

1. Planning a Resource Manager deployment	8
1.1 Welcome to Zenoss	9
1.1.1 Introduction to Control Center	10
1.1.1.1 Docker fundamentals	11
1.1.1.2 Control Center internal services	12
1.1.1.3 ZooKeeper and Control Center	13
1.1.1.4 Control Center application data storage	14
1.1.2 Introduction to Zenoss Service Dynamics	15
1.1.2.1 Key Resource Manager concepts	16
1.1.3 Installation options	17
1.2 Resource Manager virtual appliances	18
1.2.1 Resource requirements for multi-host deployments	19
1.2.2 Resource requirements for single-host deployments	20
1.3 Control Center resource requirements	21
1.3.1 Installation considerations	22
1.3.2 Compute and storage requirements	23
1.3.2.1 Master host CPU and RAM resources	24
1.3.2.1.1 Master host storage areas	25
1.3.2.2 Delegate host CPU and RAM resources	27
1.3.2.2.1 Delegate host storage requirements	28
1.3.3 Operating system requirements	30
1.3.3.1 Networking	31
1.3.3.2 Security	32
1.3.3.3 Resource Manager network ports	33
2. Installing Resource Manager	34
2.1 Adding Resource Manager to a Control Center deployment	36
2.1.1 Downloading template and image files	37
2.1.2 Installing the Resource Manager template	38
2.1.3 Importing Resource Manager images	39
2.1.4 Deploying Resource Manager in Control Center	40
2.2 Installing a master host	41
2.2.1 Deploying Resource Manager	42
2.2.2 Adding the master host to a resource pool	43
2.2.3 Configuring the Control Center master host	44
2.2.4 Creating a master host with Hyper-V	45
2.2.4.1 Configuring and starting a Hyper-V master host	46
2.2.5 Creating a master host with vSphere	47
2.3 Installing delegate hosts	49
2.3.1 Editing the /etc/hosts file	50
2.3.2 Configuring a delegate host virtual machine	51
2.3.3 Creating a delegate host with Hyper-V	52
2.3.3.1 Configuring and starting a Hyper-V delegate host	53
2.3.4 Creating a delegate host with vSphere	54
2.4 Configuring a multi-host deployment	55
2.4.1 Adding a delegate host through an SSH connection	56
2.4.2 Adding a delegate host using a file	57
2.4.3 Enabling NTP on Microsoft Hyper-V guests	58
2.4.3.1 Configuring NTP clients	59
2.4.3.2 Configuring an NTP master server	60
2.4.3.3 Configuring NTP for public time servers	61
2.5 Guidelines for resource pool permissions	62
3. Configuring Resource Manager	63
3.1 Enabling access to browser interfaces	64
3.1.1 Creating and changing public endpoints	65
3.1.1.1 Creating a port public endpoint	66
3.1.1.1.1 Configuring Zope for HTTPS and the default secure proxy server	68
3.1.1.1.2 Configuring Zope for HTTP and no proxy server	69
3.1.1.1.3 Configuring Zope for HTTP and a secure proxy server	70
3.1.1.2 Creating a virtual host public endpoint	71
3.1.2 Configuring name resolution for virtual hosts	72
3.2 Configuration procedures	73
3.2.1 Starting Resource Manager	74
3.2.2 Default server passwords	75
3.2.2.1 Changing MariaDB passwords	76
3.2.2.2 Changing the RabbitMQ server password	78
3.2.2.3 Changing the Zope authentication server password	79
3.2.3 Deleting the RabbitMQ guest user account	80
3.2.4 MariaDB database utilities	81
3.2.4.1 Installing the Percona Toolkit with internet access	82
3.2.4.2 Installing the Percona Toolkit without internet access	83
3.2.5 Optional: Assigning a virtual IP address to a resource pool	84
3.2.6 Optional: Replacing the default digital certificate	85
3.2.7 Optional: Customization management	86
3.2.7.1 Installing Quilt with internet access	87
3.2.7.2 Installing Quilt without internet access	88
3.2.8 Optional: Configuring OpenTSDB compaction	89
3.2.9 Optional: Creating a Redis cluster in a collector pool	90
3.2.10 Optional: Enabling monitoring on IPv6 networks	91
3.3 Preparing for monitoring	93

3.3.1	Preparing network devices	94
3.3.2	Preparing storage devices	95
3.3.3	Preparing server devices	96
3.3.4	Preparing hypervisor devices	97
3.3.5	Validating configuration using Inspector tool	98
3.4	External HBase configuration	99
3.4.1	Configuring OpenTSDB for an external HBase cluster	100
3.4.2	Configuring the OpenTSDB service startup command	101
3.4.3	Disabling the Resource Manager HBase cluster	102
4.	Administering Resource Manager	103
4.1	Using Resource Manager	104
4.1.1	Initial login	105
4.1.2	Interface and navigation	107
4.1.3	Administering Resource Manager dashboards	108
4.1.3.1	Portlets	111
4.1.4	Search	113
4.1.5	Navigating the event console	114
4.1.5.1	Sorting and filtering events	115
4.1.5.2	Creating an actionable view	116
4.1.5.3	Saving a custom view	117
4.1.5.4	Refreshing the view	118
4.1.5.5	Viewing event details	119
4.1.5.6	Selecting events	120
4.1.5.7	Managing events	121
4.1.6	Running a command	122
4.1.7	Visualizing your environment	123
4.1.7.1	Dynamic service view	124
4.1.7.1.1	Dynamic view of organizers	126
4.1.7.1.2	Dynamic view of devices	127
4.1.7.2	Datacenter view	131
4.1.7.2.1	Working with the List View	132
4.1.7.2.2	Working with the Custom View	133
4.1.7.2.3	Activating the auto-generated rack view	138
4.2	Preparing devices for monitoring	140
4.2.1	Configuring Linux devices to provide data through SNMP	141
4.2.2	Configuring Windows devices to provide data through SNMP	142
4.3	Working with devices in Resource Manager	143
4.3.1	Viewing the device list	144
4.3.1.1	Devices hierarchy	145
4.3.1.2	Managing multiple devices from the device list	146
4.3.2	Working with devices	147
4.3.2.1	Dependency view	149
4.3.2.2	Dynamic view	150
4.3.2.3	Events view	151
4.3.2.4	Components	152
4.3.2.4.1	Component graphs	153
4.3.2.4.2	Disabling component monitoring	154
4.3.2.5	Graphs (Performance)	155
4.3.2.6	Modeler plugins	156
4.3.2.7	Configuration properties	157
4.3.2.8	Custom properties	158
4.3.2.9	Device administration	159
4.3.2.10	Override objects	160
4.3.2.11	Software	161
4.3.3	Managing devices and device attributes	162
4.3.3.1	Clearing heartbeat events	163
4.3.3.2	Locking device configuration	164
4.3.3.3	Renaming a device	165
4.3.3.4	Re-identifying a device	166
4.3.3.5	Remodeling a device	167
4.3.3.6	Resetting the device manage IP address	168
4.3.3.7	Deleting a device	169
4.3.3.8	Exporting device list to load into another system	170
4.3.3.9	Batch loading or modifying devices	171
4.3.4	Adding and discovering devices	173
4.3.4.1	Discovering devices	174
4.3.4.1.1	Discovery mapping	175
4.3.4.1.2	Providing network or IP address range for device discovery	176
4.3.4.1.3	Discovering devices with the CLI	177
4.3.4.1.4	Classifying discovered devices	178
4.3.4.1.5	Updating device authentication details	179
4.3.4.1.6	Adding or editing information on a device record	180
4.3.4.2	Adding devices manually	181
4.3.4.2.1	Adding a single device	182
4.3.4.2.2	Adding multiple devices	183
4.3.4.2.3	Adding a Cisco UCS device	184
4.3.4.2.4	Add VMware vSphere endpoint	185
4.4	Basic monitoring	186
4.4.1	Availability monitoring	187

4.4.1.1 zenping correlation	188
4.4.1.2 Controlling ping cycle time	189
4.4.1.3 Using the predefined /Ping device class	191
4.4.1.4 Monitoring processes	192
4.4.1.4.1 Example: Creating a process class	193
4.4.1.4.2 Process class options	201
4.4.1.5 Monitoring IP services	202
4.4.1.6 Monitoring Windows Services	203
4.4.2 Monitoring using ZenCommand	204
4.4.2.1 Plugin format for ZenCommands	205
4.4.2.2 Testing ZenCommands	206
4.4.3 SNMP monitoring	207
4.4.4 Monitoring devices remotely through SSH	208
4.4.4.1 Changing Resource Manager to monitor devices remotely using SSH	209
4.4.4.2 Using the predefined /Server/Command device class	210
4.4.5 The network map page	211
4.5 Performance monitoring	213
4.5.1 Monitoring templates and performance data	214
4.5.2 Template names	215
4.5.3 Data sources	216
4.5.4 Data points	217
4.5.5 Data point aliases	218
4.5.5.1 Alias formula evaluation	219
4.5.5.2 Adding a data point alias	220
4.5.6 Thresholds	221
4.5.6.1 MinMax threshold	222
4.5.6.1.1 Editing MinMax thresholds	223
4.5.6.2 ValueChange threshold	224
4.5.6.2.1 Editing ValueChange thresholds	225
4.5.6.3 CiscoStatus threshold	226
4.5.6.3.1 Editing CiscoStatus thresholds	227
4.5.6.4 Predictive threshold	228
4.5.6.4.1 Editing Predictive Thresholds	229
4.5.6.5 Adding thresholds	230
4.5.6.6 Adding a trendline to a graph	231
4.5.7 Performance graphs	232
4.5.7.1 Graph points	233
4.5.7.1.1 Re-sequencing graph points	234
4.5.7.1.2 DataPoint graph points	235
4.5.7.1.3 Editing threshold graph points	236
4.5.8 Performance data retention	237
4.6 Distributed monitoring	238
4.6.1 Adding a hub or collector	239
4.6.2 Adding devices to collectors	240
4.6.3 Moving all devices from one collector to another	241
4.6.4 Moving individual devices between collectors	242
4.6.5 Navigating collectors and hubs	243
4.7 Monitoring Zenoss	245
4.7.1 Monitoring Control Center	246
4.7.2 Control Center components	247
4.7.3 Customizing Control Center monitoring	248
4.7.4 Monitoring file system storage	249
4.7.5 Monitoring Resource Manager	250
4.7.6 ZenossRM components	251
4.7.7 Customizing Resource Manager monitoring	253
4.8 Extending Resource Manager with ZenPacks	254
4.8.1 Preparing to install or update a ZenPack	256
4.8.2 Installing or updating a ZenPack	257
4.8.3 Removing a ZenPack	258
4.8.4 Creating a ZenPack	259
4.9 Using organizers	260
4.9.1 Device classes	261
4.9.2 Groups, Systems, and Locations	265
4.9.2.1 Setting an address for a location	268
4.9.2.2 Network links	270
4.9.2.3 Google Maps example	271
4.9.2.4 Clearing the Google Maps cache	272
4.9.3 Component groups	273
4.10 Managing background tasks	274
4.11 Using configuration properties	275
4.11.1 Configuration properties inheritance and override	276
4.11.1.1 Inheritance in the device class tree	278
4.11.2 Configuration property types	279
4.11.3 Viewing and overriding device properties	280
4.11.3.1 List of device configuration properties	281
4.11.4 Viewing and overriding event properties	290
4.11.4.1 Table of event configuration properties	291
4.11.5 Viewing and overriding network properties	292
4.11.5.1 List of network configuration properties	293

4.12 Modeling	294
4.12.1 Modeling devices	295
4.12.1.1 Modeling devices using SSH/COMMAND	296
4.12.1.2 Modeling devices using port scan	297
4.12.1.3 Using the /Server/Scan device class to monitor with port scan	298
4.12.2 About modeler plugins	299
4.12.3 Debugging the modeling process	300
4.13 About monitoring templates	301
4.13.1 Device templates	302
4.13.2 Component templates	305
4.13.3 Interface templates	306
4.14 Production states and maintenance windows	307
4.14.1 Production states	308
4.14.2 Maintenance windows	309
4.14.2.1 Maintenance window events	310
4.14.2.2 Creating and using maintenance windows	311
4.15 Event management	312
4.15.1 Event fields	313
4.15.2 De-duplication	316
4.15.3 Auto-clear correlation	317
4.15.4 Event consoles	318
4.15.5 Event sources	322
4.15.5.1 SNMP traps	323
4.15.5.1.1 Classifying SNMP traps	324
4.15.5.1.2 Example: Sending test traps	325
4.15.5.1.3 Transforming events with event mappings	326
4.15.5.1.4 Event transforms based on event class	329
4.15.5.1.5 SNMP trap filtering at the collector level	330
4.15.5.1.6 Configuring SNMP trap forwarding	331
4.15.5.2 Configuring syslog message forwarding	332
4.15.6 Creating events manually	333
4.15.7 Understanding event classes	335
4.15.8 Event mapping and transforms	336
4.15.8.1 Event class mappings	337
4.15.8.2 Event class mapping sequence	338
4.15.8.3 Event class transform	339
4.15.9 Capturing email messages as events	340
4.15.10 Event severity levels	342
4.15.11 Administering MIB files	343
4.15.11.1 Using MIB files	344
4.15.11.2 Importing pre-loaded MIB organizers and files	345
4.15.11.3 Creating a MIBs organizer	346
4.15.11.4 Installing custom or additional MIB files	347
4.16 Triggers and notifications	348
4.16.1 Working with triggers	349
4.16.2 Working with notifications	351
4.16.2.1 Defining notification content	355
4.16.2.1.1 Notification content variables	358
4.16.2.2 Setting individual notification permissions	361
4.17 Managing users in Resource Manager	362
4.17.1 Creating user accounts	363
4.17.2 Editing user accounts	364
4.17.2.1 Associating objects with specific users	365
4.17.2.2 Changing the admin account password with the CLI	366
4.17.3 User groups	367
4.17.4 Roles and permissions	368
4.17.5 Device access control lists	369
4.17.5.1 Example: Restricted user with ZenOperator role	370
4.17.5.2 Detailed restricted screen functionality	371
4.18 General administration and settings	372
4.18.1 Audit logging	373
4.18.2 User commands	374
4.18.3 Changing events database connection information	377
4.18.4 Rebuilding the events index	378
4.18.5 Support bundles	379
4.18.6 Working with the job manager	381
5. Troubleshooting Resource Manager	383
5.1 Data pipelines	384
5.1.1 Event pipeline	385
5.1.1.1 Troubleshooting event flow	387
5.1.2 Model pipeline	389
5.1.2.1 Troubleshooting modeling	390
5.1.3 Performance data pipeline	391
5.1.3.1 Troubleshooting performance collection and retrieval	392
5.2 How to read zenhub log messages	393
6. Updating Resource Manager	394
6.1 Preparing to update an appliance-based deployment	397
6.1.1 Preparing an appliance-based deployment for update	398
6.1.2 Downloading the update ISO file	399

6.1.3	Attaching an update ISO with vSphere	400
6.1.4	Attaching an update ISO with Hyper-V	401
6.1.5	Stopping Resource Manager	402
6.1.6	Removing Docker containers on the master host	403
6.1.7	Removing Docker containers on delegate hosts	404
6.2	Updating an appliance-based deployment	406
6.2.1	Updating Control Center on a delegate host	407
6.2.2	Updating Control Center on the master host	408
6.2.3	Updating delegate hosts with authentication	409
6.2.3.1	Registering a host using SSH	410
6.2.3.2	Registering a host using a file	411
6.2.4	Updating Resource Manager on the master host	412
6.2.5	Starting Resource Manager after an update	413
6.2.6	Clearing heartbeat events after updating	414
6.3	Updating non-appliance deployments	415
6.3.1	Preparing a deployment for update	416
6.3.1.1	Installing the base image, if necessary	417
6.3.2	Downloading Resource Manager image files	418
6.3.3	Importing Resource Manager image files	419
6.3.4	Stopping Resource Manager on non-appliance deployments	420
6.3.5	Updating Resource Manager on non-appliance deployments	421
6.3.6	Removing the pre-upgrade snapshot	422
6.4	Using Zenoss Toolbox	423
6.4.1	Zenoss Toolbox tools	424
6.4.2	Running Zenoss Toolbox tools	425
6.5	Common update error recovery procedures	426
6.5.1	A snapshot with the given tag already exists	427
6.6	Installing an application template	428
6.6.1	Downloading the template package	429
6.6.2	Installing the application template	430
6.7	ZenPack considerations	431
6.7.1	Alternate naming convention for LUN- and VM-specific metrics	432
7.	Release notes	433
7.1	Resource Manager 6.4.1	434
7.1.1	Updating the Dashboard ZenPack	439
7.1.2	Patching a Linux modeling issue	440
7.1.3	Removing customized MariaDB passwords for system updates	442
7.1.4	Verifying zenhubworker services updated successfully	444
7.2	Resource Manager 6.4.0	445
7.3	Resource Manager 6.3.2	446
7.4	Resource Manager 6.2.1	453
8.	Appendixes	457
8.1	Resource Manager interface reference	458
8.1.1	DASHBOARD (Resource Manager)	459
8.1.2	EVENTS	460
8.1.3	INFRASTRUCTURE	461
8.1.4	SERVICES	462
8.1.5	REPORTS	463
8.1.5.1	AWS Reports	464
8.1.5.1.1	Monitoring Costs	465
8.1.5.2	Azure	466
8.1.5.2.1	Unattached VHDs	467
8.1.5.3	Capacity Planning	468
8.1.5.3.1	Capacity Usage	469
8.1.5.4	Cisco UCS Capacity Reports	470
8.1.5.4.1	Aggregate Bandwidth Utilization	471
8.1.5.4.2	Aggregate Port Pool Utilization	473
8.1.5.4.3	Bandwidth Utilization vs Capacity	474
8.1.5.4.4	HostNetworkCapacity	475
8.1.5.4.5	Interface 95th Percentile	476
8.1.5.4.6	Interface Utilization	478
8.1.5.4.7	Interface Volume (Cisco UCS)	480
8.1.5.4.8	Port Utilization	482
8.1.5.4.9	Storage Utilization vs Capacity	483
8.1.5.5	Cisco UCS Reports	484
8.1.5.5.1	Free Memory Slots	485
8.1.5.5.2	Hardware Inventory	486
8.1.5.6	Custom Device Reports	487
8.1.5.6.1	Creating a custom device report	488
8.1.5.7	Device Reports	489
8.1.5.7.1	All Devices	490
8.1.5.7.2	All Monitored Components	491
8.1.5.7.3	Device Changes	492
8.1.5.7.4	HP Chassis List	493
8.1.5.7.5	HP Device Bay List	494
8.1.5.7.6	MAC Addresses	495
8.1.5.7.7	Model Collection Age	496
8.1.5.7.8	New Devices	497
8.1.5.7.9	Ping Status Issues	498

8.1.5.7.10	SNMP Status Issues	499
8.1.5.7.11	Software Inventory	500
8.1.5.8	Enterprise Reports	501
8.1.5.8.1	Cisco Inventory	502
8.1.5.8.2	Customized Performance Templates	503
8.1.5.8.3	Data Sources in Use	504
8.1.5.8.4	Datapoints Per Collector	505
8.1.5.8.5	Defined Thresholds	506
8.1.5.8.6	Event Time to Resolution	507
8.1.5.8.7	Interface Utilization (enterprise report)	508
8.1.5.8.8	Interface Volume	509
8.1.5.8.9	Maintenance Windows (enterprise report)	510
8.1.5.8.10	Network Topology	511
8.1.5.8.11	Notifications and Triggers by Recipient	512
8.1.5.8.12	Organizer Availability	513
8.1.5.8.13	Organizer Graphs	514
8.1.5.8.14	User Event Activity	515
8.1.5.8.15	Users Group Membership	516
8.1.5.9	Event Reports	517
8.1.5.9.1	All EventClasses (All Event Classes)	518
8.1.5.9.2	All EventMappings (All Event Mappings)	519
8.1.5.9.3	All Heartbeats	520
8.1.5.9.4	Disabled Transforms	521
8.1.5.10	Graph Reports	522
8.1.5.10.1	Creating a graph report	523
8.1.5.10.2	Working with graph reports	524
8.1.5.11	Monitoring Capabilities Reports	525
8.1.5.11.1	Installed Templates	526
8.1.5.12	Multi-Graph Reports	527
8.1.5.12.1	Adding collections	528
8.1.5.12.2	Creating a multi-graph report	529
8.1.5.12.3	Adding graph definitions	530
8.1.5.12.4	Adding graph groups	532
8.1.5.13	Performance Reports	534
8.1.5.13.1	Availability Report	535
8.1.5.13.2	CPU Utilization	536
8.1.5.13.3	Filesystem Util Report	537
8.1.5.13.4	Interface Utilization (performance report)	538
8.1.5.13.5	Memory Utilization	539
8.1.5.13.6	Threshold Summary	540
8.1.5.14	Storage	541
8.1.5.14.1	Clients	542
8.1.5.14.2	Disk Firmware	543
8.1.5.14.3	Licenses	544
8.1.5.15	vSphere	545
8.1.5.15.1	Clusters	546
8.1.5.15.2	Datastores	547
8.1.5.15.3	Hosts	548
8.1.5.15.4	LUNs	549
8.1.5.15.5	Resource Pools	550
8.1.5.15.6	VM to Datastore	551
8.1.5.15.7	VMs	552
8.1.5.15.8	VMware Utilization	553
8.1.6	ADVANCED	554
8.1.6.1	ADVANCED > Settings	555
8.1.6.1.1	Event configuration settings	557
8.1.6.1.2	User interface configuration	558
8.1.6.2	ADVANCED > Control Center	559
8.1.6.3	ADVANCED > Move Devices	560
8.1.6.4	ADVANCED > Monitoring Templates	561
8.1.6.5	ADVANCED > Jobs	562
8.1.6.6	ADVANCED > MIBs	563
8.1.6.7	ADVANCED > Licensing	564
8.2	Using the Appliance Administration menu	565
8.2.1	Configure Network and DNS	566
8.2.1.1	Editing a connection to configure static IPv4 addressing	567
8.2.1.2	Edit a connection (Docker virtual bridge)	569
8.2.1.2.1	CIDR prefix lengths for common subnet masks	571
8.2.1.3	Activate a connection	572
8.2.1.4	Setting the system hostname	573
8.2.2	Configure IPv6 Network CIDR	574
8.2.3	Configure Timezone	575
8.2.4	Change Docker Hub Credentials	576
8.2.5	Change Root Password	577
8.2.6	Change ccuser Password	578
8.2.7	Update System	579
8.2.8	Change SSL settings	580
8.2.9	Root Shell	582
8.2.10	Reboot / Poweroff System	583

8.3	Managing Zope instances	584
8.3.1	Changing the number of Zope instances	585
8.3.2	Dedicated Zope service for reporting	586
8.3.2.1	Disabling automatic start of the reporting Zope service	587
8.3.2.2	Stopping the reporting Zope service	588
8.3.3	Dedicated Zope service for Zenoss JSON API use	589
8.3.3.1	Disabling automatic start of the Zenoss JSON API Zope service	590
8.3.3.2	Stopping the Zenoss JSON API Zope service	591
8.3.4	Dedicated Zope service for debugging	592
8.3.4.1	Enabling the Zope service for debugging	593
8.4	SNMP device preparation	594
8.4.1	Net-SNMP	595
8.4.2	SNMP v3 support	596
8.4.2.1	Advanced Encryption Standard	597
8.4.3	Community information	598
8.4.4	System contact information	599
8.4.5	Extra information	600
8.5	Syslog device preparation	601
8.5.1	Forwarding syslog messages from UNIX/Linux devices	602
8.5.2	Forwarding syslog messages from a Cisco IOS router	603
8.5.2.1	Other Cisco syslog configurations	604
8.5.3	Forwarding syslog messages from a Cisco CatOS switch	605
8.5.4	Forwarding syslog messages using syslog-ng	606
8.6	TALES expressions	607
8.6.1	TALES expression examples	608
8.6.2	TALES device attributes	609
8.6.3	TALES event attributes	611
8.7	Managing multi-realm networks	612
8.7.1	Configuring IP realms	613
8.7.2	Example multi-realm system	614
8.7.3	Prerequisites and considerations	615
8.8	Monitoring large file systems	616
8.8.1	Configuring the UCD dskTable MIB	617
8.9	Integrating LDAP authentication	618
8.9.1	LDAP configuration requirements	619
8.9.2	Adding an SSL certificate	620
8.9.3	Configuring LDAP authentication	621
8.9.4	Editing LDAP configurations	623
8.9.5	Configuring local authentication as a fallback	624
8.9.6	Verifying connectivity and credentials outside of Resource Manager	625
8.10	Tuning Considerations	626
8.10.1	Analytics ETL Services	627
8.10.2	Collection Services	628
8.10.3	Event Processing Services	630
8.10.4	Infrastructure Services	631
8.10.5	Metric Services	632
8.10.6	Tuning Control Center	633
8.10.7	User Interface Services	634
9.	Glossary	635

Planning a Resource Manager deployment

- [Welcome to Zenoss](#)
- [Resource Manager virtual appliances](#)
- [Control Center resource requirements](#)

Welcome to Zenoss

This section provides an overview of Zenoss software, including introductions to Control Center and Resource Manager, and a description of the installation options.

- [Introduction to Control Center](#)
- [Introduction to Zenoss Service Dynamics](#)
- [Installation options](#)

Introduction to Control Center

Control Center is an open-source application service orchestrator based on [Docker Community Edition](#) (Docker CE, or just Docker).

Control Center can manage any Docker application, from a simple web application to a multi-tiered stateful application stack. Control Center is based on a service-oriented architecture, which enables applications to run as a set of distributed services spanning hosts, datacenters, and geographic regions.

Control Center relies on declarations of application requirements to integrate Docker containers. A service definition template contains the specifications of application services in JSON format. The definition of each service includes the IDs of the Docker images needed to run the service.

Control Center includes the following key features:

- Intuitive HTML5 interface for deploying and managing applications
- Integrated backup and restore, and incremental snapshots and rollbacks
- Centralized logging through Logstash and Elasticsearch
- Integration with database services and other persistent services
- Encrypted communications among all services and containers
- Delegate host authentication to prevent unauthorized system access
- Storage monitoring and emergency shutdown of services to minimize the risk of data corruption
- Rolling restart of services to reduce downtime of multi-instance services
- Audit logging, including application audit logging

Docker fundamentals

This section summarizes [the architecture description provided by Docker](#) as customized for Control Center. For additional information, refer to the Docker site.

Docker provides convenient tools that make use of the [control groups feature of the Linux kernel](#) to develop, distribute, and run applications. Docker internals include images, registries, and containers.

Docker images

Docker images are read-only templates that are used to create Docker containers. Images are easy to build, and image updates are change layers, not wholesale replacements.

Docker registries

Docker registries hold images. Control Center uses a private Docker registry for its own images and Zenoss application images.

Docker containers

Docker containers have everything needed to run a service instance, and are created from images. Control Center launches each service instance in its own Docker container.

Docker storage

Docker and Control Center data are stored in customized LVM thin pools that are created from one or more block devices or partitions, or from one or more LVM volume groups.

Control Center internal services

Elasticsearch

A distributed, real-time search and analytics engine. Control Center uses it to index log files and store service definitions.

Kibana

A browser-based user interface that enables the display and search of Elasticsearch databases, including the log files that Control Center monitors.

Logstash

A log file collector and aggregator that forwards parsed log file entries to Elasticsearch.

OpenTSDB

A time series database that Control Center uses to store its service performance metrics.

ZooKeeper ([Apache ZooKeeper](#))

A centralized service that Control Center uses for configuration maintenance, naming, distributed synchronization, and providing group services.

ZooKeeper and Control Center

Control Center relies on [Apache ZooKeeper](#) to distribute and manage application services. ZooKeeper maintains the definitions of each service and the list of services assigned to each host. The scheduler, which runs on the master host, determines assignments and sends them to the ZooKeeper node that is serving as the ensemble leader. The leader replicates the assignments to the other ensemble nodes, so that the other nodes can assume the role of leader if the leader node fails.

All Control Center hosts retrieve assignments and service definitions from the ZooKeeper ensemble leader and then start services in Docker containers as required. So, the Control Center configuration files of all Control Center hosts must include a definition for the `SERVICED_ZK` variable, which specifies the ZooKeeper endpoints of the ensemble nodes. Additional variables are required on ensemble nodes.

A ZooKeeper ensemble requires a minimum of three nodes, which is sufficient for most environments. An odd number of nodes is recommended and an even number of nodes is strongly discouraged. A five-node ensemble improves failover protection during maintenance windows but larger ensembles yield no benefits.

The Control Center master host is always an ensemble node. All ensemble nodes should be on the same subnet.

Control Center application data storage

Control Center uses a dedicated LVM thin pool on the master host to store application data and snapshots of application data.

- The distributed file system (DFS) of each tenant application that serviced manages is stored in separate virtual devices. The initial size of each tenant device is copied from the base device, which is created during the initial startup of serviced.
- Snapshots of tenant data, used as temporary restore points, are stored in separate virtual devices, outside of tenant virtual devices. The size of a snapshot depends on the size of the tenant device, and grows over time.

The Control Center master host requires high-performance, persistent storage. Storage can be local or remote.

- For local storage, solid-state disk (SSD) devices are recommended.
- For remote storage, storage-area network (SAN) systems have been tested. High-performance SAN systems are recommended, as is assigning separate logical unit numbers (LUNs) for each mounted path.

The overall response times of master host storage affect the performance and stability of Control Center internal services and the applications it manages. For example, ZooKeeper (a key internal service) is sensitive to storage latency greater than 1000 milliseconds.

The physical devices associated with the application data thin pool must be persistent. If removable or re-connectable storage such as a SAN based on iSCSI is used, then the Device-Mapper Multipath feature of RHEL/CentOS must be configured and enabled.

Control Center includes the `serviced-storage` utility for creating and managing its thin pool. The `serviced-storage` utility can:

- use physical devices or partitions, or LVM volume groups, to create a new LVM thin pool for application data
- add space to a tenant device at any time
- identify and clean up orphaned snapshots
- create an LVM thin pool for Docker data

Introduction to Zenoss Service Dynamics

The foundation of the Zenoss Service Dynamics product suite is Zenoss Resource Manager, the hybrid IT monitoring platform. Starting with release 5.0, Zenoss Resource Manager (or simply, Resource Manager) is deployed into a distributed architecture using Control Center. That is, Resource Manager is an application whose installation and deployment configuration is managed by Control Center.

Control Center and Resource Manager are independent and unaware of each other, although Control Center is designed for the unique requirements of Resource Manager. In particular, Resource Manager and Control Center have a different version numbering scheme and different release schedules. Some releases of Resource Manager include features that rely on specific capabilities of Control Center. A separate series of guides describes how to use Control Center.

In the Resource Manager context, a resource is a device to monitor. Devices include compute, storage, network, converged infrastructure, applications, and unified communications systems. In the Control Center context, a resource is a physical or virtual host in a Control Center deployment.

Zenoss Service Impact (Service Impact) and Zenoss Analytics (Analytics) are the final components of the Zenoss Service Dynamics product suite.

- Service Impact correlates the infrastructure Resource Manager is monitoring with critical business services. In practical terms, Service Impact is additional services in a Resource Manager deployment, **not** a separate application that Control Center manages.
- Analytics adds data warehouse ETL (extract, transform, and load) capabilities and data analysis capabilities to Resource Manager. Like Service Impact, Analytics adds services to a Resource Manager deployment, but also requires a separate, standalone host that is **not** managed by Control Center.

Key Resource Manager concepts

collector

A set of Resource Manager services that provide monitoring capabilities for a specific network location. Collector services are deployed into a Control Center resource pool through the Resource Manager browser interface.

collector pool

A Control Center resource pool that is dedicated for use by one or more Resource Manager collectors; usually, just one. Most deployments use one collector pool for each distinct network location, although each collector pool typically includes multiple Control Center hosts.

collector pool VIP

A virtual IP address that is assigned a collector pool, to float among the hosts that are members of the pool. Resource Manager services that require a permanent IP address are unaffected by the status of specific pool hosts. The virtual IP address is created and assigned through the Control Center browser interface.

master pool

A Control Center resource pool that contains only the Control Center master host. By isolating the master host in its own pool and configuring a ZooKeeper ensemble, Resource Manager services are not affected when Control Center internal services are temporarily unavailable. The recommended name for this pool is `master`.

Resource Manager pool

A Control Center resource pool that runs Resource Manager infrastructure services and other key services. This resource pool must be on the same subnet as the master pool. A Resource Manager pool includes a minimum of two Control Center hosts. The recommended name is `Resource Manager`.

Installation options

You can create an on-premise deployment of Resource Manager by following one of two mutually-exclusive installation paths:

- Install Resource Manager virtual appliances as guest systems on a VMWare vSphere or Microsoft Hyper-V hypervisor. The guest systems include Control Center and require relatively little configuration. This is the recommended deployment option.
- Install Control Center on Red Hat Linux or CentOS hosts and add Resource Manager. This the most flexible installation path and it includes support for high-availability configurations.

For more information about additional installation options, please contact your Zenoss representative.

Resource Manager virtual appliances

Zenoss Service Dynamics is available as a pair of virtual appliances, one for a Control Center master host, and another for Control Center delegate hosts. Control Center and Resource Manager are installed, the required Docker images are loaded into the local registry, and the latest supported version of CentOS is installed with all of the packages needed to begin using Resource Manager as quickly as possible. In addition, the components of Zenoss Service Impact and Zenoss Analytics that integrate with Resource Manager are included and ready for final configuration.

Virtual appliances are packaged as Open Virtual Appliance (OVA) and ISO disk image files, and their resource requirements vary by hypervisor (VMware vSphere or Microsoft Hyper-V) and role (Control Center master or delegate).

To download Zenoss Service Dynamics virtual appliance files, you need permission to gain access to <https://delivery.zenoss.com>. For more information, contact Zenoss Support.

Resource requirements for multi-host deployments

Zenoss strongly recommends a minimum of 3 hosts for all production deployments of Zenoss Service Dynamics. Most environments require more than 3 hosts. Development or testing deployments can be single-host deployments.

CPU and memory requirements

The following table provides CPU and memory requirements for hosts in the types of Control Center resource pools that a typical Resource Manager deployment uses. For more information about pool types, see [Key Resource Manager concepts](#).

Resource pool type	Host count	Host resources	Comments
Master pool	1	4 CPU cores 16GB main memory	The amount of storage required for application data varies greatly.
Resource Manager pool	2+n	8 CPU cores 32GB main memory	To support failover or additional services, more than two hosts is preferred.
Collector pool	n+1	4 CPU cores 8GB main memory	Collector services do not need DFS; remote pools can use higher connection timeout settings.

Control Center master host storage

The Control Center master host virtual machine requires a total of 7 virtual hard disks. The following table identifies the purpose and size of each disk.

Purpose	Size
Root (/)	30GB
Swap	16GB
Temporary (/tmp)	16GB
Docker data	50GB
Control Center internal services data	50GB
Application data	200GB
Application data backups	400GB

On vSphere hypervisors, the disks are created when the OVA is installed. On Hyper-V hypervisors, the disks must be created manually.

Control Center delegate host storage

Control Center delegate host virtual machines requires a total of 4 virtual hard disks. The following table identifies the purpose and size of each disk.

Purpose	Size
Root (/)	30GB
Swap	16GB
Temporary (/tmp)	16GB
Docker data	50GB

On vSphere hypervisors, the disks are created when the OVA is installed. On Hyper-V hypervisors, the disks must be created manually.

Resource requirements for single-host deployments

A single-host deployment of the Zenoss Service Dynamics virtual appliance requires a virtual machine with a minimum of 8 CPU cores and 64GB of main memory.

The disk requirements are identical to the master host disk requirements for multi-host deployments. On vSphere hypervisors, the disks are created when the OVA is installed. On Hyper-V hypervisors, the disks must be created manually.

Control Center resource requirements

This section describes the CPU, RAM, storage, and operating system requirements for creating a Control Center deployment for Resource Manager.

- [Installation considerations](#)
- [Compute and storage requirements](#)
- [Operating system requirements](#)

Installation considerations

All Resource Manager data is stored on the Control Center master host. The Resource Manager application, running on delegate hosts, accesses the data through the distributed file system, which is based on NFS.

Using hypervisor commands alone to pause or stop Resource Manager virtual machines is unsupported. Resource Manager relies on timestamps and the system clock to keep services in sync, and pausing or stopping a virtual machine by using a hypervisor command disrupts the synchronization. Zenoss recommends the following procedure for pausing or stopping Resource Manager virtual machines:

1. Log in to the Control Center browser interface.
2. Stop Resource Manager.
3. Use a hypervisor feature to shut down the virtual machine, or log in to the virtual machine as root and enter a shutdown command.

Similarly, vSphere vMotion is not supported unless all of the virtual machines in your Resource Manager deployment are paused or stopped.

vSphere hosts that run Control Center guest systems must be configured to synchronize their clocks with public or private NTP servers. Control Center guest systems synchronize their clocks with their vSphere hosts through an hourly invocation of VMware Tools. For more information about configuring a vSphere host for NTP, refer to your VMware documentation.

Multi-host deployments of Resource Manager running on Hyper-V hosts must be configured to synchronize their clocks with public or private NTP servers. The Resource Manager Installation Guide includes instructions for configuring NTP on Control Center guest systems. Hyper-V hosts do not provide the equivalent of VMware Tools so that guest systems can synchronize with the host.

Control Center includes backup and restore features for archiving and restoring Resource Manager data. The best practice is to use third party backup tools to archive the backup file.

Hypervisor backups can only be used instead of Control Center backups when Resource Manager, Control Center, and the Control Center master host are shut down cleanly and completely. Do not rely on hypervisor backups otherwise.

Compute and storage requirements

Control Center requires real or virtual master and delegate hosts that

- implement the 64-bit version of the x86 instruction set
- support Red Hat Enterprise Linux (RHEL) 7.x or CentOS 7.x
- support Advanced Encryption Standard (AES)
- include a network interface controller that supports TCP/IP and IPv4

If Control Center hosts are virtual, they should be configured **not** to fail over to other hosts or storage.

Hardware resource and storage requirements for Control Center vary by role (master or delegate host) and by the services assigned to the resource pool to which a host belongs. This section provides minimum requirements for master and delegate hosts.

Master host CPU and RAM resources

You can create a multi-host or single-host deployment of Control Center, on real or virtual hosts.

Zenoss recommends that all production deployments include one Control Center master host and two or more delegate hosts. In this configuration, the master host runs in its own, separate resource pool, and runs only Control Center internal services. Delegate hosts run application services in other resource pools.

- For multi-host deployments, the master host typically requires 4 real or virtual CPU cores and 16GB RAM. Very large multi-host deployments may require 8 CPU cores, but no additional RAM.
- For single-host deployments, the master host requires a minimum of 8 real or virtual CPU cores and 64GB RAM. Actual CPU and RAM requirements depend on application load.

An under-resourced master host cannot function properly. Deploy Control Center and Resource Manager on a master host that meets or exceeds the minimum requirements.

Master host storage areas

A Control Center master host requires the following storage areas:

Root

Size: Depends on configuration
Type: XFS file system (typically)
Mount point: /
Resource: One or more block devices or partitions, or one or more LVM physical volumes.
Preparation: None. The file system is created when the operating system is installed.

Docker temporary

Size: 10GB
Type: XFS file system (typically)
Mount point: /tmp
Resource: One or more block devices or partitions, or one or more LVM physical volumes.
Preparation: Depends on configuration. Typically, additional space is allocated for the root filesystem, but a separate real or logical partition may be used. The Control Center installation procedures include instructions to link the Docker temporary directory to /tmp.

Swap

Size: 12GB to 16GB
Type: swap
Mount point: None
Resource: One or more block devices or partitions, one or more LVM physical volumes, or a special file on the root filesystem.
Preparation: Depends on configuration.

Docker data

Size: 50GB
Type: LVM thin pool
Mount point: None
Resource: One or more block devices or partitions, or one or more LVM physical volumes.
Preparation: None. The installation procedures include steps for creating the thin pool.

Control Center internal services data

Size: 50GB
Type: XFS file system
Mount point: /opt/serviced/var/isvcs
Resource: One or more block devices or partitions, or one or more LVM physical volumes.
Preparation: None. The installation procedures include steps for formatting the resource.

Control Center audit logging

Size: 10GB (default)
Type: XFS file system
Mount point: /var/log/serviced
Resource: One or more block devices or partitions, one or more LVM physical volumes, or a remote file server that is compatible with Linux.
Preparation: Depends on configuration. Typically, additional space is allocated for the root filesystem. For more information, see [Control Center audit logging](#).

Application data

Size: 200GB suggested. The size should be twice as large as the base device, which determines the size of tenant virtual devices. The ratio of application data stored in the thin pool to the total amount of storage in the pool should always be 50%. When the amount of available space is less than double the amount of total storage, the snapshot feature is prone to failure.
Type: LVM thin pool
Mount point: None
Resource: One or more block devices or partitions, or one or more LVM physical volumes.
Preparation: None. The installation procedures include steps for creating the thin pool.

Application data backups

Size: 150GB suggested. The size should be at least 150% of the size of the base device, or a minimum of 150GB, whichever is greater. Note: For large environments, this size should be much greater. Individual backup files can be 100GB each, or more.
Type: XFS file system, or a Linux-compatible file server
Mount point: /opt/serviced/var/backups
Resource: One or more block devices or partitions, one or more LVM physical volumes, or a remote file server that is compatible with Linux. Note: When the storage for backups and Control Center internal services data use the same physical device, resource contention can cause failures during backup and restore operations. For optimal results, use separate physical devices.
Preparation: None. The installation procedures include steps for formatting or mounting the resource.

The following example shows the disk configuration of a Control Center master host.

```

# lsblk -ap --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
NAME                               SIZE TYPE FSTYPE      MOUNTPOINT
/dev/fd0                            4K disk
/dev/sda                           29.3G disk
/dev/sda1                          29.3G part xfs          /
/dev/sdb                            48.8G disk
/dev/sdb1                          48.8G part LVM2_member
  /dev/mapper/docker-docker--pool_tmeta 900M lvm
  /dev/mapper/docker-docker--pool      42.2G lvm
    /dev/mapper/docker-8:1-42048-7a049d85cad6c 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /7a049d85cad6c
    /dev/mapper/docker-8:1-42048-0f93eb8de36bb 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /0f93eb8de36bb
    /dev/mapper/docker-8:1-42048-d84939118b6de 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /d84939118b6de
    /dev/mapper/docker-8:1-42048-2daa75d92947e 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /2daa75d92947e
    /dev/mapper/docker-8:1-42048-e7277d3081955 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /e7277d3081955
    /dev/mapper/docker-8:1-42048-e91db5ccdcf21 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /e91db5ccdcf21
    /dev/mapper/docker-8:1-42048-99104ce4a2d39 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /99104ce4a2d39
    /dev/mapper/docker-8:1-42048-a29f148212846 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /a29f148212846
  /dev/mapper/docker-docker--pool_tdata 42.2G lvm
  /dev/mapper/docker-docker--pool      42.2G lvm
    /dev/mapper/docker-8:1-42048-7a049d85cad6c 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /7a049d85cad6c
    /dev/mapper/docker-8:1-42048-0f93eb8de36bb 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /0f93eb8de36bb
    /dev/mapper/docker-8:1-42048-d84939118b6de 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /d84939118b6de
    /dev/mapper/docker-8:1-42048-2daa75d92947e 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /2daa75d92947e
    /dev/mapper/docker-8:1-42048-e7277d3081955 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /e7277d3081955
    /dev/mapper/docker-8:1-42048-e91db5ccdcf21 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /e91db5ccdcf21
    /dev/mapper/docker-8:1-42048-99104ce4a2d39 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /99104ce4a2d39
    /dev/mapper/docker-8:1-42048-a29f148212846 45G dm  xfs          /var/lib/docker/devicemapper/mnt
  /a29f148212846
/dev/sdc                            15.6G disk
/dev/sdc1                          15.6G part swap      [SWAP]
/dev/sdd                            15.6G disk
/dev/sdd1                          15.6G part xfs       /tmp
/dev/sde                            48.8G disk
/dev/sde1                          48.8G part xfs       /opt/serviced/var/isvcs
/dev/sdf                            149.2G disk
/dev/sdf1                          149.2G part xfs      /opt/serviced/var/backups
/dev/sdg                            195.3G disk
/dev/sdg1                          195.3G part LVM2_member
  /dev/mapper/serviced-serviced--pool_tmeta 1.8G lvm
  /dev/mapper/serviced-serviced--pool      172.3G lvm
    /dev/mapper/docker-8:1-33639769-3314GP 90G dm  ext4         /exports/serviced_volumes_v2
  /a9hil5o55iy266s5l
    /dev/mapper/serviced-serviced--pool_tdata 172.3G lvm
    /dev/mapper/serviced-serviced--pool      172.3G lvm
    /dev/mapper/docker-8:1-33639769-3314GP 90G dm  ext4         /exports/serviced_volumes_v2
  /a9hil5o55iy266s5l
/dev/sr0                            1024M rom

```

Delegate host CPU and RAM resources

The following table identifies resource requirements for delegate hosts. For more information about pool types, see [Key Resource Manager concepts](#).

Resource pool type	Host count	Host resources	Comments
Resource Manager pool	2+n	8 CPU cores 32GB main memory	To support failover or additional services, more than two hosts is preferred.
Collector pool	n+1	4 CPU cores 8GB main memory	Collector services do not need DFS; remote pools can use higher connection timeout settings.

Delegate host storage requirements

Like master hosts, Control Center delegate hosts require high-performance, persistent storage. Storage can be local or remote.

- For local storage, solid-state disk (SSD) devices are recommended.
- For remote storage, storage-area network (SAN) systems have been tested. High-performance SAN systems are recommended, as is assigning separate logical unit numbers (LUNs) for each mounted path.

The following storage configuration is recommended for delegate hosts:

- root filesystem with a minimum of 30GB of storage, formatted with XFS
- dedicated swap area
- one or more block devices or partitions, or one or more LVM physical volumes, with a total of 50GB of space. The installation procedures include steps for using serviced-storage to create an LVM thin pool for Docker data.

The following example shows the disk configuration of a Control Center delegate host.

```
# lsblk -ap --output=NAME,SIZE,TYPE,FSTYPE,MOUNTPOINT
NAME                               SIZE TYPE FSTYPE      MOUNTPOINT
/dev/fd0                            4K disk
/dev/sda                          29.3G disk
/dev/sda1                          29.3G part xfs        /
/dev/sdb                           48.8G disk
/dev/sdb1                          48.8G part LVM2_member
  /dev/mapper/docker-docker--pool_tmeta 900M lvm
  /dev/mapper/docker-docker--pool      42.2G lvm
  /dev/mapper/docker-8:1-16786822-2a6c2771415 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/2a6c2771415
  /dev/mapper/docker-8:1-16786822-579a2fb3611 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/579a2fb3611
  /dev/mapper/docker-8:1-16786822-c95ca5f04a4 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/c95ca5f04a4
  /dev/mapper/docker-8:1-16786822-e9733ecdda1 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/e9733ecdda1
  /dev/mapper/docker-8:1-16786822-947636812ed 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/947636812ed
  /dev/mapper/docker-8:1-16786822-d65a56e4300 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/d65a56e4300
  /dev/mapper/docker-8:1-16786822-3d1fb13572c 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/3d1fb13572c
  /dev/mapper/docker-8:1-16786822-542e3e08ccc 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/542e3e08ccc
  /dev/mapper/docker-8:1-16786822-3beb7120639 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/3beb7120639
  /dev/mapper/docker-8:1-16786822-1c88db1b91c 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/1c88db1b91c
  /dev/mapper/docker-8:1-16786822-3a246ac134c 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/3a246ac134c
  /dev/mapper/docker-8:1-16786822-045933b4418 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/045933b4418
  /dev/mapper/docker-8:1-16786822-e420dfc0229 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/e420dfc0229
  /dev/mapper/docker-8:1-16786822-918a6f92734 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/918a6f92734
  /dev/mapper/docker-8:1-16786822-a7efd4f15c3 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/a7efd4f15c3
  /dev/mapper/docker-8:1-16786822-17b847994e2 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/17b847994e2
  /dev/mapper/docker-8:1-16786822-3564e748e14 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/3564e748e14
  /dev/mapper/docker-docker--pool_tdata 42.2G lvm
  /dev/mapper/docker-docker--pool      42.2G lvm
  /dev/mapper/docker-8:1-16786822-2a6c2771415 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/2a6c2771415
  /dev/mapper/docker-8:1-16786822-579a2fb3611 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/579a2fb3611
  /dev/mapper/docker-8:1-16786822-c95ca5f04a4 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/c95ca5f04a4
  /dev/mapper/docker-8:1-16786822-e9733ecdda1 45G dm  xfs        /var/lib/docker/devicemapper/mnt
/e9733ecdda1
  /dev/mapper/docker-8:1-16786822-947636812ed 45G dm  xfs        /var/lib/docker/devicemapper/mnt
```

/947636812ed					
/dev/mapper/docker-8:1-16786822-d65a56e4300	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/d65a56e4300					
/dev/mapper/docker-8:1-16786822-3d1fb13572c	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/3d1fb13572c					
/dev/mapper/docker-8:1-16786822-542e3e08ccc	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/542e3e08ccc					
/dev/mapper/docker-8:1-16786822-3beb7120639	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/3beb7120639					
/dev/mapper/docker-8:1-16786822-1c88db1b91c	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/1c88db1b91c					
/dev/mapper/docker-8:1-16786822-3a246ac134c	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/3a246ac134c					
/dev/mapper/docker-8:1-16786822-045933b4418	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/045933b4418					
/dev/mapper/docker-8:1-16786822-e420dfc0229	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/e420dfc0229					
/dev/mapper/docker-8:1-16786822-918a6f92734	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/918a6f92734					
/dev/mapper/docker-8:1-16786822-a7efd4f15c3	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/a7efd4f15c3					
/dev/mapper/docker-8:1-16786822-17b847994e2	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/17b847994e2					
/dev/mapper/docker-8:1-16786822-3564e748e14	45G	dm	xfs		/var/lib/docker/devicemapper/mnt
/3564e748e14					
/dev/sdc	15.6G	disk			
/dev/sdc1	15.6G	part	swap		[SWAP]
/dev/sdd	15.6G	disk			
/dev/sdd1	15.6G	part	xfs		/tmp
/dev/sr0	1024M	rom			

Operating system requirements

Control Center has been tested the 64-bit version of the following Linux distributions:

- Red Hat Enterprise Linux (RHEL) 7.2 and above
- CentOS 7.2 and above

Kernel versions 3.10.0-327.22.2.el7.x86_64 and higher are tested. For best performance, keep the kernel up-to-date.

The RHEL/CentOS 7 distributions provide a variety of server configurations. Control Center is tested on operating system platforms that are installed and configured with standard options. Docker and Control Center are tested on the Minimal Install configuration with the NFS and Network Time Protocol (NTP) packages installed.

Control Center relies on the system clock to synchronize its actions. The installation procedures include steps to add the NTP daemon to all hosts. By default, the NTP daemon synchronizes the system clock by communicating with standard time servers available on the internet. Use the default configuration or configure the daemon to use a time server in your environment.

Because Control Center relies on the system clock, while an application is running, do not pause a virtual machine that is running Control Center.

Networking

On startup, Docker creates the `docker0` virtual interface and selects an unused IP address and subnet (typically, `172.17.0.1/16`) to assign to the interface. The virtual interface is used as a virtual Ethernet bridge, and automatically forwards packets among real and virtual interfaces attached to it. The host and all of its containers communicate through this virtual bridge.

Docker can only check directly connected routes, so the subnet it chooses for the virtual bridge might be inappropriate for your environment. To customize the virtual bridge subnet, refer to Docker's [advanced network configuration](#) article.

The following list highlights potential communication conflicts:

- If you use a firewall utility, ensure that it does not conflict with Docker. The default configurations of firewall utilities such as [Firewalld](#) include rules that can conflict with Docker, and therefore Control Center. The following interactions illustrate the conflicts:
 - The `firewalld` daemon removes the `DOCKER` chain from iptables when it starts or restarts.
 - Under `systemd`, `firewalld` is started before Docker. However, if you start or restart `firewalld` while Docker is running, you must restart Docker.
- Even if you do not use a firewall utility, your firewall settings might still prevent communications over the Docker virtual bridge. This issue occurs when iptables INPUT rules restrict most traffic. To ensure that the bridge works properly, append an INPUT rule to your iptables configuration that allows traffic on the bridge subnet. For example, if `docker0` is bound to `172.17.42.1/16`, then a command like the following example would ensure that the bridge works.

Note: Before modifying your iptables configuration, consult your networking specialist.

```
iptables -A INPUT -d 172.17.0.0/16 -j ACCEPT
```

Additional requirements and considerations

Control Center requires a 16-bit, private IPv4 network for virtual IP addresses. The default network is `10.3/16`, but during installation you can select any valid IPv4 16-bit address space.

Before installation, add DNS entries for the Control Center master host and all delegate hosts. Verify that all hosts in Control Center resource pools can

- Resolve the hostnames of all other delegate hosts to IPv4 addresses. For example, if the public IP address of your host is `192.0.2.1`, then the host name `-i` command should return `192.0.2.1`.
- Respond with an IPv4 address other than `127.x.x.x` when ping Hostname is invoked.
- Return a unique result from the `hostid` command.

Control Center relies on Network File System (NFS) for its distributed file system implementation. Therefore, Control Center hosts cannot run a general-purpose NFS server, and all Control Center hosts require NFS.

Disabling IPv6 can prevent the NFS server from restarting, due to [an rpcbind issue](#). Zenoss recommends leaving IPv6 enabled on the Control Center master host.

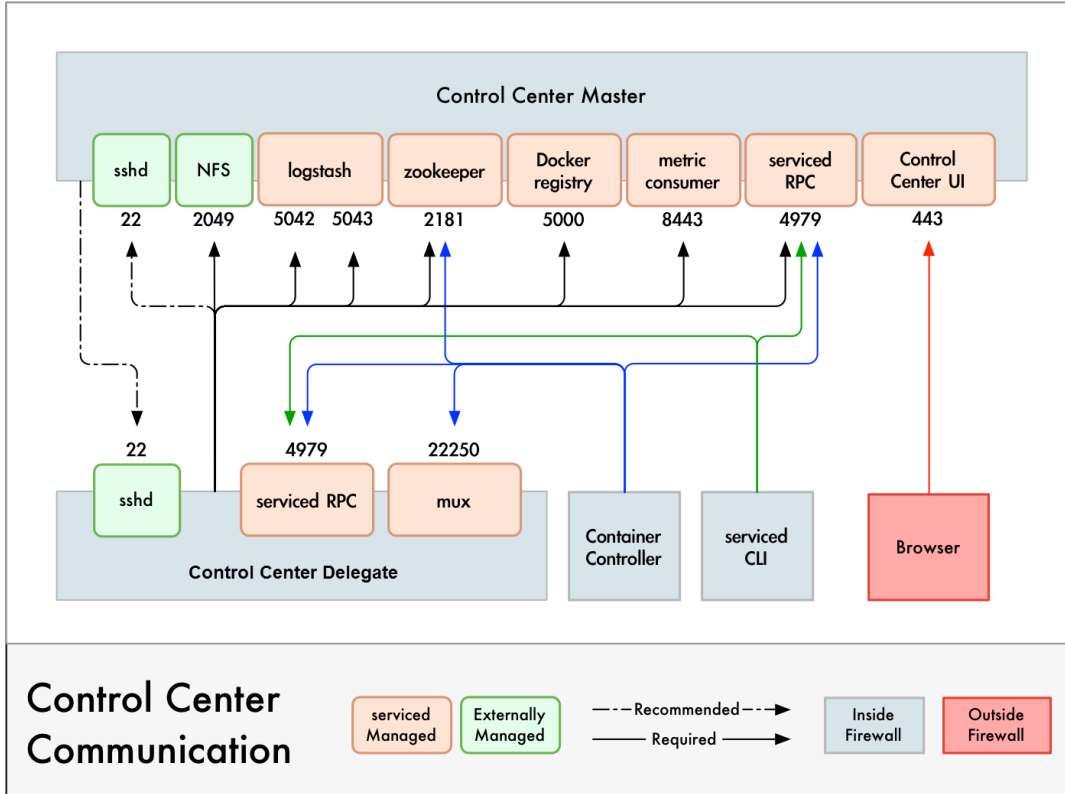
Security

During installation, Control Center has no knowledge of Resource Manager port requirements, so the installation procedure includes disabling the firewall. After both Control Center and Resource Manager are installed, you can close unused ports.

Control Center includes a virtual multiplexer (mux) that performs the following functions:

- Aggregates the UDP and TCP traffic among the services it manages. The aggregation is opaque to services, and mux traffic is encrypted when it travels among containers on remote hosts. (Traffic among containers on the same host is not encrypted.)
- Along with the distributed file system, enables Control Center to quickly deploy services to any pool host.
- Reduces the number of open ports required on a Control Center host to a predictable set.

The following figure identifies the ports that Control Center requires. All traffic is TCP. Except for port 4979, all ports are configurable.



Control Center relies on the system clock to synchronize its actions, and indirectly, `ntpd` or `chrony` to synchronize clocks among multiple hosts. In the default configuration of `ntpd`, the firewalls of master and delegate hosts must support an incoming UDP connection on port 123.

For more information about port requirements, see [Resource Manager network ports](#).


Additional requirements and considerations

- To install Control Center, you must log in as root, or as a user with superuser privileges.
- Access to the Control Center browser interface requires a login account on the Control Center master host. Pluggable Authentication Modules (PAM) is tested. By default, users must be members of the wheel group. The default group may be changed by setting the `SERVICED_ADMIN_GROUP` variable, and the replacement group does not need superuser privileges.
- The serviced startup script sets the hard and soft open files limit to 1048576. The script does not modify the `/etc/sysconfig/limits.conf` file.
- Control Center has been tested with [Security Enhanced Linux](#) enabled.

Resource Manager network ports

This section includes a network diagram of a Resource Manager deployment featuring four Control Center resource pools:

- One pool for the Control Center master host.
- One pool for most of the Resource Manager services.
- Two pools for Resource Manager collector services (collector pools).



High-availability configurations have additional port requirements. For more information, see the planning chapter in the high-availability installation guide ([Control Center 1.5.x](#)).

Installing Resource Manager

You can install Resource Manager by using virtual appliances or by adding it to an existing Control Center deployment.

Keep this page open, and open new tabs or windows for each installation procedure.

The procedures on this page do not conclude with instructions for starting Resource Manager. Instead, a link sends you along to the configuration procedures, documented in another section.

Adding Resource Manager to Control Center

Follow these steps to add Resource Manager to a Control Center deployment. For more information about creating a Control Center deployment, see [Planning a Resource Manager deployment](#).

1. [Download template and image files](#)
2. [Install the Resource Manager template](#)
3. [Import Resource Manager images](#)
4. [Deploy Resource Manager in Control Center](#)
5. [Configure resource pool permissions](#)

Installing with virtual appliances

Perform the procedures in the following sections to install Resource Manager virtual appliances, pre-configured virtual machines that are ready to deploy to your hypervisor.

Installing a master host

Follow these steps to install a Resource Manager virtual appliance as a Control Center master host. All Resource Manager deployments require a Control Center master host.

1. Create a master host virtual machine on your hypervisor:
 - vSphere: [Creating a master host with vSphere](#)
 - Hyper-V:
 - a. [Creating a master host with Hyper-V](#)
 - b. [Configuring and starting a Hyper-V master host](#)
2. Optional: [Replace DHCP with static addressing](#)
3. Optional: [Replace the default hostname](#)
4. [Configure the Control Center master host](#)
5. Multi-host deployments: [Create resource pools](#)
6. [Add the master host to a resource pool](#)
7. [Deploy Resource Manager](#)

Installing delegate hosts

Perform these steps to install Resource Manager virtual appliances as Control Center delegate hosts. You can install as many delegate hosts as you need.

1. Create a delegate host virtual machine on your hypervisor:
 - vSphere: [Creating a delegate host with vSphere](#)
 - Hyper-V:
 - a. [Creating a delegate host with Hyper-V](#)
 - b. [Configuring and starting a Hyper-V delegate host](#)
2. [Configure the delegate host](#)
3. Optional: [Replace DHCP with static addressing](#)
4. Optional: [Replace the default hostname](#)
5. Optional: [Add delegate host FQDNs to /etc/hosts](#)
6. [Configure delegate host authentication](#)

Configuring a multi-host deployment

Perform these steps to configure a multi-host deployment.

Zenoss recommends deploying at least two delegate hosts, to enable creating a ZooKeeper ensemble.

1. (Hyper-V systems only) [Enable NTP](#)
 - a. [Configuring NTP clients](#)
 - b. [Configuring an NTP master server](#)
 - c. [Configuring NTP for public time servers](#)
2. [Configure ZooKeeper](#)

3. [Configure resource pool permissions](#)

On appliance hosts, the `/etc/default/serviced` file includes only the minimum required variables to effect an installation. The complete, original contents of the configuration file are included in `/etc/default/serviced.orig`. For more information about the Control Center configuration file, see [Configuration variables](#).

Adding Resource Manager to a Control Center deployment

For the complete list of installation instructions, see [Installing Resource Manager](#).

Downloading template and image files

To perform this procedure, you need:

- A workstation with internet access
- Permission to download files from delivery.zenoss.com. Customers can request permission by filing a ticket at the [Zenoss Support](#) site.
- A secure network copy program

Use this procedure to

- download required files to a workstation
- copy the files to a Control Center master host

Follow these steps:

1. In a web browser, navigate to delivery.zenoss.com, and then log in.
2. Download the self-installing Docker image files for Resource Manager.
Replace MAJOR.MINOR with the major and minor numbers of this version (for example, 6.4) and replace VERSION with the entire version number (6.4.1):
 - `install-zenoss-hbase-LATEST.run`
 - `install-zenoss-opentsdb-LATEST.run`
 - `install-zenoss-resmgr_MAJOR.MINOR-VERSION_1.run`
 - `install-zenoss-mariadb:10.1-VERSION.run`
3. Download the Resource Manager service definition, which is distributed as an RPM file.
Replace *VERSION* with the current release number (for example, 6.4.1).

```
zenoss-resmgr-service-VERSION-1.noarch.rpm
```

4. Use a secure copy program to copy the files to the Control Center master host.

Installing the Resource Manager template

Use this procedure to install the Resource Manager service definition template on the Control Center master host.

1. Log in to the Control Center master host as root, or as a user with superuser privileges.
2. Move the RPM file to /tmp.
Replace *VERSION* with the current release number (for example, 6.4.1).

```
mv zenoss-resmgr-service-VERSION-1.noarch.rpm /tmp
```

3. Install the Resource Manager template file.
Replace *VERSION* with the current release number (for example, 6.4.1).

```
yum install /tmp/zenoss-resmgr-service-VERSION-1.noarch.rpm
```

The template file is stored in /opt/serviced/templates.

Importing Resource Manager images

Use this procedure to import the Resource Manager images into the local registry.

1. Log in to the Control Center master host as root, or as a user with superuser privileges.
2. Move the Docker image files to `/root`.

```
mv install-zenoss-*.run /root
```

3. Add execute permission to the image files.

```
chmod +x /root/install-zenoss-*.run
```

4. Change directory to `/root`.

```
cd /root
```

5. Import the images.

The images are contained in self-extracting archive files.

```
for image in install-zenoss-*.run
do
  /bin/echo -en "\nLoading $image..."
  yes | ./$image
done
```

6. List the images in the registry.

```
docker images
```

The result should include one image for each archive file.

7. Optional: Delete the archive files.

```
rm -i ./install-zenoss-*.run
```

Deploying Resource Manager in Control Center

Use this procedure to add the Resource Manager application to Control Center, and to tag application images in the local registry.

1. Log in to the Control Center master host as a user with serviced CLI privileges.
2. Add the Resource Manager application template to Control Center.

```
serviced template add /opt/serviced/templates/zenoss*.json
```

On success, the `serviced` command returns the template ID.

3. Identify the resource pool to which the host belongs.

```
serviced host list
```

4. Deploy the application.
 - Replace Template-ID with the identifier of the Resource Manager application template (returned in step 2)
 - Replace Pool with the name of the resource pool to which the master host belongs (single-host system) or to which the delegate hosts belong (multi-host system)
 - Replace Deployment with a name for this deployment (for example, Test or Production)

```
serviced template deploy Template-ID Pool Deployment
```

Control Center tags Resource Manager images in the local registry.

Resource Manager is ready to be configured for your environment. For more information, see [Configuring Resource Manager](#).

Installing a master host

For the complete list of installation instructions, see [Installing Resource Manager](#).

Deploying Resource Manager

Use this procedure to add the Resource Manager application to Control Center and tag application images in the local registry.

1. Log in to the Control Center master host as a user with serviced CLI privileges.
2. Add the Resource Manager application template to Control Center.

```
serviced template add /opt/serviced/templates/zenoss*.json
```

On success, the `serviced` command returns the template ID.

3. Identify the resource pool to which the host belongs.

```
serviced host list
```

4. Deploy the application.
 - Replace Template-ID with the identifier of the Resource Manager application template (returned in step 2)
 - Replace Pool with the name of the resource pool to which the master host belongs (single-host system) or to which the delegate hosts belong (multi-host system)
 - Replace Deployment-ID with a name for this deployment (for example, Test or Production)

```
serviced template deploy Template-ID Pool Deployment-ID
```

Control Center tags Resource Manager images in the local registry.

- If you are creating a single-host deployment, proceed to [Configuring Resource Manager](#).
- If you are creating a multi-host deployment, proceed to [Installing delegate hosts](#).

Adding the master host to a resource pool

Complete this procedure to add the Control Center master host to the `resmgr` resource pool.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Add the master host to the `resmgr` resource pool.
Replace `Hostname-Or-IP` with the hostname or IP address of the Control Center master host:

```
serviced host add --register Hostname-Or-IP:4979 resmgr
```

If you enter a hostname, all hosts in your Control Center deployment must be able to resolve the name, either through an entry in `/etc/hosts`, or through a nameserver on your network.

4. To exit the command-line session, at the command prompt, enter `exit`.

Configuring the Control Center master host

Perform this procedure immediately after creating and starting the Control Center master host. All Control Center deployments must include one master host.

1. Gain access to the console interface of the Control Center host through your hypervisor console interface.

```
YOU HAVE NOT ACTIVATED THIS APPLIANCE.  
PLEASE LOGIN TO ACTIVATE Zenoss Service Dynamics  
  
Welcome to Zenoss Service Dynamics  
  
After initial setup, the Control Center UI can be accessed by  
browsing to:  
  
https://zsd-master  
  
Ensure that zsd-master is resolvable to 10.88.121.241, either through your  
DNS system or through a HOSTS entry on the browser client. For more  
information refer to the installation notes.  
  
You can log in to this console to perform administrative tasks such  
as setting up networking and safely rebooting this system.  
  
Control Center UI login credentials are ccuser/zsd.  
To log into the console as root, use the default password 'zsd'.  
  
Linux Kernel 3.10.0-693.2.2.el7.x86_64 on an x86_64  
zsd-master login: _
```

2. Log in as the `root` user.
The initial password is provided in the console.
3. The system prompts you to enter a new password for `root`.
4. The system prompts you to enter a new password for `ccuser`.
The `ccuser` account is the default account for gaining access to the Control Center browser interface.

Creating a master host with Hyper-V

To perform this task, you need:

- Microsoft Remote Desktop Connection
- Administrator privileges on a Microsoft Hyper-V server
- Permission to download files from delivery.zenoss.com.

Use this procedure to install the Resource Manager master host appliance as a virtual machine managed by Microsoft Hyper-V.

1. Use a Microsoft Remote Desktop Connection to log in to a Hyper-V host as Administrator, or as a user with Administrator privileges.
2. In a web browser, navigate to delivery.zenoss.com, and then log in.
3. Download the Resource Manager master host ISO file.
Replace *VERSION* with the current release number (for example, 6.4.1).

```
zenoss-zsd-VERSION-1-master.x86_64.iso
```

4. Open Hyper-V Manager.
5. In the left navigation pane, choose a server to host the virtual machine.
6. From the Action menu, choose New > Virtual Machine.
The New Virtual Machine Wizard opens.
7. In the New Virtual Machine Wizard dialog, display the Specify Name and Location panel.
If the first panel displayed is the Before You Begin panel, click Next.
8. In the Specify Name and Location panel, provide a name for the virtual machine, and then click Next.
9. In the Specify Generation panel, choose Generation 1, and then click Next.
10. In the Assign Memory panel, specify the memory for the virtual machine.
 - a. In the Startup memory field, enter the amount of memory for the host.
 - For multi-host deployments, enter 16384 (16GB).
 - For single-host deployments, enter 32768 (32GB).
 - b. Optional: Check Use Dynamic Memory for this virtual machine.
Resource Manager is tested with dynamic memory enabled.
 - c. Click Next.
11. In the Configure Networking panel, choose a virtual switch, and then click Next.
12. In the Connect Virtual Hard Disk panel, specify a new disk on which to install the guest operating system.
 - a. Choose Create a virtual hard disk.
 - b. Specify a name.
 - c. In the Size field, enter 30.
 - d. Click Next.
13. In the Installation Options panel, specify the master host ISO file.
 - a. Choose Install an operating system from a bootable CD/DVD-ROM.
 - b. Choose Image file (.iso).
 - c. Specify or browse to the location of the master host ISO file.
 - d. Click Next.
14. In the Completing the New Virtual Machine Wizard panel, verify the description, and then click Finish.

Hyper-V Manager creates the new virtual machine, and then closes the wizard.

Proceed to [Configuring and starting a Hyper-V master host](#).

Configuring and starting a Hyper-V master host

To perform this task, you need:

- A Microsoft Remote Desktop Connection
- Administrator privileges on a Microsoft Hyper-V server
- The master host created in the previous procedure ([Creating a master host with Hyper-V](#))


The Resource Manager master host virtual machine requires a total of 7 virtual hard disks. More information about required disks: [Resource requirements for multi-host deployments](#).

Use this procedure to configure hardware resources, create disks 2-6, and start the master host.

1. Use a Microsoft Remote Desktop Connection to log in to a Hyper-V host as Administrator, or as a user with Administrator privileges.
2. Open Hyper-V Manager.
3. In the Hyper-V Manager Virtual Machines area, right-click the new virtual machine, and then choose Settings.
The Settings dialog displays.
4. In the Hardware area, locate the virtual hard disk created previously, and then determine whether it is attached to an IDE controller. Hyper-V guest machines can only boot from an IDE drive.
5. In the Hardware area, choose Processor, and then change the number of processors assigned to the machine.
 - a. In Number of virtual processors, enter the value for your deployment.
 - For single-host deployments, enter 8.
 - For multi-host deployments, enter 4.
 - b. Click Apply.
6. In the Hardware area on the left, choose SCSI Controller, and then create additional virtual hard disks. Repeat the following substeps to create new disks in the following sizes:
 - 50GB
 - 400GB
 - 50GB
 - 200GB
 - 16GB
 - 16GB
 - a. In the controller area on the right, choose Hard Drive, and then click Add.
 - b. In the Location field, choose an unused location number.
 - c. In the Media area, choose Virtual hard disk, and then click New.
 - d. Complete panels in the New Virtual Hard Disk Wizard as follows:
 - i. Choose Disk Format: Choose VHDX, and then click Next.
 - ii. Choose Disk Type: Choose Dynamically expanding, and then click Next.
 - iii. Specify Name and Location: Enter the disk name, and then click Next.
 - iv. Configure Disk:
 1. Choose Create a new blank virtual hard disk.
 2. Size: Enter a disk size from the list at the beginning of this step.
 3. Click Next.
 - v. Summary/New Virtual Hard Disk Wizard: Verify the description, and then click Finish.
 - e. At the bottom of the Settings window, click Apply.

When all of the disks are created, click OK.

7. In the Hyper-V Manager Virtual Machines area, right-click the new virtual machine, and then choose Start.



8. In the Hyper-V Manager Virtual Machines area, right-click the new virtual machine, and then choose Connect.
9. In the Virtual Machine Connection window, press Enter.
The appliance installation process takes about 15 minutes, and should complete with no additional input. If received, disregard the Fast TSC calibration failure message.

Proceed to [Configuring the Control Center master host](#).

Creating a master host with vSphere

To perform this task, you need:


- A VMware vSphere client
- Permission to download files from delivery.zenoss.com. Customers can request permission by filing a ticket at the Zenoss Support site.

This procedure installs Resource Manager OVA packages as a virtual machine managed by vSphere Server version 6.5.0, using VMware vSphere Web Client 6.5. The procedure might differ with other versions of VMware vSphere Client.

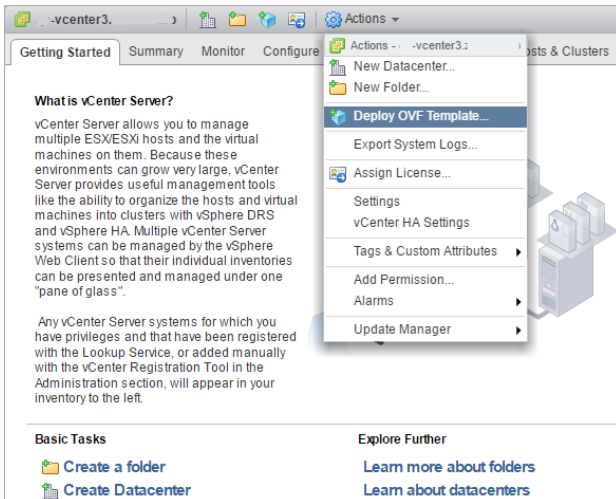
1. In a web browser, navigate to delivery.zenoss.com, and then log in.
2. Download the Resource Manager master host OVA file for the current release.
Replace *VERSION* with the current release number (for example, 6.4.1).

```
zenoss-zsd-VERSION-1-master.vmware.ova
```

3. Use the VMware vSphere Client to log in to vCenter as root, or as a user with superuser privileges, and then display the Home view.



4. Choose VMs and Templates.
5. In the top navigation bar, choose Actions > Deploy OVF Template.



6. Use panels in the Deploy OVF Template wizard to select the OVF package:
 - a. To choose the package from a drive on your workstation or network share, browse to the location and choose the OVA file. Click Next.
 - b. Select name and location: Specify a name for the OVF, select a datacenter or folder as the deployment location, and then click Next.
 - c. Select a resource: Select the host, cluster, or other resource on which to run the deployed template, and then click Next.
 - d. Review details: Verify the template details, and then click Next.
 - e. Select storage: In Select virtual disk format, choose Thin Provision, accept defaults for other fields, and then click Next.
 - f. Select networks: Accept defaults and click Next.
 - g. Ready to Complete: Verify the deployment settings, and then click Finish.
The Recent Tasks pane displays deployment progress and status information.
7. Navigate to the new virtual machine's Getting Started tab, and then click Edit virtual machine settings.
8. For a multi-host deployment, edit the Virtual Hardware settings of the virtual machine.
 - a. Change the settings.

- Reduce the number of CPUs from 8 to 4.
 - Reduce the amount of memory from 32 to 16.
- b. Click OK.
 9. On the new virtual machine's Getting Started tab, click Power on virtual machine.
 10. In the top navigation bar, choose Actions > Open console.

Proceed to [Configuring the Control Center master host](#).

Installing delegate hosts

For the complete list of installation instructions, see [Installing Resource Manager](#).

Editing the /etc/hosts file

Perform this procedure only if you use hostnames or fully-qualified domain names instead of IPv4 addresses, and only after all delegate hosts are installed and renamed.

Perform this procedure on the Control Center master host and on each delegate host.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:


```
[root@hostname ~]#
```

3. Open the /etc/hosts file in a text editor.

The following steps use the [nano](#) editor.

 - a. Start the editor.

```
nano /etc/hosts
```



- b. Optional: On delegate hosts, the file might include two entries with the same IP address. Remove the first of the two entries, which maps the IP address to the `zsd-master` hostname.
 - c. Add entries for the Control Center master host and for each delegate host.
 - d. To save, press Control-o.
 - e. To exit, press Control-x.
4. Return to the Appliance Administration menu.

```
exit
```

5. Exit the Appliance Administration menu.
 - a. Use the down-arrow key to select Exit.
 - b. Press Tab, and then press Enter.


Configuring a delegate host virtual machine

This procedure configures a new virtual machine as a delegate host.

1. Gain access to the console interface of the Control Center host through your hypervisor console interface.

```
YOU HAVE NOT ACTIVATED THIS APPLIANCE.  
PLEASE LOGIN TO ACTIVATE Zenoss Service Dynamics  
  
Welcome to Zenoss Service Dynamics  
  
After initial setup, the Control Center UI can be accessed by  
browsing to:  
  
https://zsd-delegate  
  
Ensure that zsd-delegate is resolvable to 10.88.121.216, either through your  
DNS system or through a HOSTS entry on the browser client. For more  
information refer to the installation notes.  
  
You can log in to this console to perform administrative tasks such  
as setting up networking and safely rebooting this system.  
  
Control Center UI login credentials are ccuser/zsd.  
To log into the console as root, use the default password 'zsd'.  
  
Linux Kernel 3.10.0-693.2.2.el7.x86_64 on an x86_64  
zsd-delegate login:
```

2. Log in as the root user.
The initial password is provided in the console.
3. The system prompts you to enter a new password for root.
4. The system prompts you to enter a new password for ccuser.
The ccuser account is the default account for gaining access to the Control Center browser interface.
5. In the IP field, enter the hostname, fully-qualified domain name, or IPv4 address of the master host.



Note: If you enter the hostname or fully-qualified domain name of the master host, you need an entry in the /etc/hosts file of the delegate host or a nameserver on your network that resolves the name to its IPv4 address.

6. Press Tab to select Ok, and then press Enter.
The system reboots.

Creating a delegate host with Hyper-V

To perform this task, you need:

- Microsoft Remote Desktop Connection
- Administrator privileges on a Microsoft Hyper-V server
- Permission to download files from delivery.zenoss.com.

Use this procedure to install the Resource Manager delegate host appliance as a virtual machine managed by Microsoft Hyper-V.

1. Use a Microsoft Remote Desktop Connection to log in to a Hyper-V host as Administrator, or as a user with Administrator privileges.
2. In a web browser, navigate to delivery.zenoss.com, and then log in.
3. Download the Resource Manager delegate host ISO file.
Replace *VERSION* with the current release number (for example, 6.4.1).

```
zenoss-zsd-VERSION-1-delegate.x86_64.iso
```

4. Open Hyper-V Manager.
5. In the left navigation pane, choose a server to host the virtual machine.
6. From the Action menu, choose New > Virtual Machine.
The New Virtual Machine Wizard opens.
7. In the New Virtual Machine Wizard dialog, display the Specify Name and Location panel.
If the first panel displayed is the Before You Begin panel, click Next.
8. In the Specify Name and Location panel, provide a name for the virtual machine, and then click Next.
9. In the Specify Generation panel, choose Generation 1, and then click Next.
10. In the Assign Memory panel, specify the memory for the virtual machine.
 - a. In the Startup memory field, enter 32768 (32GB).
 - b. Optional: Check Use Dynamic Memory for this virtual machine.
Resource Manager is tested with dynamic memory enabled.
 - c. Click Next.
11. In the Configure Networking panel, choose a virtual switch, and then click Next.
12. In the Connect Virtual Hard Disk panel, specify a new disk on which to install the guest operating system.
 - a. Choose Create a virtual hard disk.
 - b. Specify a name.
 - c. In the Size field, enter 30.
 - d. Click Next.
13. In the Installation Options panel, specify the delegate host ISO file.
 - a. Choose Install an operating system from a bootable CD/DVD-ROM.
 - b. Choose Image file (.iso).
 - c. Specify or browse to the location of the delegate host ISO file.
 - d. Click Next.
14. In the Completing the New Virtual Machine Wizard panel, verify the description, and then click Finish.
Hyper-V Manager creates the new virtual machine, and then closes the wizard.

Proceed to [Configuring and starting a Hyper-V delegate host](#).

Configuring and starting a Hyper-V delegate host

To perform this task, you need:


- A Microsoft Remote Desktop Connection
- Administrator privileges on a Microsoft Hyper-V server
- The delegate host created in the previous procedure ([Creating a virtual machine with Hyper-V](#))

Resource Manager delegate host virtual machines requires a total of 4 virtual hard disks. More information about required disks: [Resource requirements for multi-host deployments](#).

Use this procedure to configure hardware resources, create disks 2-4, and start the master host.

1. Use a Microsoft Remote Desktop Connection to log in to a Hyper-V host as Administrator, or as a user with Administrator privileges.
2. Open Hyper-V Manager.
3. In the Hyper-V Manager Virtual Machines area, right-click the new virtual machine, and then choose Settings.
The Settings dialog displays.
4. In the Hardware area, locate the virtual hard disk created previously, and then determine whether it is attached to an IDE controller. Hyper-V guest machines can only boot from an IDE drive.
5. In the Hardware area, choose Processor, and then change the number of processors assigned to the machine.
 - a. In Number of virtual processors, enter 8.
 - b. Click Apply.
6. In the Hardware area on the left, choose SCSI Controller, and then create additional virtual hard disks.
Repeat the following substeps to create new disks in the following sizes:
 - 50GB
 - 16GB
 - 16GB
 - a. In the controller area on the right, choose Hard Drive, and then click Add.
 - b. In the Location field, choose an unused location number.
 - c. In the Media area, choose Virtual hard disk, and then click New.
 - d. Complete panels in the New Virtual Hard Disk Wizard as follows:
 - i. Choose Disk Format: Choose VHDX, and then click Next.
 - ii. Choose Disk Type: Choose Dynamically expanding, and then click Next.
 - iii. Specify Name and Location: Enter the disk name, and then click Next.
 - iv. Configure Disk:
 1. Choose Create a new blank virtual hard disk.
 2. Size: Enter a disk size from the list at the beginning of this step.
 3. Click Next.
 - v. Summary/New Virtual Hard Disk Wizard: Verify the description, and then click Finish.
 - e. At the bottom of the Settings window, click Apply.

When all of the disks are created, click OK.
7. In the Hyper-V Manager Virtual Machines area, right-click the new virtual machine, and then choose Start.



8. In the Hyper-V Manager Virtual Machines area, right-click the new virtual machine, and then choose Connect.
9. In the Virtual Machine Connection window, press Enter.
The appliance installation process takes about 15 minutes, and should complete with no additional input. If received, disregard the Fast TSC calibration failure message.

Proceed to [Configuring a delegate host virtual machine](#).

Creating a delegate host with vSphere

To perform this task, you need:


- A VMware vSphere client
- Permission to download files from delivery.zenoss.com. Customers can request permission by filing a ticket at the [Zenoss Support](http://zenoss.com/support) site.

This procedure installs Resource Manager OVA packages as a virtual machine managed by vSphere Server version 6.5.0, using VMware vSphere Web Client 6.5. The procedure might differ with other versions of VMware vSphere Client.

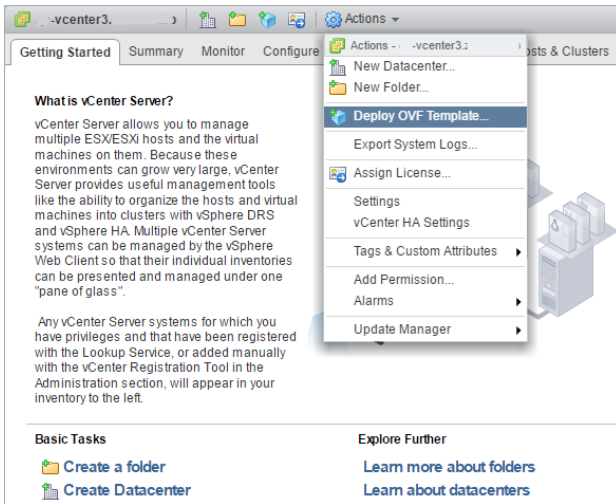
1. In a web browser, navigate to delivery.zenoss.com, and then log in.
2. Download the Resource Manager delegate host OVA file for the current release.
Replace *VERSION* with the current release number (for example, 6.4.1).

```
zenoss-zsd-VERSION-1-delegate.vmware.ova
```

3. Use the VMware vSphere Client to log in to vCenter as root, or as a user with superuser privileges, and then display the Home view.



4. In the top navigation bar, choose Actions > Deploy OVF Template.



5. Use panels in the Deploy OVF Template wizard to select the OVF package:
 - a. To choose the package from a drive on your workstation or network share, browse to the location and choose the OVA file. Click Next.
 - b. Select name and location: Specify a name for the OVF, select a datacenter or folder as the deployment location, and then click Next.
 - c. Select a resource: Select the host, cluster, or other resource on which to run the deployed template, and then click Next.
 - d. Review details: Verify the template details, and then click Next.
 - e. Select storage: In Select virtual disk format, choose Thin Provision, accept defaults for other fields, and then click Next.
 - f. Select networks: Accept defaults and click Next.
 - g. Ready to Complete: Verify the deployment settings, and then click Finish.The Recent Tasks pane displays deployment progress and status information.
6. Navigate to the new virtual machine's Getting Started tab, and then click Edit virtual machine settings.
7. On the new virtual machine's Getting Started tab, click Power on virtual machine.

Configuring a multi-host deployment

For the complete list of installation instructions, see [Installing Resource Manager](#).

Adding a delegate host through an SSH connection

To succeed, the following statements about the login account used to perform this procedure must be true:

- The account exists on both the master host and on the delegate host.
- The account has serviced CLI privileges.
- The account has either public key authentication or password authentication enabled on the master host and on the delegate host.

Use this procedure to add a delegate host to a resource pool through an SSH connection. Repeat this procedure on each delegate in your Resource Manager deployment.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Add a delegate host to a resource pool.

If the master and delegate host are configured for key-based access, the following command does not prompt you to add the delegate to the list of known hosts or to provide the password of the remote user account.

Use the hostname or IP address to identify a Control Center host. If you use a hostname, all Control Center hosts must be able to resolve it, either through an entry in `/etc/hosts` or through a nameserver on the network. In the following example, replace `Hostname-Or-IP` with the hostname or IP address of a delegate host, and replace `Resource-Pool` with the name of a resource pool.

If the host is behind a router or firewall for network address translation (NAT), include the option `--nat-address` to specify the NAT device's hostname or IP address and port of the delegate host.

```
serviced host add --register Hostname-Or-IP:4979 Resource-Pool --nat-address==NAT-Hostname-Or-IP:NAT-Port
```


Adding a delegate host using a file

Use this procedure to add a delegate host to a resource pool by using a key bundle file. Repeat this procedure on each delegate in your Resource Manager deployment.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Add a delegate host to a resource pool.

Use the hostname or IP address to identify a Control Center host. If you use a hostname, all Control Center hosts must be able to resolve it, either through an entry in `/etc/hosts` or through a nameserver on the network. In the following example, replace `Hostname-Or-IP` with the hostname or IP address of a delegate host, and replace `Resource-Pool` with the name of a resource pool.

If the host is behind a router or firewall for network address translation (NAT), include the option `--nat-address` to specify the NAT device's hostname or IP address and port of the delegate host.

```
serviced host add Hostname-Or-IP:4979 Resource-Pool --nat-address==NAT-Hostname-Or-IP:NAT-Port
```

The command creates a unique key bundle file in the local directory.

4. Use a file transfer utility such as `scp` to copy the key bundle file to the delegate host.

After you copy it to the delegate host, you can delete the key bundle file from the master host.
5. Log in to the Control Center delegate host as a user with `serviced` CLI privileges.
6. Install the key bundle.

Replace `Key-Bundle-Path` with the pathname of the key bundle file:

```
serviced host register Key-Bundle-Path
```

7. Delete the key bundle file.

The file is no longer needed on the delegate host.

Replace `Key-Bundle-Path` with the pathname of the key bundle file:

```
rm Key-Bundle-Path
```

Enabling NTP on Microsoft Hyper-V guests

Control Center requires a common time source. The procedures in this section configure [NTP](#) to synchronize the system clocks of Control Center hosts.

The procedures in this section are required only for multi-host deployments running as Microsoft Hyper-V guests. VMware vSphere guests use an hourly `cron` job to synchronize their system clocks with the host.

Configure NTP to rely on a time source as follows, depending on internet access of the Control Center hosts:

- If all hosts can access the internet, configure NTP to rely on public time servers.
- If no host can access the internet, configure NTP to rely on a private master server.

Perform the following procedures, in order:

- [Configuring NTP clients](#)
- [Configuring an NTP master server](#)
- [Configuring NTP for public time servers](#)

Configuring NTP clients

(*Hyper-V only*) This procedure configures delegates to synchronize their clocks with the NTP server on the Control Center master host. Perform this procedure only if the delegates do not have internet access. Repeat this procedure on each Control Center delegate host.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Create a backup of the NTP configuration file.

```
cp -p /etc/ntp.conf /etc/ntp.conf.orig
```

4. Edit the NTP configuration file.
 - a. Open `/etc/ntp.conf` with a text editor.
 - b. Replace all of the lines in the file with the following lines:

```
# Point to the master time server
server Master-Address

restrict default ignore
restrict 127.0.0.1
restrict Master-Address mask 255.255.255.255 nomodify notrap noquery

driftfile /var/lib/ntp/drift
```

- c. Replace both instances of `Master-Address` with the IPv4 address of the host where the NTP server is running (the Control Center master host).
 - d. Save the file and exit the editor.
5. Stop Control Center.

```
systemctl stop serviced
```

6. Synchronize the clock with the master server.

```
ntpdate -gg
```

7. Enable and start the NTP service.

```
systemctl enable ntpd && systemctl start ntpd
```

8. Start Control Center.

```
systemctl start serviced
```

Configuring an NTP master server

(*Hyper-V only*) This procedure configures an NTP master server on the Control Center master host. Perform this procedure only if the host does not have internet access.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Create a backup of the NTP configuration file.

```
cp -p /etc/ntp.conf /etc/ntp.conf.orig
```

4. Edit the NTP configuration file.
 - a. Open `/etc/ntp.conf` with a text editor.
 - b. Replace all of the lines in the file with the following lines:

```
# Use the local clock
server 127.127.1.0 prefer
fudge 127.127.1.0 stratum 10
driftfile /var/lib/ntp/drift
broadcastdelay 0.008

# Give localhost full access rights
restrict 127.0.0.1

# Grant access to client hosts
restrict Address-Range mask Netmask nomodify notrap
```

- c. Replace Address-Range with the range of IPv4 network addresses that are allowed to query this NTP server. For example, the following IP addresses are assigned to Control Center hosts:
 - 203.0.113.10
 - 203.0.113.11
 - 203.0.113.12
 - 203.0.113.13

For the preceding addresses, the value for Address-Range is 203.0.113.0.

- d. Replace Netmask with the IPv4 network mask that corresponds with the address range. For example, a valid network mask for 203.0.113.0 is 255.255.255.0.
- e. Save the file and exit the editor.

5. Stop Control Center.

```
systemctl stop serviced
```

6. Enable and start the NTP service.

```
systemctl enable ntpd && systemctl start ntpd
```

7. Start Control Center.

```
systemctl start serviced
```

Configuring NTP for public time servers

This procedure uses the default configuration of NTP to synchronize system clocks with public time servers. If all Control Center hosts can access the internet, repeat this procedure on each host, starting with the Control Center master host.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Stop Control Center.

```
systemctl stop serviced
```

4. Synchronize the system clock, and then enable and start the NTP service.
 - a. Set the system time.

```
ntpdate -gq
```

- b. Enable and start the ntpd service.

```
systemctl enable ntpd && systemctl start ntpd
```

5. Start Control Center.

```
systemctl start serviced
```

Guidelines for resource pool permissions

Zenoss recommends that multi-host Resource Manager deployments include separate Control Center resource pools for the following host or groups of hosts:

1. The Control Center master host
2. Control Center delegate hosts for the most or all Resource Manager services
3. Control Center delegate hosts for remote collectors (may be multiple pools)

Pools one and two require distributed file system (DFS) and administrative permissions. Pool three requires neither.

