



NexaVM Operational Adjustment Plan

nSSV + nSAN Management + Storage Network Change

Risk Level: **High**

Contents

1 Change Overview	3
2 Modify Step.....	3
2.1 Current Network.....	3
2.2 Change Principle and Scope	3
2.3 Change Overall Process.....	3
2.4 Modify Step.....	4
2.4.1 Modify Management IP on Cloud	4
2.4.2 Modify Storage Management IP Address	8
2.4.3 Modify the Cluster Network.....	11
2.4.4 Modify Public Network.....	12
2.4.5 Login.....	14
3 Risk Assessment	15
3.1 Vulnerability List	15
3.2 Risk Mitigation Plan	18
4. Reversion Plan.....	19
5 Emergency Response Plan	19

1 Change Overview

This document describes the modification plan for the nSSV+NSAN hyper-converged management network and Storage network. It includes change steps, risk assessment, rollback plans, emergency response plans, change schedules, and other specific information involved in the maintenance changes.

Phase 2 of the Sensory Source Construction Project for XXX - Due to the customer's network requirements, it is necessary to change the server management network and storage network involving modifications to the dual management network environment of three nodes in Cloud and NSAN environments.

2 Modify Step

2.1 Current Network

Software Information	
Business Information	Dual Management Node + Bulk Hyper-converged Deployment

2.2 Change Principle and Scope

As the global modification involves changes to the management network of nSAN distributed storage, we need to communicate with the customer to pause business operations and perform the modification offline. During the modification, the management network IP addresses of three servers will be modified, and related service configurations will also be changed.

2.3 Change Overall Process

1. Disable Global HA, backup database, business VMs are powered off, hosts and primary storages are placed into maintenance mode

2. Modify IP address for all hosts
3. Double management modifies IP address and changes time server
4. Modify IP, Primary Storage Mon Address, and Image Storage Mon Address on the Cloud Platform UI
5. Modify the management IP for the storage side
6. Modify Cluster network of storage side
7. Modify Public Network of Storage side
6. Check the status of the site environment and confirm that the resource connection is normal.

2.4 Modify Step

2.4.1 Modify Management IP

2.4.1.1 **Disabling Global HA, backup databases, business VMs are powered off, hosts and primary storages enter maintenance mode, and dual management services are stopped.**

2.4.1.2 Modify IP Address for All Hosts

```
zs-network-setting -i br_bond0 172.28.21.38 255.255.248.0 172.28.16.1 [Management Network]
```

```
zs-network-setting -i br_bond1 172.38.20.30 255.255.0.0 [storage network]
```

```
Last login: Tue Oct 22 12:41:52 from 172.28.3.10
[root@zstack-Z8 ~]# zs-network-setting -i br_bond0 172.28.21.38 255.255.248.0 172.28.16.1
RTNETLINK answers: No such process
Network Setting Successfully! Set Default Gateway 172.28.16.1.
[root@zstack-Z8 ~]# _
```

2.4.1.3 Dual-Management Adjustment of IP Address and Time Server

```
zsha2 change-ip -myip 172.28.21.37 --slave "root:password@172.28.21.36" -nic br_bond0 -gateway 172.28.16.1 -vip 172.28.21.80 --db-root-pw NexaVM@123
```

Comment:

br_bond0 identifies the NIC name of the management node where the management IP address is located

root:password@172.18.100.180 indicates the SSH information of a slave management node, including the username and password.

172.18.0.1 represents the gateway for management node arbitration

172.18.100.180 is a VIP IP address

The password for the root account of the MySQL database, `zstack.mysql.password` by default. Note that this operation will overwrite the database of the slave node. Make sure that you execute this command on the original VIP-located node. The MN service will be auto-started after modification. After the modification, check whether you need to modify the time server in the `NexaVM.properties` configuration file.

```
[root@zstack-27 ~]# zsha2 change-ip -mysql 172.28.21.37 --slave "root:password@172.28.21.36" -nic br_bond0 -gateway 172.28.16.1 -vip 172.28.21.80 --db-root-pw zstack@123
✓ checking network interface and gateway ...
✓ Task 1: checking network interface and gateway ... completed, spend: 2.13 s.
✓ setting up DB config before replication ...
✓ Task 2: setting up DB config before replication ... completed, spend: 3.76 s.
✓ creating DB user for replication ...
✓ Task 3: creating DB user for replication ... completed, spend: 2.49 s.
✓ update iptables rules ...
✓ Task 4: update iptables rules ... completed, spend: 0.27 s.
✓ starting the initial replication ...
*****
File: mariadb-bin.000018
Position: 2208
Binlog_Do_DB:
Binlog_Ignore_DB:
✓ starting the initial replication ...
✓ starting the initial replication ...
✓ starting the initial replication ...
P Starting Initial Replication ...
Local database backed up to ...
✓ starting the initial replication ...
✓ Task 5: starting the initial replication ... completed, spend: 351.33 s.
✓ wait peer slave sync status ...
|
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
Last_IO_Error:
Last_SQL_Error:
Last_Error:
Last_Error: ()
✓ Task 6: wait peer slave sync status ... completed, spend: 0.53 s.
✓ wait local DB sync status ...
*****
File: mariadb-bin.000026
Position: 360
Binlog_Do_DB:
Binlog_Ignore_DB:
✓ wait local DB sync status ...
|
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
Last_IO_Error:
Last_SQL_Error:
Last_Error:
Last_Errno: 0
✓ Task 7: wait local DB sync status ... completed, spend: 0.59 s.
✓ Re-generating keepalived config ...
```

