

3PAR StoreServ存储 RHEL6.8主机iSCSI多路径配置

主机相关 存储配置 孙清雷 2019-05-19 发表

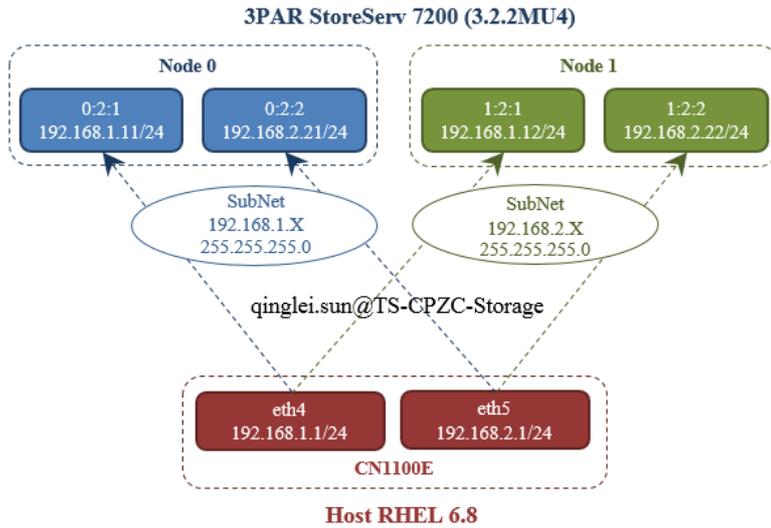
组网及说明

测试用使用StoreServ 7200, 3PAR OS版本为3.2.2MU4。

配置步骤

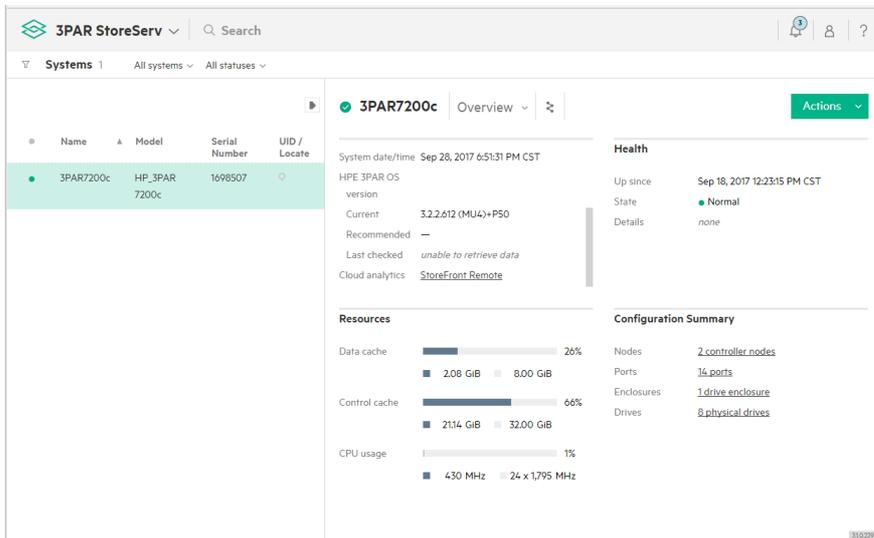
注意 主机实施务必以最新版本的《HPE 3PAR Red Hat Enterprise Linux and Oracle Linux Implementation Guide》为准。

本次测试场景为主机端未配置iface, 而将两个网卡设置不同网段, 分别连接至相同网段的存储端iSCSI端口, 实现4条iSCSI路径, 连接示意图如下



存储端信息如下

3PAR StoreServ 7200c, OS 3.2.2 MU4.



Node 0 iSCSI端口0:2:1 192.168.1.11/24, iSCSI端口0:2:2 192.168.2.21/24, 未设置网关。

3PAR StoreServ | portType:Host protocol:iSCSI | 4 matches out of 14

Ports 4 | All systems | All statuses | Host | All port states | iSCSI | Reset

Port ID (N:S:P)	Label	System	Port Type	Port State	UID / Locate
0:2:1	—	3PAR7200c	Host	Ready	
0:2:2	—	3PAR7200c	Host	Ready	
1:2:1	—	3PAR7200c	Host	Ready	
1:2:2	—	3PAR7200c	Host	Ready	

0:2:1 Settings Actions

Network Configuration

IP Address	VLAN Tag	Subnet Mask	Gateway	DHCP	MTU	iSNS Primary IP Address	iSNS TCP Port	Target Portal Group Tag
192.168.1.11	—	255.255.255.0	—	Disabled	1500	—	3205	21

Persistent Port

Partner (N:S:P) 1:2:1
 Failover state None
 IP Address 192.168.1.12
 VLAN Tag —

SFP Details

State Normal
 Max speed 10.3 Gbps
 Serial number AA1209AK7RX
 Model AFBR-703SDZ-HP1
 Tx disable No
 Tx fault No
 Rx loss No
 Rx power low No

3PAR StoreServ | portType:Host protocol:iSCSI | 4 matches out of 14

Ports 4 | All systems | All statuses | Host | All port states | iSCSI | Reset

Port ID (N:S:P)	Label	System	Port Type	Port State	UID / Locate
0:2:1	—	3PAR7200c	Host	Ready	
0:2:2	—	3PAR7200c	Host	Ready	
1:2:1	—	3PAR7200c	Host	Ready	
1:2:2	—	3PAR7200c	Host	Ready	

0:2:2 Settings Actions

Network Configuration

IP Address	VLAN Tag	Subnet Mask	Gateway	DHCP	MTU	iSNS Primary IP Address	iSNS TCP Port	Target Portal Group Tag
192.168.2.21	—	255.255.255.0	—	Disabled	1500	—	3205	2

Persistent Port

Partner (N:S:P) 1:2:2
 Failover state None
 IP Address 192.168.2.22
 VLAN Tag —

SFP Details

State Normal
 Max speed 10.3 Gbps
 Serial number AX703YC
 Model FTLX8571D3BCL-HP
 Tx disable No
 Tx fault No
 Rx loss No
 Rx power low No

Node 1 iSCSI端口1:2:1 192.168.1.12/24, iSCSI端口1:2:2 192.168.2.22/24, 未设置网关。

3PAR StoreServ | portType:Host protocol:iSCSI | 4 matches out of 14

Ports 4 | All systems | All statuses | Host | All port states | iSCSI | Reset

Port ID (N:S:P)	Label	System	Port Type	Port State	UID / Locate
0:2:1	—	3PAR7200c	Host	Ready	
0:2:2	—	3PAR7200c	Host	Ready	
1:2:1	—	3PAR7200c	Host	Ready	
1:2:2	—	3PAR7200c	Host	Ready	

1:2:1 Settings Actions

Network Configuration

IP Address	VLAN Tag	Subnet Mask	Gateway	DHCP	MTU	iSNS Primary IP Address	iSNS TCP Port	Target Portal Group Tag
192.168.1.12	—	255.255.255.0	—	Disabled	1500	—	3205	12

Persistent Port

Partner (N:S:P) 0:2:1
 Failover state None
 IP Address 192.168.1.11
 VLAN Tag —

SFP Details

State Normal
 Max speed 10.3 Gbps
 Serial number AA1209AK8M7
 Model AFBR-703SDZ-HP1
 Tx disable No
 Tx fault No
 Rx loss No
 Rx power low No

3PAR StoreServ | portType:Host protocol:ISCSI

Ports 4 | All systems | All statuses | Host | All port states | iSCSI | Reset | 4 matches out of 14

1:2:2 | Settings | Actions

Port ID (N:S:P)	Label	System	Port Type	Port State	UID / Locate
0:2:1	—	3PAR7200c	Host	Ready	
0:2:2	—	3PAR7200c	Host	Ready	
1:2:1	—	3PAR7200c	Host	Ready	
1:2:2	—	3PAR7200c	Host	Ready	

Network Configuration

IP Address	VLAN Tag	Subnet Mask	Gateway	DHCP	MTU	iSNS Primary IP Address	iSNS TCP Port	Target Port Group
192.168.2.22	—	255.255.255.0	—	Disabled	1500	—	3205	12