From the `serviced` CLI, the command that displays information about pools uses integers to summarize the permissions associated with a resource pool. The following table associates the values with the permissions they represent.

Value	Definition
0	No permissions
1	Administrative permission
2	DFS access permission
3	Both permissions

Follow these steps to set permissions on a resource pool:

1. Log in to the master host as a user with `serviced` CLI privileges.
2. Display the list of resource pools and their permissions.

```
serviced pool list -v | grep -E 'ID|Permissions'
```

Example result on a deployment with four resource pools:

```
"ID": "default",
"Permissions": 3,
"ID": "rm_pool",
"Permissions": 3,
"ID": "2",
"Permissions": 0,
"ID": "3",
"Permissions": 0,
```

3. Optional: Remove DFS access permission from a pool.

If you intend to remove both DFS access and administrative access permissions from a resource pool, you must remove DFS access permissions first!

Replace Pool-Name with the name of a resource pool:

```
serviced pool set-permission --dfs=false Pool-Name
```

4. Optional: Remove administrative permission from a pool.

If you intend to remove both DFS access and administrative access permissions from a resource pool, you must remove DFS access permissions first!

Replace Pool-Name with the name of a resource pool:

```
serviced pool set-permission --admin=false Pool-Name
```

Configuring Resource Manager

- [Enabling access to browser interfaces](#)
- [Configuration procedures](#)
- [Preparing for monitoring](#)
- [External HBase configuration](#)

Enabling access to browser interfaces

Control Center and Resource Manager have independent browser interfaces that are served by independent web servers. Both web servers are configured to use SSL/TLS communications.

The Control Center web server listens at the hostname of the Control Center master host and port 443. For a Control Center master host with the fully-qualified domain name (FQDN) `cc-master.example.com`, the hostname URL is `https://cc-master`. You can substitute an IP address for the hostname portion of the URL.

The Resource Manager web server can listen at *port public endpoints* and *virtual host public endpoints*.

- A *port public endpoint* is a combination of the IP address or hostname of the Control Center master host and a port number. The default configuration of Resource Manager does not include any port public endpoints. If the Control Center master host has more than one interface, you can configure port public endpoints with different hostnames. Also, you can disable TLS communications for a port public endpoint.

To use a port public endpoint to gain access to the Resource Manager browser interface, no additional network name resolution entries are required. The default entries for the network interfaces of the Control Center master host are sufficient.

- The default *virtual host public endpoint* is the text `zenoss5` prefixed to the hostname of the Control Center master host and port 443. For the FQDN `cc-master.example.com`, the URL of the default virtual host public endpoint is `https://zenoss5.cc-master:443`. You can change the name of the default virtual host and configure additional virtual host public endpoints.

To use a virtual host public endpoint to gain access to the Resource Manager browser interface, you must add name resolution entries for the virtual host to the DNS servers in your environment or to the hosts files of individual client systems.

The following sections provide additional information about public endpoints, and instructions for creating public endpoints and configuring virtual hostname resolution.

- [Creating and changing public endpoints](#)
- [Configuring name resolution for virtual hosts](#)

Creating and changing public endpoints

This section provides instructions for creating and changing port public endpoints and virtual host public endpoints.

The following table lists communication requirements and outlines the process for creating public endpoints. Step-by-step instructions follow this overview.


Port public endpoint	Virtual host public endpoint
Port public endpoints can communicate with or without SSL/TLS. <ol style="list-style-type: none">1. Create the endpoint.2. Configure the Zope service.	Virtual host public endpoints must use SSL/TLS communications. <ol style="list-style-type: none">1. Create the endpoint.2. Configure the Zope service.3. Configure virtual hostname resolution.

To change an existing public endpoint, create a new endpoint and then delete the existing endpoint.

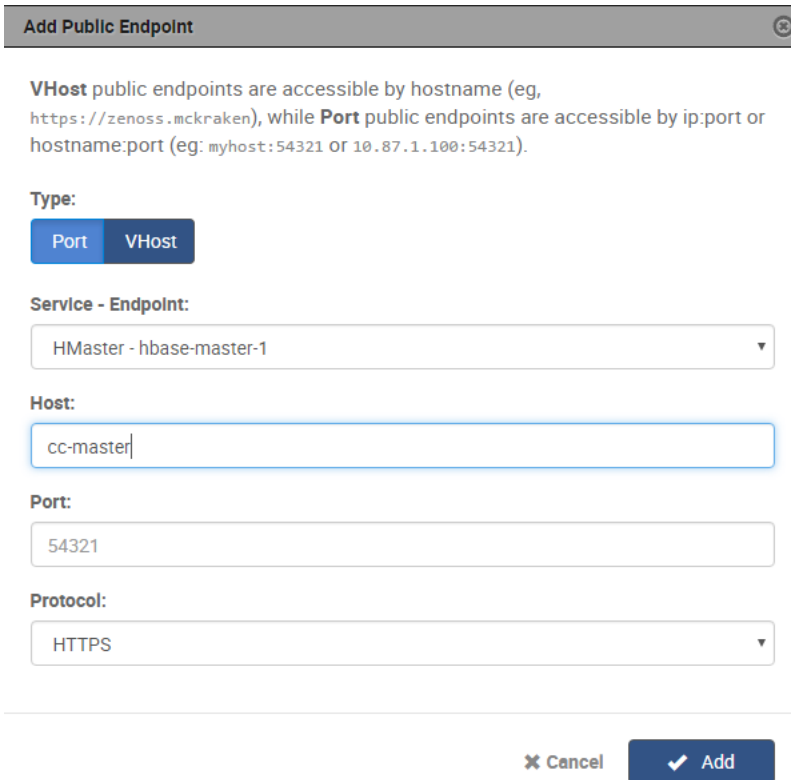
Creating a port public endpoint

Use this procedure to create a new port public endpoint. Port public endpoints can communicate with or without SSL/TLS.

1. Log in to the Control Center browser interface.
2. In the Application column of the Applications table, click the application name (Zenoss.resmgr).



3. On the right, above the Public Endpoints table, click Add Public Endpoints. The default view of the Add Public Endpoint dialog box displays the fields for creating a port public endpoint.



4. Define a new port public endpoint.
 - a. In the Type area, click Port.
 - b. From the Service - Endpoint list, select Zenoss.resmgr - zproxy.
 - c. In the Host field, enter a hostname or IP address that is assigned to a network interface on the Control Center master host. The default value is the hostname that was added with the Deployment Wizard when Resource Manager was initially deployed. If the Control Center master host has more than one network interface, you can add the hostname or IP address that is assigned to another interface.
 - d. In the Port field, enter a safe, unused port number that is greater than or equal to 1024 and less than or equal to 65535. For a list of ports that are considered unsafe, see [Unsafe ports on Chrome](#). For the list of ports that the Control Center master host uses, see [Security](#).
 - e. In the Protocol field, select HTTPS or HTTP. Optionally, you can set up a secure proxy server to handle HTTP requests that are sent to a port public endpoint.
 - f. Click Add.

Next step: Configure the Zope service to use the new port public endpoint. Choose one of the configuration options in the following table.

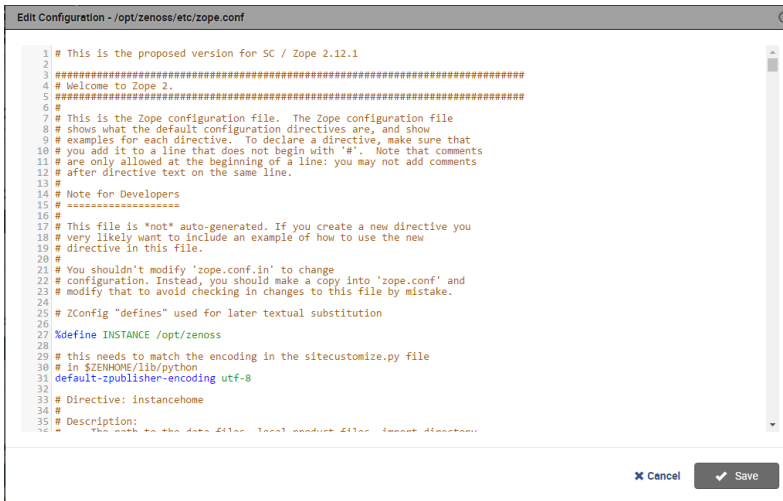
Zope configuration	Procedure
HTTPS and the default secure proxy server	Configuring Zope for HTTPS and the default secure proxy server
<p>HTTP and no proxy server</p> <p>Note that when you configure Zope for HTTP protocol and no proxy server, you can only gain access to the Resource Manager browser interface through port public endpoints that are configured for HTTP. Because virtual host public endpoints must use HTTPS protocol, any existing virtual host public endpoints stop working.</p>	Configuring Zope for HTTP and no proxy server
HTTP and a secure proxy server other than the default	Configuring Zope for HTTP and a secure proxy server

Configuring Zope for HTTPS and the default secure proxy server

Before performing this procedure, create a port public endpoint or a virtual host public endpoint to use the HTTPS protocol.

Use this procedure to configure the Zope service for SSL/TLS communications and the secure proxy server that is included in Resource Manager.

1. Log in to the Control Center browser interface.
2. In the Application column of the Applications table, click the application name (Zenoss.resmgr).
3. In the Services table, expand Zenoss > User Interface, and then click Zope.
The Zope service details page appears.
4. In the Configuration Files table, locate path /opt/zenoss/etc/zope.conf, and in the Actions column, click Edit.



5. Configure Zope for secure communications with the proxy server.
 - a. Locate the cgi-environment directive.
The directive is about one-third of the way down from the top of the file, on or near line 380.
 - b. Configure the proxy server for SSL/TLS communications:

```
<cgi-environment>
  HTTPS ON
</cgi-environment>
```

6. Configure the Beaker add-on product to use secure communications.
 - a. Locate the product-config directive.
The directive is at the bottom the file, on or near line 1122.
 - b. Set the value of the session.secure key to True.
7. Click Save.

Next steps:

- If you created a port public endpoint before performing this procedure, the endpoint is ready to use.
- If you created a virtual host public endpoint before performing this procedure, proceed to [Configuring name resolution for virtual hosts](#).

Configuring Zope for HTTP and no proxy server

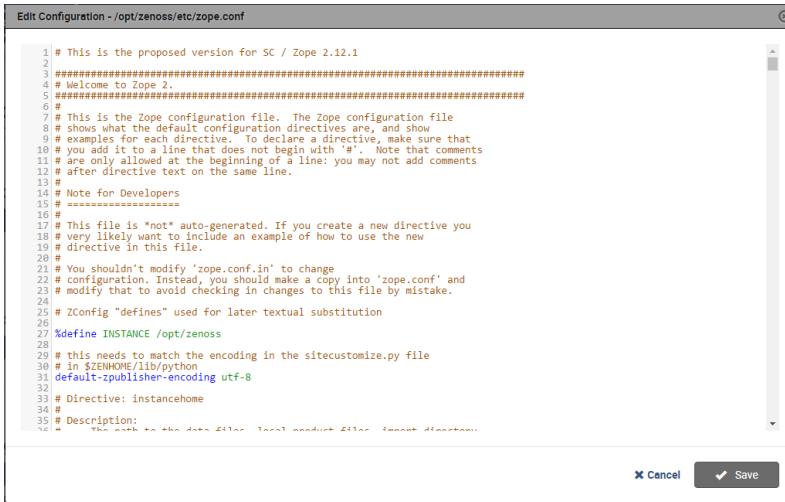
Before performing this procedure, create a port public endpoint to use the HTTP protocol. For more information, see [Creating a port public endpoint](#).

Use this procedure to configure the Zope service for insecure communications with Resource Manager browser interface clients.

When you configure Zope for insecure communications, existing virtual host public endpoints stop working.

Follow these steps:

1. Log in to the Control Center browser interface.
2. In the Application column of the Applications table, click the application name (Zenoss.resmgr).
3. In the Services table, expand Zenoss > User Interface, and then click Zope.
The Zope service details page appears.
4. In the Configuration Files table, locate path /opt/zenoss/etc/zope.conf, and in the Actions column, click Edit.



5. Configure Zope for insecure communications with the proxy server.
 - a. Locate the cgi-environment directive.
The directive is about one-third of the way down from the top of the file, on or near line 380.
 - b. Configure the proxy server for insecure communications:

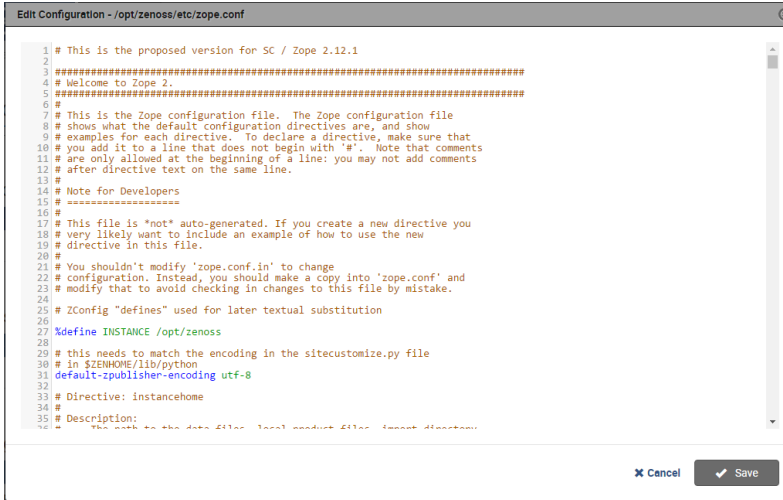
```
<cgi-environment>
  HTTPS OFF
</cgi-environment>
```

6. Configure the Beaker add-on product to use insecure communications.
 - a. Locate the product-config directive.
The directive is at the bottom the file, on or near line 1122.
 - b. Set the value of the session.secure key to False.
7. Click Save.

Configuring Zope for HTTP and a secure proxy server

Before performing this procedure, create a port public endpoint to use the HTTP protocol. For more information, see [Creating a port public endpoint](#). Use this procedure to configure the Zope service for SSL/TLS communications and a secure proxy server that is available on your network.

1. Log in to the Control Center browser interface.
2. In the Application column of the Applications table, click the application name (Zenoss.resmgr).
3. In the Services table, expand Zenoss > User Interface and then click Zope.
The Zope service details page appears.
4. In the Configuration Files table, locate path /opt/zenoss/etc/zope.conf, and in the Actions column, click Edit.



5. Configure Zope for secure communications with your proxy server.
 - a. Locate the `cgi-environment` directive.
The directive is about one-third of the way down from the top of the file, on or near line 380.
 - b. Configure the proxy server for SSL/TLS communications:


```
<cgi-environment>
  HTTPS ON
</cgi-environment>
```

6. Configure the Beaker add-on product to use secure communications.
 - a. Locate the `product-config` directive.
The directive is at the bottom the file, on or near line 1122.
 - b. Set the value of the `session.secure` key to `True`.
7. Click Save.

Creating a virtual host public endpoint

Use this procedure to create a new virtual host public endpoint. Virtual host public endpoints must use SSL/TLS communications.

1. Log in to the Control Center browser interface.
2. In the Application column of the Applications table, click the application name (Zenoss.resmgr).



3. On the right, above the Public Endpoints table, click Add Public Endpoints.
4. Define a new virtual host public endpoint.
 - a. In the Type area, click VHost.

Add Public Endpoint

VHost public endpoints are accessible by hostname (eg, `https://zenoss.mckraken`), while Port public endpoints are accessible by ip:port or hostname:port (eg: `myhost:54321` or `10.87.1.100:54321`).

Type:

Service - Endpoint:

VHost Hostname:

- b. From the Service - Endpoint list, select Zenoss.resmgr - zproxy.
- c. In the VHost Hostname field, enter a virtual hostname.

The hostname must be different from the Control Center hostname. For example, if the Control Center host is `https://zenoss.123`, then the virtual hostname cannot be `zenoss-123`.

The following strings of text are valid in this field:

 - A fully-qualified domain name (FQDN). Any string of text that includes one or more full stop characters (.) is treated as an FQDN.
 - A string of text that contains only letters and one or more hyphen characters (-). The string is prepended to the hostname of the Control Center master host, with a full stop character (.) separating the string and the hostname.
- d. Click Add.

Configuring name resolution for virtual hosts

To enable access to browser interfaces by virtual hosts, add name resolution entries to the DNS servers in your environment or to the hosts files of individual client systems.

The following line shows the syntax of the entry to add to a name resolution file:

```
IP-Address FQDN Hostname zenoss5.Hostname
```

For example, the following entry identifies a Control Center master host at IP address 192.0.2.12, hostname `cc-master`, in the `example.com` domain.

```
192.0.2.12 cc-master.example.com cc-master zenoss5.cc-master
```

Configuring name resolution on a Windows 7 system

To perform this procedure, you need Windows Administrator privileges.

1. Log in to the Windows 7 system as a user with Administrator privileges.
2. Click Start > All Programs > Accessories > Notepad.
3. Right click Notepad and then select Run as administrator.
4. Click File > Open, and then enter the following file `C:\Windows\System32\drivers\etc\hosts`.
5. At the end of the file, add a name resolution entry.
6. Save the file, and then exit Notepad.

Configuring name resolution on a Linux or OS/X system

To perform this procedure, you need superuser privileges on the client system.

1. Log in to the client system as root, or as a user with superuser privileges.
2. In a text editor, open the `/etc/hosts` file.
3. At the end of the file, add a name resolution entry.
4. Save the file, and then close the editor.

Configuration procedures

This section contains configuration procedures that you perform after Resource Manager is installed. Some of the procedures are optional, and indicated as such in the section title.

For installation and deployment instructions, see [Installing Resource Manager](#).

- [Starting Resource Manager](#)
- [Default server passwords](#)
- [Deleting the RabbitMQ guest user account](#)
- [MariaDB database utilities](#)
- [Optional: Assigning a virtual IP address to a resource pool](#)
- [Optional: Replacing the default digital certificate](#)
- [Optional: Customization management](#)
- [Optional: Configuring OpenTSDB compaction](#)
- [Optional: Creating a Redis cluster in a collector pool](#)
- [Optional: Enabling monitoring on IPv6 networks](#)

Starting Resource Manager

You can start Resource Manager from the Control Center browser interface or from the command-line interface.

Using the Control Center browser interface to start Resource Manager

To perform this procedure, you need:

- A tested client system and browser
- A user account on the Control Center master host with access privileges for the Control Center browser interface

For more information, see [User access control](#).

1. Log in to the Control Center browser interface.
2. In the Actions column of the Applications table, click Start for Zenoss.resmgr.
3. In the Start Service dialog box, click Start Service and x Children.
4. Optional: Monitor the startup.
 - a. In the Applications table, click Zenoss.resmgr.
 - b. Scroll down to the Services table and review the Health icon for each service. As services start, the Health icon changes to a check mark.

Using the command line to start Resource Manager

To perform this procedure, you need serviced CLI privileges. For more information, see [User access control](#).

1. Log in to the Control Center master host as a user with serviced CLI privileges.
2. Start Resource Manager.

```
serviced service start Zenoss.resmgr
```

3. Optional: Monitor the startup.

```
serviced service status Zenoss.resmgr
```

Default server passwords

Resource Manager adds global configuration parameters that include passwords to the run-time environments (Docker containers) of every service. The default passwords for some servers are the same in all Resource Manager distributions. To avoid security issues, Zenoss recommends changing the default passwords of the servers.

Changes to global configuration parameters persist across upgrades.

The following list associates the affected servers, their Resource Manager services, and their account information.

The list includes both account names and passwords. Zenoss recommends changing the passwords of each account and strongly discourages changing the account names. Changes to either default usernames or passwords may require an update to configuration properties related to the RMMonitor ZenPack.

MariaDB server for the events database

Service: mariadb-events
Administrator account: global.conf.zep-admin-user
Administrator password: global.conf.zep-admin-password
User account: global.conf.zep-user
User password :global.conf.zep-password

MariaDB server for the models database

Service: mariadb-model
Administrator account: global.conf.zodb-admin-user
Administrator password: global.conf.zodb-admin-password
User account: global.conf.zodb-user
User password: global.conf.zodb-password

RabbitMQ server

Service: RabbitMQ
User account: global.conf.amqpuser
User password: global.conf.amqppassword

Zope authentication server

Service: Zauth
User account: global.conf.zauth-username
User password: global.conf.zauth-password

Changing MariaDB passwords

Use this procedure to change the passwords of the MariaDB databases for event and model data.

To perform this procedure, the mariadb-events and mariadb-model child services of Resource Manager must be running.

1. Log in to the Control Center master host as root, or as a user with superuser privileges.
2. Change the passwords of the events database server.
 - a. Log in to the Docker container of the mariadb-events service as zenoss.

```
serviced service attach mariadb-events su - zenoss
```

- b. Start an interactive session.

```
export TERM=dumb; mysql -u root
```

- c. Access the administration database.

```
USE mysql
```

- d. Set the password of the root user.

Replace New-Password with a new password:

```
SET PASSWORD FOR 'root'@'127.0.0.1' = PASSWORD('New-Password');  
SET PASSWORD FOR 'root'@'localhost' = PASSWORD('New-Password');
```

Record the password for use in a subsequent step.

- e. Update the password of the zenoss user.

Replace New-Password with a new password:

```
SET PASSWORD FOR 'zenoss'@'127.0.0.1' = PASSWORD('New-Password');  
SET PASSWORD FOR 'zenoss'@'%' = PASSWORD('New-Password');
```

Record the password for use in a subsequent step.

- f. Exit the interactive session.

```
QUIT
```

The MariaDB server loads the grant tables into memory immediately when account management statements like SET PASSWORD are used, so the FLUSH PRIVILEGES statement is not necessary.

- g. Log out of the Docker container.

```
exit
```

3. Change the passwords of the model database server.

- a. Log in to the Docker container of the mariadb-model service as zenoss.

```
serviced service attach mariadb-model su - zenoss
```

- b. Start an interactive session.

```
export TERM=dumb; mysql -u root
```

- c. Access the administration database.

```
USE mysql
```

- d. Set the password of the root user.

Replace New-Password with a new password:

```
SET PASSWORD FOR 'root'@'127.0.0.1' = PASSWORD('New-Password');  
SET PASSWORD FOR 'root'@'localhost' = PASSWORD('New-Password');
```

Record the password for use in a subsequent step.

- e. Update the password of the zenoss user.

Replace New-Password with a new password:

```
SET PASSWORD FOR 'zenoss'@'127.0.0.1' = PASSWORD('New-Password');
SET PASSWORD FOR 'zenoss'@'%' = PASSWORD('New-Password');
```

Record the password for use in a subsequent step.

- f. Exit the interactive session.


```
QUIT
```

- g. Log out of the Docker container.

```
exit
```

4. Log in to the Control Center browser interface.
5. In the Applications table, click Zenoss.resmgr.
6. In the application title line, click Edit Variables.

Initially, the application title line appears immediately below the Control Center banner at the top of the page. When you scroll down the page, the application title line persists at the top of the page.



7. Update the passwords of the event database server.
 - a. In the Edit Variables dialog, locate the `global.conf.zep-password` variable.
 - b. Replace its value with the password specified previously for the zenoss user of the events database server.
 - c. Locate the `global.conf.zep-admin-password` variable.
 - d. Replace its value with the password specified previously for the root user of the events database server.
8. Update the passwords of the model database server.
 - a. Locate the `global.conf.zodb-password` variable.
 - b. Replace its value with the password specified previously for the zenoss user of the model database server.
 - c. Locate the `global.conf.zodb-admin-password` variable.
 - d. Replace its value with the password specified previously for the root user of the model database server.
 - e. At the bottom of the Edit Variables dialog, click Save Changes.
9. In the application title line, click Restart.

Changing the RabbitMQ server password

Use this procedure to change the password of the RabbitMQ server.

Changing this password will require an update to the `zRMMonRabbitUser` and `zRMMonRabbitPassword` configuration properties for this Resource Manager instance if it is being monitored by the RMMonitor ZenPack.

To perform this procedure, the mariadb-model child services of Resource Manager must be running.

1. Log in to the Control Center master host as root, or as a user with superuser privileges.
2. Change the password of the zenoss user.
 - a. Log in to the Docker container of the RabbitMQ service as root.

```
serviced service attach rabbitmq
```

- b. Change the password.
Replace `New-Password` with a new password:


```
rabbitmqctl change_password zenoss New-Password
```

Record the password for use in a subsequent step.

- c. Log out of the Docker container.

```
exit
```

3. Log in to the Control Center browser interface.
4. In the Applications table, click Zenoss.
5. In the application title line, click Edit Variables.
Initially, the application title line appears immediately below the Control Center banner at the top of the page. When you scroll down the page, the application title line persists at the top of the page.



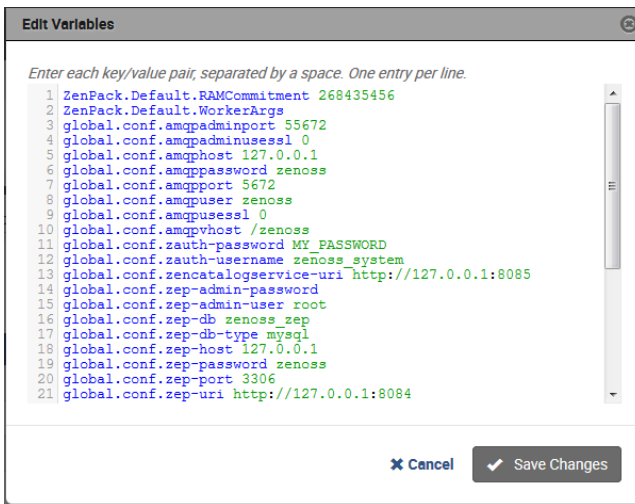
6. Change the password of the RabbitMQ server.
 - a. In the Edit Variables dialog, locate the `global.conf.amqppassword` variable.
 - b. Replace its value with the new password specified previously.
 - c. At the bottom of the Edit Variables dialog, click Save Changes.
7. Restart the RabbitMQ service.
 - a. Scroll down to the Services table, and then locate the RabbitMQ service.
 - b. In the Actions column of the service, click the Restart control.

Changing the Zope authentication server password

Use this procedure to change the password of the Zope authentication server.

To perform this procedure, the Resource Manager application must be running. During the procedure, Resource Manager must be restarted.

1. Log in to the Control Center browser interface.
2. In the applications table, click Zenoss.resmgr.
3. Stop all metricshipper, metricconsumer, and centralquery services.
4. Log in to the Resource Manager browser interface as zenoss_system.
The default password is MY_PASSWORD.
5. Click the ADVANCED tab, and then click Settings.
6. From the left column, select Users.
7. In the UserId table, click the zenoss_system link.
8. In the USER SETTINGS area, enter a new password in the Set New Password field, and then enter it again, in the Confirm New Password field.
Record the password for use in a subsequent step.
9. In the Current Password for zenoss_system field, enter the password you used to log in as zenoss_system.
10. Click Save Settings, and then log out of the browser interface.
11. In Control Center, go to Applications tab and click Zenoss.resmgr.
12. In the application title line, click Edit Variables.



13. Update the password of the zenoss_system user account.
 - a. In the Edit Variables dialog, locate the global.conf.zauth-password variable.
This variable sets the password of the Zope authentication server.
 - b. Replace its value with the password specified previously for the zenoss_system user account.
 - c. At the bottom of the Edit Variables dialog, click Save Changes.
14. In the application title line, click the Restart control.

Deleting the RabbitMQ guest user account

By default, RabbitMQ distributions include the guest user account. To prevent security issues, Zenoss recommends deleting the account.

1. Log in to the Control Center master host as a user with serviced CLI privileges.
2. Attach to the RabbitMQ container.

```
serviced service attach rabbitmq
```

3. Delete the guest user account.

```
rabbitmqctl delete_user guest
```

4. Exit the container session.

```
exit
```

5. Restart the RabbitMQ service.

```
serviced service restart rabbitmq
```


MariaDB database utilities

The [Percona Toolkit](#) is a collection of helpful utilities for MySQL and MariaDB databases. For licensing reasons, Zenoss can not distribute it. Zenoss strongly recommends that all deployments of Resource Manager install the Percona Toolkit.

- [Installing the Percona Toolkit with internet access](#)
- [Installing the Percona Toolkit without internet access](#)

Installing the Percona Toolkit with internet access

To perform this procedure, you need one of the following:

- a login account on the master host that is a member of the docker group
- the password of the root user account

For more information, see [User access control](#).

1. Log in to the Control Center master host.
2. Install the package.

```
serviced service run zope install-percona
```

At the end of the installation process, the message `Container not committed` is displayed. This is normal. The tools are installed in the distributed file system, not in an image.

Installing the Percona Toolkit without internet access

To perform this procedure, you need one of the following:

- a login account on the master host that is a member of the docker group
- the password of the root user account

In addition, you need the Percona Toolkit package file. This procedure includes steps for downloading it to a client system, and then copying it to the Control Center master host.

1. On a client system, use a web browser to download the latest version of the [Percona Toolkit](#) package.
2. Log in to the Control Center master host.
3. Prepare the package for installation.
 - a. On the Control Center master host, create a directory for the package, and then change directory.

```
mkdir /tmp/percona && cd /tmp/percona
```

- b. Copy the package to the temporary location.
You may use a file transfer utility such as [WinSCP](#).
- c. Update the access permissions of the file and directory.

```
chmod -R 777 /tmp/percona
```

4. Start a shell as the zenoss user in a Zope container.
 - a. Change directory to the location of the Percona Toolkit file.

```
cd /tmp/percona
```

- b. Start an interactive shell in a Zope container and save a snapshot named PerconaToolkit.

```
mySnap=InstallPerconaToolkit  
serviced service shell -i -s $mySnap zope bash
```

- c. Switch user to zenoss.

```
su - zenoss
```

5. Install the package and exit the Zope container.
 - a. Create a directory for the package.

```
PERCONADIR=/var/zenoss/percona  
mkdir -p $PERCONADIR
```

- b. Extract the package files.
Replace Version with the version number of the package file:

```
tar --strip-components=1 -C $PERCONADIR -xzvf /mnt/pwd/percona-toolkit-Version.tar.gz
```

- c. Exit the zenoss shell.

```
exit
```

- d. Exit the Zope container.

```
exit
```

6. Commit the named snapshot.

```
serviced snapshot commit $mySnap
```

7. Restart the zeneventserver service.

```
serviced service restart zeneventserver
```

Optional: Assigning a virtual IP address to a resource pool

The zentrap and zensyslog services are designed to receive data from devices in your environment at a specific IP address. Typically, the address is assigned to a specific host. However, if the host fails, then no data is received. To avoid this issue, you can assign a virtual IP address to a resource pool, and then Control Center can create a virtual IP interface on any host in the pool. Zenoss recommends using a virtual IP with resource pools that include Resource Manager collection services as a best practice.

To perform this procedure, you need an unused IPv4 address in the same subnet as the other hosts in the resource pool to modify. To avoid conflicts, ask your networking specialist to assign or reserve the address. In addition, all of the hosts in the resource pool to modify must have the same network interface names. Note: For additional information about network interface names, contact Zenoss Support.

1. Log in to the Control Center browser interface.
2. At the top of the page, click Resource Pools.
3. In the Resource Pool column of the Resource Pools table, click the name of the resource pool to modify.
4. At the right side of the Virtual IPs table, click Add Virtual IP.
5. In the Add Virtual IP dialog, specify the virtual IP.
 - a. In the IP field, enter an IPv4 address.
The address must be in the same subnet as the other hosts in the current resource pool.
 - b. In the Netmask field, enter an IPv4 subnet mask.
The mask must match the range of addresses in the current resource pool. The following table associates commonly-used subnet masks with the number of addresses they include.

Subnet mask	Addresses in subnet
255.255.255.192	64
255.255.255.224	32
255.255.255.240	16
255.255.255.248	8

- c. In the Interface field, enter the name of the network interface that is used on all hosts in the resource pool.
 - d. At the bottom of the Add Virtual IP dialog, click Add Virtual IP.

When you configure devices to send syslog or SNMP trap messages, use the virtual IP address assigned to a resource pool.

Optional: Replacing the default digital certificate

The default configuration of the Resource Manager web server uses a Zenoss self-signed certificate for SSL/TLS communications. Use this procedure to install your own digital certificate. Note: If your environment uses a reverse proxy, contact Zenoss Support for customized assistance.

To perform this procedure, you need:

- the certificate and key files of a digital certificate from a certificate authority or from a digital certificate created with a utility such as OpenSSL Note: Certificates that require a passphrase are not supported.
- superuser privileges on the Control Center master host

1. Log in to the Control Center master host.
2. Copy the certificate and key files of your digital certificate to /etc on the master host.
You can store the files in any location that remains unchanged during operating system upgrades.
3. Configure Control Center to use your digital certificate.
 - a. Open /etc/default/serviced with a text editor.
 - b. Locate the SERVICED_CERT_FILE declaration, and then replace its value with the absolute path of your certificate file.
 - c. Remove the number sign character (#) from the beginning of the line.
 - d. Locate the SERVICED_KEY_FILE declaration, and then replace its value with the absolute path of your key file.
 - e. Remove the number sign character (#) from the beginning of the line.
 - f. Save the file, and then close editor.
4. Reload the Control Center service.

```
systemctl reload serviced
```

Optional: Customization management

Resource Manager software is distributed as Docker images. Upgrades often replace images, so customizations of Resource Manager services are lost, unless customizations are installed with a change management system.

Quilt is a utility for managing software changes, and Zenoss recommends installing it to manage customizations.

- [Installing Quilt with internet access](#)
- [Installing Quilt without internet access](#)

Installing Quilt with internet access

To perform this procedure, you need superuser privileges on the Control Center master host.

Use this procedure to add the Quilt patch management system to Resource Manager.

1. Log in to the Control Center master host.
2. Install the Quilt package.

```
serviced service run zope install-quilt
```

Installing Quilt without internet access

To perform this procedure, you need superuser privileges on the Control Center master host and the Quilt package file. This procedure includes steps for downloading the package to a client system, and then copying it to the Control Center master host.

Use this procedure to add the Quilt patch management system to Resource Manager.

1. On a client system, use a web browser to download the latest version of [the Quilt package](#).
2. Log in to the Control Center master host.
3. Prepare the package for installation.
 - a. On the Control Center master host, create a directory for the package, and then change directory.

```
mkdir /tmp/quilt && cd /tmp/quilt
```

- b. Copy the package to the temporary location.
You may use a file transfer utility such as [WinSCP](#).
- c. Update the access permissions of the file and directory.

```
chmod -R 777 /tmp/quilt
```

4. Start a shell as the zenoss user in a Zope container.
 - a. Change directory to the location of the Quilt package file.

```
cd /tmp/quilt
```

- b. Start an interactive shell in a Zope container and save a snapshot named InstallQuilt.

```
mySnap=InstallQuilt  
serviced service shell -i -s $mySnap zope bash
```

- c. Switch user to zenoss.

```
su - zenoss
```

5. Extract the package files, and then compile and install Quilt.

- a. Extract the package files.

```
tar xzvf /mnt/pwd/quilt-*.tar.gz -C /tmp
```

- b. Compile and install the package.

```
cd /tmp/quilt-* && ./configure --prefix=/opt/zenoss/var/ext && make && make install
```

6. Exit the container.

- a. Exit the zenoss shell.

```
exit
```

- b. Exit the Zope container.

```
exit
```

7. Commit the named snapshot.


```
serviced snapshot commit $mySnap
```


Optional: Configuring OpenTSDB compaction

Resource Manager uses OpenTSDB to store the monitoring data it collects. When OpenTSDB compaction is enabled, multiple columns in an HBase row are merged into a single column, to reduce disk space. In testing, Zenoss has observed that these merges result in duplicate data points, so by default, compaction is disabled. Duplicate data points do not affect data integrity.

Enabling compaction slows performance and is not recommended. Zenoss customers may contact Zenoss Support for additional assistance.

1. Log in to the Control Center browser interface.
2. In the Applications table, click Zenoss.resmgr.
3. In the application title line, click Edit Variables.
Initially, the application title line appears immediately below the Control Center banner at the top of the page. When you scroll down the page, the application title line persists at the top of the page.



4. In the Edit Variables dialog, scroll to the bottom of the list.
5. Change the value of the `tsd.storage.enable_compaction` variable from `False` to `True`.
6. Click **Save Changes**.
7. Restart the OpenTSDB services.
 - a. Scroll down the page to the Services table, and then locate the `opentsdb` service.
 - b. In the Actions column of the `opentsdb` service, click the **Restart** control.

Optional: Creating a Redis cluster in a collector pool

Use this procedure to cluster the [Redis](#) caching service (collectorredis) in collector pools. Clustering increases the efficiency of multi-host collectors by ensuring that cached configurations are used effectively between collectorredis restarts.

1. Log in to the Control Center browser interface.
2. In the Applications table, click Zenoss.resmgr.
3. Scroll down to the Services table, locate and open the collectorredis service of a collector pool.
The DFS permission of the collector pool to modify must be disabled.
4. On the collectorredis page, click Edit Service.
5. In the Instances field, enter the number of hosts in the collector pool, and then click Save Changes.

Optional: Enabling monitoring on IPv6 networks

This procedure describes how to enable monitoring of devices that are located on an IPv6 network. The network must be reachable from the IPv4 network environment in which Control Center is deployed. Use this procedure to route an IPv6 address block to Control Center using Docker's virtual bridge interface, `docker0`. Resource Manager can monitor IPv6 devices that have addresses in the routed block.

To perform this procedure, each Control Center host needs a unique IPv6 prefix routed to it by an upstream router, and the Docker service on Control Center host needs to be configured to forward IPv6 packets.

The subnet for Docker containers should at least have a size of /80, so that an IPv6 address can end with the container's MAC address and you prevent NDP neighbor cache invalidation issues in the Docker layer.

For example, a multi-host deployment with one master host and three delegates could have the IPv6 configuration in the following table.

Control Center host	IPv6 link prefix	IPv6 routed prefix
Master	2001:DB8:ABCD:1000::500/64	2001:DB8:ABCD:2000::/64
Delegate 1	2001:DB8:ABCD:1000::501/64	2001:DB8:ABCD:2001::/64
Delegate 2	2001:DB8:ABCD:1000::502/64	2001:DB8:ABCD:2002::/64
Delegate 3	2001:DB8:ABCD:1000::503/64	2001:DB8:ABCD:2003::/64

The following example shows how to configure the static routes in the preceding table on an upstream Cisco router:

```
ipv6 route 2001:DB8:ABCD:2000::/64 2001:DB8:ABCD:1000::500
ipv6 route 2001:DB8:ABCD:2001::/64 2001:DB8:ABCD:1000::501
ipv6 route 2001:DB8:ABCD:2002::/64 2001:DB8:ABCD:1000::502
ipv6 route 2001:DB8:ABCD:2003::/64 2001:DB8:ABCD:1000::503
```

Perform the following steps on each Control Center host:

1. Log on to the Control Center host as root, or as a user with superuser privileges.
2. Configure IPv6 packet forwarding.
 - a. Open `/etc/sysctl.d/ipv6.conf` with a text editor.
 - b. Add or edit the following line:

```
net.ipv6.conf.all.forwarding=1
```

- c. Save the file, and then close the text editor.
3. Enable IPv6 packet forwarding without rebooting the host.

```
sysctl -w net.ipv6.conf.all.forwarding=1
```

4. Configure Docker for IPv6 communications.
 - a. Open `/etc/sysconfig/docker` with a text editor.
 - b. Add the following flags to the end of the `OPTIONS` declaration.
Replace `Subnet-Block` with the IPv6 subnet to route to Control Center, in CIDR notation:

```
--ipv6 --fixed-cidr-v6="Subnet-Block"
```

- c. Change the delimiter of the `OPTIONS` declaration to the apostrophe character (`'`).
The default delimiter of the `OPTIONS` declaration is the quotation mark character (`"`), which is the same delimiter used with the `--fixed-cidr-ipv6` flag.
 - d. Save the file, and then close the text editor.
5. Restart the Docker service.

```
systemctl restart docker
```

After all Control Center hosts are configured, test IPv6 by using the Docker container of the `zenping` service to ping a known address:

```
serviced service attach zenping ping6 -c 1 ipv6.google.com
```

If the ping is successful, Docker is able to resolve IPv6 addresses and you can monitor devices on the IPv6 network.

If you are unable to reach an IPv6 address, or you need help with this procedure, contact Zenoss Support.

Preparing for monitoring

Resource Manager uses standard management APIs to collect performance data, and therefore does not install proprietary agents on your infrastructure devices to collect monitoring data. However, Zenoss recommends that you review the information in this section to verify that the devices to you want to monitor are ready to respond to requests for data.

This section describes how to prepare the most common IT infrastructure. If the infrastructure you want to monitor is not described here, please refer to the corresponding ZenPack documentation in the [ZenPack catalog](#).

When your infrastructure is ready to monitor, the Resource Manager Setup Wizard guides you through the process of discovering devices on your network and adding devices by category and type.

Preparing network devices

Preparing switches and routers

To prepare a switch or router device for monitoring, verify that an SNMP agent is installed and currently running on the device.

The rest of this section describes how to prepare Cisco network devices for monitoring. For other device types, refer to the [ZenPack catalog](#) documentation.

Preparing Cisco UCS network devices

Resource Manager uses SNMP to provide customized or generalized support for many Zenoss products.

The following table associates Zenoss products with the customized Resource Manager device types that support them. Device types are listed in the Network area of the Add Infrastructure wizard, which is both part of the setup wizard and available through the Resource Manager browser interface.

The following device considerations apply:



- Some supported devices, such as the Cisco Nexus 7000 and 9000 switches, represent a large number of discrete monitoring endpoints. If you are unsure which Resource Manager deployment size supports the number of high-density devices you wish to monitor, contact your Zenoss representative.
- To monitor Cisco Nexus 9000 Series devices, you must first enable NX-API with the feature manager CLI command on the device. For detailed instructions on performing this task, refer to the [Cisco documentation](#) for the Nexus 9000.

Cisco product	Device type
Cisco Catalyst 6500 and 3560 Series Switches	Cisco 6500 (SNMP)
Cisco Nexus 5000 Series Switches	Cisco Nexus 5000 (SNMP + Netconf)
Cisco Nexus 7000 Series Switches	Cisco Nexus 7000 (SNMP + Netconf)
Cisco Nexus 1000v Series Switches	Cisco Nexus 1000V (SNMP + Netconf)
Cisco Nexus 3000 Series Switches	Cisco Nexus 3000 (SNMP + Netconf)
Cisco Nexus 9000 Series Switches	Cisco Nexus 9000 (NX-API)
Cisco Catalyst 6500 Series Virtual Switching Systems	Cisco VSS (SNMP)
Cisco MDS 9000 Series Multilayer Switches	Cisco MDS 9000 (SNMP)

In addition, Resource Manager provides two generalized device types.

Cisco product	Device type
Cisco CatOS-based switches or routers	Generic Switch/Router (SNMP)
Cisco IOS-based switches or routers	Cisco IOS (SNMP)

Preparing storage devices

This section describes how to prepare NetApp and EMC storage devices for monitoring.

For other device types, refer to the [ZenPack catalog](#) documentation.

Legacy NetApp filers

Resource Manager uses SNMP to monitor legacy NetApp filers that do not support the Data ONTAP® API (ZAPI).

The data gathered are approximate because the values for many objects (Aggregate, Volume, Plex, and RAID group) are not exposed by the NetApp MIB.

To prepare a legacy NetApp filer for monitoring, verify that SNMPv2 is installed, and then start an SNMP agent.

Recent NetApp filers

Resource Manager uses HTTP to monitor NetApp filers that support the Data ONTAP® API (ZAPI).

To prepare a recent NetApp filer for monitoring, verify the following conditions:

- The filer is running in 7-Mode or C-Mode.
- A supported version of ZAPI is installed and enabled. The minimum required version is 8.x.
- The user name and password of your account on the filer is authorized to use ZAPI.

EMC storage arrays

Resource Manager uses the Web-Based Enterprise Management (WBEM) protocol to send queries to EMC Storage Management Initiative Specification (SMI-S) providers that are associated with EMC VMAX and VNX storage arrays.

To prepare EMC arrays for monitoring:

- At least one EMC SMI-S provider must be running for each type of array to monitor. (The VMAX and VNX data models are different.)
- Before adding an SMI-S provider to Resource Manager, Zenoss recommends that you confirm that it is responding to requests.
- You need the following information:
 - user name and password for an account that is authorized to collect data on each SMI-S provider
 - IP address of each SMI-S provider
 - port number at which each SMI-S provider listens for requests
 - whether to use SSL

When statistics logging is disabled on the EMC device, graphs for component types of EMC arrays display NaN. The logging feature has a low default timeout value and must be set to a higher value or turned on again periodically.

Verifying an SMI-S provider on EMC devices

To perform this procedure, you need a Linux host that has a network path to the SMI-S providers of the arrays to monitor.

Do **not** perform this procedure on the Resource Manager host.

Perform this procedure to verify that the SMI-S providers associated with EMC arrays are configured correctly, and are responding to WBEM queries from command line tools.

1. Log in to a Linux host as root, or as a user with superuser privileges.
2. Install a WBEM command-line interface package, such as `wbemcli`.
3. Verify the SMI-S provider. Replace the variables with values that are valid in your environment.

```
wbemcli IP-Address:Port -u admin -p 'Password' -n root/emc --no-sslei('EMC_DiskDrive')
```

The expected result is a list of Disk Drive classes.

Preparing server devices

This section describes how to prepare Linux and Windows servers for monitoring. For other device types, refer to the [ZenPack catalog](#) documentation.

Preparing Linux servers for monitoring

Resource Manager can monitor Linux servers with SNMP or SSH. For SNMP monitoring, install an SNMP package on the server (for example, [Net-SNMP](#)) and start the agent.

For SSH monitoring, install an SSH server package (for example, [OpenSSH](#)) and start the SSH daemon. In addition, you must enable an account to run the `pvs`, `vgs`, `lvs`, `systemctl`, `initctl`, and `service` commands through SSH, without a TTY.

Enabling the root account to run monitoring commands

By default, `root` may only run monitoring commands locally. To enable remote privileges, follow these steps on each server to monitor:

1. Log in to the server as `root` or as a user with superuser privileges.
2. If necessary, install the `sudo` package on the server.
For more information, refer to your operating system documentation.
3. Allow `root` to execute commands through SSH, without a TTY.
 - a. Open the `/etc/sudoers` file with a text editor.
 - b. Find the line containing `root ALL=(ALL) ALL`.
 - c. Add the following line beneath it:

```
Defaults:root !requiretty
```

- d. Save the file, and then exit the editor.

Enabling a non-privileged account to run monitoring commands

To enable an account other than `root` to run monitoring commands remotely, follow these steps on each server to monitor:

1. Log in to the server as `root` or as a user with superuser privileges.
2. Create a user named `zenmonitor`.
3. If necessary, install the `sudo` package on the server.
For more information, refer to your operating system documentation.
4. Configure the `zenmonitor` user to run the commands through SSH, without a TTY.
 - a. Open `/etc/sudoers.d/zenoss` with a text editor.
If `sudoers.d` is not supported, open `/etc/sudoers`.
 - b. Add the following lines to the bottom of the file:

```
Defaults:zenmonitor !requiretty
Cmdnd_Alias ZENOSS_LVM_CMDS = /sbin/pvs, /sbin/vgs, /sbin/lvs, \
    /usr/sbin/pvs, /usr/sbin/vgs, /usr/sbin/lvs
Cmdnd_Alias ZENOSS_SVC_CMDS = /bin/systemctl list-units *, \
    /bin/systemctl status *, /sbin/initctl list, /sbin/service --status-all, \
    /usr/sbin/dmidecode
zenmonitor ALL=(ALL) NOPASSWD: ZENOSS_LVM_CMDS, ZENOSS_SVC_CMDS
```

- c. Verify that all paths in the preceding lines are correct.
- d. Save the file, and then exit the editor.

Preparing Windows servers for monitoring

Resource Manager uses SNMP or WinRM to monitor Microsoft Windows systems as follows:

- Microsoft Windows Server 2106 - WinRM only.
SNMP support does not exist for Windows Server 2106.
- Microsoft Windows Server 2012 and 2012 R2 - WinRM only.
SNMP support does not exist for Windows Server 2012.
- Microsoft Windows Server 2008 R2 - SNMP v1/v2 or WinRM.
SNMP v3 support does not exist for Windows Server 2008 R2.

To prepare a Windows 2008 system for SNMP monitoring, start the SNMP service.

To prepare a Windows system for WinRM monitoring, refer to the [support article](#) that describes the options and provides the procedures for configuring your systems.

Preparing hypervisor devices

This section describes how to prepare vSphere and Hyper-V hypervisors for monitoring. For other device types, refer to the [ZenPack catalog](#) documentation.

vSphere hosts

Resource Manager uses SOAP to monitor VMware vSphere servers running versions 4.1, 5.0, 5.1, 5.5, and 6.0.

To prepare to monitor a VMware vSphere server:

- Verify that you are running a supported version of the software.
- Obtain the user name and password of an account on the server that is authorized to use the vSphere API.
- Determine whether to use SSL.

Hyper-V hosts

Resource Manager uses WinRM to monitor the following Microsoft Hyper-V systems:

- Microsoft Hyper-V Server 2016
- Microsoft Hyper-V Server 2012 and 2012 R2
- Microsoft Hyper-V Server 2008 and 2008 R2

To prepare a Hyper-V system for WinRM monitoring, refer to the [support article](#) that describes the options and provides the procedures for configuring your systems.

Validating configuration using Inspector tool

Once you have set up your environment, you can validate your configuration using Inspector. The Inspector tool is typically installed on the Control Center master host and performs read-only checks on your environment and provides advice on resolving potential issues.

For more information on the Inspector tool, including download and installation instructions see the following knowledge base article: [Inspector: A tool to validate configuration](#).

External HBase configuration

Resource Manager can be configured to use an external HBase cluster, rather than the cluster that is included in the application.

If you do not already have an external HBase cluster, there is no need to create one. The procedures in this section are for customers who wish to use an existing HBase cluster for Resource Manager data.

The version of HBase installed in your external HBase cluster must be compatible with the version of OpenTSDB used by the Resource Manager application. The minimum supported version of HBase is 0.92.

Perform the procedures in the following sections in order.

- [Configuring OpenTSDB for an external HBase cluster](#)
- [Configuring the OpenTSDB service startup command](#)
- [Disabling the Resource Manager HBase cluster](#)

Configuring OpenTSDB for an external HBase cluster

To perform this procedure, install and start Resource Manager.

This procedure configures OpenTSDB to use an external HBase cluster, rather than the HBase cluster that is included in the Resource Manager application.

1. Log in to the Control Center browser interface.
2. In the Applications table, click Zenoss.resmgr.
3. Scroll down to the Services table and locate the OpenTSDB services reader and writer.
If you do not see reader and writer, expand the OpenTSDB service node.
4. Click reader or writer.
You will repeat the procedure for the other service.
5. On the service details page, scroll down to the Configuration Files table and in the Actions column, click Edit.
6. In the Edit Configuration dialog box, replace the value of the `tsd.storage.hbase.zk_quorum` key with the ZooKeeper quorum of the external HBase cluster.
 - a. Delete the existing value.
The default value is a [Go language template](#) expression.
 - b. Specify the ZooKeeper quorum of the external HBase cluster.
To specify a ZooKeeper quorum, create a comma-separated list of all quorum members. Specify each member of the quorum with a hostname or IP address, the colon character (:), and then the port number on which the ZooKeeper service is listening.
Note: If you use hostnames, the Control Center master host must be able to resolve them to IPv4 addresses, either through a nameserver on the network or through entries in `/etc/hosts`.
The following example shows the correct syntax for a 3-member ZooKeeper quorum:

```
zk-1.example.com:2181, zk-2.example.com:2181, zk-3.example.com:2181
```
 - c. Click Save.
7. At the top of the page, click Stop, and then click Start.
8. For the other OpenTSDB service (reader or writer), repeat the preceding steps.

Configuring the OpenTSDB service startup command

This procedure configures the OpenTSDB service to use the external HBase cluster on startup.

1. Log in to the Control Center browser interface.
2. In the Applications table, click Zenoss.resmgr.
3. Scroll down to the Services table and locate the OpenTSDB services reader and writer.
If you do not see reader and writer, expand the OpenTSDB service node.
4. Click reader or writer.
You will repeat the procedure for the other service.
5. Near the top of the service details page, click Edit Service.
6. In the Edit Service dialog box, change the value of the Startup Command field.
 - a. Delete the [Go language template](#) expression.
The expression is everything after start-opentsdb.sh.
 - b. Specify the ZooKeeper quorum of the external HBase cluster.
To specify a ZooKeeper quorum, create a comma-separated list of all quorum members. Specify each member of the quorum with a hostname or IP address, the colon character (:), and then the port number on which the ZooKeeper service is listening. Between start-opentsdb.sh and the ZooKeeper quorum list, include at least one empty space.
Note: If you use hostnames, the Control Center master host must be able to resolve them to IPv4 addresses, either through a nameserver on the network or through entries in /etc/hosts.
The following example shows the correct syntax for a 3-member ZooKeeper quorum:

```
zk-1.example.com:2181, zk-2.example.com:2181, zk-3.example.com:2181
```
7. Click Save Changes.
8. At the top of the page, click Stop, and then click Start.
9. For the other OpenTSDB service (reader or writer), repeat the preceding steps.

Disabling the Resource Manager HBase cluster

This procedure disables the HBase cluster that is included in the Resource Manager application.

1. Log in to the Control Center master host as root, or as a user with superuser privileges.
2. Stop the Resource Manager HBase cluster.

```
serviced service stop HBase
```

3. Disable automatic start of the HBase services.
 - a. Change the configuration of each service.

```
for svc in hmaster regionserver zookeeper
do
    serviced service list $svc | sed -e 's/"Launch": "auto"/"Launch": "manual"/' | serviced service
edit $svc
done
```

The serviced command displays the new configuration after each edit.

- b. Verify that each service is set to manual start.

```
for svc in hmaster regionserver zookeeper
do
    serviced service list $svc | egrep '"Launch":'
done
```

4. Remove the OpenTSDB prerequisite for the Resource Manager HBase cluster.

Depending on your version of Control Center, the OpenTSDB service is either opentsdb or two separate services, reader and writer.

 - a. Edit opentsdb, or one of reader or writer.

```
serviced service edit reader
```

The serviced command opens the service's configuration in the default text editor.

- b. Locate the Prereqs section, and then remove everything between the left square bracket ([) and the right square bracket (]) characters. The following lines show an example Prereqs section:

```
"Prereqs": [
  {
    "Name": "HBase RegionServers up",
    "Script": "{{with $rss := (child (child (parent) \"HBase\")).Instances }}"
  }
],
```

After editing, the section should look like the following example:

```
"Prereqs": [],
```

- c. Save the file, and then exit the text editor.
- d. If your version of Resource Manager includes two OpenTSDB services (reader and writer) repeat the preceding substeps for the other service.

Administering Resource Manager

- [Using Resource Manager](#)
- [Preparing devices for monitoring](#)
- [Working with devices in Resource Manager](#)
- [Basic monitoring](#)
- [Performance monitoring](#)
- [Distributed monitoring](#)
- [Monitoring Zenoss](#)
- [Extending Resource Manager with ZenPacks](#)
- [Using organizers](#)
- [Managing background tasks](#)
- [Using configuration properties](#)
- [Modeling](#)
- [About monitoring templates](#)
- [Production states and maintenance windows](#)
- [Event management](#)
- [Triggers and notifications](#)
- [Managing users in Resource Manager](#)
- [General administration and settings](#)

Using Resource Manager

- [Initial login](#)
- [Interface and navigation](#)
- [Administering Resource Manager dashboards](#)
- [Search](#)
- [Navigating the event console](#)
- [Running a command](#)
- [Visualizing your environment](#)

Initial login

The first time you log in to Resource Manager, you will immediately be taken to a startup wizard where you will perform the following tasks:

1. Set your admin password
2. Set your personal login
3. Discover devices (optional)
4. Add Infrastructure (optional)
5. Setup SMTP (optional)

To use the startup wizard, follow these steps:

1. Launch your Resource Manager application the first time by clicking on the Virtual Host Name in Control Center. You will be presented with the following page showing you the initial steps to follow:

Zenoss Installation Wizard

This wizard will guide you through the initial setup of Zenoss. Click Get Started to begin.

Step 1
Setup Users
Set the admin password and create your user account.

Step 2
Network Discovery
Discover devices to monitor.

Step 3
Add Infrastructure
Manually add the devices in your infrastructure.

Get Started »

2. Click Get Started to begin the wizard.

Step 1: Setup Users

Set admin password

The admin account has extended privileges, similar to Linux's root or Windows' Administrator. Its use should be limited to administrative tasks.

Password Must:

- Contain 8 or more characters
- Contain at least one number
- Contain at least one upper and lower case character

Admin password:

Confirm password:

Create your account

Enter information for your personal user account. You'll use this to perform most tasks.

Username:

Password:

Retype password:

Your email:

« Previous Next »

3. Set the password of the admin account and create your personal account. Click Next.

Network Discovery

Networks/Range
Enter one or more networks (such as 10.0.0.0/24) or IP ranges (such as 10.0.0.1-50):

SNMP
Community Strings:

SSH Authentication
Username:

Password:

Windows Authentication
Administrator Username:

Password:

Discover

4. Optional: Discover devices. If you are not ready to discover devices, you can skip this page and add devices later. For more information about the fields on this page see [Adding and discovering devices](#). To continue, click Next.

Step 3: Add Infrastructure

Category

- CiscoUCS
- ControCenter
- HTTP
- KVM
- Network
- Ping
- Power
- Printer
- Server
- Storage

Type

CiscoUCS

Connection Information

Enter multiple similar devices, separated by a comma, using either hostname or IP Address:

Username: admin

Password: *****

Port: 443

Use SSL?:

Add

Devices

Status	Host	Credentials	Type	Duration	Job Log	Remove	Retry
--------	------	-------------	------	----------	---------	--------	-------

Add infrastructure using the above form

Previous

Finish

5. Optional: Add devices. For each device, select the category, type and enter the connection information. If you are not ready to add devices, you can skip this page. When you are ready to continue, click Done. You will be taken to the Dashboard view of Resource Manager.

Step 4: Setup SMTP

Define SMTP server host, port, username, and password to enable email notifications

SMTP Host:

SMTP Port (usually 25):

SMTP Username (blank for none):

SMTP Password (blank for none):

From Address for Emails:

Use Transport Layer Security for E-mail?:

« Previous

✓ Finish

6. Optional: Setup SMTP. For email notifications, define your SMTP host, port, and credentials.

Note: When you launch Resource Manager in the future, you will go directly to the login screen.

Interface and navigation

After you install the system and navigate to the interface from your Web browser, the Resource Manager dashboard appears. The Dashboard provides at-a-glance information about the status of your IT infrastructure. It is the primary window into devices and events that the system enables you to monitor.

The screenshot displays the Resource Manager dashboard with the following components:

- Welcome to Zenoss!**: A sidebar menu with links to guides and training resources.
- Open Events**: A table listing system events with columns for Status, Severity, Resource, Component, Event Class, Summary, First Seen, and Last Seen.
- Production States**: A table showing the production state of various devices.
- Open Events Chart**: A bar chart showing the count of open events by severity level.

Status	Severity	Resource	Component	Event Class	Summ.	First Seen	Last Seen
⚠	Warning	ucs1.zb...	sys_rack...	iCiscoUCS...	Re...	2014-12-17 11:14:...	2014-12-17 11:14:...
⚠	Warning	ucs1.zb...	sys_rack...	iCiscoUCS...	rac...	2014-12-17 11:14:...	2014-12-17 11:14:...
⚠	Warning	oracle-s...		Status/PIn...	10...	2014-12-17 11:14:...	2014-12-17 11:14:...
⚠	Warning	ucs1.zb...	sys_mgm...	iCiscoUCS...	Fa...	2014-12-17 12:05:...	2014-12-17 12:05:...
⚠	Warning	ucs1.zb...	sys_mgm...	iCiscoUCS...	Fa...	2014-12-17 12:05:...	2014-12-17 12:05:...
⚠	Warning	ucs1.zb...	sys_mgm...	iCiscoUCS...	Fa...	2014-12-17 12:05:...	2014-12-17 12:05:...
⚠	Warning	ucs1.zb...	sys_rack...	iCiscoUCS...	Ac...	2014-12-17 01:52:...	2014-12-17 01:52:...
⚠	Warning	ucs1.zb...	org-root_...	iCiscoUCS...	iqn...	2014-12-17 11:14:...	2014-12-17 11:14:...
⚠	Warning	localhost		Unknown		2014-12-17 10:38:...	2014-12-17 10:38:...
⚠	Warning	ucs1.zb...	sys_rack...	iCiscoUCS...	eth...	2014-12-17 01:57:...	2014-12-17 01:57:...
⚠	Warning	ucs1.zb...	sys_rack...	iCiscoUCS...	eth...	2014-12-17 01:57:...	2014-12-17 01:57:...
⚠	Warning	ucs1.zb...	ESXI Rac...	iCiscoUCS...	Se...	2014-12-17 12:13:...	2014-12-17 12:13:...

Device	Production State
10.111.9.114	Production
oracle-srv.zenoss.loc	Production
test10mcal.zenoss.loc	Production
ucs1.zenoss.loc	Production

Severity	Count
Critical	5
Error	28
Warning	43
Info	36
Debug	0

For more information about the Resource Manager user interface, see [Resource Manager interface reference](#).

The Resource Manager dashboard can be customized to display a variety of information including:

- System information resources and other Web pages, such as internal portal pages
- Important error-level device events
- Geographical high-level view
- "Troubled" devices
- Devices in a certain production state. You can also create additional dashboards that can be cloned from existing dashboards and can also be restricted to
- certainly users or user groups.

Key Resource Manager dashboard and interface areas include:

- Navigation menu
- User information area
- Portlets
- System Network Map

User information area



The user information area offers information and selections:

- Search - Search area to find information within the application. Click the down arrow in the search box to manage your saved searches.
- Login ID - The ID of the user currently logged in appears at the far left of this area. Click the ID to edit user settings, such as authentication information, roles, and groups. (You also can access user settings from the **ADVANCED > Settings > Users** page.)
- Sign Out - Click to log out of the system.
- Help icon - Click to access product documentation.

Administering Resource Manager dashboards


Creating a new dashboard

A default administration dashboard is created when you launch Resource Manager. Administrators can customize this dashboard. However, the default dashboard cannot be deleted.

Users that are not administrators initially see a read-only version of the default administration dashboard. Non-administrators can create dashboards that display distinctive information or are targeted to a specific type of user or to only themselves. To customize a dashboard, select who can view it, and select and customize portlets to display the most important information. The number of customized dashboards is not limited.

To create a dashboard:

1. From the Add icon on the dashboard controls, select New Dashboard.



2. Use the following table to complete the fields in the dialog box.

Field	Description
Dashboard Name	A name for the dashboard. In the dashboard list, the name the user who creates a dashboard is displayed next to the dashboard name.
Who can view this Dashboard?	The name of the user account or user group that can view the dashboard. To specify a user group, you must be logged in as a member of the group to add.
Number of Columns	The number of columns to display in the dashboard. The default is 3.
Clone from dashboard	Create a dashboard with no portlets (the default) or copy the portlets of the current dashboard.


3. Click Create.

Adding a portlet to a dashboard

You can customize a dashboard by adding portlets. A dashboard can display multiple instance of the same portlet type. For example, several Device Chart portlets showing different device classes.

To add a portlet to a dashboard:

1. Select the dashboard to which you wish to add a portlet.
2. From the Add icon on the dashboard controls, and select Add portlet.




3. In the Portlet field, select a portlet.
The dialog box adds fields for configuring the selected portlet.
4. Edit the fields as appropriate for the selected portlet.
5. Click Add. The portlet is added at the top of the dashboard.

Editing dashboard settings

Customize a dashboard to display a different number of columns or limit access to the dashboard.

1. On the upper-right side of the dashboard, click the Action icon.



2. Use the following table to update the fields in the dialog box.

Field	Description
Dashboard Name	A name for the dashboard. In the dashboard list, the name the user who creates a dashboard is displayed next to the dashboard name.
Who can view this Dashboard?	The name of the user account or user group that can view the dashboard. To specify a user group, you must be logged in as a member of the group to add.
Number of Columns	The number of columns to display in the dashboard.
Audit logs?	Update the audit log when changes are made to the dashboard.
Lock from updates?	Prevent unauthorized edits of the dashboard.


3. Click Save.

Moving portlets in a dashboard

To arrange portlets, click the portlet header and drag the portlet to any location on the dashboard. The other portlets move to accommodate the moved portlet's position.

Portlets

Use the following buttons to control how a portlet is displayed.



Button	Description
Collapse	Shrink the portlet so that only the title appears on the dashboard.
Pop Out	Expand the portlet to occupy the entire browser window.
Edit	Edit portlet settings.
Close	Remove the portlet from the current dashboard.

In tabular portlets, you can control the display by sorting columns as well as adding and hiding columns.

- To sort based on a column, hover over the column header and click the arrow to display the sort and display options.
- To add or hide columns, hover over the Columns entry and check or clear the boxes of the columns to add or hide.

List of available portlets

Daemon Processes Down

Contains system self-monitoring information.

Device Chart

Allows the display of a graph of multiple data points for a selected device class.

Device Issues

Displays a list of devices associated with color-coded events of critical, error, or warning severity levels. To view details, click a device name. To go to the event console for the device, click an event.

Event View

Displays a list of events similar to the view on the Event console. Event management buttons are not provided in the Event View portlet.

Google Maps

Shows configured device locations and network connections.

HTML Portlet

Displays HTML content. You must use HTML markup in this portlet. If you want to populate a portlet with content from a specific URL, use the Site Window portlet instead.

Multi-Graph Report

Displays an existing Multi-Graph Report (created by using the Reports page). You can choose a specific graph group from the multi-graph report and select the time range for the portlet.

Network Map

Displays a network map for a defined network that is being monitored. You can define the refresh interval and level of depth of the map.

Open Events Chart

Displays a bar graph of the number of open events, grouped by severity. You can define the event class to be display and the number of days for which to show events.

Organizers

Choose the root organizer (devices, locations, systems, groups) and then child organizer options are enabled.

Past Events Line Chart

Displays a line chart of past events grouped by severity. You can define the event class to be used and the number of past days for which to show events.

Production States

Shows devices assigned to a particular production state. If needed, you can define multiple production states to display.

Site Window

Initially provides links to resources such as product guides, forums, and training events.

You can customize the portlet to display any web page. However, Zenoss recommends that you keep a portlet with the default URL so that you can stay up-to-date with Zenoss training and product updates.


Watch List

Allows the display of high-level status of device classes, groups, systems, event classes, and locations that you select.


Search

The Resource Manager search facility supports locating devices and other system objects, as well as events and services.

In the Resource Manager interface, the search feature is part of the user information area. Enter part or all of a name in the search box at the top right of the interface. The system displays matches, categorized by type.



To view all search results, click the indicator at the top of the list.



From here, you can display search results by category. Click in the left panel to filter search results by a selection.

You can save the search to access later.

1. Click **Save As** (at the bottom left of the Search Results page). The Save Search As dialog box appears.
2. Enter a name for the search, and then click **Submit**.

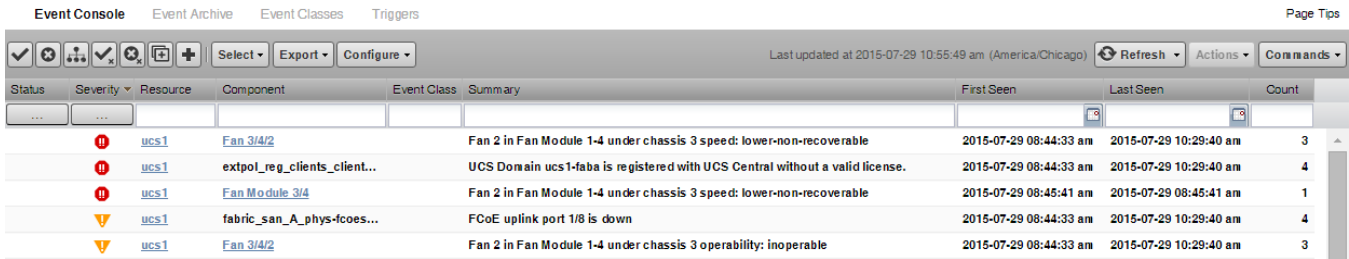
You can access saved searches from:

- Action menu located at the bottom of the Search Results page.
- Search box located at the top of the interface. Click the arrow, and then select Manage Saved Searches.

Navigating the event console

The event console is the system's central nervous system, enabling you to view and manage events. It displays the repository of all events that are detected by the system.

To access the event console, click Events in the Navigation menu.



The screenshot shows the Event Console interface with a table of events. The table has columns for Status, Severity, Resource, Component, Event Class, Summary, First Seen, Last Seen, and Count. The events listed are:

Status	Severity	Resource	Component	Event Class	Summary	First Seen	Last Seen	Count
Warning	Warning	ucs1	Fan 3/4/2	Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable	Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	3
Warning	Warning	ucs1	extpoi_reg_clients_client...	UCS Domain ucs1-faba is registered with UCS Central without a valid license.	UCS Domain ucs1-faba is registered with UCS Central without a valid license.	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	4
Warning	Warning	ucs1	Fan Module 3/4	Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable	Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable	2015-07-29 08:45:41 am	2015-07-29 08:45:41 am	1
Warning	Warning	ucs1	fabric_san_A_phys-fcoes...	FCoE uplink port 1/8 is down	FCoE uplink port 1/8 is down	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	4
Warning	Warning	ucs1	Fan 3/4/2	Fan 2 in Fan Module 1-4 under chassis 3 operability: inoperable	Fan 2 in Fan Module 1-4 under chassis 3 operability: inoperable	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	3

- [Sorting and filtering events](#)
- [Creating an actionable view](#)
- [Saving a custom view](#)
- [Refreshing the view](#)
- [Viewing event details](#)
- [Selecting events](#)
- [Managing events](#)

Sorting and filtering events

You can sort and filter events that appear in the event console to customize your view.

You can sort events by any column that appears in the event console. To sort events, click a column header. Clicking the header toggles between ascending and descending sort order.

Filter options appear below each column header. A match value can be any full string or a subset of a string, optionally with the wildcard (*) contained in the values in that column. You can also use "||" (OR), or "!!" (NOT) expressions to further target your filters. For example, typing !!windows in the Event Class filter will return all the non-Windows device events.

Status	Severity	Resource	Component	Event Class	Summary	First Seen	Last Seen	Count
<input checked="" type="checkbox"/> New		ucs1	Fan 3/4/2		Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	3
<input checked="" type="checkbox"/> Acknowledged		ucs1	extpol_reg_clients_client...		UCS Domain ucs1-faba is registered with UCS Central without a valid license.	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	4
<input type="checkbox"/> Suppressed		ucs1	Fan Module 3/4		Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable	2015-07-29 08:45:41 am	2015-07-29 08:45:41 am	1
<input type="checkbox"/> Closed		ucs1	fabric_san_A_phys-fcoes...		FCoE uplink port 1/8 is down	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	4
<input type="checkbox"/> Cleared		ucs1	Fan 3/4/2		Fan 2 in Fan Module 1-4 under chassis 3 operability: inoperable	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	3
<input type="checkbox"/> Aged		ucs1	Fabric FC SAN Port Chan...		san port-channel 5 on fabric interconnect A oper state: failed, reason: No operatio...	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	4
		ucs1	Fabric FC SAN Port Chan...		san Member 1/27 of Port-Channel 5 on fabric interconnect A is down, membership:...	2015-07-29 08:46:12 am	2015-07-29 10:29:40 am	4

You can filter the events that appear in the list in several ways, depending on the field type:

- **Resource** - Enter a match value to limit the list.
- **Component** - Enter a match value to limit the list.
- **Event Class** - Enter a match value to limit the list.
- **Summary** - Enter a match value to limit the list.
- **First Seen** - Enter a value or use a date selection tool to limit the list.
- **Last Seen** - Enter a value or use a date selection tool to limit the list.
- **Count** - Enter a value to filter the list, as follows:
 - *N* - Displays events with a count equal to *N*.
 - *:N* - Displays events with a count less than or equal to *N*.
 - *M:N* - Displays events with a count between *M* and *N* (inclusive).
 - *M* - Displays events with a count greater than or equal to *M*.

To clear filters, select Configure > Clear filters.

You also can re-arrange the display order of columns in the event console. Click-and-drag column headers to change their display.

Creating an actionable view


For users that are not Administrators, an option filters the list of events to show only those that are not read-only for the user's permission level, and enable the action buttons above the event table header.

To turn on the actionable view, click Configure and select the Only show actionable events check box. The view is changed to show only events that can have an action performed on them based on the user's permission level.

Saving a custom view

Save a custom event console view by bookmarking it for quick access.

1. Select Configure > Save this configuration.
2. In the dialog box, select the link, and then drag it to the bookmarks area of the browser window.



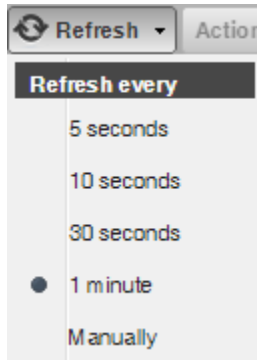
The browser adds a link to the bookmarks list.

3. Change the title of the bookmark to distinguish this event console view.

Refreshing the view

You can refresh the list of events manually or specify that they refresh automatically. To manually refresh the view, click **Refresh**. You can manually refresh at any time, even if you have an automatic refresh increment specified.


To configure automatic refresh, select one of the time increments from the Refresh list. By default, automatic refresh is enabled and set to refresh each minute.



Viewing event details

You can view details for any event in the system. To view details, double-click an event row.

Do not double-click on or near the device (resource) name, component, or event class in the row. Doing this displays details about that entity, rather than information about the event.



To see more information about the event, click Event Details.

You can use the Log field (located at the bottom of the area) to add specific information about the event. Enter details, and then click Add.


Selecting events

To select one or more events in the list:

- To select a single event, click a row.
- To select multiple events, Ctrl-click each row or Shift-click rows to select a range of events.
- To select all events, click Select > All.

Managing events

Use the event console to manage events or add an event. Click an event row and use the following tools to perform actions.



- Acknowledge the event.
- Close the event.
- Reclassify the event by associating it with a specific event class.
- Return the event to New status (revoke its Acknowledged status).
- Reopen the event.
- Add a note to the log.
- Add an event. This feature is useful for testing a specific condition by simulating an event.


Running a command

You can run commands on a single device or on a group of devices.

The system includes several built-in commands, such as `ping` and `traceroute`.

To run commands from the interface:

1. Navigate to the INFRASTRUCTURE tab.
2. In the Devices list, select one or more devices.
To select a device, click anywhere in the row except on the link.
3. Click Commands and select a command from the list.



You can resize the command output window. You also can stop automatic scrolling by de-selecting the Autoscroll option at the bottom right corner of the output window.


Visualizing your environment

- [Dynamic service view](#)
- [Datacenter view](#)

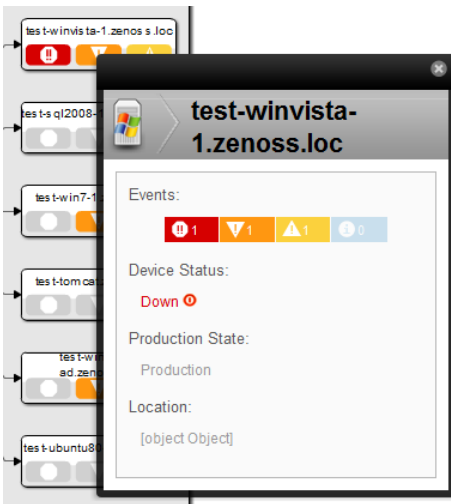
Dynamic service view

Resource Manager provides a dynamic visualization of system objects and their relationships to other objects.

You can access the dynamic view from groups, systems, and locations. Depending on the object type, different relationships are illustrated. Each dynamic view shows related objects in a graph. Each object in that graph displays its associated event information. See the following procedures for creating a particular dynamic view you are interested in. The following figures show some example dynamic service views and navigation aids.

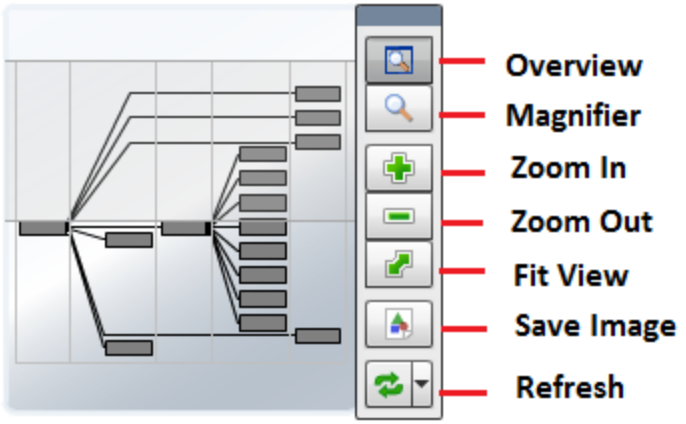


When you click an object in the graph, the "inspector" panel appears. This panel provides detailed information about the object and links directly to it. Information that appears in the inspector depends on the object type selected.



View controls appear to the right of the graph. These allow you to adjust your view:

- **Overview** - Toggles display on and off of the graph overview illustration.
- **Magnifier** - Toggles on and off the magnifier, which allows you to magnify selected portions of the graph.
- **Zoom In** - Zooms in on the graph.
- **Zoom Out** - Zooms out on the graph.
- **Fit View** - Fits the graph to the browser page.
- **Refresh** - Refreshes the graph.



Dynamic view of organizers

The dynamic view of organizers shows objects that can impact the status of the organizer, such as other organizers and devices. This view also shows relationships between devices and a virtual infrastructure, such as VMware or Cisco UCS objects monitored by the system, as well as storage information.

To access the dynamic view for an organizer (such as a group, system, or location):

1. From **INFRASTRUCTURE > DEVICES**, select the organizer in the devices hierarchy.
2. Click **Details**.
3. Select **Dynamic Service View**.

Dynamic view of devices

The dynamic view of devices shows the relationship between a device and monitored components.

To access the dynamic view for a device:

1. From INFRASTRUCTURE > DEVICES, click a device in the device list. The device overview page appears.
2. Select Dynamic View in the left panel.

Dynamic view of Cisco UCS devices

On Cisco UCS devices, the dynamic view shows the components and relationships that make up a Cisco UCS domain. The following list describes the components that are listed in the Dynamic View for various types of UCS devices:

- UCS Classic
 - UCS Domains
 - UCS Fabric Interconnects
 - UCS Fabric Ports
 - UCS IO Cards
 - UCS Fabric Extenders
 - UCS Chassis
 - UCS Blades
 - UCS Racks
 - UCS Service Profiles (only those bound to servers)
- UCS Mini
 - UCS Domains
 - UCS FI-IO Modules
 - UCS Chassis
 - UCS Blades
 - UCS Service Profiles (only those bound to servers)

Dynamic view of VMware hosts

On VMware Hosts (ESX servers), the dynamic view shows the relative VMware elements that are connected to the host, such as:

- VMs that currently are running on the Host
- Data stores that are mounted by the Host
- Clusters to which the Host belongs

Dynamic view of storage devices


Storage devices such as NetApp Filers have two dynamic views:

- **Physical Storage View** - Shows the device's storage enclosures and associated hard disks.
- **Logical Storage View** - Shows the logical storage arrangement that the storage device presents, such as file systems and RAID groups.

Datacenter view

Resource Manager provides a visual representation of devices (such as servers or blades) and device containers (such as racks or chassis).

With this feature, you can create a custom view that represents a physical space (such as a data center) by customizing the view background. You can then overlay this view with active representations of your devices and device containers.



For each device or device container, the system can generate a rack view, which diagrams the physical location of devices in a chassis or rack. Each represented device provides at-a-glance information about its status.



Before a device or sub-location can appear in Datacenter View:

- At least one organizer must be configured
- At least one device or sub-organizer must be included in a location

To see the auto-generated rack view, you must set a rack slot value for the device.

Working with the List View

The List View provides a view of your devices (or, if configured, the Rack View). Follow these steps to access the List View:

1. From the interface, select Infrastructure.
2. In the Devices hierarchy, select a location, group, or system.
3. Click Details.
4. Select Datacenter View.

After you create a Custom View, that view appears by default.

Working with the Custom View

The Custom View lets you create a visual representation of your physical space (such as a data center).

To access the Custom View, from the Diagram selection, click Custom View.

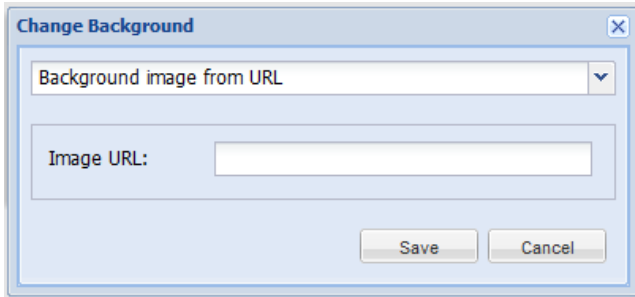
You can edit the Custom View to:

- Add or change a background image
- Move or resize device images
- Remove the view

Adding a background image to a custom view

Follow these steps to create a custom view and add a background image to the view:

1. From the Datacenter View page (accessed from the Diagram selection), click Custom View.
2. Click Edit to enable edit mode.
The Edit button highlights to indicate that it is active, and Options selections become available.
3. Select Options > Change Background.
4. In the Change Background dialog, select Background Image from URL.
5. In the Image URL field, enter an image location, and then click Save.
Any image format and size supported by your browser can be used.



Removing a custom view background image

To remove the current background image from the Custom View:

1. From the Custom View area, click Edit.
2. Select Options > Change Background.
3. In the Change Background dialog, select No background image from the list of options.
4. Click Save.
The image no longer appears in the view.

Working with devices in the custom view

Devices in the custom view can be moved and resized. To work with devices in this view, click **Edit**. You can then drag devices to a specific location in the view, and resize them to accurately represent your physical space.

You also can view device details from this view. Click the device to go to its Status page.

Note: To access device status, you cannot be in edit mode.

Removing a custom view

Removing the custom view removes the view and custom background image, if any. To remove a custom view:

1. From the Datacenter View page (accessed from the Diagram selection), click Custom View.
2. Click Edit to enable edit mode.
3. Select Options > Remove Custom View.

The custom view no longer appears by default. If you select Custom View, devices still appear in the view; however, they are reset to default positions and sizes.

Activating the auto-generated rack view

First, ensure that the device is included in a location. Then follow these steps to make devices visible in Datacenter View.

1. Select a device, from the list of Devices, click Details, and then click Edit.
2. Enter values for Rack Slot.
For more information about specifying values, see the next topic.
3. Click Save.

The device appears in Datacenter View.

- In the List View, it appears as part of a rack illustration. The rack illustration becomes the default image in the List View.
- In the Custom View, it appears as a single device image.

Note: You can customize the device image by modifying the zlcon configuration property in the device class.


Rack view specification syntax

The specification of a rack view features two required and two optional key-value pairs, which are described in the following table. The value for each key is a positive integer.

Key	Type	Description
ru	Required	The lowest rack unit used by the device
rh		The height of the device, in rack units
st	Optional	The rack slot
sc		The slot capacity (only for chassis devices)

The syntax of a rack view specification is a comma-separated list. The following example specifies a 1u device in the second rack unit for a device named `beta.zenoss.loc`:

```
ru=2, rh=1
```



In the preceding example, no rack slot value is included, because there is only one device.

Preparing devices for monitoring

Resource Manager uses collectors to gather data from a wide variety of devices, using the method that is best suited to each device type. For example, most network devices (switches, routers, and so on) support the Simple Network Management Protocol (SNMP). Windows devices support WinRM, which is a [WSMAN](#)-based mechanism for communicating with devices running Microsoft Windows.

For each device type, some end-device configuration may be required. For more information about preparing specific devices, see the [ZenPack catalog](#).

Examples:

[Microsoft Windows ZenPack Documentation: Setting up WinRM Service for Target Windows Machines](#)

[Linux Monitor ZenPack Documentation: Set Linux Server Monitoring Credentials](#)

Configuring Linux devices to provide data through SNMP

To configure a Linux machine for monitoring, it must have SNMP installed. A good Linux SNMP application is [net-snmp](#). Download, install, and configure net-snmp to then use SNMP to monitor Linux devices.

Configuring Windows devices to provide data through SNMP

To monitor Microsoft Windows Server 2008 R2 systems, Zenoss Cloud uses SNMP v1/v2 or WinRM. (There is no SNMP v3 support.) For Windows Server 2012 and Windows Server 2016, Microsoft has deprecated SNMP support.

By default, Windows may not have SNMP installed. To install SNMP on your particular version of Windows, please refer to the Microsoft documentation.

After setting up and configuring the SNMP service, you must set the zSnmpCommunity string in Resource Manager to match, to obtain SNMP data.

To set up WinRM on a Windows device, refer to the article in the [Microsoft Windows ZenPack Documentation: Setting up WinRM Service for Target Windows Machines](#)

Working with devices in Resource Manager

This section includes information and procedures for managing devices in Resource Manager.

- [Viewing the device list](#)
- [Working with devices](#)
- [Managing devices and device attributes](#)
- [Adding and discovering devices](#)

Viewing the device list

The device list shows all devices in the system. From this view, you can search for devices and perform a range of management tasks on all devices.

To access the device list, from the navigation menu, select INFRASTRUCTURE.

The screenshot shows a web-based interface for managing infrastructure. The top navigation bar includes tabs for 'Devices', 'Networks', 'Processes', 'IP Services', 'Windows Services', 'Network Map', and 'Manufacturers'. The 'Devices' tab is active, and the page title is '/Devices'. Below the navigation bar, there is a search bar and a toolbar with buttons for 'Select', 'Configure', 'Export', 'Refresh', 'Actions', and 'Commands'. A status bar at the top indicates 2 errors, 29 warnings, and 10 alerts. The main content area is a table with the following columns: Device, IP Address, Device Class, Production State, and Events. The table lists 10 devices, including various server and network configurations. A left-hand navigation menu shows a tree view of infrastructure categories like 'DEVICES (10)', 'GROUPS (0)', 'SYSTEMS (0)', 'LOCATIONS (0)', and 'COMPONENT GROUPS (0)'. The bottom of the interface shows 'DISPLAYING 1 - 10 of 10 ROWS' and '0 Jobs'.

Device	IP Address	Device Class	Production State	Events
127.0.0.1		/ZenossRM	Production	1
per72_switcb_zenoss_loc	10.87.207.10	/Network/Cisco	Production	
per72_vcenler_zenoss_loc		/vSphere	Production	6
qa-centos-6_zenoss_loc	10.88.120.104	/Server/Linux	Production	
qa-centos-7_events-ssh_zenoss_lab	10.88.120.73	/Server/SSH/Linux	Production	
qa-centos-7_events_zenoss_lab	10.88.120.72	/Server/Linux	Production	
qa-centos-7_ssh_zenoss_lab	10.88.120.74	/Server/SSH/Linux	Production	
qa-centos-7_zenoss_loc	10.88.120.105	/Server/Linux	Production	1
resmgr_hyperv_80x_zenoss_lab	10.88.120.224	/ControlCenter	Production	
ucs1_zenoss_loc	10.87.208.163	/CiscoUCS/UCS-Manager	Production	1

Devices hierarchy

Devices are organized in the tree view by:

- Devices
- Groups
- Systems
- Locations
- Component Groups

Click the indicator next to each category name to expand it and see included devices.

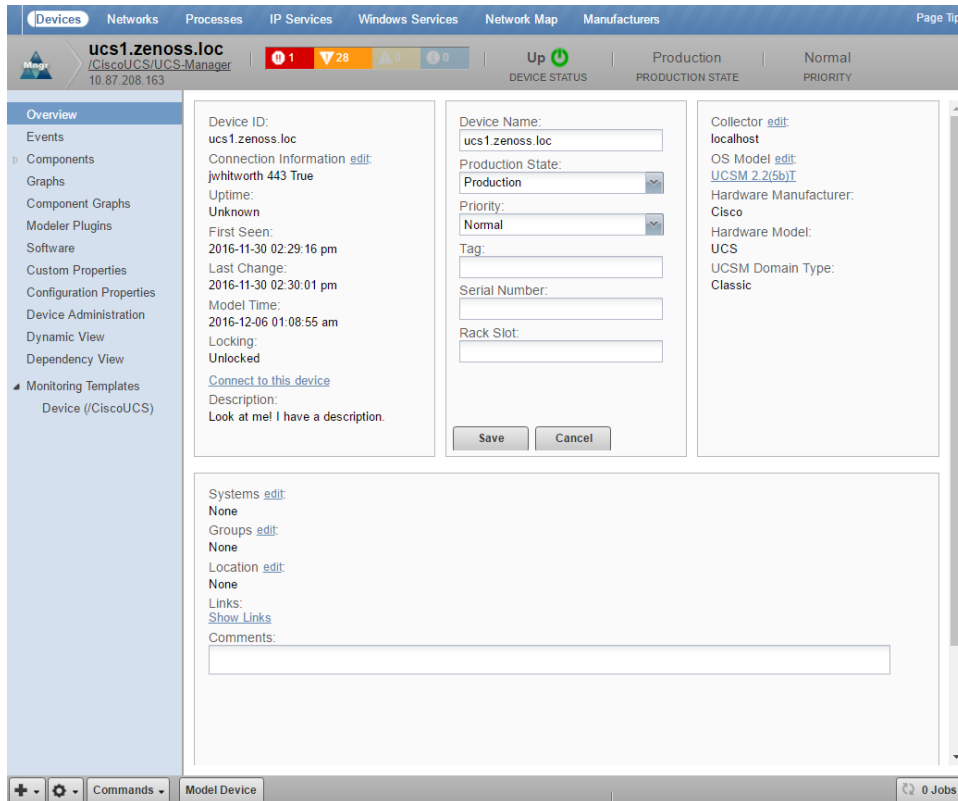
Managing multiple devices from the device list

You can perform some management tasks for more than one device at a time. You can:

- Move devices to a different class
- Assign devices to groups, systems, and locations
- Remove devices
- Perform actions such as assigning priority, production state, or collector
- Lock devices

Working with devices

To view details for a single device, click its name in the device list. The device overview page appears.



Event status is shown in the "event rainbow" at the top of the page. Other key information that appears at the top of the device overview page includes:

- Device name
- IP address used to communicate with the device
- Device status (shows the current results of a ping test)
- Production state (Pre-Production, Production, Test, Maintenance, or Decommissioned)
- Priority

When you open the page, device overview information displays. This view provides classification and status information. From here, you can edit device information (indicated by text fields or edit links). Editable fields include:

- Device Name
- Production State
- Priority
- Tag
- Serial Number
- Rack Slot
- Collector
- Hardware and software manufacturer and model
- Systems
- Groups
- Location

The Links area displays links between the device and other external systems. Click Show Links to view the links.

The left panel of the device overview page allows you to access other device management views, such as:

- Events
- Components
- Graphs (Performance)
- Component Graphs
- Modeler Plugins
- Software
- Custom Properties
- Configuration Properties
- Device Administration
- Dynamic View
- Dependency View
- Monitoring Templates

Information that appears here varies depending on device type.

