount" : 10240,
      "uname" : "Darwin MacMini.local 13.2.0 Darwin Kernel Version 13.2.0: Thu Apr 17 23:03:13 PDT 2014;
root:xnu-2422.100.13~1/RELEASE_X86_64 x86_64\n",
      "uptime" : "15:48 up 3 days, 7:08, 7 users, load averages: 2.13 2.01 1.91\n"
    }
  }
}
```

10.26. Signals API

The Signals API accepts a set of signals, encoded as JSON objects, for indexing into a signals collection.

Normally, signals are indexed just like ordinary documents, through a configured datasource and index pipeline. This API is provided for cases where it is more convenient to index signals directly.

To aggregate the signals, see the Signals Aggregator API.

10.26.1. Signal document structure

A raw signal is stored as a Solr document with the following fields, which are derived from the raw signal as follows:

Field	Description
<code>id</code> <i>Optional</i>	The signal ID. If no ID is supplied, one will be automatically generated.
<code>type</code> <i>Required</i>	<p>The signal type that is being sent. This value is used during aggregation to filter events of the same type. Types can be mixed in aggregation jobs, if needed.</p> <p>The type can consist of any string you choose. For consistency, always send events of the same type with the same type value.</p> <p>During indexing, type values will be moved to a field named <code>type_s</code>.</p>

Field	Description
<p><code>params</code> <i>Optional</i></p>	<p>The params allow flexible definition of the fields you care about and will use later for signal aggregation:</p> <ul style="list-style-type: none"> • <code>docId</code> – A unique document ID This is stored in the Solr raw signal document as field <code>doc_id_s</code>. • <code>userId</code> – A unique user ID This is stored in the Solr raw signal document as field <code>user_id_s</code>. • <code>query</code> – A query string; for example, a user’s search This is copied to the Solr raw signal document as both fields <code>query_s</code> and <code>query_t</code>. Some cleanup occurs to convert the string to lowercase, decode URL encoding, and replace white space with single space characters. The original query is saved in field <code>query_orig_s</code>. • <code>filterQueries</code> – A list of strings, such as filters on the search query This is copied to the Solr raw signal document as both <code>filters_s</code> and <code>filters_orig_ss</code>. • <code>collection</code> – The primary collection name • <code>weight</code> – A float value representing the relative weight of this signal This is saved in the field <code>weight_d</code>. • <code>count</code> – A positive integer value representing the incremented count of signals This is saved in the field <code>count_i</code>.
<p><code>timestamp</code></p>	<p>The timestamp of the signal event.</p> <ul style="list-style-type: none"> • When using the Signals API, this property is optional; it defaults to the current server time. • When using the Signal Formatter index stage, one of the following fields must be present: <code>timestamp</code>, <code>timestamp_tdt</code>, <code>timestamp_dt</code>, or <code>epoch</code>.

Here is the JSON representation of one click signal, taken from an example dataset of synthetic clickstream data:

```
{ "params": {
  "docId": "2125233",
  "filterQueries": ["cat00000","abcat0100000", "abcat0101000", "abcat0101001"],
  "query": "Televisiones Panasonic 50 pulgadas" }
"type":"click",
"timestamp": "2011-09-01T23:44:52.533000Z",
}
```

10.26.2. Examples

Send two signal events to record user clicks:

REQUEST

```
curl \
-u user:pass -X POST -H 'Content-type:application/json' -d @- \
http://localhost:8764/api/apollo/signals/docs?commit=true \
<<EOF
[
{"params": {
  "query": "Televisiones Panasonic 50 pulgadas",
  "filterQueries": ["cat00000","abcat0100000", "abcat0101000", "abcat0101001"],
  "docId": "2125233" },
"type":"click",
"timestamp": "2011-09-01T23:44:52.533000Z"
},
{"params": {
  "query": "Sharp",
  "filterQueries": ["cat00000", "abcat0100000", "abcat0101000", "abcat0101001"],
  "docId": "2009324" },
"type":"click",
"timestamp": "2011-09-05T12:25:37.420000Z"
}
]
EOF
```

A successful request results in events being added to the signals collection. For the above example, the events will be represented as follows:

```
{
  "responseHeader":{
    "status":0,
    "QTime":1,
    "params":{
      "indent":"true",
      "q":"doc_id_s: 2125233",
      "wt":"json"}
  },
  "response":{"numFound":1198,"start":0,"docs":[
    {
      "id": "7aee7b1f-5cde-4957-b73c-c15881f559ec",
      "filters_s": "abcat0100000 $ abcat0101000 $ abcat0101001 $ cat00000",
      "query_orig_s": "Televisiones Panasonic 50 pulgadas",
      "params.user_s": "000000df17cd56a5df4a94074e133c9d4739fae3",
```

