



AIC/XTREME SAS SERVER/JBOD Expander Firmware release note

Last updated: Nov-10th-2008

This guide provides information on AIC/XTREME SAS products firmware Release note for the following Models. Review this entire document before updating the firmware

Models	Brand	Firmware Revision	Remark
XJ-SA13/26-110R -x	AIC/ Xtore	xxxxxxx	1U/10 SAS JBOD Series
XJ-SA13/26-224R-x	AIC/ Xtore	xxxxxxx	2U/24 SAS JBOD Series
XJ-SA13/26-212R-x	AIC/ Xtore	xxxxxxx	2U/12 SAS JBOD Series
XJ-SA12/24-316R-x	AIC/ Xtore	xxxxxxx	3U/16 SAS JBOD Series
RSC-1A-SA1S/D-2	AIC/ Xtore	xxxxxxx	1U/10 SAS Server Series
RSC-2A-SA1S/D-2	AIC/ Xtore	xxxxxxx	2U/24 SAS Server Series
RSC-2EC-SA1S/D-2	AIC/ Xtore	xxxxxxx	2U/12 SAS Server Series
RSC-3EC-SA1S/D-2	AIC/ Xtore	xxxxxxx	3U/16 SAS Server Series
RSC-3EG-SAS1S/D-2	AIC/ Xtore	xxxxxxx	3U/16 SAS Server Series

Firmware Editions history:

- [1.14H.08 Release note \(Nov-10th -2008\)](#)
- [1.14H.03 Release note \(July-16th -2008\)](#)
- [1.14H.02 Release note \(May-27th -2008\)](#)
- [1.14H.01 Release note \(May-23rd -2008\)](#)
- [1.14F.01 Release note \(Jun-27th-2007\)](#)
- [1.09G Release note \(Oct-17th-2006\)](#)
- [1.09D Release note \(Sep-15th-2006\)](#)
- [1.07C Release note \(July-17th-2006\)](#)
- [1.06d Release Note \(April-11th-2006\)](#)

Release note for 1.14H.08 Date: Nov-10th -2008

Bug Description

1. LED control by SES Identify command (all models)
There is a bug since 1.14H.02 that causes the Blue (Identity/presence/IO-activity) LED staying Off after an SES 'Identify' command ended.
Most of our models have the same problem with most SES capable HBAs.
2. A few incorrect LED behaviors with some HBA of variant vendors.
3. For the new integrated 3U16 expander board only:



- Disk error LED of the swapped disk stays On after a RAID member drive hot-swapped.
- The Blue LED doesn't blink (a little while) during initiation of SATA disk (such as when a SATA disk inserted).

Misc. Notes:

1. Power Supply status of the new integrated 3U16 expander board may not correctly reflect the physical state.

There are two jumpers reserved for wiring the power-good status signal between the PSU and sensor; in factory default, they are not wired and the jumpers are left open to force positive power-good state.

The expander FW has no ways to tell if physical power-good signals are present; it will assume the signals are wired in and report the status as read.

2. Error LED of a hot-swapped disk incorrectly blinks (with most Adaptec RAID board)

On a system configured with hot-spared disk, If you remove a raid member disk, wait a while until RAID rebuilding kicked in, and then insert this RAID member disk back, the error LED of the swapped disk keeps blinking even after the RAID was rebuilt successfully on the hot-spared disk.

This is a bug of Adaptec FW; it forgets to reset the 'rqst_inCriticalArray' SES flag for the swapped disk if a hot-spared disk is available for the rebuilding of the degraded array. This bug was never spotted because our FW had other bugs preventing the error LED from truthfully reflects the SES state in this specific situation.

We'll contact Adaptec R&D and have this issue fixed.

Work-around:

Create a small single disk volume on the swapped disk to have the Adaptec FW turning off the error LED and then delete this volume to set free the disk for normal use.

