DDN™

DataDirect™ NETWORKS

SFA™ OS

Version 2.0.0

Product Release Notes

SFA12K-40, SFA12K-20, SFA12K-20E, SFA10K-X, and SFA7700

Revision A1
August 2013

Important Information

Information in this document is subject to change without notice and does not represent a commitment on the part of DataDirect Networks, Inc. No part of this manual may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, for any purpose other than the purchaser's personal use without the written permission of DataDirect Networks, Inc.

DataDirect Networks makes no warranties, express or implied, including without limitation the implied warranties of merchantability and fitness for a particular purpose of any products or software. DataDirect Networks does not warrant, guarantee or make any representations regarding the use or the results of the use of any products or software in terms of correctness, accuracy, reliability, or otherwise. The entire risk as to the results and performance of the product and software are assumed by you. The exclusion of implied warranties is not permitted by some jurisdictions; this exclusion may not apply to you.

In no event will DataDirect Networks, their directors, officers, employees, or agents (collectively DataDirect Networks) be liable to you for any consequential, incidental, or indirect damages, including damages for loss of business profits, business interruption, loss of business information, and the like, arising out of the use or inability to use any DataDirect product or software even if DataDirect Networks has been advised of the possibility of such damages by you. Because some jurisdictions do not allow the exclusion or limitation of liability for consequential or incidental damages, these limitations may not apply to you. DataDirect Networks liability to you for actual damages from any cause whatsoever, and regardless of the form of the action (whether in contract, tort including negligence, product liability or otherwise), is limited to the sum you paid for the DataDirect product or software.

TABLE OF CONTENTS

1.0 Overview

This document applies to the SFA OS version 2.0.0.3-17400 release for DataDirect Networks' SFA12K-40, SFA12K-20, SFA12K-20E, SFA10K-X, and SFA7700 storage systems. It details upgrade procedures, enhancements, resolved issues, known issues with workarounds, and recommended practices associated with SFA OS 2.0.0.

"SFA10K" represents SFA10K-X and SFA10K-M.

"SFA12K" represents SFA12K-40, SFA12K-20, and SFA12K-20E.

2.0 Installation

2.1 Overview

SFA OS version 2.0.0 is a new feature release for all supported systems.

2.2 Recommended Code Level Requirements

The recommended controller code levels are shown below.

SFA OS Products	BIOS	BMC	CONFIG	FPGA
SFA7700	23_001	21.86	n/a	001.009-000.035
SFA10K	GEMDV250	DDGEN050	JANUSC23	n/a
SFA12K	21.0L	21.67	n/a	n/a

2.3 Procedure to Request Firmware

To request the new firmware, contact DDN by sending an email to support@ddn.com and include the following information:

Name:
Title:
Company:
Address:
Phone:
Email:
System Type:
System Serial Number:
Firmware Upgrade Requested:

A response will be sent to the provided email with instructions on downloading the requested firmware.

2.4 Pre-Firmware-Upgrade Procedure: SFA12K-40 and SFA12K-20

SFA12K-40, SFA12K-20, and SFA12K-20E

NOTE: Before upgrading SFA OS on an SFA12K system, please read this section carefully.

SFA OS 2.0.0 has a feature that will upgrade the SFA12K Baseboard Management Controller (BMC) firmware version automatically as part of the SFA OS upgrade. The automatic update will only function with SFA12K BMC versions newer than 21.54.

Prior to upgrading to 2.0.0, please identify your current BMC version and then check these notes:

1. If you have BMC version 21.54 or older, you must perform a manual BMC firmware update before you load the SFA OS 2.0.0 firmware. Refer to the *SFA12K BIOS and BMC Firmware Update Field Bulletin*, 96-00344-001.
2. The automatic update will take as long as 10 minutes. During that time, the SFA12K controller will be unresponsive and will not show progress.

WARNING: Do NOT interrupt this upgrade process; doing so could damage the controller.

3. If the BMC version is lower than 21.67, the BMC firmware will be automatically updated.
4. If the BMC version is 21.67 or higher, then the upgrade to SFA OS 2.0.0 will skip the BMC update (and its associated delay).

To identify your current BMC version:

1. View the output from **'SHOW ENCLOSURE'** to identify the enclosures representing your SFA12K-20E chassis.

   For example, see Figure 1.

```
$ show enclosure

************************
*     Enclosure(s)     *
************************

Idx|Pos|Type      |Logical ID        |Vendor ID|      Product ID|Revision|...
-------------------------------------------------------------------------...
  0  NA CONTROLLER 0x0001ff0900160000       DDN          SFA12000   0000   ...
  1   3 DISK       0x50001ff211320000       DDN            SS7000   0502   ...
  2   2 DISK       0x50001ff2114be000       DDN            SS7000   0502   ...
  3   5 DISK       0x50001ff2114ce000       DDN            SS7000   0502   ...
  4   1 DISK       0x50001ff21156a000       DDN            SS7000   0502   ...
  5   4 DISK       0x50001ff2114ba000       DDN            SS7000   0502   ...
  6  NA CONTROLLER 0x0001ff0900100000       DDN          SFA12000   0000   ...

Total Enclosures: 7
```