```

"doc_id_s": "2125233",
"params.query_time__s": "2011-09-01T23:43:59.752Z",
"query_t": "televisiones panasonic 50 pulgadas",
"query_s": "televisiones panasonic 50 pulgadas",
"filters_orig_ss": [
  "abcat0100000",
  "abcat0101000",
  "abcat0101001",
  "cat00000"
],
"flag_s": "EVENT",
"type_s": "click",
"attr_params.filterQueries_": [
  "cat00000",
  "abcat0100000",
  "abcat0101000",
  "abcat0101001"
],
"timestamp_dt": "2011-09-01T23:44:52.533Z",
"_version_": 1478892846857584600
},
{
  "id": "6789a209-f5b5-457e-9df6-8033b8f7f317",
  "filters_s": "abcat0100000 $ abcat0101000 $ abcat0101001 $ cat00000",
  "query_orig_s": "Sharp",
  "params.user_s": "000001928162247ffaf63185cd8b2a244c78e7c6",
  "doc_id_s": "2009324",
  "params.query_time__s": "2011-09-05T12:25:01.187Z",
  "query_t": "sharp",
  "query_s": "sharp",
  "filters_orig_ss": [
    "abcat0100000",
    "abcat0101000",
    "abcat0101001",
    "cat00000"
  ],
  "flag_s": "EVENT",
  "type_s": "click",
  "attr_params.filterQueries_": [
    "cat00000",
    "abcat0100000",
    "abcat0101000",
    "abcat0101001"
  ],
  "timestamp_dt": "2011-09-05T12:25:37.42Z",
  "_version_": 1478892846859681800
}
]
}
}

```

10.27. Signals Aggregator API

The Signals Aggregator API is used to aggregate signal events, which allows faster querying for recommendations. To use recommendations, signals need to be recorded and then aggregated.

When signals are enabled for a collection, two system-level collections are created. The first is named `collection_signals`, where `collection` is the sibling collection name, and signal events are indexed to this collection. The second is named `collection_signals_agg`, and is the default location for aggregated signal events. See Signals API for more information on how to index signal events.

The aggregation process creates tuples for the fields selected when creating the aggregator job. A default tuple is applied if none is specified.

The aggregation process can remove the raw signals if desired, or keep them for other aggregation jobs.

10.27.1. Signals Aggregator Definitions Properties

Parameter	Description
<code>id</code> <i>Optional</i>	A unique identifier for this aggregator job.
<code>groupingFields</code> <i>Optional</i>	The fields that define unique tuples. The fields list is defined as a JSON array, with commas between each field name. If a set of fields is not defined, then a default tuple <code>'doc_id_s', 'query_s', 'filters_s'</code> will be used.
<code>signalTypes</code> <i>Optional</i>	The types of signals to aggregate. The type list is defined as a JSON array, with commas between each type. The types must be existing types used for events in your signals collection.
<code>aggregator</code> <i>Optional</i>	The name of the aggregator implementation. If it is not defined, this will default to click , which is an implementation optimized for aggregating signals based on user clicks. Aggregated records from this implementation will include a <code>'weight_d'</code> field which can be used in boosting clicked documents. If you are not aggregating user click events, you can choose simple . This implementation does not add a <code>'weight_d'</code> field to each record. A third option is special is described in more detail in page Aggregator Scripting.

Parameter	Description
selectQuery <i>Optional</i>	Any query to identify signal events.
timeRange <i>Optional</i>	A valid range query to select events to aggregate.
sort <i>Optional</i>	<p>Specifies ordering of raw signal events within an aggregation.</p> <p>The default ordering is by event id ("id asc"). It can be set to use other fields using the standard Solr sort expressions, e.g. "timestamp_dt asc", also multiple criteria separate by comma, e.g. "type_s asc,timestamp_dt desc".</p> <p>Note: the sorting by "id asc" is always appended as the last sort criteria in order to break ties.</p>
outputPipeline <i>Optional</i>	The name of a pipeline to use for processing aggregating events.
outputCollection <i>Optional</i>	The collection in which to store the aggregated events.
rollupPipeline <i>Optional</i>	The pipeline to use for rollups.
rollupAggregator <i>Optional</i>	The name of the aggregator implementation to use for rollups.
sourceRemove <i>Optional</i>	<p>If true, then signal events that have been aggregated will be removed from the index.</p> <p>The default is false.</p>
sourceCatchup <i>Optional</i>	<p>If true, the original time range of the aggregation will be modified to span only the period since the last successful aggregation.</p> <p>The default is false.</p>
outputRollup <i>Optional</i>	If true , the default, after performing the source data aggregation an additional aggregation step will be executed to roll-up the new aggregates with old aggregates that exist in the output collection for the same aggregation type.

Parameter	Description
aggregates <i>Optional</i>	A list of aggregation functions. Since it's possible to pass side-effects from one function to a later function in the list, the functions should be declared in the desired order of execution. The available aggregator functions are described in more detail in the section Aggregator Functions.
params <i>Optional</i>	The params allows defining aggregation job parameters. The most common use of this property is to define JavaScript scripts to customize the aggregator behavior. See the section Aggregator Scripting for more details.

Note that for large aggregation definitions, you could create a .json formatted file with the desired properties and upload it with cURL's `-d` parameter.

No output is returned when creating or updating an aggregator job.

When a job is listed, the properties returned are the same as the possible properties when defining a job.

10.27.2. Examples

Create an aggregator job for the click type of signals, with an aggregate function to provides counts by the id field:

REQUEST

```
curl -u user:pass -X POST -H 'Content-Type: application/json' -d '{"id": "1", "signalTypes": ["click"], "aggregates": [{"type": "count", "sourceFields": ["id"], "targetField": "count_d"}]}'
http://localhost:8764/api/apollo/aggregator/aggregations
```

RESPONSE

None.

Update the properties for aggregator job '1', including all the original properties plus the ones we want to add or change:

REQUEST

```
curl -u user:pass -X PUT -H 'Content-Type: application/json' -d '{"signalTypes": ["click"], "timeRange": "[NOW/-1 TO NOW]", "aggregates": [{"type": "count", "sourceFields": ["id"], "targetField": "count_d"}]}'
http://localhost:8764/api/apollo/aggregator/aggregations/1
```

RESPONSE

None.

List the properties for aggregator job '1':

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/aggregator/aggregations/1
```

RESPONSE

```
{
  "id" : "1",
  "groupingFields" : [ ],
  "signalTypes" : [ "click" ],
  "timeRange" : "[NOW/-1 TO NOW]",
  "sourceRemove" : false,
  "sourceCatchup" : false,
  "outputRollup" : false,
  "aggregates" : [ {
    "type" : "count",
    "sourceFields" : [ "id" ],
    "targetField" : "count_d",
    "params" : { }
  } ],
  "params" : { }
}
```

Start job '1' on the 'demo_signals' collection:

REQUEST

```
curl -u user:pass -X POST http://localhost:8764/api/apollo/aggregator/jobs/demo_signals/1
```

RESPONSE

The following output has been truncated to omit the aggregation job definition and only shows the other job properties that are returned on start.

```
{
  "signals" : {
    "types" : [ "click" ],
    "stats" : { }
  },
  "state" : "running",
  "job_id" : "4d69ec73358b41d38caf1eb3b378809e",
  "aggregation_time_date" : "2014-09-11T16:39:58.347Z",
  "aggregation" : {
    "id" : "r1",
    "groupingFields" : [ "doc_id_s", "query_s", "filters_s" ],
    "signalTypes" : [ "click" ],
    "selectQuery" : "*:*",
    ...
  }
  "output_collection" : "bestbuy_signals_aggr",
  "NOW" : 1410453598347,
  "NOW_date" : "2014-09-11T16:39:58.347Z",
  "collection" : "bestbuy_signals",
  "aggregation_time" : 1410453598347,
  "compound_id" : "bestbuy_signals:r1"
}
```

See the list of aggregator job items:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/history/aggregator/items
```

RESPONSE

```
[ "demo_signals:1" ]
```

Get the history of job "demo_signals:1":

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/history/aggregator/items/demo_signals:1
```

RESPONSE

```

{
  "events" : [ {
    "start" : "2014-04-16T20:45:16.582Z",
    "end" : "2014-04-16T20:45:16.781Z",
    "source" : "demo_signals:1",
    "type" : "run",
    "status" : "ok",
    "details" : {
      "signals" : {
        "click" : {
          "state" : "finished",
          "raw" : 2,
          "aggr_type_s" : "click",
          "aggr_class" : "com.lucidworks.apollo.service.aggregation.ClickSignalAggregator",
          "aggregated" : 2
        }
      }
    },
    "state" : "finished",
    "job_id" : "467bc0db-a9c9-4b48-8080-439958818907",
    "aggregation_time_date" : "2014-04-16T20:45:16.556Z",
    "aggregation" : {
      "id" : "1",
      "fields" : [ "doc_id_s", "query_s", "filters_s" ],
      "types" : [ "click" ],
      "select" : ":*:*",
      "range" : "[* TO NOW]",
      "remove" : false,
      "rolling" : false,
      "params" : { },
      "anyAggr" : false
    },
    "NOW" : 1397681116556,
    "commit" : "done",
    "NOW_date" : "2014-04-16T20:45:16.556Z",
    "collection" : "demo_signals",
    "aggregation_time" : 1397681116556,
    "compound_id" : "demo_signals:1"
  },
  "error" : null
} ]
}

```

10.28. Solr API

The Solr API is used to manage collection-level configurations.

The path for this request is:

```
/api/apollo/solr/<collectionName>/<solrRequest>
```

where *<collectionName>* is the name of an specific collection and *<solrRequest>* is the Solr command you wish to run.

Since this API proxies requests to Solr, each available method corresponds to the method in Solr. So, a GET request to Solr would use the GET method of this endpoint; a POST would use the POST method, etc.