Persistent Port

Partner (N:S:P) G:22

Failover state None

IP Address 192.168.2.21

VLAN Tag —

SFP Details

State Normal

Max speed 10.3 Gbps

Serial number ARR481Q

Model FTLX8571D38CL-HP

Tx disable No

Tx fault No

Rx loss No

Rx power low No

主机端信息如下

RHEL 6.8, eth4 192.168.1.1/24, eth5 192.168.2.1/24.

```
[root@SunQingleiTest ~]# lsb_release -a
LSB Version: :base-4.0-amd64;base-4.0-noarch;core-4.0-amd64;core-4.0-noarch;graphics-4.0-amd64;graphics-4.0-noarch;printing-4.0-amd64;printing-4.0-noarch
Distributor ID: RedHatEnterpriseServer
Description: Red Hat Enterprise Linux Server release 6.8 (Santiago)
Release: 6.8
Codename: Santiago
[root@SunQingleiTest ~]# uname -a
Linux SunQingleiTest 2.6.32-642.el6.x86_64 #1 SMP Wed Apr 13 00:51:26 EDT 2016 x86_64 x86_64 GNU/Linux
[root@SunQingleiTest ~]# ifconfig
eth0      Link encap:Ethernet  HWaddr 78:E7:01:80:05:F4
          inet addr:192.168.1.11  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::7ae7:d1ff:fe8d:d5f4/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:95837 errors:0 dropped:0 overruns:0 frame:0
          TX packets:142 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueue:len:1080
          RX bytes:6714827 (6.4 MiB)  TX bytes:21132 (20.6 KiB)

eth4      Link encap:Ethernet  HWaddr FC:15:84:44:08:B0
          inet addr:192.168.1.1  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::fe15:baff:fe44:8b0/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:15111 errors:0 dropped:0 overruns:0 frame:0
          TX packets:34 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueue:len:1080
          RX bytes:914326 (892.0 KiB)  TX bytes:3772 (3.6 KiB)

eth5      Link encap:Ethernet  HWaddr FC:15:84:44:08:B4
          inet addr:192.168.2.1  Bcast:192.168.2.255  Mask:255.255.255.0
          inet6 addr: fe80::fe15:baff:fe44:8b0/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:15131 errors:0 dropped:0 overruns:0 frame:0
          TX packets:29 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueue:len:1080
          RX bytes:918630 (897.0 KiB)  TX bytes:1818 (1.7 KiB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
          RX packets:20 errors:0 dropped:0 overruns:0 frame:0
          TX packets:20 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueue:len:0
          RX bytes:1568 (1.5 KiB)  TX bytes:1568 (1.5 KiB)

[root@SunQingleiTest ~]#
```

主机端可ping通存储端4个iSCSI端口。

```
[root@SunQingleiTest ~]# ping 192.168.1.11
PING 192.168.1.11 (192.168.1.11) 56(84) bytes of data:
64 bytes from 192.168.1.11: icmp_seq=1 ttl=64 time=0.06 ms
64 bytes from 192.168.1.11: icmp_seq=2 ttl=64 time=0.273 ms
64 bytes from 192.168.1.11: icmp_seq=3 ttl=64 time=0.284 ms
^C
--- 192.168.1.11 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2524ms
rtt min/avg/max/mdev = 0.273/0.540/1.005/0.371 ms
[root@SunQingleiTest ~]# ping 192.168.1.12
PING 192.168.1.12 (192.168.1.12) 56(84) bytes of data:
64 bytes from 192.168.1.12: icmp_seq=1 ttl=64 time=1.53 ms
64 bytes from 192.168.1.12: icmp_seq=2 ttl=64 time=0.290 ms
64 bytes from 192.168.1.12: icmp_seq=3 ttl=64 time=0.286 ms
^C
--- 192.168.1.12 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2573ms
rtt min/avg/max/mdev = 0.286/0.704/1.538/0.590 ms
[root@SunQingleiTest ~]# ping 192.168.2.21
PING 192.168.2.21 (192.168.2.21) 56(84) bytes of data:
64 bytes from 192.168.2.21: icmp_seq=1 ttl=64 time=1.62 ms
64 bytes from 192.168.2.21: icmp_seq=2 ttl=64 time=0.269 ms
64 bytes from 192.168.2.21: icmp_seq=3 ttl=64 time=0.273 ms
^C
--- 192.168.2.21 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2574ms
rtt min/avg/max/mdev = 0.269/0.721/1.622/0.637 ms
[root@SunQingleiTest ~]# ping 192.168.2.22
PING 192.168.2.22 (192.168.2.22) 56(84) bytes of data:
64 bytes from 192.168.2.22: icmp_seq=1 ttl=64 time=1.58 ms
64 bytes from 192.168.2.22: icmp_seq=2 ttl=64 time=0.288 ms
64 bytes from 192.168.2.22: icmp_seq=3 ttl=64 time=0.290 ms
^C
--- 192.168.2.22 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2751ms
rtt min/avg/max/mdev = 0.288/0.720/1.582/0.609 ms
[root@SunQingleiTest ~]#
```

存储端iSCSI端口可ping通主机网卡，网络均正常。

3PAR S Ping from 0:2:1 General

Ports 4

Port ID 4 (N.S.P)

- 0:2:1
- 0:2:2
- 1:2:1
- 1:2:2

System 3PAR7200c

Port 0:2:1

MTU 1500

Destination IP address 192.168.11

Count 1

VLAN tag optional

Results

Ping succeeded

Changed: Destination IP address to "192.168.11"

Yes, ping Cancel

atches out of 14

3102295

3PAR S Ping from 0:2:2 General

Ports 4

Port ID 4 (N.S.P)

- 0:2:1
- 0:2:2
- 1:2:1
- 1:2:2

System 3PAR7200c

Port 0:2:2

MTU 1500

Destination IP address 192.168.21

Count 1

VLAN tag optional

Results

Ping succeeded

Changed: Destination IP address to "192.168.21"

Yes, ping Cancel

atches out of 14

3102295

3PAR S Ping from 1:2:1 General

Ports 4

Port ID 4 (N.S.P)

- 0:2:1
- 0:2:2
- 1:2:1
- 1:2:2

System 3PAR7200c

Port 1:2:1

MTU 1500

Destination IP address 192.168.11

Count 1

VLAN tag optional

Results

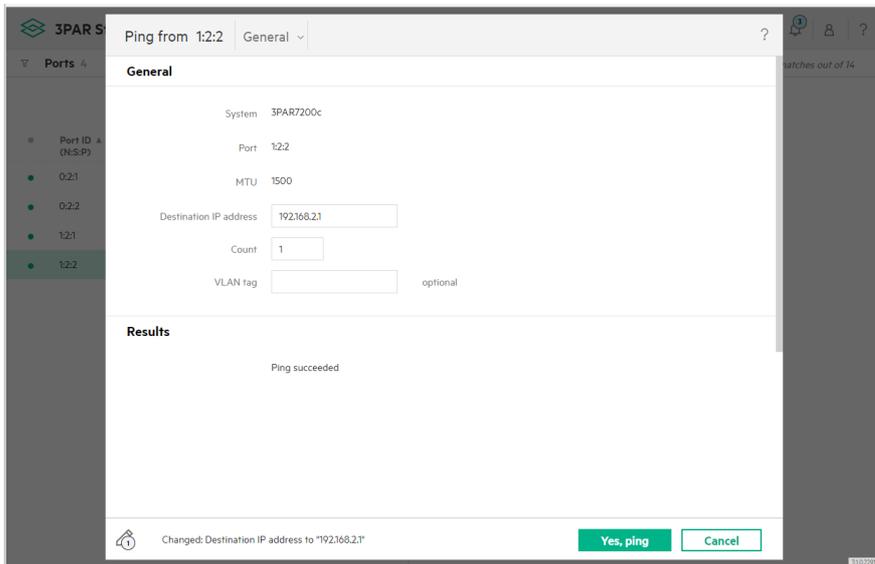
Ping succeeded

Changed: Destination IP address to "192.168.11"

Yes, ping Cancel

atches out of 14

3102295



未配置之前，主机端未见3PAR存储设备。

```
[root@SunQingLeiTest ~]# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 03 Id: 00 Lun: 00
Vendor: HP Model: P4101 Rev: 6.64
Type: RAID ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 01
Vendor: HP Model: LOGICAL VOLUME Rev: 6.64
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 00
Vendor: HP Model: LOGICAL VOLUME Rev: 6.64
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 02
Vendor: HP Model: LOGICAL VOLUME Rev: 6.64
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi1 Channel: 00 Id: 04 Lun: 00
Vendor: HP Model: Ultrium 5-SCSI Rev: Y6PW
Type: Sequential-Access ANSI SCSI revision: 06
Host: scsi1 Channel: 00 Id: 05 Lun: 00
Vendor: HP Model: Ultrium 5-SCSI Rev: Y6PW
Type: Sequential-Access ANSI SCSI revision: 06
Host: scsi1 Channel: 00 Id: 06 Lun: 00
Vendor: HP Model: MSL G3 Series Rev: 6.70
Type: Medium Changer ANSI SCSI revision: 05
[root@SunQingLeiTest ~]#
```

fdisk -l，仅见服务器本地逻辑驱动器所对应的磁盘设备，分别为sda、sdb和sdc。

```
[root@SunQingLeiTest ~]# fdisk -l
Disk /dev/sda: 146.8 GB, 146778695440 bytes
255 heads, 63 sectors/track, 17844 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x8001b02b

Device Boot      Start         End      Blocks   Id  System
/dev/sda1 *        1         64      51200    83  Linux
Partition 1 does not end on cylinder boundary.
/dev/sda2           64      17845   142825472  8e  Linux LVM

Disk /dev/sdb: 283.6 GB, 283564211200 bytes
255 heads, 32 sectors/track, 70265 cylinders
Units = cylinders of 8160 * 512 = 4177920 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x37565c53

Device Boot      Start         End      Blocks   Id  System
/dev/sdb1 *        1       70266   286681088  7   HPFS/NTFS

WARNING: GPT (GUID Partition Table) detected on '/dev/sdc': The util fdisk doesn't support GPT. Use GNU Parted.

Disk /dev/sdc: 146.8 GB, 146778695440 bytes
255 heads, 63 sectors/track, 17844 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x80008a60a

Device Boot      Start         End      Blocks   Id  System
/dev/sdc1 *        1       17845   143385590  ee  GPT

Disk /dev/mapper/vq_sunqingleitest-lv_root: 53.7 GB, 53687091200 bytes
255 heads, 63 sectors/track, 6527 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x80000000

Disk /dev/mapper/vq_sunqingleitest-lv_swap: 6257 MB, 6257901568 bytes
255 heads, 63 sectors/track, 768 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x80000000

Disk /dev/mapper/vq_sunqingleitest-lv_home: 86.3 GB, 86306193408 bytes
255 heads, 63 sectors/track, 10492 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
Disk identifier: 0x80000000
```

