



Sun StorageTek™ 5800 System Release Notes

Version 1.1

Sun Microsystems, Inc.
www.sun.com

Part No. 819-7559-10
November 2007, Revision A

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

Copyright © 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

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, Java, Netra, Solaris, Sun Ray, Sun StorageTek 5800 System and SPARC are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

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

Products covered by and information contained in this service manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

Use of any spare or replacement CPUs is limited to repair or one-for-one replacement of CPUs in products exported in compliance with U.S. export laws. Use of CPUs as product upgrades unless authorized by the U.S. Government is strictly prohibited.

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 © 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

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

Sun, Sun Microsystems, le logo Sun, Java, Netra, Solaris, Sun Ray, Sun StorageTek 5800 System et SPARC sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-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 Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

Ce produit est soumis à la législation américaine en matière de contrôle des exportations et peut être soumis à la réglementation en vigueur dans d'autres pays dans le domaine des exportations et importations. Les utilisations, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou reexportations vers les pays sous embargo américain, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exhaustive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites. L'utilisation de pièces détachées ou d'unités centrales de remplacement est limitée aux réparations ou à l'échange standard d'unités centrales pour les produits exportés, conformément à la législation américaine en matière d'exportation. Sauf autorisation par les autorités des Etats-Unis, l'utilisation d'unités centrales pour procéder à des mises à jour de produits est rigoureusement interdite.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.



Please
Recycle



Adobe PostScript

Contents

Sun StorageTek 5800 System Version 1.1 Release Notes	1
Enhancements in This Release	2
Multicell Capability	2
GUI Management	2
Tape Disaster Recovery	2
Hot-Pluggable Disk Drives	3
Upgrade Considerations	3
Issues With the Power LED on the X2100 Server Service Node	4
Issues When Operating on Secondary Switch	4
Supported Client Operating Systems	5
Supported Browsers	6
Supported WebDAV Access	6
WebDAV and Multicell Configurations	7
Resolved Issues	7
Known Issues	9
Release Documentation	23
Service Contact Information	23

Sun StorageTek 5800 System Version 1.1 Release Notes

This document contains important information about the Sun StorageTek™ 5800 System, Version 1.1. Read this document so that you are aware of issues or requirements that can affect the installation and operation of the 5800 system.

This document contains the following sections:

- [“Enhancements in This Release” on page 2](#)
- [“Upgrade Considerations” on page 3](#)
- [“Issues With the Power LED on the X2100 Server Service Node” on page 4](#)
- [“Supported Client Operating Systems” on page 5](#)
- [“Supported Browsers” on page 6](#)
- [“Supported WebDAV Access” on page 6](#)
- [“Resolved Issues” on page 7](#)
- [“Known Issues” on page 9](#)
- [“Release Documentation” on page 23](#)
- [“Service Contact Information” on page 23](#)

Enhancements in This Release

This section includes a description of the major enhancements available with Version 1.1.

Multicell Capability

Version 1.1 provides support for the following configurations:

- One half-cell system (8 storage nodes)
- One full-cell system (16 storage nodes)
- Two full-cell systems (32 storage nodes)

Users can start with a half-cell configuration, add 8 nodes to make a full-cell configuration, and then add another full cell for a configuration with two full cells.

When the cell is expanded from 8 to 16 nodes, the 5800 system rebalances data previously stored in the 8-node cell across all 16 nodes.

When an additional full-cell system is added, new data can be stored to any cell in the configuration (it is usually stored on the cell with the most capacity).

As part of multicell support, some CLI commands were changed. The `ifconfig` command is no longer available and has been replaced with a `cellcfg` and `hivecfg` command.

GUI Management

Version 1.1 includes a Graphical User Interface (GUI) for performing configuration and management functions.

Tape Disaster Recovery

Version 1.1 allows for complete restoration of data in the event of a catastrophic loss of the system using the Network Data Management Protocol (NDMP) with NetVault, Version 7.4.5, from BakBone Software.

Hot-Pluggable Disk Drives

Version 1.1 includes hot-pluggable disks that can be removed and replaced easily after issuing the `hwcfg` command to disable the disk. After the disk is replaced, you issue the `hwcfg` command again to enable the disk.