Depending on the request, the response may consist of records that match a query or output from a Schema API request.

10.28.1. Examples

Query Solr for documents in the 'docs' collection containing the term 'solr', limiting the results to only 2 records, returning only the title, and in JSON format:

REQUEST

```
http://localhost:8764/api/apollo/solr/docs/select?q=solr&rows=2&fl=title&wt=json
```

RESPONSE

```
{
  "responseHeader": {
    "status": 0,
    "QTime": 2,
    "params": {
      "fl": "title",
      "q": "solr",
      "wt": "json",
      "rows": "2"
    }
  },
  "response": {
    "numFound": 52,
    "start": 0,
    "docs": [
      {
        "title": [
          "Solr and SolrAdmin APIs - Fusion Documentation - Lucidworks"
        ]
      },
      {
        "title": [
          "Search Clusters - Fusion Documentation - Lucidworks"
        ]
      }
    ]
  }
}
```

10.29. Solr Configuration API

The Solr Configuration REST API is used to access and manage the Solr configuration files stored in ZooKeeper.

To manage Solr synonym lists, use the Synonyms API.

To manage Solr lists of stopwords, use the Stopwords API.

10.29.1. Solr Configuration Definition Properties

Path Parameters

Parameter	Description
collection	The collection that contains the configuration files to list or view.
path	The path to a specific file or nested child nodes. If the file is not nested, the filename can be entered without any path information.

Query Parameters

Parameter	Description
expand	<p>If true, the binary content of a file will be returned base64 encoded. If this is not included, only the metadata about each node will be returned.</p> <p>If you would like to see the content of the file in plain text, you can add 'Accept: text/plain' to the request header. Alternately, you can get the raw bytes by adding 'Accept: application/octet-stream' to the header.</p>
recursive	If true, children of nested ZooKeeper nodes will be returned. This can be used in conjunction with the path to show only children of a specific node.

The output will include a list of each file with the following information:

- **name:** The name of the file or node.
- **version:** The file or node version from ZooKeeper.
- **isDir:** If the node has children, this value will be true.
- **value:** Only returned if "expand=true" is added as a query parameter. This will be returned as the base64 encoding of the contents of the file.

If the header includes 'Accept: text/plain' or 'Accept: application/octet-stream', metadata about the file will not be returned and only the content as either plain text or raw bytes.

If recursive=true and expand=true are both used in the same request, the request may be a little slow depending on how much data is requested.

10.29.2. ZooKeeper Collection Configuration Definition Properties

Path Parameters

Parameter	Description
collection	The collection that contains the configuration files to list or view.
path	The path to a specific file (including nested child nodes).

Query Parameters

Parameter	Description
reload	If true, the collection will be reloaded to make the changes available to Solr immediately.

Input Content

The REST API does not restrict the type of content that can be sent to ZooKeeper, but care should be taken to make sure the the format of the file remains compatible with Solr's requirements.

As a best practice, all requests should include the Content-type in the request header.

10.29.3. Examples

Show `_rest_managed.json` for the 'docs' collection, with the base64 encoded content:

INPUT

```
curl -u user:pass http://localhost:8764/api/apollo/collections/docs/solr-config/_rest_managed.json?expand=true
```

OUTPUT

```
{
  "name" : "_rest_managed.json",
  "version" : 0,
  "isDir" : false,
  "value" : "eyJpbm10QXJncyI6e30sIm1hbmFnZWRMaXN0IjpbXX0NCg=="
}
```

Get the plain text version of `synonyms.txt`:

INPUT

```
curl -u user:pass -H 'Accept: text/plain' http://localhost:8764/api/apollo/collections/docs/solr-config/synonyms.txt
```

OUTPUT

GB,gib,gigabyte,gigabytes
MB,mib,megabyte,megabytes
Television, Televisions, TV, TVs
pixima => pixma

Show the child nodes of 'clustering':

INPUT

```
curl -u user:pass http://localhost:8764/api/apollo/collections/docs/solr-config/clustering?recursive=true
```

OUTPUT

```
{
  "name" : "clustering",
  "version" : 0,
  "isDir" : true,
  "children" : [ {
    "name" : "carrot2",
    "parent" : "clustering",
    "version" : 0,
    "isDir" : true,
    "children" : [ {
      "name" : "kmeans-attributes.xml",
      "parent" : "clustering/carrot2",
      "version" : 0,
      "isDir" : false
    }, {
      "name" : "lingo-attributes.xml",
      "parent" : "clustering/carrot2",
      "version" : 0,
      "isDir" : false
    }, {
      "name" : "stc-attributes.xml",
      "parent" : "clustering/carrot2",
      "version" : 0,
      "isDir" : false
    } ]
  } ]
}
```

Replace the contents of the existing synonyms.txt file with the contents of a file named "updated-synonyms.txt" and reload the collection:

INPUT

```
curl -u user:pass -X PUT -H 'Content-type: text/plain' -d '@updated-synonyms.txt'
http://localhost:8764/api/apollo/collections/docs/solr-config/synonyms.txt?reload=true
```

OUTPUT

None.

10.30. SolrAdmin API

The Solr Admin API lets you send commands to Solr through Fusion's proxy service. This allows you to protect your Solr instances from outside connections, and apply roles and user permissions from Fusion when running Solr commands. Requests sent to this API are subject to access restrictions above the collection level.

Note that because one searchCluster may host several collections, it's not recommended to use this with a collection-level command (such as a query, or document update).

10.30.1. Example

Request all CORE MBeans for the 'default' searchCluster, formatted in JSON:

REQUEST

```
http://localhost:8764/api/apollo/solrAdmin/default/admin/mbeans?cat=CORE&wt=json
```

RESPONSE

```
{
  "responseHeader": {
    "status": 0,
    "QTime": 1
  },
  "solr-mbeans": [
    "CORE",
    {
      "searcher": {
        "class": "org.apache.solr.search.SolrIndexSearcher",
        "version": "1.0",
        "description": "index searcher",
        "src": null
      },
      "core": {
        "class": "collection1",
        "version": "1.0",
        "description": "SolrCore",
        "src": null
      },
      "Searcher@435cf502[collection1] main": {
        "class": "org.apache.solr.search.SolrIndexSearcher",
        "version": "1.0",
        "description": "index searcher",
        "src": null
      }
    }
  ]
}
```

10.31. Spark Jobs API

This is a set of endpoints for configuring and running Spark jobs.

10.31.1. Spark job subtypes

For the Spark job type, the available subtypes are listed below.