3. Please pay more attention for testing the LED behavior with HBA from different vendors.
4. For the new integrated 3U16 expander board, you may like to utilize both 4i ports for hot connection; just issue the following console commands:



PHYAttribute set 14 1

PHYAttribute set 15 1

PHYAttribute set 16 1

PHYAttribute set 17 1

These command set the phys to subtractive routing mode immediately for connecting to host; these settings are also saved in configuration NVRAM.

You may check the new routing mode by looking at the 'Sub' column of output by the console command 'phystat' (they should have value of 1, just the same as phy 0 ~ 3).

[Release note for 1.14H.03](#) Date: July-16th -2008

Bug Description

Fixed fault LED issue with Areca SAS HBA

For now, we just ignore the SES 'rqstMissing' control bit sent from Areca HBA so the faulty LED of the drive bay won't be turned on when a drive not configured as RAID member is unplugged.

[Release note for 1.14H.02](#) Date: May-27th -2008

A serious bug in 1.14H.01 FW for all of SAS enclosures was found and fixed in 1.14H.02.

Bug Description

In a dual-modules setup, the expander module plugged in slot-B mistakenly insists on taking the role as a Master enclosure controller and refuses granting the request from expander module at slot-A for the Master controller role.

Due to design limit ion in our code, the expander module at slot-A won't be able to sync. its SAS address with the master controller module at slot-B and will have an null SAS address.

Note: The above symptom manifests only when both expander modules are present when powering up the enclosure; It won't notice the problem if only one module installed or even another module hot plugged afterward.



Release note for 1.14H.01 Date: May-23rd -2008

1.14H.01 Release Highlight:

1. Enhancement to the staggering spin-up implementation
 - Aggressively reducing disk staggering spin-up delay
 1. All models for 2.5" drive default to 3 seconds of spin-up delay for each disk group
 2. All models for 3.5" drive default to 5 seconds of spin-up delay for each disk group
 - Reorganizing and reducing spin-up disk group
 1. For all models, the number of disks grouped to be spontaneously spun up is increased (and thus spin-up groups reduced).
 - In general, the above enhancements shall drastically speed up the power on process, with some side effects:
 1. The code boot issue with LSI 8888ELP, 8708ELP, and a Marvell RAID board, caused by their much shorter time span in enumerating attached devices, is considered fixed, except those models (2U24 server, 2U12 JBOD and server) with H/W patches mandated by OEM engineer (the patches adds up 10 seconds delay to the 7153/7154 reset sequence).
 2. The increased spontaneously spun up disks will surely put more stress on the power supply units. The implication of this will be carefully analyzed and verified by our R/D.
2. Adding a console command 'tSpinupDelay' fine tuning the spin-up delay

```
tSpinupDelay [delay_in_seconds (0 - 32)]
```

the *delay_in_seconds* parameter:

can be omitted to get the current value.

can be given a value of '0' to reset to model specific system default.

3. Supporting the retrieving of full FW version stamp
We hijacked the SCSI 'String In' Diagnostic page (Page Code 04H) to report the full version stamp of the SEP FW:
 - The returned string will always be in the form of 'NN.nnX.nn', such as '1.14h.04' or '1.14H.01'.



4. Fixing several bugs in controlling drive bay LED when requested to change state of a SES ArrayDevice element
 - Correcting the LED behavior when requested to change from some states via the [RqstOk, RqstRebuilding] request.
 - Correcting the LED behavior when requested to turn off the access LED (the blue one) with the 'RqstOk' flag reset (set to 0).
 - These bugs only manifest themselves when special SES requests sequences are called, and:
 - You may spot a few LED behavior changes when running mainstream SAS management suits but all seems to make the indicators more reasonable and understandable.
 - Your home brewed SES client may or may not be affected (in positive or negative way) by the changes; carefully testing your application is advised.

 5. Fixing a SMP bug that reports SATA device presence prematurely

 6. Fixing a bug in SES 'Download Microcode' control page
Correcting the returned status code on each received data chunk and the final status code upon completion or aborting.
 - This fix should make it a lot easier for writing your own FW download client.

