Quantum

H2000 with 1.6.0 Software Release Notes

Content

November 2021 6-68747-14 Rev A

About the H2000 with 1.6.0 Software

H2000 Compatibility

See H2000 Planning > Installation Site Prerequisites

Important Notes

Upgrade dependencies

For online (internet-connected) upgrades:

- An upgrade to 1.6.0 requires 1.5.2 or 1.5.3. If your system is on a release prior to 1.5.2 you will need run the upgrade multiple times to upgrade to 1.6.0.

For offline/dark site (not internet-connected) upgrades:

- If your system is on a release prior to 1.5.2 you will need to download and run the dark site/offline upgrade multiple times to upgrade to 1.6.0. You can download and run the upgrade to 1.5.2, and then upgrade directly to 1.6.0.

Enhancements and Fixed Issues

Issue	Support #	Description
BLK-6640	n/a	Paths failover correctly when controller is pulled out from the chassis.
BLK-6537	n/a	On first system boot the "Trouble getting Alert count Error: Network Error" is no longer seen.
BLK-6453	n/a	The StorNext NTP server no longer fails when NTP services are not running on the StorNext VM.
BLK-6423	n/a	Host entries for both StorNext VMs are now added to the **`/etc/hosts`** file on both StorNext VMs.
BLK-6394	n/a	The **`HSeriesWS - Alerts - Acknowledged`** filter now works correctly when the "Closed" value is selected in the **`Status`** filter.
BLK-6328	n/a	On the **Configuration > Volume Mapping** page, unchecking all targets on a controller now allows you to add them back.
BLK-6067	n/a	The collect log now includes the OS7 **`systemd`** services information.

Known Issues and Workarounds

Issue	Support #	Description/Applicable Workarounds
BLK-5982	n/a	The **Health > System** page may display the following timeout or error message when multiple JBODs are initializing at the same time. With 1.5.x installed, you could see one of the following error messages: **Error Message #1:**  **Error Message #2:**  These messages are benign, and there is no error condition with the system. If you see one of these error messages, it just means that the H2000 UI is expecting information from other parts of the system, and is unable to display the page correctly until the arrays have completed initializing. Once the arrays have completed initializing, the error messages will clear. If you see either of these messages, you can close the message.

Issue	Support #	Description/Applicable Workarounds

BLK-5314	n/a	If you attempt to upgrade the system using the GUI upgrade page, and see a message that the upgrade has failed, try the upgrade process again using the command-line.
		
		Workaround:
		
		To upgrade the system manually on the command-line using the -force option:
		
		1. Open an SSH session using the management IP address of the controller you want to access using a PuTTY/Terminal client.
		
		Use the **admin** user name and the password for the **admin** account on the H2000 system.
		
		2. Enter:
		
		`sudo /opt/quantum/ansible/scripts/cluster_manager.py setmaintenancemode false`
		
		3. Enter:
		
		`sudo /opt/quantum/scripts/bootstrap_node_upgrade --force`
		
		Note: This method upgrades the system from the production YUM repository.

BLK-4627	n/a	When powering on the system when both controllers are in a powered off state, push the power button on one controller, and then immediately push the power button on the other controller (within 30 secs of each other) so that the NVDs for both controllers are properly recognized by the system and prevents a fenced controller state. If you did not power on both controllers within 30 seconds of one another, see How to un-fence a controller and restart the H2000 cluster below.

How to un-fence a controller and restart the H2000 cluster

See Known Issue 4627.

Step 1: Determine which controller is fenced

- When a block storage controller is fenced H2000 UI will contain an error message across the top of the page:

- Review RAS tickets
- Use the CLI to look for a controller node "reset" state, as described here:
  1. Open an SSH session using the management IP address of the controller you want to access using a PuTTY/Terminal client.