ALS Recommender	Train a collaborative filtering matrix decomposition recommender using SparkML's Alternating Least Squares (ALS) to batch-compute user recommendations and item similarities.
Aggregation	Define an aggregation job to be executed by Fusion Spark.
Co-occurrence Similarity	Compute a mutual-information item similarity model.
Random Forest Classifier Training	Train a random forest classifier for text classification.
Script	Run a custom Scala script as a Fusion Job.
Matrix Decomposition-Based Query-Query Similarity Job	Train a collaborative filtering matrix decomposition recommender using SparkML's Alternating Least Squares (ALS) to batch-compute query-query similarities.
Bisecting KMeans Clustering Job	Train a bisecting KMeans clustering model .
Logistic Regression Classifier Training Job	Train a regularized logistic regression model for text classification.
Item Similarity Recommender	Compute user recommendations based on pre-computed item similarity model.
Levenshtein	Compare the items in a collection and produces possible spelling mistakes based on the Levenshtein edit distance.
Collection Analysis	Produce statistics about the types of documents in a collection and their lengths.
Statistically Interesting Phrases (SIP)	Output statistically interesting phrases in a collection, that is, phrases that occur more frequently or less frequently than expected.
Doc Clustering	An end-to-end document clustering job that preprocesses documents, separates out extreme-length documents and other outliers, automatically selects the number of clusters, and extracts keyword labels for clusters. You can choose between Bisecting KMeans and KMeans clustering methods, and between TFIDF and word2vec vectorization methods.
Outlier Detection	Find groups of outliers for the entire set of documents in the collection.
Cluster Labeling	Attach keyword labels to documents that have already been assigned to groups.

10.31.2. Spark Configuration Properties

Fusion passes all configuration properties with prefix "spark." to the Spark master, Spark worker and each Spark application, both for aggregation jobs and custom-scripted processing.

These properties are stored in Fusion's ZooKeeper and can be updated via requests to Fusion endpoint `api/apollo/configurations` which will update the stored value without restarting the service, therefore existing jobs and SparkContexts will not be affected. The Fusion endpoint `api/apollo/configurations` returns all configured properties for that installation. You can examine spark default configurations in a Unix shell using the utilities `curl` and `grep`. Here is an example which checks a local Fusion installation running on port 8764:

```
curl -u username:password http://localhost:8764/api/apollo/configurations | grep '"spark.'
```

```
"spark.executor.memory" : "2g",  
"spark.task.maxFailures" : "10",  
"spark.worker.cleanup.appDataTtl" : "7200",  
"spark.worker.cleanup.enabled" : "true",  
"spark.worker.memory" : "2g",
```

The default SparkContext that Fusion uses for aggregation jobs can be assigned a fraction of cluster resources (executor memory and/or available CPU cores). This allows other applications (such as scripted jobs, or shell sessions) to use the remaining cluster resources even when some aggregation jobs are running. Fusion 2.3 also permits dynamic allocation for all applications. This can be overridden per application. In practice, this means that even when there's an already running SparkContext with a relatively long idle time (eg. 10 minutes) but there are no active jobs that use it, its resources (CPU cores and executor memory) will be released for use by other applications.

For scripted Spark jobs, users can specify per-job configuration overrides as a set of key / value pairs in a "sparkConfig" property element of a script job configuration, which takes precedence over values stored in ZooKeeper. The following is an example of a scripted job with a "sparkConfig" section:

```
{  
  "id": "scripted_job_example",  
  "script": "val rdd = sc.textFile(\"/foo.txt\")\nrdd.count\n",  
  "sparkConfig": {  
    "spark.cores.max": 2,  
    "spark.executor.memory": "1g"  
  }  
}
```

The following table lists those Spark configuration properties that Fusion overrides or uses in order to determine applications' resource allocations.

Property	Description
<code>spark.master.url</code>	By default, left unset. This property is only specified when using an external Spark cluster; when Fusion is using its own standalone Spark cluster, this property isn't set.
<code>spark.cores.max</code>	The maximum number of cores across the cluster assigned to the application. If not specified, there is no limit. The default is unset, i.e., an unlimited number of cores.

Property	Description
<code>spark.executor.memory</code>	Amount of memory assigned to each application's executor. The default is 2G.
<code>spark.scheduler.mode</code>	Controls how tasks are assigned to available resources. Can be either 'FIFO' or 'FAIR'. Default value is 'FAIR'.
<code>spark.dynamicAllocation.enabled</code>	Boolean - whether or not to enable dynamic allocation of executors. Default value is 'TRUE'.
<code>spark.shuffle.service.enabled</code>	Boolean - whether or not to enable internal shuffle service for standalone Spark cluster. Default value is 'TRUE'.
<code>spark.dynamicAllocation.executorIdleTimeout</code>	Number of seconds after which idle executors are removed. Default value is '60s'.
<code>spark.dynamicAllocation.minExecutors</code>	Number of executors to leave running even when idle. Default value is 0.
<code>spark.eventLog.enabled</code>	Boolean - whether or not event log is enabled. Event log stores job details and can be accessed after application finishes. Default value is 'TRUE'.
<code>spark.eventLog.dir</code>	Directory which stores event logs. Default location is <code>fusion/3.1.x/var/spark-eventlog</code> .
<code>spark.eventLog.compress</code>	Boolean - whether or not to compress event log data. Default value is 'TRUE'.
<code>spark.logConf</code>	Boolean - whether or not to log effective SparkConf of new SparkContext-s. Default value is 'TRUE'.
<code>spark.deploy.recoveryMode</code>	Default value is 'ZOOKEEPER'
<code>spark.deploy.zookeeper.url</code>	ZooKeeper connect string. Default value is <code>\$FUSION_ZK</code>
<code>spark.deploy.zookeeper.dir</code>	ZooKeeper path, default value is <code>/lucid/spark</code>
<code>spark.worker.cleanup.enabled</code>	Boolean - whether or not to periodically cleanup worker data. Default value is 'TRUE'.
<code>spark.worker.cleanup.appDataTtl</code>	Time-to-live in seconds. Default value is 86400 (24h).
<code>spark.deploy.retainedApplications</code>	The maximum number of applications to show in the UI. Default value is 50.
<code>spark.deploy.retainedDrivers</code>	The maximum number of drivers. Default value is 50.
<code>spark.worker.timeout</code>	The maximum timeout in seconds allowed before a worker is considered lost. The default value is 30.
<code>spark.worker.memory</code>	The maximum total heap allocated to all executors running on this worker. Defaults to value of the executor memory heap.

10.31.3. Fusion Configuration Properties

Property	Description
<code>fusion.spark.master.port</code>	Spark master job submission port. Default value is 8766.

Property	Description
<code>fusion.spark.master.ui.port</code>	Spark master UI port. Default value is 8767.
<code>fusion.spark.idleTime</code>	Maximum idle time in seconds, after which the application (ie. SparkContext) is shut down. Default value is 300.
<code>fusion.spark.executor.memory.min</code>	Minimum executor memory in MB. Default value 450Mb, which is sufficient to let Fusion components in application task's to initialize themselves
<code>fusion.spark.executor.memory.fraction</code>	A float number in range $(0.0, 1.0]$ indicating what portion of <code>spark.executor.memory</code> to allocate to this application. Default value is 1.0.
<code>fusion.spark.cores.fraction</code>	A float number in range $(0.0, 1.0]$ indicating what portion of <code>spark.cores.max</code> to allocate to this application. Default value is 1.0.

10.32. Stopwords API

Note	These endpoints have been deprecated since Fusion 2.4.
------	--

10.32.1. Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Update the stop words list with a new list of stop words:

REQUEST

```
curl -u user:pass -X PUT -H 'Content-type: application/json' -d '["a", "and", "of", "the"]'
http://localhost:8764/api/apollo/stopwords/docs
```

RESPONSE

None.