Figure 1

Issue **'SHOW ENCLOSURE <x> ALL'** to identify your current BMC version. For example, see Figure 2.

```
$ show enclosure 0 all

************************
*     Enclosure(s)     *
************************

Index:                  0
Enclosure:              0x50000000
Type:                   CONTROLLER
Logical ID:             0x0001ff0900160000
Vendor ID:              DDN
Product ID:             SFA12000
Revision:               0000
Is Responsive:          TRUE
Fault Status:           OK
Child Fault Status:     NON-CRITICAL
Zones:                  0
Zoning Configuration:   0
Position:               NOT APPLICABLE
Fault Indicator:        OFF
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Baseboard part number:  X9DRX+-F-DD007
Baseboard serial number:WM21S31110
Baseboard date:         19960101
BIOS version:           21.0k
BMC version:            21.67
Firmware version:       1.5.1.2

Total Enclosures: 1
```

Figure 2

2.5 Post-Firmware-Upgrade Procedure

SFA12K-40, SFA12K-20, SFA12K-20E, and SFA10K

After upgrading the firmware, perform these steps to initialize the battery life remaining feature:

1. Issue the command, **SHOW UPS * ALL_ATTRIBUTES**.

2. If you see a date in "Battery Mfg. Date", do nothing more.

3. If you see the message shown in Figure 3, then issue the command, **CLEAR CONTROLLER x UPS_HISTORY**, where "**x**" is either "**local**" or "**remote**" depending on which UPS needs to be cleared.

```
Battery Mfg. Date: NOT AVAILABLE
Battery Life Remaining: NOT AVAILABLE
```

Figure 3

4. Issue the command, **SHOW UPS * ALL_ATTRIBUTES**. You should see the message shown in Figure 4.

```
Battery Mfg. Date:  Thu Sep 8 4:10:30 2011
Battery Life Remaining: 730 days
```

Figure 4

3.0 Enhancements

SFA OS Version 2.0.0 adds the following enhancements.

- Quality Improvements
  - Ability to mirror system disks on SFA10Ks

    NOTE: For more details on adding disks to the mirror set, see section 2.19 of the *SFA OS CLUI Command Reference Guide*.
  - Ability to monitor wear level of SFA7700 mirrored system disks (SSDs)
  - Ability to monitor wear level of back-end SSDs
- Supportability Improvements
  - Ability to review all attributes of the disk channels for controllers
  - Ability to downgrade SFA firmware to an earlier supported version (1.5.3.1 for this release)
  - Ability to select an option to restart the controller automatically during a firmware upgrade
- API Enhancements
  - Ability to view/set/execute the same properties/functions available in the CLUI using the API
  - Improved ability to gather performance statistics
- Configurations Supported
  - Latent support for the SFA7700
  - Windows Server 2012 Support on SFA7700, SFA10K, and SFA12K
  - Fibre Channel 16 Support on SFA12K
- E-Platform Enhancements
  - SFA12K-20E PXE Boot Support

    Support for booting VMs over the network (PXE) with DirectMon has been added. This feature is available on SFA12K-20E when used in conjunction with DirectMon and enhanced network management (SR-IOV for management Ethernet).

    Please see the *SFA OS User Guide* for more information on this topic.
  - SFA12K-20E Enhanced Management Network MAC Address Reporting

    If you use enhanced management networking on SFA12K-20E, it is now possible to show the generated MAC address for a virtual function interface in a stack

using CLUI. The 'app show stack' command now includes this information within the IOC's output details.

- Customer Controller Replacement Support

  In the event of a controller failure, the process of replacing the failed controller can be now performed by a customer. Details on how to replace a failed controller can be found in the *SFA OS User Guide*.

- Virtual Machine Root on VD Support With SFA12K-20E

  A Virtual Machine can be configured to have its root on a virtual disk instead of a local disk when used in conjunction with DirectMon. To use this feature, separately configure virtual disks in your SFA storage pools and configure DirectMon to boot the VM over the network every time it starts up. If you use this feature, it means you do not need to associate a VM disk image with the VM.

- Increased Memory Available for VMs

  System memory usage was improved to allow additional memory for virtual machines (VMs) on E-platforms. On SFA12K-20E, the amount of memory for a four VM configuration is now 31 GB per VM.

  New installations are pre-configured by the factory to use this additional memory in each VM. Existing configurations upgrading to this release will need to reconfigure their VM to take advantage of this additional memory. To reconfigure and existing VM, use CLUI commands as shown in the following example for a SFA12K-20E:

  ```
  app set stack 0 memory_pending 31
  app set stack 1 memory_pending 31
  app set stack 2 memory_pending 31
  app set stack 3 memory_pending 31
  app set stack 32768 memory_pending 31
  app set stack 32769 memory_pending 31
  app set stack 32770 memory_pending 31
  app set stack 32771 memory_pending 31
  ```

  The new settings won't take effect until next time the VMs/stacks are started. Refer to the GRIDScaler or EXAScaler Admin Guide for further information on shutdown and restarting the VM/stack.