Upgrade Considerations

The process of upgrading from Version 1.0.1 to Version 1.1 should be performed only by Sun service personnel. Before the service person performs the upgrade, please keep these considerations in mind:

- All nodes and disks on the 5800 system must be online and operating correctly before the upgrade. The service person may have to replace nodes and/or disks before the upgrade begins.
- If you have stored metadata with trailing white spaces, queries might return unexpected results after the upgrade. If your metadata includes trailing white spaces, the query engine might have to be wiped before the upgrade. Consult with the service person to determine if a wipe is required for your situation.
- During the upgrade, the OIDs of all data objects stored in the 5800 system are automatically converted to new Version 1.1 OIDs. If you have stored the Version 1.0.1 OIDs in a location external to the 5800 system (such as a database or file on the client side), the service person will use the OID conversion tool to convert the OIDs in that external location to Version 1.1 OIDs.

OID conversion is not required if you access data exclusively via Web-based Distributed Authoring and Versioning (WebDAV). OID conversion is also not required if you store data with metadata, and access the data by querying the metadata. The query returns an OID, which the clients then retrieve. The clients never store or remember OIDs.

- Depending on how much data you have on the system, and whether or not the query engine must be wiped, the upgrade could require that normal data input and output operations be interrupted for an average of 30 hours.

Consult Sun service personnel for complete information about the upgrade process.

Issues With the Power LED on the X2100 Server Service Node

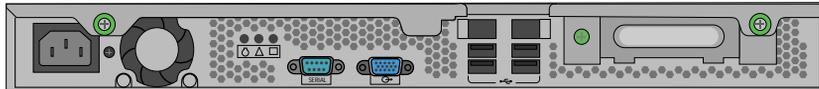
If your 5800 system uses the Sun Fire™ X2100 Server for a service node, you may notice some issues with the power LED.

To determine if your system uses the Sun Fire X2100 Server for a service node, compare the service node on your system with the front and back views of the Sun Fire X2100 Server, as shown in [FIGURE 1](#) and [FIGURE 2](#).

FIGURE 1 Front View of the Sun Fire X2100 Server



FIGURE 2 Rear View of the Sun Fire X2100 Server



If your system uses the Sun Fire X2100 Server, you can determine if the power is on or off to the service node when the power LED is lit by pushing in the CD/DVD drive eject button. If the LED on the caddy starts to blink, the service node is powered on. If the LED on the caddy does not light, then the service node is powered off.

Issues When Operating on Secondary Switch

If you receive email alerts and syslog messages indicating that one or more nodes are offline, inspect the 5800 system to determine if the secondary switch (the top switch on the system) is active. [FIGURE 3](#) shows the components on the back of the switch. If the port connection status LEDs are mostly lit for the top switch, that secondary switch is active.

FIGURE 3 Back of Switch

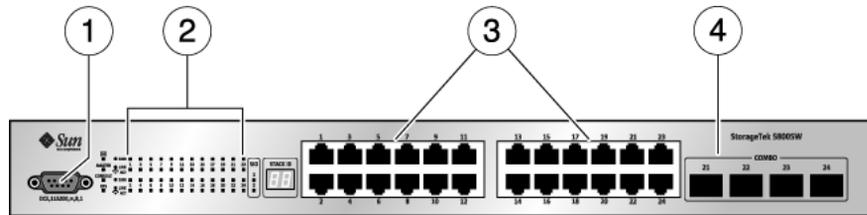


Figure Legend

-
- | | |
|---|-----------------------------|
| 1 | Serial port |
| 2 | Port connection status LEDs |
| 3 | Ethernet ports |
| 4 | Not used |
-

If the secondary (top) switch is active, contact Sun service immediately to arrange for a replacement of the primary (bottom) switch.