Dependency view

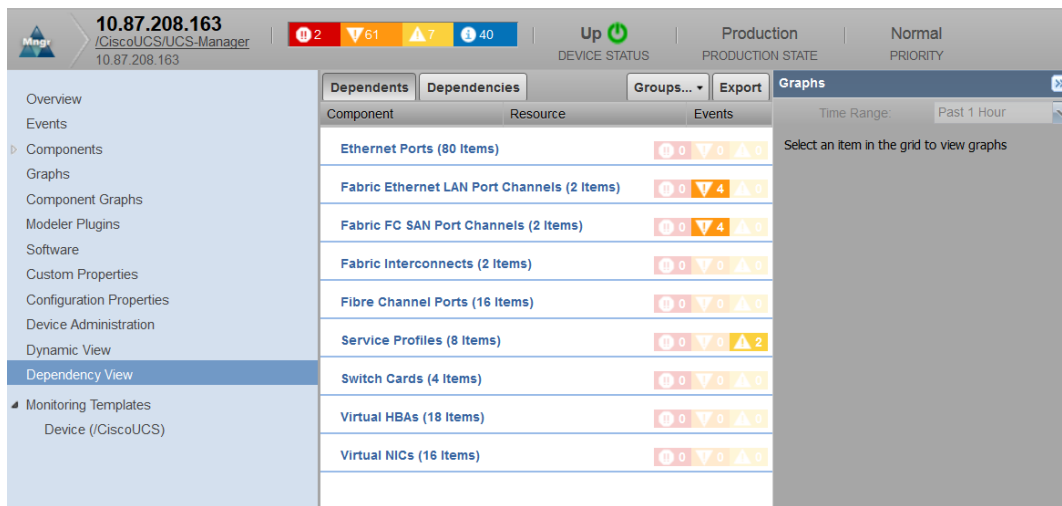
The dependency view of a device shows the resources that are dependent on the selected device as well as those resources that the device is dependent on. In this view, you can see the resource, associated component, and any events related to the component.

The following table identifies the places where a dependency view is available.

Location	How to display
Device Overview page	Navigate to the device overview page and select Dependency View
Device Component page	Navigate to the device overview page and select the component you are interested in. In the Display drop-down list, select Dependencies.
Group Details page (including Groups, Systems, and Locations)	Navigate to the Group, System or Location name and click Details. Then, select Dependency View

Regardless of how you navigate to the Dependency View, the functionality remains the same.

Sample Dependency View:



Using Dependency View:


1. Click Dependents to see the resources that are dependent on the selected device or component. Click Dependencies to see the resources that this device or component is dependent on. For example, if you click the Dependency View of a Hyper-V server, you will see that VMs, datastores, and other resources are dependent on the Hyper-V server, while the server itself is dependent on the Host CPUs, HDDs, and Network Adapters (seen after clicking Dependencies).
2. Click on a Component to expand its dependents or dependencies.
3. Click Groups to select the resources that you want to display in the view.
4. Export the data displayed by clicking the Export button. A CSV file is exported with the data as it is presented in the view.
5. Click a resource name to be taken to its overview page which displays performance graphs and other details.

All selections made with respect to the display of the view will be saved so that when a particular user selects another dependency view, Resource Manager renders the dependency view using the same parameters. This is only applicable for the same user. There is no global setting available for the dependency view.

Dynamic view

Resource Manager provides a dynamic visualization of system objects and their relationships to other objects.


You can access a dynamic view from a device overview, a group, a system, or a location. Depending on the object type, different relationships are illustrated. Each dynamic view shows related objects in a graph. Each object in that graph displays its associated event information.



When you click an object in the graph, the "inspector" panel appears. This panel provides detailed information about the object and links directly to it. Information that appears in the inspector depends on the object type selected.

View controls appear to the right of the graph. These allow you to adjust your view:

- **Overview** - Toggles display on and off of the graph overview illustration.
- **Magnifier** - Toggles on and off the magnifier, which allows you to magnify selected portions of the graph.
- **Zoom In** - Zooms in on the graph.
- **Zoom Out** - Zooms out on the graph.
- **Fit View** - Fits the graph to the browser page.
- **Save Image** - Saves the dynamic view as a PNG image.
- **Refresh** - Refreshes the graph.



Events view

Detailed information about events, scoped to the device, appears in the Events view. From here, you can:

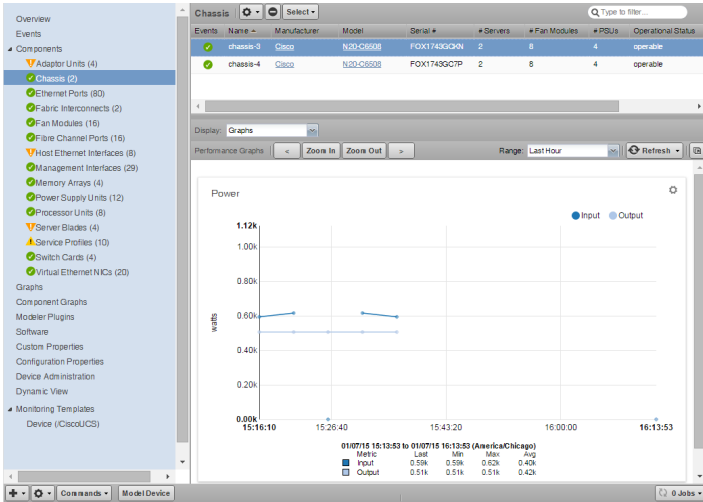
- Sort event and event archive information by a range of categories
- Classify and acknowledge events
- Filter events by severity, state, or by one of several categories

Components

The Components view provides information about the different types of device components, including:

- IPService
- WinService
- IpRouteEntry
- IpInterface
- CPU
- FileSystem

To access components information, select Components in the left panel, and then select a component type. The components available will vary based on the type of device.



The status of each device component type, as shown by the color of its indicator, is determined by the collective status of the monitored components of the same type. For example, if the IpService status is green, then all monitored IpServices on the device are functioning normally. If there is an event related to a monitored IpService, then the highest severity event associated with that component is displayed.


Note: If there is an event unrelated to a known component, then the system places it in the component type Other.

From this view, you can:

- Lock components
- Turn on or off component monitoring
- Delete components

Component graphs

To access component graphs for a device, in the left panel, choose Component Graphs. All metrics are displayed on the same graph.



You can control the following component graph options:

- Component type - From the drop-down list, choose from the available components, based on the type of monitored device.
- Graph type - From the drop-down list, choose from the available graphs, based on the selected component type.
- All on same graph - Choose whether to display all metrics on one graph or to have a separate graph for each metric.
- Zoom In/Zoom Out - Adjust the magnification of the view.
- Center the graph - Click a point inside the graph to center the graph around the selected point. If other graphs are displayed in a view, they are also affected.
- Scroll through time on the graph - Click the forward and back arrowheads. Using these controls automatically selects a custom time range.
- Time range - From the drop-down list, choose from the available ranges. For a custom range, fields appear in which you can specify start and end dates and times. For the current date and time, choose Now.
- Refresh the graph manually, or from the drop-down list, choose a refresh interval.
- Expand the graph to a full screen view.

Disabling component monitoring


There may be occasions when you want to stop monitoring certain components of your monitored resources.

To disable monitoring on one or more components:

1. On the Device overview page, select the component group.
2. In the Component list, select the components for which you want to disable monitoring.
3. Click Action > Monitoring.
4. Click NO to disable monitoring.
You may want to clean up the Event log of any events that were created by these components prior to the disabling of monitoring.

Graphs (Performance)

The Graphs view shows performance graphs that are defined for the device or component. To access graphs, in the left panel, select Graphs.



You can control the following performance graph options:

- Time range controls - To narrow or expand the size of the time range, click Zoom In/Zoom Out. To scroll through time on the graph, click the forward and back arrowheads. Clicking these controls automatically puts you into a custom time range. You can also Zoom In and center around a graph point by clicking the point inside the graph. This will center the graph around the selected point. If there are other graphs displayed in a view, they will also be affected by any change to the range.
- Range - Select the span of time that the graph displays. Depending on the range, data may be [downsampled](#). By default, the behavior is to display an average per downsampled point, but it can be changed on a per-graph basis (see below).
 - Last Hour - This is the default view and no aggregation/downsampling occurs in this view.
 - Last 24 Hours - Data is downsampled to 5-minute points.
 - Last Week - Data is downsampled in 1-hour points.
 - Last 30 days - Data is downsampled in 6-hour points.
 - Last Year - Data is downsampled in 10-day points.
 - Custom - Select the Start and End time to display. To set the end time to the current time, check Now. After changing a custom range setting, click Refresh to update the graph. Data displayed in this view may or may not be downsampled, depending on the range chosen.
- Refresh - To modify the refresh value (by default, 30 minutes), click the drop-down list. If you set the refresh rate to manual, click Refresh each time you want an updated graph.
- Pop-out - To render the current graphs in full-screen mode, click the icon in the upper right corner of the page.
- Downsample method - To set the downsample method, choose the drop-down in the upper right corner of a graph. There are four options:
 - Min - the minimum value during the downsampled range
 - Max - the maximum value during the downsampled range
 - P95 - calculates the 95th percentile during the downsampled range
 - Avg - (default) the average value during the downsampled range
- Action (gear) - To open a submenu of the following actions, click the icon in the upper right corner of the graph.
 - Definition - View the JSON definition.
 - Export to CSV - Export the datapoints as a .csv file for use in a spreadsheet. Only data contained in the defined range will be included.
 - Link to this Graph - Generate a link to this graph to save in browser bookmarks or use the URL to directly point to the graph in another Web page or dashboard. For example, to show the graph in the Dashboard, create a Site Window portlet and insert the URL to the graph.

If a graph has more than 28 graph points, you may receive an error stating "Unable to generate link, length is too great" when trying to [link to it](#). This is a known limitation and no workaround exists.

 - Expand graph - Render the current graph in full-screen mode.
- Table Legend - To highlight a particular data set, hover the pointer over a legend description. To toggle the displayed legend description, click it. A solid dot indicates that data will be displayed. A hollow dot indicates data will be hidden.

Modeler plugins

Use the Modeler Plugins view to manage plugins that are run against a device. To access plugins, select Modeler Plugins in the left panel.

Configuration properties

From the Configuration Properties view, you can set certain configuration properties for a device, and delete local properties from a device.

To access configuration properties, in the left pane, choose Configuration Properties.

Configuration Properties view showing a table of properties for a device. The table has the following columns: Is Local, Category, Name, Value, and Path.

Is Local	Category	Name	Value	Path
	Cisco	zCiscoACEUseSSL	true	/
	Cisco	zCiscoRemodeEventClassKeys		/
	Cisco UCS	zCiscoUCSCIMCEventsInterval	60	/
	Cisco UCS	zCiscoUCSCIMCPerInterval	300	/
	Cisco UCS	zCiscoUCSMManagerPort	443	/
	Cisco UCS	zCiscoUCSMManagerUseSSL	true	/
	Cisco UCS	zCiscoUCSMManagerUser	admin	/
	Modeler Controls	zCollectorClientTimeout	180	/
	Modeler Controls	zCollectorDecoding	utf-8	/
	Misc	zCollectorLogChanges	false	/
	zencommand	zCommandCommandTimeout	15	/
	zencommand	zCommandExistenceTest	test -f %s	/
	zencommand	zCommandLoginTimeout	10	/
	zencommand	zCommandLoginTries	1	/
	zencommand	zCommandPassword		/
	zencommand	zCommandPath	\$ZENHOME/libexec	/
	zencommand	zCommandPort	22	/
	zencommand	zCommandProtocol	ssh	/
	zencommand	zCommandSearchPath		/
	zencommand	zCommandUsername		/
	Control Center	zControlCenterHost	\$(here/managerip)	/
	Control Center	zControlCenterModelCycle	3600	/
	Control Center	zControlCenterPassword		/
	Control Center	zControlCenterPortCycle	onn	/

DISPLAYING 1 - 24 of 147 ROWS

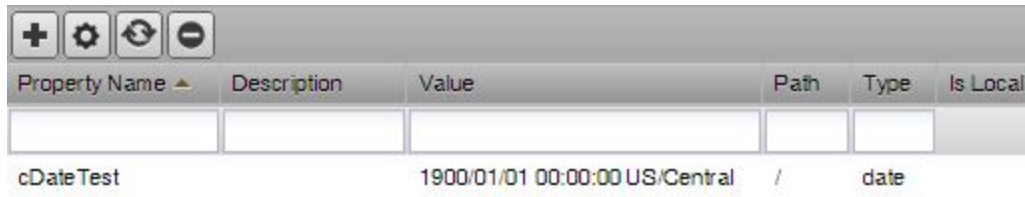
0 Jobs

Custom properties

You can edit custom property values that are associated with a device.

In the device list, click the device name. In the left panel, choose Custom Properties. You can perform the following actions on custom properties:

- Add
- Edit
- Refresh
- Delete



Property Name	Description	Value	Path	Type	Is Local
cDateTest		1900/01/01 00:00:00 US/Central	/	date	

Note: The Custom Properties view allows you to edit the value of a custom property on an individual device, but not to define new custom properties for device classes.

Device administration

Use the Device Administration view to:

- Add, delete, and run custom user commands
- Manage maintenance windows
- Determine who holds administration capabilities for the device, and their roles

To access administration options, select Device Administration in the left panel.

The screenshot displays the Device Administration interface. On the left is a navigation pane with 'Device Administration' selected. The main area is divided into three sections:

- Maintenance Windows:** A table with columns: Enabled, Name, Start, Duration, Repeat, State.

Enabled	Name	Start	Duration	Repeat	State
No	1st of Month	2014/12/01 02:00:00	01:00:00 hrs	Monthly: day of month	Maintenance
Yes	Every Thursday Night	2014/12/11 22:00:00	30:00 mins	Weekly	Maintenance
Yes	One Time Testing	2014/12/20 08:00:00	04:00:00 hrs	Never	Test
- User Commands:** A table with columns: Name, Command.


Name	Command
DNS forward	host \${device.id}
DNS reverse	host \${device.manageIp}
ping	\${device.pingCommand} -c2 \${device.mana...
snmpwalk	snmpwalk -s \${device.zSnmpVer} -c \${device.z...
traceroute	\${device.tracerouteCommand} -q -w 2 \${d...
- Administrators:** A table with columns: Name, Role, Email, Pager.

Name	Role	Email	Pager
caichrist	Manager	cgil...	

Overriden objects

Use the Overriden Objects view to see the objects that have overrides on their configuration properties. This view is available when looking at details of all devices.

To display the Overriden Objects view, navigate to the INFRASTRUCTURE > Devices pages and click Details. Then, click Overriden Objects from the left-column menu.



Select a configuration property from the drop-down list to view the overridden objects for that property. Double-click the row of the overriding object to open an edit dialog box.

Do not click the link of the overriding object. You will be taken to that object's page in the infrastructure view. Instead, double-click the clear area of the row of the overriding object to view the Edit Configuration Property dialog.

Software

The Software view lists software installed on the device. The details provided in this area depend on the method used to model the device.

Listed software links into the system's inventory of software in your IT infrastructure. You can view this inventory from the Manufacturers link on the sub-navigation menu.

To access software information, select Software in the tree view.

Managing devices and device attributes

Read the information and procedures in this section to learn about specific device management tasks, including:

- [Clearing heartbeat events](#)
- [Locking device configuration](#)
- [Renaming a device](#)
- [Re-identifying a device](#)
- [Remodeling a device](#)
- [Resetting the device manage IP address](#)
- [Deleting a device](#)
- [Exporting device list to load into another system](#)
- [Batch loading or modifying devices](#)

Clearing heartbeat events

If you have devices configured to send a recurring event that is mapped to a heartbeat class, you can clear stale heartbeat events.

To clear the heartbeat events associated with a device:

1. Log in to the Resource Manager browser interface as a user with ZenManager or Manager privileges.
2. Navigate to ADVANCED > Settings.
3. In the left panel, select EVENTS.
4. At the bottom of the Event Configuration page, click the Clear button in the Clear Event Heartbeats section.
The system displays a brief message banner.

Locking device configuration

You can lock a device's configuration to prevent changes from being overwritten when remodeling the device. Two levels of locking are available. You can lock the configuration from deletion and updates, or solely from deletion.

Note: Device locking prevents changes and deletion due to remodeling. It does not prevent manual changes and deletion.

To edit lock selections for a device configuration:

1. Navigate to the device in the device list.
2. At the bottom of the device overview page, select Locking from the Action menu. The Lock Device dialog box appears.
3. Select the type of lock you want to implement or remove.
4. To send events when actions are blocked by a lock action, select the "Send an event..." option. The lock or unlock action is implemented on the device, and the system displays a confirmation message of the action.

Renaming a device

Because the system uses the manage IP to monitor a device, the device name may be different than its fully-qualified domain name (FQDN). The device name must always be unique in the system.

To rename a device:

1. Navigate to the device in the device list. Click the device name.
2. On the device overview page, edit the Device Name field with the new device name.
3. Click Save. The system renames the device and displays a confirmation message of the action.

Re-identifying a device

Changing the device ID in the system is different from changing the device name. If you change the ID, you must associate existing performance data with the new device ID; otherwise, you lose the data for this device. However, the re-association process takes time, and during processing, metrics are not collected or graphed for this device.

You have the option of deleting existing performance data and starting fresh with the new device ID. Collection of new performance data begins immediately.

1. Navigate to the device in the device list and click the device name.
2. At the bottom of the device overview page, click the Action menu and choose Reidentify Device.
3. In the Reidentify Device dialog box, enter a new ID for the device.
4. Choose whether to re-associate existing performance data for the device or delete it.
5. Click SUBMIT.
On successful completion of the job, collection resumes automatically and no further action is required.
6. Only if the job fails, manually resume data collection for the device as follows:
 - a. Log in to the Resource Manager interface as a user with ZenManager privileges.
 - b. At the bottom of the device overview page, click the Action menu and choose Resume Collection.

Remodeling a device

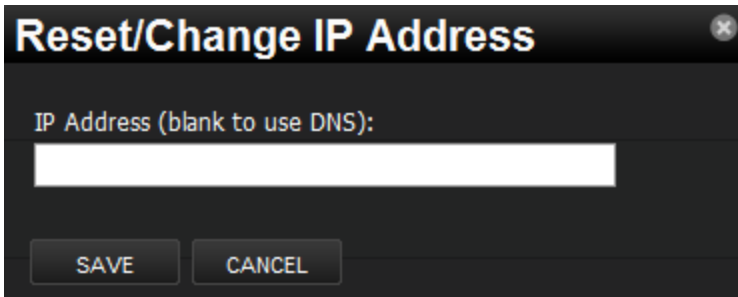
Remodeling forces the system to re-collect all configuration information associated with a device. Normally, the system models devices every 720 minutes; however, if you want to remodel a device immediately, follow these steps:

1. Navigate to the device in the device list and click on the Device name.
2. At the bottom of the Device Overview page, click the Model Device button. The system remodels the device. A dialog box appears that shows progress of the action.

Resetting the device manage IP address

You might reset the manage IP address if the IP address of a device changed and you want to maintain the historical data at the original IP address. To reset the manage IP address of a device:

1. Navigate to the device in the device list.
2. At the bottom of the device overview page, select Reset/Change IP Address from theAction menu.

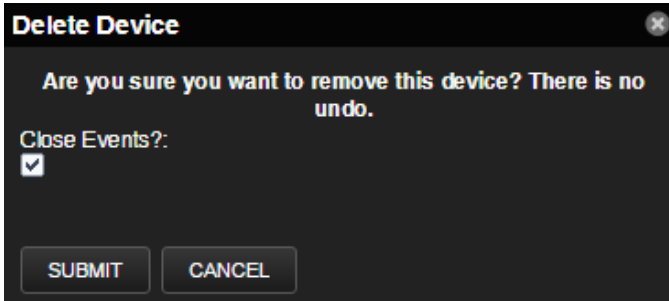
A screenshot of a dark-themed dialog box titled "Reset/Change IP Address" with a close button in the top right corner. Below the title is a text input field with the placeholder text "IP Address (blank to use DNS):". At the bottom of the dialog are two buttons: "SAVE" and "CANCEL".

3. Enter the new IP address for the device, or leave the field blank to allow the IP address to be set by DNS.
4. Click Save. The IP address for the device is reset.

Deleting a device

To delete a device from the system:

1. Navigate to the INFRASTRUCTURE page.
2. Select the device you want to remove from the system by clicking on its row. You can select multiple devices by Ctrl-clicking or Shift-clicking the devices. Be sure to click on the row in an area that is not defined by a link.



3. **Optional:** Change the selection to close current events for the device. By default, event data is removed.
4. Click Submit.

The system removes the devices and associated data (if selected), and displays a confirmation message of the action.

Exporting device list to load into another system

Use the `zenbatchdump` command to write the names of your devices, their device classes, locations, groups, and systems to a text file. You can then use the `zenbatchload` command to import the devices into another system instance.

The following command options might be helpful. For information about all options, run the following command:

```
zenbatchdump --help
```

- `--root=ROOT` - The default path is `/Devices`. Use this option to set the root device path to dump. For example,

```
/Devices/Servers
/Devices/Network/Cisco/Nexus
```

- `--prune` - Specifies whether device classes should only be dumped if they are part of the root path.
- `--noorganizers` - Specifies whether organizers, such as device classes and groups, should be dumped.

1. Log in to the Control Center host as a user with `serviced` CLI privileges.
2. Start a shell in a `zope` service container.

```
serviced service shell zope/0
```

This command mounts `/mnt/pwd` so you can use files that exist on the host.

3. Change to the `zenoss` user.

```
su - zenoss
```

4. Export the device list to the text file.

```
zenbatchdump -o mydevicelist.txt
```

This command writes the names of your devices, their device classes, locations, groups, and systems to a file named `mydevicelist.txt`.

5. If you plan to use the text file in a batch load on another system, copy the file to that host.
6. Exit from the `zenoss` user shell in the container.

```
exit
```

7. Exit from the `Zope` container shell.

```
exit
```

Batch loading or modifying devices

Use the `zenbatchload` command to load devices into the system or modify device properties. The utility creates new device classes automatically, and by default, models each device.

Before you can load or modify devices, you need a text file to use as input. You can use a file created by the batch dump utility (`zenbatchdump`), or create a text file that specifies the following information for each device:

- Resource Manager device class where the device should be loaded, for example, Linux SSH or SNMP, Windows WMI or WinRM, Network, Storage.
- Device name (resolvable hostname for the device or IP address and descriptive name).
- Optional: Provide additional data for device classes or individual devices; for example, login credentials, specific IP address if DNS resolution is not available, and the collector on which to load the device.

Configuration file syntax is as follows.

```
/Devices/Destination/Device/Class
Hostname | IP-Address [option='value' [, option='value']] ...]
Hostname | IP-Address [option='value' [, option='value']] ...]
```

The following command options might be helpful. For information about all options or a sample configuration file, run the following commands:

```
zenbatchload --help

zenbatchload --sample_configs
```

- `--nocommit` - Perform a test run of the batch load process; do not commit changes to the ZODB. If necessary, edit your configuration file.
- `--nomodel` - Do not model the remote devices. You must be able to commit changes to the ZODB.

The following example shows SNMP, SSH, and Windows monitored devices:

```
/Devices/Server zSnmpCommunity='underwriting'

/Devices/Server/Linux
linuxsnmp.hypothetical.loc

/Devices/Server/Microsoft/Windows zWinKDC='winkdc.hypothetical.loc', zWinRMPassword='Zenny456',
zWinRMUser='zenmonitor@hypothetical.loc'
windows.hypothetical.loc

/Devices/Server/SSH/Linux zCommandPassword='Zenny123', zCommandUsername='zenmonitor'
linuxssh.hypothetical.loc

/Devices/Server/Linux/WordPress zDeviceTemplates=['Device', 'Apache']
wordpress.hypothetical.loc
```

For additional ways to add and discover devices, see [Adding and discovering devices](#).

1. Log in to the Control Center host as a user with serviced CLI privileges.
2. Start a shell in a Zope service container.

```
serviced service shell zope
```

This command mounts `/mnt/pwd` so you can use files that exist on the host.

3. Change to the `zenoss` user.

```
su - zenoss
```

4. Load devices and components that are listed in a text file into Resource Manager.

```
zenbatchload mydevicelist.txt
```

5. Exit from the `zenoss` user shell in the container.

```
exit
```

6. Exit from the Zope container shell.

```
exit
```

Adding and discovering devices

Modeling is the process by which Resource Manager:

- Populates the device database
- Collects information about the devices in the system (such as operating system type or file system capacity)

Resource Manager models devices when they are added to the database automatically through the discovery process or manually.

Discovering devices

Using network or IP address range information that you provide, Resource Manager performs processing to discover your devices.

During the discovery process, Resource Manager:

- Iterates through every IP address in the networks and IP ranges that you specify.
- Adds each device that responds to a ping request.
- Adds information to any device that responds to an SNMP, WinRM, or SSH request.
- Based on configuration settings, assigns devices to device classes.

For a discovered device to be assigned to a device class that matches its hardware and operating system, the `zDiscoveryMappingOn` configuration setting on the `/Discovered` device class must be enabled. If the setting is disabled, discovered devices remain in the `/Discovered` device class, and you can manually classify them.

Discovery mapping

Discovery mapping is the process by which Resource Manager automatically assigns a discovered device to a device class based on priority values that you set. Priority values provide the device characteristics that Resource Manager uses to select the device class. Devices are categorized by hardware type or operating system, as follows:

- If the hardware field is set, the device is assigned to the device class that is associated with that hardware type in the manufacturer settings.
- If the operating system field is set, the device is assigned to the device class that is associated with that operating system in the manufacturer settings.
- If a discovered device has both the hardware and operating system fields set, classifies the device based on values of the discovery mapping configuration options. Providers with a lower priority value are tried first. Providers with the same priority value are attempted in an indeterminate order.

To change the discovery mapping configuration options, choose **ADVANCED > Settings > Discovery Mapping**. Enter numerical values for each provider to set the priority. Valid values are 0 or any negative, positive, or decimal numeric value. An empty field defaults to 0.

zenoss
SERVICE DYNAMICS

DASHBOARD EVENTS INFRASTRUCTURE REPORTS ADVANCED

Settings
Commands
Users
ZenPacks
Versions
Events
User Interface
LDAP
Support
Discovery Mapping

Discovery Mapping Configuration

Set priority for discovery mapping providers. Lower priority providers will be tried first. Providers with the same priority will be attempted in an indeterminate order.

N1KDiscoveryMappingProvider: 1

HWKeyMapZcmlDiscoveryMappingProvider: 1

OSKeyMapZcmlDiscoveryMappingProvider: 1

OidMapZcmlDiscoveryMappingProvider: 1

Save Cancel

Discovery Mapping ZenPack Overview

Providing network or IP address range for device discovery

Provide network or IP address range information so that Resource Manager can discover your devices.

1. From the navigation menu, select INFRASTRUCTURE > Devices.
2. Click the Add Devices icon and, from the drop-down list, choose Discover Networks.

Network Discovery

Networks/Range	SNMP	SSH Authentication	Windows Authentication
Enter one or more networks (such as 10.0.0.0/24) or IP ranges (such as 10.0.0.1-50): <input type="text"/>	Community Strings: <input type="text" value="public"/> <input type="text" value="private"/>	Username: <input type="text"/> Password: <input type="text"/>	Administrator Username: <input type="text"/> Password: <input type="text"/>
<input type="button" value="Discover"/>			

3. For each network or IP range in which you want the system to discover devices, enter an address or range. For example, enter a network address in CIDR notation 192.0.2.0/24 or a range of IP addresses 192.0.2.1-50. Note: Trying to add a /16 or /8 network can take a very long time, and might have unintended consequences.
4. For each network or IP range, specify the Windows, SSH, or SNMP credentials that you want Resource Manager to use on the devices that it discovers, and then click Discover.
You can enter only one of each. Resource Manager attempts to use the same credentials on each device that it discovers in the specified networks or IP ranges.
Resource Manager uses Advanced Encryption Standard (AES) with a 256-bit key size to encrypt all passwords, and stores them in the Zope object database.

Discovering devices with the CLI

To discover devices with the command-line interface, follow these steps:

1. Log in to the Control Center master host as root, or as a user with superuser privileges.
2. Attach to the zenhub service as the zenoss user.

```
serviced service attach zenhub su - zenoss
```

3. Run the zendisc command to discover devices.

```
zendisc run --now --monitor localhost --range networkOrIPRange
```

The variable `networkOrIPRange` represents the network or IP range in which you want the system to discover devices. For example, enter a network address in CIDR notation `192.0.2.0/24` or a range of IP addresses `192.0.2.1-50`.

Attempting to add a /16 or /8 network requires considerable time, and might have unintended consequences.

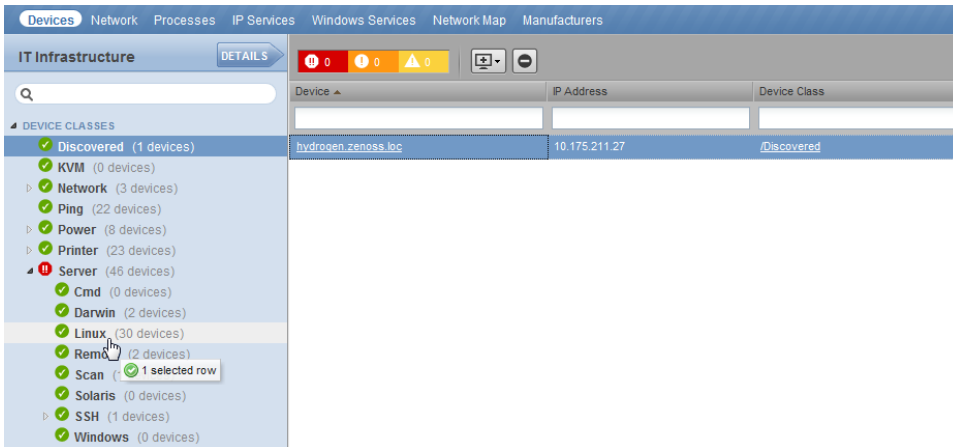
Classifying discovered devices

When discovery is complete, by default, devices are assigned to an appropriate device class in the hierarchy. If necessary, you can change the device class of discovered devices.

If the configuration property `zDiscoveryMappingOn` for the `/Discovered` device class, devices are not moved from the `/Discovered` class. You can manually change the device class of discovered devices.

Servers are organized by operating system. For example, if the system discovers Windows devices, you can relocate them to `/Server/Windows`. If you want to monitor and model using SNMP, classify discovered Linux devices in `/Server/Linux`. If you want to monitor and model using SSH, classify discovered Linux devices in `/Server/SSH/Linux`.

1. From the device list, choose one or more discovered devices (highlight one or more rows).
2. Drag the selected devices to the new device class in the tree view.



3. In the Move Devices dialog box, click OK.

The list of devices refreshes, and the devices appear in the newly selected class.

Updating device authentication details

For each device that is added to the database and set to its proper device class, Resource Manager might require additional or different authentication information before it can gather device information and monitor the device.

For example, before the system can monitor a device in the `/Server/Windows` class, you must supply your Windows user name and password.

Similarly, for a device in the `/Server/SSH/Linux` class, you must supply your SSH user name and password. Set these values in for the device with the `zCommandUsername` and `zCommandPassword` configuration properties.

Note: Resource Manager uses Advanced Encryption Standard (AES) with a 256-bit key size to encrypt all passwords, and stores them in the Zope object database.

1. In the devices list, click a device name. The device summary page appears.
2. In the left panel, choose Configuration Properties.
3. Double-click the `zWinRMUser` configuration property. The Edit ConfigProperty dialog box appears.
4. In the Value field, enter your Windows user name, and then click Submit.
5. Double-click the `zWinRMPassword` configuration property. Enter your Windows password, and then click Submit.
6. After making changes, to ensure that the authentication changes are valid, remodel the device.

Adding or editing information on a device record

1. In the devices list, click a device name. The device overview page appears.
2. Choose values to change, or click "edit" adjacent to a label. Enter or change information, and then click Save.

Adding devices manually

You can manually add one or more devices to the database. When adding devices, you choose the device class and whether Resource Manager models the device immediately.

- [Adding a single device](#)
- [Adding multiple devices](#)
- [Adding a Cisco UCS device](#)
- [Add VMware vSphere endpoint](#)

Adding a single device

When you manually add a device, information that you provide might conflict with information that the system discovers about the device. Therefore, in most cases, you can add a device by providing only the following required information.

- Hostname or IP address - Enter the network (DNS) name or IP address of the device.
- Device Class - Select a device class to which this device will belong. For example, if the new device is a Linux server, then select /Server/Linux.
- Collector - By default, this is localhost. Select a collector for the device.

The exception is adding a Cisco router in a device class other than /Network. Before adding the device, set the value of the zlfDescription configuration property to True. By default, this option is set to True for the /Network class.

1. From the navigation menu, select INFRASTRUCTURE > DEVICES.
2. Click the Add Devices icon and choose Add a Single Device.

The screenshot shows a dark-themed dialog box titled "Add a Single Device". It contains the following fields and controls:

- Hostname or IP address:** Text input field containing "dns.testserver.loc".
- Device Class:** Dropdown menu with "/Server/Linux" selected.
- Collector:** Dropdown menu with "localhost" selected.
- Title:** Text input field containing "DNS server 1".
- Production State:** Dropdown menu with "Production" selected.
- Device Priority:** Dropdown menu with "Normal" selected.
- Model Device:** A checked checkbox.
- More...:** A link below the checkbox.
- Buttons:** "ADD" and "CANCEL" buttons at the bottom.

3. Enter information or make selections to add the device.
By default, Model Device is selected. If you do not want the device to be modeled when it is added, clear the check box for this option.
4. **Optional:** To display additional fields that are specific to the chosen device class, click More. For example, on the expanded page, you can
 - Enter device-specific details.
 - Edit SNMP settings.
 - Set hardware and operating system information.
 - Add device comments.
5. Click ADD.
6. **Optional:** To view the Add Device job in progress, click View Job Log in the notification that appears when you add the device.

When the job completes, the device is added in the selected device class.

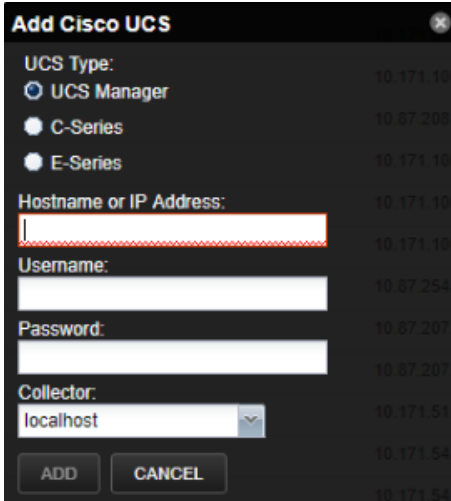
Adding multiple devices

1. From the Navigation menu, select INFRASTRUCTURE. The Devices page appears.
2. Click the Add Devices icon and select Add Multiple Devices. The Add Infrastructure page appears.
3. Select the category, type, and connection information for each device that you want to add, and then click Add.
4. After adding devices, click Done.

Adding a Cisco UCS device

The Cisco Unified Computing System (UCS) is a data center server computer product line that is composed of computing hardware, virtualization support, switching fabric, and management software. You can add a UCS Manager or a Cisco Integrated Management Controller (CIMC), which manages and monitors C-Series and E-Series rack servers, to the applicable /CiscoUCS device class.

1. From the Navigation menu, select INFRASTRUCTURE > Devices.
2. From the Add Devices icon, choose Add Cisco UCS.



3. Choose the UCS device type:
 - UCS Manager
 - C-Series
 - E-Series
4. Complete the dialog box with the hostname or IP address, the user name and password, and select the collector for this device, and then click Add.
5. **Optional:** To view the Add Device job in progress, click View Job Log in the notification that appears.

Add VMware vSphere endpoint

Follow these steps to add a vSphere endpoint:

1. From the Navigation menu, select INFRASTRUCTURE > DEVICES.
2. From the Add Devices icon, select Add VMware vSphere Endpoint.
The Add VMware vSphere Endpoint page appears with the Add Single Device tab selected.

The screenshot shows a dialog box titled "Add VMware vSphere Endpoint" with a close button in the top right corner. It has two tabs: "Add Single Device" (selected) and "Add Multiple Devices". The form contains the following fields: "Name:" (empty text box with a red dashed border), "Hostname or IP:" (empty text box), "Username:" (empty text box), "Password:" (empty text box), "Port:" (text box containing "443"), "SSL:" (checkbox checked), and "Collector:" (dropdown menu showing "localhost"). At the bottom are "ADD" and "CANCEL" buttons.

3. If you want to just add a single device, complete the dialog with the appropriate information, then click Add.
4. If you want to add multiple devices at once, click the Add Multiple Devices tab.

The screenshot shows the same dialog box but with the "Add Multiple Devices" tab selected. It features a table with one row of input fields: "Name", "Hostname or", "Username", "Password", "443", a checked checkbox, "localhost" (dropdown), "Add", and "Remove". "ADD" and "CANCEL" buttons are at the bottom right.

5. Fill out the information in the row for a device and click Add.
A new row appears and you can add another device. When you are finished be sure to click the Add button at the bottom of the dialog to add all of your devices.

Basic monitoring

- [Availability monitoring](#)
- [Monitoring using ZenCommand](#)
- [SNMP monitoring](#)
- [Monitoring devices remotely through SSH](#)
- [The network map page](#)

Availability monitoring

The availability monitoring system provides active testing of the IT infrastructure, including:

- Devices
- Network
- Processes
- Services

Availability monitoring is facilitated by:

- **Zenping** - The system's Layer-3 aware, topology-monitoring service. Zenping performs high-performance, asynchronous testing of ICMP status. The most important element of this service is that Resource Manager has built a complete model of your routing system. If there are gaps in the routing model, the power of Zenping's topology monitoring will not be available. If there are gaps, this issue can be seen in the zenping.log file. Zenping uses Nmap to build a ping tree and perform Layer 3 suppression. The service runs every 30 minutes, which is configurable.
- **Zenstatus** - Performs active TCP connection testing of remote daemons.

zenping correlation

By default, zenping attempts to correlate all zenping instances on the same collector host. If zenping instances are on different physical hosts, they must share the same redis instance. In this case, the setting `--redis-url` which controls what redis instance zenping is talking to needs to be set to point to the same redis instance.

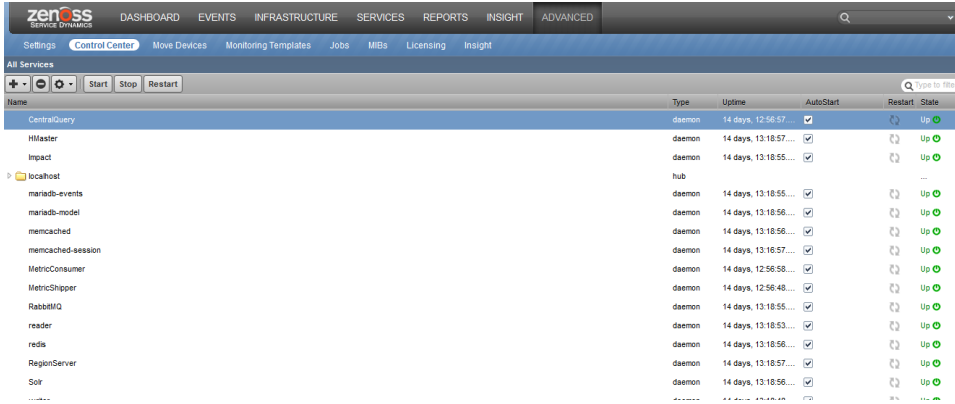
There are other correlation settings to use on the command line that you need to be aware of:

- `--disable-correlator`: This will set the ping correlation value to `False` and will turn off correlation.
- `--correlation-backend`: Default value is `distributed`, which is the default behavior when using the redis instance. If you set this to `simple`, no communication with redis is available and correlation will only be performed within the same zenping instance.
- `--connected-ip-suppress`: Defaults to disabled. This is used when zenping does not or cannot monitor all the interfaces of routers along a traceroute path. In this case, zenping can use the model information to infer if an IP along the traceroute path belongs to a device (like a router). It will use this information to determine whether the root cause for a ping down is the "connected" device.
- `--delay-count`: Defaults to 0. This means that as soon as a ping down is detected an event is sent. When using distributed ping correlation, zenping instances are not in sync. `--delay-count` can be set to 1 or higher in order to allow zenping instances to share the ping state before sending events. It can also be used to avoid sending events after N ping downs. For example, you may only want to send an event after 3 ping downs.


Controlling ping cycle time

Follow these steps to modify the ping cycle time:

1. In the Resource Manager browser interface, open **ADVANCED > Control Center**.



2. In the All Services area, select the collector to modify. Each collector has a zenping service, but the configuration option is set on the collector level, rather than the service level.



3. In the Display area, select Details.



4. Set the Ping Cycle Interval (seconds) to the desired value.

Ping Cycle Interval (seconds):

Modeler Cycle Interval (minutes):

Discovery Networks:

5. Click Save

Using the predefined /Ping device class

The /Ping device class is a configuration for devices that you want to monitor only for availability. The system does not gather performance data for devices placed in this class. You can use the /Ping device class as a reference for your own configuration; or, if you have a device that you want to monitor solely for availability, you can place it under this class.

Monitoring processes

Resource Manager provides process availability monitoring for hosts that support SNMP or SSH access. Process monitoring features include:

- Process classes, defined by Python regular expressions. Classes may generate one or more process sets, each containing one or more process instances.
- Process sets may include process instances running on multiple hosts. This captures related or redundant processes, enabling a more holistic view of data center services.
- Process set names, to replace the often-cryptic names of process instances with descriptive labels.
- Process set locking, to maintain continuity of data collection if the members of a given process set are not running during modeling.
- A testing dialog, to discover and refine the sets a class generates.

Use the INFRASTRUCTURE > PROCESSES page to create and manage process classes and process sets.

The tree view shows process class organizers (at the top) and the list of process classes in each organizer (the rest of the view). You may filter the list with the active search field, at the top of the list.

Example: Creating a process class

This section provides an example of using process availability monitoring to create a new process class, for database processes. The database runs on a Linux host, and the following output is a partial list of the results of the `ps axho args` command on the database host. Process monitoring uses the output of that command (or its equivalent) as input for regular expression matching.

```
ora_pmon_orcl ora_psp0_orcl ora_vktm_orcl ora_gen0_orcl ora_diag_orcl ora_dbrm_orcl ora_dia0_orcl ora_mman_orcl  
ora_dbw0_orcl ora_lgwr_orcl ora_ckpt_orcl ora_smon_orcl ora_reco_orcl ora_mmon_orcl ora_mmln_orcl ora_d000_orcl  
ora_s000_orcl ora_s001_orcl ora_s002_orcl ora_s003_orcl ora_s004_orcl ora_s005_orcl ora_s006_orcl ora_s007_orcl  
ora_s008_orcl ora_s009_orcl ora_p000_orcl ora_p001_orcl ora_p002_orcl ora_p003_orcl ora_p004_orcl ora_qmnc_orcl  
ora_n000_orcl ora_l000_orcl ora_l001_orcl ora_l002_orcl ora_l003_orcl
```

The following subsections provide procedures for creating a process class that captures a selection of the preceding process instances in process sets.

Test existing process classes

The existing process classes may already capture the process sets you want. Follow these steps to test the existing process classes.

1. Log in to the Resource Manager browser interface, and then navigate to INFRASTRUCTURE > PROCESSES.
2. In the lower-left corner of the tree view, click the Action menu, and select Test All Process Classes Regular Expressions.
The top-right portion of the dialog box displays all of the process classes, in the order in which they are evaluated.
3. Select the process names in the previous section, and copy them into a paste buffer.
4. In the Test Process Class Regular Expressions dialog, select all of the existing text in the Input area, and then paste the process names from the buffer.
Alternatively, you may paste the process names into an empty file, save the file on the system from which your browser is launched, and then use the Choose File button to insert the process names.
5. At the bottom-left corner of the dialog, click Test.
6. If any process sets are created, they are displayed in the list area, above the Test button.

Create an organizer

Complete the previous section, and then follow these steps to create a process class organizer.

1. Log in to the Resource Manager browser interface.
2. Navigate to INFRASTRUCTURE > PROCESSES.
3. In the lower-left corner of the tree view, click the Add menu, and select Add Process Class Organizer.
4. In the Add Process Class Organizer dialog, enter Database, and then click Submit.

Create a process class

Complete the previous section, and then follow these steps to create a process class.

1. At the top of the tree view, double-click Processes, the root organizer to open it.
2. Select Database.
3. In the lower-left corner of the tree view, click the Add menu, and then select Add Process Class.
4. In the Add Process Class dialog, enter DB daemons, and then click Submit.
5. At the top of the tree view, double-click Processes to open it, and then select Database.

Define the regular expression series of a process class

Complete the previous section, and then follow these steps to create the series of regular expressions that define a process class, and generate process sets.

1. In the list area of the tree view, select DB daemons.
2. In the Description field, enter Database daemons.
3. In the Include processes like field, replace DB daemons with a that selects the database processes. For example, `ora_[^_]{4}_orcl`. The example regular expression selects all of the processes in the sample process instance list.
4. In the Exclude processes like field, enter a regular expression to remove processes from the results of the preceding regular expression. The default entry excludes common user commands. The default entry does not exclude any of the processes in the sample process instance list.
5. The next two fields, Replace command line text and With, work together to simplify the names of process sets.
 - In the Replace command line text field, enter a regular expression containing one or more pattern groups.
 - In the With field, enter replacement text, along with the sequence number of one or more of the pattern groups defined in the previous field.

For example, to create four pattern groups for database processes, enter the following regular expression in the Replace command line text field:

```
^(ora_)([a-z])(.{4})(orcl)
```

To use two of the pattern groups in the replacement text, enter the following text in the With field:

```
DB [\4] daemons starting with [\2]
```

Each unique replacement generated by the combination of the text plus the inserted pattern sequences becomes a process set. In the case of this example, pattern group 4 does not vary, but pattern group 2 does. So the number of process sets generated by this class will equal the number of unique alphabetic characters found in the first position after the first underscore.

6. Click Save.

Test a process class

Complete the previous section, and then follow these steps to test a single process class.

1. In the lower-left corner of the tree view, click the Action menu, and select Test Process Class Regular Expressions. The top-right portion of the dialog box displays the regular expression series that defines this process class.
2. Select the process names in the previous section, and copy them into a paste buffer.
3. In the Input area, select all of the existing text, and then paste the process names from the buffer.
4. At the bottom-left corner of the dialog, click Test.
5. The Output area displays each individual match, along with the count of unique process sets. You may refine the regular expressions and retest as often as you like.
6. Click Done. Changes made to regular expressions in this dialog are copied to the process class definition page. However, the changes are not saved until you click the Save button on that page.

Test and review the process class sequence

The order in which process classes are evaluated is significant. During modeling, each time a process matches a class, the matching process is put into a process set, and then removed from the list of processes that are passed on to the next class in the sequence. New process classes are inserted into the process class sequence automatically, and may not be in the appropriate position in the sequence.

Complete the previous section, and then follow these steps to test and review the process class sequence.

1. In the lower-left corner of the tree view, click the Action menu, and select Test All Process Classes Regular Expressions.
2. Select the process names in the previous section, and copy them into a paste buffer.
3. In the Input area, select all of the existing text, and then paste the process names from the buffer.
4. At the bottom-left corner of the dialog, click Test.
5. The number of processes matched and process sets created in this test should be identical to the results of testing the process class alone. If they are not, follow these steps to adjust the process class sequence.
 - a. In the Test Process Class Regular Expressions dialog, click Done.
 - b. From the Action menu, select Change Sequence.
 - c. Scroll through the list of process classes, and then select the class to move.
 - d. Drag the class to an earlier (higher) location in the sequence.
 - e. Click Submit.
6. Re-open the Test Process Class Regular Expressions dialog, and re-test the sequence.

Test the process class on a host

Process sets are created during modeling. To test a process class, choose a device host that is configured for SNMP or SSH access, and model it manually.

Complete the previous section, and then follow these steps to test the process class on a host.

1. Navigate to **INFRASTRUCTURE > DEVICES**.
2. Select a host that is configured for SNMP or SSH access, and is running process that match the new class.
For example, the list of processes used in this section is collected from a VirtualBox virtual appliance downloaded from the [Oracle Technology Network](#).
3. Open the host's **Overview** page. From the Action menu, select **Model Device**. When modeling completes, the **OS Processes** section of the tree view is updated to include the new process sets.
4. Navigate to **INFRASTRUCTURE > PROCESSES**.
5. In the tree view, select the **DB daemons** class.
The process sets found on the host are displayed in the list at the bottom of the page.

Process class options

The process class page includes the options described in the following sections.

Process Count Threshold

You may set minimum and maximum values for the number of process instances included in a process set. The threshold values apply to all of the process sets generated by a class. The minimum and maximum values are inclusive. That is, if the minimum is 3 and the maximum is 5, then 3, 4, and 5 are all valid process instance counts.

You may define a threshold as an exclusive range. If the minimum is 5 and the maximum is 3, then 4 is an invalid process instance count.

Monitoring Options

- **Enable Monitoring (zMonitor)**

To disable monitoring for all process sets generated by this class, set the local value to No.

- **Send Event on Restart (zAlertOnRestart)**

To send an event when monitoring restarts, set the local value to Yes.

- **Failure Event Severity (zFailSeverity)**

To specify a non-default event severity for the failure of process sets generated by this class, set a local value.

Process Set Locking

Process sets are generated at modeling time. Since modeling recurs regularly, a given modeling run may result in a missing process set, due to a transient absence of one or more process instances. To prevent this from happening, set the local value of the **Lock Process Components? (zModelerLock)** field to one of the following options.

- **Lock from Deletes**

Prevent deletion of process sets generated by this process class if modeling returns empty sets.

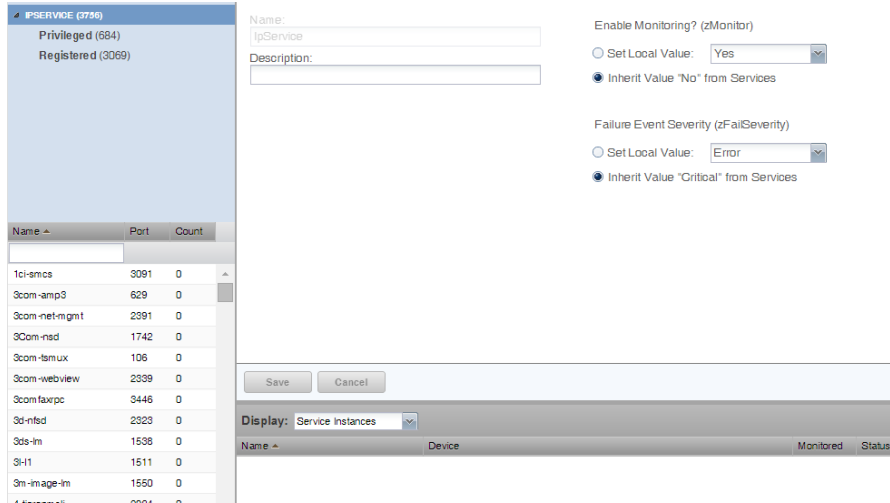
- **Lock from Updates**

Prevent updates to process sets generated by this process class if modeling returns new process sets.

The final option, **Send an event when action is blocked? (zSendEventWhenBlockedFlag)**, is used only when a process set is locked. If you lock the process sets of a class, you may set the local value of this field to Yes, and an event will be created when a process set would have been either deleted or updated during modeling.

Monitoring IP services

The INFRASTRUCTURE > IP Services page lets you manage and monitor IP services that are running on your network. Monitored IP services are tested with a simple TCP client to ensure they are listening and responding appropriately. IP service monitoring also has the ability to send a clear-text command to a TCP port (similar to doing the same in telnet manually) and evaluate the response with a regular expression pattern.



The tree view lists all monitored IP services. Filter this list by using the active search area at the top of the view.

The details area shows:

- Service class description
- TCP port
- Associated service keys

To add or change details for a service class, enter or change information, and then click **Save**.

The lower section of the page lists currently running services in this class (by device), and shows their monitoring status. You can also display Configuration Properties by selecting that from the drop-down list.

Enabling IP service monitoring

You can choose to monitor individual services or service classes.

When monitoring a service class, you can choose not to monitor one or more individual services in the class. For example, the SMTP service class is monitored by default, but may not be a critical service on some devices. In this case, you can disable its monitoring on those devices.

Note: If a service is configured to listen only on local host (127.0.0.1), then it is not monitored by default.

Note: When adding a new IP service that uses a port value higher than 1024, you need to increase the value of zIpServiceMapMaxPort to a number higher than the port you are monitoring.

To enable monitoring for a service class or service:

1. In the tree view, select the service class or service to monitor.
2. Make one or more selections:
 - **Enable Monitoring (zMonitor)** - By default, Inherit Value is selected for all services below the IPService node. When selected, the service class or service will inherit monitoring choices from its parent. If you want to individually enable monitoring choices, select the Set Local Value option, and then select a value.
 - **Failure Event Severity (zFailSeverity)** - By default, Inherit Value is selected for all services below the IPService node. When selected, the service class or service will inherit severity level choices from its parent. If you want to individually select severity levels, select the Set Local Value option, and then select a value.
3. Click Save to save your choices.

Using the predefined /Server/Scan device class


The predefined /Server/Scan device class is an example configuration for monitoring TCP services on devices using a port scan. If you have a device that you want to monitor for service availability alone, you can place it under this device class. The system will not collect performance data for devices in this class.

Monitoring Windows Services

The Windows Services page (INFRASTRUCTURE > WINDOWS SERVICES) uses the WinService monitoring template, which is included in the [Zenoss.zenpack.Microsoft.Windows](#) ZenPack.

Monitoring using ZenCommand

Resource Manager has the ability to run Nagios® and Cacti plug-ins through the ZenCommand process. ZenCommand can run plugins locally and remotely by using a native SSH transport. When run, the system tracks the return code of each plug-in and then creates events with plug-in output. Additionally, it can track performance information from a plug-in.



Plugin format for ZenCommands

Nagios® plugins are configured by using a command template. A template named "Device" will bind to all devices below the template definition. Within each template is a list of commands that will run. The commands can be any program that follows the Nagios® plug-in standard. Inputs are command line arguments; output is the first line of stdout, plus a return code.

Note: Resource Manager return codes differ from Nagios® return codes, as follows:

Value	Resource Manager	Nagios
0	Clear	OK
1	Data Source	WARNING
2	Data Source+1	CRITICAL
3	Data Source	UNKNOWN

For comprehensive information about Nagios® plugins, refer to the [Nagios® Plugins Development Guidelines](#).

A Nagios® command has several fields:

- name - Specifies the name of the command object.
- enabled - Indicates whether this command should be used on a given device.
- component - Specifies the component name to use when zencommand sends events to the system.
- event class - Specifies the event class to use when sending events to the system.
- severity - Sets the default severity to use when sending events to the system.
- cycle time - Sets the frequency a command should be run (in seconds).
- command template - Specifies the command to run.

The command template string is built by using Zope TALEX expressions. Several variables are passed when evaluating the template. They are:

- zCommandPath - Path to the zencommand plug-ins on a given box it comes from the configuration property zCommandPath. zCommandPath is automatically added to a command if a path is absent from the beginning of the command.
- devname - Device name of the device against which the command is being evaluated.
- dev - Device object against which the command is being evaluated.
- here - Context of evaluation. For a device, this is equivalent to dev for a component (such as a file system or interface). This is the component object.
- compname - If this command evaluates against a component, specifies its name as a string.
- now - Current time.

Template values are accessed like shell variables.

Testing ZenCommands

You can test ZenCommand data sources by using the `zentestcommand` shell script.

1. Log in to the Control Center host as a user with serviced CLI privileges.
2. Attach to the zenhub service as the zenoss user.

```
serviced service attach zenhub su - zenoss
```

3. Run the `zentestcommand` script.
Replace `DeviceID` with the ID of the device on which you want to run the command, and `DataSourceName` with the name of a data source on a template associated with the device.

```
zentestcommand -d DeviceID --datasource=DataSourceName
```

The `zentestcommand` script prints the results of the command to standard output.

SNMP monitoring

An object identifier (OID) represents the data points where the data for the graphs comes from. Sometimes the reason that a graph is not appearing is because the OID for the particular graph is not valid for the device. You can test this validity using the command line to see if you can return a value. To test the validity of an OID data point giving performance data:

1. Log in to the Control Center host as a user with serviced CLI privileges.
2. Attach to the Zenoss.resmgr service as the zenoss user.

```
serviced service attach Zenoss.resmgr su - zenoss
```

3. Run the snmpget command for one of the OIDs.
For example, for a server named SERVER, enter the following command:

```
snmpget -v 2c -cpublic SERVER .1.3.6.1.4.1.2021.4.14.0
```

If the OID is valid, the snmpget command returns a value.

Basic SNMP commands

Here are some basic SNMP commands to gather certain information.

- Walk a basic system MIB.

```
snmpwalk -v 2c -cpublic <device_name_or_ip_address> system
```

- Walk an interface description.

```
snmpwalk -v 2c -cpublic <device_name_or_ip_address> ifDescr
```

- Get a single value.

```
snmpget -v 2c -cpublic <device_name_or_ip_address> ifDescr.2
```

- Detailed description of an OID value.

```
snmptranslate -Td RFC1213-MIB::ifDescr
```

- Convert a name to a raw OID.

```
snmptranslate -On RFC1213-MIB::ifDescr
```

- Convert a raw OID to a short name.

```
snmptranslate -OS .1.3.6.1.2.1.2.2.1.2
```

Monitoring devices remotely through SSH

You can monitor devices remotely through SSH. Follow the steps in the following sections to set up remote monitoring.

Changing Resource Manager to monitor devices remotely using SSH

You must edit system properties for the group where you want to collect remote information using SSH.

1. Navigate to the device class path that you want to monitor remotely. You can apply this monitoring for a device or a device class path.
2. Change the configuration properties value for the group. After selecting the device class, click Details, and then select Configuration Properties.
3. On the Configuration Properties page, change the properties that are listed in the following table.
The table includes sample values set up for remote devices. These have a pre-shared key (with no password) set up from the collector to the remote boxes. It also can use password authorization if the password is entered into zCommandPassword.

Configuration properties	Sample value
zCollectorPlugins	snmp portscan
zCommandPassword	The SSH password for the remote machine
zCommandPath	The path to zenplugin.py
zCommandUsername	The SSH user name for the remote machine
zSnmpMonitorIgnore	True

4. Two passes are required for full modeling. The first pass obtains the platform type (so that the system knows which plugins to run). The second pass provides detailed data on interfaces and file systems.
 - a. Log in to the Control Center master host as a user with serviced CLI privileges.
 - b. Display the list of zenmodeler services.

```
serviced service list zenmodeler
```

On a system with multiple collectors, the result is similar to the following example:

```
Name           ServiceID           DepID/Path
zenmodeler     7itut0ryz759ua77ntrm3hi8w  1/Zenoss.resmgr/Zenoss/Collection/localhost
/localhost/zenmodeler
zenmodeler     e3bpfy6j6pyl81346xq446myk  1/Zenoss.resmgr/Zenoss/Collection/localhost
/collectorPool2/zenmodeler
zenmodeler     7dnmgcwexlqxjqko6nja0942y  1/Zenoss.resmgr/Zenoss/Collection/localhost
/collectorPool3/zenmodeler
```

- c. Select the zenmodeler service that is associated with the device to model, and then attach to it as the zenoss user.
Replace ServiceID with the container ID of a zenmodeler service. For example, 7itut0ryz759ua77ntrm3hi8w.

```
serviced service attach ServiceID su - zenoss
```

- d. Run the zenmodeler command. Replace DeviceName with the fully qualified device name.

```
zenmodeler run -d DeviceName
```

- e. Repeat the zenmodeler command to employ the plugins the command gathered on the first pass.

Using the predefined /Server/Command device class

The /Server/Command device class is an example configuration for modeling and monitoring devices using SSH. The configuration properties have been modified (as described in the previous sections), and device, file system, and Ethernet interface templates that gather data over SSH have been created.

You can use this device class as a reference for your own configuration; or, if you have a device that needs to be modeled or monitored via SSH /Command, you can place it under this device class to use the pre-configured templates and configuration properties. You must set the zCommandUsername and zCommandPassword properties to the appropriate SSH login information for each device.

Changing Resource Manager to monitor devices remotely using SSH

You must edit system properties for the group where you want to collect remote information using SSH.

1. Navigate to the device class path that you want to monitor remotely. You can apply this monitoring for a device or a device class path.
2. Change the configuration properties value for the group. After selecting the device class, click Details, and then select Configuration Properties.
3. On the Configuration Properties page, change the properties that are listed in the following table.

The table includes sample values set up for remote devices. These have a pre-shared key (with no password) set up from the collector to the remote boxes. It also can use password authorization if the password is entered into zCommandPassword.

Configuration properties	Sample value
zCollectorPlugins	snmp portscan
zCommandPassword	The SSH password for the remote machine.
zCommandPath	The path to zenplugin.py
zCommandUsername	The SSH user name for the remote machine.
zSnmpMonitorIgnore	True

4. Two passes are required for full modeling. The first pass obtains the platform type (so that the system knows which plugins to run). The second pass provides detailed data on interfaces and file systems.

Using the predefined /Server/Cmd device class

The /Server/Cmd device class is an example configuration for modeling and monitoring devices using SSH. The configuration properties have been modified (as described in the previous sections), and device, file system, and Ethernet interface templates that gather data over SSH have been created.

You can use this device class as a reference for your own configuration; or, if you have a device that needs to be modeled or monitored via SSH /Command, you can place it under this device class to use the pre-configured templates and configuration properties. You must set the zCommandUsername and zCommandPassword properties to the appropriate SSH login information for each device.

The network map page

The network map represents your network's layer 3 topology. From the map, you can quickly determine the status of each device by its highlighted color. Adobe Flash is required to view the map.

To access the network map, select **INFRASTRUCTURE > Network Map**.

The screenshot displays the Network Map interface. At the top, there is a navigation bar with tabs for Devices, Networks, Processes, IP Services, Windows Services, Network Map (selected), and Manufacturers. On the left, a configuration panel includes a text field for 'Root device or component' with the value '10.87.254.1', a 'Maximum hops from root' slider set to 3, and checkboxes for 'Show MAC addresses' and 'Show dangling connectors'. Below this is a 'Layers' section with a tree view showing 'VLANs' expanded to show 'layer3' and 'layer2', both checked. An 'Apply' button is at the bottom of this panel. The main area shows a network topology diagram with nodes and connecting lines. One node is highlighted with a purple circle. On the right, a legend shows status indicators: Critical (red), Error (orange), Warning (yellow), Info (green), Debug (light green), and Map root (purple). A 'Scale: 100%' and 'Center' button are also present. At the bottom right, a status bar shows '1 Job (0 Pending)'.

Adjusting the network map

Use the Repulsion slider to expand or contract the icons that appear on the map. Move the slider right to expand the icons, or left to contract them.

Select the Fit to Window option to bring all displayed icons into the viewable area.

Adjusting viewable hops

You can adjust the number of hops that appear on the network map. Use the Number of Hops slider, which adjusts the number of hops from 1 to 4.

Choosing the network to display

Users with Manager or ZenManager privilege can configure a default network to display in the map. The setting can be different for each Resource Manager user.

For more information, see [Managing users in Resource Manager](#).

Filtering by device type

You can filter the devices that appear on the network map. To do this, select a filter from the Device Class Filter list of options. For example, to show only Linux devices on the map, select `/Server/Linux` from the list of options, then click Refresh.

Loading link data

To load link data for a node:

1. Double-click the node on the map to focus on it, or enter the device name or IP address in the Selected Device or Network field.
2. Select the number of hops to download and display by sliding the counter.
3. Click Refresh.

Viewing device and network details

Double-click a device or network icon in the map to focus on it. Focusing on a node centers it on the map and shows links from the node, based on the number of hops selected.

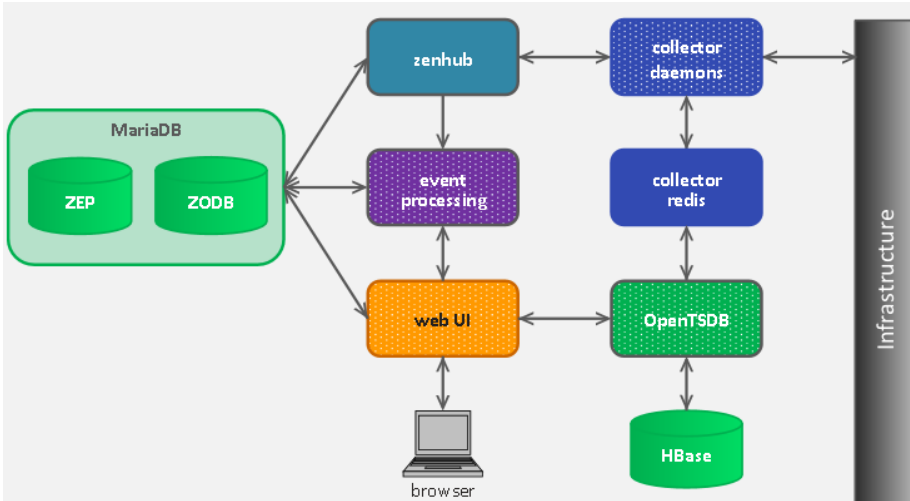
Alternatively, you can type the name or IP address of a device or network in the Selected Device or Network field, and then click Refresh to focus on that node.

To see detailed information about a device or network, select it in the map, and then click Go to Status Page.

When you select a node, the network map displays only the links that are currently loaded into the map. It does not download and display new link data.

Performance monitoring

Resource Manager stores device and daemon performance metrics directly in OpenTSDB, a time series database that runs on top of an HBase instance. Writing directly to OpenTSDB eliminates the need for RRD files to be stored on the collectors. The following image shows how the collector daemons fit into the data collection portion of the Resource Manager architecture.



Resource Manager uses the following methods to monitor performance metrics of devices and device components:

- **SNMP** - Collects data through SNMP from any device that is correctly configured for SNMP monitoring.
- **Microsoft.Windows** - ZenPack that allows performance monitoring of Windows servers.
- **ZenCommand** - Using telnet or ssh, logs in to devices and runs scripts to collect performance data.
- **Other ZenPacks** - Collect additional performance data.

Regardless of the monitoring method, Resource Manager stores performance monitoring configuration information in monitoring templates.

Monitoring templates and performance data

Monitoring templates determine how the system collects performance data for devices and device components. You can define monitoring templates for device classes and individual devices.

Templates comprise the following types of objects:

- **Data sources** specify the exact data points to collect and the collection method to use.
- **Thresholds** define expected bounds for collected data, and specify events to be created if the data does not match those bounds.
- **Graph definitions** describe how to graph the collected data on the device or device components.

Before the system can collect performance data for a device or component, it must use the template binding process to determine which monitoring templates apply.

To view monitoring templates, from the main navigation menu, choose **ADVANCED > Monitoring Templates**.

The screenshot displays the 'Monitoring Templates' configuration page. On the left is a navigation tree with categories like 'Active Directory', 'Apache', and 'Server'. The main area is split into two panels: 'Data Sources' and 'Thresholds'. The 'Data Sources' panel contains a table with one entry for 'apache'. The 'Thresholds' panel contains a table with one entry for 'CPU over 90 percent'. At the bottom right, the 'Graph Definitions' panel lists several graph types for Apache.

Data Sources			
Data Points by Data Source	Source	Enabled	T
apache	\$(dev/managerp)	true	A

Thresholds	
Name	
CPU over 90 percent	

Graph Definitions	
Name	
Apache - Requests	
Apache - Throughput	
Apache - CPU Utilization	
Apache - Slot Usage	

Template names

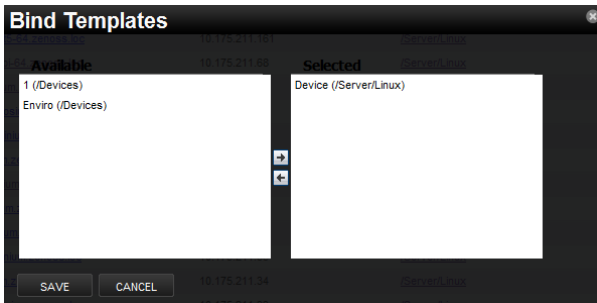
Resource Manager determines the list of template names that apply to a device or component. For device components, the list is defined by the meta type of the component (for example, FileSystem, CPU, or HardDisk). For devices, the list is defined by the zDeviceTemplates configuration property.

After defining the list, Resource Manager locates templates that match the names on the list. For each name, it searches the device and then searches the device class hierarchy. Resource Manager uses the lowest template in the hierarchy that it can locate with the correct name, ignoring others of the same name that might exist further up the device class hierarchy.

Editing templates bound to a device

Select the templates that are bound to a device.

1. From the main navigation menu, choose INFRASTRUCTURE.
2. From the device list in the left pane, choose a device.
3. Below the device list, click the context-sensitive actions menu and choose Bind Templates.
4. From the list of available templates, move one or more templates to the selected list.



5. Click Save.

Data sources

Data sources specify which data points the system collects and how to collect them. Each monitoring template comprises one or more data sources. The system provides the following built-in data source types. ZenPacks provide other data source types.

- **SNMP** - Define data to be collected via SNMP by the ZenPerfSNMP daemon. You specify which SNMP OID to collect. (Many OIDs must end in .0.) Because SNMP data sources specify only one performance metric, they contain a single data point.
- **Command** - Specify data to be collected by a shell command that is executed on the Resource Manager server or on a monitored device. The ZenCommand daemon processes COMMAND data sources. A COMMAND data source can return one or more performance metrics, and usually has one data point for each metric.

Shell commands that are used with COMMAND data sources must return data that conforms to the Nagios® plug-in output specification.

Adding a data source to a monitoring template

1. From the navigation menu, choose ADVANCED > MONITORING TEMPLATES.
2. In the tree view, select the monitoring template.
3. In the Data Sources area, click Add.
4. In the Add Data Source dialog box, enter a name for the data source, select the type, and then click Submit.
5. Double-click the new data source.

An Edit Data Source dialog box appears, specific to the data source type. For example, the following figure shows the Edit Data Source dialog box for a COMMAND data source.

Edit Data Source

Name:

Type: COMMAND

Severity:

Event Class:

Cycle Time (seconds):

Parser:

Use SSH

Component:

Event Key:

Command Template:

Test Against a Device

Device Name:

6. Enter or select values to define the data source. For example, the COMMAND data source type, Use SSH must be enabled. Otherwise, the commands only run locally on the collector that is assigned to the device to which the monitoring template is bound. In the Command Template field, enter the command to run. Note that any script that you enter in this field is first run through TALES running. You might need to escape certain characters. For more information, see [TALES expressions](#).

Data points

Data sources can return data for one or more performance metrics. Each metric that a data source retrieves is represented by a data point.

You can define data points for data sources with all source types except SNMP and VMware. Because these data source types each rely on a single data point for performance metrics, additional data point definition is unnecessary.

1. From the navigation menu, choose **ADVANCED > MONITORING TEMPLATES**.
2. In the Data Sources area, highlight the row that contains the datasource.
3. From the action menu, choose **Add Data Point**.
4. In the Add Data Point dialog box, enter a name for the data point, and then click **Submit**.

Note: For **COMMAND** data points, enter the name that the shell command uses when returning data.

5. Double-click the new data point and enter information or make selections to define it:
 - **Name** - Displays the name that you entered in the Add a New Data Point dialog box.
 - **RRD Type** - Review the following considerations, and then specify the data source type to use for storing data for this data point. Note: The following considerations apply to **COUNTER** and **DERIVE** types:
 - In previous releases, the system saved these rate data points by using their raw counters. Beginning with release 6.1.x, the system stores the actual rate value that is calculated at the collector daemon as it collects data. If a graph or API request for a **COUNTER** or **DERIVE** data point spans the date of the upgrade to 6.1.x, the system automatically queries the data correctly.

In graph legends, the minimum and maximum values are the actual calculated rates for the time period covered by the graph. However, in the graph, as you zoom out, increasing the time frame, minimum and maximum values are averaged; thus values might be different.


 - Instead of **COUNTER**, you could choose **DERIVE** and specify a minimum of 0. This approach creates the same conditions as **COUNTER**, with one exception. Because **COUNTER** is a "smart" data type, it can wrap the data when the system reaches a maximum number of values. If a reporting loss occurs, the system (when looking at **COUNTER** values) might erroneously wrap the data, thereby creating an artificial spike in the system and statistical anomalies.
 - **COUNTER** - Saves the rate of value change over a step period (interval). **COUNTER** assumes that the value is always increasing (the difference between the current and the previous value is greater than 0). Example use: Traffic counters on a router.
 - **DERIVE** - Same as **COUNTER**, but allows negative values. Example use: To see the rate of change in free disk space on a server.
 - **GAUGE** - Does not save the rate of change, but saves the actual value. There are no divisions or calculations. Example use: To see memory consumption in a server.
 - **ABSOLUTE** - Saves the rate of change, but assumes that the previous value is set to 0. The difference between the current and the previous value always equals the current value. Thus, **ABSOLUTE** stores the current value, divided by the step interval. - **Create Command** - Enter an RRD expression that is used to create the database for this data point. If you do not enter a value, then the system uses a default that is applicable to most situations.
 - **RRD Minimum** - Enter a value. This system ignores values less than this number.
 - **RRD Maximum** - Enter a value. This system ignores values greater than this number.
6. Click **Save**.

Data point aliases

Performance reports pull information from various data points that represent a metric. The report itself knows which data points it requires, and which modifications are needed, if any, to put the data in its proper units and format.

The addition of a data point requires changing the report.

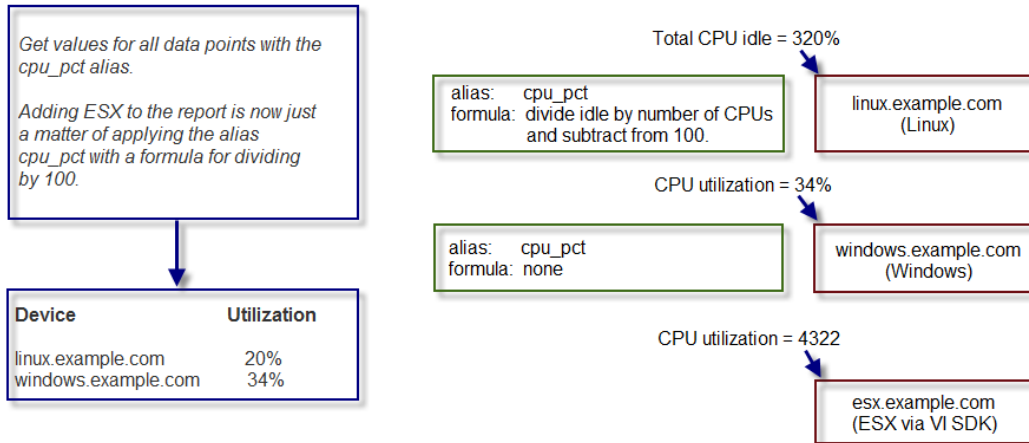
CPU Utilization Report



To allow for more flexibility in changes, some reports use *data point aliases*. Data point aliases group data points so they can be more easily used for reporting. In addition, if the data points return data in different units, then the plugin can normalize that data into a common unit.

An alias-based report looks up the data points that share a common alias string, and then uses them. This approach allows you to add data points without changing the report.

CPU Utilization Report (using aliases)



In the simplest cases, data from the target data points are returned in units expected by a report. For cases in which data are not returned in the same units, an alias can use an associated formula at the data point. For example, if a data point returns data in kilobytes, but the report expects data in bytes, then the formula multiplies the value by 1024.

Alias formula evaluation

The system evaluates the alias formula in three passes.

Pass 1: Reverse polish notation

When complete, the alias formula must resolve to a Reverse Polish Notation (RPN) formula.

In RPN, the formula to convert kilobytes into bytes is:

```
1024,*
```

For more information about RPN formulas, refer to [RPN documentation](#).

Pass 2: TALES expressions

For cases in which contextual information is needed, the alias formula can contain a TALES expression that has access to the device as context (labeled as "here"). The result of the TALES evaluation should be an RRD formula.

In TALES, the formula to convert kilobytes into bytes is:

```
${here/hw/totalMemory},/
```

Pass 3: Python expressions

You also can embed full Python code in an alias formula. The code must construct a string that results in a valid RRD formula.

To signal the system to evaluate the formula correctly, it must begin with:

```
__EVAL:
```

In Python, the formula to convert kilobytes into bytes is:

```
__EVAL:here.hw.totalMemory + ",/"
```

Adding a data point alias

To add an alias to a data point:

1. Navigate to a data source on a monitoring template.
2. Double-click a data point in the list to edit it. The Edit Data Point dialog appears.
3. Enter the alias name and the formula.
Note: If the data point returns values in the preferred units, then leave the Formula value blank.
4. Click Save.

Thresholds

Thresholds define expected bounds for data points. When the value returned by a data point violates a threshold, the system creates an event. There are several threshold types available:

- MinMax
- ValueChange
- CiscoStatus
- PredictiveThreshold

There are many thresholds already defined in the system. You can see all the defined thresholds in the Defined Thresholds report which can be accessed on the REPORTS > Enterprise Reports menu.

The following sections describe each type of threshold and how to create a new one and edit an existing one.

Additional threshold types may be provided through installed ZenPacks.

MinMax threshold

MinMax thresholds inspect incoming data to determine whether it exceeds a given maximum or falls below a given minimum. You can use a MinMax threshold to check for these scenarios:

- *The current value is less than a minimum value.* To do this, you should set only a minimum value for the threshold. Any value less than this number results in creation of a threshold event.
- *The current value is greater than a maximum value.* To do this, you should set only a maximum value for the threshold. Any value greater than this number results in creation of a threshold event.
- *The current value is not a single, pre-defined number.* To do this, you should set the minimum and maximum values for the threshold to the same value. This will be the only "good" number. If the returned value is not this number, then a threshold event is created.
- *The current value falls outside a pre-defined range.* To do this, you should set the minimum value to the lowest value within the good range, and the maximum value to the highest value within the good range. If the returned value is less than the minimum, or greater than the maximum, then a threshold event is created.
- *The current value falls within a pre-defined range.* To do this, you should set the minimum value to the highest value within the bad range, and the maximum value to the lowest value within the bad range. If the returned value is greater than the maximum, and less than the minimum, then a threshold event is created.

Editing MinMax thresholds

The threshold must be created by the Add Threshold functionality or already be defined in the system before you can edit it.

To edit a MinMax threshold:

1. Double-click the threshold in the list.

Edit Threshold

Name: High Utilization Rx

Description:

Type: MinMaxThreshold

Explanation:

Resolution:

DataPoints:

fcErrStats_crcRx fcStats_bytesRx

fcErrStats_discardRx

fcErrStats_discardTx

fcErrStats_linkFailures

fcErrStats_rx

fcErrStats_signalLosses

fcErrStats_syncLosses

fcErrStats_tooLongRx

fcErrStats_tooShortRx

fcErrStats_tx

fcStats_bytesTx

fcStats_packetsRx

fcStats_packetsTx

Severity: Warning

Enabled

Minimum Value:

Maximum Value: (here.linkSpeed or 1e10) / 8 * .9

Event Class: /Perf/Interface

Escalate Count: 0

SAVE CANCEL

2. Enter or select values to define the threshold:

- **Name**- Displays the value for the ID you entered on the Add a New Threshold dialog.
- **Description**- Description of the threshold that you entered on the Add a New Threshold dialog. The description is included in each event that is created from this threshold.
- **Type**- Type of threshold that you selected. You cannot change the type of threshold. If you want a different type of threshold, you need to create a new one and assign the correct type.
- **Explanation**- Information field where a user can enter information about what the event means. This field is included in each event that is created from this threshold
- **Resolution**- Information field where a user can enter information about what to do to resolve the event. This field is included in each event that is created from this threshold.
- **Data Points**- Select one or more data points to which this threshold will apply and click the right-arrow button to move them to the selected column.
- **Severity**- Select the [event severity level](#) of the first event triggered when this threshold is breached.
- **Enabled**- Select the check box to enable the threshold, or clear the check box to disable it.
- **Minimum Value**- If this field contains a value, then each time one of the select data points falls below this value an event is triggered. This field may contain a number or a Python expression. When using a Python expression, the variable here references the device or component for which data is being collected. For example, a 90% threshold might be specified as:

```
(here.linkSpeed or 1e10) / 8 * .9
```

The division by 8 is needed because interface speed frequently is reported in bits/second, where the performance data is bytes/second.

- **Maximum Value**- If this field contains a value, then each time one of the selected data points goes above this value an event is triggered. This field may contain a number or a Python expression.
- **Event Class**- Select the event class of the event that will be triggered when this threshold is breached.
- **Escalate Count**- Enter the number of consecutive times this threshold can be broken before the event severity is escalated by one step. A value of zero (0) indicates that the severity will not escalate.

3. Click Save to confirm the edits.

ValueChange threshold

ValueChange thresholds inspect incoming data to determine whether a status change has occurred and if so issues an event based on the defined severity.

Editing ValueChange thresholds

The threshold must be created by the Add Threshold functionality or already be defined in the system before you can edit it.

To edit a ValueChange threshold:

1. Double-click the threshold in the list.

Edit Threshold

Name: #OperStatusChange

Type: ValueChangeThreshold

Data Points:

#InErrors_#InErrors	→	#OperStatus_#OperStatus
#InOctets_#InOctets	←	
#InUcastPackets_#InUcastPackets	→	
#OutErrors_#OutErrors	←	
#OutOctets_#OutOctets	→	
#OutUcastPackets_#OutUcastPackets	←	

Severity: Info

Enabled

Event Class: #Status.Perf

SAVE CANCEL

2. Enter or select values to define the threshold:
 - **Name**- Displays the value for the ID you entered on the Add a New Threshold dialog.
 - **Type**- Type of threshold that you selected. You cannot change the type of threshold. If you want a different type of threshold, you need to create a new one and assign the correct type.
 - **Data Points**- Select one or more data points to which this threshold will apply and click the right-arrow button to move them to the selected column. When the datapoint changes status, an event will be triggered.
 - **Severity**- Select the severity level of the first event triggered when this threshold is breached.
 - **Enabled**- Select the check box to enable the threshold, or clear the check box to disable it.
 - **Event Class**- Select the event class of the event that will be triggered when this threshold is breached.
3. Click Save to confirm the edits.

CiscoStatus threshold

The CiscoStatus threshold is a special threshold that uses preconfigured maps to translate the numeric values returned by SNMP data sources to Resource Manager event severities. The following is an example OID and its mapped values.

CISCO-ENTITY-FRU-CONTROL-MIB::cefcModuleOperStatus

OID value	Description	Resource Manager severity level
1	unknown	Critical
2	OK	Clear
3	disabled	Clear
4	OK (diag failed) -	Warning
5	boot	Warning
6	self-test	Warning
7	failed	Critical
8	missing	Critical
9	mismatch w/parent	Critical
10	mismatch w/config	Critical
11	diag-failed	Critical
12	dormant	Critical
13	out of service (admin)	Info
14	out of service (environ)	Critical
15	powered down	Critical
16	powered up	Critical
17	power denied	Critical
18	power cycle	Warning
19	OK (power warning)	Warning
20	OK (power critical)	Critical
21	sync in progress	Clear
22	upgrading	Critical
23	OK (auth failed)	Critical

For the complete list of supported OIDs and values, refer to the [CiscoStatus Threshold](#) page.

For more information about event severity levels, see [Event severity levels](#).

Editing CiscoStatus thresholds

The threshold must be created by the Add Threshold functionality or already be defined in the system before you can edit it.

To edit a CiscoStatus threshold:

1. Double-click the threshold in the list.

Edit Threshold

Name: status

Type: CiscoStatus

DataPoints:

- cvnVethIfAdditionalState_cvnVethIfAdditionalState
- ifHCInBroadcastPkts_ifHCInBroadcastPkts
- ifHCInMulticastPkts_ifHCInMulticastPkts
- ifHCInOctets_ifHCInOctets
- ifHCInUcastPkts_ifHCInUcastPkts
- ifHCOutBroadcastPkts_ifHCOutBroadcastPkts
- ifHCOutMulticastPkts_ifHCOutMulticastPkts
- ifHCOutOctets_ifHCOutOctets
- ifHCOutUcastPkts_ifHCOutUcastPkts
- ifInDiscards_ifInDiscards
- ifInErrors_ifInErrors
- ifOutDiscards_ifOutDiscards
- ifOutErrors_ifOutErrors

Severity: Warning

Event Class Key: ifOperStatus

Enabled

Event Class:

SAVE CANCEL

2. Enter or select values to define the threshold:

- **Name**- Displays the value for the ID you entered on the Add a New Threshold dialog.
- **Type**- Type of threshold that you selected. You cannot change the type of threshold. If you want a different type of threshold, you need to create a new one and assign the correct type.
- **Data Points**- Select one or more data points to which this threshold will apply and click the right-arrow button to move them to the selected column. When the datapoint changes status, an event will be triggered.
- **Severity**- Select the severity level of the first event triggered when this threshold is breached.
- **Event Class Key**- The event class key from the data point you selected will be shown here.
- **Enabled**- Select the check box to enable the threshold, or clear the check box to disable it.
- **Event Class**- Select the event class of the event that will be triggered when this threshold is breached.

3. Click Save to confirm the edits.

Predictive threshold

The Predictive threshold allows you to use data from the past to project a value in the future and send an event if the future projected threshold is breached during a pre-defined time period. Resource Manager uses a linear projection algorithm for all projections. For example, you could predict total raw storage capacity of a disk or the bandwidth utilization percentage. Several predictive thresholds are already defined in the system. You can see all the defined thresholds in the Defined Thresholds report which can be accessed on the REPORTS > Enterprise Reports menu.

Predictive thresholds are also used for creating a trendline on a graph. You can add the threshold to the management of graph points so that a trendline will be added to the graph.

Editing Predictive Thresholds

The threshold must be created by the Add Threshold functionality or already be defined in the system before you can edit it.

1. Double-click the threshold in the list.

Edit Threshold

Name: Projected High Pct Rx

Description:

DataPoint: bandwidthUtilizationPctRx_bandwidthUtilizationPctf

Event Class: /Capacity

Type: PredictiveThreshold

Enabled

Severity: Info

Aggregate Function: max

Projection Algorithm: linear

Projection Algorithm Parameters:

Alerting

Amount of Data Used in Projection: 10 days

Send an Event if the Threshold is Breached in the Next: 10 days

Minimum Value:

Maximum Value: 90

SAVE CANCEL

2. Enter or select values to define the threshold:

- **Name**- Displays the value for the ID you entered on the Add a New Threshold dialog.
- **Description**- Description of the threshold that you entered on the Add a New Threshold dialog. The description is included in each event that is created from this threshold.
- **Data Point**- Select the data point to which this threshold will apply.
- **Event Class**- Select the event class of the event that will be triggered when this threshold is breached.
- **Type**- Type of threshold that you selected. You cannot change the type of threshold. If you want a different type of threshold, you need to create a new one and assign the correct type.
- **Enabled**- Select the check box to enable the threshold, or clear the check box to disable it.
- **Severity**- Select the severity level of the first event triggered when this threshold is breached.
- **Aggregate Function**- The type of function to use when analyzing the past data. For example, do you look at the peak values in the past or average values. The default value is max.
- **Algorithm**- Algorithm to use for the projection. Value is linear.
- **Projection Algorithm Parameters**- Some algorithms need additional parameters. Enter them as required.
- **Amount of Data Used in Projection**- Set the amount of historical data to be used when calculating the projection.
- **Send an Event if the Threshold is Breached in the Next**- Set the time frame in days, weeks, or months. If the threshold is breached within this time period in the future, an event will be sent.
- **Minimum Value**- If this field contains a value, then each time one of the select data points falls below this value an event is triggered. This field may contain a number or a Python expression. An expression is required when using gauge or calculated values. When using an expression, the variable here references the device or component for which data is being collected. For example, a 90% threshold for interface speed might be specified as:

```
(here.linkSpeed or 1e10) / 8 * .9
```

- The division by 8 is needed because interface speed frequently is reported in bits/second, where the performance data is bytes/second.
- **Maximum Value**- If this field contains a value, then each time one of theselected data points goes above this value an event is triggered. This field may contain a number or a Python expression. An expression is required for gauge or calculated values.

3. Click Save to confirm the edits.

Adding thresholds

To define a threshold for a data point:


1. From the Navigation menu, select **ADVANCED > MONITORING TEMPLATES**.
2. Click on the template that contains the data point you want to use in your threshold. The Data Sources window is populated with folders containing all the data points being collected.
3. Select the data point by opening the appropriate folder in the data source and clicking on the data point row.
4. In the Thresholds area, click the Add icon. The Add Threshold dialog box appears.
5. Enter a name and select the threshold type, then click Add. The threshold name and type is displayed in the Thresholds window.

Adding a trendline to a graph

Trendlines show you the projected utilization of a device or component over a specified range of time. The following procedure is tailored to be performed on a CiscoUCS device. In order to add a trendline to a graph, the threshold that you want to use must already be created in the system.

To add a trendline to an existing graph:

1. Navigate to **INFRASTRUCTURE > Devices**, then select a CiscoUCS device.
The selected device's overview page appears.
2. From the **Devices** panel, select **Components > Chassis**, and then select a chassis from the **Chassis** list.
3. Change the **Display** drop-down menu to **Templates**.
4. To the right of **Graph Definitions**, click **+** to create a new graph.
5. In the **Add Graph Definition** dialog, enter a name for the graph.
Alternatively, click the name of an existing graph to add the trendline to that graph.
6. Select the new graph, then from the gear icon, select **Manage Graph Points**.
The **Manage Graph Points** dialog box appears.
7. Click **+** > **Threshold**.



8. Select the name of the new predictive threshold and click **Submit**.
9. In the **Manage Graph Points** dialog, click **Save**.
10. To view the trendline, change the **Display** drop-down menu to **Graphs**.
11. Scroll down to the new (or existing graph) to view the newly added trendline.
Note: By default only the last 10 days of data is used in the trendline calculation. There may not be enough data to see the trendline appear on the graph. However, you can view the projected exhaustion dates in the lower right corner.

Performance graphs

You can include any data point or threshold from a monitoring template in a *performance graph*.

1. From the navigation menu, choose ADVANCED > MONITORING TEMPLATES.
2. In the left column, select the monitoring template in which you want to create a graph.
3. In the Graph Definitions area, click Add.
4. In the Add Graph Definition dialog box, enter a name for the graph, and then click Submit.
5. Double-click the new graph name and enter information or select values to define the graph:
 - **Name** - Optionally edit the name of the graph. This name appears as the title of the graph.
 - **Height** - Enter the height of the graph in pixels.
 - **Units** - Enter a label for the graph's vertical axis.
 - **Logarithmic Scale** - Activate the check box to specify that the scale of the vertical axis is logarithmic. Deactivate the check box (the default) to set the scale to linear. Example use: Use Logarithmic Scale if the data being graphed grows exponentially. Only positive data can be graphed logarithmically.
 - **Base 1024** - Activate the check box if the data that you are graphing is measured in multiples of 1024.
 - **Min Y** - Enter the bottom value for the graph's vertical axis.
 - **Max Y** - Enter the top value for the graph's vertical axis.
 - **Description** - Enter a description of the graph.
 - **Has Summary** - Activate the check box (default) to display a summary of the data's current, average, and maximum values at the bottom of the graph.
6. Click Submit.

Graph points

Graph points represent each data point or threshold that is part of a graph. You can add any number of graph points to a graph definition by adding data points or thresholds.