- Increased User Mode Networking Performance, VM Driver Change

  Performance of user mode networking has been greatly improved. User mode networking provides access to the VMs using the controller IP address and port numbers starting at 5555 (for the first VM). The guest OS is now configured to use the e1000 driver instead of the 8139 based driver. Performance is significantly better, especially for using scp to copy large files in to the VM using these special ports. The VM does not need to be re-configured for this change if it has already been installed. See known issues for potential problems with new installations when combined with enhanced network management support (sr-iov).

- SFA Controller Subsystem Shutdown Will No Longer Succeed While VMs are Running

Previously, if you issued a subsystem shutdown on the SFA controller, the VMs were shut down without warning. Now, a check is made and if VMs are running, the shutdown command will fail with an error message. This will prevent virtual machines from being shut down by accident.

- Presenting a VD to all VMs/all Hosts no Longer Fails if Some Presentations Exist

  Previously, if you had one or more presentations already set up to one or more hosts, the following command would fail with an error stating that some presentations were already defined:

  ```
  app create presentation host ALL  vd 0
  ```

  This has been resolved. Issuing the above command will now create the presentation for a given VD to all hosts even if some presentations already exist.

- Ability to See When an Image is Ready

  The CLUI **`app show image`** command now informs you if the image is ready for a stack (VM) and can be started.

  Previously, it was difficult to know if an image was ready to be used by a stack/VM. You would have to start the stack to see if it produced an error message to know if the image was fully allocated or not.

  Now this information is provided as part of the output of the **`app show image`** command—including the percent complete.

- The Mellanox firmware on the SFA12K Infiniband ICL is upgraded from 2.10.700 to 2.11.500 keeping it up to date with the front-end Mellanox HCAs on the SFA12K.

4.0 Resolved Issues

These issues have been resolved.

- **Errors on one drive cause other drives to go missing**

  An issue was found in the drive correction firmware where the wrong drive is power cycled in an attempt at recovery. This would occur when a drive was in an error state and not responding or when it was physically pulled from the enclosure. This caused subsequent incorrect power cycles to occur because the drive that was incorrectly power cycled is now not reporting. As a result, pools will fluctuate between normal, degraded, non-redundant, and AWL as drives come back from the power cycling.

- **Pool inits on systems where pools were created with GUI or API take excessive amount of time to complete**

  If the GUI or API was used to create the pools, two commands were sent to the controller instead of one. This resulted in the DIF init being a non-destructive init (read-modify-write operation instead of zeroing out the DIF field).

- **InfiniBand ICL issues can cause controller watchdog crash**

  When the ICL HCA experienced issues, a timer in the SFA OS code would go off and would attempt to correct the HCA. However, this error correction on the InfiniBand ICL was done by resetting the controller to reset the HCA. This has now been changed so that the ICL channel will be taken down and back up in an attempt to correct the issue. If this fails, that channel will be taken down and the ICL will run in a degraded fashion on the other channels.

- **Controller crash after shutdown subsystem abrupt command was issued**

A controller crash occurred when the **`shutdown subsystem abrupt`** command was issued on controller 0 and then again on controller 1. This was due to an issue where the second controller forwarded the second command to the master controller while it was trying to shutdown.

- **SFA12K controller crash with JEX_Storage_EnqueueRequest**

  This issue was due to a problem in the logging agent inside the SFA OS.

- **Latency on SSD virtual disks without corresponding physcial disk latency**

  High levels of latency were observed under certain workloads in the use of SSD pools/VDs without an associated PD latency. This was due to small reads getting blocked waiting for buffers due to the large read workload ahead of it.

- **LOG_ES_SES_COMMAND_QUERY_FAILURE repeating every 30 seconds**

  This issue was reported on an older SFA-10K with AAMUX dongles. This was due to path changes on the SES device not being registered correctly.

- **After DEM upgrade the VDs continue to all stay on one controller**

  This issue occurred when the secondary controller is in MIR state NO BACKEND DRIVES and the primary controller is activating the JIS.

- **Show subsystem summary all shows bad block numbers with negative values**

  The show subsystem summary all would show bad block numbers with negative values while the 'show vd X all' command, the bad block count shows up as zero instead of the large negative number. This was a reporting issue caused by a problem in the SFA OS.

- **Show internal_disk incorrectly reports disks as not present**

  A reporting issue with the show internal disk command was reported where disks were reported as not present when they were present. This was due to an issue in SFA OS where stale data was used by the reporting function.

- Initiator logout and subsequent attempt to abort outstanding I/O caused controllers running SFA OS 1.5.3.0 to crash. This was due to a bug in the SCSI persistent reservation code that accessed memory beyond its array size.

