

Sun Flash Accelerator F20 PCIe Card

Product Notes



Part No. 820-7267-14
July 2010, Revision A

Copyright © 2009, 2010, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related software documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2009, 2010, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.



Adobe PostScript

Contents

Preface v

Sun Flash Accelerator F20 PCIe Card Product Notes 1

System Requirements 1

Supported Servers 1

Supported Operating Systems 2

Minimum Supported Firmware Versions 3

Minimum Required System Firmware Patches 3

Solaris Performance Patch 3

FMod Alignment 4

Required ESM Replacement 4

Monitoring ESM Lifespan 5

Sun Flash Accelerator F20 ESM Monitoring Utility 6

ILOM ESM Monitoring Option 6

XSCF ESM Monitoring Option 7

Known Issues 10

Documentation Errata 11

Preface

This document contains late-breaking information about the Sun Flash Accelerator F20 PCIe card from Oracle. Read this document first. This document is written for technicians, system administrators, authorized service providers (ASPs), and users who have advanced experience troubleshooting and replacing hardware.

Note – For specific installation instructions, see your system installation guide. For late-breaking information about installation and use of the Sun Flash Accelerator F20 PCIe card on your server, see the most recent version of the server product notes.

This preface contains the following topics:

- “UNIX Commands” on page v
- “Shell Prompts” on page vi
- “Related Documentation” on page vi
- “Documentation, Support, and Training” on page vi
- “Documentation Feedback” on page vii
- “UNIX Commands” on page v
- “Shell Prompts” on page vi
- “Related Documentation” on page vi
- “Documentation, Support, and Training” on page vi
- “Documentation Feedback” on page vii

UNIX Commands

This document might not contain information on basic UNIX[®] commands and procedures such as shutting down the system, booting the system, and configuring devices. Refer to the following for this information:

- Software documentation that you received with your system

- Oracle Solaris Operating System documentation, which is at
(<http://docs.sun.com>)

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

The documents listed as online are available at:

(<http://docs.sun.com/app/docs/prod/>)

Application	Title	Part Number	Format	Location
Installation, service, and configuration	<i>Sun Flash Accelerator F20 PCIe Card User's Guide</i>	820-7265	PDF, HTML	Online
Late-breaking information	<i>Sun Flash Accelerator F20 PCIe Card Product Notes</i>	820-7267	PDF	Online

Documentation, Support, and Training

These sites provide information about the following additional resources:

- Documentation (<http://docs.sun.com>)

- Support (<http://www.sun.com/support>)
- Training (<http://www.sun.com/training>)

Documentation Feedback

Submit comments about this document by clicking the Feedback[+] link at (<http://docs.sun.com>).

Include the title and part number of your document with your feedback:

Sun Flash Accelerator F20 PCIe Card Product Notes, part number 820-7267-14.

Sun Flash Accelerator F20 PCIe Card Product Notes

This document contains late-breaking information and known issues in the Sun Flash Accelerator F20 PCIe card.

- “System Requirements” on page 1
- “FMod Alignment” on page 4
- “Required ESM Replacement” on page 4
- “Monitoring ESM Lifespan” on page 5
- “Known Issues” on page 10
- “Documentation Errata” on page 11

System Requirements

- “Supported Servers” on page 1
- “Supported Operating Systems” on page 2
- “Minimum Supported Firmware Versions” on page 3
- “Minimum Required System Firmware Patches” on page 3
- “Solaris Performance Patch” on page 3

Supported Servers

This section describes servers that support the Sun Flash Accelerator F20 PCIe card. For detailed information about using this card with your server, see the product notes for your server, available at:

<http://docs.sun.com/app/docs/prod/servers#hic>

Note – The following servers are qualified to work with the Sun Flash Accelerator F20 PCIe card at the *initial release of the card*. Check your server product notes for confirmation that your server has subsequently been qualified.

SPARC Servers	x86 Servers
Sun SPARC Enterprise® T5120	Sun Fire™ X4540
Sun SPARC Enterprise T5220	
Sun SPARC Enterprise T5140	
Sun SPARC Enterprise T5240	
Sun SPARC Enterprise T5440	
Sun SPARC Enterprise M4000	
Sun SPARC Enterprise M5000	
Sun SPARC Enterprise M8000	
Sun SPARC Enterprise M9000	

Supported Operating Systems

The following table describes which combinations of operating systems and servers are supported for use with the Sun Flash Accelerator F20 PCIe card.