Modify and Check Service Status `zsha2 status`

```

Status report from 172.28.21.36
=====
Owns virtual address:      no
Self 172.28.21.36 reachable: yes
Gateway 172.28.16.1 reachable: yes
VIP 172.28.21.80 reachable: yes
Peer 172.28.21.37 reachable: yes
Keepalived status:        active
ZStack HA Monitor:        active
MySQL status:             mysqld is alive
MN status: Running [PID:931529]
UI status: Running [PID:953957] https://172.28.21.36:5443

Slave Status:
-----
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
Last_Error:
Seconds_Behind_Master: 0
Last_IO_Error:
Last_SQL_Error:
Slave_SQL_Running_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Status report from 172.28.21.37
=====
Owns virtual address:      yes
Self 172.28.21.37 reachable: yes
Gateway 172.28.16.1 reachable: yes
VIP 172.28.21.80 reachable: yes
Peer 172.28.21.36 reachable: yes
Keepalived status:        active
ZStack HA Monitor:        active
MySQL status:             mysqld is alive
MN status: Running [PID:1442251]
UI status: Running [PID:1458607] https://172.28.21.37:5443

Slave Status:
-----
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
Last_Error:
Seconds_Behind_Master: 0
Last_IO_Error:
Last_SQL_Error:
Slave_SQL_Running_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Note: visit ZStack UI with https://172.28.21.80:5443

```

Modify Time Server

Check the current time server by running the command "`chronyc sources -v`".

```

[root@zstack-27 ~]# chronyc sources -v

.-- Source mode  '^' = server, '=' = peer, '#' = local clock.
/ -- Source state '*' = current best, '+' = combined, '-' = not combined,
| /              'x' = may be in error, '~' = too variable, '?' = unusable.

Reachability register (octal) --.
Log2(Polling interval) --.

MS Name/IP address          Stratum Poll Reach LastRx Last sample
=====
~? 172.28.21.26              0      8      0      -      +0ns[ +0ns] +/- 0ns
[root@zstack-27 ~]#

```

Modify Time Server Settings and Stop Management Service

```
zsha2 stop-node
```

```
vim /usr/local/NexaVM/apache-tomcat/webapps/NexaVM/WEB-INF/classes/NexaVM.properties
```

```

# for Spring Rest Clients
http.keepAlive = true
http.maxConnections = 1024

#Prometheus.versionMode = 1.8.2 / 2.x / 2.x-compatible / none
Prometheus.versionMode = 2.x
# support v1 / v2
InfluxDB.metadata.version = v2
Ansible.var.zstack_repo = "zstack-mn,qemu-kvm-ev-mn"
Ansible.executable = python2
ui_mode = zstack
management.server.ip = 172.28.21.37
RepoVersion.Strategy = permissive
chrony.serverIp.0 = 172.28.21.26
beanConf = zstack.xml
InfluxDB.writeProxy.on = true
Prometheus.useClusterRead = true
C3p0ConnectionTester.isValidTimeout = 5
ThreadFacade.maxThreadNum = 500
DbFacadeDataSource.maxPoolSize = 150
management.server.vip = 172.28.21.80

#ZStack MN network
#MN.network.0=172.20.16.250/32
#MN.network.1=10.86.4.0/23

