



Sun StorageTek™ 5800 System Release Notes

Sun Microsystems, Inc.
www.sun.com

Part No. 819-3813-10
May 2006, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright 2006 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Sun StorageTek, Java, AnswerBook2, docs.sun.com, et Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2006 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, États-Unis. Tous droits réservés.

Sun Microsystems, Inc. possède les droits de propriété intellectuelle relatifs à la technologie décrite dans ce document. En particulier, et sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs des brevets américains listés sur le site <http://www.sun.com/patents>, un ou les plusieurs brevets supplémentaires ainsi que les demandes de brevet en attente aux États-Unis et dans d'autres pays.

Ce document et le produit auquel il se rapporte sont protégés par un copyright et distribués sous licences, celles-ci en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Tout logiciel tiers, sa technologie relative aux polices de caractères, comprise, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit peuvent dériver des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux États-Unis et dans d'autres pays, licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Sun StorageTek, Java, AnswerBook2, docs.sun.com, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux États-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface utilisateur graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox dans la recherche et le développement du concept des interfaces utilisateur visuelles ou graphiques pour l'industrie informatique. Sun détient une licence non exclusive de Xerox sur l'interface utilisateur graphique Xerox, cette licence couvrant également les licenciés de Sun implémentant les interfaces utilisateur graphiques OPEN LOOK et se conforment en outre aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DÉCLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES DANS LA LIMITE DE LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE À LA QUALITÉ MARCHANDE, À L'APTITUDE À UNE UTILISATION PARTICULIÈRE OU À L'ABSENCE DE CONTREFAÇON.



Contents

Sun StorageTek 5800 System Release Notes 1

Features in This Release 1

List of Platform APIs 2

Known Issues 2

 Functional Limitations 3

 Installation and Initial Configuration Issues 3

 CLI Issues 4

 Client API Issues 7

 General Issues 7

Release Documentation 11

Service Contact Information 11

Sun StorageTek 5800 System Release Notes

This document contains important information about the Sun StorageTek 5800 system that was not available at the time the product documentation was published. Read this document so that you are aware of issues or requirements that can affect the installation and operation of the Sun StorageTek 5800 System.

- [“Features in This Release” on page 1](#)
 - [“List of Platform APIs” on page 2](#)
 - [“Known Issues” on page 2](#)
 - [“Release Documentation” on page 11](#)
 - [“Service Contact Information” on page 11](#)
-

Features in This Release

The Sun StorageTek 5800 system is an online storage appliance featuring a fully integrated hardware and software architecture in which the disk-based storage nodes are arranged in a symmetric cluster. The clustered and redundant design provides high availability, good performance, and exceptional data integrity.

The Sun StorageTek 5800 system provides the following features:

- Storage control, data, and metadata path operations distributed across the cluster to provide reliability and performance scaling
- Total symmetry in both hardware and software
- No single point of failure, with non-disruptive serviceability that includes fail-in-place components for deferred maintenance
- Optional storage of metadata associated with all data for easy reference and parallel search capability against an in-memory metadata index

- Automated integrity and placement algorithms that eliminate redundant array of inexpensive disks (RAID) configuration and volume management
- Fully integrated load and capacity balancing technology
- A virtual-view file system interface for auxiliary archive access using the Web-based Distributed Authoring and Versioning (WebDAV) protocol
- A Software Developers Kit (SDK) with Application Programming Interfaces (APIs), documentation, and an emulator

List of Platform APIs

This section provides an approved platform list for the Sun StorageTek 5800 system Application Programming Interfaces (APIs) including:

- Java™ library
- C library
- WebDAV

Known Issues

This section provide information about functional limitations and bugs filed against this version 1.0 product release. Note that if a recommended workaround is available for a bug, it follows the bug description.

The section contains the following topics:

- [“Functional Limitations” on page 3](#)
- [“Installation and Initial Configuration Issues” on page 3](#)
- [“CLI Issues” on page 4](#)
- [“Client API Issues” on page 6](#)
- [“General Issues” on page 7](#)

Note that if a recommended workaround is available for a bug, it follows the bug description.

