

# UCS Hyperflex Zookeeper 서비스 설명 및 확인

## 목차

### [소개](#)

### [이그지비터 서비스 상태 확인](#)

### [Zookeeper 서비스 쿼리](#)

### [라이브 설정의 Zookeeper 로그 파일](#)

### [지원 번들\(storfs\)의 Zookeeper 로그 파일](#)

## 소개

이 문서에서는 분산 시스템의 중앙 집중식 서비스인 ZooKeeper를 계층 키 값 저장소에 대해 설명합니다. 대규모 분산 시스템에 분산 구성 서비스, 동기화 서비스 및 이름 지정 레지스트리를 제공하는 데 사용됩니다. ZooKeeper의 아키텍처는 중복 서비스를 통해 고가용성을 지원합니다. 그러면 고객은 다른 동물원 사육사가 대답하지 못하면 물어볼 수 있습니다. ZooKeeper 노드는 파일 시스템이나 트리 데이터 구조와 마찬가지로 계층 네임스페이스에 데이터를 저장합니다. 클라이언트는 노드에서 읽고 노드에 쓸 수 있으며 이러한 방식으로 공유 구성 서비스를 가집니다. ZooKeeper는 업데이트가 완전히 주문되는 원자 브로드캐스트 시스템으로 볼 수 있습니다.

ZooKeeper는 다음과 같은 주요 기능을 제공합니다.

- 신뢰할 수 있는 시스템 - 노드가 실패하더라도 시스템이 계속 작동하므로 매우 안정적입니다.
- 간단한 건축 - ZooKeeper의 건축물은 꽤 간단합니다. 공유 계층 네임스페이스를 사용하여 프로세스 조정을 지원합니다.
- 빠른 처리 - ZooKeeper는 읽기 중심 워크로드에서 특히 속도가 빠릅니다.
- 확장성 - 노드를 추가하여 ZooKeeper의 성능을 향상시킬 수 있습니다.

HX에는 다음과 같은 구체적인 구현이 있습니다.

- 이그지비터라는 서비스는 Zookeeper 시작/종료를 관리합니다.
- HX 클러스터 내의 프로세스는 Zookeeper에 대한 클라이언트로, tcp 포트 2181ex storfs, stmgr 등을 통해 통신합니다.
- 노드가 5개 이상인 시스템에는 독립형 노드가 있습니다. 5개 이하의 노드가 있는 시스템은 독립형 노드가 없어야 합니다.
- 쿼럼 =  $N/2 + 1$ 에 필요한 최소 노드 수

예를 들어, 3노드 클러스터의 경우  $N/2=1.5$  Rounded off to  $1 + 1 = 2$ (노드 오류는 하나만 허용됨)

예를 들어 5노드 클러스터의 경우  $N/2=2.5$  2로 반올림됨 =  $1 + 3$ (2개의 노드 오류만 허용됨)

ZK 클러스터에 대해 5개의 노드만 수행하므로 클러스터의 노드 수에 관계없이 최대 2개의 노드 오류만 허용합니다. 컨버지드 노드에서도 마찬가지입니다.

## 이그지비터 서비스 상태 확인

```
root@SpringpathControllerMSH7NHXRFL: /var/log/zookeeper# service exhibitor status
exhibitor start/running, process 4905
```

```
root@help:/var/log/springpath# ps -aux | grep -i exhibitor
root 12519 0.0 0.2 4690592 198892 ? Ssl May19 7:19 exhibitor -cp exhibitor.jar:/etc/exhibitor/ -
Xmx256M -XX:+HeapDumpOnOutOfMemoryError -
XX:HeapDumpPath=/var/log/exhibitor_heap_dump_2019_05_19_22:19:48.hprof -
Dlog4j.configuration=file:///etc/exhibitor/log4j.properties -
Dsringpath.zkdownscript=/usr/share/springpath/storfs-misc/zkMonitor.sh -
Djava.security.egd=file:/dev/./urandom -jar exhibitor.jar --hostname 10.197.252.100 -c file --
fsconfigdir /etc/exhibitor --port 8180 --listenaddress 10.197.252.100
root@help:/var/log/springpath# pidof exhibitor
12519
```

## Zookeeper 서비스 쿼리

Zookeeper에는 상태, 연결 나열, znodes 수 등을 쿼리할 수 있는 4자 명령 구문이 있습니다.