Until the primary switch is replaced, the system might encounter the following issues while operating from the secondary switch:

- One or more storage nodes might go offline. You will receive email alerts and/or syslog messages indicating that the nodes are offline. If necessary, use the CLI `sysstat` command, or check the GUI main screen to determine which nodes are affected. Locate the affected nodes on the system and manually power cycle those nodes. (Note that the nodes are numbered from the bottom up, starting with Node 1 on the bottom.) If more than two nodes were offline, data services might be unavailable until the nodes are manually power-cycled, and the query engine might need to be recreated, which will take at least 12 hours.
- One or more storage nodes might go into a panic mode and reboot. You will receive email alerts and syslog messages indicating that the nodes are offline and then that the nodes have rejoined. While the nodes are rebooting, data services might be unavailable and the query engine might have to be recreated, which will take at least 12 hours.

Supported Client Operating Systems

This section lists the client operating systems from which you can run applications that store, retrieve, and query data on the 5800 system. The applications can be written in either the C or Java™ programming languages, using the 5800 system application programming interface (API):

- Red Hat Enterprise Linux (32-bit) - RHEL4-U4

- Red Hat Enterprise Linux (64-bit) - RHEL4-U4
 - Solaris 9 SPARC® - Solaris 9 9/05 HW OS
 - Solaris 10 SPARC and x64/x86 - Solaris 10 11/06 OS
 - Windows Server 2003 R2 32-bit
-

Supported Browsers

You can access the 5800 system GUI using the following browsers:

- Internet Explorer 6.0 or later
- Mozilla 1.7 or later
- Firefox 1.5.0.9 or later

The browser must be running Versions 1.5 or 1.6 of the Java Runtime Environment.

Supported WebDAV Access

You can access the data on the 5800 system using Web-based Distributed Authoring and Versioning (WebDAV).

You can read the data on the system using WebDAV from any Hypertext Transfer Protocol (HTTP) browser running on any system that is on the same network as the 5800 system

For full read and write access to data on the 5800 system, you can use free software cadaver (command-line WebDAV client for Unix) or neon (HTTP and WebDAV client library with C interface). Consult the following URL for more information about this cadaver and neon:

<http://www.webdav.org/>

MAC OS X allows you to mount the 5800 system as a network share and gain read and write access via WebDAV to the data on the 5800 system.

The WebDAV implementation on the 5800 system has also been tested with the KDE Konqueror, Version 3, browser, which provides full read/write WebDAV access.

WebDAV and Multicell Configurations

The use of WebDAV in multicell configurations is not supported.

Resolved Issues

This section provides information about functional limitations and bugs that were described in the Version 1.0.1 release notes, and which have been resolved in this Version 1.1 product release.

TABLE 1 Issues Resolved in Version 1.1

Bug ID	Description
6403228	Server node fails to start.
6466326	Setting invalid NTP server IP address causes nodes to behave improperly.
6470857	Authorized clients must be reconfigured after upgrade.
6247537	ssh command error.
6247537	No warning when <code>hwcfg --disable DISK-1XX:X fails</code> .
6380366	Using <code>mdconfig</code> before the database is ready causes an error.
6405506	Unrecognized port configuration.
6241900	No CLI error message when system is nonoperational.
6403938	Adding the <code>-v</code> option to certain commands returns an error.
6421305	Clarification needed for reported storage utilization.
6450643	System hangs after <code>reboot --all</code> .
6466323	<code>passwd</code> command syntax documented incorrectly.
6194366	Supported metadata type double does not work in views.
6464058	System allows a table beyond maximum size.
6464061	System does not create index for fsViews with more than 15 fields.
6395771	IllegalStateException error with a QueryPlus operation.
6427145	Limitations on the overall metadata size of stored data items.
6427141	Limitations on metadata values containing nonprintable characters.
6466803	Performance might degrade with too many authorized clients.
6468507	Configuring authorized clients seems to hang CLI.