Functional Limitations

The functional limitations of the version 1.0 release are as follows:

- Scaling to integrated multiple clusters is not enabled
- Disaster protection by way of user-level replication is the only native capability
- The maximum queryable size for a single metadata attribute is 512 bytes for data stored on a cluster and 256 bytes for data stored in the emulator. It is recommended that you adhere to a string size limit of 256 bytes in both the emulator and the cluster.
- The current behavior of indexable string attributes is as follows:
 - The first 512 bytes of the string value (256 bytes when using the emulator) are stored in the metadata database and can be queried against.
 - When using `RetrieveMetadata`, the entire data value is retrieved.
 - When using `QueryPlus`, only the queryable portion of the data value is retrieved (that is, the first 512 or 256 bytes).
 - In a query, string literals longer than the maximum queryable size are also silently truncated to the same limit.

Installation and Initial Configuration Issues

This section describes known issues and bugs related to installing and initially configuring the Sun StorageTek 5800 system.

Server Node Fails to Start

Bug 6403228 – When you attempt to bring up a server node with missing or faulty disks, that server node might fail during the startup procedure. Contact Sun Support to schedule replacement of the failed drives.

Unable to Change Service Node IP Configuration

Bug 6405531 – You can use `netcfg` and `ifconfig` to change the cluster administrative and data virtual IP addresses (VIPs). However, you cannot change the service node's IP address. Thus, if you have a conflict with the predefined factory service node IP, you will be unable to change it. Contact your Sun service representative for assistance.

CLI Issues

This section describes known issues and bugs related to using the command line interface (CLI).

ssh Command Error

Bug 6247537 – After a master failover occurs, the following message may appear on the console when you log in to the CLI:

```
WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!
```

```
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
```

```
Someone could be eavesdropping on you right now (man-in-the-  
middle attack)! It is also possible that the RSA host key has  
just been changed...
```

This message is innocuous and can be ignored.

No Warning When hwcfg --disable DISK-1XX:X Fails

Bug 6327227 – If disabling a disk fails (for example, because the disk cannot be unmounted), no warning message is displayed. Use `hwstat -v` to view the status of the disk.

Using mdconfig Before the Database Is Ready

Bug 6380366 – If you access the CLI and type the `mdconfig` command before the database is ready, the following error message is displayed:

```
Timed out waiting for the state machine
```

Workaround – Ensure that the database is operational by entering `sysstat` before using the `mdconfig` command.

Unrecognized Port Configuration

Bug 6405506 – If you configure an SMTP server and a port in the CLI using the `netcfg` command, the switch will not recognize the port you configured.

Executing Configuration Changes with Reboot

Bug 6406170 – When you make a configuration change, certain properties require a reboot to take effect. Once the change is entered, however, you can no longer determine the current value since the `netcfg` command shows the new (pending) value instead. You also cannot tell that the displayed value is a pending value and that a reboot is still required.

No CLI Error Message When System is Non-Operational

Bug 6241900 – If you log in to the system, enter the CLI `help` command, and then enter additional commands, such as `version` and `sysstat`, an immediate exit from the login session results. No error message is displayed.

Adding the -v Option to Certain Commands

Bug 6403938 – Should you inadvertently add the `-v` or `--verbose` option to certain CLI commands, such as `hwcfg`, where its use is not supported, the following error appears:

```
unknown error: Can't find resource for bundle
java.util.PropertyResourceBundle, key common.version.name
```

NullPointerException in Alert Code

Bug 6408010 – Should you use the `wipe` command, a `NullPointerException` appears on the console. The operation will continue despite the appearance of this error.

Inaccurate Disk Status Message

Bug 6409249 – Entering the commands `hwstat`, `hwstat -v`, and `sysstat -v` shows the disk status as disabled and not as initializing while the disk is being enabled. This is confusing and can lead to a duplicate enable operation being performed.

