

# DataDirect<sup>TM</sup>

## NETWORKS



## **S2A6620**

### **Recovery Procedure for Missing I/O Module in Base Enclosure**

#### **Technical Support Bulletin**

Revision B

24 September 2014

## 1.0 Overview

### 1.1 Issue Summary

The S2A6620 can get into a state where a SAS Expander Processor (SEP) in a base unit controller is reported as “NOT INSTALLED”, although it is physically present. While in this state, the S2A6620 may start to fail disks, which can put pools in a non-redundant state. If subsequent I/O to a nonredundant pool should fail, the pool will be placed in Auto Write Lock (AWL) mode, and data could become inaccessible to file systems and applications.

This document shows how to identify this state and describes the procedure for recovery.

### 1.2 Affected Systems

All S2A6620 storage systems are potentially affected by this issue.

## 2.0 Problem Description

The S2A6620 can get into a state where a SAS Expander Processor (SEP) in a base unit controller is reported as “NOT INSTALLED”, although it is physically present. The most visible symptom of this condition is a **SES Status** of “NOT INSTALLED” in the output of the `show expander all` command. For example:

```
*****
*      Expander(s)      *
*****

| Indicators |
Idx|Encl|Pos|Present|Failure|Locate|Location|SES Status
-----
 1  1  1  TRUE  OFF  OFF  IOM 1 OK
 2  1  2  FALSE OFF  OFF  NOT INSTALLED
 3  1  3  TRUE  OFF  OFF  DEM 1A OK
 4  1  4  TRUE  OFF  OFF  DEM 1B OK
 5  1  5  TRUE  OFF  OFF  DEM 2A OK
 6  1  6  TRUE  OFF  OFF  DEM 2B OK
 7  1  7  TRUE  OFF  OFF  DEM 3A OK
 8  1  8  TRUE  OFF  OFF  DEM 3B OK
 9  1  9  TRUE  OFF  OFF  DEM 4A OK
10  1 10  TRUE  OFF  OFF  DEM 4B OK
 1  2  1  TRUE  OFF  OFF  IOM 1 OK
 2  2  2  TRUE  OFF  OFF  IOM 2 OK
 3  2  3  TRUE  OFF  OFF  DEM 1A OK
 4  2  4  TRUE  OFF  OFF  DEM 1B OK
 5  2  5  TRUE  OFF  OFF  DEM 2A OK
[...]
```

This happens when one of the two SEPs in the affected controller stops generating “heartbeat” signals, making it invisible to the other SEP over the I2C bus, which is used to arbitrate requests for access to shared enclosure resources. When this occurs, Enclosure Services (ES) logs the apparent removal of the SEP from the system. For example:

```
002905 2013-06-30 17:31:53:3528713 G=62 S=1 T=1 RP=0 VP=00
LOG_ES_EXPANDER_ELEMENT_REMOVED ES Enclosure 0x50001ff127980000 expander element 2
was removed. SES status = 5, DISAB = 0, PF = 0.
```

When SFA OS subsequently attempts to communicate with the nonresponsive SEP and fails, that failure is also logged by Enclosure Services. For example:

```
476790 2013-06-30 17:32:53:7636792 G=60 S=1 T=1 RP=0 VP=00
LOG_ES_SES_COMMAND_QUERY_FAILURE ES Command sent in query state 9 to enclosure
50001ff127980000 failed. PBX STS:JS_AMPD_HW_ABORTED_CMD XPSTS:15 SCSISTS:0 XC:0 SK:0
ASCQ:0:0.
```