로컬 노드에서 Zookeeper 상태를 확인합니다. (ruok ==> Are you OK?. i목==>I am OK(정상)).

```
root@SpringpathControllerMSH7NHXRFL:/var/log/zookeeper# echo ruok|nc localhost 2181
imok
```

동물사육사가 원동체인지 종동체인지 확인합니다.

```
root@SpringpathControllerMSH7NHXRFL:/var/log/zookeeper# echo srvr | nc localhost 2181
Zookeeper version: 3.4.6--1, built on 06/16/2015 22:50 GMT
Latency min/avg/max: 0/0/101
Received: 213128515
Sent: 213164119
Connections: 6
Outstanding: 0
Zxid: 0xa000301d0
Mode: leader
Node count: 17090
```

```
root@SpringpathControllerMSH7NHXRFL:/var/log/zookeeper# echo stat | nc localhost 2181
Zookeeper version: 3.4.6--1, built on 06/16/2015 22:50 GMT
Clients:
 /192.168.5.161:56128[1](queued=0,recvd=169146196,sent=169162634)
 /192.168.5.161:38614[1](queued=0,recvd=186015,sent=186017)
 /192.168.5.164:44412[1](queued=0,recvd=184398,sent=184399)
 /192.168.5.164:44447[1](queued=0,recvd=561168,sent=563034)
 /127.0.0.1:60060[0](queued=0,recvd=1,sent=0)
 /192.168.5.161:58754[1](queued=0,recvd=39233,sent=39261)

Latency min/avg/max: 0/0/101
Received: 213109927
Sent: 213145531
Connections: 6
Outstanding: 0
Zxid: 0xa000301d0
Mode: leader
Node count: 17090
```

```
root@SpringpathControllerMSH7NHXRFL:/var/log/zookeeper# echo mntr | nc localhost 2181
zk_version      3.4.6--1, built on 06/16/2015 22:50 GMT
zk_avg_latency  0
zk_max_latency  101
zk_min_latency  0
zk_packets_received  213148668
zk_packets_sent    213184272
```

```
zk_num_alive_connections    6
zk_outstanding_requests    0
zk_server_state             leader
zk_znode_count              17090
zk_watch_count              4305
zk_ephemerals_count        20
zk_approximate_data_size    1831768
zk_open_file_descriptor_count  43
zk_max_file_descriptor_count 4096
zk_followers                3
zk_synced_followers        3
zk_pending_syncs            0
```

Zookeeper 컨피그레이션을 확인합니다.

```
root@SpringpathControllerMSH7NHXRFL:/var/log/zookeeper# echo conf | nc localhost 2181
clientPort=2181
dataDir=/var/zookeeper/version-2
dataLogDir=/var/zookeeper/version-2
tickTime=3000
maxClientCnxns=60
minSessionTimeout=6000
maxSessionTimeout=60000
serverId=3
initLimit=10
syncLimit=3
electionAlg=3
electionPort=3888
quorumPort=2888
peerType=0
```

## 라이브 설정의 Zookeeper 로그 파일

Zookeeper 서비스에 문제가 있는 경우 이러한 로그 파일은 추적을 찾는 데 도움이 됩니다.

- /var/log/zookeeper/zookeeper\* - 보관된 로그, 유용한 검색 키워드 WARN, ERROR, Goodbye, Leader 등을 유지 관리합니다.
- /var/log/springpath/zk-\*
- /var/log/springpath/exhibitor.log

```
root@SpringpathControllerMSH7NHXRFL:/var/log/zookeeper# grep -i leader
/var/log/zookeeper/zookeeper.log*
/var/log/zookeeper/zookeeper.log.7:2016-10-14 22:59:26,088 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Leader@60] - TCP NoDelay set to: true
/var/log/zookeeper/zookeeper.log.7:2016-10-14 22:59:26,099 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Leader@358] - LEADING - LEADER ELECTION TOOK - 354
/var/log/zookeeper/zookeeper.log.7:2016-10-14 22:59:26,120 [myid:3] - INFO [LearnerHandler-
/192.168.5.164:36487:LearnerHandler@522] - Received NEWLEADER-ACK message from 0
/var/log/zookeeper/zookeeper.log.7:2016-10-14 22:59:26,120 [myid:3] - INFO [LearnerHandler-
/192.168.5.163:43451:LearnerHandler@522] - Received NEWLEADER-ACK message from 1
/var/log/zookeeper/zookeeper.log.7:2016-10-14 22:59:26,120 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Leader@943] - Have quorum of supporters, sids: [ 0,1,3
]; starting up and setting last processed zxid: 0x100000000
/var/log/zookeeper/zookeeper.log.7:2016-10-14 22:59:26,272 [myid:3] - INFO
[WorkerReceiver[myid=3]:FastLeaderElection@597] - Notification: 1 (message format version), 3
(n.leader), 0x0 (n.zxid), 0x1 (n.round), LOOKING (n.state), 2 (n.sid), 0x0 (n.peerEpoch) LEADING
(my state)
/var/log/zookeeper/zookeeper.log.7:2016-10-14 22:59:26,291 [myid:3] - INFO [LearnerHandler-
/192.168.5.162:48778:LearnerHandler@486] - Sending snapshot last zxid of peer is 0x0 zxid of
leader is 0x100000000sent zxid of db as 0x100000000
/var/log/zookeeper/zookeeper.log.7:2016-10-14 22:59:26,298 [myid:3] - INFO [LearnerHandler-
```