Platform	Oracle Solaris 10 05/09 (S10U7)	Oracle Solaris 10 05/08 Update 5 (U5) for x64	Red Hat Enterprise Linux 4.6, 5.1	SUSE SLES Linux 10 SP1	Windows 2003 Server R2 SP2	Windows 2008 Server
Sun SPARC Enterprise® T5120	✓					
Sun SPARC Enterprise T5220	✓					
Sun SPARC Enterprise T5140	✓					
Sun SPARC Enterprise T5240	✓					
Sun SPARC Enterprise T5440	✓					
Sun SPARC Enterprise M4000	✓					

Platform	Oracle Solaris 10 05/09 (S10U7)	Oracle Solaris 10 05/08 Update 5 (U5) for x64	Red Hat Enterprise Linux 4.6, 5.1	SUSE SLES Linux 10 SP1	Windows 2003 Server R2 SP2	Windows 2008 Server
Sun SPARC Enterprise M5000	✓					
Sun SPARC Enterprise M8000	✓					
Sun SPARC Enterprise M9000	✓					
Sun Fire™ X4540		✓				

Minimum Supported Firmware Versions

Firmware	Version
FMod firmware	D20R
SAS/SATA controller firmware	1.27.03
System firmware version for the Sun SPARC Enterprise T5120, T5220, T5240, T5240, and T5440 Servers.	7.2.7.d
System firmware version for the Sun SPARC Enterprise M4000, M5000, M8000, and M9000 Servers.	1093

Minimum Required System Firmware Patches

System	Patch
Sun SPARC Enterprise T5120 and T5220	139439-09
Sun SPARC Enterprise T5140 and T5240	139444-08
Sun SPARC Enterprise T5440	139446-08

Solaris Performance Patch

Download and install the following performance patch on your Solaris host:

- For Solaris 10 SPARC U4-U7, 138880-01 or later with MPT patch 141736-05

- For Solaris 10 x86 U4-U7, 138881-01 or later with MPT patch 141737-05

Available for download from the following web site:

<http://sunsolve.sun.com/show.do?target=patchpage>

Note – To enable maximum throughput from the MPT driver, add `mpt_doneq_thread_n_prop=8`; to `/kernel/drv/mpt.conf` and reboot the system.

FMod Alignment

You must properly align the FMods before using the card.

Solid state flash devices have block alignments typically aligned on 4-Kbyte boundaries, not the 512-byte boundaries of conventional disks. To maximize performance, partitions need to be aligned on 4-Kbyte boundaries. For directions on aligning the FMods, see the *Sun Flash Accelerator F20 PCIe User's Guide*. For more information on performance tuning, see

<http://wikis.sun.com/display/Performance/Home#Home-Flash>

Required ESM Replacement

The Sun Flash Accelerator F20 PCIe card contains a component called the Energy Storage Module (ESM) that functions similar to a battery backup. The ESM plays a critical role protecting data during power outages and enabling optimal card performance. When the ESM is online and functioning properly, the card operates in write-back mode (providing optimal performance). When the ESM is not functioning properly, the card functions in write-through mode. While data is secure in write-through mode, performance is dramatically decreased.

Based on the expected life of the ESM, and to maintain optimal card performance, plan to replace the ESM every two years. To help you monitor the age of the ESM, Sun provides two options for monitoring the ESM, see [“Monitoring ESM Lifespan” on page 5](#). Service the ESM (F371-4650) as described in the *Sun Flash Accelerator F20 PCIe User's Guide*.



Caution – If the ESM is not replaced at the recommended service interval, the level of stored energy will continue to degrade over time. Any data stored on the card is at risk of being lost if there is not enough stored energy to complete a write operation during a power failure. To avoid this risk, replace the ESM at the recommended service interval.

Note – The long-term durability of the ESM is affected by excessive heat. Refer to your server product notes to understand any slot or thermal restrictions that apply to your server.

Product documentation is available at:

<http://docs.sun.com/app/docs/prod/servers#hic>

Monitoring ESM Lifespan

Because the onboard ESM has a two-year lifespan, Sun provides two different methods that monitor how long an ESM has been installed, and notifies you when to replace the ESM.

One option is the Sun Flash Accelerator F20 ESM Monitoring Utility, a simple script that you install on your host server to track the life of the ESM. You must use this monitoring option for F20 cards with part number 511-1500-01.

The second option uses ILOM (or XSCF for M-series servers) to monitor the F20 card. ILOM will track the ESM lifespan and notify you when to replace the ESM. You must use this monitoring option for F20 cards with part numbers 511-1500-05 or greater, and with ILOM system firmware version 7.2.7.d or greater.

