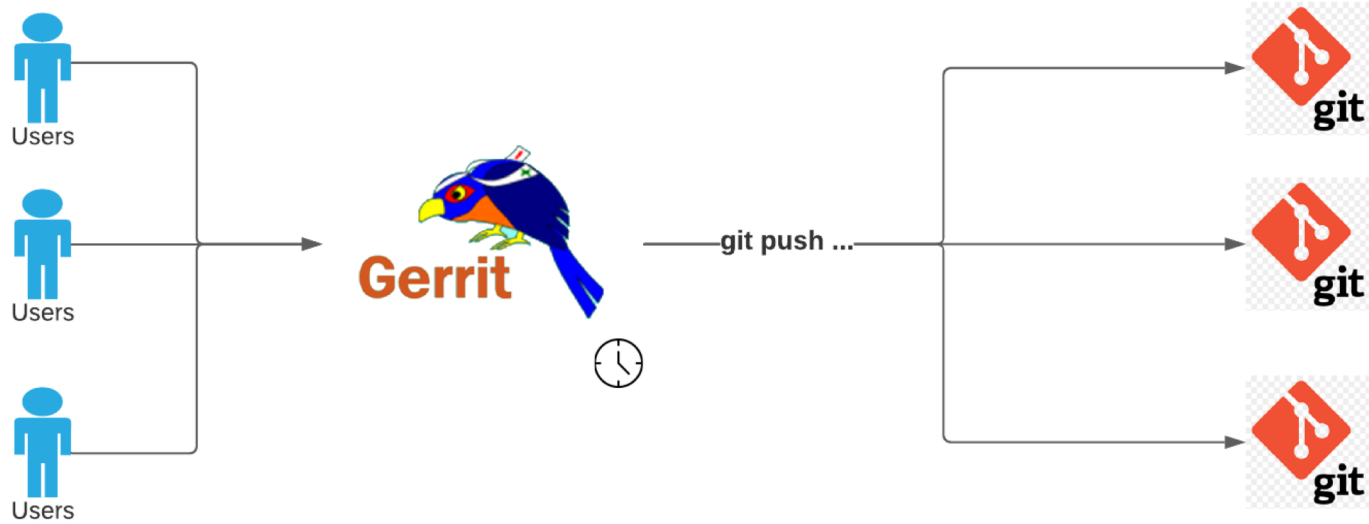


An introduction to the pull- replication plugin

Agenda:

- (push) replication and pull-replication overview
- Pull-replication configuration
- Demo
- Benchmarks
- Q&A