/192.168.5.162:48778:LearnerHandler@522] - Received NEWLEADER-ACK message from 2

```
root@SpringpathControllerMSH7NHXRFL:/var/log/zookeeper# grep -i warn
/var/log/zookeeper/zookeeper.log*
```

```
/var/log/zookeeper/zookeeper.log.7:2016-10-14 22:46:30,354 [myid:] - WARN
[main:QuorumPeerMain@113] - Either no config or no quorum defined in config, running in
standalone mode
/var/log/zookeeper/zookeeper.log.7:2016-10-14 22:52:55,238 [myid:] - WARN
[main:QuorumPeerMain@113] - Either no config or no quorum defined in config, running in
standalone mode
```

```
root@SpringpathControllerMSH7NHXRFL:/var/log/zookeeper# grep -i goodbye
/var/log/zookeeper/zookeeper.log*
```

```
/var/log/zookeeper/zookeeper.log.1:2017-01-23 03:55:50,429 [myid:3] - WARN [LearnerHandler-
/192.168.5.163:44118:LearnerHandler@646] - ***** GOODBYE /192.168.5.163:44118 *****
/var/log/zookeeper/zookeeper.log.1:2017-01-24 23:30:14,956 [myid:3] - WARN [LearnerHandler-
/192.168.5.164:44720:LearnerHandler@646] - ***** GOODBYE /192.168.5.164:44720 *****
/var/log/zookeeper/zookeeper.log.3:2016-12-01 23:45:22,510 [myid:3] - WARN [LearnerHandler-
/192.168.5.164:44051:LearnerHandler@646] - ***** GOODBYE /192.168.5.164:44051 *****
/var/log/zookeeper/zookeeper.log.3:2016-12-08 00:36:37,752 [myid:3] - WARN [LearnerHandler-
/192.168.5.162:46577:LearnerHandler@646] - ***** GOODBYE /192.168.5.162:46577 *****
/var/log/zookeeper/zookeeper.log.4:2016-11-22 23:45:30,957 [myid:3] - WARN [LearnerHandler-
/192.168.5.163:49016:LearnerHandler@646] - ***** GOODBYE /192.168.5.163:49016 *****
/var/log/zookeeper/zookeeper.log.4:2016-11-23 00:03:59,397 [myid:3] - WARN [LearnerHandler-
/192.168.5.164:45952:LearnerHandler@646] - ***** GOODBYE /192.168.5.164:45952 *****
/var/log/zookeeper/zookeeper.log.4:2016-12-01 22:51:00,538 [myid:3] - WARN [LearnerHandler-
/192.168.5.163:45284:LearnerHandler@646] - ***** GOODBYE /192.168.5.163:45284 *****
/var/log/zookeeper/zookeeper.log.5:2016-11-10 23:39:47,477 [myid:3] - WARN [LearnerHandler-
/192.168.5.163:43576:LearnerHandler@646] - ***** GOODBYE /192.168.5.163:43576 *****
/var/log/zookeeper/zookeeper.log.5:2016-11-11 00:49:39,782 [myid:3] - WARN [LearnerHandler-
/192.168.5.164:35219:LearnerHandler@646] - ***** GOODBYE /192.168.5.164:35219 *****
```