```

After adjusting two nodes, start service

```
zsha2 start-node
```

Check Time Synchronization Status

```

starting zstack management node ...
[root@zstack-27 ~]# chronyc sources -v

  -- Source mode  '^' = server, '=' = peer, '#' = local clock.
 / -- Source state '*' = current best, '+' = combined, '-' = not combined,
 | /              'x' = may be in error, '~' = too variable, '?' = unusable.

  Reachability register (octal) -.
  Log2(Polling interval) --.

  xxx[yyyy] +/- zzzz
  xxxx = adjusted offset,
  yyyy = measured offset,
  zzzz = estimated error.

MS Name/IP address         Stratum Poll Reach LastRx Last sample
=====
^* 172.28.21.36             10      6   377    30  +3663ns[+8552ns] +/- 109us
[root@zstack-27 ~]#

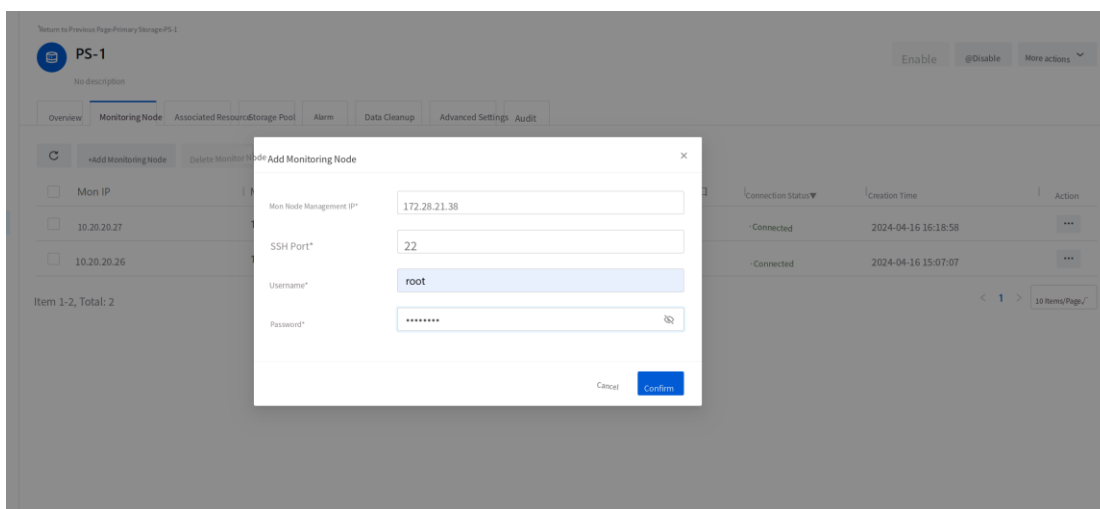
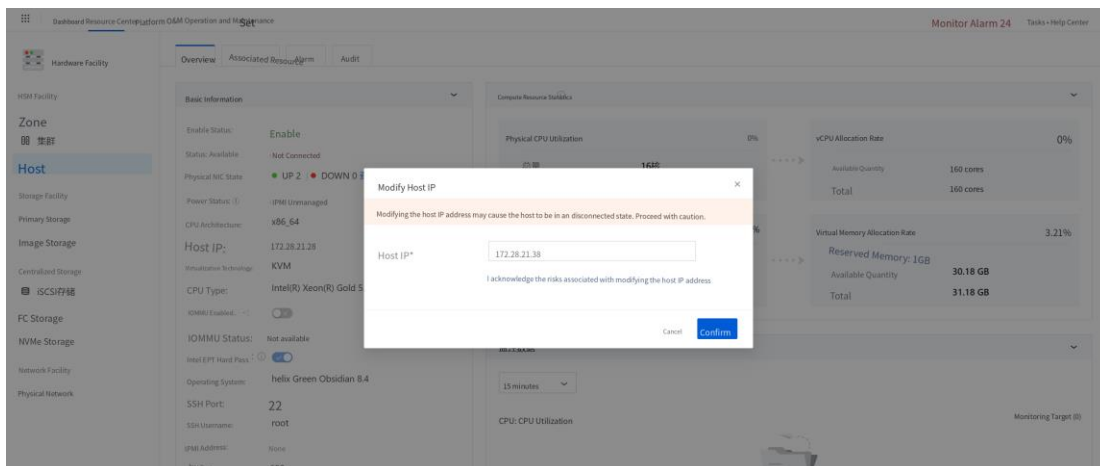
```

To modify a backup job, log into both management nodes of the corresponding cluster and adjust the job as needed. You can execute the `crontab -l` command to view the scheduled jobs.

2.4.1.4 Modify the IP address of a host, primary storage, or image storage MON on the UI

- 1) Edit Host IP on Hardware Facility-Host-Details page
- 2) On Hardware Facility - Primary Storage - Monitoring Node, delete the monitoring node and then re-add it.
- 3) On Hardware Facility-Image Storage-Monitoring Node, delete monitoring nodes and add them again.
- 4) If there are other roles that reuse the host, such as the migration server and backup server, modify these roles as needed.

Note that restarting the MN service due to bulk modifications may cause resource unable to connect. We recommend you to modify one after another, modify the compute node first and then the management node. This ensures that the primary storage and image storage attached to the MN service are in connected status. Batch modification can be achieved if the modification is made offline.



2.4.2 Modify Storage Management IP Address

2.4.2.1 Node Filter Configuration File IP Address

```
grep -r "172.28.21.26|172.28.21.27|172.28.21.50" /opt/nSAN/conf/* /opt/nSAN/watch/*
/opt/nSAN/tools/* /opt/nSAN/ui/* |grep -v log # Screen addresses that are dual management
address and VIP address
```

```
[root@zstack-26 ~]# grep -r "172.28.21.26" /opt/zstone/conf/* /opt/zstone/watch/* /opt/zstone/tools/* /opt/zstone/ui/* |grep -v log
/opt/zstone/conf/zstnmon.yml: Ips: 172.28.21.26,172.28.21.27
/opt/zstone/conf/zstone.yml: Ips: 172.28.21.26,172.28.21.27
/opt/zstone/tools/prometheus/prometheus.yml:
Binary file /opt/zstone/tools/prometheus/data2/wal/checkpoint.00001962/00000000 matches
Binary file /opt/zstone/tools/prometheus/data2/01HX6RZQJ1Y3TRRD9PRHDXR7S/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J7QEF44W84AVC3R6DSZKDV3/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAR0B6GGAPAJHX4PTAF1E6CR/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HXK69KWCMBYPYB3046PVXH/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J9XTHKXC20RJNXHT1KF3FYWF/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAM4RBAWFXN26EEVZ4V5CS13/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAS2NMV8785C2Q2KBEHF15CS/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J1R8FC0A8HKTK6P2RFV626X/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J29MMP2DY6HY4Z9XT81ZM7S/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J2V0TJMZ5PPTCHA406NWMWY51/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J3XS61GBC1YQMRWXH4TJE2WR/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J6MP2XR32HDW5CQ37KQHEKWM/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HYFZDDVMDYCSNKKYCCAB76DV/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HZJQ05F0T31FGT0EXGVC1EFH/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J88TMZ6FJ6X8JZTW9VF0ZZF/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAEBCFFWXX13NTQ876TPDEZ/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J4F5B6Z5J5QMSJRM7BYNTQ/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAS9HEHV6D1SKJWRHMKVC35H/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J3CD0JK88RA0NQC658FEV0ZS/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J639WZK506E5TT3IWM1AHF7T/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAP2HKFGZ00RARCYPDHVW4V/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HWAEMK3EN62JR04QYVDDCRZ/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HXD6ZTY9W8C56D7GARWGYA2/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JARMY7EH7Q00YF6QPT519F9/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J0NG3SX59VMDW5D0T4RB480/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J16W8V4WV2SDJYTTJR4WD5S/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J50HH7ETA0Y5GVRVYZ5872XV/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAS9HC38ANPXZHTXG658DE22/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HZ1BJG4HR0RCM6HYAXGXDW4/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J043YHN34VYSY9W3XFVWKN9/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J5HXQHZCVK570MSFH9EQ3A7/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J7629MDGSRYPMRK60CW6W7RP/index matches
/opt/zstone/tools/temporal/temporal_home/config/zstone-frontend.yaml: connectAddr: "172.28.21.26:5432"
/opt/zstone/tools/temporal/temporal_home/config/zstone-frontend.yaml: connectAddr: "172.28.21.26:5432"
/opt/zstone/tools/temporal/temporal_home/config/zstone-history.yaml: connectAddr: "172.28.21.26:5432"
/opt/zstone/tools/temporal/temporal_home/config/zstone-history.yaml: connectAddr: "172.28.21.26:5432"
/opt/zstone/tools/temporal/temporal_home/config/zstone-matching.yaml: connectAddr: "172.28.21.26:5432"
/opt/zstone/tools/temporal/temporal_home/config/zstone-matching.yaml: connectAddr: "172.28.21.26:5432"
/opt/zstone/tools/temporal/temporal_home/config/zstone-worker.yaml: connectAddr: "172.28.21.26:5432"
/opt/zstone/tools/temporal/temporal_home/config/zstone-worker.yaml: connectAddr: "172.28.21.26:5432"
```