- Inside a VM the driver for Mellanox ConnectX3 and ConnectX2 HCAs failed to load with error message "SW2HW_MPT failed."

- Previously, binding two or more Mellanox HCAs to a VM often showed one less.

- Previously, when you set a stack's CD_DVD_DEVICE attribute, you would have to unset it so that the VM doesn't attempt to re-install itself a second time. This is no longer necessary. The SFA controller now detects when the operating system ejects the DVD and then disassociates it automatically from the stack. ISO installers, including those for GRIDScaler and EXAScaler, eject the DVD at the end of installation.

5.0 Known Issues

5.1 Common

- The drive serial number reported in **SHOW PD** output may contain additional characters that are not on the disk label. The serial number on the disk label will always be contained within the potentially more verbose string reported by SFA OS.
- The SS7000 enclosure does not support the **SET SLOT xx POWER OFF** command with enclosure firmware versions prior to 05.02.01.
- In the SS7000, when removing an I/O module or cable under load, it is possible that one or more disks could be marked failed which would change the redundancy of the storage pools. This is due to in-flight I/O being timed out and retried before the IOC reports the drive as missing. There is a high probability that the newly failed drives are fine and can be returned to service by clearing the failed state and reassigning them to pools.

  Use the following procedure to clear a failed drive and reassign it back to a pool:

  show unassigned failed ←Shows which PDs have failed; get drive's <id>
  show pool ←Shows which pools are involved
  clear pd <id> failed ←<id> is from `show unassigned` output
  assign pd <id> to pool-id set ←Pool-ID is shown in the `show pool` output

  NOTE: If the pool sparing policy is set to 'SWAP', the SFA will typically auto rebuild the drives, using fractional or full rebuild depending on the duration of the interruption.

- If stack commands (CLI commands starting with 'app') are issued shortly after rebooting a controller, you may encounter communication errors displayed in the CLI, and failure of the attempted commands. Errors may include:
  - Long delays followed by:

    ```
    Communication connection failed for this command
    ```

  - Or this message:

    ```
    ASM initialization in progress
    ```

  Suggested Work Around

  - After the failed controller comes back up, do not immediately issue any CLI commands besides "`show controller`".
  - Connect to the remaining controller and issue this command:

    ```
    $ show controller
    ```

  - When "`show controller`" shows the remote controller in the output, focus on the ULA field in the output. The remote controller won't show up until it is fully booted.

- If the ULA field has the value "0000000000000000", continue to wait and not issue any CLI commands besides "show controller".

  Here is an example output of the condition where you would want to wait:

```
$ show controller

************************
*     Controller(s)     *
************************


                                    |   Up Time   |                  |Encl|               |...
Idx|Name          |Mastership|Locality|   D: H: M: S|RP|       ID       |Idx |       ULA      |...
-------------------------------------------------------------------------------------------...
  0 A              PRIMARY     LOCAL   0000:20:24:34  1 0001ff0900180000   0  00000001ff0800ac ...
  1 B              SECONDARY   REMOTE  0000:00:00:05  1 0001ff09002d0000   0  0000000000000000 ...

Total Controllers: 2
```

- When the ULA number has something other than all zeros for the remote controller, it is now safe to issue CLI commands and avoid the communication problems.

5.2 SFA12K-20E

- **VMs can be shown as running when controller hosting those VMs goes down**

  In certain situations, if one controller in a couplet goes down, the remaining controller may show the VMs in the offline controller as running. This is because the remaining online controller cannot update the status of the running VMs on the offline controller. In this scenario, the online controller continues to display a healthy running VM status. Once the offline controller comes back online, the correct state of the VMs will be visible. [DE4409]

- **SCSI DID_NO_CONNECT errors**

  When running GRIDScaler 1.6.1 (OFED 1.5.3.4.0.35), the following erorrs could be observed.

  ```
  Apr 12 16:26:42 icndcm10 kernel: sd 8:0:0:1: [sdag] Result:
  hostbyte=DID_NO_CONNECT driverbyte=DRIVER_OK
  ```

  This problem is fixed in GRIDScaler 2.0.0. There is also a field upgrade patch available. Contact DDN Support if you need this patch.

- **Clocksource tsc unstable errors [DE4620]**

  On SFA12K-20E controllers, error messages similar to the following may be reported on the running VM with dmesg or in system logging:

  ```
  Clocksource tsc unstable (delta = -519494463 ns). Enable
  clocksource failover by adding clocksource_failover kernel
  parameter.
  ```

  A workaround has been provided for GRIDScaler 2.0.0, so this error should no longer appear. For other releases, the following workaround can be applied to have the kernel fallback to HPET mode if the guest OS detects a problem with TSC. In the kernel parameter list of the VMs, add the following line to the VM kernel boot line in /boot/grub/grub.conf.

  ```
  clocksource_failover=1
  ```

- **If Installing Older DDN Solutions on SFA12K-20E, Potential Network Ordering Issues May Arise**

  If you install a version of GRIDScaler older than 2.0.0 or EXAScaler older than 1.6.1 and you are using enhanced network management in the VMs (sr-iov virtual function interfaces), you could run in to an Ethernet ordering issue during installation where eth0 may not be associated with the expected driver. This could result in failure of user mode networking (the port 555x access) or other problems.