The controller with the nonresponsive SEP may then begin aborting commands to drives in the pools that it owns. For example:

```
476791 2013-06-30 17:35:51:7633063 G=60 S=0 T=1 RP=0 VP=08
LOG_IOF_AUTO_QUIESCED IOF AUTO QUI PD:00083 STS:JS_AMPD_HW_ABORTED_CMD
476792 2013-06-30 17:35:51:7633318 G=60 S=1 T=1 RP=0 VP=08
LOG_IOF_STATE_ERR IOF ST ERR OP:0001 SRC:0000 STS:JS_IOF_STARTED
476793 2013-06-30 17:45:52:0132246 G=60 S=0 T=1 RP=0 VP=08
LOG_IOF_AUTO_QUIESCED IOF AUTO QUI PD:00004 STS:JS_AMPD_HW_ABORTED_CMD
476794 2013-06-30 17:45:52:0132499 G=60 S=1 T=1 RP=0 VP=08
LOG_IOF_STATE_ERR IOF ST ERR OP:0001 SRC:0000 STS:JS_IOF_STARTED
476795 2013-06-30 17:53:27:5126933 G=60 S=0 T=1 RP=0 VP=08
LOG_IOF_AUTO_QUIESCED IOF AUTO QUI PD:00109 STS:JS_AMPD_HW_ABORTED_CMD
476796 2013-06-30 17:53:27:5127205 G=60 S=1 T=1 RP=0 VP=08
LOG_IOF_STATE_ERR IOF ST ERR OP:0001 SRC:0000 STS:JS_IOF_STARTED
476797 2013-06-30 18:54:43:2599152 G=60 S=0 T=1 RP=0 VP=08
LOG_IOF_AUTO_QUIESCED IOF AUTO QUI PD:00069 STS:JS_AMPD_HW_ABORTED_CMD
476798 2013-06-30 18:54:43:2599423 G=60 S=1 T=1 RP=0 VP=08
LOG_IOF_STATE_ERR IOF ST ERR OP:0001 SRC:0000 STS:JS_IOF_STARTED
476799 2013-06-30 19:21:22:2584611 G=60 S=0 T=1 RP=0 VP=08
LOG_IOF_AUTO_QUIESCED IOF AUTO QUI PD:00364 STS:JS_AMPD_HW_ABORTED_CMD
476800 2013-06-30 19:21:22:2584891 G=60 S=1 T=1 RP=0 VP=08
LOG_IOF_STATE_ERR IOF ST ERR OP:0001 SRC:0000 STS:JS_IOF_STARTED
476801 2013-06-30 19:23:50:0087120 G=60 S=0 T=1 RP=0 VP=08
LOG_IOF_AUTO_QUIESCED IOF AUTO QUI PD:00107 STS:JS_AMPD_HW_ABORTED_CMD
```

Over time, the aborted commands can cause drives to fail. The affected pools become degraded, then nonredundant.

If a subsequent write request to a nonredundant pool happens to fail, the pool will be placed in Auto Write Lock (AWL) mode. This state prevents the host from executing additional writes to the pool. At this point, data in the pool can become inaccessible to file systems or applications.

When this combination of symptoms is encountered, the S2A6620 can be restored to normal operation using the procedure below.

## 3.0 Recovery Procedure

### 3.1 Recovery Summary

At a high level, you will perform the following recovery steps:

- Locate the controller where the “NOT INSTALLED” or inactive SEP resides.
- ***If you have physical access to the system:***
  - Reseat the affected controller.

- Perform a complete power cycle of the system if reseating fails.
- **If you do not have physical access to the system:**
  - Run DDN-provided, power-cycling firmware remotely on the controller.

### 3.2 Locating the Inactive SEP

The steps in this section locate the inactive SEP and prepare the S2A6620 system for recovery intervention. These steps must be followed for both onsite and remote recovery.