Modify the correct address according to the found file

After modifications, search here to ensure nothing is left unmodified

```
[root@zstack-26 ~]# grep -r "172.28.21.26" /opt/zstone/conf/* /opt/zstone/watch/* /opt/zstone/tools/* /opt/zstone/ui/* |grep -v log
Binary file /opt/zstone/tools/prometheus/data2/wal/checkpoint.00001962/00000000 matches
Binary file /opt/zstone/tools/prometheus/data2/01HX6RZQJ1Y3TRRD9PRHDXR7S/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J7QEF44W84AVC3R6DSZKDV3/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAR0B6GGAPAJHX4PTAF1E6CR/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HXK69KWCMBYPYB3046PVXH/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J9XTHKXC20RJNXHT1KF3FYWF/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAM4RBAWFXN26EEVZ4V5CS13/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAS2NMV8785C2Q2KBEHF15CS/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J1R8FC0A8HKTK6P2RFV626X/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J29MMP2DY6HY4Z9XT81ZM7S/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J2V0TJMZ5PPTCHA406NWMWY51/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J3XS61GBC1YQMRWXH4TJE2WR/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J6MP2XR32HDW5CQ37KQHEKWM/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HYFZDDVMDYCSNKKYCCAB76DV/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HZJQ05F0T31FGT0EXGVC1EFH/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J88TMZ6FJ6X8JZTW9VF0ZZF/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAEBCFFWXX13NTQ876TPDEZ/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J4F5B6Z5J5QMSJRM7BYNTQ/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAS9HEHV6D1SKJWRHMKVC35H/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J3CD0JK88RA0NQC658FEV0ZS/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J639WZK506E5TT3IWM1AHF7T/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAP2HKFGZ00RARCYPDHVW4V/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HWAEMK3EN62JR04QYVDDCRZ/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HXD6ZTY9W8C56D7GARWGYA2/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JARMY7EH7Q00YF6QPT519F9/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J0NG3SX59VMDW5D0T4RB480/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J16W8V4WV2SDJYTTJR4WD5S/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J50HH7ETA0Y5GVRVYZ5872XV/index matches
Binary file /opt/zstone/tools/prometheus/data2/01JAS9HC38ANPXZHTXG658DE22/index matches
Binary file /opt/zstone/tools/prometheus/data2/01HZ1BJG4HR0RCM6HYAXGXDW4/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J043YHN34VYSY9W3XFVWKN9/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J5HXQHZCVK570MSFH9EQ3A7/index matches
Binary file /opt/zstone/tools/prometheus/data2/01J7629MDGSRYPMRK60CW6W7RP/index matches
```

2.4.2.4 Modify Database

```
VIP Node Enters Database
# mysql -u root -p zstack.mysql.password zstone

Modify clusterVO data

# Update ClusterVO set managementNetworkCidr='51.59.130.0/24' where
managementNetworkCidr='51.59.129.0/24'; # The former network range is the new network
```

range, and the latter is the existing network range in the database

Modify hostVO data

```
update HostVO set managementIp='172.28.21.36' where managementIp='172.28.21.26'; # The first IP is the
newP, and the second IP is the original old IP in the database
update HostVO set managementIp='172.28.21.37' where managementIp='172.28.21.27'; # The first IP is the
newP, and the second IP is the original old IP in the database
update HostVO set managementIp='172.28.21.38' where managementIp='172.28.21.28'; # The first IP is the
newP, and the second IP is the original old IP in the database
```

```
MariaDB [zstone]> update HostVO set managementIp='172.28.21.36' where managementIp='172.28.21.26';
Query OK, 1 row affected (0.043 sec)
Rows matched: 1 Changed: 1 Warnings: 0

MariaDB [zstone]> update HostVO set managementIp='172.28.21.37' where managementIp='172.28.21.27';
Query OK, 1 row affected (0.010 sec)
Rows matched: 1 Changed: 1 Warnings: 0

MariaDB [zstone]> update HostVO set managementIp='172.28.21.38' where managementIp='172.28.21.28';
Query OK, 1 row affected (0.011 sec)
Rows matched: 1 Changed: 1 Warnings: 0
```