## 일부 샘플 로그 - Zookeeper 로깅 선택

```
2017-01-22 23:47:29,427 [myid:3] - INFO [Thread-2:QuorumCnxManager$Listener@504] - My election
bind port: /192.168.5.161:3888
2017-01-22 23:47:29,435 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:2181:QuorumPeer@714] - LOOKING
2017-01-22 23:47:29,438 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:2181:FastLeaderElection@815] - New election. My id = 3,
proposed zxid=0x9000a6b4d
2017-01-22 23:47:29,443 [myid:3] - INFO [WorkerReceiver[myid=3]:FastLeaderElection@597] -
Notification: 1 (message format version), 2 (n.leader), 0x800055ea0 (n.zxid), 0x1 (n.round),
FOLLOWING (n.state), 0 (n.sid), 0x9 (n.peerEpoch) LOOKING (my state)
2017-01-22 23:47:29,444 [myid:3] - INFO [WorkerReceiver[myid=3]:FastLeaderElection@597] -
Notification: 1 (message format version), 2 (n.leader), 0x800055ea0 (n.zxid), 0x1 (n.round),
FOLLOWING (n.state), 1 (n.sid), 0x9 (n.peerEpoch) LOOKING (my state)
2017-01-22 23:47:29,444 [myid:3] - INFO [WorkerReceiver[myid=3]:FastLeaderElection@597] -
Notification: 1 (message format version), 3 (n.leader), 0x9000a6b4d (n.zxid), 0x1 (n.round),
LOOKING (n.state), 3 (n.sid), 0x9 (n.peerEpoch) LOOKING (my state)
2017-01-22 23:47:29,444 [myid:3] - INFO [WorkerReceiver[myid=3]:FastLeaderElection@597] -
Notification: 1 (message format version), 2 (n.leader), 0x800055ea0 (n.zxid), 0x1 (n.round),
FOLLOWING (n.state), 1 (n.sid), 0x9 (n.peerEpoch) LOOKING (my state)
2017-01-22 23:47:29,445 [myid:3] - INFO [WorkerReceiver[myid=3]:FastLeaderElection@597] -
Notification: 1 (message format version), 2 (n.leader), 0x800055ea0 (n.zxid), 0x1 (n.round),
LEADING (n.state), 2 (n.sid), 0x9 (n.peerEpoch) LOOKING (my state)
2017-01-22 23:47:29,445 [myid:3] - INFO [WorkerReceiver[myid=3]:FastLeaderElection@597] -
Notification: 1 (message format version), 2 (n.leader), 0x800055ea0 (n.zxid), 0x1 (n.round),
FOLLOWING (n.state), 0 (n.sid), 0x9 (n.peerEpoch) LOOKING (my state)
2017-01-22 23:47:29,446 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:2181:QuorumPeer@784] - FOLLOWING
```

```
2017-01-22 23:47:29,449 [myid:3] - INFO [QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Learner@86] -
TCP NoDelay set to: true
2017-01-22 23:47:29,449 [myid:3] - INFO [WorkerReceiver[myid=3]:FastLeaderElection@597] -
Notification: 1 (message format version), 2 (n.leader), 0x800055ea0 (n.zxid), 0x1 (n.round),
LEADING (n.state), 2 (n.sid), 0x9 (n.peerEpoch) FOLLOWING (my state)
2017-01-22 23:47:29,660 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Environment@100] - Server
environment:zookeeper.version=3.4.6--1, built on 06/16/2015 22:50 GMT
2017-01-22 23:47:29,661 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Environment@100] - Server
environment:host.name=SpringpathControllerMSH7NHXRFL
2017-01-22 23:47:29,661 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Environment@100] - Server
environment:java.version=1.7.0_79
2017-01-22 23:47:29,661 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Environment@100] - Server
environment:java.vendor=Oracle Corporation
2017-01-22 23:47:29,661 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Environment@100] - Server
environment:java.home=/usr/lib/jvm/java-7-openjdk-amd64/jre
2017-01-22 23:47:29,661 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Environment@100] - Server
environment:java.class.path=/usr/share/zookeeper/bin/./build/classes:/usr/share/zookeeper/bin/./
./build/lib/*.jar:/usr/share/zookeeper/bin/./lib/slf4j-log4j12-
1.6.1.jar:/usr/share/zookeeper/bin/./lib/slf4j-api-
1.6.1.jar:/usr/share/zookeeper/bin/./lib/netty-
3.7.0.Final.jar:/usr/share/zookeeper/bin/./lib/log4j-
1.2.16.jar:/usr/share/zookeeper/bin/./lib/jline-
0.9.94.jar:/usr/share/zookeeper/bin/./zookeeper-
3.4.6.jar:/usr/share/zookeeper/bin/./src/java/lib/*.jar:/usr/share/zookeeper/bin/./conf:
2017-01-22 23:47:29,661 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Environment@100] - Server
environment:java.library.path=/usr/java/packages/lib/amd64:/usr/lib/x86_64-linux-
gnu/jni:/lib/x86_64-linux-gnu:/usr/lib/x86_64-linux-gnu:/usr/lib/jni:/lib:/usr/lib
2017-01-22 23:47:29,661 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Environment@100] - Server
environment:java.io.tmpdir=/tmp
2017-01-22 23:47:29,661 [myid:3] - INFO
[QuorumPeer[myid=3]/0:0:0:0:0:0:0:0:2181:Environment@100] - Server environment:java.compiler=
```