**Step 1.** Collect a `show sub sum` log to document the initial state of the system.

**Step 2.** Determine where the inactive SEP is physically located. In an S2A6620 system, this will be one of the two controllers in the **base** enclosure (or “head unit”); the **expansion** enclosure (if any) is not involved. Find the controller containing the inactive SEP with the command

```
show expander
```

In the command output, look for an expander with a **SES Status** of “NOT INSTALLED”. For example:

Idx	Encl	Pos	Present	Failure	Locate	Location	SES Status
1	1	1	TRUE	OFF	OFF		NOT INSTALLED
2	1	2	TRUE	OFF	OFF	IOM 2	OK
3	1	3	TRUE	OFF	OFF	DEM 1A	OK
.	.	.	....	...	...	.....	...
1	2	1	TRUE	OFF	OFF	IOM 1	OK
2	2	2	TRUE	OFF	OFF	IOM 2	OK
3	2	3	TRUE	OFF	OFF	DEM 1A	OK

The missing SEP shows a blank line in the **Location** column, rather than an I/O module (IOM) location such as “IOM 1” (the **left** IOM bay in the unit, when viewed from the back) or “IOM 2” (the **right** IOM bay in the unit, when viewed from the back). Since there are exactly two IOM locations where the affected controller can reside, you can infer the location of the controller with the “missing” SEP from the location given for the *other* SEP in the same base enclosure.

In the example above, the “NOT INSTALLED” SEP is **not** located in the “IOM 2” bay, and must therefore be located in “IOM 1”, meaning it is part of the controller on the **left** side of the base unit when viewed from the back.

**Step 3.** Get the IP address of the controller with the inactive SEP from your system administrator if you don’t already know it. By convention on S2A6620 systems:

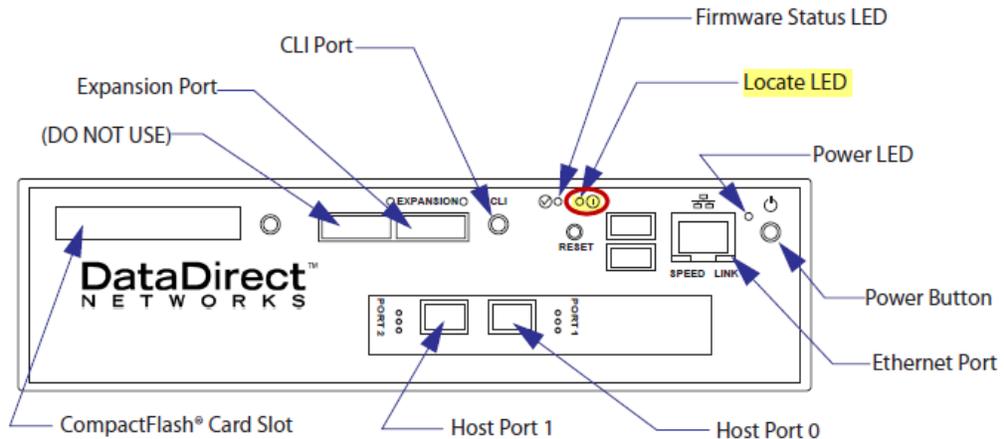
- 10.0.0.1 is the default IP address of the left controller
- 10.0.1.1 is the default IP address of the right controller

**Step 4.** Log on as user to the affected controller.

**Step 5.** If you have physical access to the system, verify that the controller you logged on to is the physically correct controller with the command

```
locate controller local
```

From the back of the S2A6620 system, verify that the **Locate** LED is lit on the physical controller in which the inactive SEP resides. (See diagram.)



If the located controller is not the one with the inactive SEP, as identified in Step 1, log out, then log in to the correct physical controller.

### 3.3 Prepare for Recovery Actions

**Step 6.** Check to see whether any virtual drives are in “INOP” status with the command

```
show vd
```

If **any** virtual drive shows as “INOP” in the **State** field of the output, it is inoperative due to pool degradation, and cache cannot be flushed. Contact DDN Technical Support to **resolve this condition before you do anything else**.