Modify licenseAppidVO

```
update LicenseAppIdVO set hostIp='172.28.21.36' where hostIp='172.28.21.26'; # The
former is the new IP and the latter is the existing IP in the database
update LicenseAppIdVO set hostIp='172.28.21.37' where hostIp='172.28.21.27'; # The
former is the new IP and the latter is the existing IP in the database
update LicenseAppIdVO set hostIp='172.28.21.38' where hostIp='172.28.21.28'; # The
former is the new IP and the latter is the existing IP in the database
```

2.4.2.5 Start Related Services of Temporal

```
systemctl start temporal-history
systemctl start temporal-matching
systemctl start temporal-frontend
systemctl start temporal-worker
```

```
[root@zstack-27 log]# systemctl start temporal-history
[root@zstack-27 log]# systemctl start temporal-matching
[root@zstack-27 log]# systemctl start temporal-frontend
[root@zstack-27 log]# systemctl start temporal-worker
```

2.4.2.6 Start zstnlet Service

```
for i in zstack-26 zstack-27 zstack-28; do ssh $i systemctl restart zstnlet; done
```

```
[root@zstack-26 ~]# for i in zstack-26 zstack-27 zstack-28; do ssh $i systemctl stop zstnlet; done
[root@zstack-26 ~]#
```

2.4.2.7 Start Dual Management nSAN Services

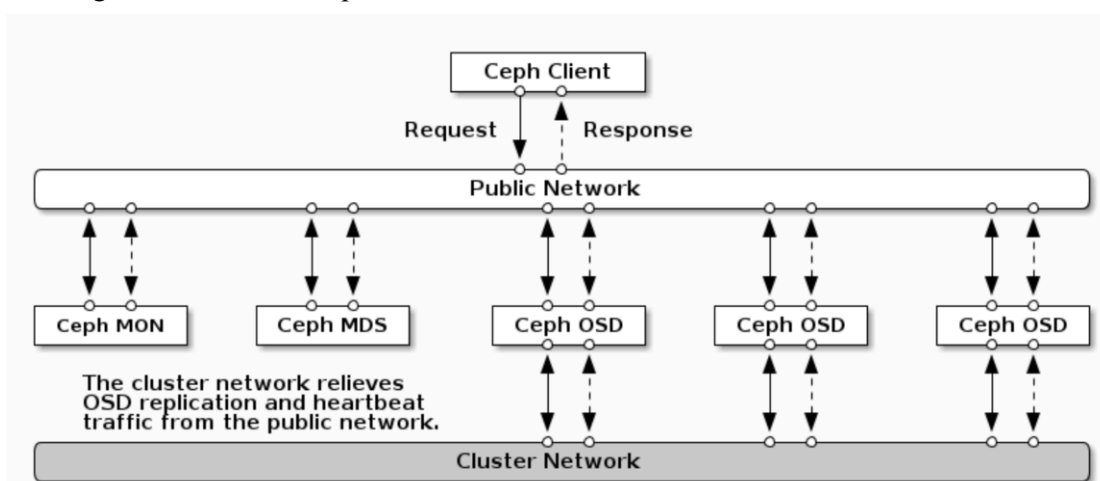
```

systemctl restart zstnmon
systemctl restart zstone-master
systemctl restart zstnwatch
systemctl restart zstone-pushgateway
systemctl restart zstone-alertmanager
systemctl restart zstone-prometheus
systemctl restart zstone-nodeExporter
systemctl restart thanos-query
systemctl restart zstone-rgwUsageExporter # If RGW exists

```

2.4.3 Modify the Cluster Network

The cluster network is primarily responsible for data replication (master-slave copying) and data migration within the Ceph cluster.



2.4.3.1 Steps

Modify the ceph.conf configuration file (only modify the commented-out parts below)

```

[global]
fsid = 2690bc72-08ea-11ee-93df-fac33c2bcc00
public_network = 172.31.10.0/24
cluster_network = 10.0.0.0/24 # Modify to new network"
mon_initial_members = node-88 node-89 node-90
mon_host = node-88 node-89 node-90
auth_cluster_required = none
auth_service_required = none
auth_client_required = none
#osd_memory_target = 4294967296 #4GB

```

...