  A workaround is to temporarily un-bind the virtual function interfaces from the stack using CLUI. Once installation is complete, re-associate the virtual function interfaces with the VMs and then configure networking.

  IOC index 01792 is associated with the first SFA controller's first Ethernet port. IOC index 01793 is associated with the first SFA controller's second Ethernet port IOC index 34560 is associated with the second SFA controller's first Ethernet port IOC index 34561 is associated with the second SFA controller's second ethernet port

  The following CLUI example would unbind the virtual function interfaces from stacks configured by default on systems leaving DDN:

```
app unbind stack 0 ioc 01792
app unbind stack 0 ioc 01793
app unbind stack 1 ioc 01792
app unbind stack 1 ioc 01793
app unbind stack 2 ioc 01792
app unbind stack 2 ioc 01793
app unbind stack 3 ioc 01792
app unbind stack 3 ioc 01793
app unbind stack 32768 ioc 34560
app unbind stack 32768 ioc 34561
app unbind stack 32769 ioc 34560
app unbind stack 32769 ioc 34561
app unbind stack 32770 ioc 34560
app unbind stack 32770 ioc 34561
app unbind stack 32771 ioc 34560
app unbind stack 32771 ioc 34561
```

  At this point, to complete the installation, start up your stacks as you normally would using CLUI.

  After installation is complete on all the VMs, you can re-bind the virtual function interfaces and shutdown/startup the VMs to make use of the interfaces. Here is an example using CLUI:

```
app bind stack 0 ioc 01792
app bind stack 0 ioc 01793
app bind stack 1 ioc 01792
app bind stack 1 ioc 01793
app bind stack 2 ioc 01792
app bind stack 2 ioc 01793
app bind stack 3 ioc 01792
app bind stack 3 ioc 01793
app bind stack 32768 ioc 34560
app bind stack 32768 ioc 34561
```

```
app bind stack 32769 ioc 34560
app bind stack 32769 ioc 34561
app bind stack 32770 ioc 34560
app bind stack 32770 ioc 34561
app bind stack 32771 ioc 34560
app bind stack 32771 ioc 34561
```

Now, shut down and start up the VMs/stacks again using CLUI. After they boot up, you can configure networking on the virtual machines.

- **Suggested Ethernet Ordering in Virtual Machines for SFA12K-20E**

  For new installations, DDN recommends the following virtual machine Ethernet ordering starting in SFA OS 2.0.0.

  - VM eth0: virtual function interface associated with controller's first port
  - VM eth1: virtual function interface associated with controller's 2nd port
  - VM eth2: user mode networking support (provides port 555x ssh functionality)

  Ethernet interface numbers larger than eth2 may be associated with Mellanox cards operating in 10G/40G mode.

  Controller eth0 (and VM eth0) are meant to be connected to the customer network and may also be used by DDN Solutions like GRIDScaler for network management. Controller eth1 (and VM eth1) are reserved for DirectMon.

  Starting with EXAScaler 1.6.1 and GRIDScaler 2.0.0, this naming and usage convention is used and configured as part of the ISO installation process.

  The *SFA OS User Guide* has more information on Ethernet interface configuration for controllers and VMs.

- ConnectX3 HCA IB read and write performance drops on PCI bus 1 when the message size is more than 128 KiB.

5.3 SFA12K and SFA10K

- When an SFA is upgraded from SFA OS v1.5.3 to v2.0.0.3, initiators running RHEL 5.x may get I/O errors. These errors can occur during a small window of time (typically less than 5 seconds) when the controller running v1.5.3 is preparing to reboot and responds 'Ready' to a Test Unit Ready inquiry from the initiator. If the host application re-issues the I/O, I/O will continue as normal.
- To improve failover time with RHEL 6.x and Qlogic cards, the following is recommended to enable the QLogic driver to react quickly to a loss of path :
  Qlogic driver: 8.04.00.06.06.0-k
  Add the parameter 'ql2xasynctmfenable=1' to /etc/modprobe.d/modprobe.conf
  Example line in modprobe.conf:
  options qla2xxx ql2xfwloadbin=1 ql2xqfullrampup=15 ql2xasynctmfenable=1

- To improve failover time and to prevent I/O errors, the following settings in multipath.conf are recommended:

  (These settings can be changed for the defaults section which will apply to all devices, or just under the SFA devices):

  checker_timeout 5
  dev_loss_tmo 10
  fast_io_fail_tmo 5

  These settings are included in the DDN multipath package version 1.5-5 and above.
- When the controller is preparing to shutdown, it will first put all its pools into write-through mode and attempt to flush all the dirty cache. On SFA platforms with multiple RAID processors (RP), in the case where one RP finishes flushing its cache before the other RP, the first RP to finish flushing its cache will not service I/O from the host until the SFA reboots. This may cause I/O errors on the host and cause applications on the hosts to hang and eventually time out. To work around this issue, reduce I/O load during planned maintenance activities such as firmware upgrades and reboots of the controllers so that the flush activities will complete quickly.
- With RHEL6.2 and OFED 1.5.4.0 in an IB switch attached environment, it is possible that a virtual disk on a controller may not be added back to the multipath device map after a failover.