---

**CAUTION** Do **not** perform this procedure with a virtual disk in INOP status.

---

**Step 7.** Quiesce all host I/O to the storage system.

**Step 8.** Ensure that all pools are in write-through mode. This flushes cache to disk.

- **For SFA OS 1.5.0 and higher**, enter the command

```
set pool * write_back_caching=false
```
- **For SFA OS 1.4.x**, turn off write-back caching for each pool individually. Enter the commands

```
set pool 0 write_back_caching=false
set pool 1 write_back_caching=false
```

```
set pool 2 write_back_caching=false
.
.
.
```

### 3.4 Remote Recovery Procedure

**Step 9.** *If you have physical access* to the S2A6620 equipment, **skip to Step 11**.

**Step 10.** If you do *not* have physical access to the equipment, you can use a script supplied by DDN Technical Support to perform a **firmware-controlled** power cycle procedure on the controller with the inactive SEP, as follows.

a. Obtain the S2A6620 power restart firmware routine from DDN Technical Support. The file name is `ddn-restart-6620-power.v6.tar.gz`.

b. Copy the power restart file to the controller with the inactive SEP using the secure copy protocol. Enter the command

```
scp ddn-restart-6620-power.v6.tar.gz firmware@<ip_address>:
```

where `<ip_address>` is the IP address of the affected controller, and the final colon is required.

For example:

```
scp ddn-restart-6620-power.v6.tar.gz firmware@10.0.0.1:
```

c. When prompted for a password, enter `Firmware` (capitalized).

d. Run the script with the command

```
update controller local file="ddn-restart-6620-  
power.v6.tar.gz"
```

---

**NOTE** The script will cause the controller to reboot. It can take **15 minutes** for the controller to come back up.

---

e. Wait for the LEDs on the back of the controller to indicate normal operation.

f. Check to see whether the “NOT INSTALLED” state has been cleared with the command

```
show expander
```

g. If the output of the `show expander` command shows the **SES Status** of all expanders as “OK”, the controller restart was successful. **Skip to Step 14**.

h. If the output of the `show expander` command still shows the SEP with a **SES Status** of “NOT INSTALLED”, the controller restart was not successful. Perform the onsite recovery procedure beginning at **Step 11**.

### 3.5 Onsite Recovery Procedure

**Step 11.** Physically reseal the controller with the inactive SEP as follows.

- a. Make sure you are still logged in to correct physical controller (see Step 1 to Step 5).
- b. Shut down the controller with the inactive SEP using the command  

```
shutdown controller local
```
- c. Wait until all LEDs on the back of the controller with the inactive SEP have gone dark (except for the Ethernet port LEDs).
- d. Ensure that you have sufficient slack in the attached cables to pull out the controller about 3 to 5 inches. Then pull both levers on the back of the controller outward gently to unseat it, and continue pulling until the controller projects about 3 to 5 inches outside the base enclosure.
- e. *Wait 2 minutes.*
- f. Insert the controller gently back into the system as far as possible, then push both levers on the back of the controller flush with the controller backplate to reseal it securely.
- g. Allow the reseated controller to power up automatically. The LEDs on the back of the controller should light and indicate normal operation.
- h. Check to see whether the “NOT INSTALLED” state has been cleared with the command  

```
show expander
```

If the output of this command shows the **SES Status** of *all* expanders as “OK”, reseating the controller was successful. *Skip to Step 14.*

**Step 12.** *If reseating the controller did not restore the system* to normal operation, power cycle the entire S2A6620 system as follows.

- a. Initiate an orderly shutdown with the command  

```
shutdown subsystem
```

---

**CAUTION** Do *not* use the power button on the back of the controller to shut down the S2A6620. *Doing so may cause data loss.*

---

- b. When all LEDs are dark on the back of **both** S2A6620 controllers in the base enclosure, power down *both* power and cooling modules (PCMs) for the unit by flipping their rear toggle switches to the **Off** position. Wait for the disks

in the base enclosure to spin down and all LEDs on the back of the enclosure to go dark. (Note that the Ethernet port LEDs may stay lit.)