Copy the configuration to all Ceph nodes

```
for i in node-1 node-2 node-3; do scp /etc/ceph/ceph.conf ${i}:/etc/ceph/ceph.conf
```

Restarting ceph.target (will interrupt all I/O)

```
for i in node-1 node-2 node-3; do systemctl restart ceph.target; done
```

Access the nSAN database in MySQL

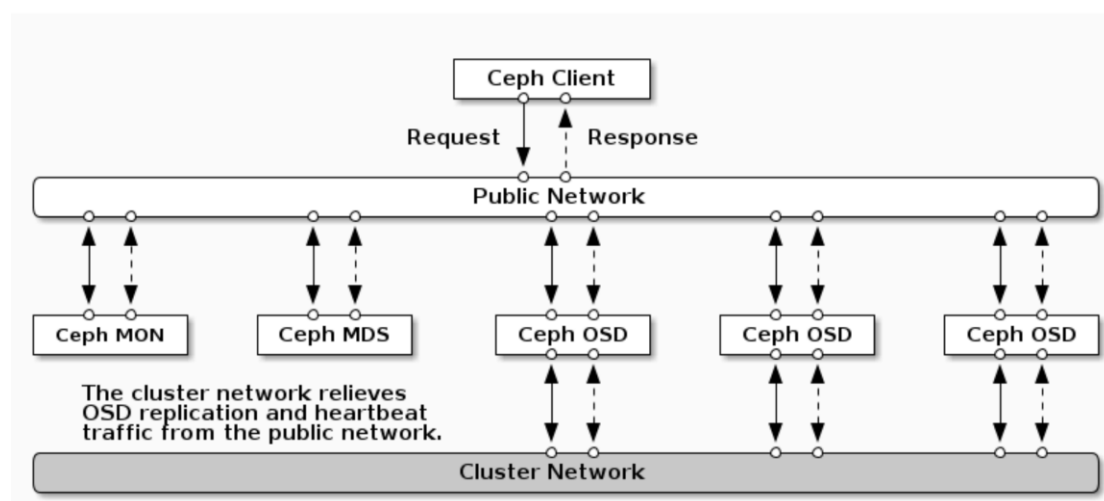
```
mysql -u root -p zstack.mysql.password zstone #password reference /opt/zstone/conf/zstone.yml
```

Modify ClusterVO and HostVO tables

```
update ClusterVO set clusterNetworkCidr='10.0.0.0/24'; # modify to new subnet
update HostVO set clusterIP='10.0.0.88' where clusterIP='172.31.10.88'; # The first IP is the newP, and the second
IP is the original old IP in the database
update HostVO set clusterIP='10.0.0.89' where clusterIP='172.31.10.89'; # The first IP is the newP, and the second
IP is the original old IP in the database
update HostVO set clusterIP='10.0.0.90' where clusterIP='172.31.10.90'; # The first IP is the newP, and the second
IP is the original old IP in the database
```

2.4.4 Modify Public Network

The public network is mainly responsible for the Ceph cluster's external communication, I/O, and heartbeats between components;



2.4.4.1 Steps

Check if /tmp/monmap exists (it exists in versions 430 and later, but not in earlier versions).

```
ll /tmp/monmap
```

If it does not exist, generate a monmap (for generality, the following command line can be used regardless of whether the current mon is available)

```
monmaptool --create --generate -c /etc/ceph/ceph.conf /tmp/monmap
```

View information in monmap

```
[root@node-89 ~]# monmaptool --print /tmp/monmap
```

```
monmaptool: monmap file /tmp/monmap
```

```
epoch 0
```

```

fsid 2690bc72-08ea-11ee-93df-fac33c2bcc00
last_changed 2023-06-12T14:42:02.070600+0800
created 2023-06-12T14:42:02.070600+0800
min_mon_release 0 (unknown)
0: v1:172.31.10.88:6789/0 mon.noname-a
1: v1:172.31.10.89:6789/0 mon.noname-b
2: v1:172.31.10.90:6789/0 mon.noname-c

```

Delete all MON

```
monmaptool --rm noname-a --rm noname-b --rm noname-c /tmp/monmap
```

Check Again

```

[root@node-89 ~]# monmaptool --print /tmp/monmap
monmaptool: monmap file /tmp/monmap
epoch 0
min_mon_release 0 (unknown)

```

Add mon according to the new IP

```
monmaptool --add node-88 172.31.10.88 --add node-89 172.31.10.89 --add node-90 172.31.10.90 /tmp/monmap
```

Check again

```

[root@node-89 ~]# monmaptool --print /tmp/monmap
monmaptool: monmap file /tmp/monmap
epoch 0
fsid 2690bc72-08ea-11ee-93df-fac33c2bcc00
last_changed 2023-06-12T14:42:02.070600+0800
created 2023-06-12T14:42:02.070600+0800
min_mon_release 0 (unknown)
0: v1:172.31.10.88:6789/0 mon.node-88
1: v1:172.31.10.89:6789/0 mon.node-89
2: v1:172.31.10.90:6789/0 mon.node-90

```

Modify the ceph.conf configuration file (only modify the commented-out parts below).

```

[root@node-89 ~]# cat /etc/ceph/ceph.conf
[global]
fsid = 2690bc72-08ea-11ee-93df-fac33c2bcc00
public_network = 172.31.10.0/24 # Modify to new network"
cluster_network = 10.0.0.0/24
mon_initial_members = node-88 node-89 node-90
mon_host = node-88 node-89 node-90
auth_cluster_required = none
auth_service_required = none
auth_client_required = none
#osd_memory_target = 4294967296 #4GB
...