List the current stop words list:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/stopwords/docs
```

RESPONSE

```
[ "a", "and", "of", "the" ]
```

Upload a stop words file:

REQUEST

```
curl -u user:pass -X PUT --form file=@stopwords.txt http://localhost:8764/api/apollo/stopwords/docs
```

RESPONSE

None.

Download a stop words list & save it as a file named "stopwords.txt":

REQUEST

```
curl -H "Accept: application/octet-stream" http://localhost:8764/api/apollo/stopwords/docs > stopwords.txt
```

RESPONSE

None.

10.33. Synonyms API

Note	These endpoints have been deprecated since Fusion 2.4. Use the Synonyms Editor API instead.
------	---

The Synonyms API manages the set of synonyms defined in Solr for a collection:

- a string of terms that will expand on the terms the user entered, like **Television**, **TV**.
- a term that should be mapped to another term, like **i-pod**⇒**ipod**.

When updating the synonyms, note that only PUT requests are supported, and any new data sent will overwrite the previous synonyms. As such, PUT requests can be seen as replacement requests.

It is only possible to have a single set of stop words when using this REST API. If you need different sets of stop words for different field types (perhaps for different languages), you will need to edit the schema.xml and manually manage the stop word files.

10.33.1. Examples

Note	Use port 8765 in local development environments only. In production, use port 8764.
------	---

Update the synonym list with a new list of synonyms for a collection named 'docs':

REQUEST

```
curl -u user:pass -X PUT -H 'Content-type: application/json' -d '{"match": [ "GB", "gib", "gigabyte", "gigabytes" ]}, {"match": [ "MB", "mib", "megabyte", "megabytes" ]}, {"match": [ "Television", "Televisions", "TV", "TVs" ]}, {"match": [ "foo", "replace": [ "bar" ]}]}'  
http://localhost:8764/api/apollo/synonyms/docs
```

RESPONSE

None.

List the current synonyms list:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/synonyms/docs
```

RESPONSE

```
[ {
  "match" : [ "GB", "gib", "gigabyte", "gigabytes" ]
}, {
  "match" : [ "MB", "mib", "megabyte", "megabytes" ]
}, {
  "match" : [ "Television", "Televisions", "TV", "TVs" ]
}, {
  "match" : [ "foo" ],
  "replace" : [ "bar" ]
}, {
  "match" : [ "i-pod" ],
  "replace" : [ "ipod" ]
} ]
```

Upload a synonyms list:

REQUEST

```
curl -u user:pass -X PUT --form file=@synonyms.txt http://localhost:8764/api/apollo/synonyms/docs
```

RESPONSE

None.

Download a synonyms list, and save it as a file named "synonyms.txt":

REQUEST

```
curl -u user:pass -H "Accept: application/octet-stream" http://localhost:8764/api/apollo/synonyms/docs >
synonyms.txt
```

RESPONSE

None.

10.34. System Admin APIs

The System Admin APIs are REST API endpoints that can be used by a system administrator to configure various Fusion processes. The REST APIs are:

- Configurations API, to check and modify global settings in Fusion.
- History API, to retrieve historic events of each service in the system.
- Nodes API, to check the status of each node in the search cluster.
- System API, to retrieve system metrics, ping the system for life, flush buffers, or view active threads.
- Usage API, to retrieve system information that is shared with Lucidworks about the size of your system.

10.34.1. Configurations API

The Configurations API allows setting global properties for Fusion. Some settings are not set by any configuration file but are reported as settings from the operating system. Those settings cannot be changed with this API.

Examples

Show the configuration items that include the pattern 'zk-connect', with verbose enabled:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/configurations?verbose=true&pattern=zk-connect
```

RESPONSE

```
{
  "com.lucidworks.apollo-admin.config/zk-connect" : [ {
    "value" : "localhost:9983",
    "location" : "system properties"
  } ]
}
```

Get the configuration items from the Connectors JVM that include the term 'connect':

REQUEST

```
curl -u user:pass http://localhost:8984/connectors/api/v1/configurations?pattern=connect
```

Get the configuration items from the Fusion JVM that start with 'com.lucidworks':

REQUEST

```
http://localhost:8764/api/apollo/configurations?prefix=com.lucidworks
```

Change the default allowed recommendation types to include 'itemsForItem':

REQUEST

```
curl -u user:pass -X PUT -H 'Content-type: application/json' -d '{"itemsForQuery", "termsForDocument", "itemsforItem"}'
http://localhost:8764/api/apollo/configurations/com.lucidworks.apollo.service.recommend.allowed.types
```

RESPONSE

None. Check the setting again to confirm the changes.

10.34.2. History API

Fusion stores history for each running service within the system. Usually this is used to log start and stop events for a service. However, the scheduler uses the history to store the results of scheduled tasks. For more information on schedule history, see the section on Schedules.

The History API provides information about the services that are running. The list of these services is provided by Introspect API, which is described in the REST API Reference.

Examples

View the history of the index-pipelines service:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/history/index-pipelines::v1
```

RESPONSE

```
{
  "events" : [ {
    "start" : "2014-05-16T14:11:48.849Z",
    "end" : "2014-05-16T14:11:48.849Z",
    "source" : "index-pipelines::v1",
    "type" : "start",
    "status" : "ok",
    "details" : null,
    "error" : null
  }, {
    "start" : "2014-05-16T14:12:48.845Z",
    "end" : "2014-05-16T14:12:48.845Z",
    "source" : "index-pipelines::v1",
    "type" : "start",
    "status" : "ok",
    "details" : null,
    "error" : null
  }
]
```

View items in the scheduler history:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/history/scheduler/items/1
```

RESPONSE

```

{
  "events" : [ {
    "start" : "2014-05-16T15:34:49.008Z",
    "end" : "2014-05-16T15:34:49.435Z",
    "source" : "scheduler",
    "type" : "execute",
    "status" : "ok",
    "details" : {
      "status" : 200,
      "entity" : "{\n  \"id\" : \"TwitterSearch\", \n  \"dataSourceId\" : \"TwitterSearch\", \n  \"state\" :
\\\"RUNNING\\\", \n  \"message\" : null, \n  \"startTime\" : 1400254489000, \n  \"endTime\" : -1, \n  \"finished\" :
false, \n  \"counters\" : { }, \n  \"exception\" : null, \n  \"running\" : true\n}"
    },
    "error" : null
  }, {
    "start" : "2014-05-16T15:38:32.536Z",
    "end" : "2014-05-16T15:38:32.559Z",
    "source" : "scheduler",
    "type" : "execute",
    "status" : "ok",
    "details" : {
      "status" : 200,
      "entity" : "{\n  \"id\" : \"TwitterSearch\", \n  \"dataSourceId\" : \"TwitterSearch\", \n  \"state\" :
\\\"RUNNING\\\", \n  \"message\" : null, \n  \"startTime\" : 1400254712000, \n  \"endTime\" : -1, \n  \"finished\" :
false, \n  \"counters\" : { }, \n  \"exception\" : null, \n  \"running\" : true\n}"
    },
    "error" : null
  }
]
}

```

10.34.3. Nodes API

The Nodes API allows users to check active services according to their hosts.

Examples

Show each service:

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/nodes/up
```

