

SFA14K™

Spontaneous Reboot or Controller Hang Every 84 Days

Issue Summary and Root Cause

Certain controllers in the SFA14K family may hang or reboot spontaneously every 84 days. The root cause of this issue is a PCIe reference counter that reaches a maximum count and then rolls over to zero. The rollover to zero may cause a reboot or a hung controller due to a null pointer reference. Based on the frequency of accesses to the driver that controls this reference counter, the rollover will occur approximately every 84 days.

Affected Systems

This issue affects SFA14K storage arrays running SFA OS versions 3.1.0.1 through 3.1.1.0. Only the following SFA14K models and configurations are affected:

- SFA14K systems with an internal PCIe inter-controller link (ICL) installed instead of an external Infiniband ICL
- All SFA14KX and SFA14KXE systems

NOTE SFA14KE systems are not affected.

Resolution

This issue is resolved in SFA OS 3.1.2, which is expected to be released for general availability by the

end of March 2017. Upgrade will be **mandatory** once this version is released.

Please contact DDN Technical Support to create an upgrade plan.

Workaround

If you have an affected SFA14K system and **system uptime is greater than 60 days**, the workaround for this issue is to restart each SFA controller on the system separately to reset the PCIe reference counter. Repeat this procedure every 60 days until SFA OS 3.1.2 can be installed.

ALERT! Do not allow system uptime to exceed 60 days.

Procedures are summarized below for:

- [Determining Storage System Uptime](#)
- [Performing an Online Controller Restart](#)

Determining Storage System Uptime

To determine the current uptime for each controller, log on to the system as user and enter the command:

```
show controller
```

Uptime is shown separately for each controller in the **UP TIME** column. This value is shown in dddd:hh:mm:ss format, where dddd is the number of days.

Performing an Online Controller Restart

This procedure resets the PCIe reference counter on each SFA controller by restarting each controller separately online. The **secondary** controller at the time the procedure begins will be restarted first, followed by the **primary** controller.

ALERT! SFA is a virtual architecture, and the physical controllers deemed **primary** and **secondary** will exchange roles during the course of the procedure. Controller **index numbers** could change as well. Throughout the procedure, be sure you are working with the correct **physical** controller, not a functional role or controller index number that may change.

This procedure makes certain assumptions that are essential for successful completion. We assume:

- **The restart process is managed entirely by SFA OS.** Do *not* use the power button on the controller to reset the equipment. Do *not* disconnect any power cords. Do *not* switch off any power distribution units in the rack. Doing any of these things may cause data loss.
- **The storage system is healthy and high-availability (HA) services are configured.** Performing an

online controller restart poses no risk to data if these requirements are met. Please verify that your SFA14K platform is healthy, your Multipath configuration fails over correctly, and the HA features of your file system are enabled and properly configured. On the SFA14KXE, be sure your virtual machine (VM) stacks are configured for automatic restart. (This is the default.)

NOTE GRIDScaler customers with embedded file systems on the SFA14KXE should ensure that the file system(s) are configured for autostart and automount. The file system must be configured with a quorum server on each controller and a tiebreaker disk.

Please refer to the *SFA OS User Guide* or contact DDN Support if you need assistance with any step.

Step 1. *Log in to the current **secondary controller*** as administrator.

NOTE The *physical* controller that *currently functions as the secondary controller* in the SFA14K couplet will be referred to as "**controller 1**" throughout this procedure. Your system may differ.

Step 2. *If the storage system is an SFA14KXE appliance* (that is, it runs embedded applications such as file system software on the storage itself), gracefully shut down the VM application stacks on **controller 1** (the current secondary controller) as follows.

a. List the VM stacks running on the subsystem with the CLUI command:

```
app show stack * all
```

All stacks are listed with their index numbers and associated controllers.

b. Gracefully shut down the VM stacks with the command:

```
app shutdown stack <stack_index>
```

where **<stack_index>** is the index number of a VM stack on **controller 1**. Repeat for each controller 1 VM stack identified by the **app show stack** output in [Step 2a](#).

NOTE There is no need to unmount or shut down the file systems on each VM individually before shutting down the VM stacks.

Step 3. *Perform a graceful shutdown and restart of controller 1* (the current secondary controller) with the CLUI command:

```
shutdown controller local restart
```

This command instructs SFA OS to pause any rebuild or initialization procedures and flush all data in cache to disk before shutting down the controller in an orderly way. These processes may take a considerable amount of time.