```

Modify the /etc/hosts file (this file needs to be modified on all nodes in the cluster, and changed to the new IP address).

```
[root@node-88 tools]# cat /etc/hosts
```

```
127.0.0.1 localhost
127.0.0.1 localhost
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
172.31.10.88 node-88 # This needs to be changed tto the new address
172.31.10.89 node-89 # This needs to be changed tto the new address
172.31.10.90 node-90 # This needs to be changed tto the new address
```

Stop the Ceph service on all nodes (this will interrupt all I/O).

```
for i in node-88 node-89 node-90; do ssh $i systemctl stop ceph.target; done
```

Inject monmap into all mon nodes (execute on each mon node separately)

```
for i in node-88 node-89 node-90; do ssh $i ceph-mon -i $i --inject-monmap /tmp/monmap; done
```

Modify permissions of the mon directory

```
for i in node-88 node-89 node-90; do ssh $i chown -HR ceph:ceph /var/lib/ceph/mon/ceph-$i; done
```

Start Ceph service

```
for i in node-88 node-89 node-90; do ssh $i systemctl start ceph.target; done
```

Access the nSAN database in MySQL

```
mysql -u root -p zstack.mysql.password zstone #password reference /opt/zstone/conf/zstone.yml
```

Modify ClusterVO and HostVO tables

```
update ClusterVO set publicNetworkCidr='172.31.10.0/24'; # Modify to a new network segment
update HostVO set publicIP='172.31.10.88' where publicIP='10.0.0.88'; # The first IP is the newP, and the second
IP is the original old IP in the database
update HostVO set publicIP='172.31.10.89' where publicIP='10.0.0.89'; # The first IP is the newP, and the second
IP is the original old IP in the database
update HostVO set publicIP='172.31.10.90' where publicIP='10.0.0.90'; # The first IP is the newP, and the second
IP is the original old IP in the database
```

2.4.5 Login

Log into the storage UI to check whether related services work as expected

```
[root@zstack-28 ~]# ceph -s
cluster:
  id:          5d19d858-fb09-11ee-8832-fa20799d3b00
  health: HEALTH_OK

services:
  mon: 3 daemons, quorum zstack-26,zstack-27,zstack-28 (age 6d)
  mgr: zstack-28(active, since 10d), standbys: zstack-26, zstack-27
  osd: 12 osds: 12 up (since 10d), 12 in (since 3M)
  rgw: 3 daemons active (10.20.20.26-7480, 10.20.20.27-7480, 10.20.20.28-7480)

task status:

data:
  pools:   8 pools, 8192 pgs
  objects: 4.26k objects, 15 GiB
  usage:   149 GiB used, 1.1 TiB / 1.2 TiB avail
  pgs:    8192 active+clean

io:
  client: 21 KiB/s rd, 169 B/s wr, 0 op/s rd, 0 op/s wr
```

2.4.5.1 Log into the Cloud platform again to create a virtual machine for related tests (including migration and attaching volumes) to ensure that everything works as expected.

2.4.5.2 Login to the Cloud and enable HA Policy after confirmation.

3 Risk Assessment

3.1 Vulnerability List

NexaVM Cloud Operations Risk Checklist				
Module	Action	Impact on Risk	Level	Involved?
Cloud Platform Detach	Detach Primary Storage from Cluster	Disable All Businesses on This Storage	High	No
	Detach iSCSI Storage from Cluster	Disable All Businesses on This Storage	High	No
	Detach L2 Network from Cluster	Global Network Disconnection of All Business VM Instances on This Cluster	High	No
	Disassociate VPC Network from VPC vRouter	Global Network Disconnection Detaches All Business VM Instances in the VPC Network	High	No
	Detach VXLAN Pool from Cluster	Global Network Disconnection of All Business VM Instances on This Cluster	High	No
	Detach VPC vRouter	Applies to all VM instances of the specified VPC network	High	No
Cloud Platform	Delete Zone	All businesses in the zone are stopped and all managed data in the zone is lost	High	No

Delete	Delete Cluster	All businesses of this cluster are stopped, including workloads running on local storage. In addition, all data on the local storage will be lost	High	No
	Delete Host	All businesses on this host are stopped. If you have attached a local storage, data of all VM instances and volumes on this host will be lost	High	No
	Delete Primary Storage	Data of all VM instances and volumes on this primary storage is lost	High	No
	Delete Image Storage	All images on this image storage are lost	High	No
	Delete VXLAN Pool	network outage means all business VM instances on this network are interrupted	High	No
	Delete L2 Network	Global Network Disconnection Disconnects All Business VM Instances on the Network	High	No
	Delete L3 Network	Global Network Disconnection Disconnects All Business VM Instances on the Network	High	No
	Delete VPC vRouter	Analyze all businesses running on the VPC vRouter when an L3 network is globally disconnected	High	No
	Delete Snapshot	The backup data of the snapshot is lost. This is an irreversible operation and you cannot recover the lost data.	High	No
	Delete Baremetal Gateway Node	Stop Baremetal services	High	No
	Delete Load Balancer	An NLB is globally interrupted	High	No
	Delete Account/Project	Expunge Resources of All Business VM Instances and Volume on This Account/project	High	No
Bottom Layer Change	Host Kernel Upgrade	Requires host reboot, with comprehensive impact on compatibility (Ceph) and business continuity, subject to thorough testing validation	High	No
	V2V Conversion Toolkit Upgrade	Has extensive impact on compatibility and requires comprehensive testing validation	High	No
	Operating System Hardening	Stops all VM instances and affects network connectivity of all VM instances in the Cloud, requiring thorough testing and validation beforehand	High	No
	Vulnerability Rectification on Cloud Platform	Has extensive impact on compatibility and requires comprehensive testing validation	High	No
	Update Firmware Driver	Changing the NIC, FC card, or RAID card firmware driver has a high impact on compatibility and business continuity and requires thorough testing and validation	High	No
	System Parameter Change	Significantly affects compatibility and business continuity and requires comprehensive testing and validation	High	No
	Execute iptables -F or ebtables -F	The execution will comprehensively affect business network connectivity and networks, requiring thorough testing to ensure its correctness.	High	No