RESPONSE

```

{
  "dynamicSchema::v1" : [ "http://localhost:8764/api/apollo/dynamicSchema" ],
  "history::v1" : [ "http://localhost:8984/connectors/v1/history", "http://localhost:8764/api/apollo/history"
],
  "query-pipelines::v1" : [ "http://localhost:8764/api/apollo/query-pipelines",
"http://localhost:8764/api/apollo/query-pipelines" ],
  "solrconfig::v1" : [ "http://localhost:8764/api/apollo/collections//solr-config" ],
  "solrAdmin::v1" : [ "http://localhost:8764/api/apollo/solrAdmin",
"http://localhost:8984/connectors/v1/solrAdmin" ],
  "stopwords::v1" : [ "http://localhost:8764/api/apollo/stopwords" ],
  "collections::v1" : [ "http://localhost:8764/api/apollo/collections" ],
  "index-stages::v1" : [ "http://localhost:8984/connectors/v1/index-stages",
"http://localhost:8764/api/apollo/index-stages" ],
  "nodes::v1" : [ "http://localhost:8984/connectors/v1/nodes", "http://localhost:8764/api/apollo/nodes" ],
  "index-profiles::v1" : [ "http://localhost:8764/api/apollo/collections//index-profiles",
"http://localhost:8984/connectors/v1/collections//index-profiles" ],
  "schema::v1" : [ "http://localhost:8764/api/apollo/collections//schema" ],
  "solr::v1" : [ "http://localhost:8984/connectors/v1/solr", "http://localhost:8764/api/apollo/solr" ],
  "query-stages::v1" : [ "http://localhost:8764/api/apollo/query-stages" ],
  "blobs::v1" : [ "http://localhost:8764/api/apollo/blobs" ],
  "connectors::v1" : [ "http://localhost:8984/connectors/v1/connectors" ],
  "recommend::v1" : [ "http://localhost:8764/api/apollo/recommend" ],
  "configurations::v1" : [ "http://localhost:8764/api/apollo/configurations",
"http://localhost:8984/connectors/v1/configurations" ],
  "system::v1" : [ "http://localhost:8764/api/apollo/system", "http://localhost:8984/connectors/v1/system" ],
  "index-pipelines::v1" : [ "http://localhost:8764/api/apollo/index-pipelines",
"http://localhost:8764/api/apollo/index-pipelines", "http://localhost:8764/api/apollo/index-pipelines",
"http://localhost:8764/api/apollo/index-pipelines", "http://localhost:8764/api/apollo/index-pipelines",
"http://localhost:8764/api/apollo/index-pipelines", "http://localhost:8764/api/apollo/index-pipelines",
"http://localhost:8984/connectors/v1/index-pipelines", "http://localhost:8984/connectors/v1/index-pipelines",
"http://localhost:8764/api/apollo/index-pipelines", "http://localhost:8764/api/apollo/index-pipelines" ],
  "query-pipeline-templates::v1" : [ "http://localhost:8764/api/apollo/pipeline-templates/query" ],
  "registration::v1" : [ "http://localhost:8984/connectors/v1/registration",
"http://localhost:8764/api/apollo/registration" ],
  "searchLogs::v1" : [ "http://localhost:8764/api/apollo/searchLogs" ],
  "scheduler::v1" : [ "http://localhost:8764/api/apollo/scheduler" ],
  "aggregator::v1" : [ "http://localhost:8764/api/apollo/aggregator" ],
  "query-profiles::v1" : [ "http://localhost:8764/api/apollo/collections//query-profiles" ],
  "usage::v1" : [ "http://localhost:8764/api/apollo/usage", "http://localhost:8984/connectors/v1/usage" ],
  "introspect::v1" : [ "http://localhost:8764/api/apollo/introspect",
"http://localhost:8984/connectors/v1/introspect" ],
  "searchCluster::v1" : [ "http://localhost:8764/api/apollo/searchCluster",
"http://localhost:8984/connectors/v1/searchCluster" ],
  "features::v1" : [ "http://localhost:8764/api/apollo/features" ],
  "index-pipeline-templates::v1" : [ "http://localhost:8984/connectors/v1/pipeline-templates/index",
"http://localhost:8764/api/apollo/pipeline-templates/index" ],
  "synonyms::v1" : [ "http://localhost:8764/api/apollo/synonyms" ],
  "reports::v1" : [ "http://localhost:8764/api/apollo/reports" ],
  "signals::v1" : [ "http://localhost:8764/api/apollo/signals" ]
}

```


10.34.4. System API

The System REST API allows you to monitor the system performance.

Examples

Metric names

Get metric names that start with 'mem.heap':

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/system/metricNames?prefix=mem.heap
```

RESPONSE

```
{
  "gauges" : [ "mem.heap.committed", "mem.heap.init", "mem.heap.max", "mem.heap.usage", "mem.heap.used" ],
  "counters" : [ ],
  "histograms" : [ ],
  "meters" : [ ],
  "timers" : [ ]
}
```

Metrics

Find all metrics that match the regular expression pattern 'com.lucidworks.apollo.pipeline.index.', format the response so it's readable, and show the sample data, if any:

REQUEST

```
curl -u user:pass
http://localhost:8764/api/apollo/system/metrics?pattern=com.lucidworks.apollo.*pipeline.index.*&pretty=true&showSamples=true
```

RESPONSE

```
{
  "version" : "3.0.0",
  "gauges" : { },
  "counters" : { },
  "histograms" : { },
  "meters" : {
    "com.lucidworks.apollo.pipeline.index.IndexPipelineCache.cacheHit" : {
      "count" : 4775,
      "m15_rate" : 0.03604340402401043,
      "m1_rate" : 0.04985610410800882,
      "m5_rate" : 0.04753263154077047,
      "mean_rate" : 0.05028487069705915,
      "units" : "events/second"
    }
  },
}
```

```

"timers" : {
  "com.lucidworks.apollo.pipeline.index.IndexPipelineCache.deserialize" : {
    "count" : 7,
    "max" : 0.078783,
    "mean" : 0.028265285714285715,
    "min" : 9.800000000000001E-5,
    "p50" : 1.94E-4,
    "p75" : 0.0699610000000001,
    "p95" : 0.078783,
    "p98" : 0.078783,
    "p99" : 0.078783,
    "p999" : 0.078783,
    "values" : [ 9.800000000000001E-5, 1.09E-4, 1.62E-4, 1.94E-4, 0.04855, 0.0699610000000001, 0.078783 ],
    "stddev" : 0.03620774742010466,
    "m15_rate" : 2.964393875E-314,
    "m1_rate" : 2.128434034679706E-46,
    "m5_rate" : 4.9195401948202935E-138,
    "mean_rate" : 7.371603924510614E-5,
    "duration_units" : "seconds",
    "rate_units" : "calls/second"
  },
  "com.lucidworks.apollo.pipeline.index.IndexStageConfigCache.deserialize" : {
    "count" : 7,
    "max" : 0.002377,
    "mean" : 4.6642857142857147E-4,
    "min" : 8.7E-5,
    "p50" : 1.16E-4,
    "p75" : 3.46E-4,
    "p95" : 0.002377,
    "p98" : 0.002377,
    "p99" : 0.002377,
    "p999" : 0.002377,
    "values" : [ 8.7E-5, 9.300000000000001E-5, 1.11E-4, 1.16E-4, 1.35E-4, 3.46E-4, 0.002377 ],
    "stddev" : 8.47267737523163E-4,
    "m15_rate" : 2.964393875E-314,
    "m1_rate" : 2.1522151745137967E-46,
    "m5_rate" : 5.086288568158318E-138,
    "mean_rate" : 7.372680363461613E-5,
    "duration_units" : "seconds",
    "rate_units" : "calls/second"
  },
  "com.lucidworks.apollo.pipeline.index.IndexStageConfigStore.deserialize" : {
    "count" : 6,
    "max" : 0.0019760000000000003,
    "mean" : 5.606666666666667E-4,
    "min" : 1.0800000000000001E-4,
    "p50" : 1.4250000000000002E-4,
    "p75" : 0.00114875,
    "p95" : 0.0019760000000000003,
    "p98" : 0.0019760000000000003,
    "p99" : 0.0019760000000000003,
    "p999" : 0.0019760000000000003,
    "values" : [ 1.0800000000000001E-4, 1.2200000000000001E-4, 1.35E-4, 1.5000000000000001E-4,
8.7300000000000001E-4, 0.0019760000000000003 ],
    "stddev" : 7.54704622131511E-4,
    "m15_rate" : 2.964393875E-314,
    "m1_rate" : 1.0220227879692082E-48,
    "m5_rate" : 7.20591046931865E-140,
  }
}

```