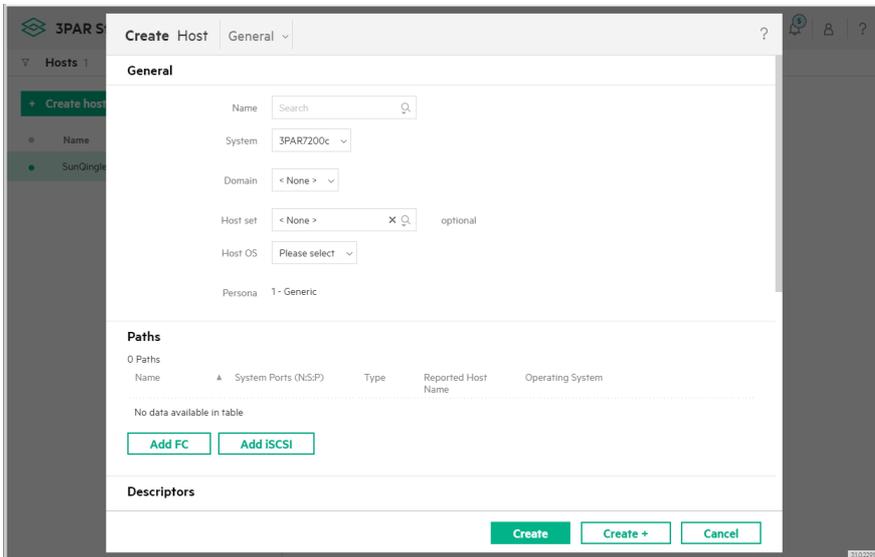
主机端已安装iSCSI initiator。

```
[root@SunQingLeiTest ~]# rpm -qa | grep iscsi-initiator*
iscsi-initiator-utils-6.2.0-373-21.el6.x86_64
[root@SunQingLeiTest ~]# rpm -qa | grep device-mapper*
device-mapper-event-libs-1.02.117-7.el6.x86_64
device-mapper-1.02.117-7.el6.x86_64
device-mapper-persistent-data-0.8.2-0.1.rc7.el6.x86_64
device-mapper-libs-1.02.117-7.el6.x86_64
device-mapper-event-1.02.117-7.el6.x86_64
[root@SunQingLeiTest ~]#
```

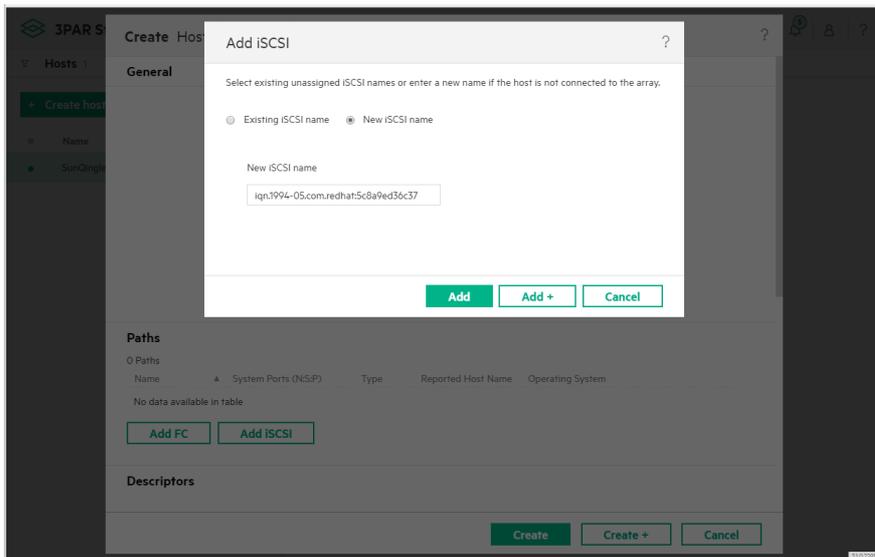
编辑iSCSI配置文件/etc/iscsi/iscsid.conf，修改内容如下

```
[root@SunQingLeiTest ~]# vim /etc/iscsi/iscsid.conf
[root@SunQingLeiTest ~]#
# Startup settings
#
# To request that the iscsi initd scripts startup a session set to "automatic".
node.startup = automatic
#
# To manually startup the session set to "manual". The default is automatic.
node.startup = automatic
node.conn[0].startup = automatic
#
# For "automatic" startup nodes, setting this to "Yes" will try logins on each
# available iface until one succeeds, and then stop. The default "No" will try
# logins on all available ifaces simultaneously.
node.leading_login = No
node.session.timeo
```

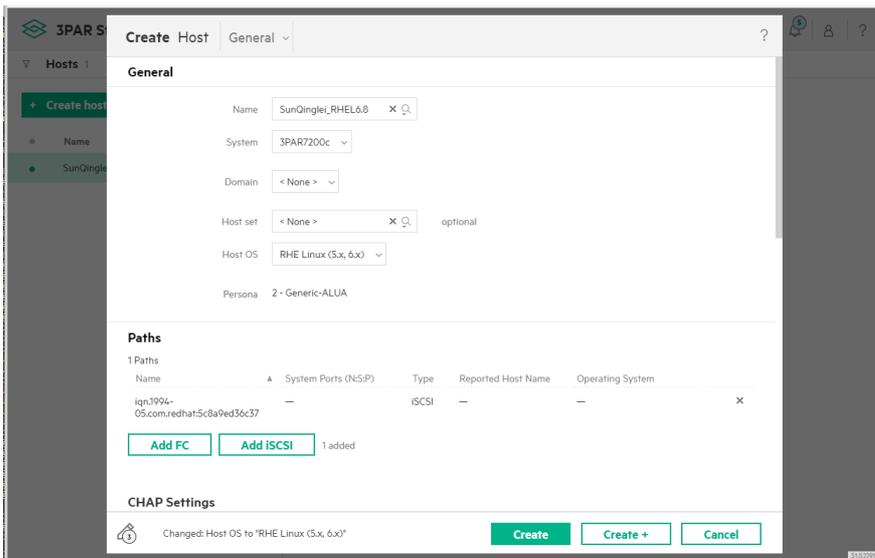
node.session.timeo



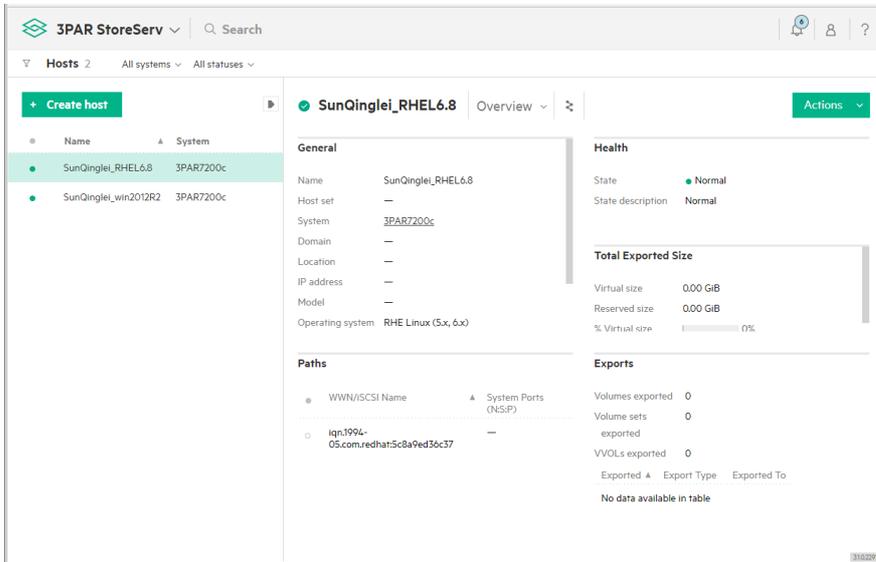
点选新iSCSI名称，输入RHEL主机的iqn信息。



按需输入主机名称，测试中主机名称为“SunQinglei_RHEL6.8”，主机OS类型选择RHE Linux (5.x, 6.x)，角色相应设置为“2 - Generic -ALUA”。其余选项暂不设置。