  To find the offline device, issue the command:

  lsscsi | awk -F/ '{print $NF}' | while read a; do printf "%s " $a; cat ,/sys/block/$a/device/state; done

  To bring the device back online, issue the command:

  #echo running > /sys/block/<*sd??*>/device/state and #multipath -r

  where you replace <*sd??*> with the appropriate sd, for example *sdaf*, found from the previous command.

  A workaround for this issue is to update these packages:

  - **device-mapper:1.02.74-10.e16**
  - **device-mapper-multipath: 0.4.9-56.el6_3.1.x86_64**

  DDN recommends that you install these as soon as possible.

- In an InfiniBand switch environment, there is a small chance that if a cable between the switch and the controller is pulled, both the physical link and the logical link will be lost. This has occurred with a Mellanox 6025F Switch and a Mellanox HCA.
- In an InfiniBand switch environment running RHEL 5.7, if a cable is pulled from either an initiator or a target, a failover occurs as expected; however, once the connection is reestablished, it does not fail back to the original controller.
  - In order to resolve this issue, you must issue the command, **# udevtrigger**.
  - The Linux man pages state that the command will simply “request kernel devices events for coldplug”. This will force udev to send a notification for newly discovered path, which will allow the multipath daemon to detect that the path has returned. After the multipath discovery takes place, I/Os can be rebalanced back to allow the preferred paths to be used.

- In an InfiniBand switch environment, an online upgrade causes an I/O error while upgrading the second controller with SLES 10 SP4. In an FC environment, there are no issues because the discovery rate is much quicker. The only option for SFA upgrades with SLES 10 SP4 is to perform the upgrade offline.

  To perform the upgrade offline,

  1. Issue the CLI command **`set subsystem offline.`**
  2. Perform the controller firmware upgrade on both controllers.
  3. Issue the CLI command **`shutdown subsystem restart`**.
  4. After both controllers come back online, issue the CLI command **`application show discovered_initiator.`** This should report "No discovered_initiators exist." This will be the state until a read request for information from those SCSI devices is issued. Once the devices are accessed, there will be a small delay as they are rediscovered by the system.

5.4 SFA12K

- Only 180 (not 256) QLogic FC16 initiators are currently supported. This restriction will be lifted in a future release. The number of FC8 initiators supported remains at 256.

5.5 SFA10K

- The SS2460 does not support SLOT POWER OFF feature and so the **SET SLOT POWER** command has no effect on drives in the SS2460.
- The Management Ethernet port (LAN4) may boot running at 10 Mb/sec half-duplex, instead of 100 Mb/sec or 1000 Mb/sec full-duplex. Typically, full Ethernet port speeds can be restored by unplugging and replugging the Ethernet cable to the controller port experiencing the slow speed. If the issue persists, please contact, support@ddn.com.
- On SS6000 enclosures the **DISPLAY POWER_SUPPLY** command may fail to display the last three characters of the power supply serial number. This is only a display issue, the command will complete successfully with no adverse effects other than the truncated results.
- The SS6000 may encounter a problem where show enclosure commands will return the following error message: "A command issued to an enclosure to get attributes for the UI timed out." This is an issue in the enclosure firmware v3.001 that is resolved by upgrading to enclosure firmware v3.010.
- Powering off drives is not supported with 6KSSI interposers in the SS6000 enclosure. If you issue the command, **SET SLOT <***encl-id slot-id***> POWER OFF**, the system erroneously displays a success message. However, the drive power is **not** affected by the command.

5.6 SFA7700 Early Access

- Ensure that each controller is connected to a network with an NTP server prior to booting. A boot without a network connection can cause a BMC hang.
- Do not remove a controller canister while it is running unless otherwise instructed to do so.
- Power Button Operation is not fully implemented yet. There may be instances where the power button on the canister does not operate as expected and may require the canister to be removed and powered down outside of the enclosure.

Known issues with the SFA7700 are presented in the following tables. Also included are the issues' frequencies, impact on customers, and resolutions.

BIOS: Build 23_001

Issue	Freq.	Customer Impact	Resolution
Warm boot hang at step B2	Rare	Controller restart required.	Manually power-cycle controller by unseating the canister from the enclosure for 60 seconds, then reseat. If this does not turn off the canister, try pressing the hidden power button.