Timeout for Invalid Schema File

Bug 6411146 – If you attempt to commit an invalid `schema.xml` file with the `mdconfig -c` command, the CLI becomes non-operational and eventually times out.

Misleading Indicator of Utilization Relative to Total Space

Bug 6421305 – Note the following with the `df` command:

- `df` reports raw storage utilization

You can use `df` effectively to view raw storage statistics for each disk in the cluster. Therefore, *used* is not equivalent to the total number of Object bytes stored in the system. For example, it includes space consumed by data parity, Object headers and footers, query indexes and so on.

- `df` statistics are refreshed every three minutes

Storage utilization statistics reflected by `df` are refreshed every three minutes.

- 15% of storage capacity is reserved for recovery

When using `df` to view storage utilization, be aware that the system reserves 15% of raw storage. This space is available so data recovery can be completed should a node or disk fail on a *full* cluster.

Misleading Data Integrity Verified Message

Bug 6421293 – If you power up the system, access the CLI, and then enter `sysstat`, a *Data Integrity Verified* indication appears. At this point, however, the Data Doctor tool has not completed a lost fragment recovery cycle and is unable to determine whether or not all data fragments are indeed accounted for.

No Node Identifiers With `sysstat --verbose`

Bug 6421314 – If you enter the `sysstat` command and its `--verbose` option, the system output that appears does not include node identifiers.

Updating the Schema During a Load

Bug 6427699 – If you attempt to update the schema using `mdconfig` while a system load is occurring, the result is that the database becomes suspended in a create schema state.

Client API Issues

This section describes known issues and bugs related to using the client API.

IllegalStateException Error With a QueryPlus Operation

Bug 6395771 – Harmless errors are printed to `stdout` during query operations with a `select` clause. For a cluster, they are displayed in the log. For the emulator, these errors are written to the shell where the emulator is launched.

Data Space Not Reclaimed in Emulator

Bug 6403951 – The emulator supports the `Delete Metadata` operation of `NameValueObjectArchive.delete` and `hc_delete_ez`. However, the emulator does not remove the underlying data file when the last metadata record is deleted. The semantics are correct, but the underlying space is not reclaimed.

Emulator Installation Directory Must Not Contain Space Characters

Bug 6407770 – If you try to install the Software Development Kit (SDK), the emulator configuration script fails if the specified directory name contains spaces.

Limitations on the Overall Metadata Size of Stored Data Items

Bug 6427145 – When using the C API, the overall metadata size of a stored data item is limited to 76384 bytes. (The exact maximum size depends on many factors and should not be relied on). This limitation does not apply to data stored using the Java API.

Limitations on Metadata Values Containing Non-Printable Characters

Bug 6427141 – The behavior of metadata values that contain non-printable characters is not guaranteed. In particular, two known limitations are that metadata values cannot contain the null character, and that metadata values that contain the `<LF>` character will have the `<LF>` values silently removed in the stored value.

General Issues

This section describes general issues related to the Sun StorageTek 5800 system.

Cryptic Error Message When Deleting a Deleted Object

Bug 6187582 – Deleting an object twice produces the following error message:

```
ERROR: failed to retrieve object: request failed with status
400: no oa ctx.
```

The error message should say:

```
ERROR: failed to retrieve object: request failed with status
400: noSuchObject
```

Supported Metadata Type

Bug 6194366 – The supported metadata type field, double, does not work in views in the version 1.0 release.

Lost Client Connection Error

Bug 6268321 – If the client connection is lost during store and delete operations, the following error is displayed:

```
INFO: request failed with IOException...
```

```
Error in parsing the status line from the response: unable to
find line starting with "HTTP"
```

Workaround – The store failed because of a network problem. Retry the operation.

Inadvertent Double Storage

Bug 6187879 – Store operations become suspended and then time out in the client. If you then try the operation again, you might store the same data twice.