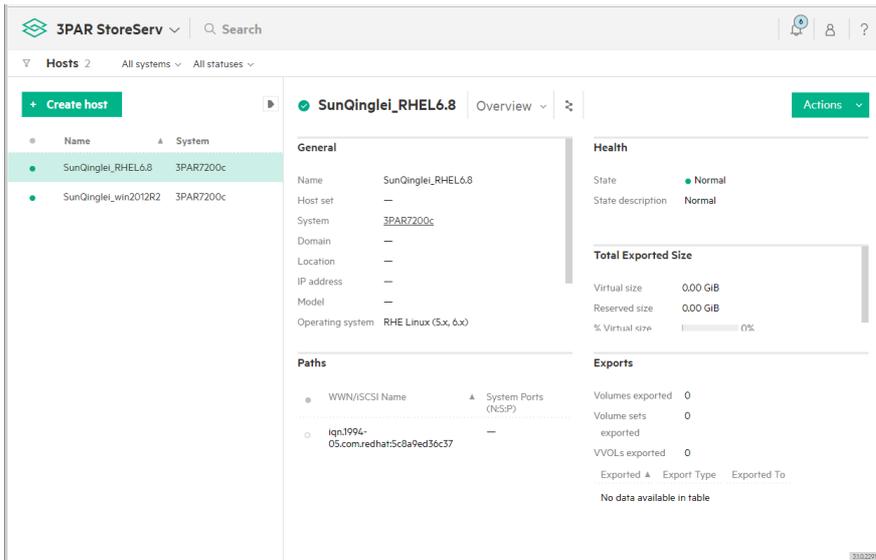
主机“SunQinglei_RHEL6.8”已创建完成，注意，由于主机端与存储端未进行iSCSI连接，故主机概况-路径信息的系统端口为空。



RHEL主机使用iscsiadm命令对存储端的4个iSCSI端口进行发现操作。
发现后，iscsiadm --mode node，可见已将存储端4个iSCSI端口识别为4个iSCSI目标。

```
[root@SunQingleiTest ~]# iscsiadm --mode discovery --type sendtargets --portal 192.168.1.11:3260
Starting iscsid: [ OK ]
192.168.1.11:3260,21 iqn.2000-05.com.3pardata:20210002ac0180cb
[root@SunQingleiTest ~]# iscsiadm --mode discovery --type sendtargets --portal 192.168.1.12:3260
192.168.1.12:3260,121 iqn.2000-05.com.3pardata:21210002ac0180cb
[root@SunQingleiTest ~]# iscsiadm --mode discovery --type sendtargets --portal 192.168.2.21:3260
192.168.2.21:3260,22 iqn.2000-05.com.3pardata:20220002ac0180cb
[root@SunQingleiTest ~]# iscsiadm --mode discovery --type sendtargets --portal 192.168.2.22:3260
192.168.2.22:3260,122 iqn.2000-05.com.3pardata:21220002ac0180cb
[root@SunQingleiTest ~]#
[root@SunQingleiTest ~]# iscsiadm --mode node
192.168.2.22:3260,122 iqn.2000-05.com.3pardata:21220002ac0180cb
192.168.1.12:3260,121 iqn.2000-05.com.3pardata:21210002ac0180cb
192.168.1.11:3260,21 iqn.2000-05.com.3pardata:20210002ac0180cb
192.168.2.21:3260,22 iqn.2000-05.com.3pardata:20220002ac0180cb
[root@SunQingleiTest ~]#
```

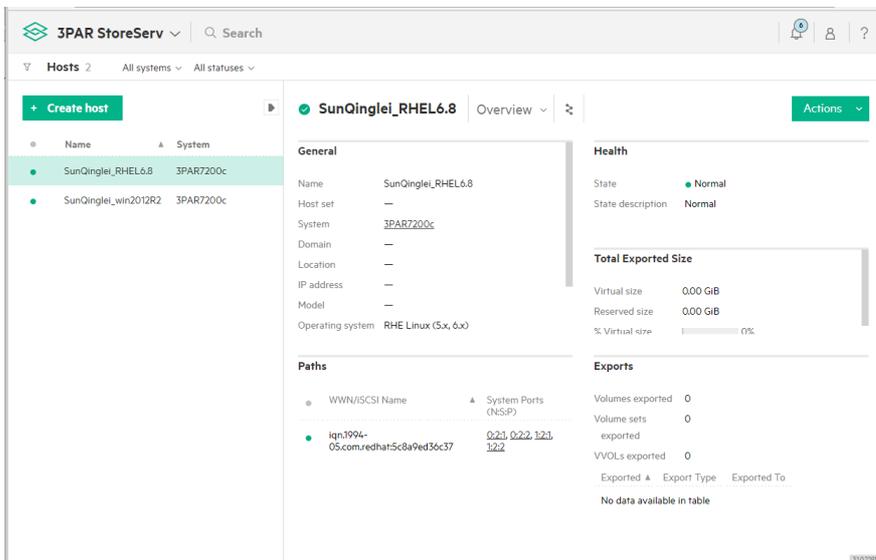
主机端仅发现存储目标，未建立iSCSI连接，存储端的主机路径信息的系统端口仍为空。



主机端对所有已发现的存储目标进行iSCSI login，建立iSCSI连接。

```
[root@SunQingleiTest ~]# iscsiadm --mode node --login all
Logging in to [iface: default, target: iqn.2000-05.com.3pardata:21220002ac0180cb, portal: 192.168.2.22,3260] (multiple)
Logging in to [iface: default, target: iqn.2000-05.com.3pardata:21210002ac0180cb, portal: 192.168.1.12,3260] (multiple)
Logging in to [iface: default, target: iqn.2000-05.com.3pardata:20210002ac0180cb, portal: 192.168.1.11,3260] (multiple)
Logging in to [iface: default, target: iqn.2000-05.com.3pardata:20220002ac0180cb, portal: 192.168.2.21,3260] (multiple)
Login to [iface: default, target: iqn.2000-05.com.3pardata:21220002ac0180cb, portal: 192.168.2.22,3260] successful.
Login to [iface: default, target: iqn.2000-05.com.3pardata:21210002ac0180cb, portal: 192.168.1.12,3260] successful.
Login to [iface: default, target: iqn.2000-05.com.3pardata:20210002ac0180cb, portal: 192.168.1.11,3260] successful.
Login to [iface: default, target: iqn.2000-05.com.3pardata:20220002ac0180cb, portal: 192.168.2.21,3260] successful.
[root@SunQingleiTest ~]#
```

存储端相应可见，主机的路径信息中系统端口变为4个iSCSI端口，说明RHEL主机端已连接至存储端的iSCSI端口。



建立iSCSI连接后，可见iSCSI会话情况。

```
[root@SunQingleiTest ~]# iscsiadm --mode session
tcp: [1] 192.168.2.22:3260,122 iqn.2000-05.com.3pardata:21220002ac0180cb (non-flash)
tcp: [2] 192.168.1.12:3260,121 iqn.2000-05.com.3pardata:21210002ac0180cb (non-flash)
tcp: [3] 192.168.1.11:3260,21 iqn.2000-05.com.3pardata:20210002ac0180cb (non-flash)
tcp: [4] 192.168.2.21:3260,22 iqn.2000-05.com.3pardata:20220002ac0180cb (non-flash)
[root@SunQingleiTest ~]#
```

对于RHEL6，可以通过service iscsi status查看iSCSI服务状态。

该命令可以方便地观察主机端已连接的iSCSI目标信息，包括发起端(主机网卡)、目标端(存储iSCSI端口)以及存储目标所包含的SCSI设备。

主机端与存储端建立了4条iSCSI会话，主机端发现了4个存储目标，以下图为例，存储目标iqn.2000-05.com.3pardata:21220002ac0180cb是由主机端网卡eth4与存储端SCSI端口1:2:2之间建立的iSCSI连接，Attached SCSI Device可见主机端识别到存储端的SES机箱设备，LUN 254。