 7. Fixing a minor bug in reporting the value of the 5V DC sensor upon SES requests
 - The returned value of DC 5V sensor was not normalized from the raw reading of the LM80 chip.
 - Because the calculation of the normalization is quite a heavy load to the 7153/7154 core, the bug is fixed with a workaround:
The description of the DC 5V sensor is changed to '1.9V'; the returned value of the nominal readings is usually in the range of 190±8.
-

Release note for 1.14F.01 Date: Jun-27th-2007

1.14F.01 Release Highlight:



1. Fixed a critical bug in replying SCSI Inquiry command for the 'Device Identification' VPD page of the enclosure SES device.
 - o In 1.14F.00 and 1.14f.xx, the SAS address of the enclosure SES device is not returned if this address is not exposed to SMP client (controlled by 'tExposeELI on/off' cmd, the default is off). This severely violates SAS specification.
 - o This bug may confuse host in enumerating disk drives in a multiple expanders environment (either cascading or not)
 - o With this fix, Adaptec 3085 RAID board, when connecting with both expander modules in one enclosure, can now list correct drive population and sustain RAID operations upon failure of a single expander module. This may also applies to other Adaptec RAID boards with multiple 4i or 4x connectors..

1.14F Release Highlight:

1. Merge in all low level libraries from Vitesse SDK 1.14.
You can use the 'stat' console command to verify the version of those low level libraries.
2. Formally added customized version for OEM models:
 1. The image files are packed separately in '1.15F-schg.tar.gz'.
Major customization for these OEM models are:
 1. The host port signal strength is boosted to support long cable (max. 6 meters)
 2. For the 2U24 server model:
 1. 'Split-Expander' feature (see Highlight 5) is activated by default.
 2. Thermal sensitive fan control is disable and SES fan element support is removed.
3. Added customized general OEM version the 2U24 JBOD/server and 2U12 JOB D models:
 1. The image files are packed separately in '1.15F-oem.tar.gz'.
 2. Major customization for these OEM models are:
 1. The host port signal strength is boosted to support long cable (max. 6 meters).
4. **Discontinuing support of older version of backplane** because new (and fancy) drive bay LED behavior for SES Array Device Element can't be implemented on those backplanes.



5. SES-2 disk device support is reverted from 'Device Element' to 'Array Device Element' (except on OEM model).
1. Most 'Array Device Element' specific control and status functions are supported.
 2. Hardware based LED pulse train for most supported status indication.

The following table depicts the LED flash pattern for each supported control/status field in SES-2 Enclosure Control/Status page (page code 0x02) for Array Device Element:

Supported Control	Supported Status	LED Pulse for Status	Priority	Note
Rqst Ok	Ok	Access LED constantly ON when: 1. this bit is on, or 2. this bit being off but hot-spare bit is on	2	'Ident' bit has 1st priority in governing Access LED.
Rqst Rsvd Device (Reserved Device)	Rsvd Device	Error LED flash. Pulse train segment: 000111000000 (in about 5 seconds)	7	
Rqst Hot Spare	Hot Spare	Same as 'Rsvd Device'	6	
Rqst Cons Check (Consistency Check)	Cons Check	Error LED flash. Pulse train segment: 010101010101 (in about 0.502 second, e.g., fast flashing)	5	Always accompanying with Access LED flashing
Rqst in Crit Array (In Critical Array)	In Crit Array	Error LED flash. Pulse train segment: 0110000000 (in about 0.418 second)	4	
Rqst in Failed Array	In Failed Array	Same as 'In Critical Array'	3	
Rqst Rebuild/Remap	Rebuild/Remap	Same as 'Cons Check'	2	Always accompanying with Access LED flashing