- c. Power down any expansion enclosure that may be installed by flipping the toggle switches on the back of **both** expansion PCMs to the **Off** position. Wait for the disks to spin down and the LEDs on the back of the enclosure to go dark. (Note that the Ethernet port LEDs may stay lit.)
- d. **Wait 5 minutes.**
- e. Power up the expansion enclosure, if any, by flipping the toggle switches on the back of **both** expansion PCMs to the **On** position. Wait for the disks to spin up in the expansion enclosure and the LEDs on the rear of the unit to light and show normal operation.
- f. Power up the base enclosure by flipping the toggle switches on the back of **both** base enclosure PCMs to the **On** position.
- g. Verify that the **Power** LED on the back of **both** controllers turns green within a few seconds, indicating that power has been applied to the controllers.
- h. Verify that the **Firmware Status** LED blinks green on the back of **both** controllers within 1 minute of power application. This indicates that the SFA firmware is loading.
- i. Verify that the **Firmware Status** LED on the back of **both** controllers stays solid green (not blinking) within two minutes of power application. This indicates that the firmware load operation is complete and commands may be issued to the S2A6620.
- j. Verify that the system came up in a healthy state, as described in the *S2A6620 User Guide*.
- k. Check to see whether the “NOT INSTALLED” state has been cleared. Enter the command  

```
show expander
```

If the output of this command shows the **SES Status** of *all* expanders as “OK”, power cycling the system was successful. **Skip to Step 14.**

**Step 13.** *If power cycling the entire system was not successful*, replace the controller with the “missing” SEP as follows.

- a. Log on as user to **the controller with the inactive SEP**.
- b. Take the affected controller offline to prevent issues with the pools it owns.

At the command line, enter  
`shutdown controller local`

- c. Contact DDN Technical Support to request an RMA replacement of the degraded controller you just shut down.
- d. ***Proceed to Step 14.***

### 3.6 Verify Data and Resume Operation

Ensure all data is written to nonvolatile storage and verify the data integrity of all storage pools potentially affected by the nonfunctional SEP. You can then resume host I/O operations, even if one controller is shut down. The verify job will proceed in the background.

**Step 14.** If any pools are in AWL status, clear that condition. This will flush any cache that could not be written to the pools in AWL mode and will force a verify operation on those pools.

- ***For SFA OS 1.5.0 and higher***, enter the command  
`clear pool * auto_write_lock`
- ***For SFA OS 1.4.x***, clear AWL status for each pool individually. Enter the commands  
`clear pool 0 auto_write_lock`  
`clear pool 1 auto_write_lock`  
`clear pool 2 auto_write_lock`  
.  
.  
.

**Step 15.** ***If both controllers are operational***, restart write-back caching on the pools normally configured for it.

- ***For all pools on SFA OS 1.5.0 and higher***, enter the command  
`set pool * write_back_caching=true`
- ***For selected pools on SFA OS 1.5.0 and higher, or for SFA OS 1.4.x***, turn on write-back caching for each pool individually. Enter the commands  
`set pool 0 write_back_caching=true`  
`set pool 1 write_back_caching=true`  
`set pool 2 write_back_caching=true`  
.  
.  
.

### 3.7 Post-Recovery Actions

Document the state of the system after all recovery actions are complete.

**Step 16.** Even if recovery was successful without a controller replacement, you should use the mark command to note the SEP error in the controller log. For example:

```
mark controller=<id> "INACTIVE SEP, RESEATED TO RECOVER"
```

where <id> is the object ID of the affected controller.

If the same controller encounters this issue again, replace the controller.

**Step 17.** Collect a show sub sum log to document the final state of the system.

## 4.0 Contacting DDN Technical Support

Contact DDN Technical Support if you need assistance at any time. Support can be reached by telephone, email, or on the web.

### Telephone

DDN Support Worldwide Directory <http://www.ddn.com/support/contact-support>

### Email

Support Email [support@ddn.com](mailto:support@ddn.com)

### Web

Support Portal <https://portal.ddnsupport.com>  
Portal Assistance [webportal.support@ddn.com](mailto:webportal.support@ddn.com)

### Bulletins

Support Bulletins <http://www.ddn.com/support/technical-service-bulletins>  
Bulletin Subscription Requests [support-tsb@ddn.com](mailto:support-tsb@ddn.com)