From the Graph Definitions area of the Monitoring Templates page:

1. From the Action menu, select Manage Graph Points. The Manage Graph Points dialog box appears.
2. From the Add menu, add a data point, threshold, or custom graph point.
3. Select values, and then click Submit. Note: Thresholds are always drawn before other graph points.

Re-sequencing graph points

To re-sequence graph points, drag a graph point row in the Manage Graph Points dialog box. (Click-and-drag from an "empty" part of the row.)

DataPoint graph points

DataPoint graph points draw the value of data points from the template on a graph.

Adding DataPoint graph points

To define a DataPoint graph point:

1. From the Add menu on the Manage Graph Points dialog, select Data Point. The Add Data Point dialog box appears.
2. Select one or more data points defined in this template. On data point graph point is created for each data point you select from the list.
3. Optional: Select the Include Related Thresholds option. If selected, then any graph points are created for any thresholds that have been applied to the selected data points as well.
4. Click Submit.

Editing DataPoint graph points

Double-click the name of the graph point to go to its edit page. Enter information or select values to edit the graph point:

- Name - This is the name that appears on the Graph Definition page. By default, it appears in the graph legend.
- Line Type - Select Line to graph the data as a line. Select Area to fill the area between the line and the horizontal axis with the line color. Select None to use this data point for custom RRD commands and do not want it to be explicitly drawn.
- Line Width - Enter the pixel width of the line.
- Stacked - If selected, then the line or area is drawn above the previously drawn data. At any point in time on the graph, the value plotted for this data is the sum of the previously drawn data and the value of this data point now. You might set this value, for example, to assess total packets if measuring packets in and packets out.
- Format - Specify the RRD format to use when displaying values in the graph summary. For more information about formatting strings, refer to its [documentation](#).
- RPN - Optionally enter an RPN expression that alters the value of the data being graphed for the data point. For example, if the data is stored as bits, but you want to graph it as bytes, enter an RPN value of "8,/" to divide by 8. For more information about RPN notation, refer to the [RPN tutorial](#).
- Limit - Optionally specify a maximum value for the data being graphed.
- Consolidation - Specify the RRD function used to graph the data point's data to the size of the graph. Most of the time, the default value of AVERAGE is appropriate.
- Color - Optionally specify a color for the line or area. Enter a six-digit hexadecimal color value with an optional two-digit hex value to specify an alpha channel. An alpha channel value is only used if 'stacked' is True.
- Legend - Name to use for the data in the graph legend. By default, this is a TALES expression that specifies the graph point name. The variables available in this TALES expression are here (the device or component being graphed) and graphPoint (the graph point itself).
- Available RRD Variables - Lists the RRD variables defined in this graph definition. These values can be used in the RPN field.

Editing threshold graph points

Threshold graph points graph the value of thresholds from the template.

To edit a threshold graph point, double-click it in the list.

You can edit values for Name, Color, and Legend for a threshold graph point.

Performance data retention

Resource Manager stores all performance data (a.k.a., metrics) in HBase using OpenTSDB. The default retention policy saves performance data for 90 days. To change the default, the time to live (TTL) must be adjusted on the OpenTSDB column families in HBase.

Note: TTL is defined in seconds.

Once a TTL value is changed, the data retention will adjust on the next major HBase compaction, which by default is once per day.

Changing the performance data retention time

To change the performance data retention time from the default value of 90 days:

1. Log in to the Control Center master host as a user with sudo and docker privileges.
2. Stop the openTSDB writer service.

```
serviced service stop opentsdb/writer
```

3. Execute the following command to list all the services and their SERVICEID values. Take note of the SERVICEID for the openTSDB reader service. It will be used as an argument in the following step.

```
serviced service list
```

4. Execute the following command, where \$id is the openTSDB reader SERVICEID and \$ttl is your TTL value, in seconds.

```
serviced service shell $id /opt/opentsdb/set-opentsdb-table-ttl.sh $ttl
```

5. Start the openTSDB writer service.

```
serviced service start opentsdb/writer
```

Distributed monitoring

Resource Manager supports distributed monitoring through the use of collectors, hubs, and resource pools; each of which is detailed in the sections below.

At installation, Resource Manager is configured with one hub and one collector (each named localhost). The initial hub and collector cannot be removed.

About collectors

A collector is a logical set of collection services that share a common group of monitored devices. Collection services are distributed across all of the hosts in the collector's resource pool. When you add a device to Resource Manager, you choose the collector to which it is assigned, which will normally be the collector that runs on the same network as (or closest network to) the device. You can move a device from one collector to another if necessary.

Collector services retrieve and accept data from your IT infrastructure for

- Monitoring
- Modeling
- Event generation

Collector data storage

Collector services initially queue collected performance metrics to a Redis key-value store on the collector, and then send the data to a single OpenTSDB (time series database) instance that runs on HBase for persistent storage.

Managing collectors

In Resource Manager, collector services appear under their collector name on the ADVANCED > Control Center page and can be started, stopped, and restarted from there. You can click on the name of a service to display its details, configuration files, and logs.

About hubs

A hub represents one or more instances of the zenhub service, through which all collector services communicate with the object and event databases. Resource Manager supports multiple hubs, although most Resource Manager deployments require only one hub. All collectors must belong to exactly one hub; however, many collectors can be associated with a single hub. All hubs (and indirectly all collectors) refer to the same object and event databases. Typically, only very large systems benefit from multiple hubs. We recommend that you contact Zenoss Support or Professional Services before adding a new hub.

About resource pools

A resource pool is a group of hosts running the Control Center software that reside on the same local area network (LAN). When you add a host to Control Center you specify the resource pool to which it will belong. This allows for better collector distribution and efficiency by minimizing the network latency between collectors and the devices they monitor.

Initially, a single resource pool named default is defined. You can create additional resource pools by identifying them with a unique pool name. You can then add hosts and assign them to the desired pool. Once the resource pool is configured, you can add a collector that runs on that resource pool.

Adding a hub or collector

Add a collector

Follow these steps

1. From the ADVANCED > Control Center page, open the Add menu and select Add a new Collector.
2. Enter the requested information:
 - Name - Name of the new collector
 - Hub - Name of the hub to be associated with
 - Resource Pool - Resource pool where the new collector's services will run
 - Clone collector from - An existing collector to copy configuration information from
3. Click SUBMIT.

Add a hub

Follow these steps:

1. From the ADVANCED > Control Center page, open the Add menu and select Add a new Hub.
2. Enter the requested information:
 - Name - Name of the new hub
 - Resource Pool - Resource pool where the new hub's services will run
 - Clone Hub from - An existing hub to copy configuration information from
3. Click SUBMIT.

Note that typically only very large deployments benefit from multiple hubs. We recommend that you consult with Zenoss Support or Professional Services before adding a new hub.

Adding devices to collectors


A device is assigned to a collector when you add the device to Resource Manager. See [Adding and discovering devices](#) for detailed information on adding devices.

Moving all devices from one collector to another

If you decide to remove a collector, you should move all of the devices on that collector to another collector before removing the collector. If you need to move one or more individual devices between collectors, please see the next section, [Moving individual devices between collectors](#).

To move all of the devices associated with one collector to another collector.

1. Navigate to ADVANCED > Move Devices.



2. Click the name of the source collector whose devices you want to move, then select a target collector from the drop-down list.
3. Click Submit to confirm. Resource Manager moves the devices to the newly selected collector.

Moving individual devices between collectors

Moving a single device from one collector to another

Follow these steps.

1. Navigate to the device overview page for the device you wish to move. You can use the search box to quickly locate a device.
2. Click the edit link to the right of the Collector header

Device ID: mysql.hypothetical.loc
Connection Information [edit](#):
underwriting
Uptime: 00d:02h:43m:49s

Device Name: mysql.hypothetical.loc
Production State: Production
Priority: Normal

Collector [edit](#): localhost
Hardware Manufacturer [edit](#): Generic
Hardware Model [edit](#): Net-SNMP Agent

3. Select the new collector for the device from the drop-down list.
4. Click SAVE.

Moving multiple devices from one collector to another

Follow these steps.

1. Navigate to the lowest-level device class or device organizer that contains all of the devices you wish to move.
2. Select the devices you wish to move by holding down the Ctrl key and clicking the desired device rows. Note: click anywhere on the row that is not over a hyperlink.
3. From the Actions menu, select Set Collector.
4. Select the new collector for the devices from the drop-down list.
5. Click OK.

You can add the current collector to the set of displayed columns by following these steps.

1. Hover on a column name and click the down arrow to the right of the name.
2. Select Columns.
3. Select Collector.

Device	IP Address	Device Class	Production State	Events
127.0.0.1			...	
windows.hypothetical.loc		crossRM	Production	
mysql.hypothetical.loc	10.0.0.11	/Se		
wordpress.hypothetical.loc	10.0.0.10	/Se		
linuxssh.hypothetical.loc	10.0.0.6	/Se		
linuxsnmp.hypothetical.loc	10.0.0.5	/Se		

Columns menu items:
 System Name
 IP Address
 Device Class
 Device Status
 Production State
 Serial Number
 Tag Number
 Hardware Manufacturer
 Hardware Model
 OS Manufacturer
 OS Model
 Collector
 Priority

Navigating collectors and hubs

1. To view and manage collectors and hubs:
2. Log in as the Resource Manager user.
3. From the navigation menu, select **ADVANCED > Control Center**. The Control Center All Services page appears.

Settings Move Devices **Control Center** Monitoring Templates Jobs

All Services

+ - [Refresh] [Start] [Stop] [Restart]

Name	Type	Uptime	AutoStart	Restart	State	Host
CentralQuery	daemon	6 days, 9:06:34.813614	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
HMaster	daemon	6 days, 9:06:26.678772	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
Imp4MariaDB	daemon		<input type="checkbox"/>		Down	
Imp4OpenTSDB	daemon		<input type="checkbox"/>		Down	
localhost	hub				...	
localhost	collector				...	
collectorredis	daemon	6 days, 9:06:15.053303	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
MetricShipper	daemon	6 days, 9:06:33.888581	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
zencommand	daemon	6 days, 9:06:17.002472	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
zenmail	daemon		<input type="checkbox"/>		Down	
zenmodeler	daemon	6 days, 9:06:15.942412	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
zenperfnmp	daemon	6 days, 9:06:15.928226	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
zenping	daemon	6 days, 9:06:16.972205	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
zenpropertymon...	daemon	6 days, 9:06:17.017761	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
zenpython	daemon	6 days, 9:06:16.987216	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
zenucsevents	daemon	6 days, 9:06:16.956418	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
zensphere	daemon	5 days, 0:15:49.037176	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
zmionion	daemon	6 days, 9:06:15.901294	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
zenhub	daemon	6 days, 9:06:15.875667	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
mysql-events	daemon	6 days, 9:06:24.785442	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
mysql-model	daemon	6 days, 9:06:24.798263	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc

The page lists existing hubs and collectors in hierarchical form. Hubs are listed at the top level; collectors are nested below the hub to which they belong.

From this page, you can:

- Add a hub or a collector
- Delete a hub (which also deletes its associated collectors) or a collector
Note: You cannot delete the default hub and collector (localhost)
- View and edit hub settings
- Configure associated monitoring and performance templates

Select a hub to display details and graphs. The Resource Pool ID for the hub appears and can be changed if needed. You can add a description and password if needed.

Settings Move Devices **Control Center** Monitoring Templates Jobs

All Services Control Center

+ - ⚙ Start Stop Restart

Name	Type	Uptime	AutoStart	Restart	State	Host
CentralQuery	daemon	6 days, 9:23:35.624274	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
HMaster	daemon	6 days, 9:23:27.476577	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
Imp4MariaDB	daemon	6 days, 9:16:15.988272	<input type="checkbox"/>		Down	ucspm-stable.zenoss.loc
Imp4OpenTSDB	daemon	6 days, 9:16:34.818722	<input type="checkbox"/>		Down	ucspm-stable.zenoss.loc
localhost	hub	6 days, 9:16:17.944580	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
localhost	collector		<input type="checkbox"/>		Down	ucspm-stable.zenoss.loc
localhost	collectorredis	6 days, 9:23:15.905571	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc

Display: Details

Resource: default

Description:

Password:

Save Cancel

Select the zenhub daemon to view the details about the daemon, its logs, and to view and edit its configuration. Use the buttons on the top of the window to start, stop, or restart it.

Settings Move Devices **Control Center** Monitoring Templates Jobs

All Services

+ - ⚙ Start Stop Restart

Name	Type	Uptime	AutoStart	Restart	State	Host
CentralQuery	daemon	6 days, 9:29:35.690465	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
HMaster	daemon	6 days, 9:29:27.560066	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
Imp4MariaDB	daemon		<input type="checkbox"/>		Down	
Imp4OpenTSDB	daemon		<input type="checkbox"/>		Down	
localhost	hub				...	
localhost	collector				...	
zenhub	daemon	6 days, 9:29:16.754306	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc
mysql-events	daemon	6 days, 9:29:25.661023	<input checked="" type="checkbox"/>		Up	ucspm-stable.zenoss.loc

Display: Details

ID: Details

btnorm: Configuration Files

Description: Logs

Zenoss ZenHub

AutoStart: true

State: RUNNING

Text: zenhub

Restarting

Uptime: 6 days, 9:27:59.782546

Save Cancel

Monitoring Zenoss

Using the Resource Manager browser interface, you can monitor health and performance of Zenoss systems.

From the same or another instance, monitor one or more instances of the Control Center application management and orchestration system and the Resource Manager system.

Troubleshoot components and services, such as RabbitMQ, Nginx, MariaDB, ZooKeeper, and so on. Review the history of a health check status to determine the cause of a health check change. Determine why an instance or component is running slowly, and whether a component or service is overloaded.

Proactively manage the performance of an instance. Determine when you need to add instances, memory, CPU, or adjust configuration settings.

- [Monitoring Control Center](#)
- [Control Center components](#)
- [Customizing Control Center monitoring](#)
- [Monitoring file system storage](#)
- [Monitoring Resource Manager](#)
- [ZenossRM components](#)
- [Customizing Resource Manager monitoring](#)

Monitoring Control Center

Resource Manager provides capabilities to monitor and manage the health and performance of Control Center resources.

The installation or upgrade process creates and configures the /ControlCenter device class for monitoring the local Control Center instance. You can add and monitor local and remote instances. Control Center must be running and listening on the target device's HTTPS port. To monitor a remote instance, you must provide the user name and password for that instance.

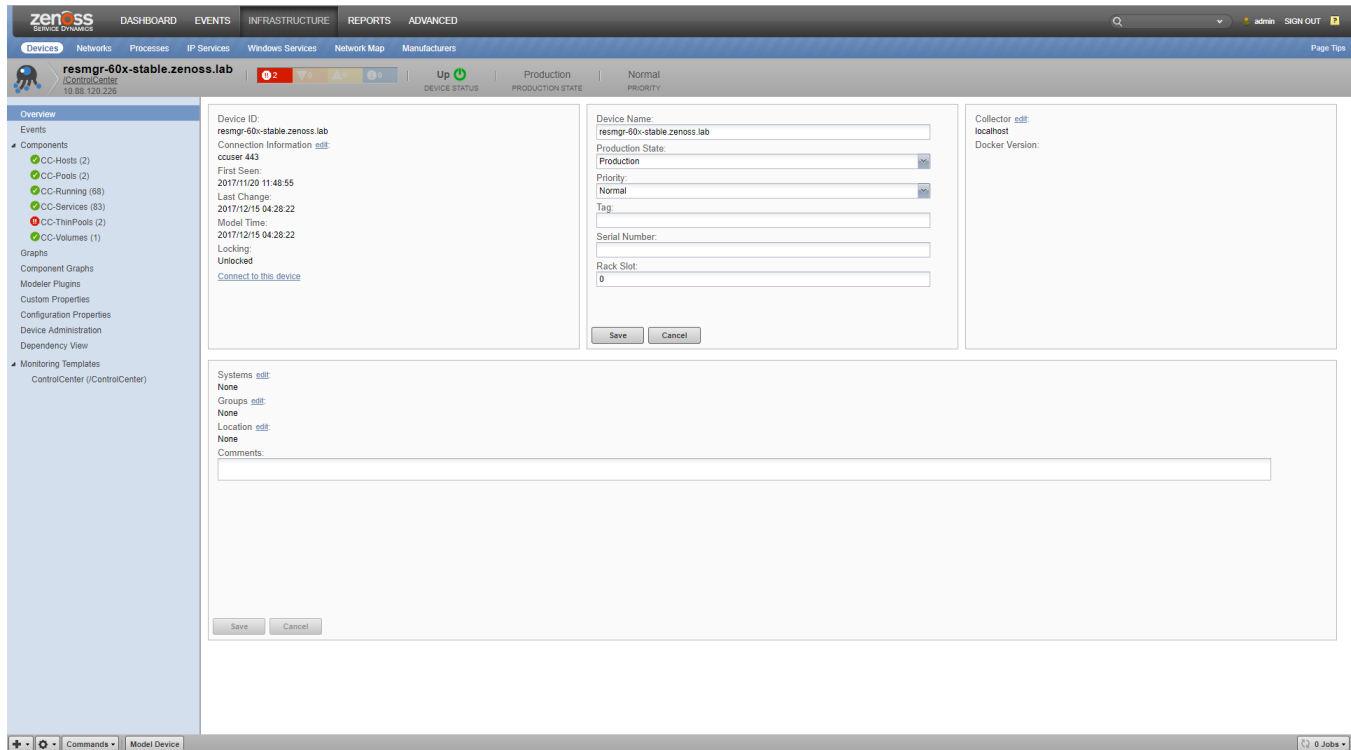
The modeler plugin is zenoss.ControlCenter. Modeling automatically discovers components of the /ControlCenter device. During normal remodeling intervals, the system updates component attributes. The system automatically generates data source types and monitoring templates for the /ControlCenter devices and device components.

Monitored Control Center resources include pools, hosts, services, running services, thin pools, and volumes. Graphs show memory usage, CPU usage, and page faults. Resource Manager generates events when a Control Center service fails or generates an error.

Pre-defined thresholds warn you and then issue an error when volume or thin pool space is low and extremely low. You can add custom events and thresholds, and edit shipped thresholds for components that have them.

Shipped threshold values are subject to change in future releases. If you edit a value, your change will be overwritten by future updates.

The following figure shows the overview page for the default Control Center device.

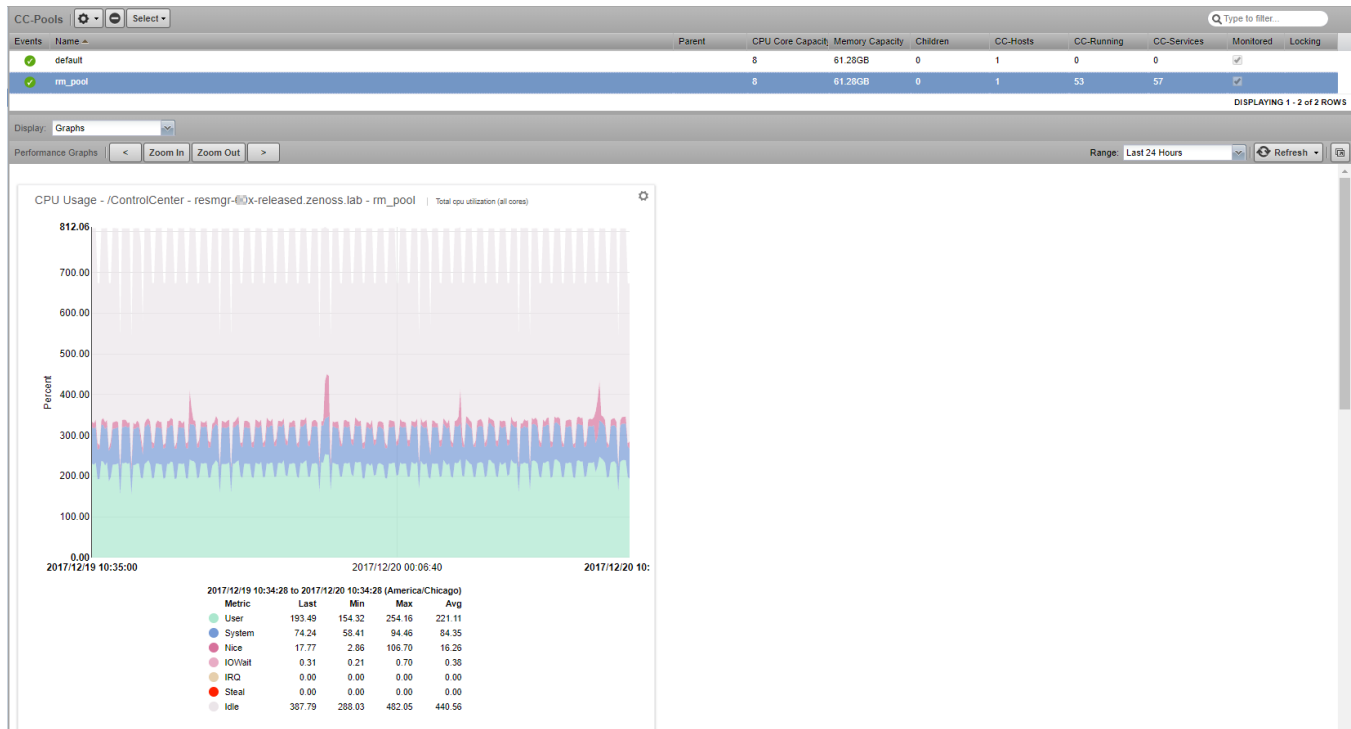


Control Center components

The /ControlCenter device includes the components in the following table.

Name	Description	Attributes	Relationships
CC-Hosts	Host systems that run Resource Manager services.	CPU cores, IP address, memory	CC-Pools, CC-Running
CC-Pools	Resource pools that are associated with Resource Manager services.	CPU core capacity, memory capacity	CC-Hosts, CC-Services, CC-Running
CC-Running	Running instances of services that are defined in CC-Services.	Instances, status, service ID, parent service ID, host ID	CC-Pools, CC-Hosts, CC-Services
CC-Services	All services that are defined for your system to run on CC-Hosts. Service definitions specify service properties and states.	Target state	Parent service, child services, CC-Pools, CC-Running
CC-Thin pools	LVM thin pools for Control Center application data. Thin pools include separate storage areas for data and metadata.	Driver, driver type, pool name, total bytes, used bytes	None
CC-Volumes	Storage volumes for the /ControlCenter device.	Total bytes, used bytes, data file, driver, status	None

The following figure shows the CPU usage graph for the resource pool named rm_pool.



Customizing Control Center monitoring

Customize monitoring of Control Center health and performance to fit your environment.

1. Add and model one or more Control Center instances as devices. For each device, specify the device class /ControlCenter, and set configuration properties that have prefix zControlCenter.
2. To determine base-level performance for the Control Center instance, study component graphs after a week, a month, and so on. If a metric for a component captures data that indicates a potential problem, take the following actions:
 - To the template for the affected component, add thresholds for the data point.
 - Create a threshold to send a warning level event that you can investigate.
 - Create another threshold to send an error level event to indicate a more serious problem.
 - Create events with event class /App/Zenoss for the threshold violations.
 - As you collect more data, make adjustments to thresholds and events as needed.

The following figure shows events for components of the default Control Center device.

The screenshot displays the Zenoss Control Center interface. The top navigation bar includes 'DASHBOARD', 'EVENTS', 'INFRASTRUCTURE', 'REPORTS', and 'ADVANCED'. The main content area shows a list of events for the device 'resmgr-53x.zenoss.lab'. The table below represents the data shown in the screenshot.

Status	Severity	Resource	Component	Event Class	Summary	First Seen	Last Seen	Count
⊗	3	resmgr-53x.zenoss.lab	CentralQuery	/App/Zenoss	CentralQuery-0: Healthchecks failed	12/21/17 10:07:50	12/21/17 10:07:50	1
⊗	3	resmgr-53x.zenoss.lab	writer	/App/Zenoss	writer-0: Healthchecks failed	12/21/17 10:07:50	12/21/17 10:07:50	1
⊗	3	resmgr-53x.zenoss.lab	RegionServer	/App/Zenoss	RegionServer-2: Healthchecks failed	12/21/17 10:07:50	12/21/17 10:07:50	1
⊗	3	resmgr-53x.zenoss.lab	RegionServer	/App/Zenoss	RegionServer-1: Healthchecks failed	12/21/17 10:07:50	12/21/17 10:07:50	1
⊗	3	resmgr-53x.zenoss.lab	RegionServer	/App/Zenoss	RegionServer-0: Healthchecks failed	12/21/17 10:07:50	12/21/17 10:07:50	1
⊗	3	resmgr-53x.zenoss.lab	HMaster	/App/Zenoss	HMaster-0: Healthchecks failed	12/21/17 10:07:50	12/21/17 10:07:50	1
⊗	3	resmgr-53x.zenoss.lab	reader	/App/Zenoss	reader-0: Healthchecks failed	12/21/17 10:07:50	12/21/17 10:07:50	1
⊗	3	resmgr-53x.zenoss.lab	CentralQuery	/App/Zenoss	CentralQuery-0: Healthchecks failed	12/21/17 08:47:50	12/21/17 08:47:50	1
⊗	3	resmgr-53x.zenoss.lab	RegionServer	/App/Zenoss	RegionServer-2: Healthchecks failed	12/21/17 08:47:50	12/21/17 08:47:50	1
⊗	3	resmgr-53x.zenoss.lab	HMaster	/App/Zenoss	HMaster-0: Healthchecks failed	12/21/17 07:27:50	12/21/17 07:27:50	1
⊗	3	resmgr-53x.zenoss.lab	RegionServer	/App/Zenoss	RegionServer-1: Healthchecks failed	12/21/17 07:27:50	12/21/17 07:27:50	1
⊗	3	resmgr-53x.zenoss.lab	CentralQuery	/App/Zenoss	CentralQuery-0: Healthchecks failed	12/21/17 06:07:50	12/21/17 06:07:50	1
⊗	3	resmgr-53x.zenoss.lab	writer	/App/Zenoss	writer-0: Healthchecks failed	12/21/17 06:07:50	12/21/17 06:07:50	1
⊗	3	resmgr-53x.zenoss.lab	reader	/App/Zenoss	reader-0: Healthchecks failed	12/21/17 06:07:50	12/21/17 06:07:50	1
⊗	3	resmgr-53x.zenoss.lab	zennreports	/App/Zenoss	zennreports-0: Healthchecks failed	12/20/17 23:15:25	12/20/17 23:15:25	1
⊗	3	resmgr-53x.zenoss.lab	MetricShipper	/App/Zenoss	MetricShipper-0: Healthchecks failed	12/20/17 22:15:33	12/20/17 22:15:33	1
⊗	3	resmgr-53x.zenoss.lab	zminion	/App/Zenoss	zminion-0: Healthchecks failed	12/20/17 22:15:33	12/20/17 22:15:33	1
⊗	3	resmgr-53x.zenoss.lab	RegionServer	/App/Zenoss	RegionServer-0: Healthchecks failed	12/20/17 21:50:25	12/20/17 21:50:25	1
⊗	3	resmgr-53x.zenoss.lab	RegionServer	/App/Zenoss	RegionServer-0: Healthchecks failed	12/20/17 18:06:15	12/20/17 18:06:15	1
⊗	3	resmgr-53x.zenoss.lab	HMaster	/App/Zenoss	HMaster-0: Healthchecks failed	12/20/17 18:06:15	12/20/17 18:06:15	1
⊗	3	resmgr-53x.zenoss.lab	RegionServer	/App/Zenoss	RegionServer-2: Healthchecks failed	12/20/17 18:06:15	12/20/17 18:06:15	1
⊗	3	resmgr-53x.zenoss.lab	MetricShipper	/App/Zenoss	MetricShipper-0: Healthchecks failed	12/20/17 17:11:19	12/20/17 17:11:19	1
⊗	3	resmgr-53x.zenoss.lab	zenneventserver	/App/Zenoss	zenneventserver-0: Healthchecks failed	12/20/17 17:11:19	12/20/17 17:11:19	1
⊗	3	resmgr-53x.zenoss.lab	MetricShipper	/App/Zenoss	MetricShipper-0: Healthchecks failed	12/20/17 15:51:19	12/20/17 15:51:19	1
⊗	3	resmgr-53x.zenoss.lab	Zope	/App/Zenoss	Zope-2: Healthchecks failed	12/20/17 15:51:19	12/20/17 15:51:19	1
⊗	3	resmgr-53x.zenoss.lab	MetricShipper	/App/Zenoss	MetricShipper-0: Healthchecks failed	12/20/17 15:51:19	12/20/17 15:51:19	1

Monitoring file system storage

Resource Manager monitors the Control Center application management and orchestration system and provides

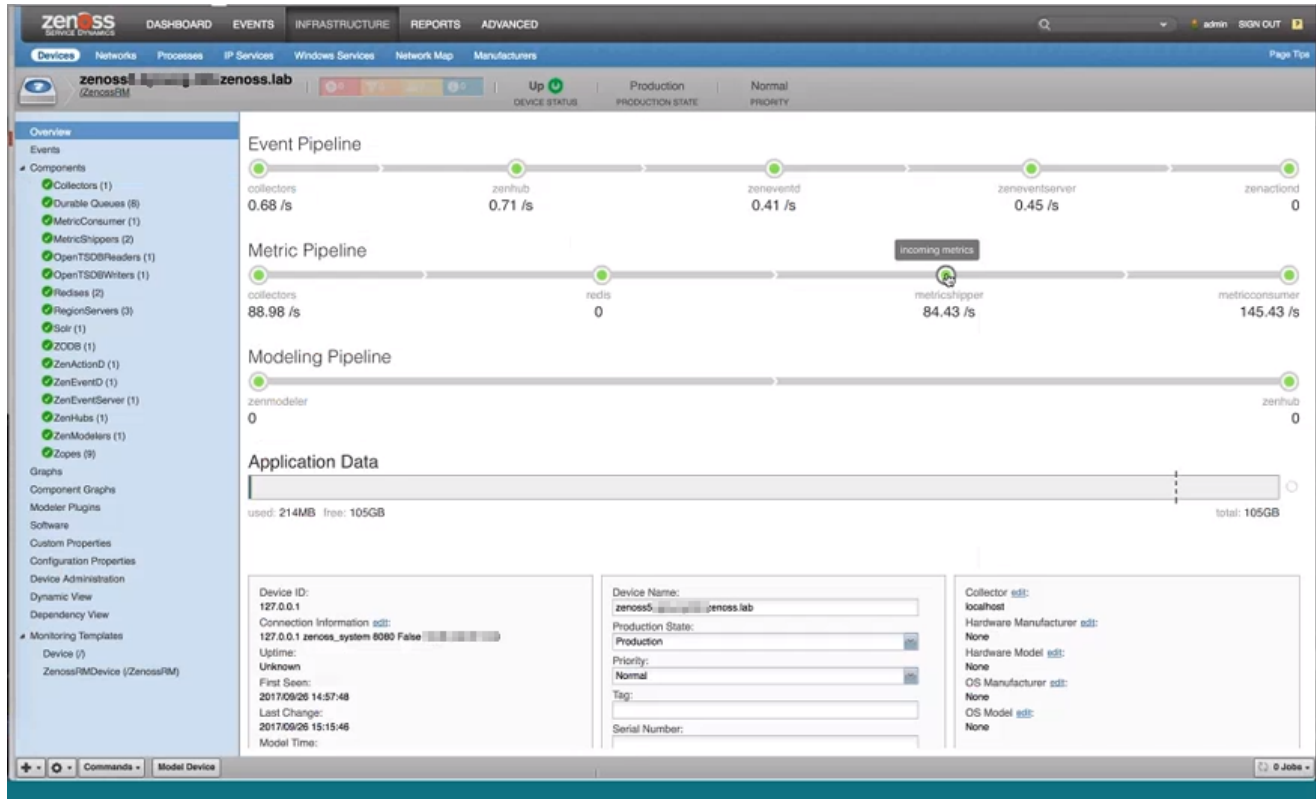
- storage usage reports for the thin pool data and metadata devices
- usage statistics in the browser interface and the command-line interface (`serviced volume status` command)
- for DFS volumes and thin pool data and metadata devices, a "space low" warning event (80% full) and a "space very low" alarm (90% full)

Set up a notification so that Resource Manager alerts you to low storage events. For more information, see [Triggers and notifications](#).

Monitoring Resource Manager

Resource Manager provides self-monitoring capabilities. A visual pipeline-based status page in the browser interface shows an overview of the monitored Resource Manager instance, including key metrics for nodes in the metric collection, event generation, and modeling processes.

The following figure shows the Resource Manager status overview page.



Resource Manager monitoring listens to events in the event class /ZenossRM. Each pipeline includes nodes and aggregate data for the associated components. Events on monitored components are used to flag a data pipeline node for your investigation. When a threshold violation occurs, the node changes from green to flashing gray, as shown by the metricshipper node in the figure. To open the detail page for the associated component, click a node in a pipeline.

Use data to troubleshoot performance issues with Resource Manager components and proactively manage performance. For example, data can help you determine when a component or service in the system is overloaded. Before issues occur, optimize system performance by adding more instances, more memory, more CPU, or adjusting configuration settings.

The Resource Manager installation or upgrade process creates device class /ZenossRM, and creates and models the local Resource Manager device with ID 127.0.0.1. Locally stored metrics for the monitored device provide detailed performance tracking for Resource Manager components.

You can add multiple Resource Manager instances as devices. Monitor the health and performance of multiple Resource Manager and Control Center instances from the same or another Resource Manager instance. You can monitor remote Resource Manager instances on the same Control Center and remote Resource Manager instances that are on a different Control Center.

ZenossRM components

Components of the /ZenossRM device are automatically discovered. Attributes are updated on the normal remodeling interval, which defaults to 12 hours.

You can add custom events and custom thresholds, and edit shipped thresholds for components that have them.

Shipped threshold values are subject to change in future releases. If you edit a value, your change will be overwritten by future updates.

Collectors

The aggregation of collector daemons and their supporting services grouped as localhost.

CollectorDaemons

The aggregation of statistics broken down by individual daemon service; for example, zenpython and zencommand.

Durable Queues

Persistent RabbitMQ message queues.

MetricConsumer

Pulls data from the Redis queue and passes it to MetricShipper.

This component ships with a default threshold. The maximum number of seconds that MetricConsumer needs to process its internal queue at the current rate is currently 300. If this value is exceeded, the MetricConsumer node on the metric pipeline turns gray and flashes.

Shipped threshold values are subject to change in future releases. If you edit a value, your change will be overwritten by future updates.

MetricShipper

Inserts metrics into OpenTSDBWriter.

This component ships with a default threshold. The maximum number of seconds MetricShipper needs to process its Redis queue at the current rate is 300. If this value is exceeded, the MetricShipper node on the metric pipeline turns gray and flashes.

Shipped threshold values are subject to change in future releases. If you edit a value, your change will be overwritten by future updates.

OpenTSDBReader

Allows queries to metric storage.

OpenTSDBWriter

Performs writes to metric storage.

Redis

Used as a cache for device configuration and metrics data.

RegionServer

OpenTSDBWriter uses this component for metric storage.

Solr

Provides the index of modeled devices and system objects.

ZODB

Object database that stores the model.

ZenActionD

Watches the event stream and sends configured notifications.

ZenEventD

Filters, enhances, and transforms events.

This component ships with a default threshold. The maximum number of seconds ZenEventD needs to process the rawevents queue at the current rate is 300. If this value is exceeded, the ZenEventD node on the metric pipeline turns gray and flashes.

Shipped threshold values are subject to change in future releases. If you edit a value, your change will be overwritten by future updates.

ZenEventServer

Performs event processing, storage, and retrieval.

This component ships with a default threshold. The maximum number of seconds ZenEventServer needs to process the zenevents queue at the current rate is 300. If this value is exceeded, the ZenEventServer node on the metric pipeline turns gray and flashes.

Shipped threshold values are subject to change in future releases. If you edit a value, your change will be overwritten by future updates.

ZenHub

The central coordinator of event generation and configuration delivery to daemons.

ZenModeler

Periodically scans modeled devices and saves attribute changes to the model.

Zope

Web server processes that serve the browser interface, reports, Zenoss JSON API requests, and debugging.

Customizing Resource Manager monitoring

Customize monitoring of Resource Manager components to fit your environment.

1. Add and model one or more Resource Manager instances as devices. For each device, specify the device class `/ZenossRM`, and set configuration properties that have prefix `zRMM`.
2. To determine base-level performance for the Resource Manager instance, study component graphs after a week, a month, and so on. If a metric for a component captures data that indicates a potential problem, take the following actions:
 - To the template for the affected component, add thresholds for the data point.
 - Create a threshold to send a warning level event that you can investigate.
 - Create another threshold to send an error level event to indicate a more serious problem.
 - Create events with event class `/ZenossRM` for the threshold violations.
 - As you collect more data, make adjustments to thresholds and events as needed.

Extending Resource Manager with ZenPacks

ZenPacks extend and add new functionality to Resource Manager. This can be as simple as adding new device classes or monitoring templates, or as complex as extending the data model and creating new collection services.

You can use ZenPacks to add:

- Monitoring templates
- Data sources
- Graphs
- Event classes
- User commands
- Reports
- Model extensions
- Product definitions

Simple ZenPacks can be created completely within the user interface. Complex ZenPacks include development of scripts or services, using Python or another programming language. ZenPacks can be distributed for installation on other Resource Manager systems.

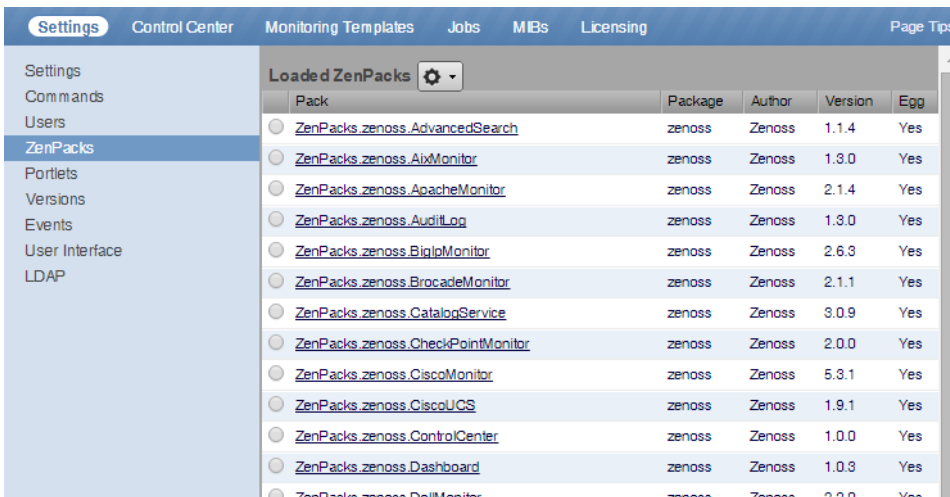
Displaying installed ZenPacks

You can display installed ZenPacks by using the browser interface or the command-line interface (CLI).

Displaying installed ZenPacks in the browser interface

To perform these steps, you must be logged in as a user with Manager or ZenManager privileges. Otherwise, the ADVANCED tab is not available.

1. Navigate to ADVANCED > Settings.
2. In the left column, click ZenPacks.



The screenshot shows the 'Settings' page in the Zenoss interface. The left sidebar has 'ZenPacks' selected. The main content area is titled 'Loaded ZenPacks' and contains a table with the following columns: Pack, Package, Author, Version, and Egg. The table lists various ZenPacks, all from the 'zenoss' package and 'Zenoss' author, with versions ranging from 1.0.0 to 3.0.9. Each row has a radio button in the 'Pack' column and a 'Yes' in the 'Egg' column.

Pack	Package	Author	Version	Egg
<input type="radio"/> ZenPacks.zenoss.AdvancedSearch	zenoss	Zenoss	1.1.4	Yes
<input type="radio"/> ZenPacks.zenoss.AixMonitor	zenoss	Zenoss	1.3.0	Yes
<input type="radio"/> ZenPacks.zenoss.ApacheMonitor	zenoss	Zenoss	2.1.4	Yes
<input type="radio"/> ZenPacks.zenoss.AuditLog	zenoss	Zenoss	1.3.0	Yes
<input type="radio"/> ZenPacks.zenoss.BigIpMonitor	zenoss	Zenoss	2.6.3	Yes
<input type="radio"/> ZenPacks.zenoss.BrocadeMonitor	zenoss	Zenoss	2.1.1	Yes
<input type="radio"/> ZenPacks.zenoss.CatalogService	zenoss	Zenoss	3.0.9	Yes
<input type="radio"/> ZenPacks.zenoss.CheckPointMonitor	zenoss	Zenoss	2.0.0	Yes
<input type="radio"/> ZenPacks.zenoss.CiscoMonitor	zenoss	Zenoss	5.3.1	Yes
<input type="radio"/> ZenPacks.zenoss.CiscoUCS	zenoss	Zenoss	1.9.1	Yes
<input type="radio"/> ZenPacks.zenoss.ControlCenter	zenoss	Zenoss	1.0.0	Yes
<input type="radio"/> ZenPacks.zenoss.Dashboard	zenoss	Zenoss	1.0.3	Yes
<input type="radio"/> ZenPacks.zenoss.DellMonitor	zenoss	Zenoss	2.2.0	Yes

Displaying installed ZenPacks in the CLI

To perform this procedure, you need a user account with `serviced` command-line interface (CLI) privileges on the Control Center master host.

1. Log in to the Control Center master host as a user with `serviced` CLI privileges.
2. Display the list of installed ZenPacks:

```
serviced service run zope zenpack list
```

ZenPack information resources

Zenoss provides a broad range of ZenPacks, described in the [ZenPack catalog](#).

You can create your own ZenPacks and have your ZenPacks and ZenPacks developed by others installed in Resource Manager. For more information, refer to the following ZenPack resources:

- [ZenPack SDK](#)
- [Zenoss Community](#), which includes the ZenPack development forum
- [Public Zenoss repositories on GitHub](#)

Preparing to install or update a ZenPack

Perform this procedure to minimize the amount of time that Resource Manager is unavailable during a ZenPack installation or update.

1. Log in to your workstation and start a web browser.
2. Download the ZenPack to install or update from the [Zenoss Support](#) site. Contact your Zenoss representative for login credentials.
3. Copy the ZenPack `egg` file to a local directory on the Control Center master host.
 - a. Create a directory for the ZenPack `egg` file. The directory must be local (not mounted). The following command creates a directory in `/tmp`:

```
mkdir /tmp/zenpack
```

- b. Use a file transfer command or utility to copy the file.
- c. Set full permissions on the directory and files:

```
chmod -R 0777 /tmp/zenpack
```

4. Optional: Install ZenPack dependencies. A ZenPack might require packages or other software not included in the ZenPack `egg` file. To ensure that the dependencies are available, perform the following substeps:
 - a. Log in to the Control Center master host as a user with Control Center CLI privileges.
 - b. Start an interactive shell in a Zope service container. In the following command, the `-s` flag saves and tags the changes that you make. Replace `MyTag` with a short name that describes the dependencies that you are installing.

```
serviced service shell -i -s MyTag zope bash
```

The `serviced` daemon starts a Bash shell and logs you in as the `root` user.

- c. Install required dependencies. For example, to install the `terminus` font for X Windows in Ubuntu Linux, enter the following command:

```
apt-get install xfonts-terminus
```

Enter any number of commands to install the required dependencies.

- d. Return to the Control Center master host shell session:

```
exit
```

- e. Create a snapshot and commit your changes:

```
serviced snapshot commit MyTag
```

- f. Restart all Resource Manager application services:

```
serviced service restart Zenoss.resmgr/Zenoss
```


Installing or updating a ZenPack

Before you begin, review the following requirements and considerations:

- Complete the steps in [Preparing to install or update a ZenPack](#).
- Do not use this procedure to install or update the Zenoss Service Impact ZenPacks, ZenPacks.zenoss.ImpactServer and ZenPacks.zenoss.Impact. For more information, see [Installing or updating Service Impact](#).

Perform these steps:

1. Log in to the Control Center master host as a user with Control Center CLI privileges.
2. Create a snapshot:

```
serviced service snapshot Zenoss.resmgr
```

On completion, the serviced command returns the ID of the new snapshot. If the installation of a ZenPack fails, you can restore the snapshot created in this step. For more information about restoring a snapshot, see [Creating snapshots and rolling back](#).

3. Change directory to the directory in which the ZenPack egg file is located.
For example:

```
cd /tmp/zenpack
```

4. Install the ZenPack:

```
serviced service run zope zenpack-manager install ZenPack-File.egg
```

Daemons that a ZenPack provides are packaged in Docker containers and installed as child services of the current instance of Resource Manager.

5. Restart all Zenoss services:

```
serviced service restart Zenoss.resmgr/Zenoss
```

Removing a ZenPack

Removing a ZenPack can have unexpected consequences! Often, the safest choice is not to remove a ZenPack.

Before you begin, review the following requirements and considerations:

- Removing a ZenPack removes all objects provided by the ZenPack and all objects that depend on code provided by the ZenPack.
- Removing a newer version of a ZenPack to install an older version fails if the newer version includes migration code.
- Removing a ZenPack that installs a device class removes the device class, any contained device classes, and all devices in that class.
- Some ZenPacks provide services upon which other ZenPacks rely. Make sure the service you remove is not needed by another ZenPack.
- Do not use this procedure to remove the Zenoss Service Impact ZenPacks, ZenPacks.zenoss.ImpactServer and ZenPacks.zenoss.Impact. For more information, see [Installing or updating Service Impact](#).
- Review the documentation of the ZenPack that you want to remove for information about classes and daemons (services) associated with it.
- Delete data sources provided by the ZenPack that you want to remove.

Perform these steps:

1. Log in to the Control Center master host as a user with Control Center CLI privileges.
2. Create a snapshot:

```
serviced service snapshot Zenoss.resmgr
```


On completion, the serviced command returns the ID of the new snapshot. If the installation of a ZenPack fails, you can restore the snapshot created in this step. For more information about restoring a snapshot, refer to the Control Center Reference Guide.

3. Obtain the exact name of the ZenPack to remove:

```
serviced service run zope zenpack list
```

The first item of each line of output is the full name of an installed ZenPack.

4. Stop services that are associated with the ZenPack.
Daemons that a ZenPack provides are packaged in Docker containers and installed as child services of the current instance of Resource Manager. For example, the `zenwebtx` service is provided by the ZenPacks.zenoss.ZenWebTx ZenPack.
 - a. Log into the Resource Manager browser interface as a user with Manager or ZenManager privileges.
 - b. From the navigation menu, select **ADVANCED > Control Center**.



- c. In the All Services area, select the name of the service to remove, and then click Stop.
For example, select the `zenwebtx` daemon.

5. Remove the ZenPack.

In the Control Center CLI, replace `ZenPack-Name` with the full name of the ZenPack to remove.

```
serviced service run zope zenpack-manager uninstall ZenPack-Name
```

The ZenPack and any daemons that it provides are removed.

6. Restart all Zenoss services:

```
serviced service restart Zenoss.resmgr/Zenoss
```

Creating a ZenPack

To perform this procedure, you need a user account with `serviced` CLI privileges on the Control Center master host.

This procedure demonstrates how to create a ZenPack for customized monitoring templates, customized event classes and mappings, device MIBs, and similar items which require no customized Python code. For more advanced ZenPack development, refer to the [ZenPack SDK](#) site.

1. Log in to the Control Center master host as a user with `serviced` CLI privileges.
2. Create a ZenPack.
Replace `ZenPacks.myOrg.myPackName` with the name of the ZenPack to create.

```
serviced service run zope zenpack-manager create ZenPacks.myOrg.myPackName
```

3. Restart the Zope service.

```
serviced service restart zope
```

Using organizers

Resource Manager includes implicit and explicit organizers for categorizing and accessing monitored devices.

- [Device classes](#)
- [Groups, Systems, and Locations](#)
- [Component groups](#)

Device classes

A device class is an implicit organizer that defines the base set of properties that characterize a group of devices. A device may belong to only one device class.


Templates and configuration properties can be inherited based on device class. These attributes can be overwritten further down the class hierarchy, all the way down to the individual component level. The device class hierarchy includes all defined and standard classes and sub-classes.

The following procedures are illustrated using device classes and sub-classes, but the same concepts apply to event classes, service classes, and product classes. When you add a device to the system, after providing the network name or IP address, specify its device class. Templates and configuration properties can be set at any level in the device class hierarchy.

Viewing device classes

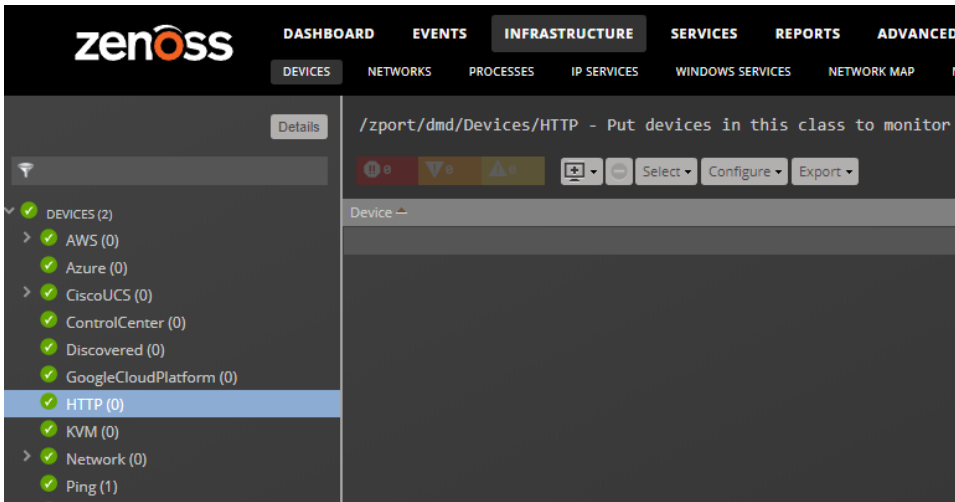
To view device classes and the devices they contain, from the navigation menu, select **INFRASTRUCTURE**.

At the top level of the device hierarchy are device classes, such as CiscoUCS. To view devices in a class or to expand the organizer to show subclasses, click a name in the tree. Severity indicators show the most severe type of event that is associated with any device in that class.




Adding a device class

1. After you add a device class, you can move devices into the class.
2. From the navigation menu, choose INFRASTRUCTURE.
3. In the left pane device class hierarchy, choose the parent device class in which to add a child class; for example, HTTP.



4. At the bottom of the left pane, click the Add icon . The Add Device Class dialog box appears.



- Specify a name and description for the new device class, and then click Submit.
- Optional: Choose one or more zProperties from the Configuration Information Properties list. The properties chosen will be available for configuration in the Add Infrastructure dialog. Their values will be available on the device overview pages of devices in this class.


The new device class appears in the hierarchy under the parent device class.

To move devices to the new class, choose the devices in the device list, and then drag them to the new class.

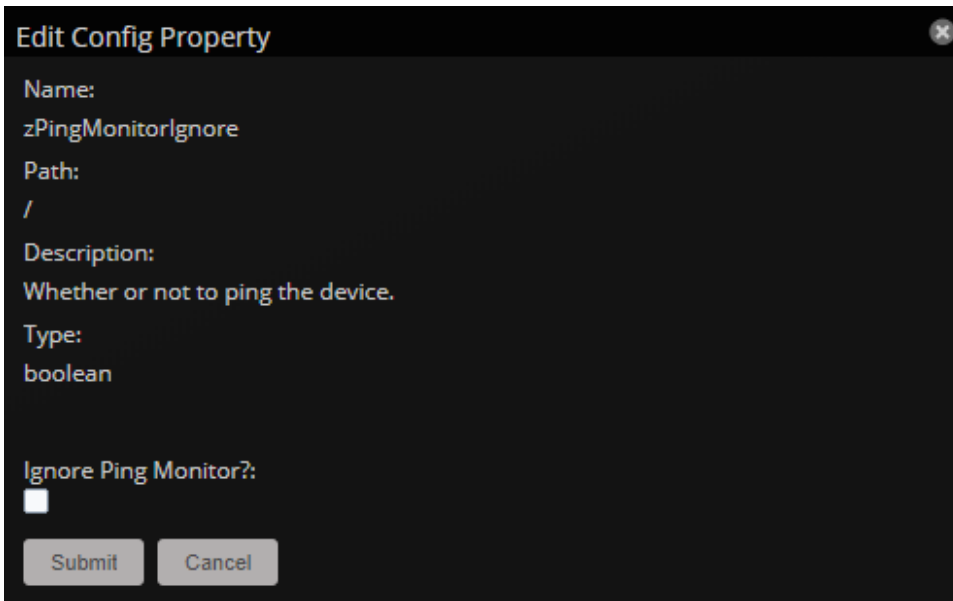
Setting configuration properties at the class level

Definitions are applied to all devices currently in the class and those added to the class (unless overridden at a lower level in the hierarchy).

- In the left pane device class hierarchy, choose Devices.
- Click Details  > Configuration Properties. The Configuration Properties page for the selected device class appears.



3. In the right pane, double-click a property to be edited.



Edit Config Property

Name:
zPingMonitorIgnore

Path:
/

Description:
Whether or not to ping the device.

Type:
boolean

Ignore Ping Monitor?:

Submit Cancel

4. In the Edit Config Property dialog box, change definitions.
5. Click SUBMIT.

Removing device classes

It is important to note that removing device classes is generally a bad idea, as they may contain templates and other valuable objects. Only remove device classes that you have added yourself; **do not remove device classes that exist by default or that have been added by ZenPacks.**

Groups, Systems, and Locations


Groups, Systems, and Locations are explicit organizers that enable arbitrary groupings of devices in the Resource Manager interface.

- **Groups** are functional divisions that allow you to assign attributes to multiple objects with similar functions. Groups can be used, for example, to arrange objects along departmental lines.
- **Systems** are intended to follow virtual setups, such as those in a network setup or systems grouped by functionality.
- **Locations** are logical groupings for physical systems that indicate the physical location of a device. A device can have one location, as general as city and state, or as specific as rack or closet. Locations appear in the Google Maps portlet.

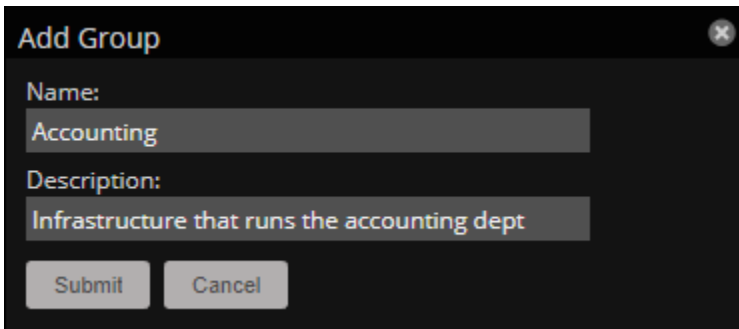
Adding a group, system, or location

To add a group:

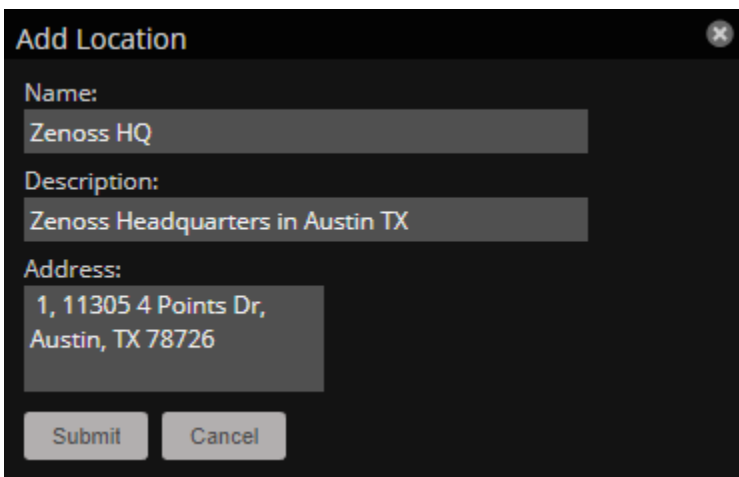
1. From the navigation menu, choose INFRASTRUCTURE.
2. In the left panel, choose the type of organizer from the heirarchy.



3. In the lower left, click the Add  icon.
4. In the Add dialog box, specify a name and description for the organizer.



Location organizers support an Address field, which is used by the Google Maps portlet.



The new organizer appears in the hierarchy. You can drag-and-drop devices from the device list to the new organizer.

Moving a group, system, or location organizer

To move a group:

1. Select the organizer in the hierarchy.
2. Drag the organizer to its new location. The Move Organizer confirmation dialog appears.
3. Click OK to confirm the action. The organizer appears at its new location in the hierarchy.


Location organizers and the Google Maps portlet

Resource Manager can map locations by using a Google Maps mashup feature that sets the location's Address property to a valid Google Maps address. The selected location appears on the map as a dot. The color of the dot represents the highest severity of any event on any device in that location.

Network connections that span locations are represented on the map by lines. Each line color matches the status of the connection it represents.


To view the Google Map for a network location:

1. From the navigation menu, choose INFRASTRUCTURE.
2. In the left panel, choose the LOCATIONS organizer and then choose a location.



3. Click **Details** above the device class tree.


4. Select Map.




You also can view the network location map from the Google Maps portlet on the dashboard.

Setting an address for a location

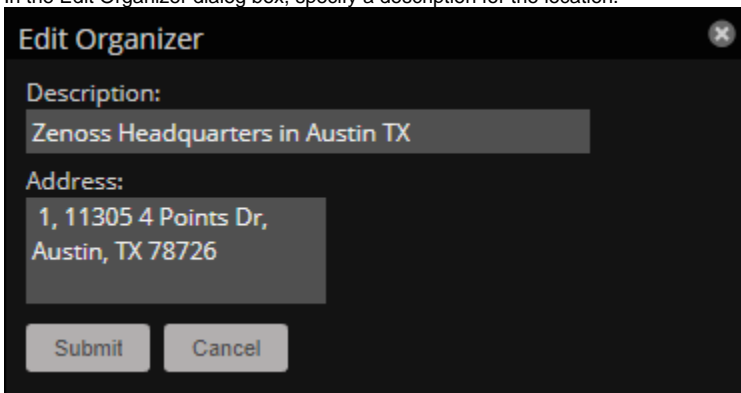
1. From the navigation menu, choose INFRASTRUCTURE.
2. In the left panel, choose the LOCATIONS organizer and then choose a location.



3. In the lower left, click the Action icon and choose Edit.



4. In the Edit Organizer dialog box, specify a description for the location.



The image shows a dark-themed dialog box titled "Edit Organizer". It has a close button in the top right corner. There are two text input fields. The first is labeled "Description:" and contains the text "Zenoss Headquarters in Austin TX". The second is labeled "Address:" and contains the text "1, 11305 4 Points Dr, Austin, TX 78726". At the bottom of the dialog box, there are two buttons: "Submit" and "Cancel".

5. In the Address field, specify the complete address with zip code, and then click Submit.
In the Address field, you can specify any information to identify the location in Google Maps. For example a country, state, city, intersection, or latitude and longitude coordinates. For more information, see the [Google Maps Help](#).

The selected address for the location is created. You must add at least one device to the location for the location "dot" to appear on the map.

Network links

If two devices in the same network are in different map-able locations, a line is on the map representing a network connection between the two. If there are multiple separate network connections between the same two locations, only one line is drawn. The color of the line represents the highest severity of any events affecting the connection. These are determined by:

- A ping down event on the device at either end of the connection; or
- Any event on the interface at either end of the connection.

Drawing map links (zDrawMapLinks configuration property)

Calculating network links on the fly is a time-intensive procedure. If you have a large number of devices that have been assigned locations, drawing those links on the map may take a long time.

To save time, you can tell the system not to attempt to draw links for specific networks. You might want to do this, for example, for a local network comprising many devices that you know does not span multiple locations.

To edit the value for this property:

1. From the navigation menu, choose INFRASTRUCTURE > NETWORKS.
2. On the Networks page, choose the network or sub-network for which you want to disable map links.
3. Display configuration properties for the network.
4. Double-click the zDrawMapLinks configuration property in the list. The Edit Config Property dialog box appears.
5. De-select the value (uncheck the box), and then click Submit.

Note: This setting will be inherited by networks or sub-networks below this selection in the hierarchy. If you have few networks for which links would be drawn, you might want to disable map links on /Networks, enabling it only on a network where you know a location-spanning WAN connection exists.

Google Maps example

This example will show you how to:

- Create and display Google map links of devices
 - Send a test event to see how map links are affected by system changes
1. Disable map links.
Refer to the procedure in Drawing Map Links for instructions.
 2. Create two locations: "New York" and "Los Angeles."
Refer to the procedure in Adding Locations for instructions.
 3. Enter the following values in the Address field of the Add Location dialog: "NewYork, NY" and "Los Angeles, CA" respectively.
 4. Set the location of a device to New York. Locate another device on the same network and set its location to Los Angeles.
 5. Select Locations in the hierarchy, click **Details**, and then select Map.
New York and Los Angeles are represented by dots on the map; however, no link is drawn between these locations.
 6. Select Networks and re-enable map linking.
 7. Select INFRASTRUCTURE > DEVICES, then select LOCATIONS in the hierarchy.
 8. Click **Details**, and then select Map.
A green line is now drawn between New York and Los Angeles.
 9. Send an event with a severity of Critical to the device in New York.
For information about creating events, see [Event Management](#). Do not specify a component.
 10. Return to the Locations map.
The dot representing New York is now red, but the link between New York and Los Angeles remains green.
 11. Navigate to the New York device and determine the ID of the component that is connected to the network shared with the Los Angeles device.
 12. Send another test event, this time specifying that component.
 13. Return to the Locations map.
The link between New York and Los Angeles is now red.

Clearing the Google Maps cache

Clearing the Geocode cache can solve issues with drawing maps and seeing the network status of locations or connections.

1. From the navigation menu, choose INFRASTRUCTURE.
2. In the left panel, choose the LOCATIONS organizer and then choose a location.
3. In the lower left, click the Action icon and choose Clear Geocode Cache.
4. Click OK.

Component groups

Component groups is an explicit organizer that enables arbitrary groupings of device components in the Resource Manager interface.

Use component groups to view or manage device component resources as a logical group. When you add device components to a component group, you can view all of the events for the group in a single location, display device component graphs with all of the device components on the same graph, and use the Dynamic and Dependencies views to see the group's dependents and dependencies. You can also monitor or lock individual device components in the group or the entire group itself.

Creating and viewing component group information

To create and view information about component groups:

1. Navigate to **INFRASTRUCTURE > DEVICES**.
2. Select **COMPONENT GROUPS**, then click the Add icon at the bottom of the page.
The Add Component Group dialog box appears.
3. Enter a suitable Name, and an optional Description, then click **SUBMIT**.
The new component group appears under **COMPONENT GROUPS**.
4. Under **COMPONENT GROUPS** at the top of the page, click the Add icon to open the Add to Component Group dialog.
5. In the dialog's search field, enter a component type, such as Blade Servers.
A list of components appears in the Search Results table.
6. Select one or more components, then click **ADD**.
7. Continue to search for and add components or click **CLOSE**.
The components are added to the group. You can now work with the new component group. For example:
8. To view Component Graphs:
 - a. Highlight the new group and click **Details** at the top of the page.
 - b. Click **Component Graphs**, then check **All on Same Graph**.
Note: When using the All on Same Graph functionality, ensure that no more than 10 items are being displayed on the same graph for best usability.
9. To view events for the component group, click **EVENTS**.
10. To disable monitoring on one or more components:
 - a. Select the components you want to disable.
 - b. Click **Action > Monitoring**.
 - c. Click **NO** to disable monitoring.

Managing background tasks

The job service runs background tasks, such as discovering a network or adding a device. When you start one of these tasks, the job service adds a job to its queue.

Not all actions are performed in the job service. Some jobs are run automatically in the foreground. Others, such as moving devices, depend on user interface configuration settings.

When running jobs in the foreground, do not navigate away from the current page until the action completes.

Viewing the job manager

In Resource Manager, navigate to **ADVANCED > JOBS**.

The jobs list displays the following information about jobs:

- **Status**—The current status.
- **Description**—A description of the job.
- **Scheduled**—The time at which the job was scheduled to begin.
- **Started / Finished**—The beginning and ending timestamps of the job.
- **Created By**—The user who created the job.

The lower section of the page displays the log of the job that you select in the list. You can also view the information in the log file.

Stopping and deleting jobs

In Resource Manager, navigate to **ADVANCED > JOBS**.

To stop a job, select it in the list, and then click **Abort**.

To remove a job from the system, select it and then click **Delete**.

Using configuration properties

Configuration properties are individual values that you can set up on major system entities.


- Devices - *Device configuration properties* control the way that devices are monitored.
- Events - *Event configuration properties* control the rules that process events as they are received by the system.
- Networks - *Network configuration properties* control options that are used when you perform network discovery.

To customize the system when you add ZenPacks, you can add configuration properties and values to your ZenPacks.

- [Configuration properties inheritance and override](#)
- [Configuration property types](#)
- [Viewing and overriding device properties](#)
- [Viewing and overriding event properties](#)
- [Viewing and overriding network properties](#)

Configuration properties inheritance and override

The following diagram illustrates a portion of the standard device class hierarchy. A *device class* is a type of organizer that manages how the system models and monitors devices. At the root of the device hierarchy is the Devices object, under which all device class configuration properties are defined. Property values at the root level provide the default values for the entire hierarchy.

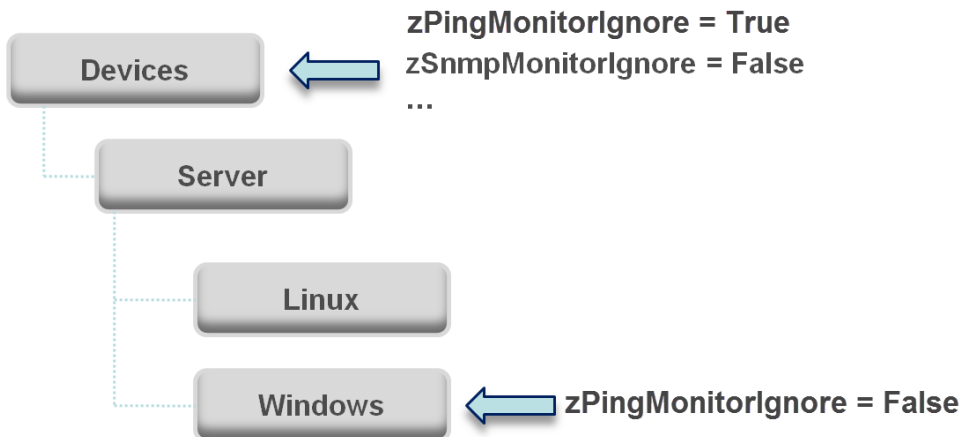


The illustration shows the following defined configuration properties:

- zPingMonitorIgnore- Turns off all daemons that use ping. By default, its value at the root of the hierarchy is False.
- zSNMPMonitorIgnore - Turns off all daemons that use SNMP. By default, its value at the root of the hierarchy is True.

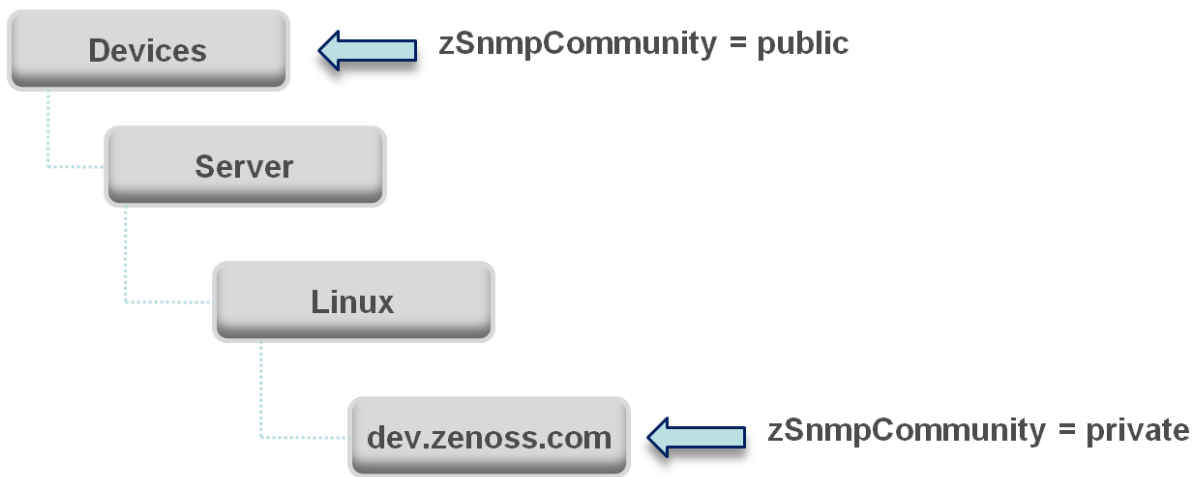
Through *inheritance*, properties that are defined at the root of the hierarchy apply to all objects beneath that node. So, at the /Devices/Server/Linux level of the device class hierarchy, the value of these two properties is the same as at /Devices, even though the property is not set explicitly at /Devices/Server/Linux. Inheritance simplifies system configuration because default values that are set at the root level apply to all devices regardless of their device class.

To further customize the system, you can change a specific configuration property at a lower level of the hierarchy without changing the definitions of other configuration properties. As shown in the following illustration, the value of zPingMonitorIgnore is changed so that ping monitoring is performed at the /Devices/Server/Windows level.



This locally defined value for zPingMonitorIgnore *overrides* the value that is set at the root of the hierarchy. No other properties at this level are affected by this local change; they continue to inherit the value that is set at the root.

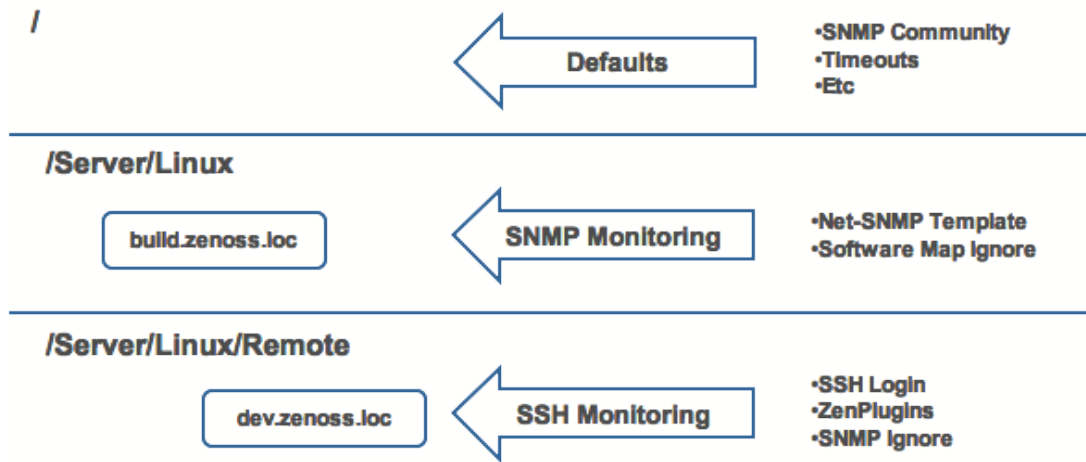
Configuration properties allow you to configure the system at a very granular level, down to a particular device. For example, in the following illustration, for the device named **dev.zenoss.com**, the value of SNMPCommunity set to private. This value overrides the root value of public.



If you change the SNMPCommunity value of **dev.zenoss.com** to public, it matches the value that is set at the root, but is still explicitly defined. Only if you remove the locally defined property does it again inherit the value of the property that is set at the root.

Inheritance in the device class tree

Inheritance is defined by how many attributes are applied to a device at different levels in the device hierarchy. The following diagram shows an example of how and where configuration properties can be set throughout the device class tree.



In this example, you can see that the default properties can be set at the highest level (/). However, as you travel further down the hierarchy, you see that you can override any of the configuration properties that are set at the root level.

The next two lines show how the device tree further defines properties for Linux servers. For example, to set up and use SNMP monitoring for all Linux servers (inclusive of) build.zenoss.loc, you could change these properties at the /Server/Linux level.

Further, if you wanted to change how you collect information for remote Linux servers, you could create a sub-group in /Server/Linux called /Server/Linux /Remote, setting these servers to use SSH monitoring and changing the associated properties for that sub-group.

All of these configuration properties and groupings co-exist, with any changes made lower in the hierarchy taking priority.

Configuration property types

Configuration properties are one of the following types:

- **String** - Text value that can be ASCII or Latin-1 encoded
- **Integer** - Whole number
- **Float** - Number that can have a decimal value
- **Boolean** - True or False
- **Lines** - List of values separated by a return. The system stores these as an array.
- **Password** - Character-masked password value

Viewing and overriding device properties

This section further illustrates the characteristics of configuration properties from the perspective of the Resource Manager interface.

1. From the navigation bar, choose INFRASTRUCTURE.
2. In the tree view, click DETAILS > Configuration Properties. The Configuration Properties page for the selected device class appears. The following figure shows device configuration properties that are defined at the root level. The zCollectorClientTimeout configuration property has a default value of 180.

Is Local	Category	Name	Value	Label	Description
Yes	Cisco UCS	zCiscoUCSCIMCPerfInterval	300	CIMC Stats Collection Interval	Seconds between CIMC s...
Yes	Cisco UCS	zCiscoUCSCIMCSSLProtocol	SSLv23	CIMC SSL/TLS Protocol	The SSL/TLS protocol use...
Yes	Cisco UCS	zCiscoUCSManagerPassword	*****	Manager Password	Password used to connect...
Yes	Cisco UCS	zCiscoUCSManagerPerfInterval	300	Manager Stats Collection Interval	The interval is used only...
Yes	Cisco UCS	zCiscoUCSManagerPort	443	Manager Port	Port used to connect to U...
Yes	Cisco UCS	zCiscoUCSManagerUseSSL	true	Manager SSL	Whether to use SSL when conn...
Yes	Cisco UCS	zCiscoUCSManagerUser	admin	Manager User	User used to connect to U...
Yes	Modeler Contr...	zCollectorClientTimeout	180	Collector Client Timeout (seconds)	Allows you to set the time...
Yes	Modeler Contr...	zCollectorDecoding	utf-8	Collector Decoding	Converts incoming chara...
Yes	Misc	zCollectorLogChanges	false	Log Collector Changes?	Indicates whether to log...
Yes	zencommand	zCommandCollectionInterval	300	Command Collection Interval	Defines, in seconds, the d...
Yes	zencommand	zCommandCommandTimeout	15	Timeout for Commands (seconds)	Specifies the time to wait...
Yes	zencommand	zCommandExistenceTest	test -f %s	Command Existence Test	
Yes	zencommand	zCommandLoginTimeout	10	Timeout for Login (seconds)	Specifies the time to wait...

In the following figure, the zCollectorClientTimeout configuration property value is set to 170 at the /Server/Linux device class. The device class value overrides the default value at this node of the hierarchy.

Is Local	Category	Name	Value	Label	Description
	Cisco UCS	zCiscoUCSCIMCPerfInterval	300	CIMC Stats Collection Interval	Seconds between CIMC statistic...
	Cisco UCS	zCiscoUCSCIMCSSLProtocol	SSLv23	CIMC SSL/TLS Protocol	The SSL/TLS protocol used to c...
	Cisco UCS	zCiscoUCSManagerPassword	*****	Manager Password	Password used to connect to U...
	Cisco UCS	zCiscoUCSManagerPerfInterval	300	Manager Stats Collection Interval	The interval is used only when t...
	Cisco UCS	zCiscoUCSManagerPort	443	Manager Port	Port used to connect to UCS M...
	Cisco UCS	zCiscoUCSManagerUseSSL	true	Manager SSL	Whether to use SSL when conn...
	Cisco UCS	zCiscoUCSManagerUser	admin	Manager User	User used to connect to UCS M...
Yes	Modeler Contr...	zCollectorClientTimeout	170	Collector Client Timeout (seconds)	Allows you to set the timeout ti...
	Modeler Contr...	zCollectorDecoding	utf-8	Collector Decoding	Converts incoming characters t...
	Misc	zCollectorLogChanges	false	Log Collector Changes?	Indicates whether to log chang...
	zencommand	zCommandCollectionInterval	300	Command Collection Interval	Defines, in seconds, the default...
	zencommand	zCommandCommandTimeout	15	Timeout for Commands (seconds)	Specifies the time to wait for a ...

- To remove the override and once again inherit the value from the root of the hierarchy:
 - Select the property in the list.
 - Click Delete Local Copy and, when prompted, click OK.

List of device configuration properties

Name	Type	Description
zAWSAutoChangeProdState	boolean	Whether to enable or disable auto change of the production state for EC2 instances. Default is true (enabled).
zAWSAutoChangeProdStateRunning	int	By default, the state is changed to <i>Production</i> (1000) for running EC2 instances. Specify a custom state ID number.
zAWSAutoChangeProdStateStopped	int	By default, the production state is changed to <i>Decommissioned</i> (-1) for stopped EC2 instances. Specify a custom state ID number.
zAWSBillingCostThreshold	int	Sets the spending threshold and triggers an event if spending exceeds the value. Default is 1000.
zAWSCloudFormationEventsAutoClear	boolean	If set to True, for each CREATE_COMPLETE and DELETE_COMPLETE, the corresponding auto-clear event is generated to clear previous CRITICAL events.
zAWSCloudWatchMaxParallel	int	Specify the number of concurrent CloudWatch calls.
zAWSCloudWatchMaxRetries	int	Specify the number of retries for CloudWatch calls.
zAWSCloudWatchSSL	boolean	Whether to use SSL when connecting to AWS CloudWatch API.
zAWSDiscover	awsdiscoverfield	To configure instance guest device discovery, for the instance, specify the <tag:value>. Guest device discovery must be configured individually for each EC2 account.
zAWSEnableSnapshotCollection	boolean	Whether to disable the collection of snapshots, thereby improving modeling performance.
zAWSGuestCollector	string	Specify the name of the collector that all discovered devices for this AWS device use.
zAWSGuestDeviceTitleTag	string	To populate the guest devices' titles based on an AWS tag from the instance, set the tag name.
zAWSGuestUsePublicIPs	boolean	Whether to use a public IP address for creating guest devices.
zAWSRegionPEM	multilinekeypath	Set the region name and path to PEM file for use in auto-discovering instance guest operating systems.
zAWSRegionToModel	string	Narrow components that are modeled. By default, all EC2 regions and child components are discovered. Specify EC2 region name or multiple names separated by commas.
zAWSRemodelEnabled	boolean	When disabled, only the instance state is updated on existing instances. When enabled, all instance properties are updated on existing instances, and new instances are added to the model.
zAWSResetGuestCollector	boolean	Applies to guest devices (not EC2 account). Specifies whether to change the collector if you have set it manually.
zAggregatorCollectionInterval	int	Aggregator collection interval in seconds. Default is 300.
zAzureClientID	string	To model and monitor resources that were created in Azure Resource Manager Deployment Model (ARM) or Azure Cloud Solution Provider (CSP), the account must be available to authenticate requests using Azure service principal. Specify the client ID number.
zAzureEAAccessKey	string	To enable billing data collection for Enterprise Accounts, you must specify the access key and enrollment number, and the AzureEABillingDataSourcePlugin data source plugin must be attached to monitoring template /Azure /AzureSubscription.
zAzureEABillingCostThreshold	multilinetreshold	Sets the spending threshold and triggers an event if spending exceeds the value.

zAzureEAEnrollmentNumber	string	To enable billing data collection for Enterprise Accounts, you must specify the access key and enrollment number, and the AzureEABillingDataSourcePlugin data source plugin must be attached to monitoring template /Azure /AzureSubscription.
zAzureMonitoringIgnore	string	Specify a Python expression that returns a boolean value, which is evaluated against each device component. If the result is True, then the component is not monitored.
zAzureSecretKey	string	To model and monitor resources that were created in Azure Resource Manager Deployment Model (ARM) or Azure Cloud Solution Provider (CSP), the account must be available to authenticate requests using Azure service principal. Specify the secret key.
zAzureTenantID	string	To model and monitor resources that were created in Azure Resource Manager Deployment Model (ARM) or Azure Cloud Solution Provider (CSP), the account must be available to authenticate requests using Azure service principal. Specify the tenant ID number.
zCiscoACEUseSSL	boolean	Whether to use SSL when connecting to ACE XML API.
zCiscoNXAPIInterval	int	Value in seconds of the rediscovery interval using the Cisco NX-API protocol. Default value is 300.
zCiscoNXAPIPort	int	Port for connecting to NX-API.
zCiscoNXAPIUseSSL	boolean	Whether to use SSL when connecting to NX-API.
zCiscoRemodelEventClassKeys	lines	Allows you to modify the list of SNMP traps that will cause the product schedule an immediate remodeling of the device from which the trap was sent.
zCiscoUCSCIMCEventsInterval	int	Event collection interval in seconds. Default is 60.
zCiscoUCSCIMCPerfInterval	int	Metric collection interval in seconds. Default is 300.
zCiscoUCSCIMCSSLProtocol	string	The SSL/TLS protocol used to connect to a CIMC device.
zCiscoUCSManagerPassword	password	Password for UCS Manager user name.
zCiscoUCSManagerPerfInterval	int	Seconds between UCS Manager statistics collections. Default value is 300.
zCiscoUCSManagerPort	int	Port used to connect to the UCS Manager or CIMC XML APIs. Default is 443 and typically should not be changed.
zCiscoUCSManagerUseSSL	boolean	Whether to use SSL when connecting to the UCS Manager or CIMC XML APIs. Default is true and typically should not be changed.
zCiscoUCSManagerUser	string	UCS Manager user name.
zCollectorClientTimeout	int	Allows you to set the timeout time of the collector client in seconds
zCollectorDecoding	string	Converts incoming characters to Unicode.
zCollectorLogChanges	boolean	Indicates whether to log changes.
zCommandCollectionInterval	int	The default collection interval (number of seconds) for command datasources.
zCommandCommandTimeout	float	Specifies the time to wait for a command to complete.
zCommandLoginTimeout	float	Specifies the time to wait for a login prompt.
zCommandLoginTries	int	Sets the number of times to attempt login.
zCommandPassword	password	Specifies the password to use when performing command logins and SSH.
zCommandPath	string	Sets the default path where ZenCommand plug-ins are installed on Resource Manager (or on a remote box where SSH is used to run the command).
zCommandPort	int	Specifies the port to connect to when performing command collection.

zCommandProtocol	string	Establishes the protocol to use when performing command collection. Possible values are SSH and telnet.
zCommandSearchPath	lines	Sets the path to search for any commands.
zCommandUserCommandTimeout	float	Specifies the number of seconds to wait for a user command to complete.
zCommandUsername	string	Specifies the user name to use when performing command collection and SSH.
zControlCenterHost	string	Control Center host name. Defaults to device name.
zControlCenterModelCycle	int	Control Center modeling interval. Defaults to 3600s.
zControlCenterPassword	password	Password for the Control Center user name.
zControlCenterPerfCycle	int	Control Center performance collection interval. Defaults to 300s.
zControlCenterPort	int	Port for Control Center. Defaults to HTTPS TCP/443
zControlCenterUser	string	Control Center user name.
zDBInstances	***	***instancecredentials. This setting is only relevant when the zenoss.winrm.WinMSSQL modeler plugin is enabled. Multiple instances can be specified to monitor multiple SQL Server instances per server. The default instance is MSSQLSERVER. Fill in the user and password to use SQL authentication. Leave the user and password blank to use Windows authentication.
zDatasourceDebugLogging	boolean	True or False to debug calculated/aggregated data sources on a single device.
zDeviceTemplates	lines	Sets the templates associated with this device. Linked by name.
zDockerMonitorContainerSize	boolean	True or false to monitor the real size and virtual size of each container once every ten minutes. The default is false (not enabled) because the command takes a long time to run when many containers, or large containers are used.
zDockerMonitorContainerStats	boolean	True or false to collect the statistics of each container once every five minutes. The default is true (enabled).
zDockerMonitorContainerStatus	boolean	True or false to check the status of the Docker daemon once per minute. The default is true (enabled).
zDiscoveryMappingOn	boolean	True or False to determine whether discovered devices are automatically moved to device classes that match their hardware and operating system. If False (disabled), devices are not moved from the /Discovered device class.
zEnablePassword	boolean	True or False to specify use of password for Cisco routers.
zFileSystemMapIgnoreNames	string	Sets a regular expression of file system names to ignore.
zFileSystemMapIgnoreTypes	lines	Do not use.
zFileSystemSizeOffset	int	SNMP typically reports the total space available to privileged users. Resource Manager (like the df command) reports capacity based on the space available to non-privileged users. The value of zFileSystemSizeOffset should be the fraction of the total space that is available to non-privileged users. The default reserved value is 5% of total space, so zFileSystemSizeOffset is preset to .95. If the reserved portion is different than 5%, then adjust the value of zFileSystemSizeOffset accordingly. The fraction should be set according to the value $(Used + Avail) / Size$ when the df -PkH command is run at the command line.
zHardDiskMapMatch	string	Regular expression that uses the disk ID in the diskstats output to filter disk activity statistics for inclusion in performance monitoring.
zHyperVDiscoverGuests	boolean	True or false to discover virtual machine guest devices. Guest device discovery must be configured individually for each Hyper-V server. The default is false.

zILOCollectSamples	boolean	True or false to collect and save raw data for debug purposes.
zILOPassword	password	The ILO authentication password. Data from ILO devices is provided by a HTTPS interface listening on port 443. The following properties control access to ILO devices: zILOUserName, zILOPassword, zILOUseSSL, and zILOPort.
zILOPort	string	The TCP port for ILO communication. Data from ILO devices is provided by a HTTPS interface listening on port 443. The following properties control access to ILO devices: zILOUserName, zILOPassword, zILOUseSSL, and zILOPort.
zILOUseSSL	boolean	True or false to use HTTPS for ILO communication. Data from ILO devices is provided by a HTTPS interface listening on port 443. The following properties control access to ILO devices: zILOUserName, zILOPassword, zILOUseSSL, and zILOPort.
zILOUserName	string	The ILO user name. Data from ILO devices is provided by a HTTPS interface listening on port 443. The following properties control access to ILO devices: zILOUserName, zILOPassword, zILOUseSSL, and zILOPort.
zIcon	lines	Specifies the icon to represent the device wherever device icon is shown, such as on the network map and device status page.
zIdiomPassword	password	IDIOM API password for Cisco IDS/IPS devices.
zIdiomUsername	string	IDIOM API username for Cisco IDS/IPS devices.
zIfDescription	boolean	Shows the interface description field in the interface list.
zIgnoreUnmounted	boolean	True or false to specify whether to ignore unmounted drives.
zInterfaceMapIgnoreDescriptions	string	Filters out interfaces based on description.
zInterfaceMapIgnoreNames	string	Filters out interfaces that should not be discovered. If you want to use an expression to define this property, note that only Python regular expressions are valid.
zInterfaceMapIgnoreTypes	string	Filters out interface maps that should not be discovered.
zIpServiceMapMaxPort	int	Specifies the highest port to scan. The default is 1024.
zJmxAuthenticate	boolean	True or False to enable/disable authentication.
zJmxManagementPort	int	Port that enables JMX management
zJmxPassword	password	JMX username password
zJmxUsername	string	JMX username
zKeyPath	string	Sets the path to the SSH key for device access.
zLDMsAutodiscover	boolean	Specify true or false for auto-discovery of LDOMs.
zLTMVirtualServerIgnoreNames	string	Regular expression that can be used to prevent matching LTM Virtual Servers from being modeled.
zLinks	string	Specifies a place to enter any links associated with the device.
zLocalInterfaceNames	string	Regular expression that uses interface name to determine whether the IP addresses on an interface should be incorporated into the network map. For instance, a loopback interface "lo" might be excluded.
zLocalIpAddresses	string	Specifies IP addresses that should be excluded from the network map (for example, 127.x addresses). If you have addresses that you reuse for connections between clustered machines, you might add those addresses.
zMaxOIDPerRequest	int	Sets the maximum number of OIDs to be sent by the SNMP collection daemons when querying information. Some devices have small buffers for handling this information, so a lower value might be appropriate.

zMySQLConnectionString	multilinecredentials	Specifies a list of properties (user, password, and port) for MySQL connections.
zMySQLPassword	password	MySQL user password
zMySQLPort	string	MySQL connection port
zMySQLTimeout	int	MySQL timeout. Default value is 30s.
zMySQLUsername	string	MySQL user name.
zNetAppNumRecordsPerRequest	int	Number of records per NetApp Monitor request. Default is 100.
zNetAppSSL	boolean	Boolean true or false to enable SSL use with the NetApp Monitor.
zNmapPortscanOptions	string	Options used on nmap when scanning ports. Used in IpServiceMap.
zNodesAutodiscover	boolean	Specify true or false for auto-discovery of nodes.
zPingMonitorIgnore	boolean	Specify whether to ping the device.
zProdStateThreshold	int	Production state threshold at which Resource Manager will begin to monitor a device.
zPropertyMonitorInterval	int	Polling interval of the configured property data sources. System-wide setting. Default is 300s.
zPythonClass	string	DO NOT USE
zRMMonCCHost	string	Host name or IP address of the Control Center instance on which the Resource Manager to monitor is running.
zRMMonCCUser	string	User name of the Control Center instance on which Resource Manager is running.
zRMMonCCPassword	password	Password for the Control Center user name.
zRMMonCCPort	int	Port for Control Center. Defaults to HTTPS TCP/443.
zRMMonRabbitPassword	password	Password for the RabbitMQ user name. Defaults to the system default.
zRMMonRabbitUser	string	User name for RabbitMQ. Defaults to the system default.
zRMMonTenantHost	string	Resource Manager host name.
zRMMonTenantPassword	password	Password for the Resource Manager user name.
zRMMonTenantPerfCycle	int	Performance collection interval. Defaults to 30s.
zRMMonTenantPort	int	Port for Resource Manager. Defaults to HTTPS TCP/443.
zRMMonTenantUseSsl	boolean	True or false value to use SSL. Defaults to true.
zRMMonTenantUser	string	Resource Manager user name.
zRouteMapCollectOnlyIndirect	boolean	Only collect routes that are indirectly connected to the device.
zRouteMapCollectOnlyLocal	boolean	Only collect local routes. (These usually are manually configured rather than learned through a routing protocol.)
zRouteMapMaxRoutes	int	Sets maximum number of routes to collect. Default value is 500.
zSnmpAuthPassword	password	The shared private key used for authentication. Must be at least 8 characters long.
zSnmpAuthType	string	Use "MD5" or "SHA" signatures to authenticate SNMP requests
zSnmpCollectionInterval	int	Defines, in seconds, how often the system collects performance information for each device.
zSnmpCommunities	lines	Array of SNMP community strings that ZenModeler uses when collecting SNMP information. When you set this property, communities are tried in order; the first in the list

		that is successful is used as zSnmCommunity. If none is successful, then the current value of zSnmCommunity is used. The default value for the entire system is "public."
zSnmCommunity	string	Community string to be used when collecting SNMP information. If it is different than what is found by ZenModeler, it will be set on the modeled device.
zSnmContext	string	Configures zSNMP context to map logical network entity such as a topology or protocol instance.
zSnmDiscoveryPorts	int	List of UDP ports to try when performing SNMP discovery. Defaults to 161 if notset.
zSnmEngineId	string	SNMPv3 engine ID for the device. Will be discovered when SNMPv3 is used.
zSnmMonitorIgnore	boolean	Whether or not to ignore monitoring SNMP on a device.
zSnmPort	int	Port that the SNMP agent listens on.
zSnmPrivPassword	password	The shared private key used for encrypting SNMP requests. Must be at least 8 characters long.
zSnmPrivType	string	"DES" or "AES" cryptographic algorithms.
zSnmSecurityName	string	The Security Name (user) to use when making SNMPv3 requests.
zSnmTimeout	float	Timeout time in seconds for an SNMP request
zSnmTries	int	Amount of tries to collect SNMP data
zSnmVer	string	SNMP version used. Valid values are v2c, v1
zSshConcurrentSessions	int	Maximum number of sessions supported by the remote device's MAX_SESSIONS parameter. Common values for AIX are 2 or 10.
zStatusConnectTimeout	float	The amount of time that the zenstatus daemon should wait before marking an IP service down.
zSysedgeDiskMapIgnoreNames	string	Regular expression used by zenoss.snmp.SysedgeDiskMap modeler plugin. Disks with matching names will not be modeled.
zTelnetEnable	boolean	When logging into a Cisco device issue the enable command to enable access during command collection.
zTelnetEnableRegex	string	Regular expression to match the enable prompt.
zTelnetLoginRegex	string	Regular expression to match the login prompt.
zTelnetPasswordRegex	string	Regular expression to match the password prompt.
zTelnetPromptTimeout	float	Time to wait for the telnet prompt to return.
zTelnetSuccessRegexList	lines	List of regular expressions to match the command prompt.
zTelnetTermLength	boolean	On a Cisco device, set term length to Zero.
zUsesManagelp	boolean	True or False to specify use of manage IP. Used to avoid setting manage IP when a device is added. Default is True for most device classes. For the /ZenossRM device class, default is False to allow modeling many Resource Manager instances per Control Center without manage IP collisions.
zUsesStandardDeviceCreationJob	boolean	True or False to specify whether to use the standard device creation job.
zVSphereEndpointHost	string	vSphere host name
zVSphereEndpointPassword	password	vSphere username password
zVSphereEndpointPort	int	Port that is used to connect to vSphere Endpoint.
zVSphereEndpointUseSsl	boolean	vSphere boolean true or false for SSL use
zVSphereEndpointUser	string	vSphere username


zVSphereHostCollectionClusterWhitelist	lines	The whitelist filters the hosts that are monitored, based on cluster names.
zVSphereHostPingBlacklist	lines	List of regular expressions to control which management IP address to ping (matches against hostname:nicname:ip). Note: zVSphereHostPingWhitelist takes precedence over zVSphereHostPingBlacklist.
zVSphereHostPingWhitelist	lines	List of regular expressions to control which management IP address to ping (matches against hostname:nicname:ip). Note: zVSphereHostPingWhitelist takes precedence over zVSphereHostPingBlacklist.
zVSphereHostSystemPassword	password	Password that is used to access ESX hosts via ssh and API.
zVSphereHostSystemUser	string	User name that is used to access ESX hosts via ssh and API.
zVSphereLUNContextMetric	boolean	Controls whether to use LUN-specific metric names when storing performance data. Note: The default value is False. Changing the value to True causes historical metrics to become inaccessible. For more information, customers may contact Zenoss Support.
zVSphereModelCache	lines	vSphere model cache
zVSphereModelIgnore	lines	vSphere model ignore
zVSphereModelMpIndexObjs	int	Advanced tuning parameter. For more information, customers may contact Zenoss Support.
zVSphereModelMpLevel	int	Advanced tuning parameter. For more information, customers may contact Zenoss Support.
zVSphereModelMpObjs	int	Advanced tuning parameter. For more information, customers may contact Zenoss Support.
zVSpherePerfDelayCollectionMinutes	int	Value of how long to lag performance data collection. Default value is 0.
zVSpherePerfMaxAgeMinutes	int	Default value is 28 minutes.
zVSpherePerfParallelQueries	int	Default value is 6.
zVSpherePerfQueryChunkSize	int	Value of how many performance requests to make at a time. Default value is 250.
zVSpherePerfQueryRaw20	boolean	Default value is true.
zVSpherePerfQueryTimeout	int	Default value is 200.
zVSpherePerfQueryVcChunkSize	int	Default value is 64.
zVSpherePerfQueryVcRaw20	boolean	Default value is false.
zVSpherePerfRecoveryMinutes	int	Default value is 240 minutes.
zVSpherePerfTimeoutRecoveryMinutes	int	When a timeout error occurs in querying a specific metric, it is "blacklisted" for this number of minutes before it is retried. This action avoids repeated errors due to attempts to query a metric that is not working properly. The default is one hour. If gaps appear in the graphs, you can safely lower the value.
zVSphereVMContextMetric	boolean	Controls whether to use VM-specific metric names when storing performance data. Note: The default value is False. Changing the value to True causes historical metrics to become inaccessible. For more information, customers may contact Zenoss Support.
zWBEMPassword	password	WBEM password
zWBEMPort	int	Value of the WBEM port number. Default value is 5989.
zWBEMUseSSL	boolean	True or false value to use SSL. Default value is true.
zWBEMUsername	string	WBEM username
zWebTxAgent	string	Default value is ZenWebTx/1.0.

zWebTxPassword	password	WebTx password
zWebTxRealm	string	WebTx realm
zWebTxUser	string	WebTx user
zWebSphereServer	string	Used by the provided template to build the xpath queries for the data to collect. You must supply a value for this field. For example: <i>serverAB</i> . There is no default value.
zWebSphereURLPath	string	Path to the PMI servlet on a WebSphere instance. The default value is the default path on a WebSphere installation: <i>wasPerTool/servlet/perfservlet</i>
zWinKDC	string	IP address or fully-qualified domain name of a valid Windows domain controller. Must be set if domain authentication is used.
zWinKeyTabFilePath	string	This property is currently used and reserved for future use when keytab files are supported.
zWinPerfmonInterval	int	Interval, in seconds, at which Windows Perfmon datapoints will be collected. Default value is 300. It is possible to override the collection interval for individual counters.
zWinRMClusterNodeClass	string	Path under which to create cluster nodes.
zWinRMEnvelopeSize	int	Used when WinRM configuration setting "MaxEnvelopeSizekb" exceeds default of 512k.
zWinRMKrb5DisableRDNS	boolean	Set to true to disable reverse DNS lookups by Kerberos. Only set at /Server/Microsoft level.
zWinRMKrb5includedir	string	Directory path for Kerberos configuration files.
zWinRMLocale	string	Communication locale to use for monitoring. Reserved for future use.
zWinRMPassword	password	Password for the user defined by zWinRMUser
zWinRMPort	int	The port on which the Windows server is listening for WinRM or WS-Management connections. Default value is 5985. It is uncommon for this to be configured as anything else.
zWinRMServerName	string	<p>This property should only be used in conjunction with domain authentication when the DNS PTR record for a monitored server's managed IP address does not resolve to the name by which the server is known in Active Directory. For example, if <i>myserver1</i> is known as <i>myserver1.ad.example.com</i> by Active Directory and is being managed by IP address 192.51.100.21, but 192.51.100.21 resolves to <i>www.example.com</i>, you will have to set <i>zWinRMServerName</i> to <i>myserver1.ad.example.com</i> for domain authentication to work.</p> <p>If many Windows servers in your environment do not have DNS PTR records that match Active Directory, it is recommended that you set the name of the Resource Manager device to be the fully-qualified Active Directory name and set <i>zWinRMServerName</i> to <code>\${here /titleOrId}</code> at the /Server/Microsoft/Windows device class. This avoids the necessity of setting <i>zWinRMServerName</i> on every device.</p> <p>It is recommended to leave <i>zWinRMServerName</i> blank if local authentication is used, or DNS PTR records match Active Directory. This allows Resource Manager to not rely on DNS resolution while monitoring, and avoids the overhead of configuring <i>zWinRMServerName</i>.</p>
zWinRMUser	string	The syntax used for <i>zWinRMUser</i> controls whether Resource Manager will attempt Windows local authentication or domain (Kerberos) authentication. If the value of <i>zWinRMUser</i> is <i>username</i> , local Windows authentication will be used. If <i>zWinRMUser</i> is <i>username@example.com</i> , domain authentication will be used. The <i>zWinKDC</i> and potentially the <i>zWinRMServerName</i> properties become important.