Rqst R/R Abort (Rebuild/Remap Abort)	R/R Aborted	Error LED flash. Pulse train segment: 10000011 (in 1 about 1 Hz)	1	
Rqst Missing	N/A	None	N/A	Event recorded without giving visual indication.
Rqst Insert	Ready to Insert	Error LED flash in 1Hz	8	
Rqst Remove	Ready to Remove	Error LED flash in 1Hz	8	
N/A	Fault Sensed	Error LED constantly ON	8	
Rqst Fault	Fault REQSTD	Error LED constantly ON	9	Usually accompanying with 'Ok' bit off and thus Access LED OFF
Rqst Ident	Ident	Access LED flash in 1Hz	1	

6. Synchronizing SES Enclosure Logical ID (ELI) of primary and secondary expander modules in a single enclosure box.
SEP of both expander modules in a dual-expander enclosure are given same SES Enclosure Logical ID that's world-wide unique. In our implementation, this ID is the SAS address of the virtual SSP PHY of the master expander module.
7. Exposing SES Enclosure Logical ID (ELI) to SMP client.
To work around the incompatibility of this feature with LSI non-RAID HBA, a new console command 'tExposeELI', is added to control this feature:
 1. **tExposeELI on**: enable it.
 2. **tExposeELI off**: disable it.
 3. **tExposeELI** (no argument) : show the activation state.
8. The 'tEnclosureState' console command is enhanced to report current SES Enclosure Logical ID.
9. Added support of the header 'EIP' bit and each element's 'bay number' field in response to SES request page 0Ah (Additional Element Status)
 1. This is necessary so that SEP client can correctly report the bay position of each drive. This has been verified with SAS management GUI suite from Adaptec and LSI.
 2. This change **may break your customized application** if your SES library doesn't take into account of the header 'EIP' bit and react accordingly.



3. The slot number field counts from 1 except on OEM models (counts from 0 per requested).

10. Added the 'Split Expander' feature.

This feature enables a single expander module to provision 2 virtual expanders to a single HBA connected with 2 host ports.

Some notes about the current implementation:

1. This feature is in fact a hack and violating SMP and SSP standard. It only works ad hoc with LSI HBA and RAID boards and it is meant to work around the issue of LSI boards not capable of aggregating SAS links across host ports.
2. Split-expander mode is disabled by default (except on specific OEM models). A new, hidden console command 't2Expanders' is added to control the activation of it:
 1. **t2Expanders SplitIt**: Activated split-expander mode on next power circle.
 2. **t2Expanders off**: Restored to normal mode on next power circle.
 3. **t2Expanders** (without parameter): Display current state.
3. The '**tZoneCfg**' console command is enhanced to cop with the special zoning requirement by split-expander mode:
 1. '**tZoneCfg -d**': display active zoning state and manual zone configuration stored in NVRAM.
This help user to have a glance on the implicit zone configuration by split-expander feature vs. explicit zone configuration manually set up .
 2. '**tZoneCfg -s**' (for manually configuring zone) is disabled when split-expander mode is active.
4. Activation status of 'split-expander' is also reported by 'tExpanderState' and 'tShowState' commands.
5. This feature is only supported by expander module with 2 or more host (up-link) ports.
6. In current implementation:
Phys for the host port #1 plus phys for the first half of drive bays; e.g., bay #1 through #12 in the case of 24-bay box, are configured into zone #1.
Phys for the host port #2 and phys for the last half of drive bays, e.g., bay #13 through #24 in the case of 24-bay box, are configured into zone #2.



This means:

1. A HBA connecting to the expander box with single 4X port will see only half of total drives.
 2. A HBA connecting to the expander box with dual 4X ports will see 2 expanders, each containing half of total drives.
7. Currently, the SES subsystem doesn't support split-expander properly due to low level library restriction. Some of the known issues:
1. The SEP sees drives in both zones and exposes all of them to the host. A HBA with dual 4X connection to the box will see 2 SES enclosures (as it should) but both enclosure are populated with duplicated drives (tested with LSI 8480E RAID controller).
Note that this is mostly a cosmetic issue since the HBA actually knows the exact total physical drives in the expander box and you are provisioned with the correct number of drives for RAID configuration, at least in LSI management suite.
 2. Hot-plugging host to expander box doesn't work reliably.
For dual-port connection, you must connect both cables before powering up the host.
11. Add more version information to the boot-up flag message and the 'tEnclosureState' command.