## LEADER ELECTION TOOK

```
root@SpringpathControllerMSH7NHXRFL:/var/log/springpath# cat zk-debug-storfs.log
```

```
2017-01-22 23:47:18,702:5866(0x7fd1f7ef5700):ZOO_INFO@check_events@1760: initiated connection to
server [192.168.5.163:2181]
2017-01-22 23:47:18,704:5866(0x7fd1f7ef5700):ZOO_INFO@check_events@1807: session establishment
complete on server [192.168.5.163:2181], sessionId=0x159165ff6310005, negotiated timeout=17001
2017-01-22 23:47:18,704:5866(0x7fd1f76f4700):ZOO_INFO@process_completions@2170: Calling a
watcher for node s], type = s
2017-01-23 01:50:16,809:5866(0x7fd1f7ef5700):ZOO_ERROR@handle_socket_error_msg@1778: Socket
[192.168.5.163:2181] zk retcode=-4, errno=112(Host is down): failed while receiving a server
response
2017-01-23 01:50:16,818:5866(0x7fd1f76f4700):ZOO_INFO@process_completions@2170: Calling a
watcher for node s], type = s
2017-01-23 01:50:16,818:5866(0x7fd1f7ef5700):ZOO_INFO@check_events@1760: initiated connection to
server [192.168.5.164:2181]
2017-01-23 01:50:16,818:5866(0x7fd1f7ef5700):ZOO_ERROR@handle_socket_error_msg@1778: Socket
```

```
[192.168.5.164:2181] zk retcode=-4, errno=112(Host is down): failed while receiving a server response
2017-01-23 01:50:17,819:5866(0x7fd1f7ef5700):ZOO_ERROR@handle_socket_error_msg@1740: Socket
[192.168.5.162:2181] zk retcode=-4, errno=115(Operation now in progress): poll refused to accept read/write from the client
```

```
root@help:/var/log/springpath# cat zkEvents.log
```

```
INFO:ZkEvents:Send changes to listeners
INFO:EventDB:Received message{"timestamp": 1559200009008, "description": "Cluster policy compliance is satisfied", "id": "ClusterPolicyComplianceSatisfiedEvent"}
DEBUG:kazoo.client:Received EVENT: Watch(type=3, state=3, path=u'/zkEvents/lastModificationTime')
DEBUG:kazoo.client:Sending request(xid=42): GetData(path='/zkEvents/lastModificationTime', watcher=
```

**Cluster is healthy**

```
root@SpringpathControllerPZTMTRSH7K:/var/log/springpath# tail exhibitor.log
```

```
05-20 05:28:52.223 INFO org.mortbay.log - Started SocketConnector@10.197.252.99:8180
05-20 05:29:20.106 INFO com.netflix.exhibitor.core.activity.ActivityLog - State: down
05-20 05:29:20.106 INFO com.netflix.exhibitor.core.activity.ActivityLog - Attempting to stop instance
05-20 05:29:20.106 INFO com.netflix.exhibitor.core.activity.ActivityLog - Attempting to start/restart ZooKeeper
05-20 05:29:20.328 INFO com.netflix.exhibitor.core.activity.ActivityLog - jps didn't find instance - assuming ZK is not running
05-20 05:29:20.347 INFO com.netflix.exhibitor.core.activity.ActivityLog - Process started via: /usr/share/zookeeper/bin/zkServer.sh
05-20 05:29:20.353 ERROR com.netflix.exhibitor.core.activity.ActivityLog - ZooKeeper Server: ZooKeeper JMX enabled by default
05-20 05:29:20.353 ERROR com.netflix.exhibitor.core.activity.ActivityLog - ZooKeeper Server: Using config: /usr/share/zookeeper/bin/./conf/zoo.cfg
05-20 05:29:21.366 INFO com.netflix.exhibitor.core.activity.ActivityLog - ZooKeeper Server: Starting zookeeper ... STARTED
05-20 05:29:50.128 INFO com.netflix.exhibitor.core.activity.ActivityLog - State: serving
```

## 지원 번들(storfs)의 Zookeeper 로그 파일

지원 번들에서는 다음 사항을 살펴봐야 합니다.

zookeeper.log	/var/log/zookeeper
zk-storfs.log	/var/log/springpath
echo_stat_ nc_localhost_2181.out	under cmds_output