Sun Flash Accelerator F20 ESM Monitoring Utility

The Sun Flash Accelerator F20 ESM Monitoring Utility is a simple tool that you install on your host server to track the life of the ESM. Once installed, the ESM Monitoring Utility runs weekly to track the age of your ESM. The utility sends messages to the console and the `/var/adm/messages` file as the ESM approaches or exceeds the two-year replacement interval. Optionally, you can use an external monitoring tool to configure an SNMP trap that sends an email alert when these messages appear.

The utility can be run manually anytime to display the current ESM replacement data on all installed cards.

Note – Installation of this utility is required on cards with part number 511-1500-01 and 511-1275-03 or less to maintain optimal performance for the life of the card. This option will not work on cards with part numbers 511-1500-02 and 511-1275-04 or higher.

Download the ESM Monitoring Utility.

1. Go to
http://www.oracle.com/technology/software/sun_az_index.html
2. Locate the “Sun Flash Accelerator F20 ESM Monitoring Utility for Solaris 1.0”.
3. Log in (if required) to download the file.

To install the utility, follow the directions in the README file.

If you have multiple Sun Flash Accelerator F20 PCIe cards of the same age installed, consider replacing the ESMs at the same time to minimize system downtime. Service the ESM (F371-4650) as described in the *Sun Flash Accelerator F20 PCIe User’s Guide* (820-7265).

ILOM ESM Monitoring Option

For later-generation F20 cards (part number 511-1500-02 and 511-1275-04 or greater), ESM lifespan is automatically monitored by the ILOM system management firmware (system firmware version 7.2.7.d or greater) installed on your host.

Note – The Sun Fire X4540 requires ILOM version SW2.2 for ESM monitoring.

ILOM monitors ESMs by recording the Total_Time_On for each installed F20 card, and then issues warning messages (to the event log and to the host Solaris syslog) as an ESM approaches the end of its two year lifespan.

For example, one week before an ESM reaches its two-year threshold, ILOM issues this warning message:

```
"/SYS/MB/RISER1/PCI4/F20CARD ESM is approaching its lifespan. Please schedule a replacement as soon as possible."
```

When an ESM reaches its two-year threshold, ILOM issues this critical event message:

```
"/SYS/MB/RISER1/PCI4/F20CARD ESM has exceeded its lifespan. Please schedule a replacement as soon as possible."
```

Note – You can configure ILOM to send these alerts by email or SNMP trap. See your ILOM documentation for more information.

Service the ESM (F371-4650) as described in the *Sun Flash Accelerator F20 PCIe User's Guide* (820-7265).

Once you have replaced your ESM, use ILOM's standard fault clearing methods to remove the fault warnings; this also resets the F20 card Total_Time_On counter to 0. For more information about using ILOM, see:

<http://docs.sun.com/app/docs/coll/ilom3.0?l=en>

XSCF ESM Monitoring Option

For later-generation F20 cards (part number 511-1500-02 and 511-1275-04 or greater), ESM lifespan is automatically monitored by the XSCF system management firmware installed on your host.

XSCF monitors ESMs by recording the Total_Time_On for each installed F20 card, and then issues fault reports when an ESM approaches the end of its two year lifespan.

To view the ESM lifespan at any time, type:

```
XSCF> ioxadm lifetime
NAC                Total Time On  (% of life)
IOU#0-PCI#1        1685           0
IOU#0-PCI#3        1685           0
XSCF> ioxadm -v lifetime
NAC                Total Time On  (% of life)  Warning Time  Fault Time
IOU#0-PCI#1        1685           0             1008000       1051200
IOU#0-PCI#3        1685           0             1008000       1051200
```

30-days before an ESM reaches its two-year threshold, XSCF posts a message, similar to the following message, to the console log:

```
Mar 25 15:35:10 burl-m4000-0 fmd: SOURCE: sde, REV: 1.16, CSN:
0000000000 EVENT-ID:
144796b5-a7e2-4285-a3f1-30ce047767f3 Refer to
http://www.sun.com/msg/SCF-8000-9X for detailed information.
```

Once you receive this message, type the following command for more detail:

```
XSCF> fmdump -m
MSG-ID: SCF-8000-9X, TYPE: Fault, VER: 1, SEVERITY: Minor
EVENT-TIME: Thu Mar 25 15:35:10 EDT 2010
PLATFORM: SPARC Enterprise M4000 , CSN: 0000000000, HOSTNAME:
burl-m4000-0
SOURCE: sde, REV: 1.16
EVENT-ID: 144796b5-a7e2-4285-a3f1-30ce047767f3
DESC: An energy storage module is approaching its lifespan.
Refer to http://www.sun.com/msg/SCF-8000-9X for more information.
AUTO-RESPONSE: No immediate action is taken
IMPACT: Backup power may not be available in event of power loss.
REC-ACTION: Schedule a repair procedure as soon as possible to
replace the ESM.
```