注* 对于SES机箱设备管理标准，此处不再展开说明。

```
[root@SunQingleiTest ~]# service iscsi status
iSCSI Transport Class version 2.0-879
version 6.2.0-873.z1.616
Target: iqn.2000-05.com.3pardata:21220002ac0180cb (non-flash)
Current Portal: 192.168.2.22:3260,122
Persistent Portal: 192.168.2.22:3260,122
*****
Interface:
*****
Interface Name: default
Interface Transport: tcp
Interface InitiatorName: iqn.1994-05.com.redhat:5c8a9ed36c37
Interface IPAddress: 192.168.2.1
Interface HAddress: <empty>
Interface Netdev: <empty>
SID: 1
iSCSI Connection State: LOGGED_IN
iSCSI Session State: LOGGED_IN
Internal iSCSI Session State: NO CHANGE
*****
Timeouts:
*****
Recovery Timeout: 10
Target Reset Timeout: 30
LUN Reset Timeout: 30
Abort Timeout: 15
*****
CHAP:
*****
username: <empty>
password: *****
username_in: <empty>
password_in: *****
*****
Negotiated iSCSI params:
*****
HeaderDigest: None
DataDigest: None
MaxRecvDataSegmentLength: 262144
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
MaxBurstLength: 262144
ImmediateData: No
InitialR2T: Yes
MaxOutstandingR2T: 1
*****
Attached SCSI devices:
*****
Host Number: 10 State: running
scsi0 channel 00 Id 0 Lun: 254
```

```
Target: iqn.2000-05.com.3pardata:21210002ac0180cb (non-flash)
Current Portal: 192.168.1.12:3260,121
Persistent Portal: 192.168.1.12:3260,121
*****
Interface:
*****
Interface Name: default
Interface Transport: tcp
Interface InitiatorName: iqn.1994-05.com.redhat:5c8a9ed36c37
Interface IPAddress: 192.168.1.1
Interface HAddress: <empty>
Interface Netdev: <empty>
SID: 2
iSCSI Connection State: LOGGED_IN
iSCSI Session State: LOGGED_IN
Internal iSCSI Session State: NO CHANGE
*****
Timeouts:
*****
Recovery Timeout: 10
Target Reset Timeout: 30
LUN Reset Timeout: 30
Abort Timeout: 15
*****
CHAP:
*****
username: <empty>
password: *****
username_in: <empty>
password_in: *****
*****
Negotiated iSCSI params:
*****
HeaderDigest: None
DataDigest: None
MaxRecvDataSegmentLength: 262144
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
MaxBurstLength: 262144
ImmediateData: No
InitialR2T: Yes
MaxOutstandingR2T: 1
*****
Attached SCSI devices:
*****
Host Number: 11 State: running
scsi1 channel 00 Id 0 Lun: 254
```

```

Target: iqn.2000-05.com.3pardata:20210002ac0180cb (non-flash)
Current Portal: 192.168.1.11:3260,21
Persistent Portal: 192.168.1.11:3260,21
*****
Interface:
*****
Iface Name: default
Iface Transport: tcp
Iface InitiatorName: iqn.1994-05.com.redhat:5c8a9ed36c37
Iface IPaddress: 192.168.1.1
Iface Hwaddress: <empty>
Iface Netdev: <empty>
SID: 5
iSCSI Connection State: LOGGED_IN
iSCSI Session State: LOGGED_IN
Internal iSCSI Session State: NO CHANGE
*****
Timeouts:
*****
Recovery Timeout: 10
Target Reset Timeout: 30
LUN Reset Timeout: 30
Abort Timeout: 15
*****
CHAP:
*****
username: <empty>
password: *****
username_in: <empty>
password_in: *****
*****
Negotiated iSCSI params:
*****
HeaderDigest: None
DataDigest: None
MaxRecvDataSegmentLength: 262144
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
MaxBurstLength: 262144
ImmediateData: No
InitialR2T: Yes
MaxOutstandingR2T: 1
*****
Attached SCSI devices:
*****
Host Number: 12 State: running
scsi12 Channel 00 Id 0 Lun: 254

```

```

Target: iqn.2000-05.com.3pardata:20220002ac0180cb (non-flash)
Current Portal: 192.168.2.21:3260,22
Persistent Portal: 192.168.2.21:3260,22
*****
Interface:
*****
Iface Name: default
Iface Transport: tcp
Iface InitiatorName: iqn.1994-05.com.redhat:5c8a9ed36c37
Iface IPaddress: 192.168.2.1
Iface Hwaddress: <empty>
Iface Netdev: <empty>
SID: 4
iSCSI Connection State: LOGGED_IN
iSCSI Session State: LOGGED_IN
Internal iSCSI Session State: NO CHANGE
*****
Timeouts:
*****
Recovery Timeout: 10
Target Reset Timeout: 30
LUN Reset Timeout: 30
Abort Timeout: 15
*****
CHAP:
*****
username: <empty>
password: *****
username_in: <empty>
password_in: *****
*****
Negotiated iSCSI params:
*****
HeaderDigest: None
DataDigest: None
MaxRecvDataSegmentLength: 262144
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
MaxBurstLength: 262144
ImmediateData: No
InitialR2T: Yes
MaxOutstandingR2T: 1
*****
Attached SCSI devices:
*****
Host Number: 13 State: running
scsi13 Channel 00 Id 0 Lun: 254
[root@SunQingle1Test ~]#

```

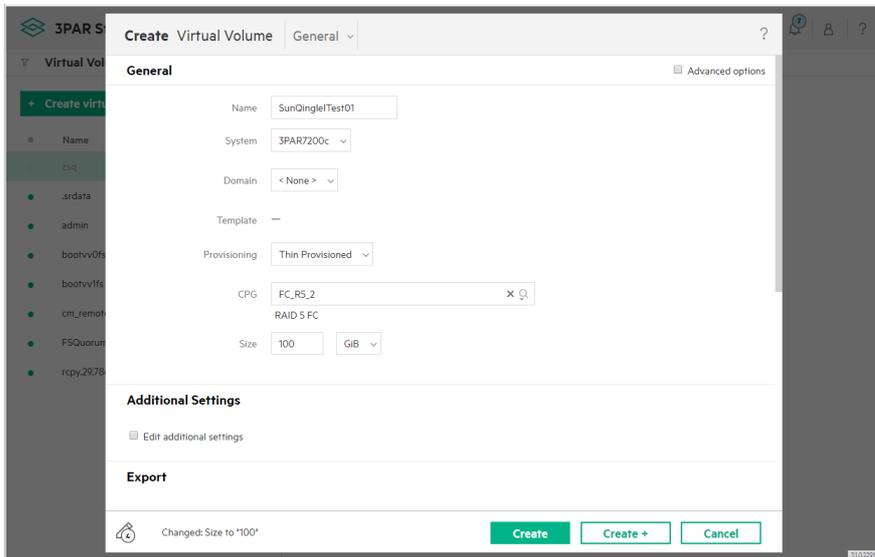
此时，RHEL主机端查看已识别的SCSI设备，已发现新增4个3PAR机箱设备，LUN 254。

```

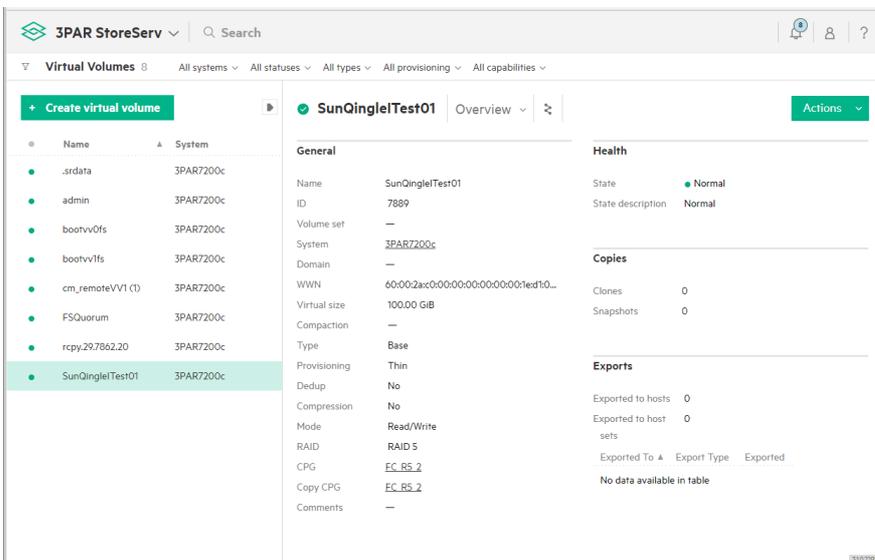
[root@SunQingle1Test ~]# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 03 Id: 00 Lun: 00
Vendor: HP Model: P4101 Rev: 6.64
Type: RAID ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 01
Vendor: HP Model: LOGICAL VOLUME Rev: 6.64
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 00
Vendor: HP Model: LOGICAL VOLUME Rev: 6.64
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 02
Vendor: HP Model: LOGICAL VOLUME Rev: 6.64
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi1 Channel: 00 Id: 04 Lun: 00
Vendor: HP Model: Ultrium 5-SCSI Rev: Y6PW
Type: Sequential-Access ANSI SCSI revision: 06
Host: scsi1 Channel: 00 Id: 05 Lun: 00
Vendor: HP Model: Ultrium 5-SCSI Rev: Y6NW
Type: Sequential-Access ANSI SCSI revision: 06
Host: scsi1 Channel: 00 Id: 06 Lun: 00
Vendor: HP Model: MSL 63 Series Rev: 6.70
Type: Medium Changer ANSI SCSI revision: 05
Host: scsi9 Channel: 00 Id: 00 Lun: 00
Vendor: HP Model: Virtual DVD-ROM Rev:
Type: CD-ROM ANSI SCSI revision: 00
Host: scsi12 Channel: 00 Id: 00 Lun: 254
Vendor: 3PARdata Model: SES Rev: 2224
Type: Enclosure ANSI SCSI revision: 06
Host: scsi10 Channel: 00 Id: 00 Lun: 254
Vendor: 3PARdata Model: SES Rev: 2224
Type: Enclosure ANSI SCSI revision: 06
Host: scsi11 Channel: 00 Id: 00 Lun: 254
Vendor: 3PARdata Model: SES Rev: 2224
Type: Enclosure ANSI SCSI revision: 06
Host: scsi13 Channel: 00 Id: 00 Lun: 254
Vendor: 3PARdata Model: SES Rev: 2224
Type: Enclosure ANSI SCSI revision: 06
[root@SunQingle1Test ~]#