For beta release:

1. Stamp of normal version will be in the form of '1.1Nx.nn', such 1.14f.05;
2. Stamp of debug version will be in the form of '1.14x.nn-debug', such as 1.14f.05-debug

For production release:

3. Stamp of normal version will be in the form of '1.1Nx[.nn]', such 1.14F and 1.14F.01;
 4. Stamp of debug version will be in the form of '1.14x.nn-debug', such as 1.14F-debug
- Note: The extra sub-version stamp doesn't apply to SMP and SES response because of the 4 bytes limitation in specification.



Release note for 1.09G Date: Oct-17th-2006

Bugs fixed:

1. With Adaptec SAS 4800/4805 RAID controllers, disk hot plug doesn't work Symptoms:
 1. Disk inserted in any slot that is empty when expander powered up won't be recognized by the mentioned HBA.
 2. However, after removing and then inserting disk in the same slot again, the disk will be correctly recognized.

Fix:

1. The fix is tentative though reliable and effective from my surface test; It's just a replacement of the offensive library (libSMP) with the one copy from Vitesse original 1.10 FW.
2. Using 'stat' console command, you can get the version stamp (01.10) of 'libSMP'

Release note for 1.09D Date: Sep-15th-2006

1.09D FW for VSC-7153x/7154x based SAS JBOD were put on-line.

1. Due to some critical issues related to fan control/monitoring, **there will be no 1.09D for server models except fan less models for OEM.**
2. Except sub-version stamp changed from 'd' to 'D', there are following differences from the beta 1.09d:
 - The syntax of 'tZoneCfg' console command was enhanced to support clearing zone configuration of a single phy. The help message for this command was also corrected.
3. For your convenience, the release note for 1.09d beta is included.

Major Improvement in 1.09D:

1. Adding support for the revised 1U-10 JBOD:
 - Fully support the redesigned fan subsystem:
 1. Total 4 fans instead of 3.
 2. Each IOM (I/O module) controls the 2 fans directly connected to it, but every module can monitor all 4 fans.
 3. 2 additional 'fan error LED' are supported to alert fan error event; one LED for each fan set.
 4. In a single IOM enclosure, a dummy I/O module mounted with 2 fans must be installed in the empty I/O module slot and both fans must work normally otherwise fan fault event will be fired (enclosure system error LED, fan error LED, enclosure beeping and SES status report).



2. Adding support for dual boot with a single EEPROM chip on revised IOM:

1. On the 4 revised IOM for 1U10 and 2U12 JBOD and server, a new 2-pin jumper header is added for selecting one of two booting FW image in a single EEPROM chip (must be in 8 Mbit size):

Jumper	FW Image Address Space Primary / Secondary	Boot Address
Open (default)	BFC0:0000 / BFC8:0000	Primary at BFC00000
Close	BFC8:0000 / BFC0:0000	Secondary at BFC80000

2. With this dual boot capability, we'll have much less danger of rendering an IOM totally non-bootable due to burning error.
When burning failed, either the Primary image is intact (not yet touched) or the Secondary image was already updated successfully; we can just blindly try booting from either image by opening or closing this jumper.
3. With the 'fwupdate -seconly' console command, we can burn different version of FW onto the Secondary image location and boot it up to test new FW before official deployment.