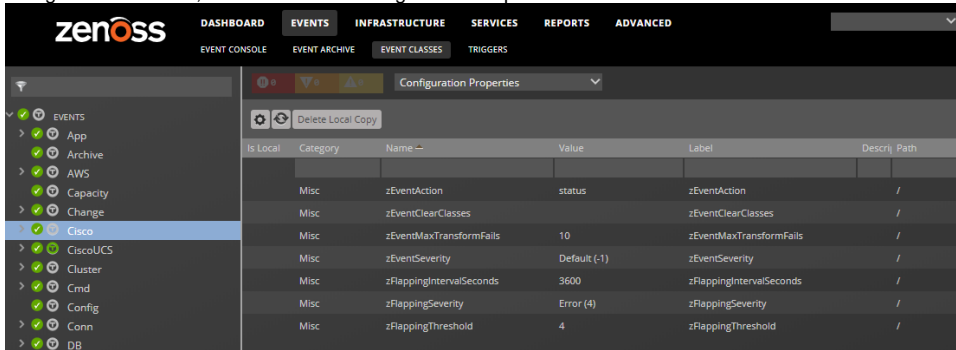
zWinRSCodePage	int	Code page used by monitoring user account.
zWinScheme	string	Must be set to http or https. Default value is http.
zWinTrustedKDC	string	Windows Trusted KDC.
zWinTrustedRealm	string	Windows Trusted Realm.
zWinUseWsmanSPN	boolean	Set to true if HTTP/HTTPS service principles are exclusively for use by a particular service account.
zWSMANPassword	password	Password for WSMAN user name.
zWSMANPort	int	The port number used to gather WSMAN information. Defaults to 443.
zWSMANUseSSL	boolean	True or false value to use SSL. Default value is true.
zWSMANUsername	string	WSMAN user name.

Viewing and overriding event properties


1. From the navigation menu, choose EVENTS > Event Classes.
2. From the drop-down list, choose Configuration Properties.



3. To view and override event configuration properties for a specific event class:
 - a. Navigate to that class, and then select Configuration Properties.



- b. Double-click the configuration property that you want to change. The presence of Yes in the Is Local column indicates an overriding value.
- c. In the Edit Config Property dialog box, make changes and then click SUBMIT.



4. To remove an override and once again inherit the value from the root of the hierarchy:
 - a. Select the property in the list.
 - b. Click Delete Local Copy and, when prompted, click OK.

Table of event configuration properties

Name	Type	Description
zEventAction	string	Specifies the database table in which an event will be stored. Possible values are: status, history and drop. Default is status, meaning the event will be an active event. History sends the event directly to the history table. Drop tells the system to discard the event.
zEventClearClasses	lines	Lists classes that a clear event should clear (in addition to its own class).
zEventMaxTransformFails	int	After the specified number of failures, disable bad transforms from executing. Default is 10.
zEventSeverity	int	Overrides the severity value of events from this class. Possible values are 5 (Critical), 4 (Error), 3 (Warning), 2 (Info), 1 (Debug), 0 (Clear), and -1 (Default).
zFlappingIntervalSeconds	int	Defines the time interval to check for event flapping (changing severity level repeatedly). Default value is 3600 seconds.
zFlappingSeverity	int	Drop-down list to set the severity to check for event flapping. If the severity level on an event changes from this value a certain number of times (zFlappingThreshold) within a certain time range (zFlappingIntervalSeconds) then an event flapping event is generated. Possible values are 5 (Critical), 4 (Error), 3 (Warning), 2 (Info), 1 (Debug), and 0 (Clear).
zFlappingThreshold	int	Number of times an event severity must flap within an interval. One of the parameters to define in order to generate event flapping events.


A good example of how the system uses the event class configuration properties is found in the /Status event class. Within the /Status event class configuration properties, zEventAction is set to "history" and zEventSeverity is set to "Default". This means that events sent with this event class are sent into the active events table with an initial state of closed, and the event severity unchanged.

For more information about event manipulation techniques, see [Event mapping and transforms](#).

Viewing and overriding network properties

The Networks page lists networks by IP address. You can view and change network configuration property values and inheritance selections.

1. From the navigation menu, choose INFRASTRUCTURE > NETWORKS.
2. From the Display drop-down menu, select Configuration Properties.



3. To remove an override and once again inherit the value from the root of the hierarchy:
 - a. Select the property in the list.
 - b. Click Delete Local Copy and, when prompted, click OK.

List of network configuration properties

Name	Type	Description
zAutoDiscover	boolean	Specifies whether the zendisc service should perform auto-discovery on this network. (When performing network discovery, this property specifies whether the system should discover devices and subnets on the network.)
zDefaultNetworkTree	lines	A network subnet is automatically created for each modeled device, based on that device's subnet mask setting. To create higher-level subnets automatically from the discovery and modeling processes, add the specific subnet mask breakpoints. For example: 8, 16. If you then model a device with, for example, an IP address of 192.0.2.0, and a subnet mask of 255.255.255.0 (corresponding to a /24 subnet), device discovery will create a 192.0.0.0/8 network containing 192.0.2.0/16, containing 192.0.2.0/24, containing your device.
zDrawMapLinks	boolean	Calculating network links "on the fly" is resource-intensive. If you have a large number of devices that have been assigned locations, then drawing those map links may take a longtime. You can use this property to prevent the system from drawing links for specific networks (for example, a local network comprising many devices that you know does not span multiple locations).
zlcon	string	Use to specify device icons that appear on the device status page, Dashboard, and network map.
zPingFailThresh	int	Specifies the number of pings sent without being returned before the zendisc service removes the device.
zPreferSnmpNaming	boolean	Specifies that when network discovery occurs, it uses the device name comes from SNMP rather than reverse DNS.
zSnmpStrictDiscovery	boolean	Specifies that if SNMP does not exist on the device during network discovery, ignore the device.

Modeling

Modeling is the process of discovering device information. It is not monitoring; modeling does not collect availability or performance data.

- [Modeling devices](#)
- [About modeler plugins](#)
- [Debugging the modeling process](#)

Modeling devices

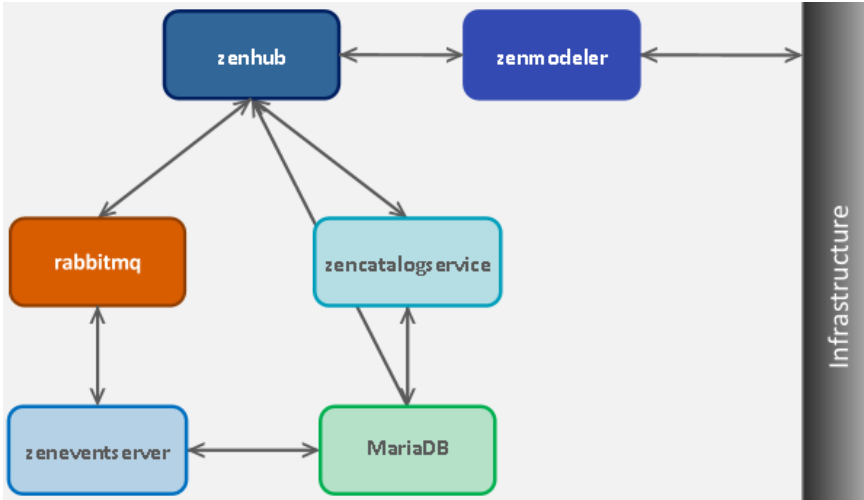
Resource Manager can use SSH, WinRM, and SNMP to model devices (SSH and WinRM are the recommended options).

The modeling method that you select depends on your environment and the types of devices that you want to model and monitor.

By default the system remodels each known device every 720 minutes (12 hours). You can change the frequency with which devices are remodeled. Edit the value of the Modeler Cycle Interval in the collector's configuration.

For larger deployments, modeling frequency might affect performance. In such environments, set the `startat` configuration setting inside the `zenmodeler.conf` file to change the scheduling of the daemon. The `startat` value only dictates the initial start time of `zenmodeler`. Each subsequent run interval is determined by the `zenmodeler` cycle time (number of minutes between runs). The cycle time is configured on the daemon settings page inside the parent's collector folder, which you can access in Control Center. For more information, refer to [this knowledgebase article](#).

The following figure shows how the `zenmodeler` daemon fits into the modeling devices portion of the Resource Manager architecture. Note that MariaDB replaced ZenDS (MySQL) in version 5.x and later and stores the object database (ZODB).



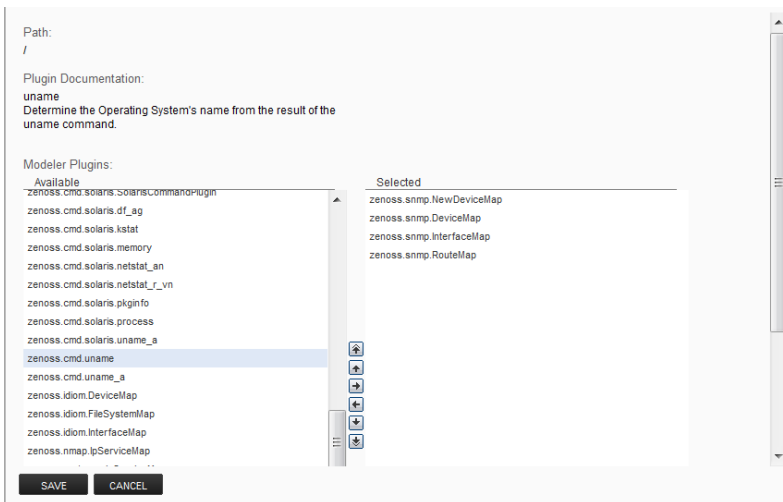
Modeling devices using SSH/COMMAND

You can gather additional information by running commands on the remote device and interpreting the results. This provides a more scalable and flexible way to gather information that may not be available through any other means.

Each built-in modeling command plugin is differentiated by the platform on which it runs. To determine the platform for the device you want to model, run the `uname` command in a shell on the device.

To model a device using command plugins, first add the device by using the protocol "none," and then choose the plugins you want to apply:

1. From the Navigation menu, select Infrastructure.
2. Click the Add Devices icon and select Add a Single Device from the drop-down list. The Add a Single Device window appears.
3. Enter values for Name or IP and Device Class.
4. Clear the Model Device option.
5. Click Add.
6. After adding the device, select the device name in the devices list. The Device Overview page appears.
7. In the left panel, select Configuration Properties.
8. If necessary, set the values of the `zCommandUsername` and `zCommandPassword` configuration properties to the user name and password of the device (or set up authentication by using RSA/DSA keys.)
Note: If using RSA keys for a device or device class, change the value of the `zKeyPath` configuration property to `"~/.ssh/id_rsa"`.
9. In the left panel, select Modeler Plugins. The list of plugins appears. The left column displays available plugins; the right column shows those currently selected.
10. Select `zenoss.cmd.uname` from the Available list, and then use the right arrow control to move it to the Selected list on the right. Use the controls to place it at the top of the list.
11. Use the left arrow control to move the other Selected plugins from the Selected list to the Available list.



12. Click Save.
13. Model the device by clicking the Model Device button.

Modeling devices using port scan

You can model IP services by doing a port scan, using the [Nmap Security Scanner](#).

To model a device using a port scan:

1. Select the device in the device list.
2. In the left panel, select Modeler Plugins.
3. Select the zenoss.nmap.ipServiceMap plugin in the list of Available plugins, and then use the right arrow control to move it to the list of Selected plugins.
4. Click Save.
5. Remodel the device by clicking the Model Device button.

Using the /Server/Scan device class to monitor with port scan

The /Server/Scan device class is an example configuration for modeling devices by using a port scan. You can use this device class as a reference for your own configuration; or, if you have a device that will use only a port scan, you can place it under this device class and remodel the device.


About modeler plugins

Resource Manager uses plug-in maps to map real world information into the standard model. Input to the plug-ins can come from a number of different sources (SNMP, SSH, WinRM, etc). Selection of plug-ins to run against a device is done by matching the plug-in name against the bound modeler plugins.

Viewing and editing modeler plugins

To view a list of plugins for any device:

1. Click the device name in the devices list.
2. In the device summary page, select Modeler Plugins.



The Modeler Plugins page appears.

Modeler plugin actions

Action	Steps
Add	To add a plugin: <ol style="list-style-type: none"> 1. Use the right arrow control to move one or more plugins from the Available list (on the left) to the Selected list (on the right). 2. Click Save.
Reorder	Plugins run in the order in which they are listed. To re-order plugins, use the up and down arrow controls, and then click Save .
Delete	To delete a plugin from a device, use the left arrow control to move the plugin from the Selected list to the Available list.

Bound modeler plugins can be viewed and managed by browsing to a device or device class, and choosing the Modeler Plugins option on the left panel.

Some standard modeler plugins are:

- **DeviceMap**– Collects basic information about a device, such as its OS type and hardware model.
- **InterfaceMap**– Collects the list of network interfaces on a device.
- **RouteMap**– Collects the network routing table from the device.
- **IpServicesMap**– Collects the IP services running on the device.
- **FileSystemMap**– Collects the list of file systems on a device.

Debugging the modeling process

You can run the modeler from the command line against a single device. This feature is useful when debugging issues with a plugin.

By passing the `--collect` command to the modeler, you can control which modeler plugins are used. For example, the following command runs only the interface plugin against the `build.zenoss.loc` device:

1. Log in to the Control Center host as a user with serviced CLI privileges.
2. Display the list of zenmodeler services.

```
serviced service list zenmodeler
```

On a system with multiple collectors, the result is similar to the following example:

Name	ServiceID	DepID/Path
zenmodeler /zenmodeler	7itut0ryz759ua77ntrm3hi8w	1/Zenoss.resmgr/Zenoss/Collection/localhost/localhost
zenmodeler /zenmodeler	e3bpfy6j6pyl8l346xq446myk	1/Zenoss.resmgr/Zenoss/Collection/localhost/collectorPool2
zenmodeler /zenmodeler	7dnmgcwexlqxjqko6nja0942y	1/Zenoss.resmgr/Zenoss/Collection/localhost/collectorPool3

3. Select the zenmodeler service that is associated with the device to debug, and then attach to it as the zenoss user. Replace `ServiceID` with the container ID of a zenmodeler service. For example, `7itut0ryz759ua77ntrm3hi8w`.

```
serviced service attach ServiceID su - zenoss
```

4. Run the zenmodeler command.

```
zenmodeler run -v10 --collect=IpInterface -d build.zenoss.loc
```

If the command returns any stack traces, forward the following details to Zenoss Support for assistance:

- Command you ran
- Stack trace or stack traces returned
- Version of your Resource Manager instance
- OS version and patch level for the remote device

About monitoring templates

Resource Manager stores performance configuration data in *templates*. Templates contain other objects that define where and how to obtain performance data, thresholds for that data, and data graphs.

You can define a template anywhere in the device class hierarchy, or on an individual device.

Templates are divided among three types:

- Device
- Component
- Interface

Creating templates

You can create a template by overriding an existing template. To override a template:

1. Navigate to the template you want to copy.
2. From the Action menu, select Copy/Override Template. The Copy/Override dialog box appears.
3. Select the bound template to override, and then click Submit. The copied template appears in the list of templates as locally defined.

Renaming templates

To rename an existing template:

1. Select ADVANCED > MONITORING TEMPLATES.
2. Expand the organizer containing the template to be renamed, and then the class containing the template.
3. From the Action menu, select View and Edit Details. The Edit Template Details dialog box appears.
4. Enter a new name in the Name field.
5. Click Submit.

Template binding

The determination of which templates apply to what objects is called *binding*. Templates are bound in different ways, depending on the objects to which they are bound.

- [Device templates](#)
- [Component templates](#)
- [Interface templates](#)

Device templates

Device templates are applied to devices, one to each device. The system employs a single rule to bind device templates to devices: the value of the zDeviceTemplates property. For most device classes, this is "Device."

Common device templates are:


- Device
- MySQL
- Apache
- Active Directory
- MSExchangeIS
- MSSQLServer
- IIS

For the Server/Linux/MySQL device class, the zDeviceTemplates property might contain, for example, "Device" and "MySQL." The system would collect CPU and memory information by using the Device template, and MySQL-specific metrics by using the MySQL template.


Binding a device template

To bind a device template to a device class or device:

1. From the devices list, select a device class or device.
2. On the Overview page, select Bind Templates from the Action menu.



The Bind Templates dialog box appears.

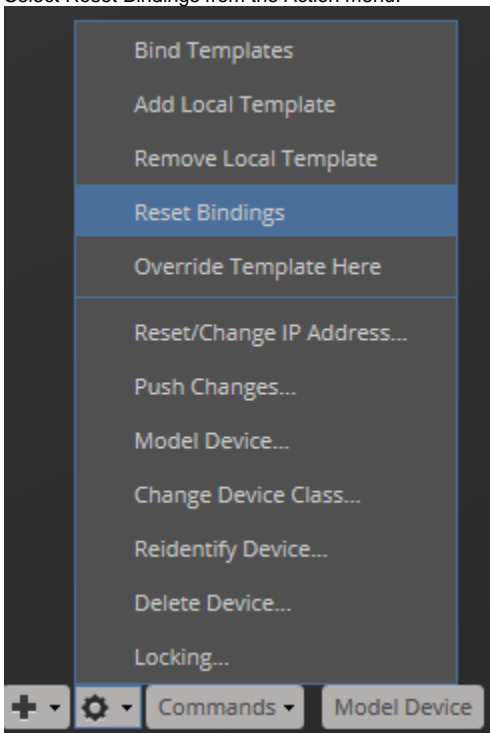


3. Move templates between the Available and Selected lists using the arrows.
4. Click Save.


Resetting bindings

Resetting template bindings removes all locally bound templates and uses the default template values. To reset bindings for a selected device or device class:

1. Select Reset Bindings from the Action menu.



The Reset Template Bindings dialog box appears.



2. Click Reset Bindings to confirm the action.

Component templates

Component templates are named exactly according to the name of the underlying class that represents a component. For example, the `FileSystem` template is applied to file systems. Component templates can be applied multiple times to each device, depending on how many of the device's components match the template. Configuration properties do not control the application of component templates.

Note: Do not manually bind component templates.

Common component templates are:

- `FileSystem`, `HardDisk`, `IPService`, `OSProcess`, `WinService`
- `Fan`, `PowerSupply`, `TemperatureSensor`
- `LTMVirtualServer`, `VPNTunnel`

Interface templates

Most interface templates are applied to network interfaces by using a special type of binding. Instead of using the name of the template's underlying target class, the system looks for a template with the same name as the interface type. You can find this type in the details information for any network interface.

For these standard interfaces when a Collection Zone cannot locate a template that matches the interface type, then it uses the `ethernetCsmacd` template.

Some vendors provide expanded monitoring capabilities on their interfaces. In some cases, ZenPacks provide custom templates specific to those types that can override this default behavior. A common example would be the `Ethernet Interfaces` interface type built to monitor these expanded capabilities for Cisco network interfaces. In this case the template's target class is used to bind the `EthernetInterface` template.

For more information about these expanded interface capabilities, refer to the [Cisco Monitor ZenPack page](#).

Production states and maintenance windows

Production state determines the level of monitoring and alerting applied to an individual device. Typically, alerting rules specify that the system will monitor and create events for devices that are in the "Production" production state.

Maintenance windows are planned time periods used to temporarily modify alerting rules so that event-generated alerts are temporarily halted during the window.

Production states

Production state determines whether a device is monitored, and can be used to control several elements of the event system, such as whether an event will produce a remote alert (email or page).

Choose a production state for a device based on whether you want:

- The device to be monitored
- The device to appear on the dashboard
- Alerting to occur

The following table lists production states and their characteristics.


Production state	Devices monitored?	Appear on dashboard?
Production	yes	yes
Pre-Production	yes	no
Test	yes	no
Maintenance	yes	may appear
Decommissioned	no	no

When you add a device to the system, its default state is Production. You may want to add triggers and notifications to alert you to various conditions that occur in the system, such as production state changes or a severity level being reached. For example, you can set up a trigger when a device is in either a production or a maintenance state and has a severity of Error or higher. You can then notify users when this trigger condition is met.

Setting the production state for a device

To set the production state for a device:

1. Click a device name in the list of devices. The Device Overview page appears.
2. Select a production state from the list of options, and then click Save.



Maintenance windows

Maintenance windows allow scheduled production state changes of a device or all devices in a system, group, or location. You might want to set up a maintenance window, for example, to change a device's production state while you perform configuration changes or reboot a device.

Instead of using a maintenance window, you can set the production state for a device manually at the time you want to make changes.

When the maintenance window starts, the production state of the device is set to the value of Start Production State (for example, Maintenance). When the maintenance window closes, the production state of the device reverts to the value of Stop Production State (the state the device was in prior to Maintenance).

Maintenance windows do not prevent notifications from being triggered on the device. If you want to define the notifications you receive during the maintenance window, you will need to set up an appropriate trigger for the device production state that you set during your maintenance window.

For more information, see [Working with triggers](#).

Maintenance window events

When a maintenance window starts, an event is created with the following information:

- depuid - zenactions | Resource |MaintenanceWindowName | TargetOrganizerOrDevice
- prodState - StartProductionState
- severity - Info
- summary/message - Maintenance window starting MaintenanceWindowName for TargetOrganizerOrDevice
- eventClass - /Status/Update
- eventClassKey - mw_change
- maintenance_devices - TargetOrganizerOrDevice
- maintenance_window - MaintenanceWindowName

When a maintenance window stops, an event is created with the following information:

- severity - Clear
- summary/message - Maintenance window stopping MaintenanceWindowName for TargetOrganizerOrDevice
- prodState - -99 (meaning "unknown.")

Maintenance window events auto-clear, meaning that stop events clear start events.

Creating and using maintenance windows

You can create a maintenance window for an individual device or group of devices (all devices, a device class, group, system, or location) in the devices hierarchy.

Create a maintenance window for a single device

Use this procedure to create a maintenance window for a device.

1. Navigate to INFRASTRUCTURE > DEVICES.
2. In the content area, click the name of the device.
3. In the sidebar, click Device Administration.
4. In the Maintenance Window toolbar, click Add.
5. In the Add New Maintenance Window dialog box, specify the attributes of the new maintenance window.
The attributes allow you to specify whether or not to enable the window, the time of day and date when the maintenance window starts (in local time zone format), the duration of the window, and whether and how the window repeats.
The Window Production State field allows you to categorize the production state of the device during the maintenance window. When the window ends, the device is returned to the production state it was in when it entered the maintenance window.
6. At the bottom of the Add New Maintenance Window dialog box, click SUBMIT.

Create a maintenance window for a group of devices


Use this procedure to create a maintenance window for a group of devices.

1. Navigate to INFRASTRUCTURE > DEVICES.
2. In the content area, select a group of devices, and then click DETAILS, located at the top of the sidebar.
3. In the sidebar, click Device Administration.
4. In the Maintenance Window toolbar, click Add.
5. In the Add New Maintenance Window dialog box, specify the attributes of the new maintenance window.
The attributes allow you to specify whether or not to enable the window, the time of day and date when the maintenance window starts (in local time zone format), the duration of the window, and whether and how the window repeats.
The Window Production State field allows you to categorize the production state of the device during the maintenance window. When the window ends, the device is returned to the production state it was in when it entered the maintenance window.
6. At the bottom of the Add New Maintenance Window dialog box, click SUBMIT.

Managing maintenance windows

Once you have created maintenance windows for your devices or groups of devices, you can quickly manage these instances on the Maintenance Windows screen.

1. Navigate to the Maintenance Window screen. This is the same place where you initially created the maintenance window (Device Administration link on Device Overview page). On this screen you can perform any of the following by clicking the appropriate icon:



- Add a new maintenance window
 - Delete the selected maintenance window
 - Edit the selected maintenance window (can also double-click a maintenance window row)
 - Toggle the selected maintenance window from enabled to disabled and vice-versa. The Enabled column will switch values.
2. Ensure that your changes are reflected in the Maintenance Window screen.

Event management

[Events](#), and the graphs generated from performance monitoring, are the primary operational tools for understanding the state of your environment.

- [Event fields](#)
- [De-duplication](#)
- [Auto-clear correlation](#)
- [Event consoles](#)
- [Event sources](#)
- [Creating events manually](#)
- [Understanding event classes](#)
- [Event mapping and transforms](#)
- [Capturing email messages as events](#)
- [Event severity levels](#)
- [Administering MIB files](#)

Event fields

To enter the event management system, an **event** must contain values for the device, severity, and summary fields. Resource Manager rejects events that are missing any of these fields.

Basic event fields are as follows:

- Summary
- Device
- Component
- Severity
- Event Class Key
- Event Class
- Collector

Events include numerous other standard fields. Some control how an event is mapped and correlated; others provide information about the event.

Device field

The device field is a free-form text field that allows up to 255 characters. Resource Manager accepts any value for this field. If the device field contains an IP address or a hostname, then the system will automatically identify and add the event to the corresponding device.

Resource Manager automatically adds information to incoming events that match a device. Fields added are:

- prodState - Specifies the device's current production state.
- Location - Specifies the location (if any) to which the device is assigned.
- DeviceClass - Classifies the device.
- DeviceGroups - Specifies the groups (if any) to which the device is assigned.
- Systems - Systems (if any) to which the device is assigned.
- DevicePriority - Priority assigned to the device.

Status field

The Status field defines the current state of an event. This field is often updated after an event has been created. Values for this numeric field are 0-6, defined as follows:

Number	Name	Description
0	New	Initial state upon creation
1	Acknowledged	A user has seen and marked the event
2	Suppressed	A transform has suppressed the event
3	Closed	A user action has closed the event
4	Cleared	A corresponding clear event has cleared the event
5	Dropped	A transform has dropped an event, so the event it not persisted
6	Aged	Automatically closed because of the severity and last seen time values

Severity field

The following table maps event severity levels to their labels and colors.

Level	Label	Color
5	Critical	Red
4	Error	Orange
3	Warning	Yellow
2	Info	Blue
1	Debug	Grey
0	Clear	Green

Summary and message fields

The summary and message fields are free-form text fields. The summary field allows up to 255 characters. The message field allows up to 4096 characters. These fields usually contain similar data.

The system handles these fields differently, depending on whether one or both are present on an incoming event:

- If only summary is present, then the system copies its contents into message and truncates summary contents to 128 characters.
- If only message is present, then the system copies its contents into summary and truncates summary contents to 128 characters.
- If summary and message are both present, then the system truncates summary contents to 128 characters.

As a result, data loss is possible only if the message or summary content exceeds 65535 characters, or if both fields are present and the summary content exceeds 128 characters.

To ensure that enough detail can be contained within the 128-character summary field limit, avoid reproducing information in the summary that exists on other fields (such as device, component, or severity).

Other fields

The following table lists additional event fields.

Field	Description
dedupid	Dynamically generated fingerprint that allows the system to perform de-duplication on repeating events that share similar characteristics.
component	Free-form text field (maximum 255 characters) that allows additional context to be given to events (for example, the interface name for an interface threshold event).
eventClass	Name of the event class into which this event has been created or mapped.
eventKey	Free-form text field (maximum 128 characters) that allows another specificity key to be used to drive the de-duplication and auto-clearing correlation process.
eventClass Key	Free-form text field (maximum 128 characters) that is used as the first step in mapping an unknown event into an event class.
eventGroup	Free-form text field (maximum 64 characters) that can be used to group similar types of events. This is primarily an extension point for customization. Currently not used in a standard system.
stateChange	Last time that any information about the event changed.
firstTime	First time that the event occurred.
lastTime	Most recent time that the event occurred.
count	Number of occurrences of the event between the firstTime and lastTime.
prodState	Production state of the device, updated when an event occurs. This value is not changed when a device's production state is changed; it always reflects the state when the event was received by the system.
agent	Typically the name of the daemon that generated the event. For example, an SNMP threshold event will have zenperfsnmp as its agent.
DeviceClass	Device class of the device that the event is related to.
Location	Location of the device that the event is related to.
Systems	Pipe-delimited list of systems that the device is contained within.
DeviceGroups	Pipe-delimited list of systems that the device is contained within.
facility	Only present on events coming from syslog. The syslog facility.
priority	Only present on events coming from syslog. The syslog priority.
ntevvid	Only present on events coming from Windows event log. The NT Event ID.
ownerid	Name of the user who acknowledged this event.
clearid	Only present on events in the archive that were auto-cleared. The evid of the event that cleared this one.
DevicePriority	Priority of the device that the event is related to.
eventClass Mapping	If this event was matched by one of the configured event class mappings, contains the name of that mapping rule.
monitor	In a distributed setup, contains the name of the collector from which the event originated.

In addition to the standard fields, the system also allows events to add an arbitrary number of additional name/value pairs to events to give them more context.

De-duplication

Resource Manager uses an event "de-duplication" feature, based on the concept of an event's fingerprint. Within the system, this fingerprint is the "dedupid." All of the standard events that the system creates as a result of its polling activities are de-duplicated, with no setup required. However, you can apply de-duplicating to events that arrive from other sources, such as syslog, SNMP traps, or a Windows event log.

The most important de-duplication concept is the *fingerprint*. An event's fingerprint (or dedupid) is composed of a pipe-delimited string that contains these event fields:

- device
- component (can be blank)
- eventClass
- eventKey (can be blank)
- severity
- summary (omitted from the dedupid if eventKey is non-blank)

When the component and eventKey fields are blank, a dedupid appears similar to:

```
www.example.com|/Status/Web||4|WebTx check failed
```

When the component and eventKey fields are present, a dedupid appears similar to:

```
router1.example.com|FastEthernet0/1|/Perf/Interface|threshName|4
```

When a new event is received by the system, the dedupid is constructed. If it matches the dedupid for any active event, the existing event is updated with properties of the new event occurrence and the event's count is incremented by one, and the lastTime field is updated to be the created time of the new event occurrence. If it does not match the dedupid of any active events, then it is inserted into the active event table with a count of 1, and the firstTime and lastTime fields are set to the created time of the new event.

The following illustration depicts a de-duplication scenario in which an identical event occurs three times, followed by one that is different in a single aspect of the dedupid fingerprint.

	Event	dedupid	count
Warning Threshold	device: router1 eventClass: /Perf/CPU eventKey: cpuThresh severity: 3	router1 Perf/CPU cpuThreshold 3	1
Warning Threshold		router1 Perf/CPU cpuThreshold 3	2
Warning Threshold		router1 Perf/CPU cpuThreshold 3	3
Error Threshold	device: router1 eventClass: /Perf/CPU eventKey: cpuThresh severity: 4	router1 Perf/CPU cpuThreshold 4	1

If you want to change the way de-duplication behaves, you can use an event transform to alter one of the fields used to build the dedupid. You also can use a transform to directly modify the dedupid field, for more powerful cross-device event de-duplication.

Auto-clear correlation

The auto-clearing feature is similar to the de-duplication feature. It also is based on the event's fingerprint. The difference is which event fields make up the fingerprint, and what happens when a new event matches an existing event's fingerprint.

All of the standard events created as a result of polling activities do auto-clearing by themselves. As with de-duplication, you would invoke auto-clearing manually only to handle events that come from other sources, such as syslog, a Windows event log, or SNMP traps.

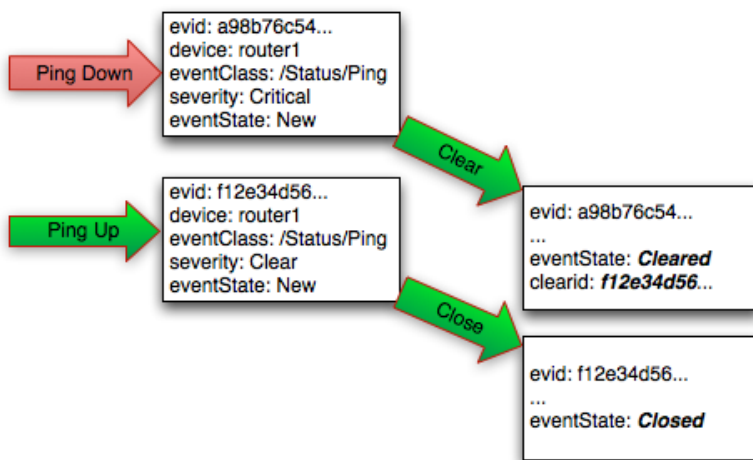
If a component has been identified for the event, then the auto-clear fingerprint consists of these fields:

- If component UUID exists:
 - component UUID
 - eventClass (including zEventClearClasses from event class configuration properties)
 - eventKey (can be blank)
- If component UUID does not exist:
 - device
 - component (can be blank)
 - eventKey (can be blank)
 - eventClass (including zEventClearClasses from event class configuration properties)

When a new event comes into the system with a special 0 (Clear) severity, Resource Manager checks all active events to see if they match the auto-clear fingerprint of the new event. All active events that match the auto-clear fingerprint are updated with a Cleared state, and the clearid field is set to the UUID of the clear event. After a configurable period of time, all events in a closed state (Closed, Cleared, and Aged) are moved from the active events table to the event archive.

If an event is cleared by the clear event, it is also inserted into the active events table with a status of Closed; otherwise, it is dropped. This is done to prevent extraneous clear messages from filling your events database.

The following illustration depicts a standard ping down event and its associated clear event.



If you need to manually invoke the auto-clearing correlation system, you can use an event transform to make sure that the clear event has the 0 (Clear) severity set. You also need to ensure that the device, component, and eventClass fields match the events you intend to clear.

To prevent inadvertently clearing a wider range of events than intended, avoid making clear events too generic.

Event consoles


Resource Manager features multiple event consoles that allow you to view and manage events. Each console shows different events subsets, depending on your current context.

- The master event console enables you to view and manage events. It displays the repository of all events that have been collected. To access this console, click EVENTS on the Navigation menu.
- Contextual event consoles are found throughout the system. Each time you see an Events selection for a device, device organizer, component, or event class, you can view event information that has been automatically filtered to show events specific to the current context.

Customizing

You can add or delete data columns to customize your event console view. The order of the selected column names determines the left-to-right display on the Event Console.

1. Navigate to EVENTS > EVENT CONSOLE.
2. Click the Configure button and select Adjust columns from the drop-down list.



3. To select a column, double-click the name in the Available list to move it to the Selected list.
4. In the Selected list, to re-order the columns, use the arrow keys.
5. Click Submit.

Selecting events

To select one or more events in the event console, you can:

- Click a row to select a single event
- Ctrl-click rows to select multiple events, or Shift-click to select a range of events

Sorting and filtering events

You can sort and filter events by any column that appears in the master event console.

To sort events, click a column header. Clicking the header toggles between ascending and descending sort order. Alternatively, hover over a column header to display its control, and then select Sort Ascending or Sort Descending.

Filter options appear below each column header.

Status	Severity	Resource	Component	Event Class	Summary	First Seen	Last Seen	Count
<input checked="" type="checkbox"/> New		ics1	Fan 3/4/2		Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	3
<input checked="" type="checkbox"/> Acknowledged		ics1	extpol_reg_clients_client...		UCS Domain ucs1-faba is registered with UCS Central without a valid license.	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	4
<input type="checkbox"/> Suppressed		ics1	Fan Module 3/4		Fan 2 in Fan Module 1-4 under chassis 3 speed: lower-non-recoverable	2015-07-29 08:45:41 am	2015-07-29 08:45:41 am	1
<input type="checkbox"/> Closed		ics1	fabric_san_A_phys-foes...		FCoE uplink port 1/8 is down	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	4
<input type="checkbox"/> Cleared		ics1	Fan 3/4/2		Fan 2 in Fan Module 1-4 under chassis 3 operability: inoperable	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	3
<input type="checkbox"/> Aged		ics1	Fabric FC SAN Port Chan...		san port-channel 5 on fabric interconnect A oper state: failed, reason: No operatio...	2015-07-29 08:44:33 am	2015-07-29 10:29:40 am	4
		ics1	Fabric FC SAN Port Chan...		san Member 1/27 of Port-Channel 5 on fabric interconnect A is down, membership:...	2015-07-29 08:46:12 am	2015-07-29 10:29:40 am	4

You can filter the events that appear in the list in several ways, depending on the field type. Date fields (such as First Seen and Last Seen) allow you to enter a value or use a date selection tool to limit the list. For other fields, such as Device, Component, and Event Class, enter a match value to limit the list.

The Count field allows you to filter the list when compared to a value. To search on count:

- *N* - Displays events with a count equal to *N*.
- *<N* - Displays events with a count less than or equal to *N*.
- *M:N* - Displays events with a count between *M* and *N* (inclusive).
- *M:* - Displays events with a count greater than or equal to *M*.

To clear filters, select **Configure > Clear filters**.

Working with live search

By default, Resource Manager uses a "live search" feature to help you locate information. From the event console, you can search for information by:

- **Device** (name) and **Component** - Device name and Component searches:
 - Are case-insensitive.
 - Are tokenized on whitespace (meaning that any searches that span whitespace and do not start with a complete token will return no results).
 - If quoted, return only exact matches.
- **Summary** - Summary searches:
 - Are case-insensitive.
 - Are tokenized on whitespace (meaning that any searches that span whitespace and do not start with a complete token will return no results).
- **Event class** - Event class searches:
 - Are case-insensitive.
 - Are tokenized on / (slash). If the search begins with a slash, and ends with a slash or asterisk, then event classes are searched by using a "starts with" approach. If a search starts with a slash and ends with any other character, then event classes are searched by using an exact match for the event class. If a search does not begin with a slash, then event classes are searched by using a sub-string match on each event class.
- **IP Address** - IP address searches (for IPv4 and IPv6 values):
 - Are tokenized by . (period) and : (colon). For example, the following searches would return a result of 129.168.1.100:
 - 168
 - 168.1
 - 129.16*
 - *29
- **Time fields**
 - **First Seen** - This is always the time of the first occurrence of the event and does not change.
 - **Last Seen** - This is the most recent occurrence of the event, and is updated each time the event occurs.
 - **State Change** - This is the time that the event state was modified, most commonly when the event is closed.

Entering a datetime in one of these filters formatted as YYYY-MM-DDHH:MM:SS displays events that have a timestamp that is equal to, or newer than, the input datetime. Note that while the input field accepts a 24-hour format, the system displays it in 12-hour format by default (using am/pm).

Additionally you can configure a time range to display events by using the following format 'startdatetime TO end datetime': "YYYY-MM-DD HH:MM:SS TO YYYY-MM-DDHH:MM:SS". An example might look like: "2017-07-21 12:00:00 TO 2017-07-22 12:00:00". This would include all events that the timestamp occurred within a 24 hour period between 12:00:00 on July 21st through 12:00:00 on July 22nd.


With live search enabled (the default behavior), the system filters available information immediately. It presents increasingly refined information with each character you type in the search window. When disabled, search responds only after you enter one or more characters and then press Enter.

Saving an event console view

Save a custom event console view by bookmarking it for quick access.

1. Select **Configure > Save this configuration**.

2. In the dialog box, select the link, and then drag it to the bookmarks area of the browser window.




The browser adds a link to the bookmarks list.

3. Change the title of the bookmark to distinguish this event console view.

Refreshing the view of events

You can refresh the list of events manually or specify that they refresh automatically. To manually refresh the view, click Refresh. You can manually refresh at any time, even if you have an automatic refresh interval specified.


To set up automatic refresh, select one of the time increments from the Refresh list.



Viewing event details

You can view details for any event in the system.

1. To view details, double-click an event row.



2. To display the event information in a new window, click the pop-out icon.
3. To see more information about the event, click the link for Event Management, Device State, Event Data, or Event Details.
4. In the log area, enter information about the event, and then click Add.

Acknowledging events

You may want to mark an event as "acknowledged" to indicate, for example, that you have taken action to remedy a problem. To mark events as acknowledged:

1. Select one or more events in the event console view.
2. Click the Acknowledge Events icon. A check mark appears for each acknowledged event.

Returning events to new status

Returning a previously acknowledged event to "new" status revokes its "acknowledged" status.

1. Select one or more events in the event console view.
2. Click the Unacknowledge Events icon. A check mark no longer appears in the event row, and the event is returned to "new" status.

Classifying events

Classifying events lets you associate events shown as /Unknown with a specific event class. To classify an unknown event, an event class key must be specified for the event.

1. Select one or more /Unknown events in the event console view.
You can also classify events from the event archive.
2. Click the Reclassify an Event icon. The Classify Events dialog appears.
3. Select an event class from the list of options, and then click Submit.

Closing events

When you no longer want to actively monitor an event (after you acknowledge it, for example), you can specify to close the event and move it to the event archive according to a configured event archive interval. To do this:

1. Select one or more events in the event console view.
2. Click the Close Events icon. The selected events are closed and moved to the archive at the specified interval.

To view events in the event archive, select EVENTS > Event Archive.

Note: Users with no assigned role can view all events in the archive.

3. Click the Refresh icon to update the event list.
The closed events are removed from the display in the event console view.

Reopening events

You can reopen events in the active event console that are in the Closed, Cleared, or Aged state.

You cannot re-open a closed event if another active event with the same fingerprint exists. Before you can re-open the closed event, you must close the new event.

1. Select one or more Closed, Cleared, or Aged events.
2. Click the Reopen Events icon. The selected events are returned to active status.

Exporting event data

You can export data from the event console to a comma-separated value (.csv) or XML file. You can select individual events (to export only those events), or make no selections (to export all events that match the current filter criteria).

1. Select one or more events.
2. Select Export > CSV or Export > XML. By default, the exported file is named events.csv or events.xml.

Event sources

Events enter the system as follows:

- *Generated events* are created as a result of active polling.
- *Captured events* are transmitted by external actions into the system.

Generated events

The following standard services generate events. They automatically perform appropriate de-duplication and auto-clearing.

- **zenping** - Ping up/down events
- **zenstatus** - TCP port up/down events
- **zenperfsnmp** - SNMP agent up/down events, threshold events
- **zencommand** - Generic status events, threshold events
- **zenprocess** - Process up/down events, threshold events
- **zenwin** - Windows service up/down events

Captured events

Captured events are those events that the system does not specifically know will occur in advance. De-duplication is performed on these events, but might require tuning. By default, no auto-clearing is done on captured events. Event transforms must be used to create the auto-clear correlations.

The following services collect captured events:

- **zensyslog**- Events created from syslog messages.
- **zentrap**- Events created from SNMP traps and informs.

ZenPacks that you install might include their own services.

SNMP traps

An SNMP trap is a message that is initiated by a network element and sent to the network management system. Often, traps indicate a failure of some sort, such as a router message indicating a power supply failure, or a printer message indicating an "out-of-ink" condition.

If an SNMP trap enters the system, and Resource Manager cannot identify the event (the event is classified as "/Unknown"), then you can classify the event so that the system handles it consistently.

Classifying SNMP traps

By default, most SNMP traps will appear in the /Unknown event class. To map them to a more meaningful event class, you can re-classify them with an event mapping.

To classify an SNMP trap event:

1. From the Event Console, select the unknown event or events.
2. Click the Reclassify an event icon. The Classify Events dialog appears.
3. Select /App, and then click Submit.

To edit this classification:

- a. From the Navigation area, select Events > Event Classes.
- b. Ensure Mapping Instances appears.
- c. Select the event map you created.
- d. In the left panel, select Edit from the Action icon.

The edit page appears. This page contains rules used to map the event to the /App category. This rule, since it matches the trap by a specific OID, is all that is needed.

In the Transform area, you can enter code to modify the summary. For example, if you want to set the summary string to "Spam Filter Detects Virus," then you can enter:

```
evt.summary = "Spam Filter Detects Virus"
```

A trap has a header with some standard information, followed by a sequence of attribute/values.

You have indicated you want the value for the OID ".1.3.6.1.4.1.9789.1500.2.5" as the summary. If you had the MIB loaded, you could do this:

```
evt.summary = evt.spamFilterDetectsVirus
```

However, the OID and the data is still in there. Instead, use the slightly more cryptic:

```
evt.summary = getattr(evt, ".1.3.6.1.4.1.9789.1500.2.5", "Unexpected missing OID")
```

The "device" object for the event has been made available, as well:

```
evt.summary = getattr(evt, ".1.3.6.1.4.1.9789.1500.2.5", "Unexpected missing OID") + " from device " + device.getId()
```

Resource Manager uses MIBs to translate SNMP traps that contain raw OID values. Loading a MIB into the system allows it to translate numeric OIDs such as .1.3.6.1.2.1.1.6 into descriptive phrases like "sysLocation". It also makes it easier to manipulate the events in an event mapping.

Following is a small demonstration MIB.

```
NOTIFICATION-TEST-MIB DEFINITIONS ::= BEGIN
IMPORTS
ucdavis FROM UCD-SNMP-MIB
NOTIFICATION-TYPE FROM SNMPv2-SMI
;
demonotifs OBJECT IDENTIFIER
::= { ucdavis 991 }
demo-notif NOTIFICATION-TYPE
OBJECTS { sysLocation }
STATUS current
DESCRIPTION "Just a test notification"
::= { demonotifs 17 }
END
```

Example: Sending test traps

To send an SNMP trap:

1. As root, attach to the zope service.

```
serviced service attach zope/0
```

2. From the command line, enter the following command:

```
snmptrap -v 2c -c public localhost '' 1.3.6.1.4.1.2021.991 .1.3.6.1.2.1.1.6 s "Device in Austin"
```

3. Save this demonstration MIB into a file.
4. Send the trap.
5. Open the Event Console and find the trap you sent.
6. Send this event to the event archive.
7. Load some MIBs into the system so that this OID is translated into a better format. As a pre-requisite to this procedure, see [Administering MIB files](#).
 - a. Copy the demonstration MIB into `/opt/serviced/var/volumes/<instance_id>/zenoss-var-ext/uploadedMIBs`.
 - b. Run the following command:

```
serviced service run zope zenmib run -v10 \  
/opt/zenoss/var/ext/uploadedMIBs/<MIB_filename> \  
--path=<device_class> --removemiddlezeros \  
--mibdepsdir=/opt/zenoss/var/ext/uploadedMIBs
```

The `--removemiddlezeros` flag is needed when MIBs use notification OIDs with embedded zeros.

- c. Send the trap a second time:

```
snmptrap -v 2c -c public localhost '' 1.3.6.1.4.1.2021.13.991 .1.3.6.1.2.1.1.6 s "Device in Austin"
```

- d. Check the event. Make sure the count is 1. If the count is 2, send the event to the event archive and send the trap again. Look at the Details tab. Now you should see something like this:

```
sysLocation Device in Austin
```

You should also see that the event summary changes from:

```
snmp trap 1.3.6.1.4.1.2021.13.991 from localhost
```

to:

```
snmp trap ucdExperimental from localhost
```

Transforming events with event mappings

To modify events as they arrive, create an event map through the user interface:

1. Create an event class.
2. Go to the event console and create an event mapping in this class from the existing event.
3. Edit the map.
4. In the Transform area, update the event with detail data. The entry field allows you to insert Python scripts. The event is provided as "evt" and the device as "dev." In this case, extract the sysLocation event detail and make it the summary with:

```
evt.summary = evt.sysLocation
```

5. Save the event mapping.

If you move the event to the event archive and resend the trap, the summary for the trap should now read the device name in the location you assigned.

If you encounter problems with the transform, check the zentrap.logfile for errors that occurred.

Mapping SNMP variables to events

Some SNMP traps can include variables (varbind objects), which are ordered implicitly. The ordering requirement takes the form of *Name.Number*—like `someVar.0`—and in many cases there will be a series of varbind objects with different numbers on the same name. The following tables provide an example variable and varbind objects.

OID	Value
1.2.1.1.3.0	Message0
1.2.1.1.3.1	Message1

Assuming a MIB (imported into Resource Manager) specifies the name `someVar` (1.2.1.1.3) then the event details would be as follows:

Name	Value
someVar.0	Message0
someVar.1	Message1
someVar.sequence	0,1

The following tables illustrate how the implicit ordering is encoded in event details.

Example trap with an SNMP varbind object

OID	Value
1.3.6.1.2.1.2.2.1.1.143	143
1.3.6.1.2.1.2.2.1.7.143	1
1.3.6.1.2.1.2.2.1.8.143	1
1.3.6.1.2.1.2.2.1.2.143	"F23"
1.3.6.1.2.1.31.1.1.1.18.143	""

Event details for example trap

Name	Value
ifIndex.143	143
ifIndex.sequence	143
ifAdminStatus.143	1
ifAdminStatus.sequence	143
ifOperStatus.143	1
ifOperStatus.sequence	143
ifDescr.143	F23
ifDescr.sequence	143
ifAlias.143	
ifAlias.sequence	143

The event details are repetitive, but an event transform can parse and process sequenced varbind objects.

For example, the following event transform concatenates the `someVar` parts into the event's summary attribute:

```
seq = getattr(evt, "someVar.sequence", None)
if seq is not None:
    values = []
    for idx in str(seq).split(','):
        value = getattr(evt, "someVar." + idx, '')
        values.append(value)
    evt.summary = ' '.join(values)
```


Event transforms based on event class

When an event arrives in the system, you can change values (such as severity). For example, you can make the summary more informative, or change severity according to text within the summary.

Each event class allows for a short Python script to be executed when an event arrives.

For example, a user may want full file system threshold events on /data to be critical. Add the following Python script in the Threshold Transform of /Events/Perf/Filesystem:

```
if evt.component == '/data' and evt.severity != 0: evt.severity = 5
```

Like event mappings for event class keys, "evt" and "dev" objects are available in the script of the transform.

SNMP trap filtering at the collector level

The zentrap service supports SNMP trap filtering. For details on the filter definition, see the `/opt/zenoss/etc/zentrap.filter.conf` file in the zentrap service definition. For more information about enabling this feature, refer to [How To Setup A Filter For SNMP Traps At The Collector Level \(Zentrap\) In Resource Manager 5.1.1](#) (also applies to later versions).

Configuring SNMP trap forwarding

For multi-host systems, you must enable the zentrap service container to move between delegate hosts in the resource pool. Otherwise, the container is bound to a specific host, and if that host fails, you cannot receive traps. For resource pools with a single host, following this best practice allows for expansion without changing or moving IP addresses.

Before performing this task, obtain IP addresses for zentrap on the same network as your delegate hosts.

1. Log in to the Control Center browser interface.
2. In the main navigation menu, click Applications.
3. In the Applications table, click Zenoss.resmgr.
4. In the IP Assignments table under Resource Pool, choose the resource pool that contains the zentrap service.
5. In the Virtual IPs table, click Add Virtual IP.
6. In the Add Virtual IP dialog box, provide information about your IP address, and then click Add Virtual IP.
7. In the main navigation menu, click Applications.
8. In the Applications table, click Zenoss.resmgr.
9. In the Services table, expand the Zenoss service and subservices, locate zentrap and click the service name.
The page for the zentrap service opens.
10. In the IP Assignments table under Actions, click Assign.
11. In the Assign IP dialog box, provide information about the virtual IP address that you added to the pool, and then click Assign IP.
12. On the service page, click Restart.
13. When forwarding SNMP traps to Resource Manager, use this IP address.

Configuring syslog message forwarding

For multi-host systems, you must enable the zensyslog service container to move between delegate hosts in the resource pool. Otherwise, the container is bound to a specific host, and if that host fails, you cannot receive syslog messages. For resource pools with a single host, following this best practice allows for expansion without changing or moving IP addresses.

Before performing this task, obtain IP addresses for zensyslog on the same network as your delegate hosts.

1. Log in to the Control Center browser interface.
2. In the main navigation menu, click Applications.
3. In the Applications table, click Zenoss.resmgr.
4. In the IP Assignments table under Resource Pool, choose the resource pool that contains the zensyslogservice.
5. In the Virtual IPs table, click Add Virtual IP.
6. In the Add Virtual IP dialog box, provide information about your IP address, and then click Add Virtual IP.
7. In the main navigation menu, click Applications.
8. In the Applications table, click Zenoss.resmgr.
9. In the Services table, expand the Zenoss service and subservices, locate zensyslog and click the service name.
The service page for zensyslog opens.
10. In the IP Assignments table under Actions, click Assign.
11. In the Assign IP dialog box, provide information about the virtual IP address that you added to the pool, and then click Assign IP.
12. On the service page, click Restart.
13. When forwarding syslog messages to Resource Manager, use this IP address.

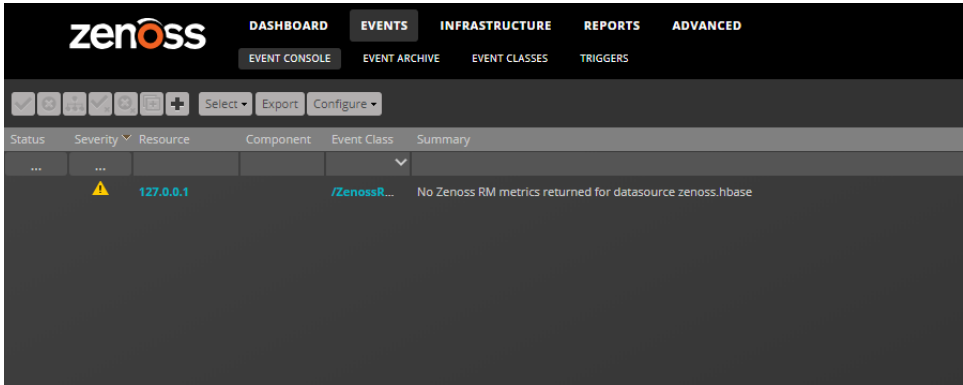
Creating events manually

Manually-created events are useful for testing event mappings, event transforms, and triggers/notifications.

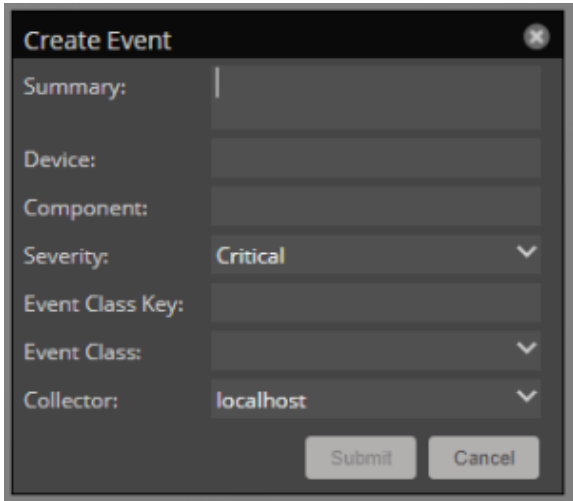
Creating events in the browser interface

Follow these steps:

1. Navigate to EVENTS > EVENT CONSOLE.



2. Click the Add an event icon.



3. In the Create Event dialog box, add event information.

Field	Required?	Description
Summary	Yes	A text message describing the event.
Device	Yes	The IP address or hostname of a device; the subject of the event.
Component	No	A device component contained in the subject of the event.
Severity	Yes	The event severity level .
Event Class Key	No	The event class key.
Event Class	No	The event class.
Collector	Yes	The Resource Manager collector that contains the subject of the event (the device).

4. Click Submit.

Event class mappings are applied only for events that do not already have an event class.

Creating events with the CLI

To send events from the command line, use the `zensendevent` command.

Common options include:

- `-d DEVICE, --device=DEVICE`
- `-i IPADDRESS, --ipAddress=IPADDRESS`
- `-y EVENTKEY, --eventkey=EVENTKEY`
- `-p COMPONENT, --component=COMPONENT`
- `-k EVENTCLASSKEY, --eventclasskey=EVENTCLASSKEY`
- `-o OTHER, --other=OTHER`
- `-s SEVERITY, --severity=SEVERITY`
- `-c EVENTCLASS, --eventclass=EVENTCLASS`
- `-f INPUT_FILE, --file=INPUT_FILE`

Follow these steps:

1. Log in to the Control Center host as a user with `serviced` CLI privileges.
2. Attach to the zenhub service as the `zenoss` user.

```
serviced service attach zenhub su - zenoss
```

3. Run the `zensendevent` command.

```
zensendevent Options "Event_Summary_Text"
```

For example, the following command shows how to simulate a ping down event:

```
zensendevent -d router1.example.com -s Critical -c /Status/Ping "Router down"
```

Understanding event classes

Event classes are a simple organizational structure for the different types of events that the system generates and receives. This organization is useful for driving alerting and reporting. You can, for example, create an alerting rule that sends you an email or pages you when the availability of a Web site or page is affected by filtering on the /Status/Web event class.

Following is a subset of the default event classes. You can create additional event classes as needed.

- /App - Application-related events.
- /Change - Events created when the system finds changes in your environment.
- /Perf - Used for performance threshold events.
 - /Perf/CPU - CPU utilization events
 - /Perf/Memory - Memory utilization or paging events
 - /Perf/Interface - Network interface utilization events
 - /Perf/Filesystem - File system usage events
- /Status - Used for events affecting availability.
 - /Status/Ping - Ping up/down events
 - /Status/Snmp - SNMP up/down events
 - /Status/Web - Web site or page up/down events
- /ZenossRM - Resource Manager system events, including key metrics for nodes in the metric collection, event generation, and modeling processes

For more information, see [Table of event configuration properties](#).


Event mapping and transforms

You can map or transform events to perform a wide range of operations, from altering the severity of certain events to altering nearly every field on an event, based on complex rules.

You cannot alter the following fields through event transformation. (This is because they are set after transformation has been performed.)

- evid
- firstTime
- lastTime
- count

The following illustration shows the path followed by an incoming event in the event mapping system.



The process begins with the "eventClass field exists" decision. This also is one of the more important differentiators in how you must handle a particular type of event.

Event class mappings

To view event class mappings, select EVENTS > EVENT CLASSES, and then select Mapping Instances in the drop-down list. This allows you to see all event class mappings in a single location. The ID column shows the mapping's event class.

You can create event class mappings directly from the event classes, but this requires that you know the event class key. A simpler way to create event class mappings is through the event console:

1. Select an event that you want to match in the event console.
2. Click the Reclassify an Event icon. The Classify Events dialog appears.
3. Select the event class to which you want to map the event, and then click Submit. This creates the event class mapping with the correct event class key, and example text against which you can develop your regular expression.

When editing an event class mapping, you can control which events it will match, as well as other properties:

- **Matching tab**
 - **Event Class Key**- Must match the incoming event's Event Class Key field for this mapping to be considered as a match for events.
 - **Rule**- Provides a programmatic secondary match requirement. It takes a Python expression. If the expression evaluates to True for an event, this mapping is applied.
 - **Regex**- The regular expression match is used only in cases where the rule property is blank. It takes a Perl Compatible Regular Expression (PCRE). If the regex matches an event's message field, then this mapping is applied.
 - **Explanation**- Free-form text field that can be used to add an explanation field to any event that matches this mapping.
 - **Resolution**- Free-form text field that can be used to add a resolution field to any event that matches this mapping.
- **Transforms tab**- Takes Python code that will be executed on the event only if it matches this mapping. For more details on transforms, see the section titled "Event Class Transform."
- **Configuration Properties tab**- Listing of Configuration Properties defined for this event class.
- **Sequence tab**- Sequence number of this mapping. This number determines the order in which mappings with the same event class key are evaluated.

Mappings have the same configuration properties as event classes. Any configuration property set locally on a mapping will override the same property set on the event class. This works in the same hierarchical, most specific match, concept that device class and device configuration properties work.

When a captured event occurs, it will not have a pre-defined event class. For this type of event, you must create an event class mapping if you want to affect the event. If a captured event occurs and none of the event class mappings in the system match it, its event class will be set to /Unknown, and it will retain all of the default properties with which it began.

The next step of evaluation for events without an event class is to check the Event Class Key field. This controls which event class mapping the event will match. If the event has a blank event class key, or its event class key does not match any event class mappings in the system, the special "defaultmapping" event class key is searched for instead. This provides for a way to map events even if they have a blank or unpredictable event class key.

Event class mapping sequence

The sequence area of an event class mapping (select Sequence in the left panel) allows you to provide more than one mapping for the same event class key. In this case, the sequence is evaluated in ascending order until a full (rule or regex) match is found.

For example, suppose a router is sending in unclassified events that need to be mapped to two event classes:

- /Events/Router/fanDown
- /Events/Router/fanUnknown

The event class key for both has been sent to "router", but one has a message of "Fan Down" and the other has no message at all. The mapping on /Events/Router/fanDown has an event class key of "router" and a regex of "Fan Down." The mapping on /Events/Router/fanUnknown has only an event class key of "router" and (in this example) no regex. Because the fanUnknown mapping matches the fanDown events, the evaluation of fanDown needs to occur first.

You can modify the evaluation of mappings with the same event class key in the Sequence area of any of those event class mappings. In the previous example, you could go to either mapping, select Sequence, and both mappings would be displayed. You can set one to 0, and the other to 1. (You can enter other values, but they will be changed to the shortest list of integers, starting with 0.) Setting fanDown to 0 and fanUnknown to 1 will ensure that the events will be mapped properly.

Event class transform

When a generated event occurs, it has an event class assigned to it. This causes the event class mapping step to be skipped. The only way to affect the fields of one of these events is through the configuration properties and transform of the event class.

To access the transform for an event class:

1. Navigate to the event class from Events > Event Classes.
2. From the drop-down list, select Transforms.
3. Enter information into the dialog box (as Python code), and then click the Save button in the upper-right corner. As you develop your transform, you can revert back to the last saved state by clicking the Revert this Transform button.

The objects available in this Python context are evt (the event); and, if the event matches a device that exists in the system database, a device object.

The following example shows how you can validate that a device object exists before using it to drop events from a particular location.

```
if device and "Hawaii" in device.getLocationName(): evt._action = "drop"
```

Capturing email messages as events

ZenMail and ZenPop3 can capture email messages as events. This capability can be useful for situations in which embedded systems (such as WAPs, NAS devices, or RAID controllers) rely on email notification for events.

ZenMail

ZenMail serves as an SMTP server that you can bind to a specific TCP port. You can then configure your embedded system to send mail to the Resource Manager server explicitly by using the server's IP address as the relay.

ZenMail supports two configuration alternatives:

- Bind to port 25 on all ports and listens for email messages to arrive (default). Ignore the TO field in the email and use the FROM address as the device IP address.
- Bind to a specific port. Useful in situations in which an SMTP server is already running on the Resource Manager server and you do not want to interfere with the existing mail delivery system. Semantics are the same as the no argument version (FROM address is used as the device IP).

ZenPop3

ZenPop3 allows you to retrieve event email from a POP server. ZenPop3 supports these configuration directives:

Directive	Description
--usesssl	Issue the STARTTLS command to the POP server and attempt to transfer email messages using SSL encryption. This is required if retrieving mail from Google.
--nodelete	Do not issue the DELE command after retrieving all messages. Typically this is used during initial testing so that you do not have to resend test messages to the POP account. Some email systems (such as Google) do not actually delete messages when the DELE command is issued.
--pophost	The hostname or IP address of the POP server from which to retrieve messages.
--popport	The TCP port the POP server listens on. Defaults to 110. Used in situations where the POP provider listens on another port (for example, Google on port 995).
--popuser	The user name that contains email messages to retrieve.
--poppass	The password to use for the user name provided.
--cycletime	The time to sleep between polls. After all email is retrieved, ZenPop3 sleeps for this amount of time before waking up and attempting to pull new email.

Translating message elements to the event

Resource Manager translates various message elements to the event, as follows:

Field	Description
FROM	If the FROM field is an IP address, then the system associates the event with the device with the same IP address. If the FROM field is a fully qualified domain name, then the system resolves it to an IP address, and then performs the device association using the resolved IP address. The resolution of hostname uses "A" records rather than "MX" records.
TO	The system ignores the TO field in the email message. ZenMail accepts email to any user and domain name combination. ZenPop also drops the TO field, and uses only the FROM field.
SUBJECT	ZenMail and ZenPop use the SUBJECT as the event summary.
Message body	ZenMail and ZenPop use the first mime attachment as the event details. The system ignores secondary message bodies (typically HTML-encoded versions of the message). It also ignores attachments (such as files).

Event severity levels

The following table maps event severity levels to their labels and colors.

Level	Label	Color
5	Critical	Red
4	Error	Orange
3	Warning	Yellow
2	Info	Blue
1	Debug	Grey
0	Clear	Green

Administering MIB files

Administering MIB files

A Management Information Base (MIB) is defined as a database of management information called managed objects that represent a specific resource. A MIB is used and maintained by a network management protocol such as SNMP. The values of the MIB object can be changed or retrieved using SNMP commands. MIBs can also be used to translate values from SNMP traps.

Object Identifiers (OIDs) are used to identify managed objects within a MIB hierarchy. Each OID is a unique value that consists of a long sequence of numbers separated by periods (.). OIDs follow strict structure, similar to that of the directory structure of a file system.


- The top level is always **root**.
- The level below **root** is ISO, represented by the number 1.
- The level below **ISO** is ORG, represented by the number 3.
- The level below **ORG** is DOD, represented by the number 6.

Top level MIB OIDs belong to various standards organizations. Vendor define their own private branches of OIDs for their products. The following table shows the OID hierarchical structure:

OID value	Level
1	iso
1.3	org
1.3.6	dod
1.3.6.1	internet
1.3.6.1.1	directory
1.3.6.1.2	mgmt
1.3.6.1.3	experimental
1.3.6.1.4	private
1.3.6.1.4.1	enterprise
1.3.6.1.5	security
1.3.6.1.6	SNMPv2
1.3.6.1.7	mail
1.3.6.1.8	features

Using MIB files

Resource Manager uses MIB files to translate SNMP traps that contain raw OID values. Loading a MIB into the system allows it to translate numeric OIDs (such as .1.3.6.1.2.1.1.6) into descriptive phrases such as `sysLocation`. It also makes it easier to manipulate the events in an event mapping or transform. The following figure shows the SNMP trap interactions:



Importing pre-loaded MIB organizers and files

Import the following pre-loaded MIB organizers and files (IANA, IETF, IRTF, SITE, and TUBS):

1. Log in to the Control Center master host as `root`, or as a user with superuser privileges.
2. Navigate to `/opt/serviced/var/volumes/<instance_id>/zenoss-var-ext` and determine whether the `uploadedMIBs` directory exists. Replace `<instance_id>` with the appropriate ID for your instance. If it does not exist, create it and change the owner and permissions as follows. (You only need to perform this action once.)
 - a. Ensure you are in the `/opt/serviced/var/volumes/<instance_id>/zenoss-var-ext` directory.

```
cd /opt/serviced/var/volumes/<instance_id>/zenoss-var-ext
```

- b. Create the `uploadedMIBs` directory.

```
mkdir uploadedMIBs
```

- c. Change the owner and permissions of the `uploadedMIBs` directory by executing the following commands.

```
chown 1337:1206 uploadedMIBs  
chmod 775 uploadedMIBs
```

3. Copy the files from the `/usr/share/mibs` directory to the `uploadedMIBs` directory. (You only have to do this once.)
 - a. Attach to the `zope` service.

```
serviced service attach zope/0
```

- b. Change to the `/opt/zenoss/var/ext/uploadedMIBs` directory.

```
cd /opt/zenoss/var/ext/uploadedMIBs
```

- c. Copy the MIB files to that directory.

```
cp -Rp /usr/share/mibs/* .
```

Creating a MIBs organizer

Create a MIBs organizer in which to load MIBs.

1. In the Resource Manager browser interface, choose **ADVANCED > MIBS**.
2. In the lower left corner, click the Action menu and choose **Add MIB Organizer**.
3. In the Create MIB Organizer dialog box, specify the identifier, and then click **SUBMIT**.

The new organizer appears in the left panel.

Installing custom or additional MIB files

Before you install custom or additional MIB files,

- Create an uploadedMIBs directory as described in [Importing pre-loaded MIB organizers and files](#).
- If you want to load MIBs into a MIBs organizer that does not exist, create the organizer as described in [Creating a MIBs organizer](#).

Perform these steps:

1. Log in to the Control Center master host as `root`, or as a user with superuser privileges.
2. Copy any MIBs and their dependencies that you want to install into the `/opt/serviced/var/volumes/<instance_id>/zenoss-var-ext/uploadedMIBs` directory on the Control Center master host. Replace `<instance_id>` with the appropriate ID for your instance. After those files are copied to that directory, they are immediately available inside containers at `/opt/zenoss/var/ext/uploadedMIBs` with no restart needed.
3. Execute the following command to install the MIB. Include `--removemiddlezeros` if MIBs use notification OIDs with embedded zeros. Include `--path` to specify an existing MIBs organizer in which to load the MIBs.

```
serviced service run zope zenmib run -v10 /opt/zenoss/var/ext/uploadedMIBs/<mib_filename> \  
--path=<mibs_organizer> --removemiddlezeros --mibdepsdir=/opt/zenoss/var/ext/uploadedMIBs
```

4. Restart the `zope` service and any instances of the `zentrapp` service.

Triggers and notifications

The event processing service uses rules ([Boolean expressions](#)) to examine each [event](#) it receives. If an event matches a rule, the service initiates the actions associated with the rule.

- [Triggers](#) contain the rules used to examine events. The rules of a trigger can be simple or complex, and most event properties are available for examination.
- [Notifications](#) associate triggers with an action. Notifications can send email messages to specific users, create SNMP traps, run arbitrary commands, or perform other actions.

A trigger may be associated with multiple notifications. A notification may include multiple triggers.

For more information about triggers and notifications, take a look at the following video.

Working with triggers

Defining a trigger involves the following steps:

1. [Create the trigger](#). This step defines the name of a trigger.
2. [Edit the trigger](#). This step defines the rules of the trigger and optionally, the users who can edit the trigger.

All new triggers include the following rules:

- The device production state equals `Production`.
- The `event severity` is greater than or equal to `Error`.

Most event properties can be used in rules, including:

- Device and device component
- Event summary
- Event count
- Device location

For more information about event properties, see [Event fields](#).

Rules can use the following operators for comparisons:

- Equals / Does not equal
- Contains / Does not contain
- Greater than / Greater than or equal to
- Less than / Less than or equal to

Text comparisons are case-sensitive.

Best practice recommendations

Create triggers that match their target events with the simplest possible set of rules. Even so, the simplest possible set of rules may be complex.

To further refine trigger and notification processing, use event mappings to transform events. The event processing service performs transforms before examining triggers.

Typically, a trigger should include the following rules:

- The device production state equals `Production`.
- The `event severity` equals (some value).
- The event status equals `New`.

If a rule uses the Count property, choose one of the following comparisons instead of the equals comparison:

- Greater than / Greater than or equal to
- Less than / Less than or equal to

Otherwise, when a notification that uses the trigger includes a delay, an equals comparison may never match, because the count could change before the delay expires.


By default, Resource Manager supports 200 triggers without performance implications. If you have more than 200 triggers, please contact Zenoss Support.

For more information about triggers, take a look at [this video](#).


Creating a trigger

Follow these steps:

1. Navigate to **EVENTS > TRIGGERS**.



2. Click the Add icon.

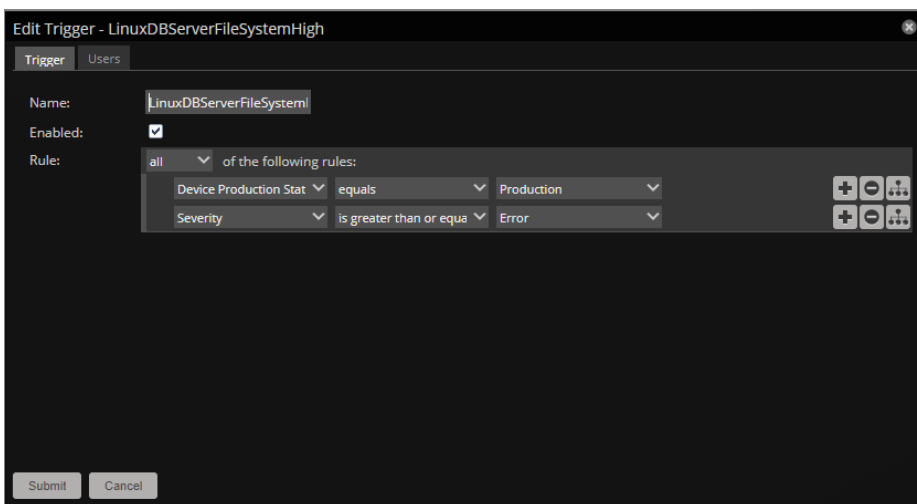


3. In the Id field, enter a name for the trigger.
Use descriptive names. For example, LinuxDBServerFileSystemHigh.
Only uppercase and lowercase letters, digits, spaces, and underscores are allowed in trigger names.
4. Click Submit.

Editing a trigger

Follow these steps:

1. Navigate to EVENTS > TRIGGERS.
2. Double-click the name of the trigger to edit.



3. Modify the default rules, add new rules, or add subordinate rules.
To identify the properties to include in a rule, navigate to EVENTS > EVENT CONSOLE and review examples of the type of event for which the trigger is targeted.
4. Click Submit.

Working with notifications

Defining a notification involves the following steps:

1. [Create the notification](#). This step defines the name and the type of a notification .
2. [Edit the notification](#). This step associates triggers, specifies the action, and for email actions, specifies the list of recipients.
3. [Set a schedule for the notification](#), if desired. This optional step defines a time window during which a notification is active.

Best practice recommendations

Use the Count property in triggers and the delay and repeat options in notifications to calibrate the sensitivity of a notification.


- Typically, configuring a notification to initiate its action the very first time one of its triggers matches an event results in false alarms. Services send events in response to temporary conditions frequently, so a delay, a count, or both can reduce frustrating false alarms.
- Set delays at a minimum of one polling cycle plus 10%. This provides enough time to determine whether the observed conditions are temporary or symptoms of an issue that needs attention. The default polling cycle is 300 seconds (5 minutes) and is defined by [the zPropertyMonitorInterval property](#).
- The repeat option has priority over other options, except delay. So even if the Send only on Initial Occurrence option is checked, the notification action is repeated until the trigger no longer matches an event.
- The Send only on Initial Occurrence option is useful for all events except intermittent, out-of-band events.

For more information about notifications, take a look at [this video](#).


Creating a notification

Follow these steps:

1. Navigate to EVENTS > TRIGGERS.



2. In the left column, click Notifications.
3. In the NOTIFICATIONS area, click the Add icon.



4. In the Id field, enter a name for the notification.

Use descriptive names. If possible, use a name that identifies the trigger that the notification contains. For example, `LinuxDBServerFileSystemHigh_Email`.

Only uppercase and lowercase letters, digits, spaces, and underscores are allowed in notification names.

5. In the Action field, choose an action type.

Action	Description
AWS Email Host	Send a plain-text or HTML email message through Amazon SES . You must have an Amazon SES account to use this action.
Command	Invoke a Linux shell command or script on a remote host. Use this action to correct issues automatically or integrate with external systems.

Email	Send a plain-text or HTML email message through the default Resource Manager SMTP service.
Syslog	Send a message to a <code>syslog</code> host.
SNMP Trap	Send an SNMP trap. Event details are converted to OIDs .
SNMP Trap w /Impact	Send an SNMP trap with additional event details that are specific to Service Impact.
WinCommand	Invoke a Windows CMD or PowerShell command on a remote host.


The list of actions in a Resource Manager instance may include additional options, because ZenPacks can introduce new actions.

6. Click Submit.

Editing a notification

Follow these steps:


1. Navigate to **EVENTS > TRIGGERS**.
2. In the left column, click **Notifications**.
3. Double-click the name of the notification to edit.



4. On the Notification tab, configure the notification.

Option	Description
Enabled	Activate or deactivate the notification. (Use a notification schedule to override this option for specific periods of time.)
Send Clear	Perform the action when the problem is resolved by a clearing event.
Send only on Initial Occurrence?	Perform the action only the first time a trigger matches an event.
Delay (seconds)	The number of seconds to wait before performing the action.
Repeat (seconds)	The number of seconds to wait before performing the action again. The action repeats until the triggering event is cleared.

5. On the Notification tab, add one or more triggers.
 - a. From the drop-down list in the TRIGGERS area, choose a trigger.




- b. Click Add.
 - c. Repeat the previous substeps as desired.
6. On the Content tab, configure the action. For more information, see one of the following sections:
 - a. [AWS Email Host](#)
 - b. [Command](#)

- c. Email
- d. Syslog
- e. SNMP Trap
- f. SNMP Trap w/Impact
- g. WinCommand

7. For AWS Email Host and Email actions only: On the Subscribers tab, specify the recipients.

To simplify maintenance, use groups rather than individual recipients.

- a. From the drop-down list in the SUBSCRIBERS area, choose a recipient.



- b. Click Add.
- c. Repeat the previous substeps as desired.

8. Click Submit.

Setting a notification schedule

A notification schedule defines the time window when a notification is active. When a notification is inactive, the event processing service does not initiate its action. Typical use cases for notification schedules include scheduling production hours, managing service level agreements, and configuring after-hours procedures.

Defining a notification schedule includes the following steps:

1. [Create the notification schedule](#). This step associates a specific notification with a schedule object.
2. [Edit the notification schedule](#). This step configures the schedule.


Creating a notification schedule

Follow these steps:

1. Navigate to EVENTS > TRIGGERS.
2. In the left column, click Notifications.
3. In the NOTIFICATIONS area, click the name of the notification to schedule.

NOTIFICATIONS				NOTIFICATION SCHEDULES			
Enabled	ID	Trigger	Action	Subscribers	Enabled	ID	Start
No	ExampleServiceEmailNotification	Example Servic...	email	0			
No	LinuxDSErveFileSystemHigh_Email		email	0			

4. In the NOTIFICATION SCHEDULES area, click the Add icon.



5. In the Id field, enter a name for the schedule.

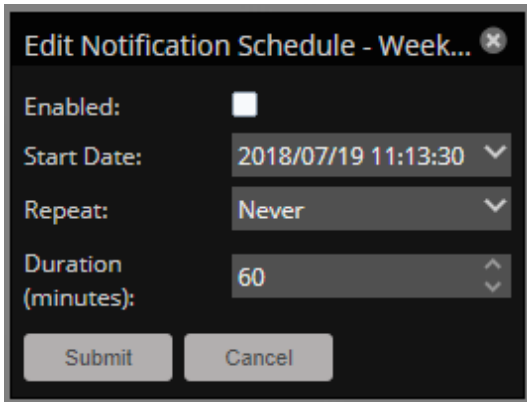
Schedules are bound to notifications, so the name can just describe the schedule.

6. Click Submit.

Editing a notification schedule

Follow these steps:

1. Navigate to EVENTS > TRIGGERS.
2. In the left column, click Notifications.
3. In the NOTIFICATION SCHEDULES area, double-click the name of the schedule to edit.



4. In the Edit Notification Schedule dialog box, configure the schedule.

Field	Description
Enabled	Activate or deactivate the schedule.
Start Date	The first date and time at which the schedule is active.
Repeat	The schedule frequency, one of the following options: <ul style="list-style-type: none">• Never• Daily• Every Weekday• Monthly• First Sunday of the Month
Duration (minutes)	The length of time for the schedule.

5. Click Submit.

Defining notification content

This page describes the options on the Content tab of the Edit Notification dialog box, which vary depending on the notification's action.

Notification emails provide hyperlinks to the Resource Manager UI. The address in those hyperlinks is driven by the `zopeurl` configuration directive in the `zenactiond` service. To edit this option, locate the `zenactiond` service in Control Center and edit the `zenactiond.conf` configuration file.

Email content

This section applies to both the AWS Email Host and the Email actions.

- [Common email fields](#) applies to both the AWS Email Host and the Email actions.
- [SMTP fields](#) applies to the Email action only.
- [Amazon SES fields](#) applies to the AWS Email Host action only.

Common email fields

The fields in the following table are present in all email actions.

Field	Description
Body Content Type	The email message format, HTML or plain text.
Message (Subject) Format	The Subject line of the email message. The default value includes content variables . Best practice: Include the event severity variable (<code>\${evt/severity}</code>) in this field.
Body Format	The body of the email message. The default value includes content variables . Best practice: To facilitate debugging, include variables for the name of the notification (<code>\${notification/name}</code>) and the name of the trigger that matched the event (<code>\${trigger/name}</code>) in this field.
Clear Message (Subject) Format	The Subject line of the clearing email message. The default value includes content variables . Best practice: Include the clearing event severity variable (<code>\${clearEvt/severity}</code>) in this field.
Body Format	The body of the clearing email message. The default value includes content variables . Best practice: To facilitate debugging, include variables for the name of the notification (<code>\${notification/name}</code>) and the name of the trigger that matched the clearing event (<code>\${trigger/name}</code>) in this field.
Skip fails in TALES evaluation ?	Mark this check box to ensure that the email message is sent even if the TALES expressions included in the message do not resolve.
From Address for Emails	The email address to use in the From line of the email message.

SMTP fields

The fields in the following table are present only for the Email action.

When the SMTP fields are empty, Resource Manager uses its default SMTP server to send email. Provide values only to use an alternate SMTP server.

Field	Description
SMTP Host	The hostname or IP of an SMTP server host to use instead of the default Resource Manager SMTP host.
SMTP Port (usually 25)	The port on which an SMTP server is listening for email to send.
Use TLS?	Mark this check box to encrypt emails.
SMTP Username (blank for none)	The username of an account on the non-default SMTP server.
SMTP Password (blank for none)	The password of the account on the non-default SMTP server.

Amazon SES fields

The fields in the following table are present only for the AWS Email Host action.

Both the sender and the recipients of emails sent through this action must be verified SES users in the target AWS region.

Field	Description
AWS Account Name	The name of an authorized Amazon SES account.
AWS Region	The AWS region to use.
AWS Access Key	The ID of the authorized account's AWS access key.
AWS Secret Key	The secret key of the authorized account's AWS access key.

Command content

This section applies to the Command action.

Use the Command action sparingly. Frequent use can degrade the performance of the event processing service.

Field	Description
Command Timeout (seconds)	The amount of time to wait for the command to return an exit value.
Command	The shell command to run on a remote host, through an SSH connection. This field supports content variables . The connection information and the name of the target host must be present in the properties of the triggering event, and the host must accept incoming SSH connections.
Clear Command	The shell command to run on a remote host, through an SSH connection, when a clearing event is received. This field supports content variables . The connection information and the name of the target host must be present in the properties of the clearing event, and the host must accept incoming SSH connections.
Environment variables	Environment variables to pass to the shell that performs the command, if any. This field supports content variables .

Syslog content

Use the Syslog action to forward event information to a receiving `syslog` system.

Field	Description
Syslog Host	The hostname or IP address of a <code>syslog</code> server host in your environment.
Syslog Port (usually 514)	The port on which the <code>syslog</code> server is listening.
Protocol	The network protocol to use, UDP or TCP.
Syslog Facility	The <code>syslog</code> facility code to include in the message.

SNMP trap content

This section applies to the SNMP Trap and SNMP Trap w/Impact actions. Use these actions to forward event information to an SNMP network management station (NMS).

The OIDs that are used in the SNMP traps that Resource Manager sends are defined in the [Zenoss MIB](#). When the action is SNMP Trap w/Impact, [additional OIDs](#) are included.

Field	Description
SNMP Trap Destination	The hostname or IP address of an SNMP NMS in your environment.
SNMP Community	The SNMP trap community string to include.

SNMP Version	The SNMP version to use. SNMPv1 (v1) and SNMPv2 (v2c, the default) are supported.
SNMP Port	The port on which the NMS in your environment is listening for traps. The default is 162.

WinCommand content

This section applies to the WinCommand action.

Field	Description
Windows CMD Command	<p>The standard CMD or PowerShell command to run on a remote host. This field supports content variables.</p> <p>The connection information and the name of the target host must be present in the properties of the triggering event, and the host must accept incoming connections.</p>
Clear Windows CMD Command	<p>The standard CMD or PowerShell command to run on a remote host when a clearing event is received. This field supports content variables.</p> <p>The connection information and the name of the target host must be present in the properties of the clearing event, and the host must accept incoming connections.</p>

Notification content variables

The content of email and command notifications can include information from events in the following form:

```
'${objectname/objectattribute}'
```

Do not escape event command messages and event summaries. For example, write this command as: `${evt/summary}` (rather than `echo '$evt/summary'`).