(push) Replication plugin overview



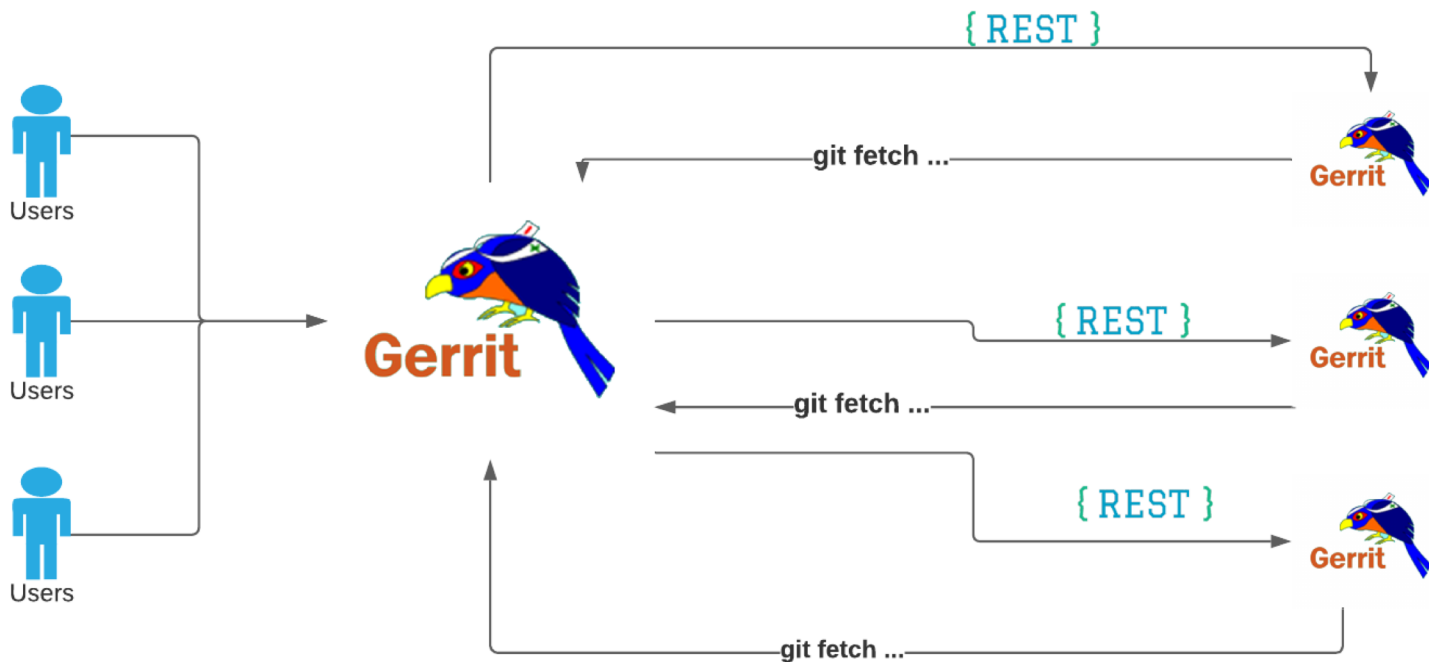
Git push vs Git Fetch

Whenever a client is issuing a push, all the references are send from the server to the client, even if the client was pushing a single ref!

```
$ export GIT_TRACE_PACKET=1
$ git push --progress ssh://admin@primary-1:/gerrit-pull-replication refs/changes/99/324499/9
23:17:24.208142 pkt-line.c:80      packet:      push< 54bf8ddf9bcb4ffc748dd6fead98eeca95c98e5 refs/heads/infra/config\0 side-
band-64k delete-refs report-status quiet atomic ofs-delta push-options agent=JGit/v5.12.0.202106070339-r-49-g60b81c5a9
23:17:24.208159 pkt-line.c:80      packet:      push< fdb3d0578926e3533d719507db4c9f4eb97d184f refs/heads-master
23:17:24.208164 pkt-line.c:80      packet:      push< ed2d5cedd8d79ee224fcc2280a6f11e8175fc2b0 refs/heads/stable-2.10
23:17:24.208168 pkt-line.c:80      packet:      push< 57bb9785c8eb51568a9bfc84b29fcb351d843a17 refs/heads/stable-2.11
23:17:24.208172 pkt-line.c:80      packet:      push< 44ef71035a53f527a0821d6ea195d32caeaafe90 refs/heads/stable-2.12
23:17:24.208175 pkt-line.c:80      packet:      push< f43a203a8650c4f139b57730e720d7cb56b2a2e7 refs/heads/stable-2.13
...
```

Thanks to git protocol v2 for fetch operation refs advertisement is stripped out (since the ls-refs command fills that role)

Pull-replication plugin overview



Configuration: primary - replica scenario

- Primary node
 - etc/replication.config

```
[replication]
  instanceLabel = primary
```

- etc/replication/replica-1.config

```
[remote]
  url = git://replica-1:29500/${name}.git
  apiUrl = http://replica-1:8080
  fetch = +refs/*:refs/*
  mirror = true
  replicationDelay = 0
  replicateHiddenProjects = true
```

- etc/replication/replica-2.config

```
[remote]
  url = git://replica-2:29500/${name}.git
  apiUrl = http://replica-2:8080
  fetch = +refs/*:refs/*
  mirror = true
  replicationDelay = 0
  replicateHiddenProjects = true
```

- Replica nodes
 - etc/replication/**primary**.config

```
[remote]
  url = git://primary:29500/${name}.git
  apiUrl = http://primary:8080
  fetch = +refs/*:refs/*
  mirror = true
  replicationDelay = 0
  replicateHiddenProjects = true
```

Configuration: primary - replica scenario

- Primary node
- etc/replication.config

```
[replication]
  instanceLabel = primary
```