	Change of underlying hardware	Modifying the CPU, memory, motherboard, and NIC may affect business continuity	High	No
Distributed Storage High Risk	Change with Multiple Disks	This operation will cause data loss. Please modify data carefully and balance it as needed.	High	No
	Change Across Multiple Nodes	This operation will cause data loss. Please modify data carefully and balance it as needed.	High	No
	Remove Multiple Disks	This operation will cause data loss. Please modify data carefully and balance it as needed.	High	No
	Remove Multiple Nodes	This operation will cause data loss. Please modify data carefully and balance it as needed.	High	No
	Remove upon Alarm Trigger	This operation will cause data loss. Please implement this operation when the VM is in good health.	High	No
	Format Disk	Executing this operation will cause data loss. Please exercise with caution.	High	No
	Reduce Storage Repliation Count	Executing this operation will cause data loss. Please exercise with caution.	High	No
	Execute cleanup script	Executing this operation will cause data loss. Please exercise with caution.	High	No
	Delete Resource from Storage UI	Executing this operation will cause data loss. Please exercise with caution.	High	No
	Volume Online Compression	Executing this operation will cause data loss. Please exercise with caution.	High	No
	High-risk command execution	Using third-party tools such as fio, dd, vdbench, Intel MAS, and others to operate on volumes or disks may directly cause data loss	High	No
	Initialize Object Storage with Underlying Objects	Executing this operation will cause data loss. Please exercise with caution.	High	No
	Change Disk	Changing system, cache, and data disks may cause data loss	High	No
Distributed Storage Medium Risk	Storage Performance Tuning	May increase IO latency, comprehensively affecting business continuity	Zhong	No
	Adjust Storage Network	Has extensive impact on business continuity, for example, public network and gateway network changes	Zhong	Yes
	Data Rebalance	High-speed balancing will increase IO latency and affect business continuity	Zhong	No
	Online Upgrade of Storage	Restarting OSD processes one by one may cause increased I/O latency, which could affect business continuity	Zhong	No
	Reboot or Power Off upon Alarm	This operation may cause severe failures that affect business continuity	Zhong	No
	Storage Expansion Node	Add no more than 5 VMs at a time. You need to scale out in batches. Otherwise, the operation may affect the business continuity	Zhong	No

	Storage Server Role Update	Modifying the role might affect business continuity	Zhong	No
	Restart Node or IO Process	MON, OSD, XDC, and RGW may affect business continuity	Zhong	No
	IO Stresstest on Production Environment	High I/O pressure increases I/O latency and affects business continuity	Zhong	No
	Modify Time Synchronization for Low Storage Version	3.2 onward (exclusive), modifying time directly may affect business continuity	Zhong	No
	Switch VIP for File Storage	Affects Business Continuity	Zhong	No
	OSD Data Imbalance Adjustment	Improving I/O latency and affecting business continuity by making the capacity utilization of these LUNs less than or more than the specified threshold	High	No
Network	Underlying Network Change	Modifying bond, IP, routing, and iptables may affect business continuity	High	No
	Network Performance Tuning	Latency, packet loss, and forwarding performance may affect business continuity	High	No
	Modify Switch	Updates, aggregations, patches, and other adjustments can significantly affect business continuity	High	No
	Modify Switch Schedule	Modifying VLANs, VPN instances, gateways, routing tables, treating loops, and configuring STP will affect your business continuity	High	No
	Object Routing Change	Completely affects business continuity	High	No
Others Medium High Risk	VM High Availability Switch Not Disabled During O&M Period	may affect business continuity and data security	High	No
	Database of the O&M Cloud Platform Has Not Been Backed Up	Accidental operations may affect all management data on the Cloud	High	No
	Adjust Host IP while Business Running	May affect business continuity and data security	High	No
	Run Business and Adjust Switch Configuration	may affect business continuity and data security	High	No
	Business Running, VPC HA Switch	May affect the network connectivity of all VM instances in the VPC	High	No
	Business Running, VPC Kernel Upgrade	May affect the network connectivity of all VM instances in the VPC	High	No
	rm -f	Irreversible operation, data lost	High	No
	systemctl/service Changes	Stopping or restarting key services may affect business continuity	High	No

3.2 Risk Mitigation Plan

To mitigate the impact of risks associated with this change, implement the following steps in advance:

1. Communicate well with your customers on the change time window before you perform modification operations. In general, set the time window to a period that has the least impact on business as possible.

2. Before making changes, check whether the Cloud and storage health status are normal. Test the change plan as needed.

4. Reversion Plan

Recover the management network of the Cloud and the storage platform following the previous steps.

5 Emergency Response Plan

If a critical failure occurs while the NexaVM Technical Support team is performing maintenance, they will report it to their supervisor immediately. Within 15 minutes, the supervisor will coordinate with technical experts to form an emergency response team. This team will organize fault analysis, continue to complete the change as needed, or implement rollback based on customer requirements. After rollback, the technical experts will further analyze the cause of the failure and submit a fault analysis report.