The **restart** parameter specifies that an **automatic** restart of the controller should then occur without manual intervention. This process will also take some time.

You can monitor the progress of the restart by logging in to **controller 0** (the current **primary** controller) and running the command:

```
show controller +30
```

The **+30** parameter repeats the **show controller** command every 30 seconds. As the shutdown and restart operations proceed, you will see controller 1 (the current secondary controller) disappear from the controller list, then reappear.

Wait until you see log messages indicating the restart is successful and host communication ports have been brought up on controller 1. At this point, SFA OS is operational and the PCIe reference counter on this controller has been reset.

ALERT! Do not proceed further until the restart cycle has completed and **controller 1** (the current secondary controller) returns to a fully operational state.

Step 4. *If the storage system is an SFA14KXE appliance*, verify that the VMs have restarted on controller 1 (the current secondary controller) by running the command:

```
app show stack * all
```

Also log in to each VM on controller 1 and verify that all EXAScaler object storage servers (OSSs) or GRIDScaler network shared disk (NSD) servers are running. Refer to the *SFA OS User Guide* for VM IP addresses.

NOTE If the VMs were configured for automatic startup as advised above, there should be no need to manually restart the VMs or to start file system software on each VM individually after the controller resumes operation.

ALERT! Do not proceed further until you have verified that the VMs have restarted on controller 1 and the file systems are fully operational.

Step 5. *Log in to the current primary controller* as administrator.

NOTE The **physical** controller that **currently functions as the primary controller** in the SFA14K couplet will be referred to as "**controller 0**" throughout this procedure. Your system may differ.

Step 6. *If the storage system is an SFA14KXE appliance*, gracefully shut down the VM application stacks on **controller 0** (the current primary controller) with the CLUI command:

```
app shutdown stack <stack_index>
```

where **<stack_index>** is the index number of a VM stack on **controller 0**. Repeat for each controller 0 VM stack shown in the **app show stack** output obtained in [Step 2a](#).

NOTE There is no need to unmount or shut down the file systems on each VM individually before shutting down the VM stacks.

Step 7. **Perform a graceful shutdown and restart of controller 0** (the current primary controller) with the CLUI command:

```
shutdown controller local restart
```

As with the secondary controller, this operation may take some time.

Monitor the progress of the restart by logging into controller 1 (the current **secondary** controller, but soon to become the primary controller), then running the command:

```
show controller +30
```

As the shutdown and restart operations proceed, you will see controller 0 disappear from the controller list while controller 1 becomes the primary controller. Controller 0 will later reappear as the new secondary controller.

Wait until you see log messages indicating the restart is successful and host communication ports have been brought up on controller 0. At this point, SFA OS is operational and the PCIe reference counter on this controller has been reset.

ALERT! Do not proceed further until the restart cycle has completed and **controller 0** (the current secondary controller) returns to a fully operational state.

Step 8. **If the storage system is an SFA14KXE appliance**, verify that the VMs have restarted on controller 0 (the current secondary controller) by running the command:

```
app show stack * all
```

Also log in to each VM on controller 0 and verify that all EXAScaler object storage servers (OSSs) or GRIDScaler network shared disk (NSD) servers are running. Refer to the *SFA OS User Guide* for VM IP addresses.

NOTE If the VMs were configured for automatic startup as advised above, there should be no need to manually restart the VMs or to start file system software on each VM individually after the controller resumes operation.

Step 9. *Verify that the system remains healthy* by running the following commands:

```
show sub sum
```

```
app show sub sum
```

Contact DDN Support if you have any questions.

Contacting DDN Technical Support

Please contact DDN Technical Support at any time if you have questions or require assistance. Support can be reached by phone, by email, or on the web as listed below.

Web

*DDN Community Support Portal
Portal Assistance*

<https://community.ddn.com/login>
webportal.support@ddn.com

Telephone

DDN Support Worldwide Directory

<http://www.ddn.com/support/contact-support>

Email

Support Email

support@ddn.com

Bulletins

Support Bulletins

End-of-Life Notices

Bulletin Subscription Requests

<http://www.ddn.com/support/technical-support-bulletins>

<http://www.ddn.com/support/end-of-life-notices>

support-tsb@ddn.com