     Use the **admin** user name and the password for the **admin** account on the H2000 system.
  2. SSH to the QCSP (block storage) VM. Enter:

```
[admin@hseries-2 ~]$ ssh 10.17.21.10 -l root
```

  3. Run the **sbd_list.sh** shell script:

```
[root@hseries-2 ~]# sbd_list.sh
```

     If the system is fenced, you will see results similar to:

```
Device '/dev/disk/by-id/nvme-eui.304e31000000000000000000000000-
part2':
0       1106e9-1        reset   node-1106e92
1       1106e9-2        clear
Device '/dev/disk/by-id/nvme-eui.314e31000000000000000000000000-
part2':
0       1106e9-1        reset   node-1106e92
1       1106e9-2        clear
```

     If any node is fenced, the state keyword “reset” (in green) will be shown in the third column (the second column shows the node name that is affected, in the case of the example above, the first

block storage device (displayed under the "Device" as "0", in green) or Controller A, or the first controller node, is fenced. The fourth column (right-most) shows the node that last wrote to that message block.

Step 2: Log into the "fenced" controller and un-fence it

1. Open an SSH session using the management IP address of the controller you want to access using a PuTTY/Terminal client.

   Use the **admin** user name and the password for the **admin** account on the H2000 system.

2. SSH to the QCSP (block storage) VM. Enter:

```
[admin@hseries-2 ~]$ ssh 10.17.21.10 -l root
```

3. Enter:

```
[root@hseries-23bf46-2 ~]# sbd_unfence.sh
```

4. Enter:

```
[root@hseries-23bf46-2 ~]$ cluster_start.sh
```

5. Run the **sbd_list.sh** shell script again, to verify the controller is no longer fenced:

```
[root@hseries-2 ~]# sbd_list.sh
```

   If the system is NOT fenced, you will see results similar to:

```
Device '/dev/disk/by-id/nvme-eui.304e31000000000000000000000000-
part2':
0       1106e9-1        clear   node-1106e92
1       1106e9-2        clear
Device '/dev/disk/by-id/nvme-eui.314e31000000000000000000000000-
part2':
0       1106e9-1        clear   node-1106e92
```

```
1       1106e9-2        clear
```

If the controller is no longer fenced, the state keyword “clear” (in green) will be shown in the third column. In the example above, controller A ("node 0"), is "clear" or is no longer in a fenced state.

6. Exit the ssh session for the QCSP (block storage) VM. Enter:

```
[root@hseries-23bf46-2 ~]# exit
```

7. Exit the PuTTY/Terminal session:

```
[admin@hseries-2 ~]$ exit
```

H2000 Management Network DNS IP Workaround

To resolve this situation:

1. Open an SSH session using the management IP address of the controller you want to access using a PuTTY/Terminal client.

   Use the **admin** user name and the password for the **admin** account on the H2000 system.

2. Log in as root. Enter:

```
[admin@hSeries-x2 ~]$ sudo su -
```

Note: While you will not see a message confirming that you are now logged in as root, notice that in the CLI prompt, you will see the user change from **admin**, as shown above, to **root** as shown below (in green):

```
[root@hSeries-x2 ~]$
```

3. Open the **/etc/resolv.conf** file using the **vi** command-line editor. Enter:

```
vi /etc/resolv.conf
```

4. Edit the **/etc/resolv.conf** file. Make sure the IP address(es) shown here for each configured nameserver (in Linux each of these is a DNS server) is the IP address of a DNS server that is active and serving as a DNS server (items highlighted in green). In this example, there are two DNS servers, each with its own unique IP address. Edit the IPs as needed.

   Example:

   ```
   search mdc.local
   nameserver 10.20.84.240
   nameserver 10.20.84.241
   ```

5. Save your changes and exit **vi**. Enter:

   ```
   :wq!
   ```

6. Close the PuTTY/Terminal session for this controller node. Enter:

   ```
   exit
   ```

7. Repeat all these steps for the other controller node to edit and save the correct DNS IPs used for the H2000 management network.

Contacting Quantum

Contacts

For information about contacting Quantum, including Quantum office locations, go to:

https://www.quantum.com/aboutus/contactus/index.aspx

For further assistance, or for training opportunities, contact the Quantum Customer Support Center:

Region	Support Contact
North America	1-800-284-5101 (toll free) +1-720-249-5700

Region	Support Contact
EMEA	+800-7826-8888 (toll free) +49 6131 324 185
Asia Pacific	+800-7826-8887 (toll free) +603-7953-3010
For worldwide support: https://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support	

Comments

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

doc-comments@quantum.com