TABLE 1 Issues Resolved in Version 1.1 (*Continued*)

Bug ID	Description
6471588	SDK example scripts fail on csh.
6472509	RetrieveMetadata SDK example application fails with FILE argument.
6187582	Cryptic error message when deleting a deleted object.
6268321	Lost client connection error.
6291970	Concurrent delete operations cause problems.
6355668 and 6403926	Extra time needed for store after idle time.
6398940	Lack of detail in email alerts.
6392770	Unable to start or shut down the database.
6413587	Node BIOS issue.
6402478	Switch failover does not produce email alert.
6422741	Potential disk replacement issues.
6424800	Disk write cache enabled may cause data loss during power outage.
6423238	Potential disk query unavailable after five disk failures.
6465815	Full system cannot heal itself after six disk failures.
6422739	Query engine not re-created under heavy recovery load.
6481942	Intermittent failures to store or retrieve after reboot.
6473958	Multiple messages indicating nodes leaving cluster and rejoining
6481952	Reboot or shutdown during client access may corrupt query engine.
6483145	Reboot or shutdown with nonoperational mail server may corrupt query engine.
6450643	CLI may appear to hang if mail server is nonoperational.

Known Issues

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

Data Services Unavailable After Quorum Regained

Bug 6331523 – If transient node failures have caused the system to fall below quorum, data services might remain unavailable even after quorum is regained.

Workaround – Reboot the system to bring data services back online.

Data Space Not Reclaimed in Emulator

Bug 6403951 – The emulator supports the `Delete Record` 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.

Configuration Values Misleading Before 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 `cellcfg` 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.

Disks Unrecovered Count Might Be Inaccurate

Bug 6407787 – Even after the system has healed a disk, that disk may still be included in the *disks unrecovered* count displayed by the `sysstat` command.

Workaround – When the system is rebooted, the *disks unrecovered* count is reset to an accurate number.

Query Engine Remains in Starting or Stopped State

Bug 6450745 – In some rare cases, the query engine might get hung in starting or stopped state.

Workaround – Try rebooting the system to create the query engine and repopulate it with metadata. The process could take 12 to 48 hours.

Reboot Shutdown Message Not Clear

Bug 6451150 – Sometimes when you issue the CLI commands `shutdown` or `reboot`, the system returns the messages “It is not safe to shut down the system” or “It is not safe to reboot the system.” These messages indicate that the system is in the process of initializing the query engine.

Workaround – Although you can continue with the shutdown or reboot process, for best performance, wait until the query engine is fully initialized before proceeding.

Virtual Views Might Fail With Parse Errors

Bug 6458160 – The use of some characters in the file name specification for a virtual file system view might cause parsing errors.

Workaround – Do not use the `-` character or other Unicode characters to specify the file name of a virtual file system view.

System Must Be Run on Secure, Internal Network

Bug 6458653 – To ensure the integrity of data on the 5800 system, the system must be operated only on a secure, internal network.

Making Field Queryable Does Not Update Existing Data

Bug 6464055 – In the schema definition file, you can specify a metadata field as `queryable = false`. If you later change the schema definition file to indicate that `queryable = true` for that field, any data that you add to the system after the change includes that field as a queryable field. However, data that was previously stored on the system is not updated and is not queryable with that field.

Not Possible to Clear Schema After Configuration

Bug 6464866 – It is not possible to clear metadata schema after it is configured.

Workaround – If you need to clear fields from the metadata schema, either wipe all hive data (which clears the schema as a side effect), or contact Sun support for assistance.

Queries Might Cause Out of Memory Errors

Bug 6481476 – The system may respond to some queries with an out-of-memory error message.

Workaround – When developing queries using the Java API, set `maxFetchsize` in the range of 2000 – 5000.

Data Operations May Fail Even When CLI Reports Data Services Online

Bug 6489627 – When the system first starts up, data operations may fail, even if the CLI reports that data services online.

Workaround – Wait for all disks to come online (run the `sysstat` or `hwstat -v` commands to determine the number of disks online). If the problem persists after all working disks are online, retry the operations from the client following the best practices described in the *Sun StorageTek 5800 System Client API Reference Guide*. If necessary, contact Sun services to replace bad drives.

Multiple Concurrent Deletes Might Bring System Offline

Bug 6491877 – If clients attempt a large number of concurrent deletes, the system might go offline.

Workaround – Avoid large numbers of concurrent deletes; if the system goes offline, reboot to bring it back online.

Deletes Might Fail After Disk or Node Failure

Bug 6495883 – After a disk or node fails, delete operations to the system might fail for up to three minutes.