BMC: Build 21.86

Issue	Freq.	Customer Impact	Resolution
Canister fault light is on but there is no fault.	Rare	Canister fault light is on, but SHOW Temp, SHOW Fan does not show a fault.	Check CPU0_Temp. If zero, this is triggering the fault (but is not a failure). Contact DDN Support.
After a controller shutdown the BMC WebUI indicates that the power is on.	100%	When using the BMC WebUI remote power feature, the WebUI will indicate host power is still on after a SHUTDOWN CONTROLLER/SUBSYSTEM is issued.	Switch to another page in the WebUI and then back to the remote power control page, this will update the page correctly.
BMC hangs during boot if not connected to a NTP server.	100%	Canister boot hangs for 6 min. during Linux boot if primary Ethernet is not connected to a network with an NTP server.	Connect primary Ethernet to a network with a NTP server.

FPGA: Build 001.009-000.035

Issue	Freq.	Customer Impact	Resolution
Reset pushbutton function cannot be disabled	100%	A short (1 second) pushbutton push will cause the controller to restart without saving cache. If only one controller is running then data is lost.	Be careful not to push the pushbutton when the controller is running.
Canister does not automatically power off when removed from the disk enclosure.	100%	Canister continues to run on battery when removed from the disk enclosure.	Pressing the power button may turn off the canister. If the power button does not work then unplugging the battery pack will power down the canister.

Issue	Freq.	Customer Impact	Resolution
Restarting a controller does not disable the battery pack.	Rare	If a power failure occurs during a controller restart (SHUTDOWN RESTART or SFA OS Update) then the controller will run on battery power until it completely drains the battery pack.	Pressing the power button may turn off the canister. If the power button does not work then unplugging the battery pack will power down the canister.
Spurious FPGA failure	Rare	When the FPGA crashes SFA OS will report that the controller's UPS interface has failed. When in this condition, SFA OS will not detect a power failure and save its cache, so a power failure may result in data loss.	Perform a SHUTDOWN SUBSYSTEM to force SFA OS to save its cache to the back-end disks. Then, a hard power cycle is required to reset the FPGA. Since the FPGA is crashed, there may be no way to do a hard power off without disconnecting the battery.
FPGA warm boot hang.	Rare	The controller fails to restart on a SHUTDOWN CONTROLLER RESTART or SHUTDOWN SUBSYSTEM RESTART. In this case, no BIOS boot step is displayed.	Manually power cycle the controller by pulling and reseating the canister.
SHUTDOWN does not shut down controllers (continue to run on battery).	~50%	After a SHUTDOWN CONTROLLER or SHUTDOWN SUBSYSTEM, the canister appears to shut down but if enclosure power is removed the canister continues to run.	The power button may turn off the canister. If not, manually power cycle the controller by pulling and reseating the canister.
IPMI Power Off does not always turn the canister completely off.	~30%	After a SHUTDOWN CONTROLLER OR SHUTDOWN SUB-SYSTEM, the canister fan may restart even though the canister is in the shutdown or standby state.	This condition is benign except that it makes it confusing whether the canister is shutdown or not. If all five canister LEDs are off, then it did indeed shut down.
A long pushbutton push does not force the canister to power off.	100%	A long pushbutton push is supposed to abruptly power off the canister, but this does not work.	Pull the canister and, if necessary, use a short pushbutton push.
FPGA update may restart	Rare	After a BMC/FPGA	Reseat canister or power off

Issue	Freq.	Customer Impact	Resolution
controller instead of holding it in reset		update the FPGA version reported will be incorrect (0.255-0.255)	disk enclosure for 60 seconds.
SEP reset signals false power failure.	100%	SHUTDOWN ENCLOSURE 1 RESTART results in a false power failure which causes both canisters to save cache and restart. This is only a problem for the *head disk enclosure*	SHUTDOWN SUBSYSTEM and manually power cycle the head disk enclosure.

5.7 Red Hat 6.2

This section applies to the SFA12K-40, SFA12K-20, SFA10K-X, and SFA10K-M.

- Updating the kernel and multipath packages to the latest version is recommended. As of this release, minimum DDN qualified versions include:
  - the kernel version - 2.6.32-279.2.1.el6.x86_64
  - the multipath version - 0.4.9-56.el6_3.1.x86_64.
- The kernel modules that come with OFED 1.5.4 will cause I/O threads to hang when the data path between the controller and switch is lost. A workaround for this issue is to use the Red Hat supplied InfiniBand solution. For a switched InfiniBand environment:
  - Install the in-box OFED drivers that come with RHEL 6 Update 2
  - Update the multipath-tools package to a minimal version of 0.4.9-56.el_3.1
  - Do not use user-friendly names in multipath.conf
  - Configure RHEL 6 Update 2 to a minimal kernel version of 2.6.32-279.2.1.el6
- These packages are required for InfiniBand support on Red Hat. (Use the command, **`yum groupinstall infiniband`** to install all IB packages):
  - ibutils-1.5.7-6.el6.x86_64
  - infiniband-diags-1.5.12-3.el6.x86_64
  - opensm-libs-3.3.13-1.el6.x86_64
  - opensm-3.3.13-1.el6.x86_64
  - libmlx4-1.0.2-5.el6.x86_64
  - srptools-0.0.4-15.el6.x86_64
- To start IB lun discovery, issue these commands:
  1. **modprobe ib_umad**
  2. **modprobe ib_uverbs**
  3. **modprobe mlx4_ib**