Object names can be `evt`, `evtSummary`, or `urls`; or for clearing event context, `clearEvt` and `clearEventSummary`. For each object name, the following lists show valid attributes (for example, `'${evt/DevicePriority}'`):

Attributes of `evt` and `clearEvt`

Value	Description
DevicePriority	value of the priority of the device
agent	Typically the name of the daemon that generated the event. For example, an SNMP threshold event has <i>zenperfsnmp</i> as its agent.
clearid	id of the event this clear event will clear
component	component this event is related to
count	how many times this event occurred
created	when the event was created
dedupid	dynamically generated fingerprint that allows the system to perform de-duplication on repeating events that share similar characteristics
device	device this event is related to
eventClass	class of this event
eventClassKey	Free-form text field that is used as the first step in mapping an unknown event into an event class.
eventGroup	Free-form text field that can be used to group similar types of events. This is primarily an extension point for customization. Currently not used in a standard system.
eventKey	Free-form text field that allows another specificity key to be used to drive the de-duplication and auto-clearing correlation process.
eventState	state of the event
eid	unique id for the event
facility	the syslog facility
firstTime	First time that the event occurred.
ipAddress	IP address
lastTime	Most recent time that the event occurred.
manager	value of manager
message	a message communicated by the event
ntheid	windows event id
ownerid	owner id
priority	syslog priority
prodState	The production state of the device.
severity	The integer that identifies the event severity level .
severityString	The descriptive label that identifies the event severity level .

stateChange	The last time that the event status changed.
status	The status of the event.
summary	A brief message summarizing the event.

Attributes of eventSummary and clearEventSummary

Some of the values in the following table are direct duplicates of evt attributes. For example, uuid -> evt.evid.

Value	Description
uuid	evt.evid
occurrence	evt.count
status	evt.eventState
first_seen_time	evt.firstTime
status_change_time	evt.stateChange
last_seen_time	evt.lastTime
count	evt.count
current_user_uuid	UUID of the user who acknowledged this event
current_user_name	name of the user who acknowledged this event
cleared_by_event_uuid	UUID of the event that cleared this event (for events with status == CLEARED)
notes	event notes
audit_log	event audit log
update_time	last time a modification was made to the event
created_time	evt.lastTime
fingerprint	evt.dedupid
event_class	evt.eventClass
event_class_key	evt.eventClassKey
event_class_mapping_uuid	If this event was matched by one of the configured event class mappings, it contains the UUID of that mapping rule.
actor	event actor
summary	evt.summary
message	evt.message
severity	evt.severity
event_key	evt.eventKey
event_group	evt.eventGroup
agent	evt.agent
syslog_priority	evt.priority
syslog_facility	evt.facility
nt_event_code	evt.ntevid
monitor	evt.monitor
tags	event tags

Attributes of urls

Value	Description
ackUrl	URL for acknowledging the event
closeUrl	URL for closing the event
reopenUrl	URL for reopening the event
eventUrl	URL for viewing the event
eventsUrl	URL for viewing events for the relevant device, or all events

ZenPacks can define additional notification actions and can extend the context that is available to notifications to add objects or attributes.

Setting individual notification permissions

You can grant permissions to individual users or groups.

- **Write** - Select this option to grant the user or group permission to update the notification.
- **Manage** - Select this option to grant the user or group permission to manage the notification.

You can manually enter in the name of a user or group, or select one from the list of options.

Managing users in Resource Manager

Every user within Resource Manager has a unique user ID, which allows an administrator to assign group permissions and alerting rules that are unique to each user. Unique IDs also help ensure secure access to the system.

To create and manage user accounts, you must be logged in as a user with Manager or ZenManager privileges.

- [Creating user accounts](#)
- [Editing user accounts](#)
- [User groups](#)
- [Roles and permissions](#)
- [Device access control lists](#)

Creating user accounts

To create a user account:

1. From the Navigation menu, select **ADVANCED**. The Settings page appears.
2. In the left panel, select **Users**. The users and groups administration page appears.
3. From the Action icon, select **Add New User**. The Add User dialog appears.
4. In the Username field, enter a unique name for the account.
5. In the Email field, enter the user account email address. Any alerts that you set up for this user will be send to this address.
6. Click **OK**. The user appears in the User List.

After creating the account, edit the account to provide a password and additional user details.

Editing user accounts

Follow these steps to edit user account information:

1. In the Users list, click the name of the user to edit.

The screenshot shows the 'ZenUsers > admin' page. At the top, there is a 'Reset Password' button with the text 'Automatically generate a new password and send it to the email listed below.' Below this is the 'USER PREFERENCES' section with a 'Reset Preferences' button and the text 'Reset all preferences such as grid columns and filters to their default values.' The main section is 'USER SETTINGS', which includes a 'Roles' dropdown menu (currently showing 'Manager', 'ZenManager', 'ZenOperator', 'ZenUser'), 'Groups', 'Email', 'Pager', 'Default Page Size' (set to 40), 'Default Admin Role' (set to 'ZenUser'), 'Network Map Start Object', 'Time Zone' (set to 'America/Chicago'), 'Date format' (set to 'year/month/day'), and 'Time format' (set to '24h'). At the bottom, there are fields for 'Set New Password', 'Confirm New Password', and 'Current Password for admin', along with a 'Save Settings' button.

2. On the Edit page, modify settings as desired.


Option	Description
Reset Password	Allows users to reset their own password.
User Preferences / Reset Preferences	Revert all settings to their default values.
Roles	Assign one or more roles to the user. To edit or assign roles, you must be Manager or ZenManager privileges. For more information, see Roles and permissions .
Groups	Add the user to one or more Resource Manager groups.
Email	The email address of the user.
Pager	The phone number of the user's pager device.
Default Page Size	The default number of rows to include in Resource Manager tables. The default is 40 rows.
Default Admin Role	The default role for administered objects associated with the user.
Network Map Start Object	The IP address of a network, or the hostname or IP address of a device to display in the network map .
Time Zone	The time zone to use on all Resource Manager charts and graphs.
Date format	The date format to use on all Resource Manager charts and graphs.
Time format	The time format to use throughout Resource Manager.
Set New Password / Confirm New Password	The user password.

3. In the Current Password for Admin field, enter the administrative user's password, and then click Save Settings.

Associating objects with specific users

You can associate any object in the system with a particular user, for monitoring or reporting purposes. Once associated with a user, you can then assign the user a specific role that applies to his privileges with respect to that object. To create an object association:

1. From ADVANCED > Settings, select Users in the left panel.
2. Click the name of a user.
3. From the Edit page, select Administered Objects in the left panel.



4. Select an object type from the Administered Objects Action menu. You can add:
 - Device
 - Device class
 - System
 - Group
 - Location
5. Specify the component you want to add as an administered object, and then click OK. The object appears in the Administered Devices list for the user.




6. **Optional:** Change the role that is associated for this user on this object. Note: The default role assigned to the user for an administered object is specified by the Default Admin Role field on the Edit page.
7. Click Save to save changes.

Adding administrators

You also can associate an object with a user by adding an administrator to the object. Perform the following:

1. Navigate to the object you want to add to the user's list of administered objects.
2. Select Device Administration.



3. Click the Add Administrator icon in the Administrators area. The Add Administrator dialog box appears.
4. Select an administrator from the list and change the role if desired, then click SUBMIT.

The administrator appears in the object's Administrators list. The object is added to the administrator's Administered Objects list.

Changing the admin account password with the CLI

You can update the password of the Resource Manager admin account with the command-line interface.

Follow these steps:

1. Log in to the Control Center master host as a user with `serviced` CLI privileges.
2. Start an interactive session in a Zope container as the `zenoss` user.

```
serviced service attach zope/0 su - zenoss
```

3. Start the `zenpass` utility, and then answer the prompts with the new password:

```
(zenoss) [zenoss@f3323bee63f7 ~]$ zenpass  
Enter the new password for the Zenoss "admin" user:
```

User groups

User groups allow you to aggregate rules and apply them across multiple user accounts.

Viewing user groups

To view user groups, select **ADVANCED > Settings**, and then select **Users** from the left panel.

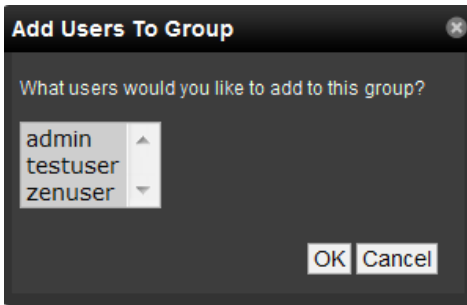
The groups area shows each user group and the users assigned to that group.

Creating user groups

You can create user groups to aggregate rules and apply them across multiple user accounts.

To create a user group:

1. Navigate to **ADVANCED > Settings**.
2. In the left panel, select **Users**. The **Users** page appears.
3. From the **Groups** area Action menu, select **Add New Group**. The **Add Group** dialog box appears.
4. In the **Group** field, enter a name for this user group, and then click **OK**. The group name appears in the **Groups** list.
5. Click the name of the group you created. The **Users in Group** page appears.
6. From the Action menu, select **Add User**.



7. From the **User** list of selections, select one or more users you want to add to the group, and then click **OK**. The user or users you select appear in the list of users for this group.

You also can choose administered objects and alerting rules for this user group. These alerting rules will apply to all users in the group. The user's original alerting rules and objects will also apply.

Roles and permissions

Roles represent specific permissions that you can assign to users or groups of users.

The following table shows the available roles, from lowest to highest privilege level.

Role	Permissions
ZenUser	Global read-only access to system objects.
ZenOperator	Event management access. Combine ZenOperator and ZenUser to grant read-only access, but also allow users to acknowledge and close events, move events to history, and add log messages to events. You can associate the ZenOperator role with an individual device, a device class, or a group of devices.
ZenManager	Global read-write access to system objects.
Manager	Global read-write access to system objects, and read-write access to the Zope object database.

Device access control lists

Resource Manager supports fine-grained security controls. For example, this control can be used to give limited access to certain departments within a large organization or limit a customer to see only his own data. A user with limited access to objects also has a more limited view of features within the system. As an example, most global views, such as the network map, event console, and all types of class management, are not available. The device list is available, as are the device organizers: systems, groups, and locations. A limited set of reports can also be accessed.

Permissions and roles

Actions in the system are assigned permissions. For instance to access the device edit screen you must have the "Change Device" permission. Permissions are not assigned directly to a user; instead, permissions are granted to roles, which are then assigned to a user. A common example is the ZenUser role. Its primary permission is "View," which grants read-only access to all objects. ZenManagers have additional permissions such as "Change Device," which grants them access to the device edit screen. When you assign a role to a user using the Roles field (on the Edit page), it is global.

Administered objects

Device ACLs provide limited control to various objects within the system. Administered objects are the same as the device organizers: Groups, Systems, and Locations and Devices. If access is granted to any device organizer, it flows down to all devices within that organizer. To assign access to objects for a restricted user, you must have the Manager or ZenManager roles. The system grants access to objects is granted using the user's or user group's administered objects. To limit access, you must not assign a "global" role to the user or group.

Users and groups

Users and user groups work exactly as they would normally. See the section in the User Management section of this guide dealing with users and groups.

Assigning administered object access

For each user or group there is an Administered Objects selection, which lets you add items for each type of administered object. After adding an object you can assign it a role. Roles can be different for each object, so a user or group might have ZenUser on a particular device but ZenManager on a location organizer. If multiple roles are granted to a device through direct assignment and organizer assignment the resulting permissions will be additive. In the example above, if the device was within the organizer the user would inherit the ZenManager role on the device.

Portlet access control

In Resource Manager, portlet access can be controlled. This is important for device ACLs.

Viewing events for restricted mode users

A user in restricted mode does not have access to the global event console. The available events for the user can be seen under his organizers.

Example: Restricted user with ZenOperator role

The ZenUser role from the previous section allows read-only access to devices. By adding the ZenOperator role to specific devices, device classes, or groups of devices, a user will be able to acknowledge and close events, move events to history, and add log messages to events.

To add the ZenOperator role to specific devices, device classes, or groups of devices:

1. Select the user name whose role must be changed on certain devices.
2. In the left-hand pane, click Administered Objects.
3. Click the Action icon and choose the device, device class, or other device organizer to which you want to grant the ZenOperator role.
4. Select the ZenOperator role from the drop-down menu for the newly selected device, device class, or device organizer.

The user now has the ZenUser role for all devices in this instance, with the exception of the devices selected above which function under the ZenOperator role.

Detailed restricted screen functionality

Dashboard

By default, the dashboard is configured with three portlets:

- Object Watch List
- Device Issues
- Production State

These have content that will be restricted to objects for a given user.

Device list

The device list is automatically filtered to devices of a restricted user scoped to accessible devices. No menu items are available.

Device organizers

Device organizers control groups of devices for a restricted user. Every device added to the group will be accessible to the user. Permissions will be inherited down multiple tiers of a device organizer.

Reporting

Reports are limited to device reports and performance reports.

General administration and settings

Use the information and procedures in this section for troubleshooting and performance improvement purposes.

- [Audit logging](#)
- [User commands](#)
- [Changing events database connection information](#)
- [Rebuilding the events index](#)
- [Support bundles](#)
- [Working with the job manager](#)

Audit logging

Resource Manager tracks user actions in audit log files and maintains logged information in a format that is optimized for searching and reporting.

Audit logging information is written to a flat file located on the Control Center master host which is located at `/var/log/serviced/application-audit.log`. The information is also sent to Elasticsearch, and is may be viewed through Kibana from the Control Center browser interface on the Logs tab.

The retention rules of the `application-audit.log` file are governed by the `logrotate` configuration file located at `/opt/serviced/etc/logrotate.conf`. For more information about editing the `logrotate` configuration file, see [Control Center audit logging](#).

Examples

The following examples show logged messages for various user interactions.

Change device class

```
2017-07-14 20:17:52 INFO user=admin action=ChangeDeviceClass
  kind=Device device=/Devices/Server/Linux/devices/emailsrv05 device_name=emailsrv05
  deviceclass=/Devices/Server/SSH/Linux old_deviceclass=/Devices/Server/Linux
```

In this example, the admin user moved device `emailsrv05` from device class `/Server/Linux` to `/Server/SSH/Linux`.

Change threshold value

```
2017-07-14 20:53:19 INFO user=admin action=Edit kind=Threshold
  threshold="/Devices/Server/Microsoft/rrdTemplates/Device/thresholds/CPU Utilization"
  maxval=95 old_maxval=90 thresholdtype=MinMaxThreshold
```

In this example, the admin user edited the max value of threshold "CPUUtilization" on Microsoft servers from 90 to 95.

The zensendaudit utility

You can send custom log messages using the `zensendaudit` script.

1. Log in to the Control Center host as a user with `serviced` CLI privileges.
2. Attach to the Zope service as the `zenoss` user.

```
serviced service attach zope/0 su - zenoss
```

3. Send a message with the `zensendaudit` script.
Replace `MESSAGE` with the text to send to the audit log.

```
zensendaudit MESSAGE
```


For example, invoking `"zensendaudit Hello world"` results in the following audit log entry:

```
2017-07-14 19:10:18 INFO action=Log kind=Comment comment="Hello world"
  source=Shell process=zensendaudit logname=zenoss parameters="Hello world"
```

User commands

User commands allow you to execute arbitrary shell commands from Resource Manager. A user command is executed in its appropriate container rather than the remote device unless the command explicitly uses SSH to connect to the remote device.

You can define and run user commands on a device or organizer (device class, system, group, location, or component group). You also can define commands globally. The User Commands menu bar shows the various functions that can be used in the User Commands screen.



Defining global user commands

Global commands appear in the Commands list of options located at the top of the Devices page.


To define global user commands:

1. From the navigation menu, choose Advanced > Settings.
2. In the left panel, choose Commands.
3. In the Define Commands area, click the action menu and choose Add User Command.
4. In the Add User Command dialog box, specify a name for the command, and then click OK.
Only letters, numbers, and underscores are allowed in command names. Spaces are not allowed.
5. In the Define Commands page, specify a description of what the command will do.
6. In the Command section, enter the TALES expression-based command you want to run on the device.
7. Enter your system account password for confirmation, and then click Save.
The command is saved and added to the commands menu.

Global commands also can be edited from a specific device. Changes to a global command from a device are not limited to that device.

Running global user commands

To run a global user command, select one or more devices in the devices list, and then select a command from the Commands list of options.



Defining user commands for a single device

To define a user command for a device:

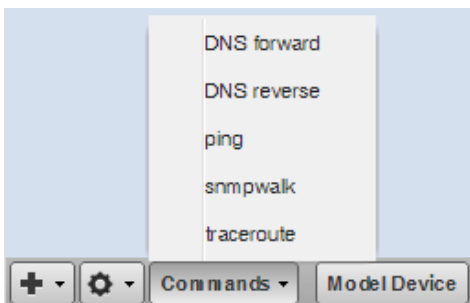
1. From the navigation menu, choose INFRASTRUCTURE.
2. In the device list, click a device name.
3. In the left panel, choose Device Administration.
4. In the User Commands area, click the Add a User Command icon.
5. In the Add New User Command dialog box, specify the following information about the user command:
 - **Name** - Name of the user command.
 - **Description** - Description of what the command will do.
 - **Command** - TALES expression-based command you want to run.

- **Confirm Password** - Enter your system account password for confirmation.
6. Click Submit. The command is saved and added to the user commands menu.
 7. **Optional:** Test the command by selecting the command from the list and clicking the Run icon.

Running user commands for a single device

To run a command defined for a single device:

1. Navigate to the INFRASTRUCTURE > DEVICES page.
2. Click the device name in the device list to open the Device Overview page.
3. Select the command from the Commands list of options located at the bottom of the page.



Defining user commands for all devices in an organizer

To define a user command for all devices in an organizer:

1. On the INFRASTRUCTURE page, select a device organizer in the devices hierarchy; for example, /Server/Linux.
2. Click Details.
3. In the left panel, select Device Administration.
4. In the User Commands area, click the Add a User Command icon.

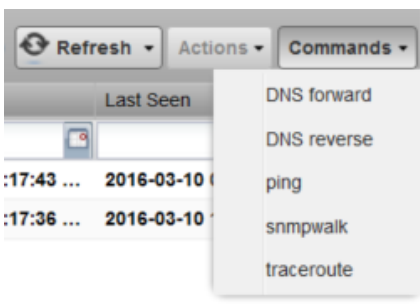
5. Enter the following information about the user command:
 - **Name** - Name of the user command.
 - **Description** - Description of what the command will do.
 - **Command** - TALES expression-based command you want to run.
 - **Confirm Password** - Enter your system account password for confirmation.
6. Click **Submit**. The command is saved and added to the user commands menu.
7. **Optional:** Test the command by selecting the command from the list and clicking the Run icon.

Running user commands for devices in an organizer

To run a command defined for devices in an organizer:

1. On the INFRASTRUCTURE page, select a device organizer.
2. Select one or more devices in the filtered view.

3. Select the command from the Commands list located at the top of the page.



Example

This example shows how to create an echo user command. You can see the use of TALES expressions in the definition of this command.

1. Add a command called "echoDevice"
2. In the command definition, echo the name and IP address of the device:

```
echo name = ${here/id} ip = ${here/manageIp}
```

In a TALES expression, here is the object against which the expression is executed. Some TALES expressions in the system have other variables (such as evt for event, and dev or device for the device). For more information, see [TALES expressions](#).

3. Select a device and then run the command.
4. Edit the command to add more information:

```
echo name = ${here/id} ip = ${here/manageIp} hw = ${here/getHWProductName}
```

5. Run the command against a group of devices and view the command outputs.

Changing events database connection information

To edit events database connection settings, make changes in the `zeneventserver.conf` file. You can edit the file directly, or run a configuration script.

Configurable database connection settings are:

- JDBC Hostname (`zep.jdbc.hostname`) - Specify the IP address of the host.
- JDBC Port (`zep.jdbc.port`) - Specify the port to use when accessing the events database.
- JDBC Database Name (`zep.jdbc.dbname`) - Specify the database name.
- JDBC Username (`zep.jdbc.username`) - Specify the user name for the database.
- JDBC Password (`zep.jdbc.password`) - Specify the password for the database.
- To edit these values, run the `zeneventserver` configuration script, as follows:

```
zeneventserver-config -u zep.jdbc.Name=Value
```

Where `Name` is the partial setting name and `Value` is the value you want to specify for the setting.

Rebuilding the events index

If you encounter inconsistent search results, you can rebuild the events index.

1. Log on to the Control Center master host as a user with `serviced` CLI privileges.
2. Stop `zeneventserver`:

```
serviced service stop zeneventserver
```

3. Delete the index data:

```
export SERVICE_ID=$(serviced service status Zenoss.resmgr | sed -n '2p' | awk {'print $2'})
export SVCROOT=/opt/serviced/var/volumes/$SERVICE_ID
rm -rf $SVCROOT/zeneventserver/index
```

4. Start `zeneventserver`:

```
serviced service start zeneventserver
```

Depending on the number of events in the database, it may take a significant amount of time for indexing to complete. Until every event is indexed, the number of events shown in the event console may be inconsistent.

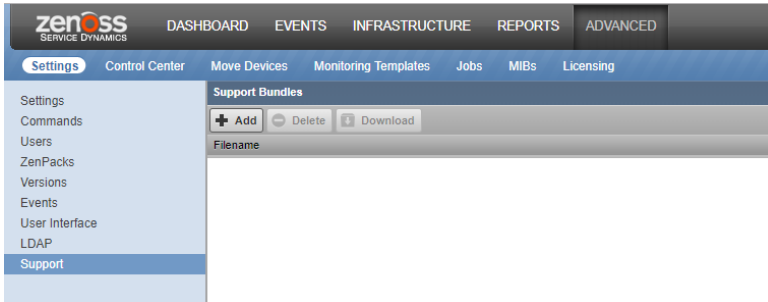
Support bundles

If you contact Zenoss Support about an issue with your Resource Manager system, you might be asked to send a support bundle or serviced and docker logs to help the technicians find a solution to your issue.

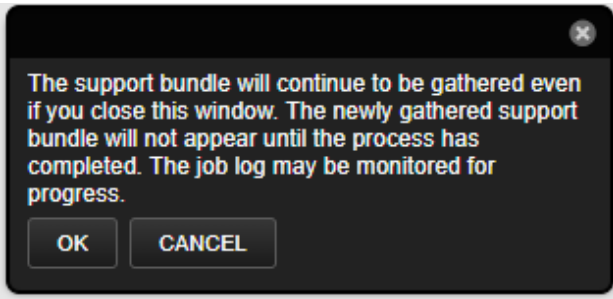
Creating a support bundle

Follow these steps:

1. As an administrator, navigate to ADVANCED > Support.



2. Click Add.



3. Click OK.



4. Click Close.
The support bundle will continue to be gathered. You can monitor the job log for progress and continue working while the bundle is being assembled.
5. When the completed support bundle appears in the Support Bundles window, click on its row and click Download to send the support bundle ZIP file to your browser's downloads folder.

Deleting support bundles

Occasionally clean up your system by deleting old support bundles that are no longer needed for troubleshooting purposes.

1. As an administrator, navigate to **ADVANCED > Support**.
2. In the **Support Bundles** window, click the row of the support bundles to be deleted. Be sure that you have selected the correct bundles because no confirmation window will appear.
3. After verifying your selections, click **Delete**.

Gathering Control Center and Docker logs

You may need to gather Control Center and Docker logs separately to help troubleshoot an issue, since these logs cannot be collected by the support bundle process.

To get Control Center and Docker logs:

1. Log in to a Control Center host as root or as a user with superuser privileges.
2. Dump the Control Center (`serviced`) log to a file:

```
journalctl -lu serviced > /tmp/serviced.log
```

3. Dump the Docker log to a file:

```
journalctl -lu docker > /tmp/docker.log
```

4. Create a compressed archive of the two log files:

```
tar czf /tmp/serviced-docker-logs.tgz /tmp/serviced.log /tmp/docker.log
```

5. Send the archive file to Zenoss Support.

Working with the job manager

The job manager runs background tasks, such as discovering a network or adding a device. When you ask the system to perform one of these tasks, it adds a job to the queue. Jobs are run by the zenjobs daemon.

Not all actions are performed in the job manager. Some jobs are run automatically in the foreground. Others, such as moving devices, depend on user interface configuration settings.

When running jobs in the foreground, do not navigate away from the current page until the action completes.

Viewing the job manager

In the browser interface, select **ADVANCED > Settings > Jobs**.

The screenshot shows the 'Background Jobs' table with the following data:

Status	Description	Scheduled	Started	Finished	Created By
Success	Create 10.171.100.14 under /Network/Cisco/6500	2 days ago	2 days ago	2 days ago	admin
Failure	Discover and model device 10.171.100.14 as /Networ...	2 days ago	2 days ago	2 days ago	admin
Success	Create 10.171.100.92 under /Network/Cisco/Nexus/70...	2 days ago	2 days ago	2 days ago	admin
Failure	Discover and model device 10.171.100.92 as /Networ...	2 days ago	2 days ago	2 days ago	admin
Success	Create 10.171.100.88 under /Network/Cisco/Nexus/10...	2 days ago	2 days ago	2 days ago	admin
Failure	Discover and model device 10.171.100.88 as /Networ...	2 days ago	2 days ago	2 days ago	admin
Success	Create 10.171.54.9 under /Network/Cisco	2 days ago	2 days ago	2 days ago	admin
Success	Discover and model device 10.171.54.9 as /Network/...	2 days ago	2 days ago	2 days ago	admin
Success	Create 10.171.100.107 under /Network/Cisco/Nexus/6...	2 days ago	2 days ago	2 days ago	admin
Failure	Discover and model device 10.171.100.107 as /Netwo...	2 days ago	2 days ago	2 days ago	admin
Success	Create 10.171.53.1 under /Network/Cisco	2 days ago	2 days ago	2 days ago	admin
Success	Discover and model device 10.171.53.1 as /Network/...	2 days ago	2 days ago	2 days ago	admin

DISPLAYING 1 - 13 of 33 ROWS

Job Log

Log file: [/opt/zenoss/log/jobs/c4583592-2506-4466-b909-854ebfa20258.log](#)

2015-08-04 17:59:20,643 INFO zen.Job: Job c4583592-2506-4466-b909-854ebfa20258 (Products.ZenModel.ZDeviceLoader.CreateDeviceJob) received

2015-08-04 17:59:20,657 INFO zen.Job: Starting job c4583592-2506-4466-b909-854ebfa20258 (Products.ZenModel.ZDeviceLoader.CreateDeviceJob)

2015-08-04 17:59:21,051 INFO zen.Job: Job c4583592-2506-4466-b909-854ebfa20258 finished with result /zport/dmd/Devices/Network/Cisco/Nexus/7000/devices/10.171.100.92

The jobs list displays the following information about the jobs:

- Status - Shows the current job status. Status options are Pending (waiting for zenjobs to begin running), Running, Succeeded, and Failed.
- Description - Provides a description of the job.
- Scheduled - Shows when the job was scheduled to begin.
- Started / Finished - Provide information about the time period in which the job ran.
- Created By - User that created the job.

The lower section of the page displays the job log for the job that you select in the list. You can also view the information in the log file.

Stopping and deleting jobs

To stop a job, select it in the list, and then click **Abort**. The zenjobs daemon will not run the job.

To remove a job from the system, select it and then click **Delete**.

Configuring jobs

When you move devices, you can choose whether the action is performed immediately or as a job. By default, if you select five or more devices, the move action is performed as a job. To adjust this setting:

1. In the browser interface, select **ADVANCED > Settings > User Interface**.
2. Enter a value for **Device Move Job Threshold**, and then click **Save**.

Running the zenjobs daemon

You can stop and start the zenjobs daemon from the command line, and from Advanced > Settings (Daemons selection).

Troubleshooting Resource Manager

This [documentation category](#) is a work in progress; be sure to check back regularly for more content.

This section of documentation contains information designed to help Zenoss Resource Manager and Zenoss Community Edition Administrators troubleshoot problems they may encounter with those applications.

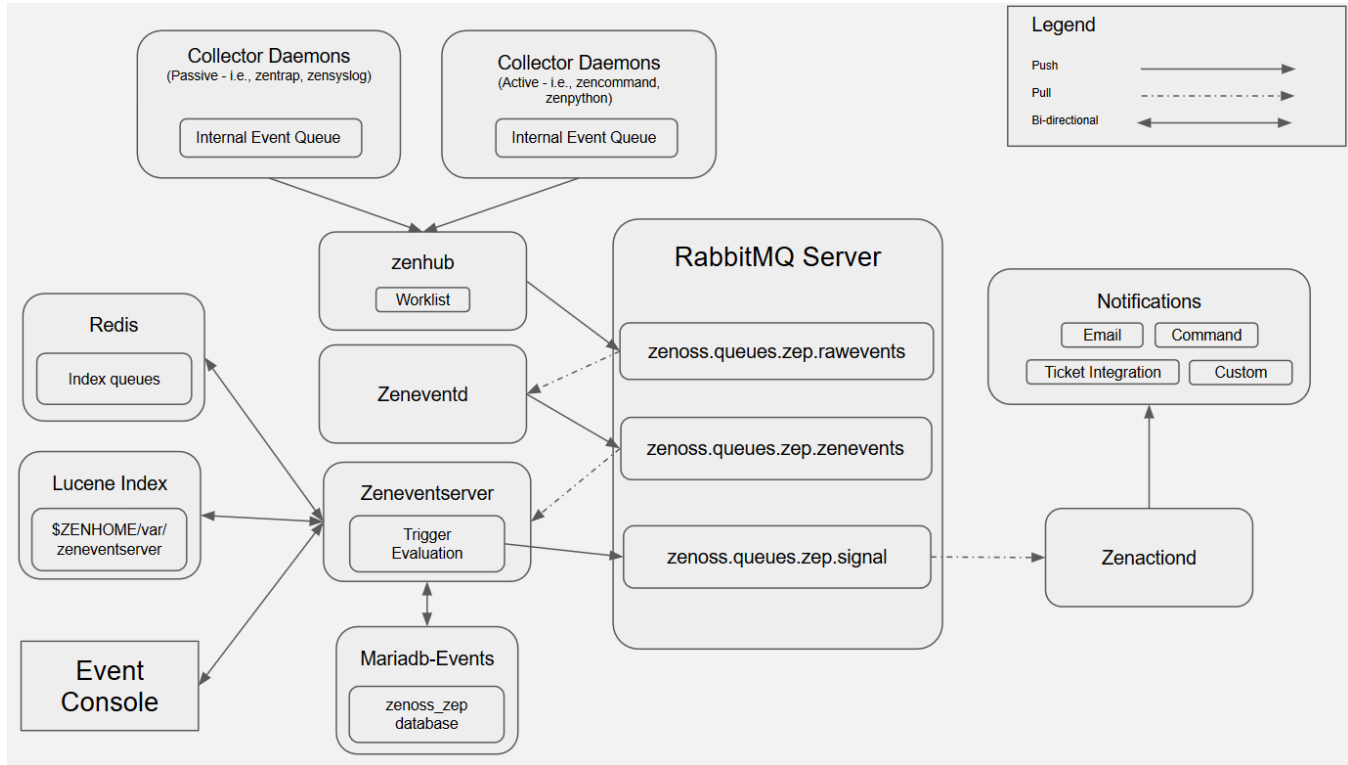
- [Data pipelines](#)
 - [Event pipeline](#)
 - [Troubleshooting event flow](#)
 - [Model pipeline](#)
 - [Troubleshooting modeling](#)
 - [Performance data pipeline](#)
 - [Troubleshooting performance collection and retrieval](#)
- [How to read zenhub log messages](#)

Data pipelines

Zenoss Resource Manager is an application that moves data. It collects data from infrastructure, manipulates, and moves that data through a series of daemons that make up what we refer to as pipelines. Understanding how these pipelines work will make troubleshooting application failures much easier.

- [Event pipeline](#)
 - [Troubleshooting event flow](#)
- [Model pipeline](#)
 - [Troubleshooting modeling](#)
- [Performance data pipeline](#)
 - [Troubleshooting performance collection and retrieval](#)

Event pipeline



Event Pipeline Overview Video

Event Processing Explained

- Events are generated by collector daemons
 - Passive collection is performed by daemons like zentrap, zensyslog, etc.
 - Active collection is performed by daemons like zenpython, zencommand, etc.
- Queued events are sent to ZenHub
 - If zenhub is not available to receive queued events, they will be held internally
 - Collector daemons keep an internal event queue
 - Not persistent
 - Defaults to 5000
 - Rotating buffer; holds only the 5000 most recent events
 - Adjust by via maxqueueelen setting in collector daemon config
- Zenhub acts as an aggregator for events from collectors
 - zenhub's parent thread receives events from collectors
 - Events are validated; they must contain at least a device (which doesn't necessarily have to be added into RM for monitoring), a severity, and a message or else they're dropped
 - Event processing tasks are queued internally in zenhub's worklist
 - zenhub workers consume tasks from the worklist
 - **SendEvents** tasks are worked by publishing events to the **zenoss.queues.zep.rawevents** queue in RabbitMQ
- Zeneventd consumes events from zenoss.queues.zep.rawevents
 - Maps events from /Unknown to appropriate class based on mappings
 - Event enrichment
 - Provides 'device' and 'component' contexts in transforms
 - Populates device data in events
 - ProductionState
 - Groups/Systems/Locations
 - Device Priority
 - etc
 - Applies transform code based on event class membership
 - Executes zeneventd post-plugins
 - Publishes contextualized events to **zenoss.queues.zep.zenevents**
 - Zeneventserver consumes events from zenoss.queues.zep.zenevents

- Serves as the heart of the event processing engine and has multiple roles
 - Stores events in the zenoss_zep database
 - MariaDB-Events container
 - Indexes stored events in Lucene indexes
 - \$ZENHOME/var/zeneventserver
 - Handles event aging and archiving
 - Provides the back-end for the event API
 - Serves the event console and archive UI elements
 - Serves the event 'rainbow' functionality in the UI
 - Maintains a list of ping-down devices which zenhub uses when it builds configs
 - Handles heartbeat monitoring
 - Evaluates triggers and publishes to zenoss.queues.zep.signal when a match is found
- Zenactiond consumes events from zenoss.queues.zep.signal queue
 - Processes notifications, which can:
 - Email a user
 - Run a command to perform corrective action
 - Send a syslog message
 - Send an SNMP trap
 - Generate or update a support ticket
 - Integrate with other custom solutions
 - ServiceNow
 - RemedyITSM
 - etc

Troubleshooting event flow

Identifying Bottlenecks

If you suspect you're having a problem with the event pipeline, check rabbitmq first:

- rabbitmq public endpoint
- **rabbitmqctl** in the rabbitmq container
 - `rabbitmqctl list_queues -p /zenoss`
 - `rabbitmqctl list_queues -p /zenoss messages consumers name`

Rawevents Troubleshooting

Backups in the rawevents queue are typically:

- An event flood
 - Look at collector performance graphs (event queues graph) to identify the source of the flood
 - Look in the event console for new syslog or snmp trap messages that are rapidly incrementing in count
 - As a last-ditch effort you can turn off zensyslog/zentrap temporarily to stop a flood
- Slowness in zeneventd
 - Look in the zeneventd logs for long-running transform messages
 - You may need to optimize your event transforms
 - You may need more zeneventd workers or instances
- Throttling in RabbitMQ (really only an issue in 4.x)
 - Make sure the RabbitMQ container has at least 1GB of memory available
 - Make sure that /var in the RabbitMQ container has at least 1GB of disk space available

Zenevents Troubleshooting

Backups in the zenevents queue are typically

- An event flood or RabbitMQ throttling
 - Use the steps in the Rawevents Troubleshooting section above
- A problem in zeneventserver
 - Check the zeneventserver log
 - If there's an error coming from the database, it will normally show up as a jdbc exception here
 - MariaDB/MySQL Error codes are often easy to diagnose with google-fu
 - If there's an error indexing events, it may be necessary to rebuild the Lucene indexes
 - Non-graceful shutdowns and other causes of data corruption can lead to index corruption, which may require rebuilding the zep Lucene indexes
- Slowness in Mariadb-Events
 - Check database tuning
 - Make sure your **innodb_buffer_pool_size** is adequate for the size of your zep db
 - Make sure there's not a cpu/memory/disk IO constraint
 - Use **zencheckdbstats** to identify tuning issues
- Slowness in zeneventserver
 - Do you have more than 200 triggers enabled?
 - You may need to increase your trigger cache size in zeneventserver.conf
 - Does zeneventserver need more memory?
 - You can increase the heapsize by increasing the RAM commitment


Signal Troubleshooting

Backups in the signal queue are typically

- An event flood
 - Use previously-defined steps for troubleshooting floods
- An external failure

- Put zenactiond in debug mode
- Check the zenactiond log
 - Slowness or failures to send notifications will show up here
- Too many notifications for your zenactiond workers to keep up
 - Add more workers to zenactiond and restart it, or add more instances of zenactiond

Model pipeline



Modeling Process Explained

- Zenmodeler performs most modeling tasks
- On startup, and every configuration flush interval (6 hours) zenmodeler gets a ModelerService config from zenhub
- zenhub pulls device data (ip address, credentials, etc) and modeler plugin bindings from ZODB (and leveraging solr) and generates a ModelerService config
- zenhub serves that config to zenmodeler
- Zenmodeler uses the ModelerService config as a list of tasks it needs to complete
- Zenmodeler executes modeler plugins
- Modeler plugins retrieve data from devices
- Modeler plugins format that data and construct an ObjectMap
 - an ObjectMap is a key:value object that contains component data
- Modeler plugins return the ObjectMap to zenmodeler
- Zenmodeler builds a DataMap using ObjectMaps from one or more modeler plugins and returns it to zenhub
- zenhub executes an applyDataMap task and commits model data to ZODB and the global catalog via solr

Troubleshooting modeling

Failure to collect data & Failure to process data

- Run zenmodeler in debug mode to see what data is being returned
 - First, connect to the appropriate zenmodeler container

```
serviced service attach zenmodeler
```

- Switch to the 'zenoss' user

```
su - zenoss
```

- Run zenmodeler in debug mode against the device

```
zenmodeler run -v10 -d $DEVICENAME --monitor=$COLLECTORNAME
```

- Optionally, you can specify the 'collect' switch to run a specific plugin

```
zenmodeler run -v10 -d $DEVICENAME --monitor=$COLLECTORNAME --collect=zenoss.snmp.IpInterface
```

Failure to connect to zenhub?

- Is zenhub running? Is it overloaded?
 - Check health checks
 - Look in the zenhub container; put it in debug mode and see if there are any messages like "All workers are busy" as this could indicate that you need more zenhub workers or instances

Failure to load a modeler plugin?

- Is this a new ZenPack?
 - Have you installed a ZenPack without restarting zenoss.resmgr services?

Failure to collect data - Auth/access issues?

- Check credentials and ZenPack documentation for end-device configuration

Failure to collect data - Incorrect modeler plugin?

- Does this modeler plugin support the device you're trying to model?


Failure to parse data

- Look for tracebacks - normally things like KeyError exceptions in modeler plugins are indicative of a failure to parse data
 - These types of failures are often indicative of either something wrong with the data the device is returning to us (potentially an end-device compatibility problem), or a potentially incorrect assumption we've made in code

Failure to apply data to the model

- Look at zenhub; it may be overloaded
- Bad relationships
 - Try deleting and re-adding the device
- Make sure the device was created with the correct zPythonClass
 - If you're running a WinRM modeler plugin on a device that was created in a class outside the /Server/Microsoft device class, for example, the device object's class may not possess the required attributes to accept that model data

Performance data pipeline



Performance Collection and Storage

Collector daemons gather metrics via their specific protocols

- Metrics are written as a collection of key-value pairs and tags in the following format:

```
2017-02-15 19:45:39,225 DEBUG zen.MetricWriter: publishing metric 10.90.36.137/cpu_ssCpuSystemPerCpu 162450
1487187939 {'device': '10.90.36.137',
'contextUUID': '12d97430-f7bc-4073-91f8-6743d3ae94a1', 'key': 'Devices/10.90.36.137'}
```

- Included in the message are: device and datapoint, value, and timestamp
- Collected metrics are sent to CollectorRedis which acts as a queue.
 - All collector daemons on a collector share a CollectorRedis instance.
- MetricShipper consumes from CollectorRedis and publishes to MetricConsumer
 - Traffic is proxied through zproxy
 - Authentication for zproxy is managed by zauth
- MetricConsumer acts as the aggregator for metrics from ALL collectors
- It then forwards metrics to the OpenTSDBWriter
- OpenTSDBWriter is responsible for actually writing the metrics to HBase

Performance Retrieval

- User interacts with RM UI through a Zope instance
- API request for performance data is made to CentralQuery
 - Request is proxied through zproxy
 - zproxy makes a request to zauth if the request doesn't have an authentication cookie
 - Request includes device(s), component(s), datapoint(s), a range of timestamps, and (sometimes) a cookie
- Central Query forwards request to OpenTSDB Reader
- OpenTSDB Reader retrieves data from HBase

Troubleshooting performance collection and retrieval

Performance Collection/Storage Troubleshooting

Sometimes performance collection can fail

- Check for events; they may point you to a failure
 - Verify credentials, connectivity, required target device configuration
- Run collector daemon in debug mode
 - See if values are being collected
 - If they are, check for failures to send data to redis
 - If they're not, check for error messages that may indicate where it's failing
 - If you're getting 500 errors, check zauth and see if it's overloaded; if it is, you may need to add more instances
- Ensure performance data pipeline services are running and passing all health-checks
 - Failing that, look at failing services' logs to identify the source of the problem
 - It's possible that you have some sort of data corruption or the service didn't start correctly for some reason; the logs should help you identify the source of the problem

Performance Retrieval Troubleshooting

Sometimes performance retrieval (graph rendering) can fail

- If you have a graphing failure localized to a single device, check its associated collector daemon
- If you have graphing failure localized to an entire remote collector, check CollectorRedis and MetricShipper
- Make sure that your graph configuration is valid
 - Has it worked before?
 - Is the problem with all graphs or one graph?
- Make sure values are being collected and stored
 - You can check the OpenTSDB public endpoint directly for graph data
- Check health checks
 - Failing that, look at failing services' logs to identify the source of the problem
- Make sure that zauth isn't overloaded
 - Instances of zauth can be added to help distribute load

How to read zenhub log messages

The log of the zenhub service contains tags that are easily searched (but are not indexed by logstash). These tags are:

- **worker:** Identifies the worker in the form of *worklist-name_instance-id*. For example, "default_0" which identifies the "zenhubworker (default)" service instance 0.
- **service:** Names the service that will be executed.
- **method:** Names the method on the service that will be executed.
- **id:** Identifies the specific request to execute the service method.
- **status:** Either *success* or *failed*.
- **duration:** Number of seconds taken to execute the method.
- **lifetime:** Number of seconds from the point the call was received until the result was returned to the original caller (typically a collection service).
- **waited:** Number of seconds the call waited in the worklist before executing.

The following, common log messages are associated with service methods:

- A "Receive task" message is written when a request to execute a service method is received by zenhub.
- The "Begin task" and "Retry task" messages are written when zenhub instructs a zenhubworker to execute a service method. The "Retry task" message also indicates that a previous attempt to execute the service method did not complete (that is, failed for reasons unrelated to the service method itself).
- The "Completed task" message is written a zenhubworker returns a result from the service method execution. The status tag will indicate whether the method succeeded or failed.

Example log message:

```
2019-06-12 12:38:53,966 INFO zen.zenhub.server.WorkerPoolExecutor: Completed task service=ModelerService
method=getThresholdClasses id=9a47fba04d424c15b756ff599fea24a9 worker=default_0 status=success duration=0.02
lifetime=0.03
```

Updating Resource Manager

This section includes update instructions for updating the following versions of Resource Manager to 6.4.1:

6.4.0, 6.3.2, 6.2.1, 6.2.0, 6.1.x, 6.0.x, 5.3.3

Tips for a successful upgrade:



- Before proceeding, review the [release notes](#); in particular, the [Considerations and workarounds](#) and [Known issues](#) sections.
- For optimum results, review update procedures before performing an update.
- Keep this page open, and open new tabs or windows for each update procedure.

Updating single-host, appliance-based deployments

Use the instructions in this column to update single-host deployments of Resource Manager based on virtual appliances. You can update Control Center and Resource Manager independently. The software includes updates of the virtual machine operating system.

Preparing to update

1. [Prepare your deployment for update](#)
2. [Download the update ISO file](#)
3. [Attach the update ISO](#)
 - [Attaching an update ISO with vSphere](#)
 - [Attaching an update ISO with Hyper-V](#)
4. [Modify MariaDB passwords, if necessary](#)
5. [Update the Dashboard ZenPack](#)
6. [Stop Resource Manager](#)
7. [Remove Docker containers on the master host](#)

Updating the master host

1. [Update Control Center](#)
2. [Update Resource Manager](#)

Post-update procedures

1. [Remove the pre-upgrade snapshot](#)
2. [Verify zenhubworker services updated successfully](#)
3. [Start Resource Manager](#)
4. [Optional: Install a patch for Linux modeling](#)
5. [Clear heartbeat events](#)

Updating multi-host, appliance-based deployments

Use the instructions in this column to update multi-host deployments of Resource Manager based on virtual appliances. You can update Control Center and Resource Manager independently. The software includes updates of the virtual machine operating system.

Preparing to update

1. [Prepare your deployment for update](#)
2. [Download the update ISO file](#)
3. [Attach the update ISO](#)
 - [Attaching an update ISO with vSphere](#)
 - [Attaching an update ISO with Hyper-V](#)
4. [Modify MariaDB passwords, if necessary](#)
5. [Update the Dashboard ZenPack](#)
6. [Stop Resource Manager](#)
7. [Remove Docker containers on the master host](#)

Updating delegate hosts

1. [Remove Docker containers](#)
2. [Update Control Center](#)

Updating the master host

1. [Update Control Center](#)
2. [Update delegate hosts with authentication tokens](#)
3. [Update Resource Manager](#)

Post-update procedures

1. [Remove the pre-upgrade snapshot](#)
2. [Start Control Center on all hosts](#)
3. [Verify zenhubworker services updated successfully](#)
4. [Start Resource Manager](#)
5. Optional: [Install a patch for Linux modeling](#)
6. [Clear heartbeat events](#)

Updating non-appliance deployments

Use the instructions in this column to update Resource Manager deployments that are not based on the Resource Manager virtual appliance. These instructions do not include updates for Control Center or the host operating system.

When a Control Center update accompanies a Resource Manager release, Control Center must be updated first. See [Updating Control Center](#).

Preparing to update

1. [Prepare your deployment for update](#)
2. [Install the base image, if necessary](#)
3. [Download Resource Manager image files](#)
4. [Import Resource Manager image files](#)
5. [Modify MariaDB passwords, if necessary](#)
6. [Update the Dashboard ZenPack](#)

Updating software

1. [Stop Resource Manager](#)
2. [Update Resource Manager](#)

Post-update procedures

1. [Remove the pre-upgrade snapshot](#)
2. [Verify zenhubworker services updated successfully](#)
3. [Start Resource Manager](#)
4. Optional: [Install a patch for Linux modeling](#)

5. [Clear heartbeat events](#)
6. Optional: [Install an application template](#)

Preparing to update an appliance-based deployment

For the complete list of update instructions, see [Updating Resource Manager](#).

Preparing an appliance-based deployment for update

Before you update your deployment, perform this procedure.

1. Verify that Resource Manager is operating normally.
 - a. Log in to the Control Center browser interface as `ccuser`.
 - b. In the Application column of the Applications table, click `Zenoss.resmgr`.
 - c. Verify that all of the services in the IP Assignments table have valid IP addresses.
 - d. Scroll down to the Services table, and then verify that none of the child services have failing health checks.
2. Check the integrity of Resource Manager databases.

For more information, see [Using Zenoss Toolbox](#).
3. Create a backup of Resource Manager.

Backups can be created with the Control Center backup feature or with the backup feature of your hypervisor. Before starting a hypervisor backup, Resource Manager, the Control Center service, and the Control Center master host virtual machine must be shut down cleanly and completely.

Downloading the update ISO file

To perform this procedure, you need:

- A workstation with internet access.
- Permission to download files from delivery.zenoss.com. Zenoss customers can request permission by filing a ticket at the [Zenoss Support](#) site.
- A secure network copy program.

Use this procedure to download the required files to a workstation and then copy the files to a location your hypervisor can use.

Perform these steps:

1. In a web browser, navigate to delivery.zenoss.com, and then log in.
2. Download the Zenoss Service Dynamics (ZSD) update file.
Replace VERSION with the entire version number of this release (for example, 6.4.1):

```
update-zenoss-zsd-VERSION-1.x86_64.iso
```

3. Use a secure copy program to copy the file to a location your hypervisor can use.

Attaching an update ISO with vSphere

To perform this task, you need:

- A VMware vSphere client
- On your workstation, a copy of the Resource Manager update ISO file

Use this procedure to attach a Resource Manager update ISO file to the CD/DVD drive of a Control Center host. Perform the procedure on each Control Center host in your deployment.

1. Use the VMware vSphere Client to log in to vCenter as root, or as a user with superuser privileges.
2. With the View menu, enable Toolbar and Show VMs in Inventory.
3. In the Inventory list, select the name of the Control Center host.
4. Enable the CD/DVD drive of the Control Center master host.
 - a. On the Summary page, click Edit Settings.
 - b. In the Hardware table of the Virtual Machine Properties dialog, select CD/DVD drive 1.
 - c. In the Device Type area, click Client Device.
 - d. At the bottom of the Virtual Machine Properties dialog, click OK.

Note: If the CD/DVD drive is configured to connect when the virtual machine starts, you may have to turn off the virtual machine, reconfigure the drive, and then turn on the virtual machine, before proceeding to the next step. Before turning off the virtual machine, stop Resource Manager. For more information, see [Stopping Resource Manager](#).

5. Attach the update ISO file to the Control Center master host.
 - a. In the toolbar, click the CD/DVD drive icon.
 - b. From the menu, select CD/DVD drive 1 > Connect to ISO image on local disk.
 - c. In the Open dialog, select the update ISO file, and then click Open.

Attaching an update ISO with Hyper-V

To perform this task, you need:

- Microsoft Remote Desktop Connection
- On the Hyper-V host, a copy of the Resource Manager update ISO file

Use this procedure to attach a Resource Manager update ISO file to the CD/DVD drive of a Control Center host. Perform the procedure on each Control Center host in your deployment.

1. Use Microsoft Remote Desktop Connection to log in to a Hyper-V host as Administrator, or as a user with Administrator privileges.
2. Start Hyper-V Manager.
3. In the Virtual Machines area of Hyper-V Manager, select the Control Center host, and then right-click to select Settings.
4. In the Hardware area of the Settings dialog, select IDE Controller 1 > DVD Drive.
5. In the Media area, click the Image file radio button, and then click Browse button.
6. In the Open dialog, select the update ISO file, and then click the Open button.
7. At the bottom of the Settings dialog, click the OK button.

Stopping Resource Manager

Use this procedure to stop Resource Manager.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Check the status of Resource Manager.

```
serviced service status --show-fields 'Name,ServiceID,Status'
```

- If the status of all services is stopped, this procedure is complete. Continue to the next procedure.
- If the status is running, perform the remaining steps.

4. Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

5. Check the status of Resource Manager.

```
serviced service status --show-fields 'Name,ServiceID,Status'
```

Repeat the command until the status of all services is stopped.

Removing Docker containers on the master host

Occasionally, stopping the Control Center service leaves Docker containers in the local registry. Perform this procedure on the Control Center master host to ensure no containers remain.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Stop the Control Center service.

```
systemctl stop serviced
```

4. Ensure that no containers remain in the local repository.
 - a. Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

- If the command returns no result, stop. This procedure is complete.
- If the command returns a result, perform the following substeps.

- b. Remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- c. Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

- If the command returns no result, stop. This procedure is complete.
- If the command returns a result, perform the remaining substeps.

- d. Disable the automatic startup of serviced.

```
systemctl disable serviced
```

- e. Reboot the host.

```
reboot
```

- f. Log in to the master host as root, or as a user with superuser privileges.
- g. Enable the automatic startup of serviced.

```
systemctl enable serviced
```

To return to the Appliance Administration menu, enter the exit command.

Removing Docker containers on delegate hosts

Occasionally, stopping the Control Center service leaves Docker containers in the local registry. Perform this procedure on each Control Center delegate host to ensure no containers remain.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Stop the Control Center service.

```
systemctl stop serviced
```

4. Ensure that no containers remain in the local repository.
 - a. Display the identifiers of all containers, running and exited.

```
docker ps -qa
```

- If the command returns no result, proceed to the next step.
- If the command returns a result, perform the following substeps.

- b. Remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, the most likely cause is an NFS conflict. Perform the following substeps.

- c. Stop the NFS and Docker services.

```
systemctl stop nfs && systemctl stop docker
```

- d. Start the NFS and Docker services.

```
systemctl start nfs && systemctl start docker
```

- e. Repeat the attempt to remove all remaining containers.

```
docker ps -qa | xargs --no-run-if-empty docker rm -fv
```

- If the remove command completes, proceed to the next step.
- If the remove command does not complete, perform the remaining substeps.

- f. Disable the automatic startup of serviced.

```
systemctl disable serviced
```

- g. Reboot the host.

```
reboot
```

- h. Log in to the delegate host as root, or as a user with superuser privileges.

- i. Enable the automatic startup of serviced.

```
systemctl enable serviced
```

5. Dismount all filesystems mounted from the Control Center master host.

This step ensures no stale mounts remain when the storage on the master host is replaced.

 - a. Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts | grep -v '/opt/serviced/var/isvcs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- b. Force the filesystems to dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts | grep -v '/opt/serviced/var/iscvs')
do
    umount -f $FS
done
```

- c. Identify filesystems mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts | grep -v '/opt/serviced/var/iscvs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the following substeps.

- d. Perform a lazy dismount.

```
for FS in $(awk '/serviced/ { print $2 }' < /proc/mounts | grep -v '/opt/serviced/var/iscvs')
do
    umount -f -l $FS
done
```

- e. Restart the NFS service.

```
systemctl restart nfs
```

- f. Determine whether any filesystems remain mounted from the master host.

```
awk '/serviced/ { print $1, $2 }' < /proc/mounts | grep -v '/opt/serviced/var/iscvs'
```

- If the preceding command returns no result, stop. This procedure is complete.
- If the preceding command returns a result, perform the remaining substeps.

- g. Disable the automatic startup of serviced.

```
systemctl disable serviced
```

- h. Reboot the host.

```
reboot
```

- i. Log in to the delegate host as root, or as a user with superuser privileges.
j. Enable the automatic startup of serviced.

```
systemctl enable serviced
```

To return to the Appliance Administration menu, enter the exit command.

Updating an appliance-based deployment

For the complete list of update instructions, see [Updating Resource Manager](#).

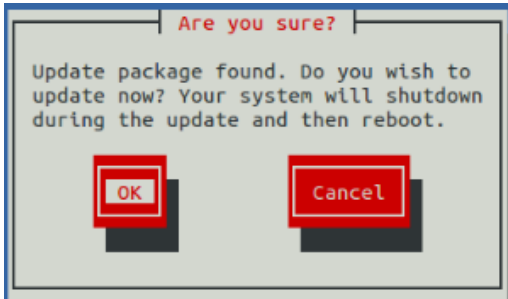
Updating Control Center on a delegate host

Before performing this procedure, attach the update ISO file to the Control Center hosts. For more information, see one of the following topics:

- [Attaching an update ISO with vSphere](#)
- [Attaching an update ISO with Hyper-V](#)

Use this option to update Control Center on a delegate host. Note: The update software creates a record of its work in `/tmp/upgrade-zsd-cc-*.log`.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Log in as the root user.
3. In the Appliance Administration menu, use the down arrow key to select Update System, and then press Enter.



- To update the system, press Enter.
 - To cancel the update, press **Tab**, and then press Enter.
4. In the Upgrade Options menu, choose Upgrade to Control Center v.r.m and then press Enter. The Appliance Administration menu is replaced by progress messages as the Control Center update begins. The following list identifies the major steps of the update, which takes about 20 minutes:
 - Install the Zenoss repository mirror. All subsequent steps use the mirror.
 - Install a new version of Control Center.
 - Install a new version of Docker.
 - Install images into the local Docker repository.
 - Install utility packages.
 - Update the host virtual machine to CentOS 7.4.

Note: During the upgrade, you might see a message about a loopback thinpool device, along with a prompt to abort or continue. Ignore the message, and at the prompt, press any key.

When the upgrade completes, the following message is displayed:

```
Complete!
The host-type host update attempt succeeded.
Control Center v.r.m is installed.
The update log is /tmp/upgrade-zsd-cc-*.log.

Follow these steps to complete the update:
1. On each delegate host, attach the update ISO, and then update Control Center.
2. On the master host, generate and distribute authentication tokens to each delegate host.
   (Control Center v.r.m requires authentication tokens for all delegate communications.)
3. On the master host, update Zenoss Service Dynamics to v.r.m.
For more information, refer to the Zenoss Service Dynamics Upgrade Guide.

Press any key to reboot...
```

5. Press any key.

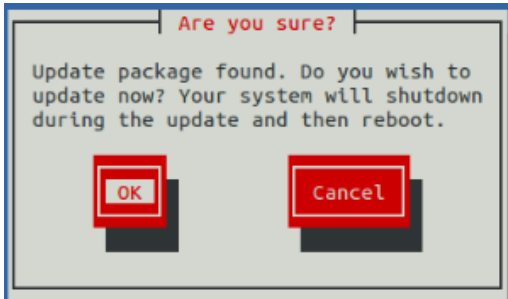
Updating Control Center on the master host

Before performing this procedure, attach the update ISO file to the Control Center hosts. For more information, see one of the following topics:

- [Attaching an update ISO with vSphere](#)
- [Attaching an update ISO with Hyper-V](#)

Use this option to update Control Center. Note: The update software creates a record of its work in /tmp/upgrade-zsd-cc-*.log.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Log in as the root user.
3. In the Appliance Administration menu, use the down arrow key to select Update System, and then press Enter.



- To update the system, press Enter.
 - To cancel the update, press **Tab**, and then press Enter.
4. In the Upgrade Options menu, choose Upgrade to Control Center v.r.m and then press Enter. The Appliance Administration menu is replaced by progress messages as the Control Center update begins. The following list identifies the major steps of the update, which takes about 20 minutes:
 - Install the Zenoss repository mirror. All subsequent steps use the mirror.
 - Install a new version of Control Center.
 - Install a new version of Docker.
 - Install images into the local Docker repository.
 - Install utility packages.
 - Update the host virtual machine to CentOS 7.4.

When the update completes, the following message is displayed:

```
Complete!
The host-type host update attempt succeeded.
Control Center v.r.m is installed.
The update log is /tmp/upgrade-zsd-cc-*.log.

Follow these steps to complete the update:
1. On each delegate host, attach the update ISO, and then update Control Center.
2. On the master host, generate and distribute authentication tokens to each delegate host.
   (Control Center v.r.m requires authentication tokens for all delegate communications.)
3. On the master host, update Zenoss Service Dynamics to v.r.m.
For more information, refer to the Zenoss Service Dynamics Upgrade Guide.

Press any key to reboot...
```

5. Press any key.

Updating delegate hosts with authentication

Starting with version 1.3.0, Control Center requires authentication tokens for all delegate communications. The tokens are based on RSA key pairs created by the master serviced instance. When you create a key pair for a delegate, serviced bundles its public key with the delegate's private key. The serviced instance on the delegate installs the credentials and uses them to sign messages with the required unique tokens.

Credentials are installed by using an SSH connection or a file.

- The command to create a key pair can initiate an SSH connection with a delegate and install credentials. This option is the most secure, because no file is created. However, it requires either public key authentication or password authentication between the master and delegate hosts.
- When no SSH connection is requested, the command to create a key pair creates a file containing the credentials. You can move the credentials file to the delegate host with any file transfer method, and then install it on the delegate.

The procedures in the following sections demonstrate how to create credentials and install them on a delegate:

- [Registering a host using SSH](#)
- [Registering a host using a file](#)

Registering a host using SSH

To succeed, the following statements about the login account used to perform this procedure must be true:

- The account exists on both the master host and on the delegate host.
- The account has serviced CLI privileges.
- The account has either public key authentication or password authentication enabled on the master host and on the delegate host.

Use this procedure to create the authentication credentials that delegate hosts require, and to register the credentials through an SSH connection. Perform the following steps on the Control Center master host, and then repeat the steps on each delegate host in your deployment.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Display the host IDs of all Control Center hosts.

```
serviced host list | cut -c-85
```

The host ID is in the first column of the output.

4. Create authentication credentials for a delegate host and register the credentials.

If the master and delegate host are configured for key-based access, the following command does not prompt you to add the delegate to the list of known hosts or to provide the password of the remote user account.

Replace Host-ID with the host ID of a delegate host:

```
serviced key reset --register Host-ID
```

5. For each delegate host in your deployment, repeat the preceding step to create and register authentication credentials.

Registering a host using a file

Use this procedure to create the authentication credentials that hosts require, and to register the credentials by using a file. Start this procedure on the Control Center master host:

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Start a command-line session as root.
 - a. In the Appliance Administration menu, select Root Shell.
 - b. Select Run, and then press Enter.

The menu is replaced by a command prompt similar to the following example:

```
[root@hostname ~]#
```

3. Display the host IDs of all Control Center hosts.

```
serviced host list | cut -c-85
```

The host ID is in the first column of the output.

4. Create and distribute credentials files for delegate hosts. Repeat the following substeps for each delegate host in your deployment.
 - a. Create authentication credentials for a delegate host.

Replace Host-ID with the host ID of a delegate host identified in the preceding step:

```
serviced key reset Host-ID
```

The command creates a unique credentials file in the local directory.

- b. Copy the credentials files to each delegate host. Use a file transfer utility such as scp to copy the files. Once copied to a delegate host, the credentials file is not needed on the master host and can be deleted.

5. Install the credentials on delegate hosts.

On each delegate host in your deployment, complete the following substeps:

- a. Log in to a delegate host as root, or as a user with superuser privileges.
- b. Use the **Down Arrow** key to select Root Shell, and then press **Enter**.

The menu is replaced by a command prompt similar to the following example:

```
[root@resmgr ~]#
```

- c. Install the credentials.

Replace Credentials-File with the pathname of the credentials file:

```
serviced host register Credentials-File
```

- d. Delete the credentials file.

The file is no longer needed on the host.

Replace Credentials-File with the pathname of the credentials file:

```
rm Credentials-File
```

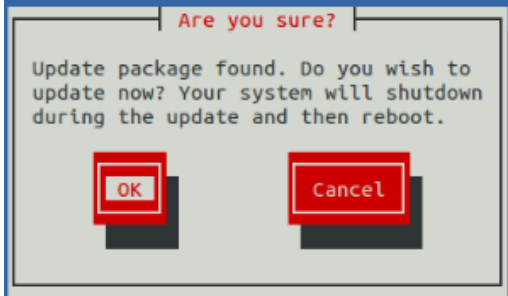
To return to the Appliance Administration menu, enter the exit command.

Updating Resource Manager on the master host

Before performing this procedure, update Control Center.

Use this procedure to update Resource Manager on a master host. Note: The update software creates a record of its work in `/tmp/upgrade-xxx-product-*.log`.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Log in as the root user.
3. In the Appliance Administration menu, use the down arrow key to select Update System, and then press Enter.



- To update the system, press Enter.
 - To cancel the update, press **Tab**, and then press Enter.
4. In the Upgrade Options menu, choose Upgrade to Zenoss Service Dynamics v.r.m and then press **Enter**. The Appliance Administration menu is replaced by progress messages as the update begins. The following list identifies the major steps of the update, which takes about 30 minutes:
 - Install or update the Zenoss mirror repository. All subsequent steps use the mirror.
 - Install images into the local Docker registry.
 - Add or update ZenPacks for Zenoss Analytics and Zenoss Service Impact.
 - Install a new application template for Resource Manager.
 - Start the Resource Manager update script.

The following message displays when the update completes:

```
Zenoss Service Dynamics upgraded successfully.  
Upgrade log stored at /tmp/upgrade-xxx-product-*.log.  
Press any key to reboot...
```

5. Press any key.

Starting Resource Manager after an update

Use this procedure to start Resource Manager and verify the update.

1. Display the login page of the Control Center browser interface.
Replace HostName with the name or IP address of the Control Center master host:

```
https://HostName
```

2. Log in as ccuser.
The Applications table shows the updated application, and the Application Templates table includes the old and new templates. Note: Do not attempt to add either the old or the new application template. The new template is already added and deployed.
3. In the Actions column of the Applications table, click the Start control of Zenoss.resmgr.
4. In the Start Service dialog, click Start Service and NN Children.

Clearing heartbeat events after updating

If you are using the Daemon Process Down portlet, zencatalogservice may be listed as down immediately after updating to this release. The status is incorrect and can be corrected by using [this procedure](#).

Updating non-appliance deployments

For the complete list of update instructions, see [Updating Resource Manager](#).

Preparing a deployment for update

Before you update your deployment, perform this procedure.

1. Verify that Resource Manager is operating normally.
 - a. Log in to the Control Center browser interface as `ccuser`.
 - b. In the Application column of the Applications table, click `Zenoss.resmgr`.
 - c. Verify that all of the services in the IP Assignments table have valid IP addresses.
 - d. Scroll down to the Services table, and then verify that none of the child services have failing health checks.
2. Check the integrity of Resource Manager databases.
For more information, see [Using Zenoss Toolbox](#).
3. Create a backup of Resource Manager with the Control Center backup feature.

Installing the base image, if necessary

The Resource Manager update script requires content from the base image of the current version. If you have flattened the layers of the base image since installation, the image may not have the expected label, and the upgrade script will not find it.

Follow these steps to determine whether the base image is present, and to download and install it, if necessary:

1. Log in to the Control Center host as `root`, or as a user with superuser privileges.
2. Search for the base image.

```
docker images | grep -E 'zenoss/resmgr'
```

- If the base image is present, the command returns a result. Continue to the next upgrade procedure.
- If the base image is absent, the command returns no result. Follow the steps in [Downloading Resource Manager image files](#) to download the base image of the current version. Then, [import the image](#).

Downloading Resource Manager image files

To perform this procedure, you need:

- A workstation with internet access
- Permission to download files from delivery.zenoss.com. Customers can request permission by filing a ticket at the [Zenoss Support](#) site.
- A secure network copy program

Use this procedure to

- download required files to a workstation
- copy the files to a Control Center master host

Follow these steps:

1. In a web browser, navigate to the download site, and then log in.
The download site is delivery.zenoss.com.
2. Download the self-installing Docker image files for Resource Manager.
Replace MAJOR.MINOR with the major and minor numbers of this version (for example, 6.4) and replace VERSION with the entire version number (6.4.1):
 - `install-zenoss-hbase-LATEST.run`
 - `install-zenoss-opentsdb-LATEST.run`
 - `install-zenoss-resmgr_MAJOR.MINOR-VERSION_1.run`
 - `install-zenoss-mariadb:10.1-VERSION.run`
3. Use a secure copy program to copy the files to the Control Center master host.

Importing Resource Manager image files

Use this procedure to import the Resource Manager image from self-installing archive files.

1. Log in to the master host as root, or as a user with superuser privileges.
2. Copy or move the archive files to /root.
3. Add execute permission to the files.

```
chmod +x /root/*.run
```

4. Change directory to /root.

```
cd /root
```

5. Import the images.

```
for image in install-zenoss-*.run
do
  /bin/echo -en "\nLoading $image..."
  yes | ./$image
done
```

6. List the images in the registry.

```
docker images
```

The result should include one image for each archive file.

7. **Optional:** Delete the archive files, if desired.

```
rm -i ./install-*.run
```

8. Copy the update scripts from the new Resource Manager image to a directory in /root.
Replace MAJOR.MINOR with the major and minor numbers of this version (for example, 6.4) and replace VERSION with the entire version number (6.4.1):

```
docker run -it --rm -v /root:/mnt/root zenoss/resmgr_MAJOR.MINOR:VERSION_1 rsync -a /root/MAJOR.MINOR.x /mnt/root
```

For example, for release 6.4.1, the command is:

```
docker run -it --rm -v /root:/mnt/root zenoss/resmgr_6.4:6.4.1_1 rsync -a /root/6.4.x /mnt/root
```

Stopping Resource Manager on non-appliance deployments

Use this procedure to stop Resource Manager.

1. Log in to the Control Center master host as a user with serviced CLI privileges.
2. Check the status of Resource Manager.

```
serviced service status Zenoss.resmgr --show-fields 'Name,ServiceID,Status'
```

- If the status of all services is stopped, this procedure is complete. Continue to the next procedure.
- If the status is running, perform the remaining steps.

3. Stop Resource Manager.

```
serviced service stop Zenoss.resmgr
```

4. Check the status of Resource Manager.

```
serviced service status Zenoss.resmgr --show-fields 'Name,ServiceID,Status'
```

Repeat the command until the status of all services is stopped.

Updating Resource Manager on non-appliance deployments

Use this procedure to update Resource Manager.

1. Log in to the Control Center master host as root, or as a user with superuser privileges.
2. Start the update script.
Replace MAJOR.MINOR with the major and minor numbers of this version (for example, 6.4):

```
/root/MAJOR.MINOR.x/upgrade-resmgr.sh
```

The update process begins. If you encounter errors, see [Common update error recovery procedures](#).

The following message appears frequently and can be ignored:

```
sync.go:85: ERR Failed to publish events caused by: EOF
```

3. Restart Resource Manager.
Some Resource Manager services are started during the update, and they must be restarted.

```
serviced service stop Zenoss.resmgr && serviced service start Zenoss.resmgr
```

Removing the pre-upgrade snapshot

The Resource Manager update script uses Control Center to create and tag a snapshot of the system before it begins the update process. Tagged snapshots persist until they are explicitly removed, and grow over time. When you are satisfied the new release is working properly, remove the pre-update snapshot.

1. Log in to the Control Center master host as a user with serviced CLI privileges.
2. Display a list of all Control Center snapshots, with their tags.

```
serviced snapshot list -t
```

Example result, truncated to save space:

Snapshot	Description	Tags
xm5mtezbyo2_20160211-220535.480		preupgrade-resmgr-5.2.0

The snapshot identifier is shown in the first column.

3. Remove the pre-update snapshot.
Replace Snapshot-ID with the identifier of the pre-update snapshot returned in the previous step:

```
serviced snapshot remove Snapshot-ID
```

Using Zenoss Toolbox

This section provides an introduction to the Zenoss Toolbox, which is included in Resource Manager. For up-to-date information, refer to the Zenoss Toolbox [KnowledgeBase article](#).

- [Zenoss Toolbox tools](#)
- [Running Zenoss Toolbox tools](#)

Zenoss Toolbox tools

The Zenoss Toolbox tools examine key Resource Manager components for common issues affecting data integrity. Zenoss recommends running the following tools, in order, before updating Resource Manager:

1. Always run `zodbscan` first.

The script is purely an analysis tool which checks the object references in the persistent lists that ZODB uses.

2. If, and only if, `zodbscan` tells you to, run `findposkeyerror` with the fix flag (`findposkeyerror -f`).

There's no reason to run `findposkeyerror` **without** the fix flag—doing so just repeats the analysis that `zodbscan` performs.

When `findposkeyerror` completes, run `zodbscan` again, to verify that no dangling references remain.

3. Run `zenrelationscan` after the previous step or steps.

The `zenrelationscan` script performs incremental checks that `zodbscan` does not. Rerun it with with the fix flag (`zenrelationscan -f`) if it tells you to.

If the script fails when you run it with the fix flag, contact Zenoss Support.

4. If, and only if, all runs of the previous scripts exit cleanly, run `zencatalogscan`.

This script determines whether catalog entries resolve to objects correctly (`getObject` calls) and fixes stale/incorrect entries. Rerun it with the fix flag (`zencatalogscan -f`) if it tells you to.

The script does not fix issues that arise when objects exist but are missing from the catalog. The fix for such issues requires specific `zendmd` actions or rebuilding the entire catalog. For assistance, contact Zenoss Support.

The tools are run inside a Zope container, and the log files for each command are found in `$ZENHOME/log/toolbox`.

Running Zenoss Toolbox tools

1. Log in to the Control Center master host as a user with serviced CLI privileges.
2. Start an interactive session in a Zope container as the zenoss user.

```
serviced service attach zope/0 su - zenoss
```

3. Run the [Zenoss Toolbox tools](#), in order.
4. Exit the Zope container.

```
exit
```

Common update error recovery procedures

This section describes common error messages during updates, and provides procedures for recovering and continuing.

- [A snapshot with the given tag already exists](#)

A snapshot with the given tag already exists

When an update attempt fails, the update script does not remove the snapshot it creates at the beginning of the update process. Use this procedure to remove the tag of the pre-update snapshot and restart the update. Untagged snapshots are removed when their time-to-live (TTL) expires. The TTL value is defined by the `SERVICED_SNAPSHOT_TTL` variable in the Control Center configuration file.

1. Log in to the Control Center master host as a user with serviced CLI privileges.
2. Create a variable for the identifier of the tenant application.

```
myTenant=$(serviced service list Zenoss.resmgr --format='{{.ID}}')
```

3. Display a list of all Control Center snapshots, with their tags.

```
serviced snapshot list -t
```

Example result, truncated to save space:

Snapshot	Description	Tags
xm5mtezbyo2_20160211-220535.480		preupgrade-resmgr-5.2.0

The snapshot identifier is shown in the first column.

4. Remove the tag of the pre-update snapshot.
Replace Tag-Name with the name of the pre-update snapshot that was displayed in the previous step:

```
serviced snapshot untag ${myTenant} Tag-Name
```

5. Restart the update script.
Replace MAJOR.MINOR with the major and minor numbers of this version (for example, 6.4):

```
/root/MAJOR.MINOR.x/upgrade-resmgr.sh
```

Installing an application template

This section includes procedures for downloading and installing the most recent Resource Manager application template. The latest template is not needed to perform an update, and is included in appliance updates. Use the procedures in this section if you plan to delete a deployed application template and then redeploy with a newer version of the template. (For example, in a development or staging environment.)

- [Downloading the template package](#)
- [Installing the application template](#)

Downloading the template package

To perform this procedure, you need:

- A workstation with internet access.
- Permission to download files from the delivery.zenoss.com site. Zenoss customers may request permission by filing a ticket at the [Zenoss Support](#) site.
- A secure network copy program.

Use this procedure to download the required files to a workstation and then copy the files to the Control Center master host.

Perform these steps:

1. In a web browser, navigate to the delivery.zenoss.com site.
2. Log in with the account provided by Zenoss Support.
3. Download the template package file.

```
zenoss-resmgr-service-VERSION-1.noarch.rpm
```

4. Use a secure copy program to copy the package file to the Control Center master host.

Installing the application template

Use this procedure to install the Resource Manager template on a Control Center master host.

1. Log in to the target host as root, or as a user with superuser privileges.
2. Change directory to the directory in which the template package file is located.
3. Install the template.
Replace VERSION with the entire version number of this release (for example, 6.4.1):

```
yum install ./zenoss-resmgr-service-VERSION-1.noarch.rpm
```

The template file is stored in /opt/serviced/templates.

ZenPack considerations

This section describes special considerations for updating ZenPacks. For more information about ZenPacks, see the [ZenPack catalog](#).

- [Alternate naming convention for LUN- and VM-specific metrics](#)

Alternate naming convention for LUN- and VM-specific metrics

This information applies to the VMware vSphere ZenPack. By default, LUN- and VM-specific metrics are written under the following naming convention:

```
<device-id>/<metric-name>
```

To improve retrieval speed of metrics (graphing speed) for vSphere instances with thousands of LUNs or virtual machines, you can enable the following alternate naming convention:

```
<device-id>/<component-id>/<metric-name>
```

Note: Changing the naming convention causes historical metrics to become inaccessible. The following configuration properties support the alternate convention:

- `zVSphereLUNContextMetric` - Controls whether to use LUN-specific metric names when storing performance data.

For example, with this configuration property enabled, the metric name `diskReadRequests_diskReadRequests` changes as follows:

Old name:

```
sol-vcenter/diskReadRequests_diskReadRequests
```

New name:

```
sol-vcenter/HostSystem_host-134_naa.600508e001da816aa59a72903/diskReadRequests_diskReadRequests
```

- `zVSphereVMContextMetric` - Controls whether to use VM-specific metric names when storing performance data.

For example, with this configuration property enabled, the metric name `cpuUsageAvg_cpuUsageAvg` changes as follows:

Old name:

```
vxchnge-vcenter/cpuUsageAvg_cpuUsageAvg
```

New name:

```
vxchnge-vcenter-02/VirtualMachine_vm-2556/cpuUsageAvg_cpuUsageAvg
```

For more information, see the **VMware vSphere ZenPack** in the [ZenPack catalog](#).

Enabling the alternate metric naming convention

This procedure requires a restart the `zenvsphere` daemon.

1. Log in to the Resource Manager browser interface.
2. Navigate to the Infrastructure page.
3. In the left pane, choose vSphere.
4. In the main pane, click the name of the VMware vSphere device.
5. In the left pane, choose Configuration Properties.
6. In the main pane, Name column, double-click the vSphere context metric name.
7. In the Edit Config Property dialog box, check Store x metrics by context, and then click Submit. The property value changes to true.
8. Restart the `zenvsphere` daemon.

Release notes

The following table shows the recent Resource Manager releases, along with their accompanying Control Center releases.

Release Date	Resource Manager	Control Center
August 2019	6.4.1	1.6.5
Withdrawn	6.4.0	1.6.5
Jan 2019	6.3.2	1.6.3
Aug 2018	6.2.1	1.5.1
Jun 2018	6.2.0	1.5.1
Mar 2018	6.1.2	1.5.0
Feb 2018	6.1.1	1.5.0
Dec 2017	6.1.0	1.5.0
Nov 2017	6.0.1	1.5.0
Nov 2017	6.0.0	1.5.0

Resource Manager 6.4.1

This section contains important information about release 6.4.1 of Zenoss Resource Manager (Resource Manager). Release 6.4.0 was withdrawn. Zenoss encourages all customers to update their deployments to release 6.4.1.

- [New features](#)
- [Considerations and workarounds](#)
 - [Compatibility with Zenoss Service Impact](#)
 - [Regular expression matching limitation](#)
 - [NFS client 4.1 is not supported](#)
 - [CentralQuery maximum memory allocation pool](#)
 - [LDAP authentication](#)
 - [High-availability documentation](#)
 - [Upgrade considerations](#)
 - [Installation considerations](#)
- [Fixed issues](#)
- [Known issues](#)
- [Installed ZenPacks](#)
 - [Packaged ZenPacks](#)
 - [Obsolete ZenPacks](#)

New features

This release includes the following modeling pipeline improvements:


- The Zope Object Database (ZODB) is backed by an upgraded version of MariaDB.
- The method for processing large models (for example, the models produced by supervisors) has been refactored.
- The efficiency of model relationship building has been improved.
- The performance of Solr has been increased for model pipeline tasks.
- In nearly two dozen ZenPacks, the edge providers for Service Impact have been consolidated into the DynamicView ZenPack, boosting performance and reducing inconsistencies.

The zenhub service includes the following enhancements:

- The zenhubworker service has been refactored to include a service dedicated to processing large models:
 - The `zenhubworker` service processes normal models.
 - The `zenhubadmworker` service processes large models.
- Invalidation workers have been refactored as the `zenhubiworker` service, to enable scaling without restarting the zenhub service.
- Log messages now include [searchable tags](#) (not indexed by logstash) that facilitate debugging.

Also, the following, general improvements are included in this release:

- New self-monitoring functions that draw on our experience operating Zenoss Cloud have been incorporated.
- Device details pages include a new option, for performing debugging from the browser interface.



- From the Infrastructure tab, you can select a device class or organizer and model all of the selected devices.
- All Docker containers are upgraded to use CentOS 7.6, and the download site now includes repository mirrors for RHEL/CentOS 7.5 and 7.6.

Considerations and workarounds

Compatibility with Zenoss Service Impact

This version of Resource Manager is compatible with Zenoss Service Impact version 5.5.1 or later. All users are strongly encouraged to upgrade Service Impact to the latest version.

Regular expression matching limitation

(ZEN-29376) Resource Manager uses the Apache Solr search platform, which tokenizes fields. Currently, searching a tokenized field with a regex matches only a single token, limiting complex regex searches.

For example, on the INFRASTRUCTURE page, IP Address is a tokenized field. Regex search `^192.168` will not work to filter only IP addresses that start with 192.168. The Device column is untokenized; regex search on the device name works as expected.

NFS client 4.1 is not supported

A file locking defect might exist in NFS 4.1 with RHEL/CentOS 7.4, which could cause zeneventserver to crash and other DFS-related problems. For more information, see [Configuring NFS 4.0](#) or knowledge base article [Potential Issues Running With RHEL 7.4 Or CentOS 7.4](#).

CentralQuery maximum memory allocation pool

(ZEN-15907) The size of the CentralQuery maximum memory allocation pool is set by the RAMCommitment variable in the CentralQuery service definition. The default value is 1024MB. Do not change the value to anything less than 1024MB.

LDAP authentication

(ZEN-31194) This release includes [LDAPAuthenticator](#) v3.3.3, which no longer includes an option to skip SSL certificate verification during installation or upgrade. If you are using a self-signed certificate, upload it before upgrading. For more information, see [Adding an SSL certificate](#).

High-availability documentation

Customers who have deployed Control Center and Resource Manager in a high-availability configuration are still supported in this release. For upgrade instructions, refer to the [Control Center Upgrade Guide for High-availability Deployments](#).

Upgrade considerations

- (ZEN-29807) Upgrading from Resource Manager 5.x to 6.x includes a required step to reindex all cataloged objects into the new Solr index. To speed this process, before upgrading, you can increase the CPU core count to use more workers. For more information, see knowledge base article [Reducing 5.X To 6.X Upgrade Times By Increasing Workers For Indexing](#).
- Beginning with release 6.2.1, support for upgrading versions earlier than 5.3.3 was removed from the upgrade instructions. For assistance with planning for upgrading versions prior to 5.3.3, please contact Zenoss Support. In short, if you are on a version prior to that, we'll recommend you should plan to upgrade to version 5.3.3 and then to the latest 6.x version from there.
- The upgrade process might reinstall an older version of the Catalog Service ZenPack. To avoid this issue, manually remove older versions of the Catalog Service ZenPack .egg file from the ZenPack backups directory.
- (ZEN-28375) Beginning with Resource Manager 6.1.0, for counter/derive (rate) data points, the calculated rate value is stored instead of raw counters. The rate is calculated at the collector daemon as it collects data. If a graph (or API) request for a derive of counter data point spans the upgrade data, the system automatically queries the data correctly. However, slight anomalies might occur in the data that is captured immediately before and after an upgrade from Resource Manager 6.0.1 or earlier.
- (ZEN-29100) No data is collected for the ZenossRM device after an upgrade if the localhost collector is moved off the master host. If you have moved the localhost collector off the Control Center master host, to enable data collection, set properties for the ZenossRM device as follows:
 1. Navigate to the ZenossRM device overview page and select Configuration Properties.
 2. Set the following zProperties:
 - zRMMonCCHost - Enter the IP address of the Control Center master host.
 - zRMMonCCUser - Enter the name of the ccuser account, which is the default account for gaining access to the Control Center browser interface.
 - zRMMonCCPassword - Enter the password of the ccuser account.
- The following issues require special procedures, all of which are integrated into [the update process](#):
 - (ZEN-32191) The version of the [Dashboard ZenPack](#) that is packaged with Resource Manager 6.4.1 makes the URLs of dashboard portlets invalid. To avoid this issue, [follow this procedure](#).
 - (ZEN-32036) Update fails when custom passwords for MariaDB are in use. To prevent a failure, [follow this procedure](#).
 - (ZEN-32092) Update occasionally fails to migrate the zenhubworker service definition. [Follow this procedure](#) prior to starting Resource Manager.
 - (ZEN-32175) In some environments, Linux modeling runs can fail with tracebacks. [Follow this procedure](#) after updating to 6.4.1.

Installation considerations

(ZEN-32175) In some environments, Linux modeling runs can fail with tracebacks. [Follow this procedure](#) after installing and starting 6.4.1.

Fixed issues

ID	Description
ZEN-28773	Graphs can load very slowly or not at all when viewing 30 days (or more) worth of data
ZEN-29093	Duplicate names for metrics in some ZenPacks
ZEN-29354	Device names can be discovered without authentication
ZEN-31004	Event console filters and column changes do not persist
ZEN-31030	Event console hangs when accessing directly

ZEN-31081	Event class mappings cannot be resequenced if their names contain spaces
ZEN-31164	Critical events on /Status show devices as down unless the event is on /Status/Ping
ZEN-31166	While processing large models, certain queries create locks, causing Zope/Zauth to be unresponsive for long periods
ZEN-31235	In the event transform display, the scroll bar is missing for large transforms
ZEN-31299, ZEN-31704	Link to this graph is not working
ZEN-31368	Maintenance task can become stuck after a disconnection from zenhub
ZEN-31467	Add a test to the self-monitoring ZenPack for browser interface responsiveness
ZEN-31476	The events console does not refresh automatically, requires a manual refresh
ZEN-31506	Upgrade Docker to address CVE-2019-5736
ZEN-31513	The zenhub service log includes a spurious error message
ZEN-31540	The Zope health check fails to restart Zope when a deadlock occurs
ZEN-31574	PythonCollector failing to build configurations
ZEN-31605	Unable to run zensendevent commands when attached to a zenactiond container
ZEN-31641	Hyperlinks in event messages do not work as expected
ZEN-31696	Processing large models stalls after an error
ZEN-31700	Add a metric to the self-monitoring ZenPack for the large model processing rate
ZEN-31705, ZEN-31709	zProperty fields store and display passwords in clear text
ZEN-31706	HTML portlets do not prevent XSS attacks
ZEN-31707	Test groups of devices for monitoring/collection/auth
ZEN-31710	The Google Maps portlet does not handle invalid addresses gracefully
ZEN-31739	Modeling is blocked when requests for device status from the event processor do not return normally
ZEN-32098	The zenping service detects devices as down but does not generate events or update their status
ZEN-32147, ZEN-31933	vSphere modeling fails during incremental modeling
ZEN-32128, ZEN-32077	The zenhub stats command does not return statistics as expected

Known issues

The following table describes the known issues in this release.

ID	Description
ZEN-30973	The zenrelationscan tool does not detect certain KeyErrors, which disables monitoring of particular devices
ZEN-32036	Update fails when custom passwords for MariaDB are in use
ZEN-32092	Update occasionally fails to migrate the zenhubworker service definition
ZEN-32175	In some environments, Linux modeling runs can fail with tracebacks
ZEN-32197	The Site Window portlet does not show graphs.

Installed ZenPacks

This section lists ZenPacks that are automatically installed, those that are packaged but not installed, and those that are obsolete and should be uninstalled. For more information about ZenPacks, see the [ZenPack Catalog](#).

This release of Resource Manager installs the following ZenPacks at the current version listed in the table.

ZenPack	Current	Previous
---------	---------	----------

	version	version
AdvancedSearch	2.0.1	Same
AixMonitor	2.2.3	Same
ApacheMonitor	2.1.4	Same
AuditLog	1.4.1	Same
AWS	4.1.1	4.0.2
CalculatedPerformance	2.5.2	2.5.1
CiscoMonitor	5.10.0	5.9.0
CiscoUCS	2.8.1	2.8.0
ComponentGroups	1.8.0	1.7.0
ControlCenter	1.6.2	Same
Dashboard (special circumstances)	1.3.8	1.3.3
Dell.PowerEdge	3.0.0	2.0.4
Diagram	1.3.2	Same
DistributedCollector	3.1.9	3.1.7
DnsMonitor	3.0.1	Same
Docker	2.0.4	Same
DurationThreshold	2.0.5	Same
DynamicView	1.7.1	1.6.2
EMC.base	2.1.0	Same
EnterpriseCollector	1.9.0	1.8.3
EnterpriseReports	2.5.0	Same
EnterpriseSecurity	1.2.0	Same
EnterpriseSkin	3.3.5	Same
HP.Proliant	3.3.2	Same
HttpMonitor	3.1.0	3.0.4
IBM.Power	1.1.3	1.1.2
InstalledTemplatesReport	1.1.2	1.1.1
JuniperMonitor	2.1.1	Same
LDAPAuthenticator	3.3.3	Same
LDAPMonitor	1.4.2	Same
Licensing	0.3.0	Same
LinuxMonitor	2.3.3	2.3.2
Microsoft.Azure	2.0.0	1.3.3
Microsoft.Windows	2.9.4	2.9.2
MySqlMonitor	3.1.0	Same
NetAppMonitor	4.0.0	3.6.0
NtpMonitor	3.0.0	Same
PredictiveThreshold	1.3.0	Same
PropertyMonitor	1.1.1	Same
PythonCollector	1.10.1	Same

RMMonitor	1.1.2	Same
SolarisMonitor	2.5.2	2.5.1
StorageBase	1.4.3	Same
SupportBundle	1.1.2	Same
vSphere	4.0.2	4.0.0
WBEM	2.1.0	Same
WSMAN	1.0.4	1.0.3
ZenDeviceACL	2.3.0	Same
ZenJMX	3.12.1	Same
ZenMail	5.1.0	Same
ZenOperatorRole	2.2.0	Same
ZenPackLib	2.1.1	Same
ZenSQLTx	2.7.1	Same
ZenWebTx	3.0.3	Same

Packaged ZenPacks

The following ZenPacks are packaged with Resource Manager, but not automatically installed:

- [BigIpMonitor](#)
- [Ceph](#)
- [CiscoAPIC](#)
- [DB2](#)
- [DatabaseMonitor](#)
- [FtpMonitor](#)
- [GoogleCloudPlatform](#)
- [HpxMonitor](#)
- [JBossMonitor](#)
- [Kubernetes](#)
- [Layer2](#)
- [Memcached](#)
- [Microsoft.Exchange](#)
- [Microsoft.Lync](#)
- [Microsoft.MSMQ](#)
- [NSX](#)
- [NetScaler](#)
- [NetScreenMonitor](#)
- [OpenStack](#)
- [OpenStackInfrastructure](#)
- [OpenvSwitch](#)
- [PostgreSQL](#)
- [RabbitMQ](#)
- [TomcatMonitor](#)
- [XenServer](#)

Obsolete ZenPacks

[ZenMailTX](#) is obsolete. If this ZenPack is installed at your site, uninstall it.

Updating the Dashboard ZenPack

The version of the [Dashboard ZenPack](#) that is packaged with Resource Manager 6.4.1 makes the URLs of dashboard portlets invalid (ZEN-32191). To avoid this issue, upgrade the Dashboard ZenPack to version 1.3.9 before upgrading your deployments to 6.4.1.

To perform this procedure, you need the updated Dashboard ZenPack, which is available on delivery.zenoss.com. Customers can request permission to download files by filing a ticket at the [Zenoss Support](#) site.

Follow these steps:

1. Log in to your workstation and start a web browser.
2. Download the ZenPack to install or update from delivery.zenoss.com.
The ZenPack egg file is in the folder that contains Resource Manager 6.4.1 files.
3. Copy the ZenPack egg file from your workstation to a directory on the Control Center master host.
4. Log in to the Control Center master host as a user with `serviced` CLI privileges.
5. Place the file in a temporary directory on the master host.
 - a. Create a directory for the ZenPack egg file.
The directory must be local (not mounted). The following command creates a directory in `/tmp`:

```
test -d /tmp/zenpack || mkdir /tmp/zenpack
```

- b. Move or copy the ZenPack file to the new directory.
- c. Set global read-write-execute permissions on the directory and its contents:

```
chmod -R 0777 /tmp/zenpack
```

6. Create a snapshot:

```
serviced service snapshot Zenoss.resmgr
```

On completion, the `serviced` command returns the ID of the new snapshot. If the installation of a ZenPack fails, you can restore the snapshot created in this step. For more information about restoring a snapshot, see [Creating snapshots and rolling back](#).

7. Change directory to the directory in which the ZenPack egg file is located.