```

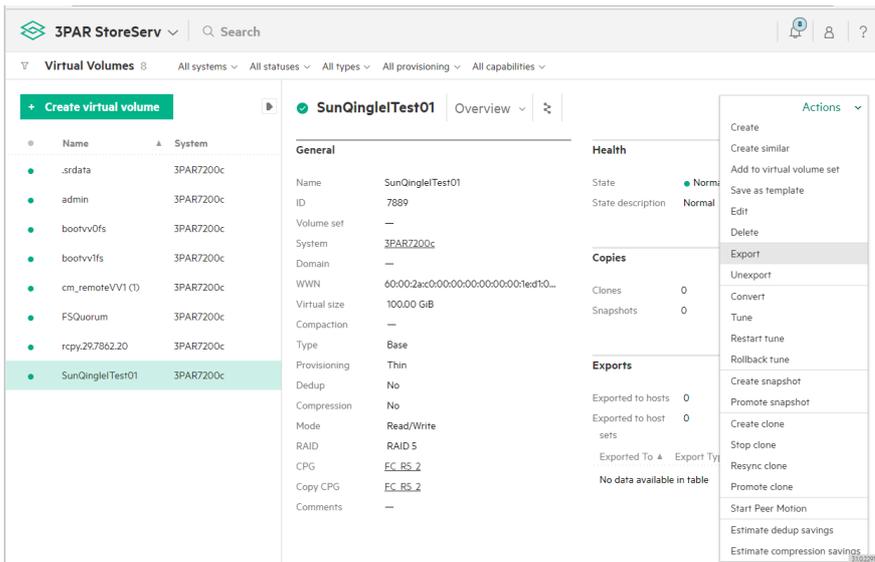
存储端创建存储卷“SunQingle1Test01”，精简模式，大小100GiB。

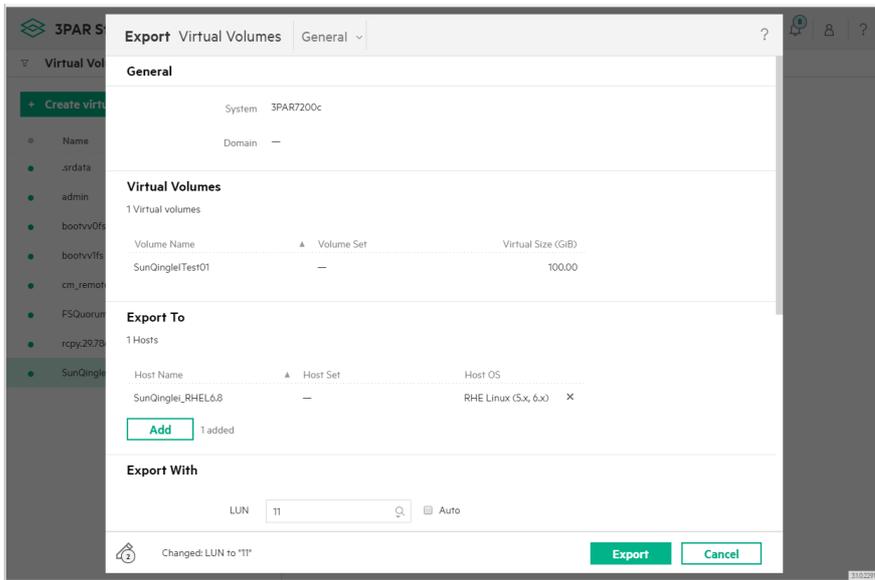


存储卷“SunQingleiTest01”创建完成。



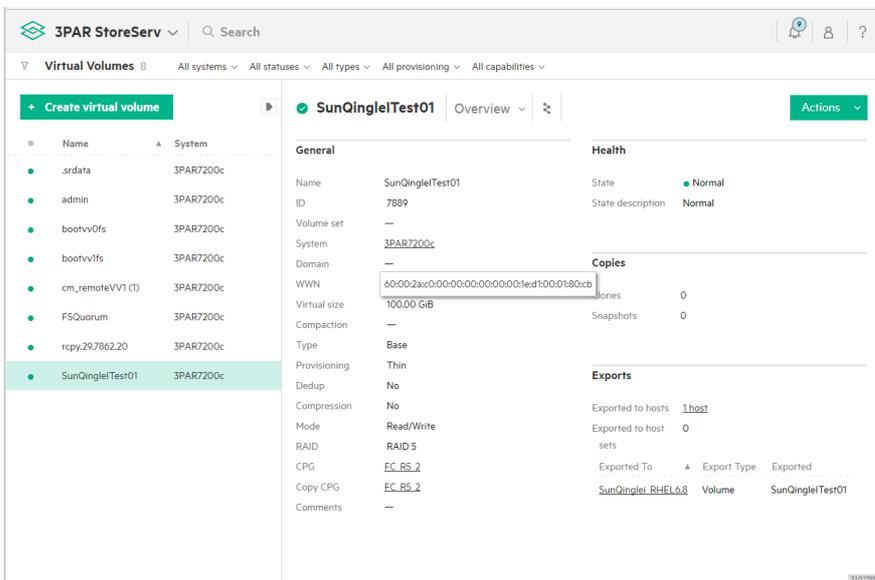
存储卷“SunQingleiTest01”导出至主机“SunQinglei_RHEL6.8”，为方便区分，LUN值设置为11。





导出完成。

另，SSMC界面可见存储卷的WWN信息。



存储端导出存储卷至主机后，RHEL主机不能自动识别。

```
[root@SunQingleTest ~]#
[root@SunQingleTest ~]# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 03 Id: 00 Lun: 00
  Vendor: HP          Model: P410i          Rev: 6.64
  Type: RAID         ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 01
  Vendor: HP          Model: LOGICAL VOLUME Rev: 6.64
  Type: Direct-Access ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: HP          Model: LOGICAL VOLUME Rev: 6.64
  Type: Direct-Access ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 02
  Vendor: HP          Model: LOGICAL VOLUME Rev: 6.64
  Type: Direct-Access ANSI SCSI revision: 05
Host: scsi1 Channel: 00 Id: 04 Lun: 00
  Vendor: HP          Model: Ultrium 5-SCSI  Rev: Y6PW
  Type: Sequential-Access ANSI SCSI revision: 06
Host: scsi1 Channel: 00 Id: 05 Lun: 00
  Vendor: HP          Model: Ultrium 5-SCSI  Rev: Y6NW
  Type: Sequential-Access ANSI SCSI revision: 06
Host: scsi1 Channel: 00 Id: 06 Lun: 00
  Vendor: HP          Model: MSL G3 Series   Rev: 6.70
  Type: Medium Changer ANSI SCSI revision: 05
Host: scsi1 Channel: 00 Id: 00 Lun: 00
  Vendor: HP          Model: Virtual DVD-ROM Rev:
  Type: CD-ROM         ANSI SCSI revision: 00
Host: scsi12 Channel: 00 Id: 00 Lun: 254
  Vendor: 3PARdata Model: SES Rev: 3224
  Type: Enclosure     ANSI SCSI revision: 06
Host: scsi10 Channel: 00 Id: 00 Lun: 254
  Vendor: 3PARdata Model: SES Rev: 3224
  Type: Enclosure     ANSI SCSI revision: 06
Host: scsi11 Channel: 00 Id: 00 Lun: 254
  Vendor: 3PARdata Model: SES Rev: 3224
  Type: Enclosure     ANSI SCSI revision: 06
Host: scsi13 Channel: 00 Id: 00 Lun: 254
  Vendor: 3PARdata Model: SES Rev: 3224
  Type: Enclosure     ANSI SCSI revision: 06
[root@SunQingleTest ~]#
```

对于iSCSI连接方式，RHEL主机需要重新查找新增设备。3PAR存储是MLPT类型，故使用iscsiadm --mode session --rescan，重新扫描当前iSCSI会话即可。

注* 根据存储产品的类型，重新扫描方式不同。对于SLPT类型的设备(如LeftHand)，则需要重新发现存储目标。

```
[root@SunQingleTest ~]# iscsiadm --mode session --rescan
Rescanning session [sid: 1, target: iqn.2000-05.com:3pardata:21220002ac0180cb, portal: 192.168.2.22,3260]
Rescanning session [sid: 2, target: iqn.2000-05.com:3pardata:21210002ac0180cb, portal: 192.168.1.12,3260]
Rescanning session [sid: 3, target: iqn.2000-05.com:3pardata:20210002ac0180cb, portal: 192.168.1.11,3260]
Rescanning session [sid: 4, target: iqn.2000-05.com:3pardata:20220002ac0180cb, portal: 192.168.2.21,3260]
[root@SunQingleTest ~]#
```