4. **opensm -g <***port GUID***>** (where *port GUID* is the port GUID from ibstat output
5. **run_srp_daemon -R 20 -T 10 -nce -i <***ca_name***> -p <***port***>** (where *ca_name* is the channel adapter name and port is the *port* number)

6.0 Recommended Practices

The recommended best practices for SFA OS version 2.0.0 include:

- **SS6000 DEM and I/O module replacement.**
  For SS6000 DEM and I/O module replacement, see the *SFA OS User Guide* for the procedures.
- **Reduce host I/O load during background initializations.**
  When a storage pool is in the process of being initialized, a significant internal I/O load is generated. To avoid poor performance, DDN recommends that heavy host I/O loads be avoided until pool initialization has completed.
- **On hosts that are directly connected, if the configuration is cleared, the Fibre Channel initiators may not log back into the subsystem.**
  This results in no initiators being available to assign to new presentations. To correct this, do a rescan from the host before setting up your presentations.
  This does not apply to the SFA12K-40 and SFA12K-20.
- **Disk firmware upgrades.**
  Before performing a disk firmware upgrade, scp the .DDN file to both controllers. This avoids a situation where a controller attempts to perform the upgrade but can't find the .DDN file.
- **Firmware upgrades should be done during low usage periods whenever possible.**
  As upgrading your system will interrupt I/O operations as paths fail over, be aware of the possibility of increased I/O latencies during the upgrade operation. Prior to upgrading your firmware, DDN recommends reviewing the system logs to ensure the system is running cleanly. DDN also recommends allowing any rebuilds and forced verify jobs to complete and pausing any running verify jobs.

  When a controller is shut down, all cached data must be flushed before the controller will shut down. This operation can take a considerable amount of time depending upon the amount of cache and if the cached data is sequential vs. random, large block vs. small block, number of VDs, and so on. Should the shutdown fail, an event will be generated. If you do not see the event, the controller is still shutting down. Please allow it to complete to avoid possible data loss issues.
- **Update the secondary controller first.**
  When performing an SFA OS upgrade, upgrade the secondary controller first. Doing this requires primary to move once instead of twice.

7.0 Compatibility Guide

The *SFA OS Compatibility Guide* lists the client operating systems, HBA and HCAs, enclosures, and disk drives that are supported by SFA OS. The guide is available at http://www.ddn.com/en/support/product-downloads-and-documentation.

NOTE: The Compatibility Guide applies to SFA12K-40, SFA12K-20, SFA10K-X, SFA10K-M, and SFA7700.

8.0 Support

To submit debug information to DDN, use the **`diag tgz`** command.

`Diag` is a captive account on SFA OS systems that is used for performing operations on a controller when there's no CLI access or the CLI isn't appropriate. Perform diag operations by using ssh to login to the diag account.

The command syntax is:

`ssh diag@xxx.xxx.xxx.xxx` *`diag-command`*

(where xxx.xxx.xxx.xxx is the controller's IP address and diag-commands are listed in Table 1.)

Table 1 Valid diag commands and descriptions

Command	Description
help	Displays the help information.
clear	Clears all old crash files (it will clear a recursive crash).
reboot	Used to restart the controller if the RAID stack is not running. If the RAID stack is running, this command will fail and you should use the CLI to perform this operation.
shutdown	Used to shutdown the controller if RAID stack is not running. If the RAID stack is running, this command will fail and you should use the CLI to perform this operation.
tgz	Creates a .tgz file of all the crash information. The file is output to stdout and since it's a binary file, having it display on your terminal isn't desirable. Redirect stdout to your workstation with a command like: `ssh diag@xxx.xxx.xxx.xxx tgz > crash.tgz`

Please contact DataDirect Networks Support at any time for assistance. Support can be reached by these methods:

Web

Support Portal	https://portal.ddnsupport.com
Portal Assistance	webportal.support@ddn.com
Technical Support Bulletins	http://www.ddn.com/en/technical-support-bulletins

Email

Support Email	support@ddn.com
Mailing List Subscriptions	https://ddntsr.com/mailman/listinfo/tsr-announce

Telephone

DDN Worldwide Directory	http://www.ddn.com/en/support/contact-support

World Headquarters 9351 Deering Avenue • Chatsworth, CA 91311 | ddn.com | Phone: +1.818.700.7600 Fax: +1.818.700.7601