```
cd /tmp/zenpack
```

8. Install the ZenPack:

```
serviced service run zope zenpack-manager install ZenPacks.zenoss.Dashboard-1.3.9-py2.7.egg
```

9. Restart all Zenoss services:

```
serviced service restart Zenoss.resmgr/Zenoss
```

Patching a Linux modeling issue

In a small number of environments, Zenoss Support has observed that Linux modeling runs can fail with tracebacks that end in `KeyError: 'rebuild'` (ZEN-32175). Installing this patch is optional, because you may not experience the failure. You can install it at any time.

Perform this procedure after installing (or updating to) Resource Manager 6.4.1.

This patch is only for Resource Manager 6.4.1. Do not install it in prior or subsequent releases.

- a. Download the `ZEN-32175.patch` file from the Resource Manager 6.4.1 folder on delivery.zenoss.com. Customers can request permission to gain access to delivery.zenoss.com by filing a ticket at the [Zenoss Support](#) site.
- b. Copy the file to a directory on the Control Center master host.
- c. Log in to the Control Center master host as a user with `serviced` CLI privileges.
- d. Create a directory for the patch file, if one does not already exist.

```
test -d /tmp/rmpatch || mkdir /tmp/rmpatch
```

- e. Move the `ZEN-32175.patch` file into the new directory.
- f. Set global read-write-execute permissions on the directory and its contents.

```
chmod -R 0777 /tmp/rmpatch
```

- g. Change directory to the `rmpatch` directory.

```
cd /tmp/rmpatch
```

- h. Start an interactive shell in a new Zope container as user `zenoss`. The following invocation also saves the Zope container with the name `ZEN-32175`:

```
serviced service shell -i -s ZEN-32175 zope su - zenoss
```

- i. Change directory to `/opt/zenoss`.

```
cd /opt/zenoss
```

- j. Copy the patch file to `/opt/zenoss`.

```
cp /mnt/pwd/ZEN-32175.patch .
```

- k. Patch the system.

```
patch -p0 < ZEN-32175.patch
```

On success, the patch utility displays the following output:

```
patching file Products/DataCollector/ApplyDataMap/reporter.py
patching file Products/DataCollector/ApplyDataMap/tests/test_reporter.py
patching file Products/ZenEvents/ZenEventClasses.py
```

If you do not see the preceding output, do not complete this procedure. Contact Zenoss Support.

1. Exit the Zope container.

```
exit
```

2. Commit the Zope container.


```
serviced snapshot commit ZEN-32175
```

3. Restart Resource Manager.

```
serviced service restart Zenoss.resmgr
```

Removing customized MariaDB passwords for system updates

During updates to release 6.4.1, customers who are using customized MariaDB passwords experience failures (ZEN-32036). Use this procedure to restore the default passwords, perform the update, and then customize the passwords again.

1. Log in to the Control Center master host as a user with `serviced` CLI privileges.
2. Stop Resource Manager services.

- a. Check the status of Zenoss services.

```
serviced service status Zenoss.resmgr/Zenoss --show-fields 'Name,ServiceID,Status'
```

- If the status of Zenoss services is stopped, continue to the next step.
- If the status is started, perform the next substep.

- b. Stop Zenoss services.

```
serviced service stop Zenoss.resmgr/Zenoss
```

3. Start Infrastructure services.

- a. Check the status of Infrastructure services.

```
serviced service status Zenoss.resmgr/Infrastructure --show-fields 'Name,ServiceID,Status'
```

- If the status of Infrastructure services is started, continue to the next step.
- If the status is stopped, perform the next substep.

- b. Start Infrastructure services.

```
serviced service start Zenoss.resmgr/Infrastructure
```

4. Store the customized passwords in a text file.

```
serviced service list Zenoss.resmgr | grep -E 'global\.conf\.z[eo].*password' > origDBpasswords.txt
```

5. Follow [the procedure for changing MariaDB passwords](#), replacing the customized passwords with the defaults.

- You need to use the customized password to log in:

```
mysql -uroot -pPASSWORD  
Example: mysql -uroot -p26D72357829DF94B817311
```

- Use the following commands to set the passwords for both mariadb-model and mariadb-events:

```
SET PASSWORD FOR 'root'@'127.0.0.1' = PASSWORD('');  
SET PASSWORD FOR 'root'@'localhost' = PASSWORD('');  
SET PASSWORD FOR 'zenoss'@'127.0.0.1' = PASSWORD('zenoss');  
SET PASSWORD FOR 'zenoss'@'%' = PASSWORD('zenoss');
```

- Likewise, the values in `global.conf` need to be reset:

```
global.conf.zep-admin-password  
global.conf.zep-password zenoss  
global.conf.zodb-admin-password  
global.conf.zodb-password zenoss
```

- If you are unable to edit `global.conf` through the Control Center browser interface, use the following command to edit the JSON file:

```
serviced service edit Zenoss.resmgr
```

6. Perform [the update to release 6.4.1](#).
7. Follow [the procedure for changing MariaDB passwords](#), restoring your customized passwords.
8. Stop and restart Resource Manager.

```
serviced serviced stop Zenoss.resmgr && serviced service start Zenoss.resmgr
```

Verifying zenhubworker services updated successfully

In some cases, the update to Resource Manager 6.4.1 fails to update the `zenhubworker` service completely. If your update succeeded, your deployment includes a `zenhubworker (adm)` service and a `zenhubworker (default)` service.

<code>zenhub</code>	Started
<code>zenhubiworker</code> ●	Started
<code>zenhubworker (adm)</code> ●	Started
<code>zenhubworker (default)</code>	Started

If your deployment does not include these new services, follow this procedure before using Resource Manager 6.4.1:

1. Log in to the Control Center master host as a user with `serviced` CLI privileges.
2. Start an interactive session in a Zope container as user `zenoss`.

```
serviced service attach zope/0 su - zenoss
```

3. Apply the `zenhubworker` service update migration.

```
zenmigrate --step=AddADMZenHubWorkerService
```

4. Restart the `Zenoss.resmgr` service.

```
serviced service restart Zenoss.resmgr
```

Resource Manager 6.4.0

This release was withdrawn. [Resource Manager 6.4.1](#) is available and should be used instead.

Resource Manager 6.3.2

This section contains important information about release 6.3.2 of Zenoss Resource Manager (Resource Manager).

- [New features](#)
- [Considerations and workarounds](#)
- [Fixed issues](#)
- [Known issues](#)
- [Installed ZenPacks](#)
- [Zenoss JSON API changes](#)

New features

Zenhub scaling enhancements

(ZEN-30832, originally identified as defect ZEN-29812). Previously the zenhub service delegated work to subprocesses, which limited its scale to the resources available on a single host. Now, the zenhub service can delegate (non-invalidation related) work to a separate service, zenhubworker, and you can run as many zenhubworker services as required, on any host in the resource pool. The upgrade process will automatically reconfigure zenhub appropriately for horizontally scaling. However, if you have previously had a "superhub" configuration (characterized by having single zenhub instance in a dedicated Control Center resource pool) put in place by Zenoss Services or Support to help mitigate scaling issues this enhancement now addresses, you may also wish to adjust its resource pool configuration manually, after the upgrade. For more information about updating "superhub" configurations, see [Migrating zenhub service pool configuration](#).


Zope deadlock improvements

This release, along with the Control Center 1.6.3 release, provide the following enhancements to Zope-based services:

- (CC-3559) Automatic restart when a service fails 3 health checks in a row. This feature ensures deadlocked services are restarted after 3 failed health checks (2.5 minutes). The `zauth`, `zenreports`, `zope`, and `zenapi` services benefit from this update.
- (ZEN-30871) The default number of `zauth` service instances is now two, so that services which depend on authentication are not blocked when `zauth` is deadlocked.
- (ZEN-29781) A supporting library was upgraded to avoid a thread lock issue in a previous version.

Debug option for device modeling

(ZEN-30344) The browser interface now includes an option to model a device with debug messages enabled.



Configurable query limit

(ZEN-29877) The Solr index service now returns a default limit of 10,000 results. This can be changed by editing `global.conf`, if desired. The variable name is `global.conf.solr-search-limit`.

Default "HostPolicy" for several RM services has been changed to PREFER_SEPARATE

(ZEN-22178) The default host policy for zenhub, zeneventd, and zope services has been change from the default of "none" to PREFER_SEPERATE. The default policy for the new "zenhubworker" service added by ZEN-30832, has also been set to PREFER_SEPARATE. This HostPolicy setting has no effect for zenhub since it instance maximum was set by ZEN-30832 to one, but has been set this way explicitly for future proofing in case that maximum is raised in the future. This change is also automatically applied by upgrades if the HostPolicy is not already set to PREFER_SEPARATE.

Considerations and workarounds

Compatibility with Zenoss Service Impact

This version of Resource Manager is compatible with Zenoss Service Impact version 5.2.3 or later.

Regular expression matching limitation

(ZEN-29376) Resource Manager uses the Apache Solr search platform, which tokenizes fields. Currently, searching a tokenized field with a regex matches only a single token, limiting complex regex searches.

For example, on the INFRASTRUCTURE page, IP Address is a tokenized field. Regex search `^192.168` will not work to filter only IP addresses that start with 192.168. The Device column is untokenized; regex search on the device name works as expected.

Load time for component graphs

(ZEN-29300) Load time might exceed 10 seconds for more than 200 component graphs when you have activated the All on same graph check box.

NFS client 4.1 is not supported

A file locking defect might exist in NFS 4.1 with RHEL/CentOS 7.4, which could cause zeneventserver to crash and other DFS-related problems. For more information, see [Configuring NFS 4.0](#) or knowledge base article [Potential Issues Running With RHEL 7.4 Or CentOS 7.4](#).

CentralQuery maximum memory allocation pool

(ZEN-15907) The size of the CentralQuery maximum memory allocation pool is set by the RAMCommitment variable in the CentralQuery service definition. The default value is 1024MB. Do not change the value to anything less than 1024MB.

High-availability documentation

Customers who have deployed Control Center and Resource Manager in a high-availability configuration are still supported in this release. For upgrade instructions, refer to the [Control Center Upgrade Guide for High-availability Deployments](#).

Upgrade considerations

- (ZEN-29807) Upgrading from Resource Manager 5.x to 6.x includes a required step to reindex all cataloged objects into the new Solr index. To speed this process, before upgrading, you can increase the CPU core count to use more workers. For more information, see knowledge base article [Reducing 5.X To 6.X Upgrade Times By Increasing Workers For Indexing](#).
- Beginning with release 6.2.1, support for upgrading versions earlier than 5.3.3 was removed from the upgrade instructions. For assistance with planning for upgrading versions prior to 5.3.3, please contact Zenoss Support. In short, if you are on a version prior to that, we'll recommend you should plan to upgrade to version 5.3.3 and then to the latest 6.x version from there.
- The upgrade process might reinstall an older version of the Catalog Service ZenPack. To avoid this issue, manually remove older versions of the Catalog Service ZenPack .egg file from the ZenPack backups directory.
- (ZEN-28375) Beginning with Resource Manager 6.1.0, for counter/derive (rate) data points, the calculated rate value is stored instead of raw counters. The rate is calculated at the collector daemon as it collects data. If a graph (or API) request for a derive of counter data point spans the upgrade data, the system automatically queries the data correctly. However, slight anomalies might occur in the data that is captured immediately before and after an upgrade from Resource Manager 6.0.1 or earlier.
- (ZEN-29100) No data is collected for the ZenossRM device after an upgrade if the localhost collector is moved off the master host. If you have moved the localhost collector off the Control Center master host, to enable data collection, set properties for the ZenossRM device as follows:
 1. Navigate to the ZenossRM device overview page and select Configuration Properties.
 2. Set the following zProperties:

zRMMonCCHost - Enter the IP address of the Control Center master host.

zRMMonCCUser - Enter the name of the ccuser account, which is the default account for gaining access to the Control Center browser interface.

zRMMonCCPassword - Enter the password of the ccuser account.

Migrating zenhub service pool configuration

As a result of enhancement ZEN-30832, the zenhub service can now scale horizontally; that is, zenhub is no longer limited by the resources available to a single Control Center delegate. Your current configuration will be automatically migrate to take advantage of this capability.

If you have previously had a "superhub" configuration (characterized by having single zenhub instance in a dedicated Control Center resource pool) put in place by Zenoss Services or Support to help mitigate scaling issues this enhancement now addresses, you should plan to manually adjust the pool this service is in. In short, there is no longer a need to a separate pool, and you should follow the follow manual process post upgrade to revert such configuration to make best use of deploy resources.

1. In the Control Center browser interface, edit the "zenhub" service and note the pool it is configured for. Typically this is something like "zenhub".
2. In the Control Center browser interface, note the pool the top level service is configured for (edit service right at the top of the Zenoss application page". Typically this is something like "default" or "resource manager"
3. Determine whether the pool in step 2 is currently properly resourced to accept the load from zenhub. Check the pool as a whole can support the increased typical memory usage for zenhub, and that the pool as a whole can support the additional CPU load for zenhub (approximately 1+ (number of zenhubworker instances)+(0.5 * number of invalidation workers listed in zenhub.conf) cores of work).
4. Stop the zenhub and zenhubworker services.

5. In the Control Center browser interface, edit the zenhub service and change the pool configuration from the pool noted in step 1, to the pool noted in step 2.
6. Repeat step 5 for the zenhubworker service.
7. If you need to move the existing delegate from the pool in step 1 to the pool in step 2, based on the results of the analysis in step 3, do so now.
8. Start the zenhub and zenhubworker services.
9. If you need did NOT need to move the existing delegate from the pool in step 1 to the pool in step 2, you can deregister the delegate from the pool now.
10. Delete the pool identified in step 1.

Fixed issues

ID	Description
ZEN-31101	The zenhub service is performing tasks that should be performed by the zenhubworker service.
ZEN-22178	Default "HostPolicy" for several RM services has been changed to PREFER_SEPARATE
ZEN-30966	On the INFRASTRUCTURE page, an error occurs when filtering devices by priority.
ZEN-30945	When making a local copy of a template, the original template is displayed.
ZEN-30892	Modeling jobs initiated in the browser interface are not given priority and may be delayed.
ZEN-30871	The default number of Zauth instance is too low. See Zope deadlock improvements .
ZEN-30866	No debug data is written to the ZODB <code>connection_info</code> table.
ZEN-30835	SNMP discovery is not parallel.
ZEN-30832 ZEN-29857	The zenhub service cannot scale horizontally. See Zenhub scaling enhancements .
ZEN-30803	The LDAP plugin does not support adding trusted certificates.
ZEN-30795	SNMP traps sent through a NAT do not include source IP address.
ZEN-30790	Devices added with the default device class (/) can't be found by class.
ZEN-30789	Zenbatchload displays managelp warning on all devices.
ZEN-30774	The Ping Status Issues and SNMP Status Issues reports don't work on new installations.
ZEN-30695	The modeler plugins for the /Server/SSH/Linux class are SNMP, not command.
ZEN-30579	The defaults of the Add maintenance window dialog box do not match its most frequent use case.
ZEN-30540	Reports run automatically when selected.
ZEN-30438	Modeling and monitoring with SNMP v1 does not work properly.
ZEN-30437	API method allows non-privileged users to view privileged data.
ZEN-30388	Unable to lock Dashboard portlets from updates.

ZEN-30363 ZEN-29514	The version of OpenSSL included in Resource Manager (1.0.1) is no longer supported.
ZEN-30361	The warning messages from the zenmib utility are not readable (one character per line).
ZEN-30344	Unable to view debug information about a modeling attempt from the browser interface. See Debug option for device modeling .
ZEN-30325	Moving devices from /Discovered to another class fails, leaving devices in an inconsistent state.
ZEN-30024	The zenmib utility is unable to process overlapping ranges.
ZEN-29982	The reportmail feature does not work for graph or multi-graph reports.
ZEN-29978	AWS account password shown in plain text on Firefox browser.
ZEN-29929	The password field in ServiceNow notifications loses the entries in the password fields when other fields are subsequently filled.
ZEN-29927 ZEN-29898	The zenperfsnmp daemon does not close sockets after collections and eventually stops working.
ZEN-29878	The Solr service logs too many INFO messages.
ZEN-29877	The Solr service does not limit the number of results that can be returned for a query. See Configurable query limit .
ZEN-29856	The zencatalogscan utility times out and eventually fails.
ZEN-29838	In browser interfaces, flares that should include hyperlinks have only HTML elements.
ZEN-29796	The Edit Trigger dialog does not resize properly.
ZEN-29781	Zope services can become deadlocked. See Zope deadlock improvements .
ZEN-29775	The Zope self-monitoring script hangs when the instance is unresponsive but still listening on TCP. See Zope deadlock improvements .
ZEN-29764	Default configuration of maximum connections for MariaDB does not match recommended configuration.
ZEN-29762	The default minimum RAM requirement for the memcached service is too small.
ZEN-29702	Upgrade script reads the FROM version from the wrong location after an initial upgrade failure.
ZEN-29654	The zenmail service is unable to parse the hostname from an email address.
ZEN-29646	Silent failure when re-identifying a device that has no performance.
ZEN-29553	When the SSH password is incorrect, monitoring fails and the resulting event flaps (auto-clear and re-appear).
ZEN-29542	Asynchronous jobs for moving devices fail with <code>global name 'log' is not defined</code> message.
ZEN-29469	(MultiRealm ZenPack) The IP addresses associated with a device do not move to a new realm until the device is remodeled.
ZEN-29452	When adding multiple devices, the duration field starts at a large negative number instead of zero.
ZEN-	Monitoring template name handling does not support case-sensitive names.

29430	
ZEN-29393	The default minimum RAM requirement for the zenhub service is too small.
ZEN-29308	The upgrade log does not include a summary of the migration scripts that were run during an upgrade.
ZEN-29264	Large multi-graph reports can't open in new tab due to long URLs.
ZEN-29220	MariaDB does not perform well with the default memory allocator <code>glibc</code> ; switch to <code>jemalloc</code> .
ZEN-29062	Changes made to the LDAP user object classes value outside of the browser interface are not maintained when the LDAP properties form is saved in the browser interface.
ZEN-28887	No option to disable the removal of HTML/XML tags from plaintext email notifications.
ZEN-28676	The page notification type is unavailable, so integrations with external paging systems is not possible.
ZEN-28619	Time values are used and displayed inconsistently in the browser interface.
ZEN-27106	The <code>zeneventserver</code> log can get flooded with messages about events being dropped due to missing fields.
ZEN-25930	Unable to specify fields to include/exclude when exporting events.

Known issues

The following list describes the known issues in this release:

- (ZPS-4986, [ZenPacks.zenoss.HttpMonitor](#), v3.0.4) PythonCollector crashes when the IP Address or Proxy Address field is blank. After the crash, the `zenpython` debug log includes content similar to the following example:

```
2018-12-14 17:31:09,797 INFO zen.zenpython.twisted: Stopping protocol <twisted.names.dns.DNSDatagramProtocol object at 0x94599fd0>
2018-12-14 17:31:09,798 INFO zen.zenpython.twisted: DNSDatagramProtocol starting on 57548
```

- To avoid the crashes, provide content for the appropriate field(s).
- (ZPS-4473, [ZenPacks.zenoss.Microsoft.Windows](#), v2.9.1) When a Windows device reboots, data from PerfMon data sources do not resume collection. The workaround is to remodel the device.
- (ZEN-31194) This release includes `ZenPacks.zenoss.LDAPAuthenticator` v3.3.3, which no longer includes an option to skip SSL certificate verification during installation or upgrade. If you are using a self-signed certificate, upload it before upgrading. For more information, see [Adding an SSL certificate](#).
- (ZEN-30457) When processing messages from large devices, zenhub invalidation workers can get disconnected from Rabbit MQ due to heartbeat timeouts. If you are seeing the invalidations queue grow without also coming back down, or the number of consumers on the invalidations queue less than the number of running workers, please contact Zenoss Support.
- This release includes an update of the Python `urllib3` library, from version 1.10.2 to version 1.22. The new version checks TLS/SSL certificates by default. If you are using a self-signed certificate, upload it before updating Resource Manager or ZenPacks that rely on `urllib3`. This change also affects Zenoss Analytics (ZEN-31334). For more information, refer to the [Zenoss Support Knowledgebase](#).

Installed ZenPacks

This section lists ZenPacks that are automatically installed, those that are packaged but not installed, and those that are obsolete and should be uninstalled. For more information about ZenPacks, see the [ZenPack Catalog](#).

This release of Resource Manager installs the following ZenPacks at the current version listed in the table.

ZenPack	Current version	Previous version
ZenPacks.zenoss.AdvancedSearch	2.0.1	2.0.0
ZenPacks.zenoss.AixMonitor	2.2.3	Same
ZenPacks.zenoss.ApacheMonitor	2.1.4	Same
ZenPacks.zenoss.AuditLog	1.4.1	Same
ZenPacks.zenoss.AWS	4.0.2	4.0.1

ZenPacks.zenoss.CalculatedPerformance	2.5.1	2.5.0
ZenPacks.zenoss.CiscoMonitor	5.9.0	Same
ZenPacks.zenoss.CiscoUCS	2.8.0	2.7.0
ZenPacks.zenoss.ComponentGroups	1.7.0	1.6.0
ZenPacks.zenoss.ControlCenter	1.6.2	Same
ZenPacks.zenoss.Dashboard	1.3.3	1.2.9
ZenPacks.zenoss.Dell.PowerEdge	2.0.4	Same
ZenPacks.zenoss.Diagram	1.3.2	1.3.1
ZenPacks.zenoss.DistributedCollector	3.1.7	3.1.6
ZenPacks.zenoss.DnsMonitor	3.0.1	Same
ZenPacks.zenoss.Docker	2.0.4	2.0.3
ZenPacks.zenoss.DurationThreshold	2.0.5	Same
ZenPacks.zenoss.DynamicView	1.6.2	1.6.1
ZenPacks.zenoss.EMC.base	2.1.0	Same
ZenPacks.zenoss.EnterpriseCollector	1.8.3	Same
ZenPacks.zenoss.EnterpriseReports	2.5.0	Same
ZenPacks.zenoss.EnterpriseSecurity	1.2.0	Same
ZenPacks.zenoss.EnterpriseSkin	3.3.5	Same
ZenPacks.zenoss.HP.Proliant	3.3.2	Same
ZenPacks.zenoss.HttpMonitor	3.0.4	Same
ZenPacks.zenoss.IBM.Power	1.1.2	Same
ZenPacks.zenoss.InstalledTemplatesReport	1.1.1	Same
ZenPacks.zenoss.JuniperMonitor	2.1.1	Same
ZenPacks.zenoss.LDAPAuthenticator	3.3.3	3.3.1
ZenPacks.zenoss.LDAPMonitor	1.4.2	Same
ZenPacks.zenoss.Licensing	0.3.0	0.2.1
ZenPacks.zenoss.LinuxMonitor	2.3.2	2.2.7
ZenPacks.zenoss.Microsoft.Azure	1.3.3	1.3.0
ZenPacks.zenoss.Microsoft.Windows	2.9.2	2.9.0
ZenPacks.zenoss.MySqlMonitor	3.1.0	Same
ZenPacks.zenoss.NetAppMonitor	3.6.0	Same
ZenPacks.zenoss.NtpMonitor	3.0.0	Same
ZenPacks.zenoss.PredictiveThreshold	1.3.0	1.2.2
ZenPacks.zenoss.PropertyMonitor	1.1.1	Same
ZenPacks.zenoss.PythonCollector	1.10.1	Same
ZenPacks.zenoss.RMMonitor	1.1.1	1.1.0
ZenPacks.zenoss.SolarisMonitor	2.5.1	Same
ZenPacks.zenoss.StorageBase	1.4.3	Same
ZenPacks.zenoss.SupportBundle	1.1.2	Same
ZenPacks.zenoss.vSphere	4.0.0	3.7.2
ZenPacks.zenoss.WBEM	2.1.0	2.0.1

ZenPacks.zenoss.WSMAN	1.0.3	1.0.1
ZenPacks.zenoss.ZenDeviceACL	2.3.0	Same
ZenPacks.zenoss.ZenJMX	3.12.1	Same
ZenPacks.zenoss.ZenMail	5.1.0	Same
ZenPacks.zenoss.ZenOperatorRole	2.2.0	Same
ZenPacks.zenoss.ZenPackLib	2.1.1	Same
ZenPacks.zenoss.ZenSQLTx	2.7.1	Same
ZenPacks.zenoss.ZenWebTx	3.0.3	Same

Packaged ZenPacks

The following ZenPacks are packaged with Resource Manager, but not automatically installed:

- ZenPacks.zenoss.BigIpMonitor
- ZenPacks.zenoss.BrocadeMonitor
- ZenPacks.zenoss.Ceph
- ZenPacks.zenoss.CheckPointMonitor
- ZenPacks.zenoss.CiscoAPIC
- ZenPacks.zenoss.DatabaseMonitor
- ZenPacks.zenoss.DB2
- ZenPacks.zenoss.HpuxMonitor
- ZenPacks.zenoss.JBossMonitor
- ZenPacks.zenoss.Memcached
- ZenPacks.zenoss.Microsoft.Exchange
- ZenPacks.zenoss.Microsoft.Lync
- ZenPacks.zenoss.Microsoft.MSMQ
- ZenPacks.zenoss.NetScaler
- ZenPacks.zenoss.NetScreenMonitor
- ZenPacks.zenoss.NSX
- ZenPacks.zenoss.OpenStack
- ZenPacks.zenoss.OpenStackInfrastructure
- ZenPacks.zenoss.OpenvSwitch
- ZenPacks.zenoss.PostgreSQL
- ZenPacks.zenoss.RabbitMQ
- ZenPacks.zenoss.TomcatMonitor
- ZenPacks.zenoss.XenServer

Obsolete ZenPacks

ZenPacks.zenoss.ZenMailTX is obsolete. If this ZenPack is installed at your site, uninstall it.

Zenoss JSON API changes

The Zenoss JSON API includes the following changes from version 6.2.0:

- (ZEN-29558) The `getDevices` method of ModelQuery router is updated to support offsets. The method requires additional fields, and the default limit on the number of devices it returns is 200.
- The `exportDevices` and `importDevices` methods of the DeviceDumpLoadRouter class now require ZenManager privileges.
- The DeviceRouter class includes the following enhancements:
 - The `getInfo` method now includes improved security for password fields.
 - The `setCollector` method includes improved handling of debug messages.
 - The `addDevice` method includes a bug fix in the duplicate device message.

Resource Manager 6.2.1

This section contains important information about release 6.2.1 of Zenoss Resource Manager (Resource Manager).

- [Considerations and workarounds](#)
- [Fixed issues](#)
- [Known issues](#)
- [Installed ZenPacks](#)

Considerations and workarounds

New download site

Downloads for Zenoss customers are now available on delivery.zenoss.com. Leapfile is no longer used.

Compatibility with Zenoss Service Impact

This version of Resource Manager is compatible with Zenoss Service Impact version 5.2.3 or later.

Regular expression matching limitation

(ZEN-29376) Resource Manager uses the Apache Solr search platform, which tokenizes fields. Currently, searching a tokenized field with a regex matches only a single token, limiting complex regex searches.

For example, on the INFRASTRUCTURE page, IP Address is a tokenized field. Regex search `^192.168` will not work to filter only IP addresses that start with 192.168. The Device column is untokenized; regex search on the device name works as expected.

Load time for component graphs

(ZEN-29300) Load time might exceed 10 seconds for more than 200 component graphs when you have activated the All on same graph check box.

NFS client 4.1 is not supported

A file locking defect might exist in NFS 4.1 with RHEL/CentOS 7.4, which could cause zeneventserver to crash and other DFS-related problems. For more information, see [Configuring NFS 4.0](#), or Knowledge Base article [Potential Issues Running With RHEL 7.4 Or CentOS 7.4](#).

CentralQuery maximum memory allocation pool

(ZEN-15907) The size of the CentralQuery maximum memory allocation pool is set by the RAMCommitment variable in the CentralQuery service definition. The default value is 1024MB. Do not change the value to anything less than 1024MB.

Upgrade considerations

- (ZEN-29807) Upgrading from Resource Manager 5.x to 6.x includes a required step to reindex all cataloged objects into the new Solr index. To speed this process, before upgrading, you can increase the CPU core count to use more workers. For more information, see knowledge base article [Reducing 5.X To 6.X Upgrade Times By Increasing Workers For Indexing](#).
- Beginning with release 6.2.1, support for upgrading versions earlier than 5.3.2 was removed from the Zenoss Resource Manager Upgrade Guide. For assistance upgrading an earlier version, please contact Zenoss Support.
- The upgrade process might reinstall an older version of the Catalog Service ZenPack. To avoid this issue, manually remove older versions of the Catalog Service ZenPack .egg file from the ZenPack backups directory.
- (ZEN-28375) Beginning with Resource Manager 6.1.0, for counter/derive (rate) data points, the calculated rate value is stored instead of raw counters. The rate is calculated at the collector daemon as it collects data. If a graph (or API) request for a derive of counter data point spans the upgrade data, the system automatically queries the data correctly. However, slight anomalies might occur in the data that is captured immediately before and after an upgrade from Resource Manager 6.0.1 or earlier.
- (ZEN-29100) No data is collected for the ZenossRM device after an upgrade if the localhost collector is moved off the master host. If you have moved the localhost collector off the Control Center master host, to enable data collection, set properties for the ZenossRM device as follows:
 1. Navigate to the ZenossRM device overview page and select Configuration Properties.
 2. Set the following zProperties:
 - zRMMonCCHost - Enter the IP address of the Control Center master host.
 - zRMMonCCUser - Enter the name of the ccuser account, which is the default account for gaining access to the Control Center browser interface.
 - zRMMonCCPassword - Enter the password of the ccuser account.

Fixed issues

ID	Description
ZEN-30495	Modeling and monitoring with SNMP v1 does not work properly.
ZEN-30324	Moving devices between classes increases resource consumption and leaves devices in an inconsistent state.
ZEN-30308, ZEN-30216	SNMP library for Python increases timeout values exponentially.
ZEN-30217	Upgrade script for v6.1.2 tags snapshots incorrectly.
ZEN-29542	Moving a device succeeds but the job reports failure (global name 'log' not defined).
ZEN-29391	Using zenbatchload on a v6.1.0 system with a file generated by zenbatchdump on a v5.3.3 system results in warnings about a missing managelp function.
ZEN-29062	LDAP configuration screen does not include a field for the user object class.
ZPS-3867	HTTP Monitor ZenPack version 3.0.3 does not include support for regular expressions, so page content can not be checked.

Known issues

ID	Description
ZEN-26802	OS model link in Device Detail page points to wrong manufacturer entry
ZEN-27499	Error message regarding dropped Events displayed during Resource Manager upgrade.
ZEN-28138	objectGUID is not available to be selected in Login Name Attribute combo box in LDAP configuration options.
ZEN-28519	Error is displayed when a correct date/time is entered in the Date Range field.
ZEN-28716	On Events page, the Show only actionable events check box is not displayed for the ZenOperator role.
ZEN-28725	On the Dashboard page, a ZenManager can see a dashboard even though that user is part of a group with a restriction.
ZEN-28956	Cisco UCS reports organizer is hidden when using ZenOperator role.
ZEN-29100	No data collected for the RM device after an upgrade if the localhost collector is moved off of the master host. For detailed instructions, see Upgrade considerations .
ZEN-29109	Disable Transforms report is not working.
ZEN-29120	Error flare messages appear intermittently on the Advanced > Control Center subtab.
ZEN-29376	Unable to use regex matching in tokenized fields. For more information, see Considerations and workarounds.
ZEN-29807	Leverage compute resources available on master to speed upgrades. For a workaround, see knowledge base article Reducing 5.X To 6.X Upgrade Times By Increasing Workers For Indexing .

Installed ZenPacks

This section lists ZenPacks that are automatically installed, those that are packaged but not installed, and those that are obsolete and should be uninstalled. For more information about ZenPacks, see the [ZenPack Catalog](#).

This release of Resource Manager installs the following ZenPacks at the current version listed in the table.

ZenPack	Current version	Previous version
ZenPacks.zenoss.AdvancedSearch	2.0.0	Same
ZenPacks.zenoss.AixMonitor	2.2.3	Same
ZenPacks.zenoss.ApacheMonitor	2.1.4	Same
ZenPacks.zenoss.AuditLog	1.4.1	Same

ZenPacks.zenoss.AWS	4.0.1	Same
ZenPacks.zenoss.CalculatedPerformance	2.5.0	Same
ZenPacks.zenoss.CiscoMonitor	5.9.0	5.8.1
ZenPacks.zenoss.CiscoUCS	2.7.0	2.6.2
ZenPacks.zenoss.ComponentGroups	1.6.0	Same
ZenPacks.zenoss.ControlCenter	1.6.2	Same
ZenPacks.zenoss.Dashboard	1.2.9	Same
ZenPacks.zenoss.Dell.PowerEdge	2.0.4	Same
ZenPacks.zenoss.Diagram	1.3.1	Same
ZenPacks.zenoss.DistributedCollector	3.1.6	Same
ZenPacks.zenoss.DnsMonitor	3.0.1	Same
ZenPacks.zenoss.Docker	2.0.3	Same
ZenPacks.zenoss.DurationThreshold	2.0.5	2.0.4
ZenPacks.zenoss.DynamicView	1.6.1	Same
ZenPacks.zenoss.EMC.base	2.1.0	2.0.0
ZenPacks.zenoss.EnterpriseCollector	1.8.3	1.8.1
ZenPacks.zenoss.EnterpriseReports	2.5.0	Same
ZenPacks.zenoss.EnterpriseSecurity	1.2.0	Same
ZenPacks.zenoss.EnterpriseSkin	3.3.5	3.3.4
ZenPacks.zenoss.HP.Proliant	3.3.2	3.3.1
ZenPacks.zenoss.HttpMonitor	3.0.4	3.0.3
ZenPacks.zenoss.IBM.Power	1.1.2	Same
ZenPacks.zenoss.InstalledTemplatesReport	1.1.1	Same
ZenPacks.zenoss.JuniperMonitor	2.1.1	Same
ZenPacks.zenoss.LDAPAuthenticator	3.3.1	Same
ZenPacks.zenoss.LDAPMonitor	1.4.2	Same
ZenPacks.zenoss.Licensing	0.2.1	Same
ZenPacks.zenoss.LinuxMonitor	2.2.7	Same
ZenPacks.zenoss.Microsoft.Azure	1.3.0	Same
ZenPacks.zenoss.Microsoft.Windows	2.9.0	Same
ZenPacks.zenoss.MySqlMonitor	3.1.0	Same
ZenPacks.zenoss.NetAppMonitor	3.6.0	Same
ZenPacks.zenoss.NtpMonitor	3.0.0	Same
ZenPacks.zenoss.PredictiveThreshold	1.2.2	Same
ZenPacks.zenoss.PropertyMonitor	1.1.1	Same
ZenPacks.zenoss.PythonCollector	1.10.1	Same
ZenPacks.zenoss.RMMonitor	1.1.0	1.0.4
ZenPacks.zenoss.SolarisMonitor	2.5.1	Same
ZenPacks.zenoss.StorageBase	1.4.3	Same
ZenPacks.zenoss.SupportBundle	1.1.2	Same
ZenPacks.zenoss.vSphere	3.7.2	3.6.3

ZenPacks.zenoss.WBEM	2.0.1	Same
ZenPacks.zenoss.WSMAN	1.0.1	Same
ZenPacks.zenoss.ZenDeviceACL	2.3.0	Same
ZenPacks.zenoss.ZenJMX	3.12.1	Same
ZenPacks.zenoss.ZenMail	5.1.0	Same
ZenPacks.zenoss.ZenOperatorRole	2.2.0	Same
ZenPacks.zenoss.ZenPackLib	2.1.1	2.0.9
ZenPacks.zenoss.ZenSQLTx	2.7.1	Same
ZenPacks.zenoss.ZenWebTx	3.0.3	3.0.2

Packaged ZenPacks

The following ZenPacks are packaged with Resource Manager, but not automatically installed:

- ZenPacks.zenoss.BigIpMonitor
- ZenPacks.zenoss.BrocadeMonitor
- ZenPacks.zenoss.Ceph
- ZenPacks.zenoss.CheckPointMonitor
- ZenPacks.zenoss.CiscoAPIC
- ZenPacks.zenoss.DatabaseMonitor
- ZenPacks.zenoss.DB2
- ZenPacks.zenoss.HpuxMonitor
- ZenPacks.zenoss.JBossMonitor
- ZenPacks.zenoss.Memcached
- ZenPacks.zenoss.Microsoft.Exchange
- ZenPacks.zenoss.Microsoft.Lync
- ZenPacks.zenoss.Microsoft.MSMQ
- ZenPacks.zenoss.NetScaler
- ZenPacks.zenoss.NetScreenMonitor
- ZenPacks.zenoss.NSX
- ZenPacks.zenoss.OpenStack
- ZenPacks.zenoss.OpenStackInfrastructure
- ZenPacks.zenoss.OpenvSwitch
- ZenPacks.zenoss.PostgreSQL
- ZenPacks.zenoss.RabbitMQ
- ZenPacks.zenoss.TomcatMonitor
- ZenPacks.zenoss.XenServer

Obsolete ZenPacks

ZenPacks.zenoss.ZenMailTX is obsolete. If this ZenPack is installed at your site, uninstall it.

Appendixes

- [Resource Manager interface reference](#)
- [Using the Appliance Administration menu](#)
- [Managing Zope instances](#)
- [SNMP device preparation](#)
- [Syslog device preparation](#)
- [TALES expressions](#)
- [Managing multi-realm networks](#)
- [Monitoring large file systems](#)
- [Integrating LDAP authentication](#)
- [Tuning Considerations](#)

Resource Manager interface reference

This section provides brief descriptions of the tabs and pages that are provided in the Resource Manager browser interface.

- [DASHBOARD \(Resource Manager\)](#)
- [EVENTS](#)
- [INFRASTRUCTURE](#)
- [SERVICES](#)
- [REPORTS](#)
- [ADVANCED](#)

DASHBOARD (Resource Manager)

The DASHBOARD tab displays status information about your Resource Manager environment, in the *dashboard* that is assigned to, or chosen by, the current user.

A Resource Manager dashboard is a grid of *portlets*, which are customizable subwindows that are designed to display specific types of information. The number of columns in a dashboard grid can be customized, and portlets can be sized and arranged as desired within the grid. You can create multiple dashboards and share them with other Resource Manager users.

Portlets are available for a variety of internal and external information, including:

- Resource Manager device events
- Google Maps views integrated with information about your infrastructure
- Resource Manager reports
- External web pages

For more information about portlets, see [Portlets](#).

For more information about dashboards, see [Administering Resource Manager dashboards](#).

EVENTS

The EVENTS tab features the pages described in the following sections.

Event Console

The EVENTS > Event Console page displays all current events, and is the default page displayed when you click the EVENTS tab.

For more information, see [Event management](#).

Event Archive

The EVENTS > Event Archive page displays closed events (events that are closed, cleared, or aged).

For more information, see [Event management](#).

Event Classes

The EVENTS > Event Classes page displays the [event classes](#) that are installed in Resource Manager.

For more information, see [Understanding event classes](#).

Triggers

The EVENTS > Triggers page displays the [triggers](#) and [notifications](#) that are used by the event processing service.

For more information, see [Triggers and notifications](#).

INFRASTRUCTURE

The INFRASTRUCTURE tab features the pages that are described in the following sections.

Devices

The INFRASTRUCTURE > Devices page provides access all of the devices in your environment, and is the default page displayed when you click the INFRASTRUCTURE tab.

Networks

The INFRASTRUCTURE > Networks page displays all of the IPv4 and IPv6 networks that have been modeled on your network.

Processes

The INFRASTRUCTURE > Processes page provides tools for monitoring processes on devices. For more information, see [Monitoring processes](#).

IP Services

The INFRASTRUCTURE > IP Services page provides tools for monitoring IP services on devices. For more information, see [Monitoring IP services](#).

Windows Services

The INFRASTRUCTURE > Windows Services page provides tools for monitoring Windows Services on devices. For more information, see the [Configuring Services Monitoring](#) entry on the Microsoft Windows ZenPack Catalog page.

Network Map

The INFRASTRUCTURE > Network Map page displays a representation of your network's layer 3 topology. For more information, see [The network map page](#).

Manufacturers

The INFRASTRUCTURE > Manufacturers page displays information about the products, by manufacturer, that are present in your environment.

SERVICES

The SERVICES tab features the pages described in the following sections.

DYNAMIC SERVICES

The SERVICES > DYNAMIC SERVICES page displays health summaries of all services, and is the default page displayed when you click the SERVICES tab.

LOGICAL NODES


The SERVICES > LOGICAL NODES page displays...

REPORTS

The REPORTS tab provides access to the reports that are available in Resource Manager. This section describes the reports that are available in all Resource Manager installations.

When you install a ZenPack that includes a report, the report is added to the REPORTS tab. For more information about ZenPack-specific reports, refer to the [ZenPacks page](#).

You can organize reports and the display order of report organizers by drag-and-drop within the tree view.



Organizing reports

You can organize reports by creating organizers and moving reports into them. You can create report organizers at multiple levels, even within another organizer. To create a report organizer:

1. Select an existing organizer or the top of the reports hierarchy, and then click Add.
2. Click Add Report Organizer.
3. In the Create Report Organizer dialog box, enter the name of the new report organizer, and then click Submit. The report organizer appears in the tree view.
4. Move reports into the organizer, or create new reports.

Scheduling reports

By default, all reports run on demand, presenting information in the interface when you run the report. You can also schedule a report to be run at a specific time for a list of recipients. For more information, please contact Zenoss Support.

Troubleshooting problems with report generation

If you experience stair-stepping in graphs, consider changing the reporting collection interval in Resource Manager. For example, setting the reporting collection interval to 60 minutes tells Resource Manager to update the API-driven reporting data at that interval, which is different from the native collection interval.

If browser interface responsiveness slows during report generation, you may need additional instances of the Zope web application server. By default, Resource Manager provides one dedicated instance for reporting. For more information, please contact Zenoss Support.

AWS Reports

Monitoring Costs

This report provides a detailed breakdown of API calls and the estimated cost per monitoring template on each monitored EC2 account.

Azure

Unattached VHDs

This report shows the unattached virtual hard drives, if any, in your Microsoft Azure environment.

Capacity Planning

Capacity Usage

The Capacity Usage report is configurable by resource type, and includes options for sorting by capacity metrics.

Cisco UCS Capacity Reports

The reports in this section highlight Cisco UCS capacity calculations.

- In reports that have a graphical display, you can click Printable to generate a printable version of the graphs in the report.
- In reports that are tabular, you can click Export all to generate a CSV file showing the data in spreadsheet form.

These reports are only available if [CiscoUCS](#) version 3.0.0 or higher is installed.

Reports:

- [Aggregate Bandwidth Utilization](#)
- [Aggregate Port Pool Utilization](#)
- [Bandwidth Utilization vs Capacity](#)
- [HostNetworkCapacity](#)
- [Interface 95th Percentile](#)
- [Interface Utilization](#)
- [Interface Volume \(Cisco UCS\)](#)
- [Port Utilization](#)
- [Storage Utilization vs Capacity](#)

Aggregate Bandwidth Utilization

A summary of the aggregate throughput shown per chassis, per fabric extender, and per I/O module for each domain. Three graphs are shown per domain. You can see the individual components of these graphs on the particular device's component page.

For example, to see the throughput on fex-2 of a particular device, you would click INFRASTRUCTURE, then click the device name and select Fabric Extenders from the Components list. Click fex-2 to display the graph of the throughput in the bottom pane. This represents this fabric extender's part in the total throughput displayed in this report.

Report filtering

The screenshot shows a dark-themed report filtering interface. It includes the following fields and options:

- Start Date:** 2019/05/28 (with a 'select' button)
- End Date:** 2019/06/04 (with a 'select' button)
- Summary Type:** Avg (dropdown)
- Domain Filter:** (empty text input)
- Fabric Interconnect:** ANY (dropdown)
- Traffic Type:** ANY (dropdown)
- Chassis:** Chassis 1, Chassis 2, Chassis 3, Chassis 4 (multi-select list)
- FEX:** FEX 1, FEX 2, FEX 3, FEX 4 (multi-select list)
- IO Module:** IO Module 1, IO Module 2, IO Module 3, IO Module 4 (multi-select list)
- FI-IO Module:** FI-IO Module 1, FI-IO Module 2, FI-IO Module 3, FI-IO Module 4 (multi-select list)
- Trendline Type:** Linear (dropdown)
- Buttons:** Update, Printable

The following fields filter the results.

Start Date / End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Summary Type

The type of summary that is to be displayed. Average is based on the 15-minute data collection, while Max is only taken once a day.

Domain Filter

A substring search of domains to include. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Fabric Interconnect

Select the FI to filter by.

Traffic Type

Type of traffic to display. Choose from ANY, Tx or Rx.

Chassis, FEX, IO Module, Shared Adaptor, FI-IO Module

Select the items to filter by. You can multi-select by Ctrl-clicking your selections.

Trendline Type

The calculation used to determine the trendline. Only available selection is Linear.

To generate or refresh the report, click Update.

To generate a printable version of the report, click Printable, then use the print functionality on your browser.

The following are some example scenarios that may be seen in graphs along with their corresponding description:

- If a graph shows 16.48 G (Current value) for Chassis-3, it is the aggregate (or sum) of Receiving (Rx) and Sending (Tx) bits of Chassis-3
- If a graph shows 1.89 M (Max value) for Fex-2, it is the aggregate (or sum) of the Receiving (Rx) and Sending (Tx) bits that happens on the ports connected to Fex-2.
- If a graph shows 17.15 G (Average value) for an IO module (for example, Chassis-1_slot-1), it is the aggregate (or sum) of Receiving (Rx) and Sending (Tx) bits that are transferred specifically to the IO module.

Aggregate Port Pool Utilization

A summary of the aggregate port pool throughput per domain. Each domain is displayed in a separate graph with a separate line for each server port. You can see the individual components of these graphs on the particular device's component page

Report filtering

The screenshot shows a dark-themed filter panel with the following fields and controls:

- Start Date:** 2019/05/28 (with a 'select' button)
- End Date:** 2019/06/04 (with a 'select' button)
- Port Filter:** ANY (dropdown menu)
- Trendline Type:** Linear (dropdown menu)
- Summary Type:** Avg (dropdown menu)
- Domain Filter:** (empty text input field)
- Traffic Type:** ANY (dropdown menu)
- Update** button
- Printable** button

The following fields filter the results.

Start Date / End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Summary Type

The type of summary that is to be displayed. Average is based on the 15-minute data collection, while Max is only taken once a day.

Domain Filter

A substring search of domains to include. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Port Filter

Select the type of ports to filter by. Options include ANY, Port Channels, Server Ports, and Stand Alone Ports.

Traffic Type

Type of traffic to display. Choose from ANY, Tx or Rx.

Trendline Type

The calculation used to determine the trendline. Only available selection is Linear.

To generate or refresh the report, click Update.

The following are some example scenarios that may be seen in graphs along with their corresponding description:

- If Switch-A to Fex-2 server ports shows a current value of 1G, the server ports that connect Switch A and Fex-2 transmit a total of 1Gbps, which includes both receiving and sending.
- If Switch-A to Chassis-1 server ports shows an average value of 2G, the server ports that connect Switch-A and Chassis-1 transmit an average of 2Gbps, which includes both receiving and sending.
- If Switch-B to Rack Server-1 server ports shows a maximum value of 5G, the server ports that connect Switch-B and Rack Server-1 transmit a total of 5Gbps, which includes both receiving and sending.

Bandwidth Utilization vs Capacity

A summary of the total bandwidth utilization compared with capacity for each domain. Each domain has its own graph. The graph shows the percentage of utilization for both Rx (receiving or incoming) and Tx (transmitting or outgoing) network traffic. You can view individual components of this graph by viewing the Fabric Interconnect Capacity Utilization graph which can be found on the Fabric Interconnect component of your device.

Report filtering

The screenshot shows a dark-themed report filtering interface. It contains the following fields and controls:

- Start Date:** 2019/05/28 with a 'select' button.
- End Date:** 2019/06/04 with a 'select' button.
- Traffic Type:** ANY (dropdown menu).
- Trendline Type:** Linear (dropdown menu).
- Summary Type:** Avg (dropdown menu).
- Domain Filter:** (empty text input field).
- Fabric Interconnect:** ANY (dropdown menu).
- Update** and **Printable** buttons.

The following fields filter the results.

Start Date / End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Summary Type

The type of summary that is to be displayed. Average is based on the 15-minute data collection, while Max is only taken once a day.

Domain Filter

A substring search of domains to include. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Traffic Type

Type of traffic to display. Choose from ANY, Tx or Rx.

Fabric Interconnect

Select the FI to filter by.

Trendline Type

The calculation used to determine the trendline. Only available selection is Linear.

To generate or refresh the report, click Update.

The following is an example scenarios that may be seen in graphs along with its corresponding description:

- If sys_switch-A shows a value of 10% for Rx, it uses only 10% of this fabric interconnect switch, e.g., if the fabric interconnect has a capacity of 800Gbps for Rx, then only 80Gbps is used for Rx. A similar calculation can be made for Tx.

HostNetworkCapacity

A summary of the network capacity of various host types in all domains. The graph shows the utilization for both Rx (receiving or incoming) and Tx (transmitting or outgoing) network traffic.

Report filtering

The following fields filter the results.

Start Date / End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Host Type

The type of Host to display. Choose from ANY, HyperV, vSphere or other.

Summary Type

The type of summary that is to be displayed. Average is based on the 15-minute data collection, while Max is only taken once a day.

Domain Filter

A substring search of domains to include. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Traffic Type

Type of traffic to display. Choose from ANY, Tx or Rx.

Chassis Filter / FEX Filter / Blade Server Filter / Rack Server Filter / Cartridge Server Filter

Select the item or items to include in the report. You can select multiple items with CTRL-click.

Trendline Type

The calculation used to determine the trendline. Only available selection is Linear.

To generate or refresh the report, click Update.

To generate a printable version of the report, click Printable, then use the print functionality on your browser.

Interface 95th Percentile

The Interface 95th Percentile report shows both the average input and output traffic rate as well as the 95th percentile value of the input and output traffic rate. The report is generated in a tabular form with the calculations shown by each interface.

Report filtering

Device Class: /Devices/Network
Device Filter:
Start Date: 2019/05/28 select
End Date: 2019/06/04 select
Trendline Type: Linear
Update
Interface 95th Percentile

The following fields filter the results.

Device Class

The device class associated with the interfaces to include in the report. The default is /Devices/Network, all network devices.

Device Filter

A complete or partial name of the interfaces to include in the report. Letter case is ignored. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Start Date / End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

To generate or refresh the report, click Update.

Report contents

Column	Content
Device	The name of the device that contains the interface, with a link to its overview page.
Interface	The name of the interface, with a link to its overview page.
Description	A description of the interface.
Speed	The maximum transfer rate the interface supports, per second.
In Avg	The average input traffic through the interface, per second.
Out Avg	The average output traffic through the interface, per second.
95% In	The input traffic rate of the 95th percentile.
95% Out	The output traffic rate of the 95th percentile.
Forecasted In Exhaustion	Projected date of exhaustion for input traffic.
Forecasted Out Exhaustion	Projected date of exhaustion for output traffic.

The following shows the navigation bar at the bottom of the report. To quickly navigate to other pages, select the device name from the popup list. In this example, there are 5 pages of results from device 10.171.100.107 followed by one page of results from device 10.171.100.109, and so on.

10.171.100.107	Ethernet108 1 46	23.3Kb	1.4b	1.4b	3.6b	3.4b
10.171.100.107	Ethernet108 1 47	22.5Kb	1.4b	1.4b	3.4b	3.2b
10.171.100.107	Ethernet108 1 48	22.8Kb	1.4b	1.3b	3.0b	3.4b
10.171.100.107	10.171.100.107	23.8Kb	1.4b	1.3b	3.0b	2.9b
10.171.100.107	10.171.100.107	22.8Kb	1.4b	1.4b	3.4b	3.3b
10.171.100.107	10.171.100.107	22.3Kb	1.3b	1.4b	3.2b	3.3b
10.171.100.107	10.171.100.109	23.7Kb	1.4b	1.3b	3.0b	3.0b
10.171.100.107	10.171.100.14	22.7Kb	1.4b	1.3b	3.4b	3.0b
10.171.100.107	10.171.100.14	23.8Kb	1.4b	1.4b	3.3b	3.2b
10.171.100.107	10.171.100.14	22.9Kb	1.3b	1.4b	3.0b	3.4b
10.171.100.107	10.171.100.88	23.3Kb	1.4b	1.4b	3.3b	3.3b
10.171.100.107	perf2-switch	23.6Kb	1.4b	1.3b	3.1b	2.8b

1 of 1017 |<< 10.171.100.107 ▾ >>| show all export all Page Size 100 ok

Click Export all to export the report as a CSV file.

Interface Utilization

The Interface Utilization report shows the average, maximum, and minimum input and output traffic rate. The report is generated in a tabular form with the calculations shown by each interface.

Report filtering

Root Organizer: /Devices/Network ▾ Device Filter:

Start Date: 2019/05/28 End Date: 2019/06/04

Trendline Type: Linear ▾

Show All Interfaces

Interface Utilization

The following fields filter the results.

Device Class

The device class associated with the interfaces to include in the report. The default is /Devices/Network, all network devices.

Device Filter

A complete or partial name of the interfaces to include in the report. Letter case is ignored. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Start Date / End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

To generate or refresh the report, click Update. If you want to show all interfaces on the report, check the Show All Interfaces box before clicking Update.

Report contents

Column	Content
Path	The path of the device class.
Device	The name of the device that contains the interface, with a link to its overview page.
Interface	The name of the interface, with a link to its overview page.
Description	A description of the interface.
Speed	The maximum transfer rate the interface supports, per second.
In Avg	The average input traffic through the interface, per second.
Out Avg	The average output traffic through the interface, per second.
In Max	The maximum input traffic through the interface, per second.
Out Max	The maximum output traffic through the interface, per second.
In Min	The minimum input traffic through the interface, per second.
Out Min	The minimum output traffic through the interface, per second.
95% In	Projected date of exhaustion for input traffic.
95% Out	Projected date of exhaustion for output traffic.

In order for the In Max and Out Max values to be calculated, you need to have at least 8 hours worth of data, otherwise a N/A value will be returned.

The following shows the navigation bar at the bottom of the report. To quickly navigate to other pages, select the device name from the popup list. In this example, there are 5 pages of results from device 10.171.100.107 followed by one page of results from device 10.171.100.109, and so on.

10.171.100.107	Ethernet108 1 46	23.3Kb	1.4b	1.4b	3.6b	3.4b
10.171.100.107	Ethernet108 1 47	22.5Kb	1.4b	1.4b	3.4b	3.2b
10.171.100.107	Ethernet108 1 48	22.8Kb	1.4b	1.3b	3.0b	3.4b
10.171.100.107	10.171.100.107	23.8Kb	1.4b	1.3b	3.0b	2.9b
10.171.100.107	10.171.100.107	22.8Kb	1.4b	1.4b	3.4b	3.3b
10.171.100.107	10.171.100.107	22.3Kb	1.3b	1.4b	3.2b	3.3b
10.171.100.107	10.171.100.107	23.7Kb	1.4b	1.3b	3.0b	3.0b
10.171.100.107	10.171.100.109	22.7Kb	1.4b	1.3b	3.4b	3.0b
10.171.100.107	10.171.100.14	23.8Kb	1.4b	1.4b	3.3b	3.2b
10.171.100.107	10.171.100.14	22.9Kb	1.3b	1.4b	3.0b	3.4b
10.171.100.107	10.171.100.14	23.3Kb	1.4b	1.4b	3.3b	3.3b
10.171.100.107	10.171.100.88	23.6Kb	1.4b	1.3b	3.1b	2.8b
	perf2-switch					

1 of 1017 |<< 10.171.100.107 ▾ >>| show all export all Page Size 100 ok

Click Export all to export the report as a CSV file.

Interface Volume (Cisco UCS)

The Interface Volume report shows the total input and output traffic for the report period, along with a calculation of the input and output traffic per day. The report is generated in a tabular form with the calculations shown by each interface.

Report filtering

The screenshot shows a dark-themed interface for report filtering. It contains the following elements:

- Device Class:** A dropdown menu showing "/Devices/Network".
- Device Filter:** An empty text input field.
- Start Date:** A date input field showing "2019/05/28" with a "select" button.
- End Date:** A date input field showing "2019/06/04" with a "select" button.
- Trendline Type:** A dropdown menu showing "Linear".
- Update:** A button to refresh the report.
- Interface Volume:** A label at the bottom left.
- Search:** A search icon at the bottom right.

The following fields filter the results.

Device Class

The device class associated with the interfaces to include in the report. The default is /Devices/Network, all network devices.

Device Filter

A complete or partial name of the interfaces to include in the report. Letter case is ignored. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Start Date / End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

To generate or refresh the report, click Update. If you want to show all interfaces on the report, check the Show All Interfaces box before clicking Update.

Report contents

Column	Content
Path	The path of the device class.
Device	The name of the device that contains the interface, with a link to its overview page.
Interface	The name of the interface, with a link to its overview page.
Description	A description of the interface.
Speed	The maximum transfer rate the interface supports, per second.
In Vol	The input traffic through the interface for the time period of the report.
In Vol/day	The input traffic through the interface, per day.
Out Vol	The output traffic through the interface for the time period of the report.
Out Vol/day	The output traffic through the interface, per day.
Total Vol	Total volume of traffic through the interface for the time period of the report.

The following shows the navigation bar at the bottom of the report. To quickly navigate to other pages, select the device name from the popup list. In this example, there are 5 pages of results from device 10.171.100.107 followed by one page of results from device 10.171.100.109, and so on.

10.171.100.107	Ethernet108 1 46	23.3Kb	1.4b	1.4b	3.6b	3.4b
10.171.100.107	Ethernet108 1 47	22.5Kb	1.4b	1.4b	3.4b	3.2b
10.171.100.107	Ethernet108 1 48	22.8Kb	1.4b	1.3b	3.0b	3.4b
10.171.100.107	10.171.100.107	23.8Kb	1.4b	1.3b	3.0b	2.9b
10.171.100.107	10.171.100.107	22.8Kb	1.4b	1.4b	3.4b	3.3b
10.171.100.107	10.171.100.107	22.3Kb	1.3b	1.4b	3.2b	3.3b
10.171.100.107	10.171.100.107	23.7Kb	1.4b	1.3b	3.0b	3.0b
10.171.100.107	10.171.100.109	22.7Kb	1.4b	1.3b	3.4b	3.0b
10.171.100.107	10.171.100.14	23.8Kb	1.4b	1.4b	3.3b	3.2b
10.171.100.107	10.171.100.14	22.9Kb	1.3b	1.4b	3.0b	3.4b
10.171.100.107	10.171.100.14	23.3Kb	1.4b	1.4b	3.3b	3.3b
10.171.100.107	10.171.100.88	23.6Kb	1.4b	1.3b	3.1b	2.8b
	perf2-switch					

1 of 1017 | << 10.171.100.107 >> | show all export all Page Size 100 ok

Click Export all to export the report as a CSV file.

Port Utilization

The report displays the total bandwidth broken down by LAN uplink, fiber channel uplink, appliance, and server ports.

Report filtering

Start Date:	2019/05/28	select	Summary Type:	Avg	▼
End Date:	2019/06/04	select	Domain Filter:		
Traffic Type:	ANY	▼			
Trendline Type:	Linear	▼			
Update		Printable			

The following fields filter the results.

Start Date / End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Summary Type

The type of summary that is to be displayed. Average is based on the 15-minute data collection, while Max is only taken once a day.

Domain Filter

A substring search of domains to include. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Traffic Type

Type of traffic to display. Choose from ANY, Tx or Rx.

Trendline Type

The calculation used to determine the trendline. Only available selection is Linear.

To generate or refresh the report, click Update.

The following is an example scenarios that may be seen in graphs along with its corresponding description:

- If the Server ports current value is 10G, then all the ports that have the Server port interface role are transmitting data at a combined rate of 10Gbps.
- If the LAN Uplink ports maximum value is 200M, then the ports connected from the Fabric Interconnect switch to the LAN network transfers data at 200Mbps.

Storage Utilization vs Capacity

A summary of the total storage utilization compared with capacity for each storage device. Each device has its own group of graphs and may display some of the following information: Storage Utilization, Bytes Per Cycle, Managed Space Usage, Time Utilization, Data Throughput, Raw Space Usage, Datastore IO Rate, and Datastore IO Disk Usage. The graphs display both the current utilization and capacity along with projected values all on the same graph. Use the following report parameters to filter the report to the devices and type of information you are most interested in.

Report filtering

The screenshot shows a dark-themed filter panel with the following fields and values:

- Start Date: 2019/05/28 (with a 'select' button)
- End Date: 2019/06/04 (with a 'select' button)
- Storage Class: ANY (dropdown)
- Storage Structure: ANY (dropdown)
- Storage Space Usage: ANY (dropdown)
- Summary Type: Avg (dropdown)
- Device Filter: (empty text input)
- Trendline Type: Linear (dropdown)
- Storage IO: ANY (dropdown)

At the bottom, there are two buttons: 'Update' and 'Printable'.

The following fields filter the results.

Start Date / End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Summary Type

The type of summary that is to be displayed. Average is based on the 15-minute data collection, while Max is only taken once a day.

Device Filter

A substring search of devices to include. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Storage Class

Select the storage device class to filter by. You can select all devices class by selecting ANY.

Trendline Type

The calculation used to determine the trendline. Only available selection is Linear.

Storage Structure

Type of storage structure. Choose from ANY, Disk or Inodes.

Storage IO

Type of storage process. Choose from ANY, Read or Write.

Storage Space Usage

Type of storage space usage. Choose from ANY, Total, Used, Subscribed, Managed, Uncommitted, Raw, or Capacity.

To generate or refresh the report, click Update.

Cisco UCS Reports

For each of the following Cisco UCS reports, you can click Export all to generate a CSV file showing the data in spreadsheet form.

Free Memory Slots

This report lists the number of free memory slots in each component grouped by domain. No filtering capability is provided in this report.

Report contents

Column	Content
Unnamed column	The domain name.
Server	The component being reported on.
Free Slots	The number of free memory slots on the corresponding component. At the bottom of this column, a list of total free memory slots is displayed.

To export the report as a CSV file, click [Export all](#).

Hardware Inventory

This reports lists the inventory of the UCS devices that are being monitored.

Report contents

Column	Content
UCS Manager	The name of the UCS Manager. Clicking the link takes you to the overview page.
Component	Lists components and related sub-components of the device. Clicking a link takes you to the appropriate device component page.
Manufacturer	The manufacturer of the component.
Model	The model number of the component.
Serial #	The serial number of the component.
Description	Detailed information about the component.

To export the report as a CSV file, click Export all.

Custom Device Reports

Creating a custom device report

To create a custom device report:

1. Navigate to REPORTS > Custom Device Reports
2. Click Add and select Add Custom Device Report from the popup menu.
3. Enter a name for the report in the Create Custom Device Report dialog box and click SUBMIT.
4. Define the following report parameters:
 - Name: Edit the report name if needed.
 - Title: Enter the report title. This title is displayed in the report and is distinct from the report name.
 - Path: Specify the path in the hierarchy where you want the system to store the report.
 - Query: Specify the actual query string for the report. For example, if you want to limit the report to just those devices with a serial number, you can set the query value to:

```
here.hw.serialNumber != ""
```

- Sort Column: Specify the column on which you want to sort the report by default.
- Sort Sense: Specify the sense that the system uses to sort:asc (ascending sort) or desc (descending sort)
- Columns: Specify the data to be retrieved and displayed in the report. For example you could specify:
 - getID: Gets the name of any device.
 - getManagelp: Gets the IP addresses of the devices.
 - getHWSerialNumber: Gets the serial numbers of the devices.

Device Reports

All Devices

A summary of each device that Resource Manager is monitoring.

Column	Content
Name	The name of the device.
Class	The Resource Manager device class associated with the device.
Product	The hardware model information associated with the device, which is provided by the device's SNMP MIB, or entered manually. If the value in this column is an SNMP OID, the Resource Manager database does not include a definition of the object.
State	The device's production state. Valid states include Production, Pre-Production, Test, and Maintenance.
Ping	The result of the most recent ping of the device.
SNMP	The result of the most recent attempt to gather data through the device's SNMP agent.

All Monitored Components

A summary of each component Resource Manager is monitoring.

Column	Content
Device	The name of the device which contains the component, with a link to its overview page.
Component	The name of the component, with a link to its overview page.
Type	The class associated with the component.
Description	A description of the component; typically, the component's name.
Status	The state of the component as of the most recent attempt to gather monitoring data.

Device Changes

A summary of the devices in which changes were detected during the most recent collection of model data.

Column	Content
Name	The name of the changed device, with a link to its overview page.
Class	The Resource Manager device class associated with the device.
First Seen	The timestamp of the initial collection of modeling data for the device.
Collection	The timestamp of the collection in which a change was detected before the most recent collection.
Change	The timestamp of the most recent collection in which a change was detected.

HP Chassis List

This report provides a list of the devices in HP Proliant rack/blade servers.

HP Device Bay List

This report provides a list of the devices in HP Proliant blade servers.

MAC Addresses

A list of the unique device name, interface ID, and MAC address combinations in the Resource Manager database.

Column	Content
Device	The name of a device, with a link to its overview page.
Interface ID	The ID of a network interface, with a link to its overview page.
MAC address	A MAC address.

Model Collection Age

A summary of devices that were not available for modeling data collection during the most recent 48 hour period.

Column	Content
Name	The name of the changed device, with a link to its overview page.
Class	The Resource Manager device class associated with the device.
First Seen	The timestamp of the initial collection of modeling data for the device.
Collection	The timestamp of the most recent collection of modeling data.
Change	The timestamp of the most recent change in modeling data for the device.

New Devices

The list of devices that were discovered and added to Resource Manager recently.

Column	Content
Name	The name of the changed device, with a link to its overview page.
Class	The Resource Manager device class associated with the device.
First Seen	The timestamp of the initial collection of modeling data for the device.
Collection	The timestamp of the most recent collection of modeling data.
Change	The timestamp of the most recent change in modeling data for the device.

Ping Status Issues

A list of the devices which were down during the most recent collection of monitoring data.

Column	Content
Name	The name of the device.
Class	The Resource Manager device class associated with the device.
Product	The hardware model information associated with the device, which is provided by the device's SNMP MIB, or entered manually. If the value in this column is an SNMP OID, the Resource Manager database does not include a definition of the object.
State	The device's production state. Valid states include Production, Pre-Production, Test, and Maintenance.
Ping	The result of the most recent ping of the device.
SNMP	The result of the most recent attempt to gather data through the device's SNMP agent.

SNMP Status Issues

A list of the devices for which no SNMP agent responded during the most recent collection of monitoring data.

Column	Content
Name	The name of the device.
Class	The Resource Manager device class associated with the device.
Product	The hardware model information associated with the device, which is provided by the device's SNMP MIB, or entered manually. If the value in this column is an SNMP OID, the Resource Manager database does not include a definition of the object.
State	The device's production state. Valid states include Production, Pre-Production, Test, and Maintenance.
Ping	The result of the most recent ping of the device.
SNMP	The result of the most recent attempt to gather data through the device's SNMP agent.

Software Inventory

A list of the software installed in the devices which Resource Manager monitors.

Column	Content
Manufacturer	The name of the company that makes the software product.
Product	The name of the software product.
Count	The total number of devices or components on which the software is installed.

Enterprise Reports

For each of the following Enterprise reports, you can click Export all to generate a CSV file showing the data in spreadsheet form.

Cisco Inventory

This report lists all the Cisco devices being monitored.

Report filtering

Device Class:	<input type="text" value="/Network/Cisco"/>
Group:	<input type="text" value="All"/>
<input type="button" value="Update"/>	

Device Class

The device class to use for filtering. The default is /Network/Cisco.

Group

The specific group to consider when running the report. The default is All. The group could be an internal department, a location, a customer, and so on.

To generate or refresh the report, click Update.

Report contents

Column	Content
Name	The name of the Cisco device. Clicking the link takes you to the device overview page.
IP Address	Lists the IP address of the Cisco device.
Model	The model of the device. Clicking the link takes you to the Manufacturers Overview page for that product.
Serial #	The serial number of the Cisco device.
Type	The type of Cisco product; for example, Device.

Customized Performance Templates

Customized Performance Templates	
Target	Template
www.zenoss.com	Ping Local Copy
github.com	Device
apache1.zenoss.loc	HttpMonitor
localhost.localdomain	HttpMonitor
oracle1.zenoss.loc	Device
export all	

Data Sources in Use

This reports lists the data sources defined in the system.

Report contents

Column	Content
Device Class	The device class of the data source in use.
Template	The template associated with the data source in use.
Data Source Name	The name of the data source.
Data Source Type	The type of connection for the data source.
ZenPack	The ZenPack related to the data source.

Datapoints Per Collector

This report shows the number of devices and data points per collector, which is useful for gauging how much monitoring load is on each collector.

Report contents

Column	Content
Hub	Name of the hub the collector is on
Collector	Name of the collector the devices are on
Hostname	Hostname location
Device Count	Total number of devices on the collector
Datapoint Count	Total number of data points being generated on the collector

Defined Thresholds

The Defined Thresholds report provides details about all thresholds that are defined in the system. The report links to the target of each threshold. The target can be a device class, individual device, or individual component.

This report is useful for administering the system. You can use it to quickly identify which threshold events can occur within the system, and the severity of those events.

Report contents

Column	Content
Target	The device class of the defined threshold.
Template	The template associated with the defined threshold.
Threshold	The name of the defined threshold. Clicking the link takes you to the Performance Template where the threshold is defined.
Severity	The severity level assigned to the alert when the threshold is reached.
Enabled	The enabled status of the threshold.

Event Time to Resolution

This report shows, for each user, the total time taken to acknowledge or clear events. Results are organized by user. It is helpful for tracking response time SLAs in a NOC-type environment.

Report contents

Column	Content
User	The user responsible for taking action on an event.
Severity	The event severity level .
Time to Ack	The amount of time to acknowledge the event.
Time to Archive	The amount of time to clear and archive the event.

Interface Utilization (enterprise report)

The Interface Utilization report shows the average, maximum, and minimum input and output traffic rate. The report is generated in a tabular form with the calculations shown by each interface.

Report filtering

Root Organizer:
 Device Filter:
 Start Date:
 End Date:
 Show All Interfaces

The following fields filter the results.

Device Class

The device class associated with the interfaces to include in the report. The default is /Devices/Network, all network devices.

Device Filter

A complete or partial name of the interfaces to include in the report. Letter case is ignored. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

To generate or refresh the report, click Update. If you want to show all interfaces on the report, check the Show All Interfaces box before clicking Update.

Report contents

Column	Content
Path	The path of the device class.
Device	The name of the device that contains the interface, with a link to its overview page.
Interface	The name of the interface, with a link to its overview page.
Description	A description of the interface.
Speed	The maximum transfer rate the interface supports, per second.
In Avg	The average input traffic through the interface, per second.
Out Avg	The average output traffic through the interface, per second.
In Max	The maximum input traffic through the interface, per second.
Out Max	The maximum output traffic through the interface, per second.
In Min	The minimum input traffic through the interface, per second.
Out Min	The minimum output traffic through the interface, per second.
95% In	Projected date of exhaustion for input traffic.
95% Out	Projected date of exhaustion for output traffic.

Note: For the In Max and Out Max values to be calculated, you must have at least eight hours of data; otherwise, a the value N/A is returned.

The following image shows the navigation bar at the bottom of the report. To quickly navigate to other pages, select the device name from the popup list. This example has five pages of results from device 10.171.100.107, and one page of results from device 10.171.100.109, and so on.

To export the report as a CSV file, click Exportall.

Interface Volume

The Interface Volume report shows the total input and output traffic for the report period, along with a calculation of the input and output traffic per day. The report is generated in a tabular form with the calculations shown by each interface.

Report parameters

Root Organizer: Device Filter:
 Start Date: End Date:
 Show All Interfaces

The following fields filter the results.

Device Class

The device class associated with the interfaces to include in the report. The default is /Devices/Network, all network devices.

Device Filter

A complete or partial name of the interfaces to include in the report. Letter case is ignored. A partial name matches all interface names that include it. You may also use a regular expression for this filter.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

To generate or refresh the report, click Update.

Report contents

Column	Content
Device	The name of the device that contains the interface, with a link to its overview page.
Interface	The name of the interface, with a link to its overview page.
Description	A description of the interface.
Speed	The maximum transfer rate the interface supports, per second.
In Vol	The input traffic through the interface for the time period of the report.
In Vol/day	The input traffic through the interface, per day.
Out Vol	The output traffic through the interface for the time period of the report.
Out Vol/day	The output traffic through the interface, per day.
Total Vol	Total volume of traffic through the interface for the time period of the report.

The following image shows the navigation bar at the bottom of the report. To quickly navigate to other pages, select the device name from the popup list. This example has five pages of results from device 10.171.100.107, and one page of results from device 10.171.100.109, and so on.

10.171.100.107	Ethernet108_1_46	23.3Kb	1.4b	1.4b	3.6b	3.4b
10.171.100.107	Ethernet108_1_47	22.5Kb	1.4b	1.4b	3.4b	3.2b
10.171.100.107	Ethernet108_1_48	22.8Kb	1.4b	1.3b	3.0b	3.4b
10.171.100.107	10.171.100.107	23.8Kb	1.4b	1.3b	3.0b	2.9b
10.171.100.107	10.171.100.107	22.8Kb	1.4b	1.4b	3.4b	3.3b
10.171.100.107	10.171.100.107	22.3Kb	1.3b	1.4b	3.2b	3.3b
10.171.100.107	10.171.100.107	23.7Kb	1.4b	1.3b	3.0b	3.0b
10.171.100.107	10.171.100.109	22.7Kb	1.4b	1.3b	3.4b	3.0b
10.171.100.107	10.171.100.14	23.8Kb	1.4b	1.4b	3.3b	3.2b
10.171.100.107	10.171.100.14	22.9Kb	1.3b	1.4b	3.0b	3.4b
10.171.100.107	10.171.100.14	23.3Kb	1.4b	1.4b	3.3b	3.3b
10.171.100.107	10.171.100.88	23.6Kb	1.4b	1.3b	3.1b	2.8b
10.171.100.107	perf2-switch	23.6Kb	1.4b	1.3b	3.1b	2.8b

1 of 1017 | << 10.171.100.107 >> | show all export all Page Size 100 ok

To export the report as a CSV file, click Export all.

Maintenance Windows (enterprise report)

This report shows all defined windows that are active during a selected time period. To change the reporting time period, enter Start and End dates (or click **Select** to select dates from a calendar). Click **Update** to refresh the report.

Report contents

Column	Content
Name	Name of the maintenance window.
Target	Name of the device.
Start Date	Date the maintenance window becomes active.
Duration	Length of time for the maintenance window to be in effect.
Repeat	Repeat interval of the maintenance window.

Network Topology

This report shows the layout of the network according to the routes that Resource Manager understands, starting from the collector and ending at the remote devices that are associated with the collector.

Troubleshooting this report

- The report does not return data if the host on which the Resource Manager collector is running does not have a device created in the device management database (DMD) object that stores the basic model of the network in the Zope database (ZODB).
To resolve this issue, create a device that represents the collector in the DMD, and then run the report again.
- An invalid route entry (for example, Missing link here value in the Route column) indicates that Resource Manager cannot determine how to route from one device to another.
To resolve this issue:
 1. Add a network interface to the model (no new hardware is required).
 2. Add a new route entry from the last device in the route to the device (the IP address shown at the far right of the table).

Report filtering settings

Report Settings

Collectors: localhost ▾

Show valid routes?:

Show invalid routes?:

[Update](#)

Collectors

Select the collector on which to base the report. Connections from this collector to associated devices is shown in the report.

Show valid routes?

Select this check box to show valid connections between the selected collector and remote devices.

Show invalid routes?

Select this check box to show invalid connections between the selected collector and remote devices. An invalid route indicates that a route from the collector to the device could not be determined.

To generate or refresh the report, click Update.

Report contents

Column	Content
Collector	Name of the collector.
Route	Route from the collector to the device.
Device IP Address	IP address of the device.
Device Name	Name of the device.
Repeat	Repeat interval of the maintenance window.

Notifications and Triggers by Recipient

This report lists all of the notifications that the system has sent.

Column	Content
Name	The recipient name.
Address	The address associated with the recipient.
Notification	The content of the notification.
Triggered by	The trigger that initiated the notification.

Organizer Availability

This report provides the availability percentage of all network organizers in the system. It can be filtered by organizer, event class, component, event severity, and date.

You can report on the availability of device classes, locations, systems, or groups within a defined time frame. This report offers two reporting modes:

- **Averaged** - Defines the organizer as available for the average availability time for all devices contained in it.
- **Coalesced** - Defines the organizer as available only if all devices are available during a certain time period.

Report filtering

Organizer Availability Filtering

Root Organizer:	<input type="text" value="/Devices"/>	Component:	<input type="text"/>	Event Class:	<input type="text" value="/Status/Nagios"/>
Start:	<input type="text" value="11/25/2016"/> <input type="button" value="select"/>	Severity:	<input type="text" value="Critical"/>		<input type="text" value="/Status/Ntp"/>
End:	<input type="text" value="12/02/2016"/> <input type="button" value="select"/>	Summation:	<input type="radio"/> Averaged <input checked="" type="radio"/> Coalesced		<input type="text" value="/Status/OSProcess"/>
<input type="button" value="Update"/>					<input type="text" value="/Status/Perf"/>
					<input type="text" value="/Status/Ping"/>

Root Organizer

The device class to use for filtering. The default is /Devices.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Component

The specific group to consider when running the report. The default is All. The group could be an internal department, a location, a customer, and so on.

Severity

The [event severity level](#) to filter by. The default is Critical. If another level is wanted, select it from the drop-down list.

Summation

Select between Averaged and Coalesced depending on how you want to define the organizer as available. The default is Coalesced.

Event Class

The event class to use for filtering. The default is /Status/Ping.

To generate or refresh the report, click Update.

Note: If you export the report, be sure to format the percentage columns to show percentages instead of decimal values.

Report contents

Column	Content
Name	Name of the device class based on the root organizer selected.
Availability	Percent availability of the selected event class.
Total	Total availability of the selected event class.

Organizer Graphs

This reports shows graphical data about a specific organizer. The information varies depending on the type of organizer. For example, selecting /Devices /Network/Cisco as an organizer displays graphs on CPU and Memory Utilization, Throughput, Errors, and so on.

Report filtering

Organizer Graphs Filtering

Organizer: Filter: Start: End:

Organizer

The class to use for filtering. The default is /Groups.

Filter

Additional text filter to refine results. The default is .* which returns all graphs available for the selected organizer.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

To generate or refresh the report, click Update.

User Event Activity

A list of users or groups and the number of events each has acknowledged and archived.

Report filtering

Start:	<input type="text" value="11/25/2016"/>	<input type="button" value="select"/>	End:	<input type="text" value="12/02/2016"/>	<input type="button" value="select"/>	Group By:	<input checked="" type="radio"/> User <input type="radio"/> Group	<input type="button" value="Update"/>
--------	---	---------------------------------------	------	---	---------------------------------------	-----------	---	---------------------------------------

The following fields define the report.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Group By

The type of report to create, User or Group.

To generate the report, click Update.

Report contents

Column	Content
User or Group	The column label and content depend on the report selection. User The name of a user, with a link to the user's USER SETTINGS page. Group The name of a group, with a link to the group's Users in Group page.
# Acknowledged	The total number of events acknowledged by the user or group during the reporting period.
# Archived	The total number of events archived (cleared) by the user or group during the reporting period.

Users Group Membership

A list of Resource Manager users and the groups to which they belong.

Column	Content
User	The name of a user, with a link to the user's USER SETTINGS page.
Groups	The list of groups to which the user belongs. Each name includes a link to the group's Users in Group page.

Event Reports

All EventClasses (All Event Classes)

A list of each item in the event hierarchy in Resource Manager. Each item (class) includes the total number of subclasses, instances, and events associated with the class.

Column	Content
Name	The event class name.
Sublasses	The total number of subclasses associated with the event class.
Instances	The total number of instances of the class and its subclasses.
Events	The total number of events associated with the class.

All EventMappings (All Event Mappings)

A list of each item in the event mapping hierarchy in Resource Manager. Each item (event mapping) includes its key and example text, and a count of the events associated with the event mapping.

Column	Content
Name	The name of the event mapping, which includes its location in the event class hierarchy, and its key.
EventClassKey	The unique identifier of the event mapping.
Evaluation	A portion of the example associated with the event mapping.
Events	The total number of events associated with the event mapping.

All Heartbeats

A list of all Resource Manager daemons, showing the number of seconds elapsed since each daemon sent a heartbeat event.

Column	Content
Device	The device on which the daemon is running.
Component	The name of the daemon.
Seconds	The number of seconds elapsed since the daemon sent a heartbeat event.

Disabled Transforms

This report shows all of the event classes in which transforms are disabled. Some classes are configured to disable transforms when too many errors occur; this report identifies those classes quickly.

Graph Reports

Graph reports allow you to assemble graphs from devices and device components into a single report. Graph reports only display those graphs that already exist on devices or components in the system. You cannot define or alter graphs in a graph report.

Graph reports are available in two views: a "normal" view (similar to the graph views for devices and device classes) and a print view.

Creating a graph report

To create a graph report:

1. Navigate to REPORTS > Graph Reports
2. Click Add and select Add Graph Report from the popup menu.
3. Enter a name for the report in the Create Graph Report dialog box and click SUBMIT.
4. In the Edit Report screen, verify or edit the following information in the Graph Report section:
 - Name: The name of the report as defined in the Create Graph Report dialog box.
 - Title: Enter a descriptive title to display in the list of reports for the report organizer.
 - Number of Columns: Specify the number of columns (1-3) in which graphs will be displayed on the report.
 - Comments: Enter comments to display at the top of the printable version of the report. This is a TALES evaluated string that can contain HTML formatting. The variables available to the TALES expression are:
 - now (current date and time)
 - report (report object)
5. Click Save to save the new graph report.
6. In the Add New Graph section, fill in the appropriate information to add a graph to the report:
 - Device: Select the one or more devices. You can narrow the list of devices by entering a search string and clicking Filter.
 - Component: Optionally, select one or more components from the Component list. This list displays the names of all components defined on at least one of the selected device. If you want to see the complete component path in the system to help with identification, select the Show component path check box.
7. Select one or more graphs from the Graph list. This list displays the names of all the graphs available for the selected devices; or, if you have selected one or more components, the graphs available for the components.
8. Click Add Graph to Report. The selected graphs are displayed in the Graphs section. You can resequence or delete graphs using the Action icon. Graph reports maintain a static list of graphs. This list does not automatically change when graphs are added or deleted from monitoring templates. For example, you have two devices with associated graphs:
 - DeviceA: Has a single graph (Graph1)
 - Device B: Has two graphs (Graph1 and Graph2)If you select DeviceA and DeviceB on the Graph Report edit page, the list of graphs will include Graph1 and Graph2. If you select both graphs and add them to the report, you will see three graphs:
 - DeviceA - Graph1
 - Device B - Graph1
 - Device B - Graph2

If, at a later time, you create a second graph (Graph2) on DeviceA's monitoring templates, that new graph will not automatically appear on the graph report. You must edit the report to add it. Similarly, if you later remove a graph from DeviceB's template (or even delete DeviceB from the system), you must manually remove the graph (or device) from the graph report.

Working with graph reports

You can customize the display of the graphs that are contained in any graph report that you created. You can change the graph name by clicking on the graph name in the Graphs section of the Edit Report screen. You can also edit the text that appears with the graph when viewing the report:

- Summary: Displays above the graph in the normal report view. It may contain TALES expressions with these variables:
 - dev - Device
 - comp - Component
 - graph - Graph
- Comments: Displays to the left of the graph in the printable view. May contain TALES expressions with these variables:
 - dev - Device
 - comp - Component
 - graph - Graph

On the Graph report, you can control the information displayed by selecting the Range of data to display or by clicking the Zoom In/Zoom Out controls. Using this latter method, automatically adjusts the data range to Custom and you can fine-tune your date range.

Monitoring Capabilities Reports

Installed Templates

This report lists all of the monitoring templates installed on Resource Manager.

Monitoring Capabilities Report Export Sort by:

Templates: Table of Contents	ACEContext
ACEContext	
ACERealServer	
ACEServerFarm	
ACEServicePolicy	
AIXAme	

Device Class: /Network/Cisco/ACE

ZenPack Name: ZenPacks.zenoss.CiscoMonitor

Data Sources:

DataSource: ace; SSH Enabled: True; Type: COMMAND

 DataPoint: bandwidth-current, Type: GAUGE

 DataPoint: bandwidth-denied, Type: GAUGE

Templates are listed alphabetically in the left column and details about the templates are provided on the right.

Multi-Graph Reports

Multi-graph reports combine data from different devices and components into a single report. You can create a graph definition and have it drawn once for each of a group of devices and components that you define. Alternatively, you can combine the data for those graphs into a single graph.

The groups of devices and components you assemble are called "collections". Specifying the graph definition to apply to collections is done through graph group objects. Multi-graph reports include their own graph definitions, and thus do not use the graph definitions that are defined in monitoring templates. To create a report that includes graphs defined on templates, use a Graph report instead.

Adding collections

A collection comprises one or more collection items. A collection item can be a list of device classes, systems, groups, locations, or specific devices or components. A single collection may contain as many collection items as desired. A multi-graph report must contain at least one collection. Collections are shown in the Collections area of the report's Edit page.

To create a collection:

1. In the Collections area of the multi-graph report Edit page, click the Action icon and select Add Collection. The Add a Collection dialog box appears.
2. Enter a name for the collection, then click OK.
The Multi-Graph Report Collection dialog box appears.
3. In the Add To Collection area, select collection items to add to the collection:
 - a. Select a value for Item Type. If you select either Device Class, System, Group, or Location, then you can select one or more of the organizers to include in the collection. If you select Specific Device/Component, you will be able to choose from a list of all the devices in the system. You can use the Filter field to narrow the selection process. Selecting one or more devices will display a list of component names that apply to the selected devices.
 - b. Select a value for Include Suborganizers?. If true, the collection will also include all organizers recursively beneath the selected organizer. These collection items are dynamic, when devices are added or removed from the organizers, they will appear or disappear from the report.
4. Click Add to Collection to create a new collection item for each of the selected organizers or specific device. The collection item appears in the Collection Items area. If desired, you can re-order collection items. Their listed order determines the order in which the graphs are drawn, or the order that data is drawn on a combined graph.

Creating a multi-graph report

1. Navigate to REPORTS > Multi-Graph Reports
2. Click Add and select Add Multi-Graph Report from the popup menu.
3. Enter a name for the report in the Create Multi-Graph Report dialog box and click SUBMIT.
4. In the report edit page, enter or select values for the following:
 - Name: The name of the report as defined in the Create Multi-Graph Report dialog box.
 - Title: Enter a descriptive title to display in the list of reports for the report organizer.
 - Number of Columns: Specify the number of columns (1-3) in which graphs will be displayed on the report.
5. Click Save.
6. Add one or more of the following to define the source you want to graph:
 - Collections: Contain the devices and components you want to graph.
 - Graph Definitions: Describe the graphs you want on the report.
 - Graph Groups: Specify the collections and graph definition to use.

Adding graph definitions

In the context of multi-graph reports, graph definitions are very similar to those in monitoring templates. Settings on the graph definition define basic parameters. Graph points are added to specify which data should be drawn.

The most significant differences between graph definitions in the two contexts is how datapoint graph points and threshold graph points are added. When adding a data point graph point to a graph definition in a performance template, you can select from a list of datapoints that are defined on that template. In the context of a multi-graph report, there are no graph point definitions listed. You must enter the name of the data point on the datapoint graph point dialog.

1. Click the Action icon in the Graph Definitions area of the Graph edit page and select Add Graph.
The Add a New Graph dialog box appears.
2. Enter a name for the graph, then click OK.
The Edit Graph Definition page appears.
3. Make any changes to the graph definition values displayed, then click Save.
4. Click the Action icon in the Graph Points section to perform the following:
 - Add data points (see next topic)
 - Add thresholds
 - Add a custom graph point
 - Delete a graph point
 - Re-sequence graph points

Adding data points

To add a data point to a graph definition:

1. Ensure that you have created a graph definition as part of your multi-graph report.
2. On the Graph Definition page, click the Action icon in the Graph Points section and select Add DataPoint.
3. Enter a DataPoint Name. The field will auto-populate based on your entry. Click OK to save the name.
4. Click the name of the graph point you want to define.
The Edit screen appears.
5. Edit the fields based on your datapoint and the way you want data displayed. You can enter a custom RPN expression on this screen if needed.
6. Click Save.

Adding graph groups

Graph groups combine a graph definition with a collection to produce graphs for the report. In order for the report to show graphs, at least one graph group must be created.

To create a graph group:

1. Click the Action icon in the Graph Groups area of the Graph edit page and select Add Group.
The Add a New Graph Group dialog box appears.
2. Enter a name for the graph group, then click OK.
The Edit Graph Group Definition page appears.
3. Make any changes to the graph group values displayed:
 - Name: Identifies the graph group on the multi-graph report page. It does not appear on the report.
 - Collection: Select a collection that has been defined for this report.
 - Graph Definition: Select a graph definition that has been defined for this report.
 - Method: Choose between having the graph drawn once for each device and component in the collection or combining the data from all devices and components into a single graph. Options are:
 - Separate graph for each device: The graph definition is used to draw one graph for each device and component in the collection. Graphs will appear in the list in the same order they are specified in the collection.
 - All devices on a single graph: Draws one graph with the data from all devices and components included.
4. Click Save to save the graph group.

Re-sequencing graph group order

Graph groups are drawn in the order in which they are listed on the multi-graph report Edit page. To change the order of the graph groups:

1. On the Multi-Graph Edit Report page, edit the sequence order numbers (0,1, 2, and so on) beside the graph groups that you have defined.
2. From the Action icon, select Re-sequence items.
The page refreshes and displays the graph groups in the re-sequenced order. Note: If a graph group results in multiple graphs, the graphs are drawn in the order that the collection items are listed in the corresponding collection. If a collection item specifies a device organizer, the order of devices drawn from that collection item is indeterminate.

Performance Reports

Availability Report

Shows the percentage of time that a device is considered available. You can filter this report on a variety of criteria, including by a time period.

Report filtering

Device Class:	/	Systems:	/
Groups:	/	Locations:	/
Device Filter:		Severity:	Error
Start Date:	07/17/2017 <input type="button" value="select"/>	End Date:	07/24/2017 <input type="button" value="select"/>
Event Class:	/Status/Ping		
<input type="button" value="Generate"/>			

Device Class

The device class to use for filtering. The default is / (all device classes).

Systems

Select the systems to filter by. The default is / (all systems).

Groups

Select the groups to filter by. The default is / (all groups).

Locations

Select the locations to filter by. The default is / (all locations).

Device Filter

Enter the name of the device to filter by.

Severity

The severity level used in the availability calculation described below. The default is Critical. If another level is wanted, select it from the drop-down list.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Event Class

The event class to use for filtering. The default is /Status/Ping.

To generate or refresh the report, click Generate.

Note: If you export the report by clicking Export all, be sure to format the percentage columns to show percentages instead of decimal values.

The percent availability value is calculated by first summing the duration of all events of a particular class with a production state of Production and with a severity greater than or equal to a specified severity in the filter criteria. This sum is then divided by the total duration of the time range, and then subtracted from 1 and multiplied by 100 to get the percent available, as in the following equation:

$$1 - ((\text{Total event down time}) / (\text{total duration})) * 100$$

Note: Events whose firstime and lasttime fields are the same are not used in the calculation. These could represent an event that occurs and is subsequently cleared by the next event, or an event that has happened only once in the specific date range.

Report contents

Column	Content
Device	Name of the device based on the filter parameters selected.
Systems	Systems name if applicable to the filter.
Availability	Total availability of the selected devices.

CPU Utilization

Shows monitored devices, load averages, % utilization, and forecasted exhaustion. You can customize start and end dates.

Report filtering

Root Organizer:	<input type="text" value="/Devices"/>	Device Filter:	<input type="text"/>
Start Date:	<input type="text" value="07/17/2017"/> <input type="button" value="select"/>	End Date:	<input type="text" value="07/24/2017"/> <input type="button" value="select"/>
Summary Type:	<input type="text" value="Average"/>	Consolidation:	<input type="text" value="Average"/>
Trendline Type:	<input type="text" value="Linear"/>		

Root Organizer

The device class to use for filtering. The default is /Devices.

Device Filter

Enter the name of the device to filter by.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Summary Type

Possible values include: Average, Maximum, Minimum, and Last.

Consolidation

Possible values include: Average and Max.

Trendline Type

Projection algorithm used in Forecasted % Util Exhaustion calculation. Possible value: Linear

To generate or refresh the report, click Generate.

Note: If you export the report by clicking Export all, be sure to format the percentage columns to show percentages instead of decimal values. This report uses data point aliases. To add data points to a report, add the alias, and then ensure the values return in the expected units.

Alias	Expected Units
loadAverage5min	Processes
cpu_pct	Percent

Report contents

Column	Content
Device	Name of the device based on the filter parameters selected.
Load Avg	Average load on the device.
% Util	% CPU utilization on the device
Forecasted % Util Exhaustion	The amount of time before the exhaustion threshold will be breached.

Filesystem Util Report

Shows mount point, total bytes, used bytes, free bytes, and percentage of utilization for each device. You can customize start and end dates and summary type.

Report filtering

Root Organizer:	<input type="text" value="/Devices"/>	Device Filter:	<input type="text"/>
Start Date:	<input type="text" value="07/17/2017"/> <input type="button" value="select"/>	End Date:	<input type="text" value="07/24/2017"/> <input type="button" value="select"/>
Summary Type:	<input type="text" value="Average"/>		
<input type="button" value="Generate"/>			

Root Organizer

The device class to use for filtering. The default is /Devices.

Device Filter

Enter the name of the device to filter by.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Summary Type

Possible values include: Average, Maximum, Minimum, and Last.

To generate or refresh the report, click Generate.

Note: If you export the report by clicking Export all, be sure to format the percentage columns to show percentages instead of decimal values.

This report uses data point aliases. To add datapoints to a report, add the alias, and then ensure the values return in the expected units.

Alias	Expected Units
usedFilesystemSpace__bytes	bytes

Report contents

Column	Content
Device	Name of the device based on the filter parameters selected.
Mount	File systems mount point. Click the link to be taken directly to the device's components page.
Total bytes	Amount of total bytes
Used bytes	Amount of used bytes
Free bytes	Amount of free bytes
% Util	Percent utilization

Interface Utilization (performance report)

Shows the traffic through all network interfaces monitored by Resource Manager.

Report filtering

Root Organizer:	<input type="text" value="/Devices"/>	Device Filter:	<input type="text"/>
Start Date:	<input type="text" value="07/17/2017"/> <input type="button" value="select"/>	End Date:	<input type="text" value="07/24/2017"/> <input type="button" value="select"/>
Summary Type:	<input type="text" value="Average"/>		
<input type="button" value="Generate"/>			

Root Organizer

The device class to use for filtering. The default is /Devices.

Device Filter

Enter the name of the device to filter by.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Summary Type

Possible values include: Average, Maximum, Minimum, and Last.

To generate or refresh the report, click Generate. Note: If you export the report by clicking Export all, be sure to format the percentage columns to show percentages instead of decimal values.

This report uses data point aliases. To add datapoints to a report, add the alias, and then ensure the values return in the expected units.

Alias	Expected Units
inputOctets__bytes	bytes/sec
outputOctets__bytes	bytes/sec

Report contents

Column	Content
Device	Name of the device based on the filter parameters selected.
Interface	Name of interface. Click the link to be taken to the Device Components page.
Speed	The interface's rated bandwidth, in bits per second
Input	Average traffic coming in to the interface, in bits per second
Output	Average traffic going out of the interface, in bits per second
Total	Total average traffic across the interface, in bits per second
% Util	Percentage of the interface's bandwidth consumed

Memory Utilization

Provides system-wide information about the memory usage for devices in Resource Manager.

Report filtering

Root Organizer:	<input type="text" value="/Devices"/>	Device Filter:	<input type="text"/>
Start Date:	<input type="text" value="07/17/2017"/> <input type="button" value="select"/>	End Date:	<input type="text" value="07/24/2017"/> <input type="button" value="select"/>
Summary Type:	<input type="text" value="Average"/>	Consolidation:	<input type="text" value="Average"/>
Trendline Type:	<input type="text" value="Linear"/>		

Root Organizer

The device class to use for filtering. The default is /Devices.

Device Filter

Enter the name of the device to filter by.

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Summary Type

Possible values include: Average, Maximum, Minimum, and Last.

Consolidation

Possible values include: Average and Max.

Trendline Type

Projection algorithm used in Forecasted % Util Exhaustion calculation. Possible value: Linear

To generate or refresh the report, click Generate. Note: If you export the report by clicking Export all, be sure to format the percentage columns to show percentages instead of decimal values.

The report uses data point aliases. To add datapoints to the report, add the alias, and then ensure the values return in the expected units.

Alias	Expected Units
memoryAvailable__bytes	bytes
memoryBuffered__bytes	bytes
memoryCached__bytes	bytes

Report contents

Column	Content
Device	Name of the device based on the filter parameters selected.
Total	Amount of total memory
Available	Amount of available memory
Cache Memory	Amount of cache memory
Buffered Memory	Amount of buffered memory
% Util	% memory utilization on the device
Forecasted % Util Exhaustion	The amount of time before the exhaustion threshold will be breached.

Threshold Summary

Provides information about the devices that are approaching or exceeding their thresholds.

Report filtering

Start Date: 07/17/2017	select	End Date: 07/24/2017	select	Event Class: /Perf	▼
Generate					

Start Date

End Date

The first and last dates of the range of dates to include in the report. To select a date from a calendar, click select. The default range is the week ending with the current date.

Event Class

The device class to use for filtering. The default is /Perf.

To generate or refresh the report, click Generate.

Note: If you export the report by clicking Export all, be sure to format the percentage columns to show percentages instead of decimal values.

Report contents

Column	Content
Device	Name of the device based on the filter parameters selected.
Component	Network interface component name, if applicable.
Event Class	Event class that had a threshold breach.
Count	Number of times a threshold was breached.
Duration	Amount of time that a threshold was breached.
%	Percentage of time that the threshold was breached for the report's time period.

Storage

Clients

This report shows the devices in Resource Manager that use storage devices.

Disk Firmware

This report allows you to select a storage device and display the disk firmware information associated with the device.

Licenses

This report shows the storage devices and installed licenses.

vSphere

Clusters

This report shows all clusters, with the count of virtual machines (total and powered on), hosts, and CPU/memory utilization within each cluster.

Datstores

This report shows all datstores, with the number of connected virtual machines (total and powered on) and the disk space available and consumed on each datstore.

Hosts

This report shows all hosts, with the count of virtual machines (total and powered on), hosts, and CPU/memory reservation and utilization on each host.

LUNs

This report shows all hosts and their logical unit numbers (LUNs), with LUN details.

Resource Pools

This report shows all resource pools, with their CPU and memory allocations and usage.

VM to Datastore

The VM to Datastore report shows the datastores associated with each virtual machine Resource Manager is monitoring.

VMs

This report shows all virtual machines, their operating system, CPU and memory utilization, and the host/cluster in which they reside.

VMware Utilization

This report provides a summary of virtual machines, CPU, memory, and disk utilization over a specified time interval, organized by host.

ADVANCED

The ADVANCED tab features the administration pages described in this section.

- [ADVANCED > Settings](#)
- [ADVANCED > Control Center](#)
- [ADVANCED > Move Devices](#)
- [ADVANCED > Monitoring Templates](#)
- [ADVANCED > Jobs](#)
- [ADVANCED > MIBs](#)
- [ADVANCED > Licensing](#)

You must have Manager or ZenManager privileges to view this tab.

ADVANCED > Settings

The ADVANCED > Settings page is the default page displayed when you click the ADVANCED tab.

You must have Manager or ZenManager privileges to view this page.

Settings

The ADVANCED > Settings > Settings page provides information about Resource Manager and global configuration options.

zenoss SERVICE DYNAMICS DASHBOARD EVENTS INFRASTRUCTURE REPORTS ADVANCED

Settings Control Center Move Devices Monitoring Templates Jobs MIBs Licensing

Settings State at time: 2019/01/03 09:19:52

Commands Instance Identifier Zenoss

Users Hostname https://zenoss5.resmgr-6.

ZenPacks SMTP Host

Versions SMTP Port (usually 25) 25

Events SMTP Username (blank for none)

User Interface SMTP Password (blank for none)

LDAP From Address for Emails

Support Use TLS?

Page Command

Dashboard Production State Threshold 1000

Dashboard Priority Threshold 2

State Conversions Production:1000
Pre-Production:500
Test:400
Maintenance:300
Decommissioned:-1

Priority Conversions Highest:5
High:4
Normal:3
Low:2
Lowest:1
Trivial:0

Authentication You will have to log off for this setting to take effect.
 Cookie-based Authentication
 Session-based Authentication (More Secure)

Enter current password to confirm changes

Save

Commands

The ADVANCED > Settings > Commands page includes tools for managing the commands that can be run from various pages in Resource Manager.

zenoss SERVICE DYNAMICS DASHBOARD EVENTS INFRASTRUCTURE REPORTS ADVANCED

Settings Control Center Move Devices Monitoring Templates Jobs MIBs Licensing

Settings Define Commands

Name	Description	Command
<input type="checkbox"/> forward	Name to IP address lookup	host \${device.id}
<input type="checkbox"/> reverse	IP address to name lookup	host \${device.manageIp}
<input type="checkbox"/> ping	Is the device responding to ping?	\${device.pingCommand} -c 2 \${device.manageIp}
<input type="checkbox"/> snmpwalk	Display the OIDs available on a device	snmpwalk -S \${device.zSnmpVer} -c \${device.zSnmpCommunity} \${device.snmpwalkPrefix} \${device.manageIp} system
<input type="checkbox"/> snmpwalk v3	snmp version v3. Display the OIDs available on a device	snmpwalk -S \${device.zSnmpVer} -d authNoPriv -a \${device.zSnmpAuthType} -x \${device.zSnmpPrivType} -A \${device.zSnmpAuthPassword} -X \${device.zSnmpPrivPassword} -u \${device.zSnmpSecurityName} \${device.snmpwalkPrefix} \${device.manageIp} \${device.zSnmpPort} system
<input type="checkbox"/> traceroute	Show the route to the device	\${device.tracerouteCommand} -q 1 -w 2 \${device.manageIp}

Users

The ADVANCED > Settings > Users page provides tools for managing user accounts. For more information, see [Managing users in Resource Manager](#).

ZenPacks

The ADVANCED > Settings > ZenPacks page displays a list of the installed ZenPacks. For more information, see [Extending Resource Manager with ZenPacks](#).

Versions

The ADVANCED > Settings > Versions page displays the version numbers of key software components.

Software Component Versions	
Zenoss	Zenoss 6.4.0 r307
OS	Linux (x86_64) 3.10.0 (Linux c5979f693c5e 3.10.0-693.11.6.el7.x86_64 #1 SMP Thu Jan 4 01:06:37 UTC 2018 x86_64)
Zope	Zope 2.13.13
Python	Python 2.7.5
Database	MySQL 5.5.60 (5.5.60-MariaDB)
Twisted	Twisted 16.1.1
RabbitMQ	RabbitMQ 3.3.5
Erlang	Erlang 5.10.4
NetSnmp	NetSnmp 5.7.2
Control Center	ControlCenter 1.6.3

Uptimes	
Zope	11 hours 4 min 44 sec

Check For Updates	
Available Zenoss Version	6.2.0 (View Licenses)
Last Checked	2019-01-08 04:30:43
Server Key	9799305e-12fc-11e9-9839-0242ac110010
Check Zenoss Version Now	

Events

The ADVANCED > Settings > Events page displays the event aging properties. For more information, see [Event configuration settings](#).

User Interface

The ADVANCED > Settings > User Interface page displays the user interface configuration settings. For more information, see [User interface configuration](#).

LDAP

The ADVANCED > Settings > LDAP page provides tools for configuratio LDAP authentication. For more information, see [Integrating LDAP authentication](#).

Support

The ADVANCED > Settings > Support page provides tools for creating log files bundles to send to Zenoss Support. For more information, see [Support bundles](#).

Discovery Mapping

The ADVANCED > Settings > Discovery Mapping page provides configuration options for the discovery mapping feature. For more information, see [Discovery mapping](#).

Event configuration settings

The properties on the Event Configuration page (ADVANCED > Settings > Events) configure the event aging feature. These properties define the conditions under which Resource Manager changes the status of events to Aged. After the configured event archive interval, all Closed, Aged, and Cleared events are moved to the event archive.

- Don't Age This Severity and Above - Options are Age All Events, Critical, Error, Warning, Info, Debug, and Clear. By default, this value is set to Error, meaning that all events with a status of Error or Critical are not aged.
- Event Aging Threshold (minutes) - Set the time value, in minutes, that an event must reach before it is aged. By default, this is 240 minutes.
- Event Aging Interval (milliseconds) - The interval when events are scanned to perform autoaging. By default, this is 60000 milliseconds (60 sec).
- Event Aging Limit - The maximum number of events to age in each interval. The limit should be kept relatively low to prevent large database transactions. By default, this is 1000 events.
- Event Archive Threshold (minutes) - Specify the number of minutes since a closed event was last seen before it is moved to the event archive. The minimum value is 1; the maximum value is 43200.
- Event Archive Interval (milliseconds) - The interval when events are scanned for moving to the archive. By default, this is 60000 milliseconds (60 sec).
- Event Archive Limit - The maximum number of events to archive in each interval. The limit should be kept relatively low to prevent large database transactions. By default, this is 1000 events.
- Delete Archived Events Older Than (days) - The number of days that events in the event archive are saved. By default, they are kept in the archive for 90 days. The minimum value is 1 and the maximum value is determined by the range of event archive partitions. With the default configuration, the maximum value is 1000 days.
- Default Syslog Priority - Specify the default severity level assigned to an event coming from zensyslog if no priority can be determined from the event.
- Default Availability Report (days) - Enter the number of days to include in the automatically generated Availability Report. This report shows a graphical summary of availability and status.
- Max Event Size in Bytes - The maximum size of an event that will be processed in bytes. Events that are too large will be logged and dropped. Events that will become too big will have their details overwritten with new details. By default, this is 32768 bytes.
- Summary Index Interval (milliseconds) - The default indexing interval of the event summary in milliseconds. By default, this is 1000 milliseconds (1 sec).
- Archive Index Interval (milliseconds) - The default indexing interval of the event archive in milliseconds. By default, this is 30000 milliseconds (30 sec).
- Index Limit - The number of events to index in each index interval. By default, this is 1000 events.
- Event Time Purge Interval (days) - The number of days that event occurrence time are kept. By default, they are kept for 7 days. The minimum value is 1 and the maximum value is determined by the range of event time partitions. With the default configuration, the maximum value is 7 days.
- Enable Event Flapping Detection - Select this check box if you wish to enable event flapping detection. If an event is created and then cleared flapping_threshold times in event_flapping_interval time then an event of event flapping event class is created.
- Event Flapping Event Class - The event class under which generated flapping events belong.
- Clear Event Heartbeats - Click Clear to clear the event heartbeats.

User interface configuration

You can configure options as follows to control how data loads, how much data is loaded, and filter and search options. To access the options, choose ADVANCED > Settings > User Interface. After setting options, click Save.

- Enable Hyperlinks in Event Fields - Click to enable this option to show hyperlinks in event fields. By default, this feature is disabled.
- Enable Infinite Grids for Events - This enables scrolling through multiple pages of events. Disabling may provide a performance increase in larger environments. By default, this feature is enabled.
- Enable Infinite Grids for Components - This enables scrolling through multiple pages of components. Disabling may provide a performance increase in larger environments. By default, this feature is enabled.
- Enable Live Filters - Disable this option to turn off the live filters feature. If you disable this feature, you will need to press enter on all search and filter inputs.
- Enable Incremental Tree Loading on the Infrastructure Page - Enable this option to load the infrastructure tree one node at a time. If disabled (the default), then the infrastructure tree is loaded all at once. You might enable this option if you have a complex hierarchy of organizers and device classes and want to improve your UI load response time.
- Show Tree Event Severity Icons - Disabling this option may speed up page loading.
- Load Infrastructure Page's Device Tree from Catalog - Enable this option to load the device tree directly from the catalog.
- Enable Tree Filters - If disabled, then tree filters (the text input area that allows you to filter the information displayed) are hidden on all pages. By default, this option is enabled.
- Show Page Statistics - Enable this option to show debugging information. By default, this option is not enabled.
- Default Time Range - Specify the default time range for all graphs. The default amount is Last Hour. Other values are Last 24 Hours, Last Week, Last 30 Days, and Last Year.
- Number of Graph Columns - The number of graph columns to display on graphs in the Device Overview page. The default value is Auto, which means that the number of columns increases as the browser's width increases. You can change the setting to the number of columns (1, 2, or 3).
- Device Grid Buffer Size - Specify the number of device data rows to fetch from the server for each buffer request. The default buffer size is 100 rows.
- Component Grid Buffer Size - Specify the number of component data rows to fetch from the server for each buffer request. The default buffer size is 25 rows.
- Event Console Buffer Size - Specify the number of event rows to fetch from the server for each buffer request. The default buffer size is 200 rows.
- Device Move Job Threshold - Specify the limit at which devices are moved immediately. If the number of devices to be moved exceeds this threshold, then the move occurs in a job; otherwise, they are moved immediately. The default value is 5 devices.
- Job Notification Refresh Interval - Specify the refresh interval, in seconds, for the job notification dialog box. The default time is 10 seconds.
- Job Grid Buffer Size - Specify the number of job data rows to fetch from the server for each buffer request. The default buffer size is 100 rows.

ADVANCED > Control Center

The ADVANCED > Control Center page enables Control Center administration from the Resource Manager browser interface.

You must have Manager or ZenManager privileges to view this page.

The screenshot shows the Zenoss Service Dynamics Control Center interface. The top navigation bar includes 'zenoss SERVICE DYNAMICS', 'DASHBOARD', 'EVENTS', 'INFRASTRUCTURE', 'REPORTS', and 'ADVANCED'. The 'ADVANCED' tab is active, and the 'Control Center' sub-tab is selected. Below the navigation bar, there are buttons for 'Start', 'Stop', and 'Restart', along with a search bar and a 'Refresh' button. The main area displays a table of services with columns for Name, Type, Uptime, AutoStart, Restart, State, and Host. The 'Capacity' service is selected, and its details are shown in a pane below the table.

Name	Type	Uptime	AutoStart	Restart	State	Host
Capacity	daemon	11:14:15.701...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...
CentralQuery	daemon	11:21:01.642...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...
HMaster	daemon	11:24:43.523...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...
localhost	hub				...	
localhost						
mariadb-events	daemon	11:23:05.455...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...
mariadb-model	daemon	11:23:05.466...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...
memcached	daemon	11:23:34.545...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...
memcached-session	daemon	11:23:35.556...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...
MetricConsumer	daemon	11:21:01.631...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...
MetricShipper	daemon	11:20:47.609...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...
RabbitMQ	daemon	11:23:35.445...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...
reader	daemon	11:21:05.489...	<input checked="" type="checkbox"/>		Up	resmgr-64x-ho...

Display: Details

ID:
aswduk4kmhlcystu8hmrvgla
Description:
Aggregates, projects, and indexes capacity data.
AutoStart:
true
State:
RUNNING
Text:
Capacity
 Restarting
Uptime:
11:04:17.058371

Save Cancel

For more information about administering Control Center, see [Administering Control Center](#).

ADVANCED > Move Devices

The ADVANCED > Move Devices page enables moving all of the devices assigned to one collector to another collector.

You must have Manager or ZenManager privileges to view this page.

To move devices:

1. Navigate to ADVANCED > Move Devices.
2. In the SOURCE COLLECTOR area, select the collector that contains the devices to move.
3. In the MOVE DEVICE(S) TO NEW COLLECTOR area, select the target collector.
4. Click Submit.

ADVANCED > Monitoring Templates

The ADVANCED > Monitoring Templates page displays all of the monitoring templates that are available in Resource Manager. Monitoring templates determine how the system collects performance data for devices and device components.

You must have Manager or ZenManager privileges to view this page.

For more information, see [About monitoring templates](#).

ADVANCED > Jobs

The ADVANCED > Jobs page displays a list of the current and past background processing jobs of Resource Manager.

You must have Manager or ZenManager privileges to view this page.

For more information, see [Managing background tasks](#).

ADVANCED > MIBs

The ADVANCED > MIBs page displays all of the SNMP [management information bases](#) (MIBs) that are installed in Resource Manager.

You must have Manager or ZenManager privileges to view this page.

To add a MIB to Resource Manager, see [Administering MIB files](#).

ADVANCED > Licensing

The ADVANCED > Licensing page displays your Resource Manager license report.

You must have Manager or ZenManager privileges to view this page.

The screenshot shows the Zenoss web interface for the 'Licensing' page. The top navigation bar includes 'DASHBOARD', 'EVENTS', 'INFRASTRUCTURE', 'REPORTS', and 'ADVANCED'. The 'Licensing' page is active, displaying the Zenoss logo and 'Own IT.' tagline. The page content is organized into several sections:

- Instance Details:** Server ID: 17199295e-126c-11e9-8039-0242ac110010; Metrics Updated: 2019-01-08T04:28:12:164306 (UTC/GMT) - reported from local instance.
- License Details:** Report Generated: 2019-01-07 22:39:43 288741 (EST); No license associated with this Zenoss instance's server key. Send this report to cam@zenoss.com.
- Software Metrics:** Zenoss: Zenoss 6.4.0
- Monitored Infrastructure:** A list of counts for various components:
 - Device Count: 1
 - GuestVM Count: 0
 - Linked Guest Count: 0
 - Locations: 0
 - Systems: 0
 - Groups: 0
 - CPU Cores: 0
 - Hubs: 1
 - Decommissioned Devices: 0
 - Templates: 507
 - Reports: 52
 - User Count: 2
 - Collectors: 1
 - Event Count: 3
 - Evt Mappings: 1012
- MR Calculation BreakDown:** MR Calculation: 1 * 1 (Dev Count) + 0 (GuestVM Count) + 0 (Linked Guests) + 0 (Decom Devices)

Using the Appliance Administration menu

This section describes the curses-based Appliance Administration menu, a text user interface (TUI).

- [Configure Network and DNS](#)
- [Configure IPv6 Network CIDR](#)
- [Configure Timezone](#)
- [Change Docker Hub Credentials](#)
- [Change Root Password](#)
- [Change ccuser Password](#)
- [Update System](#)
- [Change SSL settings](#)
- [Root Shell](#)
- [Reboot / Poweroff System](#)

Configure Network and DNS

The Configure Network and DNS option invokes `nmtui`, the `NetworkManager` text user interface (TUI) tool. The `nmtui` utility provides submenus for editing and activating network connections, and for changing the hostname.

Zenoss recommends using only the Configure Network and DNS option to change connection properties or the hostname, and always rebooting after making changes.

- [Editing a connection to configure static IPv4 addressing](#)
- [Edit a connection \(Docker virtual bridge\)](#)
- [Activate a connection](#)
- [Setting the system hostname](#)


Editing a connection to configure static IPv4 addressing

The default configuration for network connections is DHCP. To configure static IPv4 addressing, perform this procedure. To navigate in the text user interface (TUI):


- To move forward or backward through options, press the arrow keys.
- To display a menu or choose an option, press Enter.

Follow these steps:


1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Log in as the `root` user.



3. In the Appliance Administration menu, select Configure Network and DNS, and then press Enter. This step starts the NetworkManager TUI.

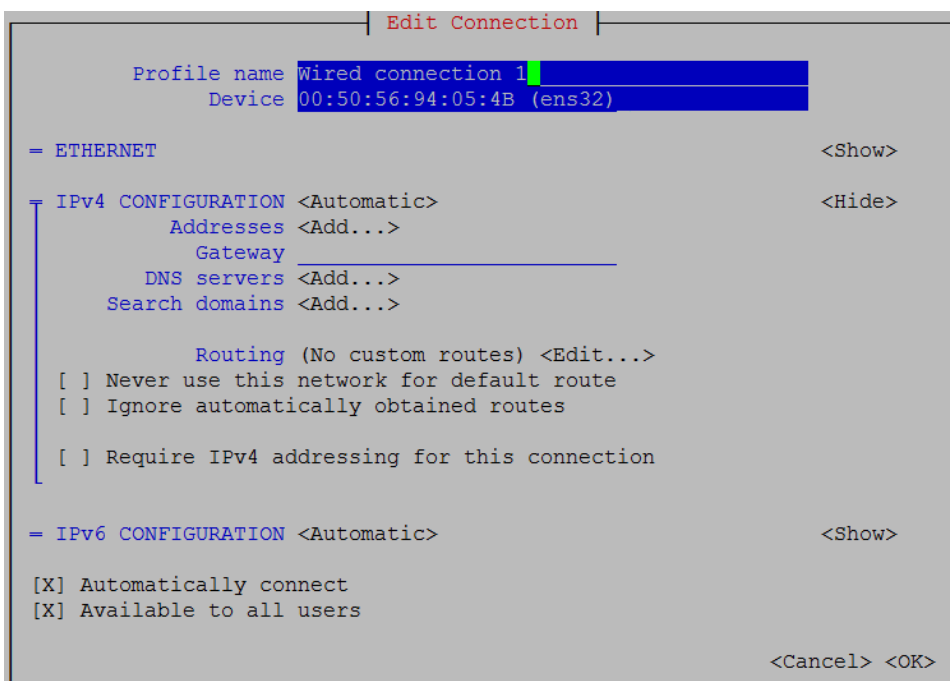


4. On the NetworkManager TUI menu, select Edit a connection, and then press Enter. The TUI displays the connections that are available on the host.




Note: Do not use this procedure to modify the `docker0` connection.

5. Select the virtual connection, and then press Enter.



6. Optional: If the IPv4 CONFIGURATION area is not visible, select its display option (<Show>), and then press Enter.
7. In the IPv4 CONFIGURATION area, select <Automatic>, and then press Enter.




8. Configure static IPv4 networking as follows:
 - a. Select Manual, and then press Enter.
 - b. Beside Addresses, select <Add>, and then press Enter.
 - c. In the Addresses field, enter an IPv4 address for the virtual machine, and then press Enter.
 - d. Repeat the preceding two steps for the Gateway and DNS servers fields.
9. Tab to the bottom of the Edit Connection screen to select OK, and then press Enter.
10. On the NetworkManager TUI screen, select Quit, and then press Enter.
This step returns control to the Appliance Administration menu.
11. Reboot the operating system:
 - a. In the Appliance Administration menu, select Reboot / Poweroff System.
 - b. Select Reboot.
 - c. Select OK, and then press Enter.


Edit a connection (Docker virtual bridge)

The default IP address space of the Docker virtual bridge is 172.17.0.1/16. To configure a different address space, perform this procedure.

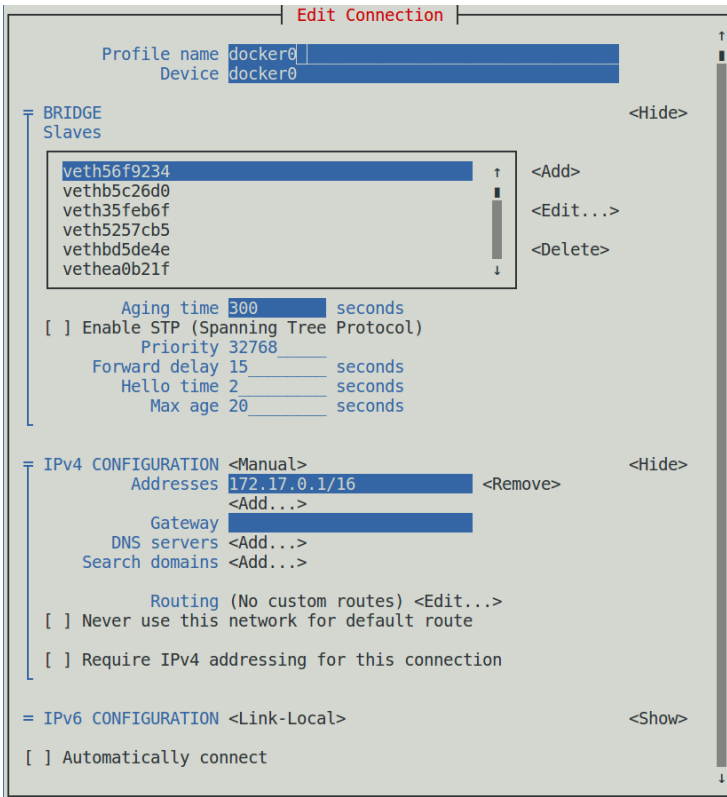
1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Log in as the root user.
3. In the Appliance Administration menu, select Configure Network and DNS, and then press Enter. This step starts the NetworkManager TUI.



4. On the NetworkManager TUI menu, select Edit a connection, and then press Enter. The TUI displays the connections that are available on this host.



5. Use the down-arrow key to select docker0, and then press Enter.



Use the Tab key and the arrow keys to navigate among options in the Edit Connection screen, and use Enter to toggle an option or to display a menu of options.

6. If the BRIDGE area is visible, select its display option (<Hide>), and then press Enter.

Do not edit any of the entries in the BRIDGE area.

7. If the IPv4 CONFIGURATION area is not visible, select its display option (<Show>), and then press Enter.
8. In the IPv4 CONFIGURATION area, navigate to Addresses, and then enter a new IPv4 address in CIDR notation.

```

Edit Connection
-----
Profile name: docker0
Device:      docker0

= BRIDGE <Show>
= IPv4 CONFIGURATION <Manual> <Hide>
  Addresses: 172.17.0.1/16 <Remove>
             <Add...>
  Gateway:   <Add...>
  DNS servers: <Add...>
  Search domains: <Add...>

  Routing (No custom routes) <Edit...>
  [ ] Never use this network for default route
  [ ] Require IPv4 addressing for this connection

= IPv6 CONFIGURATION <Link-Local> <Show>
  [ ] Automatically connect
  [X] Available to all users

<Cancel> <OK>

```

9. Use Tab or the Down Arrow key to select the <OK> option at the bottom of the Edit Connection screen, and then press Enter.
10. In the available connections screen, use Tab to select the <Quit> option, and then press Enter.
11. Reboot the operating system:
 - a. In the Appliance Administration menu, select Reboot / Poweroff System.
 - b. Press Tab to select OK, and then press Enter.

CIDR prefix lengths for common subnet masks


Subnet mask	CIDR prefix length
255.255.0.0	/16
255.255.128.0	/17
255.255.192.0	/18
255.255.224.0	/19
255.255.240.0	/20
255.255.248.0	/21
255.255.252.0	/22
255.255.254.0	/23
255.255.255.0	/24
255.255.255.128	/25
255.255.255.192	/26
255.255.255.224	/27
255.255.255.240	/28
255.255.255.248	/29

Activate a connection

The Activate a connection submenu provides options for activating and deactivating network connections.

Do not deactivate the `docker0` connection.

On selection, the Activate a connection submenu displays the available connections. The asterisk character (*) at the beginning of a connection name indicates that the connection is active.




Use the arrow keys to select a connection, and then use Tab to navigate the options at the right side of the list. Use Enter to choose an option.

Always reboot after activating or deactivating a connection.


Setting the system hostname

The default hostname is `zsd-master` for the master host and is `zsd-delegate` for delegate hosts. To change the default hostname, perform this procedure.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. In the Appliance Administration menu, select Configure Network and DNS, and then press Enter.
This step starts the NetworkManager TUI.



3. Display the hostname entry field.
 - a. In the NetworkManager TUI menu, select Set system hostname.
 - b. Select OK, and then press Enter.
4. In the Hostname field, enter the hostname or a fully-qualified domain name.
5. Press Tab twice to select OK, and then press Enter.




6. In the confirmation dialog box, press Enter.
7. On the NetworkManager TUI screen, select Quit, and then press Enter.
This step returns control to the Appliance Administration menu.
8. Reboot the operating system as follows:
 - a. In the Appliance Administration menu, select Reboot / Poweroff System.
 - b. Select Reboot.
 - c. Select OK, and then press Enter.

Configure IPv6 Network CIDR

The version of Docker included in the Resource Manager virtual appliance needs to know at startup the address prefix of the IPv6 network it will use. To enable monitoring of devices that use IPv6, perform this procedure on the Control Center master host, and all delegate hosts.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Log in as the root user.
3. In the Appliance Administration menu, select the Configure IPv6 Network CIDR option.




4. In the IPv6 CIDR screen, enter the address prefix of your IPv6 network in the CIDR field.
5. Use Tab to select the Ok button, and then press Enter.
The Docker daemon restarts, and the Appliance Administration disappears briefly before returning. This is normal.

Configure Timezone

The default timezone of the Resource Manager virtual appliance is UTC. This procedure changes the timezone setting of a single host. All hosts in a multi-host deployment must use the same timezone.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as PuTTY.
2. Log in as the root user.
3. In the Appliance Administration menu, select the Configure Timezone option.



4. Use the Down Arrow key to select the desired timezone.
5. Press Tab to highlight Select, and then press Enter.

Always reboot after changing the timezone.

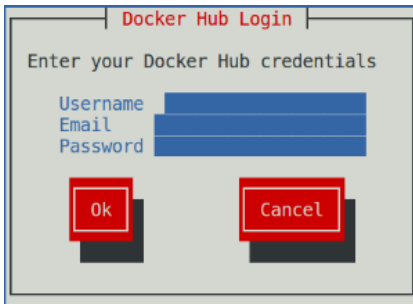
Change Docker Hub Credentials

As of release 1.4.0, Control Center does not pull Resource Manager images over the internet. This procedure should only be used on releases earlier than 1.4.0.

To perform this procedure, you need the username, email address, and password of a Docker Hub user account that is authorized to pull Resource Manager images. The information is provided by Zenoss Support when necessary.

Control Center master host can obtain software updates over the internet from the Docker Hub registry. This option obtains a JSON web token from Docker Hub and stores it at `/root/.docker/config.json`.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as PuTTY.
2. Log in as the root user.
3. In the Appliance Administration menu, select the Change Docker Hub Credentials option.



The image shows a terminal window titled "Docker Hub Login". Inside the window, the text "Enter your Docker Hub credentials" is displayed. Below this text are three input fields labeled "Username", "Email", and "Password". Each field has a blue cursor. At the bottom of the window, there are two red buttons: "Ok" and "Cancel".

4. Enter the username, email address, and password of the authorized user account in the fields. Use Tab to advance to the next field.
5. Press Tab to highlight OK, and then press Enter.

Change Root Password

This option invokes the `passwd` command to change the password of the root user account.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as PuTTY.
2. Log in as the root user.
3. In the Appliance Administration menu, select the Change Root Password option.
The Appliance Administration menu disappears, and the system prompts for a new password:

```
Changing password for user root.  
New password:
```

Passwords must include a minimum of eight characters, with at least one character from three of the following character classes: uppercase letter, lowercase letter, digit, and special.

4. Enter a new password, and then press Enter.
5. Enter the password again, and then press Enter.
The Appliance Administration menu reappears.

Change ccuser Password

This option invokes the `passwd` command to change the password of the `ccuser` user account.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as PuTTY.
2. Log in as the root user.
3. In the Appliance Administration menu, select the Change Root Password option.
The Appliance Administration menu disappears, and the system prompts for a new password:

```
Changing password for user ccuser.  
New password:
```

Passwords must include a minimum of eight characters, with at least one character from three of the following character classes: uppercase letter, lowercase letter, digit, and special.

4. Enter a new password, and then press Enter.
5. Enter the password again, and then press Enter.
The Appliance Administration menu reappears.

Update System


This option updates the Control Center and Resource Manager software on a host. For more information, see [Updating Resource Manager](#).

Change SSL settings


To perform this procedure, you need to be able to display the contents of the SSL certificate and key files that you want to install on the Control Center master host, and you need a copy of the root certificate file (rootCA.pem). In addition, Zenoss recommends logging in to the master host through SSH, rather than the hypervisor console, so that you can copy and paste content.

This option allows you to provide new content for SSL certificate and key files.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as PuTTY.
2. Log in as the root user.
3. Use the Down Arrow key to select Change SSL settings, and then press Enter.



4. When you are ready to add the contents of your SSL certificate and key files to the Control Center master host, press Enter.




5. Press Enter.
The Appliance Administration menu is replaced with the nano text editor.
6. Enter the contents of your SSL key file, and then save the file and exit the editor.
 - a. Press Ctrl-O.
 - b. Press Ctrl-X.
 - c. Press y, and then press Enter.



7. Press Enter.
The Appliance Administration menu is replaced with the nano text editor.
8. Enter the contents of your SSL certificate file, and then save the file and exit the editor.
 - a. Press Ctrl-O.
 - b. Press Ctrl-X.

c. Press `y`, and then press Enter.



9. Restart the Control Center daemon (`serviced`) now or later.

Restarting `serviced` pauses Resource Manager services briefly.

- To restart `serviced` now, press Enter.
- To restart `serviced` later, press Tab to select No, and then press Enter.

10. Install the root certificate into browser clients.

The procedures for installing a root certificate into a browser client varies by browser and client operating system. For more information, refer to your browser documentation or articles such as [this one](#).

Root Shell

This option starts a command-line session as the root user.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as [PuTTY](#).
2. Log in as the root user.
3. Use the Down Arrow key to select Root Shell, and then press Enter.
The menu is replaced by a command prompt similar to the following example:

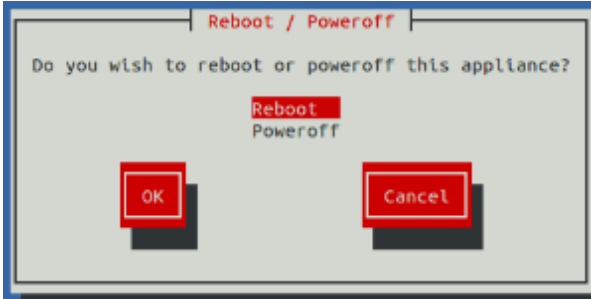
```
[root@resmgr-master ~]#
```

To return to the Appliance Administration menu, enter the `exit` command.

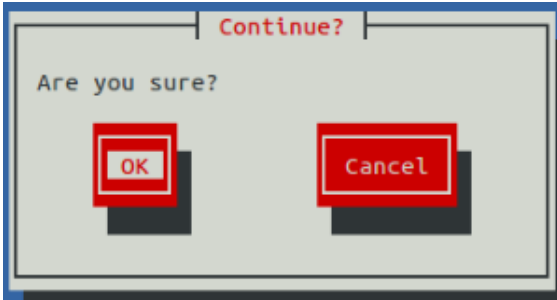
Reboot / Poweroff System

This option reboots or shuts down and powers off a Control Center host.

1. Gain access to the Control Center host, through the console interface of your hypervisor, or through a remote shell utility such as PuTTY.
2. Log in as the root user.
3. In the Appliance Administration menu, select the Reboot / Poweroff System option.



4. Use the Down Arrow key to select Reboot or Poweroff System.
5. Press Tab to highlight OK, and then press Enter.



6. Use Tab to select OK or Cancel, and then press Enter.

The system reboots or shuts down and powers off.

Managing Zope instances

Zope is the web application server that Resource Manager uses.

For general processing, the Resource Manager zope service provides six instances of the Zope web application server. Resource Manager browser interface performance might slow in deployments with many concurrent users, during frequent report generation, or with Zenoss JSON API use. To maintain responsiveness, Resource Manager services provide dedicated Zope instances for reporting, Zenoss JSON API use, and debugging.

You can enable or disable the services and change the number of Zope instances that Resource Manager uses. (If you increase the number of instances, monitor RAM usage.)

- [Changing the number of Zope instances](#)
- [Dedicated Zope service for reporting](#)
- [Dedicated Zope service for Zenoss JSON API use](#)
- [Dedicated Zope service for debugging](#)

Changing the number of Zope instances

If you need more or fewer dedicated or general-purpose Zope instances, edit the appropriate service to change the number. Debugging (the zendebug service) is limited to one instance.

1. Log in to the Control Center browser interface.
2. In the Applications table under Application, click the application name (Zenoss.resmgr).
3. In the Services table, click Zenoss.
4. In the Services table, click User Interface.
5. In the Services table, click the appropriate service as follows:
 - zenapi
 - zenreports
 - zope
6. On the service page, click Edit Service.
7. In the Edit Service dialog box, change the value in the Instances field.
8. Click Save Changes.
9. On the service page, click Restart.

Dedicated Zope service for reporting

The zenreports service provides a Zope instance that is dedicated to report generation.

No configuration or action is required to use the zenreports service. By default, the service starts when you launch Resource Manager and is automatically selected to generate Resource Manager reports.

If you need more than one dedicated instance for report generation, edit the zenreports service to increase the number, as described in [Changing the number of Zope instances](#).

Under the following conditions, you can disable the zenreports service:

- You have a small deployment or few concurrent users of the browser interface.
- You infrequently generate Resource Manager reports.

The following topics provide instructions:

- [Disabling automatic start of the reporting Zope service](#)
- [Stopping the reporting Zope service](#)

Disabling automatic start of the reporting Zope service

By default, the zenreports service starts when you launch Resource Manager. You can disable automatic start of the service.

1. Log in to the Control Center master host as a user with Control Center CLI privileges.
2. Disable automatic start of the zenreports service.
 - a. Edit the configuration of the service.

```
serviced service edit zenreports
```

The command opens the service's configuration file in the default text editor.

- b. Locate the line "Launch": "auto", and replace "auto" with "manual" as follows:

```
"Launch": "manual"
```

- c. Save the file, and then exit the text editor.

Stopping the reporting Zope service

You can disable use of the dedicated Zope service so that Resource Manager uses general-purpose Zope instances for report generation.

1. Log in to the Control Center browser interface.
2. In the Applications table under Application, click the application name (Zenoss.resmgr).
3. Turn off rules that route report requests to the zenreports service:
 - a. In the Configuration Files table, locate file zproxy-nginx.conf, and under the Actions column, click Edit.
 - b. In the configuration file, locate the following line:

```
include zopereports-proxy.conf
```

- c. Add the number sign character (#) to the beginning of the line.

```
#include zopereports-proxy.conf
```

- d. Click Save.
 4. In the Instances table under the Actions column, click Restart. This action restarts the zproxy service.
 5. Stop the zenreports service:
 - a. In the Services table, click Zenoss.
 - b. In the Services table, click User Interface.
 - c. In the Services table, locate the zenreports service, and under the Actions column, click Stop.

Dedicated Zope service for Zenoss JSON API use

The zenapi service provides a Zope instance that is intended for requests sent through the Zenoss JSON API.

The Zenoss JSON API enables advanced users to customize Resource Manager and automate tasks such as adding and removing devices and managing events.

By default, the zenapi service starts when you launch Resource Manager and provides one Zope instance. If you need more than one dedicated instance, edit the zenapi service to increase the number, as described in [Changing the number of Zope instances](#).

Under the following conditions, you can disable the service:

- You have a small deployment or few concurrent users of the browser interface.
- You infrequently use the Zenoss JSON API.

The following topics provide instructions:

- [Disabling automatic start of the Zenoss JSON API Zope service](#)
- [Stopping the Zenoss JSON API Zope service](#)

Disabling automatic start of the Zenoss JSON API Zope service

By default, the zenapi service starts when you launch Resource Manager. You can disable automatic start of the service.

1. Log in to the Control Center master host as a user with Control Center CLI privileges.
2. Disable automatic start of the zenapi service.
 - a. Edit the configuration of the service.

```
serviced service edit zenapi
```

The command opens the service's configuration file in the default text editor.

- b. Locate the line "Launch": "auto", and replace "auto" with "manual" as follows:

```
"Launch": "manual"
```

- c. Save the file, and then exit the text editor.

Stopping the Zenoss JSON API Zope service

You can disable use of the dedicated Zope service so that Resource Manager uses the general-purpose Zope instances for Zenoss JSON API processing.

1. Log in to the Control Center browser interface.
2. In the Applications table under Application, click the application name (Zenoss.resmgr).
3. In the Public Endpoints table, locate URL <https://zenapi.hostname>, and under the Actions column, click Stop.
4. Turn off rules that route Zenoss JSON API requests to the zenapi service:
 - a. In the Configuration Files table, locate file zproxy-nginx.conf, and under the Actions column, click Edit.
 - b. In the configuration file, locate the following line:

```
~*zenapi apizopes;
```

- c. Add the number sign character (#) to the beginning of the line.

```
#~*zenapi apizopes;
```

- d. Click Save.
5. In the Instances table under the Actions column, click Restart. This action restarts the zproxy service.
6. Stop the zenapi service:
 - a. In the Services table, click Zenoss.
 - b. In the Services table, click User Interface.
 - c. In the Services table, locate the zenapi service, and under the Actions column, click Stop.

Dedicated Zope service for debugging

The zendebg service enables the Zenoss Support team to start a separate Zope instance to use for troubleshooting and debugging purposes.

By default, zendebg is stopped. You can enable the service when instructed to do so by Zenoss Support. Do not remove zendebg from the Resource Manager service definition.

- [Enabling the Zope service for debugging](#)

Enabling the Zope service for debugging

If instructed to do so by Zenoss Support, enable use of the dedicated Zope service for debugging. Resource Manager then uses all general-purpose Zope instances for other processing. The maximum number of debugging service instances is one.

1. Log in to the Control Center browser interface.
2. In the Applications table under Application, click the application name (Zenoss.resmgr).
3. In the Public Endpoints table, locate URL <https://zendebug.hostname>, and under the Actions column, click Start.
4. Start the zendebug service:
 - a. In the Services table, click Zenoss.
 - b. In the Services table, click User Interface.
 - c. In the Services table, locate the zendebug service, and under the Actions column, click Start.

SNMP device preparation

This section provides information about SNMP support and lists Net-SNMP configuration settings that are required by the system.

- [Net-SNMP](#)
- [SNMP v3 support](#)
- [Community information](#)
- [System contact information](#)
- [Extra information](#)

Net-SNMP

By default, Net-SNMP does not publish the full SNMP tree. Check to see if that is currently the case on a device and configure it correctly.

1. Confirm snmpd is running:

```
> snmpwalk -v 2c -cpublic <your device name> system
```

2. Retrieve the IP table for the device with snmpwalk:

```
> snmpwalk -v 2c -cpublic <your device name> ip
```

Typical SNMP View:

```
view systemview included .1 view systemview included .1.3.6.1.2.1.25.1 access notConfigGroup "" any  
noauth exact systemview none none
```

SNMP v3 support

Resource Manager provides support for SNMP v3 data collection.

The following configuration properties control the authentication and privacy of these requests:

- **zSnmpAuthType** - Use "MD5" or "SHA" signatures to authenticate SNMP requests. If only zSnmpAuthType and zSnmpAuthPassword are set, then the message is sent with authentication but no privacy.
- **zSnmpAuthPassword** - Shared private key used for authentication. Must be at least 8 characters long.
- **zSnmpPrivType** - "DES" or "AES" cryptographic algorithms. If zSnmpPrivType and zSnmpPrivPassword are set, then the message is sent with privacy and authentication. You cannot set a PrivType and PrivPassword without also setting an AuthType and AuthPassword. If neither Priv nor Auth values are set, then the message is sent with no authentication or privacy.
- **zSnmpPrivKey** - Shared private key used for encrypting SNMP requests. Must be at least 8 characters long.
- **zSnmpSecurityName** - Security Name (user) to use when making SNMPv3 requests.

If monitoring SNMPv3 devices, make sure that msgAuthoritativeEngineID (also known as snmpEngineID or Engine ID) is not shared by two devices. It must be unique for each device.

Advanced Encryption Standard

SNMPv3 encryption using the Advanced Encryption Standard (AES) algorithm is supported only if the host platform net-snmp library supports it.

You can determine whether your platform supports AES by using the following test:

```
$ snmpwalk -x AES 2>&1 | head -1
```

If the response is:

```
"Invalid privacy protocol specified after -x flag: AES"
```

then your platform does not support AES encryption for SNMPv3.

If the response is:

```
"No hostname specified."
```

Then your platform supports AES.

Community information

Add these lines to your snmp.conf file.

This line will map the community name "public" into a "security name":

```
# sec.name source community
```

```
com2sec notConfigUser default public
```

This line will map the security name into a group name:

```
# groupName securityModel securityName
```

```
group notConfigGroup v2c notConfigUser
```

This line will create a view for you to let the group have rights to:

```
# Make at least snmpwalk -v 1 localhost -c public system fast again.
```

```
# name incl/excl subtree mask(optional)
```

```
view systemview included .1
```

This line will grant the group read-only access to the systemview view.

```
# group context sec.model sec.level prefix read write notif access notConfigGroup "" any noauth exact  
systemview none none
```

System contact information

It is also possible to set the sysContact and sysLocation system variables through the snmpd.conf file:

```
syslocation Unknown (edit /etc/snmp/snmpd.conf)
```

```
syscontact Root <root@localhost> (configure /etc/snmp/snmp.local.conf)
```

```
# Added for support of bcm5820 cards. pass .1 /usr/bin/ucd5820stat
```

Extra information

For more information, see the `snmpd.conf` manual page, and the output of the `snmpd -H` command.

```
trapcommunity public
```

```
trapsink default
```


Syslog device preparation

- [Forwarding syslog messages from UNIX/Linux devices](#)
- [Forwarding syslog messages from a Cisco IOS router](#)
- [Forwarding syslog messages from a Cisco CatOS switch](#)
- [Forwarding syslog messages using syslog-ng](#)

Forwarding syslog messages from UNIX/Linux devices

Resource Manager has its own syslog server (zensyslog). Managed devices should point their syslog daemons to the system.

To do this, edit the `/etc/rsyslog.conf` file and add an entry, where 1.2.3.4 is the zensyslog IP:

1. Log in to the target device as a super user.
2. Open the `/etc/rsyslog.conf` file with a text editor (such as vi).
3. Enter `*.debug`, and then press the Tab key.
4. Enter the host name or IP address of the server. For example:

```
*.debug @192.168.X.X
```

5. Save the file and exit the file editor program.
6. Restart the Syslog service using the command below:

```
/etc/init.d/syslog restart
```

Forwarding syslog messages from a Cisco IOS router

Here are some links to Cisco commands to turn on syslog. Typically, it is easier to use syslog than SNMP traps from network devices. The most basic IOS command to send syslog messages is:

```
logging 1.2.3.4
```

Other Cisco syslog configurations

Following are additional configurations for other Cisco devices. To set up these configurations:

1. Log in to the target router.
2. Type the command enable at the prompt.
3. Once you are prompted for a password, enter the correct password.
4. Type the command config at the prompt.
5. Type the command terminal at the configuration prompt.
6. At the prompt, set the Syslog forwarding mechanism. See example below:

```
logging <IP address of the server>
```

7. Exit out all the prompts to the main router prompt.

Catalyst

```
set logging server enable set logging server 192.168.1.100 set logging level all 5 set logging server severity 6
```

Local Director

```
syslog output 20.5 no syslog console syslog host 192.168.1.100
```

PIX Firewalls

```
logging on logging standby logging timestamp logging trap notifications logging facility 19 logging host inside 192.168.1.100
```

Forwarding syslog messages from a Cisco CatOS switch

To forward a syslog message from a Cisco CatOS switch:

1. Log in to the target switch.
2. Type the command enable at the prompt.
3. Enter the password when prompted.
4. Set the Syslog forwarding mechanism; for example:

```
set logging server <IP address of the server>
```

5. You can set the types of logging information that you want the switch to provide with the commands below as examples:

```
set logging level mgmt 7 default set logging level sys 7 default set logging level filesys 7 default
```

Forwarding syslog messages using syslog-ng

Here is an example for FreeBSD and Linux platforms.

1. Log in to the target device as a super user.
2. Open `/etc/syslog-ng/syslog-ng.conf` file with a text editor.
3. Add source information to file. See the following examples:
FreeBSD:

```
source src { unix-dgram("/var/run/log"); internal ();};
```

Linux: (will gather both system and kernel logs)

```
source src { internal(); unix-stream("/dev/log" keep-alive(yes) max-connections(100)); pipe("/proc  
/kmsg"); udp();};
```

4. Add destination information (in this case, the server). For example:

```
log { source(src); destination(zenoss);};
```

TALES expressions

Use TALES syntax to retrieve values and call methods on Resource Manager objects. Several areas accept TALES syntax; these include:

- Command templates
- User commands
- Notifications
- zLinks

Commands (those associated with devices and those associated with events) can use TALES expressions to incorporate data from the related devices or events. TALES is a syntax for specifying expressions that let you access the attributes of certain objects, such as a device or an event.

For additional documentation on TALES syntax, see the TALES section of the [Zope Page Templates Reference](#).

Depending on context, you may have access to a device, an event, or both. Following is a list of the attributes and methods you may want to use on device and event objects. The syntax for accessing device attributes and methods is `${dev/attributename}`. For example, to get the `manageIp` of a device you would use `${dev/manageIp}`. For events, the syntax is `${evt/attributename}`.

A command to ping a device might look like this. (The `${ . . }` is a TALES expression to get the `manageIp` value for the device.)

```
ping -c 10 ${device/manageIp}
```

TALES expression examples

DNS Forward Lookup (assumes device/id is a resolvable name)

```
host ${device/id}
```

DNS Reverse Lookup

```
host ${device/manageIp}
```

SNMP Walk

```
snmpwalk -v 2c -c${device/zSnmpCommunity} ${device/manageIp} system
```

To use these expressions effectively, you must know which objects, attributes, and methods are available, and in which contexts. Usually there is a device that allows you to access the device in a particular context. Contexts related to a particular event usually have event defined.

TALES device attributes

The following table lists available device attributes.

Attribute	Description
getId	The primary means of identifying a device within the system
getManagementIp	The IP address used to contact the device in most situations
productionState	The production status of the device: Production, Pre-Production, Test, Maintenance or Decommissioned. This attribute is a numeric value, use getProductionStateString for a textual representation.
getProductionStateString	Returns a textual representation of the productionState
snmpAgent	The agent returned from SNMP collection
snmpDescr	The description returned by the SNMP agent
snmpOid	The oid returned by the SNMP agent
snmpContact	The contact returned by the SNMP agent
snmpSysName	The system name returned by the SNMP agent
snmpLocation	The location returned by the SNMP agent
snmpLastCollection	When SNMP collection was last performed on the device. This is a DateTime object.
getSnmpLastCollectionString	Textual representation of snmpLastCollection
rackSlot	The slot name/number in the rack where this physical device is installed
comments	User entered comments regarding the device
priority	A numeric value: 0 (Trivial), 1 (Lowest), 2 (Low), 3 (Normal), 4 (High), 5 (Highest)
getPriorityString	A textual representation of the priority
getHWManufacturerName	Name of the manufacturer of this hardware
getHWProductName	Name of this physical product
getHWProductKey	Used to associate this device with a hardware product class
getOSManufacturerName	Name of the manufacturer of this device's operating system.
getOSProductName	Name of the operating system running on this device.
getOSProductKey	Used to associate the operating system with a software product class
getHWSerialNumber	Serial number for this physical device
getLocationName	Name of the Location assigned to this device
getLocationLink	Link to the system page for the assigned Location
getSystemNames	A list of names of the Systems this device is associated with
getDeviceGroupNames	A list of names of the Groups this device is associated with
getLastChangeString	When the last change was made to this device
getLastPollSnmpUptime	Uptime returned from SNMP
uptimeStr	Textual representation of the SNMP uptime for this device
getPingStatusString	Textual representation of the ping status of the device
getSnmpStatusString	Textual representation of the SNMP status of the device

TALES event attributes

The following table lists available event attributes.

Attribute	Description
agent	Collector name from which the event came (such as zensyslog or zentrap).
component	Component of the associated device, if applicable. (Examples: eth0, httpd.)
count	Number of times this event has been seen.
dedupid	Key used to correlate duplicate events. By default, this is: device, component, eventClass, eventKey, severity.
device	ID of the associated device, if applicable.
DeviceClass	Device class from device context.
DeviceGroups	Device systems from device context, separated by .
eventClass	Event class associated with this device. If not specified, may be added by the rule process. If this fails, then will be /Unknown.
eventClassKey	Key by which rules processing begins. Often equal to component.
eventGroup	Logical group of event source (such as syslog, ping, or nteventlog).
eventKey	Primary criteria for mapping events into event classes. Use if a component needs further de-duplication specification.
eventState	State of event. 0 = new, 1 = acknowledged, 2 = suppressed.
evid	Unique ID for the event.
facility	syslog facility, if this is a syslog event.
firstTime	UNIX timestamp when event is received.
ipAddress	IP Address of the associated device, if applicable.
lastTime	Last time this event was seen and its count incremented.
Location	Device location from device context.
message	Full message text.
monitor	Collector name from which this event came. Note: It is not the FQDN.
nteid	nt event ID, if this is an nt eventlog event.
priority	syslog priority, if this is a syslog event.
prodState	prodState of the device context.
severity	The event severity level .
severityString	the severity of the event expressed as a string (Clear, Debug, Info, Warning,Error, or Critical)
stateChange	Time the MySQLrecord for this event was last modified.
summary	Text description of the event. Limited to 255 characters.
suppid	ID of the event that suppressed this event.
Systems	Device systems from device context, separated by .

Configuration properties and custom properties also are available for devices, and use the same syntax as shown in the previous sections.

Managing multi-realm networks

Resource Manager supports overlapping IP namespaces through the optional [MultiRealmIP](#) ZenPack.

With this ZenPack, Resource Manager can prefix a realm identifier to the IP addresses of a network, enabling unified monitoring. The primary use cases for multi-realm IP management are as follows:

- A large company that manages multiple locations has defined the same network spaces across the locations, and as a result, created multiple overlapping IP spaces. Resource Manager needs a way to identify each separate IP space in the system.
- Service providers that are responsible for monitoring multiple customers that have created independent networks and IP spaces that are unique to their location, but not unique to the service provider.

Configuring IP realms

To set up a system, create the IP realms, and then associate them with a collector. The associations between IP realms and actual devices are made automatically by the device's association with the collector. All devices on a collector are associated with the realm for that collector.

The following procedures demonstrate how to set up a Resource Manager system that is similar to the example. If you do not have overlapping IP spaces, use collectors in the same network. Add a machine multiple times, once per collector, changing the device name as you add it.

Deleting networks and adding IP realms

1. Navigate to INFRASTRUCTURE > Networks.
2. Delete all Resource Manager networks.
The networks are automatically recreated with the realm associations that [MultiRealmIP](#) adds.
3. From the Add menu, select Add IP Realm.
4. Add IP realm `r1`.
5. Repeat the steps to add IP realm `r2`.

Adding collectors to realms

1. Navigate to ADVANCED > Control Center.
2. From the Add menu, select Add a new collector.
3. Name the collector `coll_r1` and specify IP realm `r1`.
4. Repeat the steps to add collector `coll_r2` and specify IP realm `r2`.

Adding devices to realms

When you add devices to the system, you can add the same device twice to simulate a multi-realm setup.


1. Navigate to INFRASTRUCTURE > Devices.
2. Add a device called `A.test` and select a remote collector. Do not select `localhost`.
For instructions, see [Adding a single device](#).
3. Rename the device.
For instructions, see [Renaming a device](#).
4. Add the device a second time, selecting a different remote collector; again do not select `localhost`.
5. In the device list, click the new device name to open the device overview page.
6. In the left panel, choose Model Device.
For instructions, see [Remodeling a device](#).
7. In the left panel, choose Software and follow the network link on one of the interfaces.
The network was created beneath the realm that you created.

As networks are discovered, they are created within each realm. Monitoring occurs on each representation of the device from the different collectors in different overlapping realms.

(Optional) Navigate to INFRASTRUCTURE > Devices. Search for the devices by IP address. Two devices are returned, one in each realm.

Example multi-realm system

The following diagram shows an example system.



The system contains network `10.10.10.0/24`, which has a central Resource Manager server, and is therefore the *default network*. The default network is treated exactly the same as a Resource Manager system without `MultiRealmIP` installed.

The system also contains network `r1` and network `r2`. These networks are behind a firewall and have the same IP space, `192.168.0.0/24`. Each realm has a distributed collector. The Resource Manager server accesses the collector by using an IP translation from the firewall to map the address that is accessible from in front of the firewall to an address that is behind the firewall. Remote collectors in a multi-realm setup must be accessible from the central server using SSH.

Prerequisites and considerations

Review the following information about multi-realm networks.

- Download and install [MultiRealmIP](#) and [DistributedCollector](#).
- On certain server configurations, if a distributed collector is configured, a "zenpack command failed" error occurs when you install [MultiRealmIP](#). If you encounter this error, run the following grant as MySQL root. Replace user and passwd with the user account and password that Resource Manager uses to access MySQL.

```
grant super on *.* to 'user'@'{FQDN_of_Zenoss_host}' identified by 'passwd';
```

- Before you set up the network, you must delete all Resource Manager networks. The networks are automatically recreated with the realm associations that [MultiRealmIP](#) adds.
- Under multi-realm IP networks, device names must be unique even though the IP addresses overlap.
- If an event contains the unique name of a device, assigning it the proper device is simple. However, if only the IP address is sent, the event is assigned by looking up the IP within the context of the realm.
- If a device is moved between realms, the device must be remodeled so that its IPs are placed in the correct location.
- The Network Map only displays the default realm.

Monitoring large file systems

By default, Resource Manager uses the Host Resources MIB to monitor file systems. A defect in the implementation of the Host Resources MIB in net-snmp causes file systems larger than 16TB to report incorrect utilization metrics, such that you might observe file system utilization values greater than 100%.

Note: Resource Manager uses the Host Resources MIB in the ethernetCsmacd template (rather than the UCD dskTable MIB) by default because most of the systems Resource Manager monitors do not have the UCD dskTable MIB enabled.

To work around this deficiency, Resource Manager can instead use the UCD dskTable MIB to monitor file system utilization.

Configuring the UCD dskTable MIB

To use the UCD dskTable MIB, you must modify the configuration of your data sources, thresholds, and graphs in your FileSystem template:

1. Create an SNMP data source named dskPercent.
2. Set the OID of the new data source to:

```
1.3.6.1.4.1.2021.9.1.9
```

3. Modify your thresholds to use the dskPercent data point. Remove any calculations that were associated with the Host Resource MIB data point. The dskPercent data point is reported as an integer from 0 to 100.
4. Modify your graphs to use the dskPercent data point and thresholds.
5. Enable the UCD dskTable MIB on your managed hosts:
 - a. Add the following line to your /etc/snmp/snmpd.conf file:

```
includeAllDisks 0%
```

- b. Restart the snmpd daemon.

The UCD dskTable MIB associates different indexes with the file systems than the Host Resources MIB. As a result, you must remodel the devices to fetch the UCD dskTable indexes, and to begin plotting data on the dskPercent graph.

Integrating LDAP authentication

You can use your existing LDAP authentication infrastructure, such as Active Directory or OpenLDAP, to enable single sign-on to the Resource Manager browser interface. This capability saves you from having to manually create user accounts and separately maintain passwords.

For those LDAP properties that are mapped, changes you make in LDAP are updated in Resource Manager. To propagate property changes immediately, clear your browser cache or log out and back in.

- [LDAP configuration requirements](#)
- [Adding an SSL certificate](#)
- [Configuring LDAP authentication](#)
- [Editing LDAP configurations](#)
- [Configuring local authentication as a fallback](#)
- [Verifying connectivity and credentials outside of Resource Manager](#)

LDAP configuration requirements

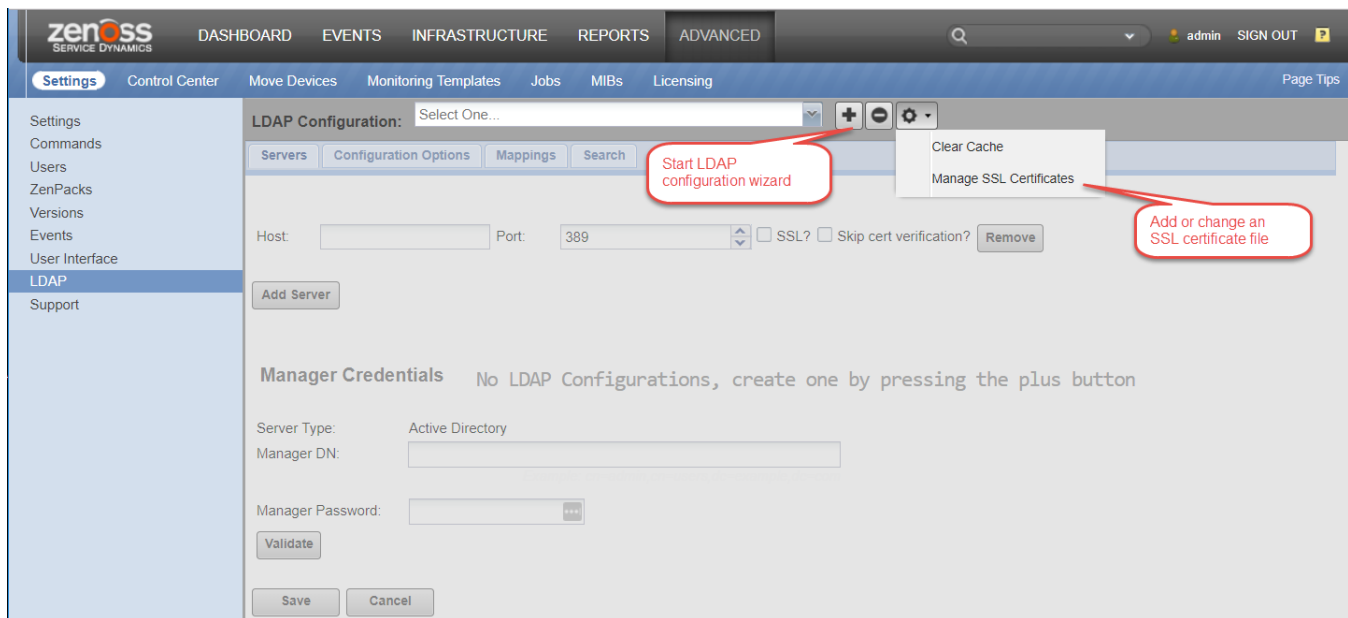
Before configuring LDAP authentication, gather the following information from your LDAP or Active Directory administrator.

- For Active Directory authentication, the host name or IP address of an Active Directory global catalog server.
- For other LDAP server authentication, the host name or IP address of an LDAP server.
- Distinguished name of a manager user in the domain administrators group (Manager DN).
- Password for the Manager DN (Manager password).
- User's base distinguished name (Users Base DN)
- The distinguished name for the branch of your LDAP database that contains group records (Groups Base DN). These group records are of the LDAP class "groupOfUniqueNames," and the entry CN attribute constitutes the group name.
- Optionally, Active Directory groups to map to Resource Manager roles.

Accessing the LDAP configuration wizard

You can configure LDAP authentication at initial setup, or from the Settings area of the Resource Manager browser interface. Use one of the following methods to start the LDAP configuration wizard:

- While in the setup wizard, at Step 2: Specify or Discover Devices to Monitor, click LDAP Setup (located at the bottom right of the wizard panel).
- From the browser interface, choose ADVANCED > LDAP and then click Add.

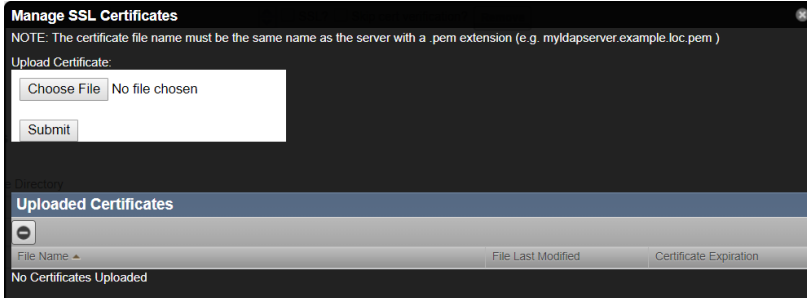


Adding an SSL certificate

The default configuration of the Resource Manager web server uses a Zenoss self-signed certificate for SSL/TLS communications. For information about installing your own digital certificate on the master host, see [Optional: Replacing the default digital certificate](#).

To upload a certificate file for this instance of Resource Manager, complete the following steps.

1. From the Resource Manager browser interface, choose **ADVANCED > LDAP**.
2. Click the gear icon and choose **Manage SSL Certificates**.



3. On the Manage SSL Certificates page, choose the certificate file with extension `.pem`, and then click **Submit**.

Configuring LDAP authentication

Use pages of the LDAP configuration wizard to configure authentication.

1. On the LDAP configuration wizard Add LDAP Servers page, specify the host and manager credentials.

New LDAP Configuration

1. Add LDAP Servers

Server Type: Active Directory Other LDAP

Host: Port: SSL? Skip cert verification?

Manager Credentials

Server Type: Active Directory Other LDAP

Manager DN:
Example: cn=admin,cn=users,dc=example,dc=com

Manager Password:

- Host - Enter the host name or IP address of an Active Directory global catalog server (for Active Directory authentication) or an LDAP server (for other LDAP server types).
 - Port - Optionally, change the server port number. By default, the port number is 389.
 - SSL - Choose if you are using SSL. When you choose this option, the default port number adjusts to 636.
 - Skip cert verification? - If you are using a self-signed certificate, choose this check box to skip its verification. Requires OpenLDAP 2.4 or later.
2. Optional: To add another LDAP server, click Add Server. To remove a server from the list, click Remove.
 3. In the Manager Credentials area, provide the following information:
 - Server Type
 - Manager DN- Enter the distinguished name of a manager user in the domain administrators group. For example, the user's base DN:

```
cn=admin,cn=users,dc=example,dc=com
```
 - Manager Password
 4. To ensure that your setup is valid, click Validate.
 5. Click Next.
 6. On the Configure LDAP Plugin page, the configuration ID field is populated with the host name that you provided. Specify user and group information.

New LDAP Configuration

2. Configure LDAP Plugin

LDAP Configuration ID:

Login Name Attribute:

Users Base DN:
Example: dc=Users,dc=example,dc=com

Groups Base DN:
Example: dc=Groups,dc=example,dc=com

User Filter:
Example: (cn=Organization.)*

Group Filter:
Example: (cn=IT Admins)

Default User Roles:

- Login Name Attribute - Choose the LDAP record attribute that is used as the user name. To add attributes, use the Mappings page of the LDAP configuration area (ADVANCED > LDAP).
- Users Base DN - For example, if your domain is ad.example.com, then your users base DN might be:

```
dc=Users,dc=example,dc=com
```

- Groups Base DN
 - User Filter and Group Filter - Using correct LDAP search filter syntax, specify free-form LDAP filter expressions to be added to the default user and the default group search filters. The default search filters and the additional search filters are combined as an AND expression. For the searches to return a record, the record must satisfy both filters.
 - Default User Roles - From the drop-down list, select roles to be given to all users that are authenticated from your LDAP tree. Zope expects all users, anonymous and authenticated, to have the role Anonymous.
7. Click Next.
 8. On the Map LDAP Groups to Local Groups page, provide group and role information.

New LDAP Configuration

3. Map LDAP Groups to Local Groups

Map LDAP Groups to Roles?

Group: Role:

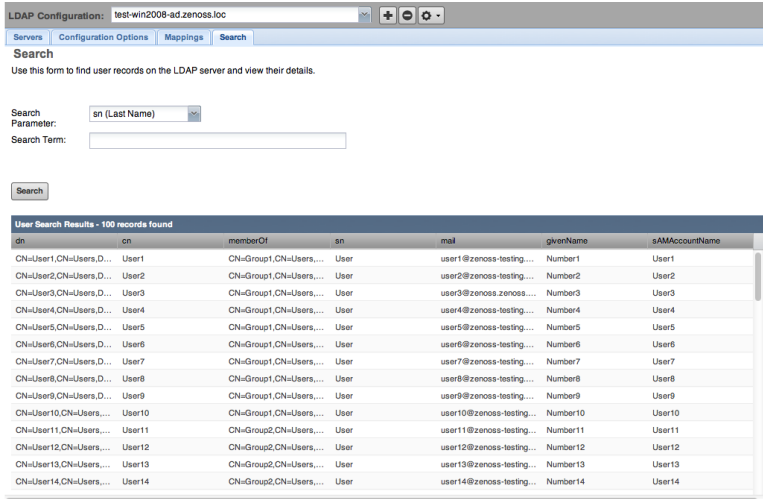
- Map LDAP Groups to Roles? - Choose this option if you want to control user roles within the Resource Manager browser interface by using Active Directory groups, instead of controlling the roles directly from within the system. Add the following groups to LDAP:
 - Resource Manager Managers
 - Resource Manager Users
 - Group - Choose the LDAP group to map to a Resource Manager role.
 - Role - Choose the Resource Manager role to map the LDAP group.
9. To map another group, click Add Group Mapping. To remove a mapped group, click Remove.
 10. Click Finish.

After setup, you can edit LDAP configuration settings from the Settings, Configuration Options, and Mappings tabs.

Editing LDAP configurations

After initial setup, you can search for and edit LDAP configuration settings.

1. From the Resource Manager browser interface, choose **ADVANCED > LDAP**.



LDAP Configuration: test-win2008-ad.zenoss.loc

Servers Configuration Options Mappings Search

Search

Use this form to find user records on the LDAP server and view their details.

Search Parameter:

Search Term:

Search

User Search Results - 100 records found

dn	cn	memberOf	sn	mail	givenName	sAMAccountName
CN=User1,CN=Users,D...	User1	CN=Group1,CN=Users...	User	user1@zenoss-testing...	Number1	User1
CN=User2,CN=Users,D...	User2	CN=Group1,CN=Users...	User	user2@zenoss-testing...	Number2	User2
CN=User3,CN=Users,D...	User3	CN=Group1,CN=Users...	User	user3@zenoss-testing...	Number3	User3
CN=User4,CN=Users,D...	User4	CN=Group1,CN=Users...	User	user4@zenoss-testing...	Number4	User4
CN=User5,CN=Users,D...	User5	CN=Group1,CN=Users...	User	user5@zenoss-testing...	Number5	User5
CN=User6,CN=Users,D...	User6	CN=Group1,CN=Users...	User	user6@zenoss-testing...	Number6	User6
CN=User7,CN=Users,D...	User7	CN=Group1,CN=Users...	User	user7@zenoss-testing...	Number7	User7
CN=User8,CN=Users,D...	User8	CN=Group1,CN=Users...	User	user8@zenoss-testing...	Number8	User8
CN=User9,CN=Users,D...	User9	CN=Group1,CN=Users...	User	user9@zenoss-testing...	Number9	User9
CN=User10,CN=Users...	User10	CN=Group1,CN=Users...	User	user10@zenoss-testing...	Number10	User10
CN=User11,CN=Users...	User11	CN=Group2,CN=Users...	User	user11@zenoss-testing...	Number11	User11
CN=User12,CN=Users...	User12	CN=Group2,CN=Users...	User	user12@zenoss-testing...	Number12	User12
CN=User13,CN=Users...	User13	CN=Group2,CN=Users...	User	user13@zenoss-testing...	Number13	User13
CN=User14,CN=Users...	User14	CN=Group2,CN=Users...	User	user14@zenoss-testing...	Number14	User14

2. On the Settings, Configuration Options, and Mappings tabs, edit settings as needed.
3. On the Search tab, locate user records on the LDAP server. Choose from the list of search parameters, enter a search term, and then click Search.

Search results return on the lower portion of the page.


Configuring local authentication as a fallback

Use local authentication as a fallback in the event that the LDAP server is unavailable. The local authentication plugin is called userManager.

1. Verify that the userManager plugin is available.
 - a. In a web browser, navigate to the Zope Management Interface. Replace Zenoss-Host with the hostname or IP address of your Resource Manager server:

```
https://Zenoss-Host/zport/acl_users/manage
```

- b. In the Name column, click plugins.
- c. Click Authentication Plugins.



- d. Verify that your LDAP plugin is first in the list of active plugins, and that the userManager plugin is second.
2. Create one or more user accounts.

The accounts created in this step enable access when the LDAP server is unavailable. If you use the same account name in this step as the LDAP account name, the user need only remember the "fallback" password.

 - a. In the Resource Manager browser interface, choose ADVANCED > Users > Add > New User.
 - b. Create user accounts.

This account must be created before the account with the same name is authenticated by the LDAP server. The passwords that are defined when creating accounts in Resource Manager are only valid when the LDAP server is unavailable.

Verifying connectivity and credentials outside of Resource Manager

Verify that your credential information is valid from the Resource Manager server by using the `ldapsearch` command. To install this command, use the following commands for RPM-based systems:

```
yum install openldap-clients
```

As the `zenoss` user on the Resource Manager server:

```
ldapsearch -LLL -x -b 'BaseDN' -D 'Bind DN' -W -H ldap://LDAP_server-name "sAMAccountName=*" member
```

Tuning Considerations

Tuning Zenoss Control Center and Resource Manager is a tricky subject due to the many factors that must be considered. For example:



- CPU core counts, including speed, hyper-threading, etc.
- RAM capacity and speed
- Disk size, speed, and technology
- [Managed Resource](#) count
- The number of components per Managed Resource
- The number of datasources on each device and component
- Event counts from active device polling, passive collection (zentrap, etc.), and event collection (i.e., Windows and VMWare event logs)

As such, what follows is to be treated as a series of best-effort recommendations the suitability of which should be considered suspect until tested. The following chart provides some example deployment sizes that we will reference in upcoming pages.

Size	Managed Resource Count
Small	< 1000
Medium	1000 - 5000
Large	5000 - 10000
XL	10000 - 20000
XXL	> 20000

When in doubt, please contact Zenoss Support or Professional Services for guidance.

Analytics ETL Services

Daemon	Instances	Memory (Small/Default)	Memory (Med)	Memory (Large)	Config Options
Zen Event ETL	1	500 MB	500 MB	500 MB	
Zen Model ETL	1	1 GB	1 GB	1 GB	
Zen Perf ETL	1	1 GB	4 GB	4 GB	If the zenperf logs indicate that the service is waiting for the cache (i.e., zenperf batches don't start 15 minutes after the hour), you can add instances of the service to assist in invalidation processing. Note, only instance 0 performs the ETL process; there is no way to make anything go faster.

These three services are all written in Python. As such, the RAM requested value in Control Center will only be used to determine which delegate host in a resource pool to assign the service to.

Collection Services

The following tuning recommendations have been tested by Zenoss Professional Services deployment engineers and have been judged as safe values for most use-cases. Due to differences between instances, however, these recommendations may be "overkill" or inadequate for your infrastructure's requirements.

Please keep the following in mind:



1. The Device Classes column provides example classes and is not exhaustive. Due to user customization and extending the product through ZenPacks, there are classes not accounted for in this table.
2. Certain Device Classes make use of multiple datasource types. For instance, devices in the /Network/Cisco/Nexus sub-classes make use of the SNMP, Command, and Python datasource types.
3. The recommendations given are based on default cycle times, except where noted. Increasing or decreasing cycle times will have an effect on instance count requirements.

Daemon	Instances	Memory	Sample Device Classes (count devices)	Config Options
zenhub	1	2 GB		<p>1. Confirm the service definition HostPolicy is configured to PREFER_SEPARATE. If upgrading from 6.2.x to 6.3.x, confirm again after upgrade.</p> <p>2. For 6.2 and earlier, set the following in zenhub.conf:</p> <p>workers: S: 2 M: 4 L: 8 XL: 12 XXL: 16 (anything L or bigger should be a SuperHub)</p> <p>3. For 6.3 and earlier, set the following in zenhub.conf:</p> <p>invalidationworkers: S: 2 M: 3 L: 4 XL: 6 XXL: 8 worker-call-limit: S: 200 (the default) M: 200 L: 500 XL: 500 XXL: 1000</p>
zenhubworker	S: 2 M: 4 L: 8 XL: 12 XXL: 16	1 GB		Introduced in 6.3. Replaces manually setting the workers option in zenhub.conf. In zenhubworker.conf, set the following: call-limit: S: 200 (default) M: 200 L: 500 XL: 500 XXL: 1000
zenhubadmworker	1	1 GB		Introduced in 6.4. Processes large applyDataMap tasks, freeing normal zenhubworkers for other tasks.
zenhubiworker	S: 2 M: 3 L: 4 XL: 6 XXL: 8	1 GB		Introduced in 6.4. Replaces manually setting the invalidationworkers option in zenhub.conf.
zenmodeler	1	1 GB	All	<p>1. In zenmodeler.conf, set parallel to 10</p> <p>2. In zenmodeler.conf, set cycletime to 1440</p> <p>3. Configure staggered modeling per https://support.zenoss.com/hc/en-us/articles/215003503</p>
zenping	1 instance / 1000 devices	1 GB	All	Instance recommendation is based on default cycle time and timeout values. Longer timeout values and shorter cycle times will require additional instances.
zencommand	1 instance / 200 devices	1 GB	/Server/SSH /Network/Cisco/Nexus	
zenperfsnmp	1 instance / 200 devices	1 GB	/Network /Server (minus /Server/SSH, minus /Server/Microsoft/Windows) /Server/Windows /Storage	
zenpython	1 instance / 200 devices	1 GB	/Server/Microsoft /Storage	
zenprocess	1 instance / 200 devices	1 GB	/Server	To get a device count, navigate to Infrastructure > Processes and click the top-most organizer ("Processes"). The number of devices being monitored for processes will list in the lower-right corner. For example, "Displaying 1 - 16 of 341 rows."
zenstatus	1	1 GB		
zenjmx	1	1 GB		
zenucsevents	1	1 GB		
zenvsphere	1	1 GB		May require additional instances depending on size of vSphere instances and count. Check vSphere health report for details.
zenproptymonitor	1	256 MB	/vSphere	

zenwe btx	1	1 GB		
zensy slog	Never increase to more than 1	1 GB		
zentrap	Never increase to more than 1	1 GB		
zenma il	0	1 GB		Disabled by default. If you plan to use this service, you will need to start it (and configure autostart) manually. See Advanced > Control Center to configure autostart.
zenpo p3	0	1 GB		Disabled by default. If you plan to use this service, you will need to start it (and configure autostart) manually. See Advanced > Control Center to configure autostart.
zenma iltx	0	1 GB		Deprecated in 6.x
zminion	1	2 5 6 MB		
collect orredis	5.x: 1 6.x: Set to EXACTLY the number of hosts in the collector pool	2 GB		This is for resiliency and quicker collection start up using cached configs, not for scale.
Metric Shipp er	1	1 GB		

Event Processing Services

Daemon	Instances	Memory (Small/Default)	Memory (Med)	Memory (Large)	Config Options
zenactiond	2	1 GB	1 GB	1 GB	If impact installed proactively set maintenance-window-batch-size to 10 in zenactiond.conf and if not, leave at default 200.
zeventd	S - 2; M - 4; L - 8	1 GB	1 GB	1 GB	
zeventserver	1	4 GB	6 GB	8 GB	
zenimpactstate	1	1 GB	1GB	1 GB	If additional instances of zenimpactstate are required, disable pre-fetching in zenimpactstate.conf by setting prefetch 1. Failure to do so may result in messages being processed out of order, leading to inconsistent service state information (i.e., a clear event can be processed before its associated fault event).

Infrastructure Services

Daemon	Instances	Memory (Small /Default)	Memory (Med)	Memory (Large)	Config Options/Notes
Region Server	3	1 GB	2 GB	8 GB	
Zookeeper	3	1 GB	1 GB	1 GB	
Impact	1	1 GB	4 GB	8 GB	
MariaDB Events	1	2 GB	4 GB	8 GB	Also increase max_connections to 1000 if it has not already been done
MariaDB Model	1	4 GB	8 GB	20 GB	Also increase max_connections to 1000 if it has not already been done. Memory allocation should be roughly the size of ZODB on disk (du -sh on ZODB DFS directory).
Memcached	1	4 GB	8 GB	20 GB	Memory allocation should be set to match MariaDB model.
memcached-session	1	1 GB	1 GB	1 GB	
OpenTSDB Reader	1	1 GB	1 GB	1 GB	May need more instances but shouldn't need more memory.
OpenTSDB Writer	1	1 GB	1 GB	1 GB	May need more instances but shouldn't need more memory.
RabbitMQ	1	256 MB	256 MB	256 MB	
Redis	1	1 GB	1 GB	1 GB	
Solr or zencatalogservice	1	1 GB	2GB	4GB	
HMaster	1	1 GB	1 GB	1 GB	

Metric Services

Daemon	Instances	Memory (Small/Default)	Memory (Med)	Memory (Large)	Config Options
CentralQuery	1	1GB	1GB	1GB	Increase memory to 4GB if Analytics is installed.
MetricConsumer	1	2G	2G	2G	
MetricShipper	1	256 MB	256 MB	256 MB	

Tuning Control Center

Tuning for all instance sizes

- In `/etc/default/serviced` on the Control Center master, set `GOMAXPROCS` to half of the value of the CPU core count.
- In `/etc/nfsmount.conf` on the Control Center master, set `Defaultvers=4.0`

Tuning for medium and large instances

- Confirm that the Control Center master has at least 8 CPU cores and 32 GB of RAM. If so, make the following changes proactively:
 - Immediate change: `echo 1 > /proc/sys/vm/swappiness`
 - Permanent change (requires restart to take effect): set `vm.swappiness = 1` in `/etc/sysctl.conf`
- Raise the Java memory allocations for the 3 logging related services which all default to 1GB by editing `/etc/default/serviced`
 - `SERVICED_ISVCS_ENV_0=elasticsearch-logstash:ES_JAVA_OPTS=-Xmx4g`
 - `SERVICED_ISVCS_ENV_1=elasticsearch-serviced:ES_HEAP_SIZE=2g`
 - `SERVICED_ISVCS_ENV_2=logstash:LS_HEAP_SIZE=2g`
- Set the Timezone (`TZ=`) in `/etc/default/serviced` config on all Control Center delegates (and master) based on the master time zone.

User Interface Services

Daemon	Instances	Memory (Small /Default)	Memory (Med)	Memory (Large)	Config Options
Zauth	2	1 GB	1 GB	1 GB	Set to restart at 150%. Set PREFER_SEPARATE in service definition.
zenjobs	1	1 GB	1 GB	1 GB	Increase instances to scale as needed.
zenjserver	1	1 GB	1 GB	1 GB	
zope	6	2GB	2 GB	2GB	One Zope for 10 concurrent users (never less than 6 instances). Set to restart at 150%.
zenapi (API Zope)	1	1 GB	1 GB	1 GB	Set to restart at 150%.
zproxy	1	1 GB	1 GB	1 GB	
zenreports (Reporting Zope)	1	1 GB	1 GB	1 GB	Set to restart at 150%.
zendebug (Debug Zope)	0	1G	1GB	1GB	Off By Default. Set to restart at 150%.
UI Settings (Advanced> User Interface)					
Incremental Tree Loading		Enable	Enable	Enable	
Load Device Tree From Catalog		Disable	Enable	Enable	

Glossary

application

In Control Center, one or more software [services](#) packaged in Docker containers. Resource Manager is the only application that Zenoss supports in Control Center.

collector

A logical group of Resource Manager services that gather availability, performance, or model data about [devices](#).

Collectors use generic protocols (ICMP, SNMP) or customized protocols (API calls) to gather data.

data point

The value of a [metric](#) at a specific time.

data source

A query that uses a specific protocol to request one or more [metrics](#) from a [device](#) or application.

For example, to gather system uptime from a Linux host, an SNMP data source requests a specific OID, and an SSH data source runs the `uptime` command.

delegate host

A Control Center host that runs the [application services](#) scheduled for the [resource pool](#) to which it belongs. A host can be configured as delegate or [master](#).

device

A physical or virtual system, or a virtual or cloud infrastructure, from which Resource Manager gathers availability, performance, and model data.

Example systems include servers, routers, storage systems, and switches. Example infrastructures include Amazon Web Services, VMware vCloud, and Microsoft Azure.

See also: [device component](#)

Synonym: managed resource

device class

The base set of properties that characterize a group of [devices](#). A [device](#) may belong to only one device class.

Device class properties identify the data to collect and the protocol for collecting data.

device component

Resources that belong to a specific [device](#), or processes that run on a specific [device](#).

For example, file systems, applications, and network interfaces are device components.

discovery

The process of identifying [device components](#).

The results of discovery are used to populate the [model](#).

event

A message that is sent to the event processing service.

External events are solicited (for example, through ZenPacks) or unsolicited (through SNMP traps or syslog). Internal events are created in response to failed pings, [threshold](#) violations, and other occurrences.

event class

A named set of properties that define a category of [events](#), and a set of [event rules](#).

event rule

A Python expression that the event processing service uses to determine whether to change the class of an event.

For more information, see [Event class mappings](#).

graph

A diagram of the relationships among [data points](#), [thresholds](#), or both, over a specific period of time.

master host

The Control Center host that runs the [application services](#) scheduler, the Docker registry, the distributed file system, and other internal services, including the HTTP server for the Control Center browser interface and application browser interface. A system can be configured as [delegate](#) or master. Only one Control Center host can be the master.

metric

A measurement that quantifies an [entity](#) property or phenomenon.

Example metrics include input/output operations per second (IOPS), CPU utilization, network interface throughput, and the amount of free space on a block device.

See also: [data point](#)

model

A schematic description of a [device](#) and its [components](#).

monitoring template

A named object that associates [data sources](#), [thresholds](#), and [graph definitions](#) with a [device](#) or [device component](#). For devices, a monitoring template also includes a [device class](#).

Monitoring templates define what and how to monitor, and how to present [data points](#).

notification

A named object that associates one or more [triggers](#) with an action.

Example actions include sending email to a list of recipients and running a command.

organizer

An implicit or explicit way to group [devices](#).

[Device classes](#) are implicit organizers. Explicit organizers enable arbitrary groupings of devices and device components in the Resource Manager interface.

For more information, see [Using organizers](#).

resource pool

A collection of one or more Control Center hosts, each with its own compute, network, and storage resources. All of the hosts in a resource pool must have identical hardware resources, and must be located in the same data center and on the same subnet. If a resource pool host is a hypervisor guest system, all of the hosts in the resource pool must be guests of the same hypervisor host system.

service

A process and its supporting files that Control Center runs in a single container to provide specific functionality as part of an [application](#).

serviced

The name of the Control Center daemon and a command-line client for interacting with the daemon.

tenant

An [application](#) that Control Center manages.

threshold

A named object that specifies the boundary criteria for the [data points](#) of a [metric](#).

Resource Manager services consider thresholds when [data points](#) are received and create [events](#) when a [data point](#) violates one or more criteria.

trigger

A named set of rules ([Boolean expressions](#)) that the event processing service uses to examine [events](#).

See also: [notification](#)

ZODB

Zope object database. For more information, refer to [ZODB](#).