重新扫描SCSI会话后，可见每个存储目标所包含的SCSI设备中出现磁盘设备，可见LUN值为11，即存储卷'SunQingleTest01'。

```

[root@SunQangle1Test ~]# service iscsi status
iSCSI Transport Class version 2.0-870
version 6.2.0-873.21.el6
Target: iqn.2000-05.com.3pardata:21220002ac0180cb (non-flash)
Current Portal: 192.168.2.22:3260,122
Persistent Portal: 192.168.2.22:3260,122
*****
Interface:
*****
Iface Name: default
Iface Transport: tcp
Iface Initiatorname: iqn.1994-05.com.redhat:5c8a9ed36c37
Iface IPaddress: 192.168.2.1
Iface Hwaddress: <empty>
Iface Netdev: <empty>
SID: 1
iSCSI Connection State: LOGGED IN
iSCSI Session State: LOGGED_IN
Internal iscsid Session State: NO CHANGE
*****
Timeouts:
*****
Recovery Timeout: 10
Target Reset Timeout: 30
LUN Reset Timeout: 30
Abort Timeout: 15
*****
CHAP:
*****
username: <empty>
password: *****
username_in: <empty>
password_in: *****
*****
Negotiated iSCSI params:
*****
HeaderDigest: None
DataDigest: None
MaxRecvDataSegmentLength: 262144
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
MaxBurstLength: 262144
ImmediateData: No
InitialR2T: Yes
MaxOutstandingR2T: 1
*****
Attached SCSI devices:
*****
Host Number: 10 State: running
scsi0 Channel 00 Id 0 Lun: 11
Attached scsi disk sdd State: running
scsi0 Channel 00 Id 0 Lun: 254

```

```

Target: iqn.2000-05.com.3pardata:21210002ac0180cb (non-flash)
Current Portal: 192.168.1.12:3260,121
Persistent Portal: 192.168.1.12:3260,121
*****
Interface:
*****
Iface Name: default
Iface Transport: tcp
Iface Initiatorname: iqn.1994-05.com.redhat:5c8a9ed36c37
Iface IPaddress: 192.168.1.1
Iface Hwaddress: <empty>
Iface Netdev: <empty>
SID: 2
iSCSI Connection State: LOGGED IN
iSCSI Session State: LOGGED_IN
Internal iscsid Session State: NO CHANGE
*****
Timeouts:
*****
Recovery Timeout: 10
Target Reset Timeout: 30
LUN Reset Timeout: 30
Abort Timeout: 15
*****
CHAP:
*****
username: <empty>
password: *****
username_in: <empty>
password_in: *****
*****
Negotiated iSCSI params:
*****
HeaderDigest: None
DataDigest: None
MaxRecvDataSegmentLength: 262144
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
MaxBurstLength: 262144
ImmediateData: No
InitialR2T: Yes
MaxOutstandingR2T: 1
*****
Attached SCSI devices:
*****
Host Number: 11 State: running
scsi1 Channel 00 Id 0 Lun: 11
Attached scsi disk sde State: running
scsi1 Channel 00 Id 0 Lun: 254

```

```

Target: iqn.2000-05.com.3pardata:20210002ac0180cb (non-flash)
Current Portal: 192.168.1.11:3260,21
Persistent Portal: 192.168.1.11:3260,21
*****
Interface:
*****
Iface Name: default
Iface Transport: tcp
Iface Initiatorname: iqn.1994-05.com.redhat:5c8a9ed36c37
Iface IPaddress: 192.168.1.1
Iface Hwaddress: <empty>
Iface Netdev: <empty>
SID: 3
iSCSI Connection State: LOGGED IN
iSCSI Session State: LOGGED_IN
Internal iscsid Session State: NO CHANGE
*****
Timeouts:
*****
Recovery Timeout: 10
Target Reset Timeout: 30
LUN Reset Timeout: 30
Abort Timeout: 15
*****
CHAP:
*****
username: <empty>
password: *****
username_in: <empty>
password_in: *****
*****
Negotiated iSCSI params:
*****
HeaderDigest: None
DataDigest: None
MaxRecvDataSegmentLength: 262144
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
MaxBurstLength: 262144
ImmediateData: No
InitialR2T: Yes
MaxOutstandingR2T: 1
*****
Attached SCSI devices:
*****
Host Number: 12 State: running
scsi2 Channel 00 Id 0 Lun: 11
Attached scsi disk sdf State: running
scsi2 Channel 00 Id 0 Lun: 254

```

```

Target: iqn.2000-05.com.3pardata:20220002ac0100cb (non-flash)
Current Portal: 192.168.2.21:3260,22
Persistent Portal: 192.168.2.21:3260,22
*****
Interface:
*****
Iface Name: default
Iface Transport: tcp
Iface InitiatorName: iqn.1994-05.com.redhat:sc8a9ed36c37
Iface IPAddress: 192.168.2.1
Iface HwAddress: <empty>
Iface Netdev: <empty>
SID: 4
iSCSI Connection State: LOGGED_IN
iSCSI Session State: LOGGED_IN
Internal iSCSI Session State: NO CHANGE
*****
Timeouts:
*****
Recovery Timeout: 10
Target Reset Timeout: 30
LUN Reset Timeout: 30
Abort Timeout: 15
*****
CHAP:
*****
username: <empty>
password: *****
username_in: <empty>
password_in: *****
*****
Negotiated iSCSI params:
*****
HeaderDigest: None
DataDigest: None
MaxRecvDataSegmentLength: 262144
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
MaxBurstLength: 262144
ImmediateData: No
InitialR2T: Yes
MaxOutstandingR2T: 1
*****
Attached SCSI devices:
*****
Host Number: 13 State: running
scsi3 Channel 00 Id 0 Lun: 11      Attached scsi disk sda
scsi3 Channel 00 Id 0 Lun: 254     State: running
scsi3 Channel 00 Id 0 Lun: 254
[root@SunQingleiTest ~]#

```

查看SCSI设备，可见已识别到4个3PAR存储磁盘设备，LUN 11。

```

[root@SunQingleiTest ~]# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
Vendor: HP Model: P4101 Rev: 6.64
Type: RAID ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 01
Vendor: HP Model: LOGICAL VOLUME Rev: 6.64
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 02
Vendor: HP Model: LOGICAL VOLUME Rev: 6.64
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi0 Channel: 00 Id: 00 Lun: 02
Vendor: HP Model: LOGICAL VOLUME Rev: 6.64
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi1 Channel: 00 Id: 04 Lun: 00
Vendor: HP Model: Ultrium 5-SCSI Rev: Y6PW
Type: Sequential-Access ANSI SCSI revision: 06
Host: scsi1 Channel: 00 Id: 05 Lun: 00
Vendor: HP Model: Ultrium 5-SCSI Rev: Y6WV
Type: Sequential-Access ANSI SCSI revision: 06
Host: scsi1 Channel: 00 Id: 06 Lun: 00
Vendor: HP Model: MSL G3 Series Rev: 6.70
Type: Medium Changer ANSI SCSI revision: 05
Host: scsi5 Channel: 00 Id: 00 Lun: 00
Vendor: HP Model: Virtual DVD-ROM Rev:
Type: CD-ROM ANSI SCSI revision: 00
Host: scsi12 Channel: 00 Id: 00 Lun: 254
Vendor: 3PARdata Model: SES Rev: 3224
Type: Enclosure ANSI SCSI revision: 06
Host: scsi10 Channel: 00 Id: 00 Lun: 254
Vendor: 3PARdata Model: SES Rev: 3224
Type: Enclosure ANSI SCSI revision: 06
Host: scsi11 Channel: 00 Id: 00 Lun: 254
Vendor: 3PARdata Model: SES Rev: 3224
Type: Enclosure ANSI SCSI revision: 06
Host: scsi13 Channel: 00 Id: 00 Lun: 254
Vendor: 3PARdata Model: SES Rev: 3224
Type: Enclosure ANSI SCSI revision: 06
Host: scsi10 Channel: 00 Id: 00 Lun: 11
Vendor: 3PARdata Model: VV Rev: 3224
Type: Direct-Access ANSI SCSI revision: 06
Host: scsi11 Channel: 00 Id: 00 Lun: 11
Vendor: 3PARdata Model: VV Rev: 3224
Type: Direct-Access ANSI SCSI revision: 06
Host: scsi12 Channel: 00 Id: 00 Lun: 11
Vendor: 3PARdata Model: VV Rev: 3224
Type: Direct-Access ANSI SCSI revision: 06
Host: scsi13 Channel: 00 Id: 00 Lun: 11
Vendor: 3PARdata Model: VV Rev: 3224
Type: Direct-Access ANSI SCSI revision: 06
[root@SunQingleiTest ~]#

```

fdisk -l，可见将存储卷对应的重复磁盘设备识别为sdd、sde、sdf和sdg。

```

Disk /dev/sdd: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 16384 bytes / 16777216 bytes
Disk identifier: 0x00000000

Disk /dev/sdf: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 16384 bytes / 16777216 bytes
Disk identifier: 0x00000000

Disk /dev/sde: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 16384 bytes / 16777216 bytes
Disk identifier: 0x00000000

Disk /dev/sdg: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 16384 bytes / 16777216 bytes
Disk identifier: 0x00000000
[root@SunQingleiTest ~]#

```

查看sdd~sdg的WWID，均相同且与存储端查看的存储卷WWID相同。至此，已完成RHEL主机对3PAR存储卷的配置、识别。

```

[root@SunQingleiTest ~]#
[root@SunQingleiTest ~]# scsi_id -g -u /dev/sdd
360002ac00000000000001ed1000180cb
[root@SunQingleiTest ~]# scsi_id -g -u /dev/sdf
360002ac00000000000001ed1000180cb
[root@SunQingleiTest ~]# scsi_id -g -u /dev/sde
360002ac00000000000001ed1000180cb
[root@SunQingleiTest ~]# scsi_id -g -u /dev/sdg
360002ac00000000000001ed1000180cb
[root@SunQingleiTest ~]#