Workaround – Retry the deletions after three minutes.

Emulator Does Not Delete Data Link

Bug 6500528 – If a record is deleted from the 5800 system emulator using DeleteRecord, a WebDAV view might still show a link to the data, although not the data itself.

Workaround – Stop and restart the emulator.

Time Stamp Issues With SDK Example Program Stores, Retrieves, and Queries

Bug 6501640 – Time stamps on stores, retrieves, and queries in the SDK example programs might seem inconsistent.

Workaround – When planning stores, retrieves, and queries using the SDK Java example programs, be aware of the following:

- StoreFile interprets the time as local time zone unless the T..Z format (indicating UTC) is used. (For example, 1952-10-27T00:30:29.999Z.)
- RetrieveMetadata always displays the time in the time zone of the shell running the program.
- Query requires the T..Z UTC format.

Attributes in Nonextensible Namespaces Allowed to Be Changed

Bug 6502605 – The system erroneously allows you to change attributes such as queryable for nonextensible namespaces.

Workaround – Do not change attributes for nonextensible namespaces.

API Core Dump if Free Calls Are Out of Order

Bug 6507353 – C API core dumps if a query result is freed after the session is freed.

Workaround – Do not call `hc_session_free()` before the resultset is freed with `hc_qrs_free()`.

Restore Operation Might Fail After Reboot

Bug 6516036 – The first attempt to do a restore operation after the 5800 system is rebooted might fail with a message of *Connection Refused*.

Workaround – Retry the restore operation; it is expected to work on the second try.

During Back Up, Alert Messages Might Indicate Nodes Joining and Leaving

Bug 6518738 – While the system is performing a backup operation, it might generate multiple alert messages about nodes joining and leaving the system.

Workaround – You can safely ignore these messages.

Emulator Might Indicate Java Errors on Third Startup

Bug 6520374 – If you stop and then restart the system emulator soon after, it might fail with Java errors.

Workaround – Try restarting the emulator again.

Deleted File Might Still Appear in WebDAV View

Bug 6522009 – After you delete a file in a WebDAV view, the file might still appear to be present.

Workaround – Wait approximately five minutes and the file should no longer be displayed.

Cannot Access CLI From Linux

Bug 6531153 – You might not be able to access the CLI on a 5800 system from a system running Linux with a Kernel version greater than 2.6.17.

Workaround – Disable windows scaling on the Linux system using the following command:

```
echo 0 > /proc/sys/net/ipv4/tcp_window_scaling
```

Or, use `sysctl` to turn off window scaling.

Queries With Fields From More than One Table Might Fail

Bug 6533145 – A query to the data on the 5800 system that includes metadata fields that are stored in more than one table might fail.

Workaround – Make sure that fields that are queried together are grouped in the same table.

Query With Many Large Metadata Fields Might Fail

Bug 6535947 – A query to the data on the 5800 system that includes a high number (more than 40, for example) of large string metadata attributes might cause the query to fail.

Workaround – Limit the number of large string metadata attributes in a query to fewer than 40.

Innocuous WARN!! EOF Messages From Emulator

Bug 6538378 – The 5800 system emulator might display a number of WARN!! EOF error messages.

Workaround – These innocuous messages can be safely ignored.

CLI/GUI Might Erroneously Show Node as Disabled

Bug 6539494 – The CLI sensors command and the GUI Environmental Status Panel might erroneously indicate that nodes 1, 3, and 13 are disabled.

Workaround – Check the CLI `hwstat` command and the GUI Cell Summary Panel to determine if the nodes are actually disabled.

CLI/GUI Might Display Sensor Data for Offline Node

Bug 6539500 – For the first 5 or 10 minutes after a node goes offline, the CLI sensors command and the GUI Environmental Status Panel might erroneously report active voltages, temperatures, and fan speeds for the node.

Workaround – Wait a few minutes and the CLI and GUI should report sensor data as disabled for the offline node.

Add/Delete Metadata Operation Might Fail

Bug 6541837 – Rarely, an add or delete metadata operation for an object might fail if the system has not released the lock put on that object during a previous operation.

Workaround – If the symptom persists for more than 30 minutes, reboot the system.