```

    "mean_rate" : 6.318494243486903E-5,
    "duration_units" : "seconds",
    "rate_units" : "calls/second"
  },
  "com.lucidworks.apollo.pipeline.index.IndexStageConfigStore.getItem" : {
    "count" : 6,
    "max" : 0.003128,
    "mean" : 0.002295,
    "min" : 0.0017770000000000002,
    "p50" : 0.0020715,
    "p75" : 0.002774,
    "p95" : 0.003128,
    "p98" : 0.003128,
    "p99" : 0.003128,
    "p999" : 0.003128,
    "values" : [ 0.0017770000000000002, 0.002066, 0.002066, 0.0020770000000000003, 0.0026560000000000004,
0.003128 ],
    "stddev" : 4.989869737778733E-4,
    "m15_rate" : 2.964393875E-314,
    "m1_rate" : 1.0220227879692082E-48,
    "m5_rate" : 7.20591046931865E-140,
    "mean_rate" : 6.318494244418448E-5,
    "duration_units" : "seconds",
    "rate_units" : "calls/second"
  }
}
}
}

```

The above output has been truncated for space to remove metrics with no data or with very long value lists.

Threads

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/system/threads
```

RESPONSE (truncated to a single thread)

```
[ {
  "id" : 2,
  "native" : false,
  "name" : "Reference Handler",
  "locks" : {
    "waiting" : {
      "identity" : "0x7257d934",
      "class" : "java.lang.ref.Reference$Lock"
    },
    "locking" : {
      "identity" : "0x7257d934",
      "class" : "java.lang.ref.Reference$Lock"
    }
  },
  "state" : "WAITING",
  "suspended" : false,
  "stackTrace" : [ {
    "methodName" : "wait",
    "fileName" : "Object.java",
    "lineNumber" : -2,
    "className" : "java.lang.Object",
    "nativeMethod" : true
  }, {
    "methodName" : "wait",
    "fileName" : "Object.java",
    "lineNumber" : 503,
    "className" : "java.lang.Object",
    "nativeMethod" : false
  }, {
    "methodName" : "run",
    "fileName" : "Reference.java",
    "lineNumber" : 133,
    "className" : "java.lang.ref.Reference$ReferenceHandler",
    "nativeMethod" : false
  } ]
} ]
```

Buffers

REQUEST

```
curl -u user:pass -X PUT http://localhost:8764/api/apollo/system/buffers
```

Ping

REQUEST

```
curl -u user:pass http://localhost:8764/api/apollo/system/ping
```

RESPONSE

```
pong
```

10.34.5. Usage API

The Usage API sends system usage data to Fusion. See the section System Usage Monitor for details.

Examples

Get the latest data that will be sent to Fusion on the next run:

INPUT

```
curl -u user:pass http://localhost:8764/api/apollo/usage
```

OUTPUT

```
[ {
  "type" : "GAUGE",
  "name" : "solr.clusters",
  "value" : 3
}, {
  "type" : "GAUGE",
  "name" : "solr.nodes",
  "value" : 2
}, {
  "type" : "RATE",
  "name" : "recommender.itemsForQuery.timer.rate",
  "value" : 0
}, {
  "type" : "RATE",
  "name" : "recommender.queriesForItem.timer.rate",
  "value" : 0
}, {
  "type" : "GAUGE",
  "name" : "nodes",
  "value" : 2
}, {
  "type" : "RATE",
  "name" : "recommender.queriesForItem.counter.rate",
  "value" : 0
}, {
  "type" : "RATE",
  "name" : "recommender.itemsForItem.timer.rate",
  "value" : 0
}, {
  "type" : "RATE",
  "name" : "recommender.itemsForQuery.counter.rate",
  "value" : 0
}, {
  "type" : "RATE",
  "name" : "recommender.itemsForItem.counter.rate",
  "value" : 0
} ]
```

10.35. Tasks API

The Tasks API allows you to view and define Fusion tasks, which can be run using the Scheduler.

10.36. Taxonomy API

The Taxonomy API is used to manage taxonomies, such as a hierarchy of product categories used for an eCommerce site.

The information in a taxonomy is meta-information about the categories used to classify a set of things. For an eCommerce site, the set of things are items in the product catalog. Fusion uses this meta-information to enhance search. In Fusion, these products are stored in a Fusion collection, (the primary collection), and the taxonomy is stored in an auxiliary collection. Naming conventions are used to relate the primary and auxiliary collections: given a primary collection named "COLL", the auxiliary taxonomy is stored in a collection named "COLL_stored_parameters".

10.36.1. Taxonomy structure

In information-theoretic terms, the structure of a taxonomy is either a tree (or trees), or at least a directed acyclic graph (DAG) or set of DAGS. Each category object is a vertex in the tree or DAG, and the "children" property of each category object lists the set of outgoing edges.

Each category object has a unique ID. In addition, category objects have an optional property called a "parameter", which is a list of named attributes associated with that category. These parameters values from the taxonomy are added to search queries over the primary collection. In addition, every category object has a property "version" which is managed by Fusion and used to ensure that concurrent updates to the same category do not conflict with each other. The value of this property corresponds to the system clock time in milliseconds. This value is set and updated by Fusion and need only be specified overtly when sending category update requests.

Given the above specification, the set of property: values for a category object is:

- id: a unique string identifier.
- children: a list of category objects.
- parameters: a list of named attributes associated with that category where each attribute is specified as a pair of strings key:value.
- version: a string corresponding to the integer value of some system clock in milliseconds.