Workaround – Delete the second instance of the stored object.

Avoid Concurrent Delete Operations

Bug 6291970 – Known issues arise when you perform concurrent delete operations.

Workaround – Until this issue is resolved, perform delete operations from a single client connection at a time.

Extra Time Needed for Store After Idle Time

Bugs 6355668 and 6403926 – Initial input/output (I/O) operations on an otherwise idle cluster may be slower than usual.

Lack of Detail in Email Alerts

Bug 6398940 – With the exception of the mail header itself, email alerts do not provide cluster-specific information to enable you to distinguish between systems.

Workaround – Read the mail headers to determine where the alert originated. It will be the Administrative VIP configured for the cluster.

Shutting Down the Service Node

Bug 6402543 – When moving a rack, there is no way to power all the components down, since the CLI shutdown command powers down only the nodes.

Unable to Start or Shut Down the Database

Bug 6392770 – Starting up or shutting down the cluster without all nodes being online may force the query indexes to be rebuilt. Until the rebuilding process is finished, query results may be incomplete. See the `sysstat` command for the status of the query engine.

Inconsistent Date Reporting

Bug 6408658 – If you attempt to determine when a file or directory was last modified or created, be aware that there are inconsistencies in date reporting. For example, the directory listing, `getlastmodified`, might show a significant offset from Coordinated Universal Time (UTC), while the `creationdate` is about the same time without the offset. These times should be similar or very close.

Node BIOS Issue

Bug 6413587 – When the Sun StorageTek 5800 node BIOS is booting up, the following two conditions may occur:

- The BIOS fails to boot successfully and becomes non-operational.
- The BIOS detects a false fan failure on one or more fans and shuts down.

Workaround – Power on the node again through the front panel power switch. If this does not fix the problem, remove the power cord, wait 30 seconds, and then connect the power cord once more.

Switch Failover Does Not Produce Switch Failover Alert

Bug 6402478 – Despite the fact that a proper failover alert was not sent, the following alert appears indicating that you should perform a cluster assessment:

```
Cluster is booting or master failed over
```

Other than by visually inspecting the cluster, there is no way to determine if the switch has failed over.

Potential Disk Replacement Issues

Bug 6422741 – If you use a good disk from another node during a disk replacement, it causes the node to assume the identity of the previous disk owner.

Workaround – Use only new disks when performing a disk replacement.

Disk Write Cache Enabled

Bug 6424800 – With disk write cache enabled, there is a very small probability that a complete power outage or three simultaneous catastrophic disk failures can result in data loss for recently stored files. Testing of the failure scenarios presented here has not yet resulted in any data loss.

Rebooting Nodes While Operating on the Secondary Switch

Bug 6425530 – If your primary switch is not functioning and you are operating on the secondary switch, rebooting a node repeatedly fails to bring up the data VIP at start-up. In addition, the system configures some unusual interfaces, eventually escalates, and then reboots once more. Rebooting the entire cluster while running on the secondary switch causes the cluster to become non-operational.

Potential Disk Query Unavailable After Five Disk Failures

Bug 6423238 – In some instances, five disk failures will cause the query engine to become disabled.

Workaround – Reboot the cluster and the query engine will automatically repair itself in approximately 12 hours.

Release Documentation

Subject	Title	Part Number
Installation and initial configuration instructions	<i>Sun StorageTek 5800 System Getting Started Guide</i>	819-3808-10
System administration information	<i>Sun StorageTek 5800 System Administration Guide</i>	819-3810-10
System regulatory and safety information	<i>Sun StorageTek 5800 System Regulatory and Safety Compliance Manual</i>	819-3809-10

Note – In addition, there are two programmer guides available for the Sun StorageTek 5800 System, including a Client API Reference (part no. 819-3811-10) and an SDK Developer’s Guide (part no. 819-5501-10). For copies of these documents, contact your local sales or support representative.

Service Contact Information

If you need help installing or using this product, go to:

<http://www.sun.com/service/contacting>