- etc/replication/replica-1.config

```
[remote]
  url = git://replica-1:29500/${name}.git
  apiUrl = http://replica-1:8080
  fetch = +refs/*:refs/*
  mirror = true
  replicationDelay = 0
  replicateHiddenProjects = true
```

- etc/replication/replica-2.config

```
[remote]
  url = git://replica-2:29500/${name}.git
  apiUrl = http://replica-2:8080
  fetch = +refs/*:refs/*
  mirror = true
  replicationDelay = 0
  replicateHiddenProjects = true
```

- Replica nodes
- etc/replication/**primary**.config

```
[remote]
  url = git://primary:29500/${name}.git
  apiUrl = http://primary:8080
  fetch = +refs/*:refs/*
  mirror = true
  replicationDelay = 0
  replicateHiddenProjects = true
```

- REST Api call:

POST http://replica-
...:8080/a/projects/**foo**/pull-replication~fetch

```
{"label": "primary", "ref_name":  
"refs/changes/01/1/meta",  
"async": false}
```

Configuration: multi-site scenario

- Primary-1 node

- etc/replication.config

```
[replication]
    instanceLabel = primary-1
    excludeRefs = refs/multi-
site/version
```

- etc/replication/primary-2.config

```
[remote]
    url = git://primary-2:29500/${name}.git
    apiUrl = http://primary-2:8080
    fetch = +refs/*:refs/*
    mirror = true
    replicationDelay = 0
    replicateHiddenProjects = true
```

- Primary-2 node

- etc/replication.config

```
[replication]
    instanceLabel = primary-2
    excludeRefs = refs/multi-
site/version
```

- etc/replication/primary-1.config

```
[remote]
    url = git://primary-1:29500/${name}.git
    apiUrl = http://primary-1:8080
    fetch = +refs/*:refs/*
    mirror = true
    replicationDelay = 0
    replicateHiddenProjects = true
```


Apply object REST API call

```
POST http://replica-...:8080/a/projects/foo/pull-replication~apply-object
```

```
{ "label": "primary-1",  
  "ref_name": "refs/changes/81/1981/meta",  
  "revision_data": { "commit_object": { "type": 1, "content": "dHJlZSA0YjgyNWRjNjQyY2I2ZWI5YTA2MGU1NGJmOGQ2OTI4OGZiZWU0OTA0CnBhcmVudCA3ZTFkNTM5NGQwMDhiOGM2ZmU1NTI4Y2VlOTAyODJiMjQxYjA3ZThmCmF1dGhvciBHZXJyaXQgVXNlciAxMDAwMDAwIDwxMDAwMDAwQDY5ZWZmZGZlLTM1MGUtNGQ5Yy05NmQ0LWJjOTU2ZjJmYWZhYz4gMTYxMTY4MDU0NyArMDEwMApjb21taXR0ZXIgr2Vycml0IENvZGUgUmV2aWV3IDxyb290QG1hY3plY2gtWFBTLTE1PiAxNjExNjgwNTQ3ICswMTAwCgpVcGRhdGUgcGF0Y2ggc2V0IDEKClBhdGNoIFNldCAxOjBDb2RlLVJldmllldysyCgpUaGlzIGNoYW5nZSBpcyByZWZkeSBmb3IgcmlvZmV3LgoKUGF0Y2gtc2V0OjAxClJldmllld2VyOiBHZXJyaXQgVXNlciAxMDAwMDAwIDwxMDAwMDAwQDY5ZWZmZGZlLTM1MGUtNGQ5Yy05NmQ0LWJjOTU2ZjJmYWZhYz4KTGFhZlZw6IEVvZGUtUmV2aWV3PSsyClcvcmstaW4tcHJvZ3Jlc3M6IGZhbnHNlCg\u003d\u003d"}, "tree_object": { "type": 2, "content": "" }, "blobs": [] },  
  "async": false }
```

Deployment scenarios:

- Standalone mode
- (push) Replication plugin + pull-replication plugin

DEMO

Benchmark:

Gerrit repository ~45 000 refs:

- push replication ~ 10 seconds
- pull replication fetch ~ 1.5 second
- pull replication apply object ~ 5 ms

Q&A