When an ESM reaches its two-year threshold, XSCF posts this message to the console log:

```
Feb 17 12:49:24 burl-m4000-0 fmd: SOURCE: sde, REV: 1.16, CSN:
0000000000 EVENT-ID:
98198f1d-2e66-4635-90dd-5381b2bf2f1f Refer to
http://www.sun.com/msg/SCF-8000-AE for detailed information.
```

Once you receive this message, type the following command for more detail:

```
XSCF> fmdump -m
MSG-ID: SCF-8000-AE, TYPE: Fault, VER: 1, SEVERITY: Major
EVENT-TIME: Wed Feb 17 12:49:24 EST 2010
PLATFORM: SPARC Enterprise M4000 , CSN: 0000000000, HOSTNAME:
burl-m4000-0
SOURCE: sde, REV: 1.16
EVENT-ID: 98198f1d-2e66-4635-90dd-5381b2bf2f1f
DESC: An energy storage module has exceeded its lifespan.
Refer to http://www.sun.com/msg/SCF-8000-AE for more information.
```

```
AUTO-RESPONSE: No immediate action is taken
IMPACT: Backup power may not be available in event of power loss.
REC-ACTION: Schedule a repair procedure as soon as possible to
replace the ESM.
```

Note – You can configure XSCF to send these alerts by email or SNMP trap. See your XSCF documentation for more information.

Service the ESM (F371-4650) as described in the *Sun Flash Accelerator F20 PCIe User's Guide* (820-7265).

Once you replace your ESM, use the following command to clear the fault warnings and reset the F20 card Total_Time_On counter to 0:

```
XSCF> ioxadm lifetime -z IOU#x-PCI#y
```

Use of `ioxadm lifetime -z` requires fieldeng privileges.

For more information,, see the XSCF product documentation:

<http://docs.sun.com/app/docs/doc/819-6202-16>

Known Issues

This section describes hardware, firmware, and software issues known to exist at this release of Oracle's Sun Flash Accelerator F20 PCIe card.

CR ID	Description	Workaround
6868341	<p>In some cases, the server fails to recognize the card after a SATA controller firmware upgrade. The upgrade fails with an error message similar to the following:</p> <pre>MPT BIOS Fault 0Dh encountered at adapter PCI(02h,00h,00h)</pre> <p>In addition, the card does not appear in the server device tree.</p>	<p>Reboot the server a second time. The card is recognized after the second reboot.</p>
6881281	<p>FMod support clips may break if pressed too far during FMod removal.</p>	<p>Loosen the clip screws when replacing an FMod, as described in the <i>Sun Flash Accelerator F20 PCIe User's Guide</i>.</p>
6809711	<p>Under heavy I/O load in Solaris, a large number of messages may be seen indicating I/O retries (such as incomplete read- retrying). These messages occur due to a hardware bug in the LSI SAS controller that may incorrectly detect underrun conditions and report them to the driver. Overall data integrity is not compromised, but a performance impact may be observed due to the messages logged and retries required.</p>	<p>Ignore the retry messages or use an external means to throttle I/O throughput down to a level where these messages are not being produced in great numbers. The method required to throttle I/O will be very configuration and workload specific.</p>

CR ID	Description	Workaround
n/a	If there is a significant unexpected drop in card performance, the supercap module on the ESM should be checked.	Look at the back panel of the card to determine if the amber ESM Service Required LED is lit. Service the ESM as described in the <i>Sun Flash Accelerator F20 PCIe User's Guide</i> .
6881856	On occasion, the system BIOS may not give the card firmware enough time to load. When this situation occurs, the card cannot be seen by host server.	Reboot the server.
6951778	Repeated power cycling with minimum traffic causes DOM to zombie.	Do not run continuous power cycle tests with minimum traffic to the DOM. If power cycling is needed for qualification either ensure there is no traffic to the DOM including SMART transactions or allow the SuperCap to fully charge and run transactions to the DOM. Because the pre-existing state of a DOM is unknown, it could have been power cycled previously, the number of power cycles cannot be specified.

Documentation Errata

The label on the card faceplate might read "SuperCap" instead of "ESM".