10.36.2. JSON example of a taxonomy

The following example is part of a taxonomy for an online pet supply store. In this example, we have the following hierarchy of category id, label pairs:

- 1 : "Pet Supplies"
 - 11 "Cat Supplies"
 - 111 "Cat Food"
 - 112 "Cat Toys"
 - 12 "Dog Supplies"
 - 121 "Dog Food"
 - 122 "Dog Toys"

We annotate these category objects with a set of parameters, which will be used organize and enrich user searches over the product catalog:

```

{ "id": "1",
  "label": "Pet Supplies",
  "children": [
    { "id": "11",
      "label": "Cat Supplies",
      "children": [
        { "id": "111",
          "label": "Cat Food",
          "parameters": [
            { "key": "facet.field", "value": "brand" },
            { "key": "facet.field", "value": "health_benefit" },
            { "key": "facet.field", "value": "size" }
          ]
        },
        { "id": "112",
          "label": "Cat Toys",
          "parameters": [
            { "key": "facet.field", "value": "price" }
          ]
        }
      ]
    },
    { "id": "12",
      "label": "Dog Supplies",
      "children": [
        { "id": "121",
          "label": "Dog Food",
          "parameters": [
            { "key": "facet.field", "value": "life_stage" },
            { "key": "facet.field", "value": "health_benefit" },
            { "key": "facet.field", "value": "flavor" }
          ]
        },
        { "id": "122",
          "label": "Dog Toys",
          "parameters": [
            { "key": "facet.field", "value": "price" },
            { "key": "facet.field", "value": "dog_size" },
            { "key": "facet.field", "value": "activity_type" }
          ]
        }
      ]
    }
  ]
}

```

The version property on these category objects isn't overtly specified here as it is managed by Fusion.

10.37. ZooKeeper Import/Export API

The ZooKeeper Import/Export API provides methods to upload or download information from Fusion's ZooKeeper service. This service provides an alternative to the ZooKeeper clients `zkCli.sh` and `zk-shell` which are part of the Apache Zookeeper distribution included as part of the Fusion distribution. It was introduced as part of the Fusion 2.0 release.

The `ZKImportExport` service may be used to export ZooKeeper data for any Fusion release. It can be used to import configuration data into the ZooKeeper service for a new or existing Fusion deployment. Note that since the Fusion 3.0 release all ZooKeeper paths vary according to the version of Fusion that you are running.

- For details on using this script during the Fusion upgrade procedure, see [Upgrading Fusion](#).
- For details on using this script to migrate Fusion configurations from one deployment to another, see [Migrating Fusion data](#).

The REST API only supports requests to export ZooKeeper configurations. The Fusion distribution includes a utility script `zkImportExport.sh` which can be used to import ZooKeeper configuration as well as to export it from arbitrary Fusion instances.

10.37.1. ZooKeeper

[Apache ZooKeeper](#) is a distributed configuration service, synchronization service, and naming registry. Fusion uses ZooKeeper to configure and manage all Fusion components in a single Fusion deployment.

- **znode:** ZooKeeper data is organized into a hierarchal name space of data nodes called znodes. A znode can have data associated with it as well as child znodes. The data in a znode is stored in a binary format, but it is possible to import, export, and view this information as JSON data. Paths to znodes are always expressed as canonical, absolute, slash-separated paths; there are no relative reference.
- **ephemeral nodes:** An ephemeral node is a znode which exists only for the duration of an active session. When the session ends the znode is deleted. An ephemeral znode cannot have children.
- **server:** A ZooKeeper service consists of one or more machines; each machine is a server which runs in its own JVM and listens on its own set of ports. For testing, you can run several ZooKeeper servers at once on a single workstation by configuring the ports for each server.
- **quorum:** A quorum is a set of ZooKeeper servers. It must be an odd number. For most deployments, only 3 servers are required.
- **client:** A client is any host or process which uses a ZooKeeper service.

See the official [ZooKeeper documentation](#) for details about using and managing a ZooKeeper service.

10.37.2. Utility script `zkImportExport.sh`

This script is located in the top-level Fusion `scripts` directory. The script takes the following command-line arguments:

<code>-c,--cmd <arg></code>	Command, one of: 'export', 'import', 'update', 'delete'.
<code>-e,--encode <arg></code>	Type of encoding for znodes. Valid options: 'none', 'utf-8', 'base64', default is 'base64'. Option 'none' will not return any data from the znodes.
<code>-ep,--exclude <arg></code>	Exclude znode paths, followed by list of paths. Can only be used to exclude nodes one level below the root node.
<code>-eph,--ephemeral</code>	Include ephemeral nodes while exporting znodes, boolean, default false.
<code>-f,--filename <arg></code>	Name of file containing import/export data.
<code>-h,--help</code>	Display help page.
<code>-ip,--include <arg></code>	Include znode paths to include, followed by a list of paths. Can only be used to include nodes one level below the root node.
<code>-nr,--non-recursive</code>	Do not perform recursive operations on znodes.
<code>-o,--overwrite</code>	Overwrite data for existing znodes. Valid only with 'update' command.
<code>-p,--path <arg></code>	Path from ZooKeeper root node, e.g. '/lucid/query-pipelines'.
<code>-r,--recursive</code>	Perform recursive operations on znodes.
<code>-z,--zkhost <arg></code>	ZooKeeper Connect string, required.

Required arguments are: * `-c, --cmd` : operation to perform. * `-z, --zkhost` : the ZooKeeper Connect string.

Examples

Export all data from a local single-node ZooKeeper service, save data to a file:

```
zkImportExport.sh -zkhost localhost:9983 -cmd export -path / -filename znode_dump.json
```

Export all Fusion configurations from a local single-node ZooKeeper service, save data to a file:

```
zkImportExport.sh -zkhost localhost:9983 -cmd export -path /lwfusion -filename znode_lucid_dump.json
```

Export Fusion user databases, groups, roles, and realms configurations from a local single-node ZooKeeper service, save data to a file:

```
zkImportExport.sh -zkhost localhost:9983 -cmd export -path /lwfusion/3.1.2/proxy/user -filename znode_lucid_admin_dump.json
```

Initial import of saved Fusion configuration into a new ZooKeeper:

```
zkImportExport.sh -zkhost localhost:9983 -cmd import -filename znode_lucid_dump.json
```

Note that the above command will fail if there is conflict between existing znode structures or contents between the ZooKeeper service and the dump file.

Update information for Fusion's ZooKeeper service:

```
zkImportExport.sh -zkhost localhost:9983 -cmd update -filename znode_lucid_dump.json
```

Remove a znode from Fusion's ZooKeeper service:

```
zkImportExport.sh -zkhost localhost:9983 -cmd delete -path /lwfusion/test
```

10.37.3. Fusion REST API service ZKImportExport

The Fusion REST API can only be used to download information from ZooKeeper, via the 'GET' method with the following configuration:

```
"zk-import-export::v1" : {
  "name" : "com.lucidworks.apollo.resources.ZKImportExportResource",
  "uri" : "/zk/export",
  "methods" : [ {
    "uri" : "/zk/export/{path:.*}",
    "name" : "getNodeInfo",
    "verb" : "GET",
    "pathParams" : [ {
      "name" : "path",
      "type" : "String"
    } ],
    "queryParams" : [ {
      "name" : "recursive",
      "type" : "Boolean"
    } ], {
      "name" : "excludePaths",
      "type" : "List"
    }, {
      "name" : "includePaths",
      "type" : "List"
    }, {
      "name" : "encodeValues",
      "type" : "String"
    } ],
    "hasEntity" : false
  } ]
}
```

GET data from path ``/lwfusion/3.1.2/core/query-pipelines``


```
curl -u user:pass -X GET http://localhost:8764/api/apollo/zk/export/lwfusion/3.1.2/core/query-  
pipelines?recursive=true&encodeValues=utf-8&excludePaths=/lwfusion/3.1.2/core/query-pipelines/default  
{  
  "path" : "/lwfusion/3.1.2/core/query-pipelines",  
  "data" : ""  
}
```