3. A lot of improvement in fan control and monitor:

1. A bug in setting fan speed on most server models was found and fixed:
The bug caused under-run of fans and may be the cause of overheat problem of most of our server models. The fix has been verified OK but currently all fans are still forced to run in max speed (In previous versions, all fans run only at about 50% to 70% of its max capacity even we set it to run at max speed because the base clock was set too slow).
2. We now have more specific PWM parameters based on combination of fan spec, IOM and backplane revision. With later version of IOM and backplane, we should now be able to get correct reading of fan RPM. With early buggy IOM/backplane, we are more confident in determining its running or stop state.
3. Fan status monitoring and reporting is added back to 3U server with A2 version backplane.

4. Some modification to AIC add-on console command set:

1. The 'tShowState' functionality has now been distributed to two new commands ('tShowState' is still kept):
 1. 'tEnclosureState': display thermal related states.
 2. 'tExpanderState': display expander states (phy related)
2. 'tUptime' cmd is added to tell how long the IOM has been up and running.
* Output of 'tEnclosureState', 'tExpanderState' and 'tShowState' also has one



line reporting uptime information.

3. 'tZoneConfig' cmd is now really functional and its help message revised to better explain usage of each sub-command.

- **Note: Because of the large scale of modification, we must test each model intensively and carefully; bug is to be expected.**

Release note for 1.07C Date: July-17th-2006

Major Improvement:

1. Fixed the 'PHY error 16 (24 in case of 7154 based expander) ... ' error message while login Windows with LSI 8408/8480 RAID.
2. Fix the problem of BIOS utility of LSI 3442 and 3800 HBA not able to display disks attached to our expander or even failed to display SAS topology.
3. Added support for the SCSI 'Report LUNS' command.
 1. The added support only replies one Logical Unit (0:0 of the SEP).
 2. This should fix the hidden problem of LSI 3442/3800 taking long time to inquiry up to 255 non-existent logical units and thus should speed up OS boot up time.

FW Filename Naming Convention:

1. For 7153 based model, the FW filename has this pattern:
e*T53C-VER*-bp*n**NN*.s3r
 - o '*T*' : could be '*d*' for dual expander daughter card model or '*m*' for single expander module
 - o '*C*' : denotes the 7153 chip version; '*b*' for 7153b, '*c*' for 7153-02.
 - o '*VER*' : denote FW version; such as '**1.07B**'.
 - o '*n*' : denotes the backplane revision; such as '0' or '1' or '2'.
 - o '*NN*' : denotes the supported drive slots; such **10**, **12** or **24**.

Example:

- o ed7153b-1.07B-bpA2_16d.s3r: for daughter card supporting 16 drive based on 7153B chip.
 - o em7153c-1.07B-bpA1_12d.s3r: for module supporting 12 drive based on 7153-02 chip.
2. For 7154 based model:
Similar but simplified since all of them support 24 drive.

5. Release note for 1.06d Date: April-11th-2006



- Added support for new models:
 - ecm7153c_a1_12d-1.06d.s3r, ecdc7153c_a1_12d-1.06d.s3r: for 2U 12bay JBOD/Server
 - ecm7153c_a1_10d-1.06d.s3r, ecdc7153c_a1_10d-1.06d.s3r: for 2U 12bay JBOD/Server
- Continuing improvement to SES support:
 - Fixed some bugs to better comply to SES and SSP transport spec..
 - Partially support of controlling enclosure indicator/alarm via SES:
 - Fully functional on all 7153 based models
 - No action for all 7154 based models due to hardware limitation. This will be addressed in later version
 - Fully support of controlling access LED (blue) on disk tray for all models
 - * Support for error LED (red) control will be addressed in later version.

Some corrections to last release announcement:

- The release version in the subject line was incorrect; it should be 1.06d, not 1.06a.
- For those models with two version of backplane (most of them, except the new 2U12 and 1U10), the support of SES control on the disk tray access LED **only works with version 2 of backplane**; the old backplane may never be supported since it's too difficult to make it right and reliable. Except this, I also recommend all upgrading their old backplane for performance issue and to enjoy planned enhancement to support different LED flash patterns in different conditions.

As said, this is a correction to the release information. The firmware remains intact; you **don't** have to upgrade again if you had already done.