```

RHEL主机已安装device mapper multipath。

开启多路径服务。

```
[root@SunQingLeiTest ~]# mpathconf
multipath is enabled
find_multipaths is disabled
user_friendly_names is enabled
dm_multipath module is not loaded
multipathd is chkconfig'd off
[root@SunQingLeiTest ~]# mpathconf --find_multipaths y
[root@SunQingLeiTest ~]# mpathconf
multipath is enabled
find_multipaths is enabled
user_friendly_names is enabled
dm_multipath module is not loaded
multipathd is chkconfig'd off
[root@SunQingLeiTest ~]# service multipathd reload
Reloading multipathd: [FAILED]
[root@SunQingLeiTest ~]# service multipathd reload
Reloading multipathd: [FAILED]
[root@SunQingLeiTest ~]# service multipathd restart
ux_socket connect: No such file or directory [FAILED]
Stopping multipathd daemons: [ OK ]
Starting multipathd daemon: [ OK ]
[root@SunQingLeiTest ~]#
```

multipath -ll, 显示多路径设备, 可见多路径软件已自动将3PAR存储卷所对应的重复磁盘聚合为多路径设备mpathd.

```
[root@SunQingLeiTest ~]# multipath -ll
mpathd (360002ac00000000000001ed1000100cb) dm-3 3PARdata,VV
size=100G Features=0 hwhandlers=0 wprw=
-- policy=round-robin 0 priol status=active
|- 10:0:0:11 sdd 8:48 active ready running
|- 12:0:0:11 sdf 8:80 active ready running
|- 13:0:0:11 sde 8:96 active ready running
|- 11:0:0:11 sde 8:64 active ready running
[root@SunQingLeiTest ~]#
```

```
Disk /dev/sdd: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 16384 bytes / 16777216 bytes
Disk identifier: 0x00000000

Disk /dev/sdf: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 16384 bytes / 16777216 bytes
Disk identifier: 0x00000000

Disk /dev/sde: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 16384 bytes / 16777216 bytes
Disk identifier: 0x00000000

Disk /dev/sdg: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 16384 bytes / 16777216 bytes
Disk identifier: 0x00000000

Disk /dev/mapper/mpathd: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 16384 bytes / 16777216 bytes
Disk identifier: 0x00000000
[root@SunQingLeiTest ~]#
```

编辑多路径配置文件/etc/multipath.conf.

multipath.conf默认内容如下

```
## This is a basic configuration file with some examples, for device mapper
## multipath.
## For a complete list of the default configuration values, see
## /usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf.defaults
## For a list of configuration options with descriptions, see
## /usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf.annotated
##
## REMEMBER: After updating multipath.conf, you must run
## service multipathd reload
## for the changes to take effect in multipathd
##
## By default, devices with vendor = "IBM" and product = "S/390.**" are
## blacklisted. To enable multipathing on these devices, uncomment the
## following lines.
blacklist_exceptions {
# device {
#     vendor "IBM"
#     product "S/390.**"
# }
}

## Use user friendly names, instead of using WWIDs as names.
defaults {
    find_multipaths yes
    user_friendly_names yes
}

## Here is an example of how to configure some standard options.
##
#defaults {
#    udev_dir                /dev
#    polling_interval        10
#    path_selector            "round-robin 0"
#    path_grouping_policy    multibus
#    getuid_callout          "/lib/udev/scsi_id --whitelisted --device=/dev/%n"
#    prio                    alua
#    path_checker             readsector0
#    rr_min_io                100
#    max_fds                  8192
#    rr_weight                priorities
#    fallback                 immediate
#    no_path_retry            fail
#    user_friendly_names     yes
#}

## The wwid line in the following blacklist section is shown as an example
## of how to blacklist devices by wwid. The 2 devnode lines are the
## compiled in default blacklist. If you want to blacklist entire types
## of devices, such as all scsi devices, you should use a devnode line.
## However, if you want to blacklist specific devices, you should use
## a wwid line. Since there is no guarantee that a specific device will
## not change names on reboot (from /dev/sda to /dev/sdb for example)
## devnode lines are not recommended for blacklisting specific devices.
```

```

##
blacklist {
#
# wwid 26353908f02796769
# devnode "*/ram|raw|loop|fd|md|dm-|sr|scd|st|[0-9]**"
# devnode "~hd[a-z]"
#
}
multipath {
#
# multipath {
#
# wwid 3600500b4000156d700012000000b000
# alias yellow
# path_grouping_policy multibus
# path_checker readsector0
# path_selector "round-robin 0"
# fallback manual
# rr_weight priorities
# no_path_retry 5
#
# }
# multipath {
#
# wwid IDEC_____321810758474
# alias red
#
# }
#
}
devices {
#
# device {
#
# vendor "COMPAQ "
# product "H5110 (C)COMPAQ"
# path_grouping_policy multibus
# getuid_callout "/lib/udev/scsi_id --whitelisted --device=/dev/kn"
# path_checker readsector0
# path_selector "round-robin 0"
# hardware_handler "0"
# fallback 15
# rr_weight priorities
# no_path_retry queue
#
# }
# device {
#
# vendor "COMPAQ "
# product "H5A1000 "
# path_grouping_policy multibus
#
# }
#
}
blacklist
}

```

按照3PAR存储Linux实施手册进行优化配置，编辑defaults、devices信息如下

```

## Use user friendly names, instead of using WWIDs as names.
defaults {
#
# find_multipaths yes
# user_friendly_names yes
# polling_interval 10
#
}

```

```

}
blacklist {
}
devices {
#
# device {
#
# vendor "3PARdata"
# product "vpe"
# path_grouping_policy group_by_prio
# path_selector "round-robin 0"
# path_checker tur
# features "0"
# hardware_handler "1 alua"
# prio alua
# fallback immediate
# rr_weight uniform
# no_path_retry 15
# rr_min_io_rq 1
# detect_prio yes
# fast_io_fail_tmo 10
# dev_loss_tmo 14
#
# }
#
}
-- INSERT --

```

编辑配置文件后，重新加载配置文件，service multipathd reload。

校验配置文件，multipath -v2。注* 务必执行校验命令，如果配置文件中的内容有误或者与设备类型不符，执行校验命令后会输出提示信息。

重启多路径服务，service multipathd restart。

查看多路径信息，multipath -ll，可见多路径磁盘设备命名已按照配置文件修改为WWID显示，并可见修改后的路径处理策略。

群组优先级变为50。

```

[root@SunQingleiTest ~]# vim /etc/multipath.conf
[root@SunQingleiTest ~]# service multipathd reload
Reloading multipathd: [ OK ]
[root@SunQingleiTest ~]# multipath -v2
[root@SunQingleiTest ~]# service multipathd restart
ok
Stopping multipathd daemon: [ OK ]
Starting multipathd daemon: [ OK ]
[root@SunQingleiTest ~]# multipath -ll
mpathd (3600022a00000000000010001000b) dm-3 3PARdata,VV
size=1006 features=1 queue_if_no_path' hwhandler='1 alua' wprsrw
+- policy='round-robin 0' prio=50 status=active
|- 10:0:0:11 sdd 8:48 active ready running
|- 12:0:0:11 sdf 8:50 active ready running
|- 13:0:0:11 sde 8:96 active ready running
|- 11:0:0:11 sde 8:64 active ready running
[root@SunQingleiTest ~]#

```

对聚合后的磁盘设备进行格式化分区、写入文件系统、挂载至目录，不再赘述。

至此，已完成多路径配置。

```

[root@SunQingleiTest ~]# fdisk /dev/mapper/mpathd
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel with disk identifier 0x00bc022f.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.

Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
switch off the mode (command 'c') and change display units to
sectors (command 'u').

Command (m for help): n
Command action
e extended
p primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-13054, default 3):
Using default value 3
Last cylinder, +cylinders or +size(K,M,G) (3-13054, default 13054):
Using default value 13054

Command (m for help): t
Selected partition 1
Hex code (type L to list codes): 83

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 22: Invalid argument.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.

```

```

[root@SunQingleiTest ~]# fdisk -l /dev/mapper/mpathd
Disk /dev/mapper/mpathd: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 16384 bytes / 16777216 bytes
Disk identifier: 0x8b8c92f1

   Device Boot      Start         End      Blocks   Id  System
/dev/mapper/mpathd1  3              13054    104840190  83  Linux
Partition 1 does not start on physical sector boundary.
[root@SunQingleiTest ~]# mkfs -t ext4 /dev/mapper/mpathd1
mkfsfs 1.41.12 (17-May-2010)
/dev/mapper/mpathd1 alignment is offset by 15360 bytes.
This may result in very poor performance, (re)-partitioning suggested.
Discarding device blocks: done
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=4 blocks, Stripe width=4096 blocks
6553600 inodes, 26210047 blocks
1310502 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
800 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872

Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 24 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
[root@SunQingleiTest ~]#

```

配置关键点