Cannot Receive Email Alerts on Mail System That Requires Authentication

Bug 6542247 – You cannot use an SMTP server that requires authentication to receive system alert emails.

Workaround – Configure the 5800 system with an SMTP server that does not require authentication.

Switch 1 Might Fail With No Email Alert

Bug 6554457 – In some cases, switch 1 might fail over to switch 2, but no email alert might be sent to indicate this.

Transient Network Cable Failures Might Cause Failover and Missing Heartbeat Messages

Bug 6557612 – If the network cable connection on a node is experiencing transient failures, the 5800 system might log missing heartbeat messages as well as messages indicating that the switch has failed over.

Workaround – Report the symptom to Sun service, and schedule a replacement of the node or network cable.

A Large Object Stored When System Is Nearing Capacity Might Generate Error Messages and Might Not Be Removable

Bug 6558322 – If a client stores a large (greater than 1000 MB) object as the 5800 system is nearing capacity, the system might generate warning message about being unable to store the object. Also, the healing processes on the system might not be able to remove any fragments of the object that were successfully stored.

Workaround – Do not store objects larger than 1000 MB when the system is nearing capacity. A cell has reached capacity when any one of its disks has reached 80% raw utilization. To display the raw utilization of the disks in a cell, issue the CLI command `df -p`.

Unicode Supplemental Characters Erroneously Accepted in Schema

Bug 6562925 – The system does not reject the metadata schema file when the metadata name field contains one or more Unicode supplemental characters.

Workaround – Do not use Unicode supplemental characters for the metadata name field.

Second Attempt at wipe Might Fail

Bug 6566083 – If the `wipe` command fails and then you try to issue it again immediately, it might fail again.

Workaround – If the `wipe` command fails, reboot the system and then try the command again.

Node Might Continuously Reboot

Bug 6570304 – Some hardware or software failures can cause a node to reboot repeatedly. Such a situation will be accompanied by email alerts and/or external syslog messages indicating that the node is leaving and joining.

Workaround – Power the node down by holding down the power button on the node that is exhibiting this behavior. Call Sun service to arrange to have the node replaced.

reboot --all Fails While System Is on Secondary Switch

Bug 6570324 – The `reboot -all` command fails if the system is running on the secondary switch. The command requires both switches to be online.

Store Operations Might Fail During Cell Expansion

Bug 6573144 – During the process of expanding a cell from 8 to 16 nodes, store operations to the system may time out.

Workaround – When programming applications, use retry loops within API calls to handle timeouts during cell expansion. One immediate retry should be sufficient in the great majority of cases.

Lower-Case “e” Not Recognized as Exponent

Bug 6577783 – The 5800 system does not recognize a lower-case “e” as a symbol for exponent in queries.

Workaround – Use an upper-case ‘E’ to symbolize exponents in queries.

Issue With Backups Made From 8-Node System Used to Restore to 16-Node System

Bug 6580181 – You cannot use a backup made of an 8-node system to restore data to a 16-node system until at least one backup session is created from the 16-node system.

Workaround – After you expand the system from 8 nodes to 16 nodes, perform one backup of any length from the 16-node system. You can then use this and any previous backups of the 8-node system to restore data to the 16-node system.

Configuration Changes Might Not Trigger Alerts

Bug 6582274 – When multiple system parameter are changed using the `cellcfg` command, all of the changes might not issue alerts.

Workaround – If you receive an alert indicating that a parameter has been changed using `cellcfg`, keep in mind that other parameters might also have been changed. Use the `cellcfg` command to check the current settings of all parameters.

Client Connection Might Fail With `java.net.ConnectionException`

Bug 6582486 – A connection attempt from a client to the 5800 system might fail with a `java.net.ConnectionException` error.

Workaround – Retry the connection.

wipe Commands in Succession Might Cause System to Disable Disks

Bug 6584310 – If you issue the `wipe` command and then retry the command without waiting for the wipe to be complete, the system might disable disks.

Workaround – Wait for one wipe action to complete before retrying the command again. If the system has already started disabling disks, reboot the system, and then reenables the disabled disks.

Before Restore Operation, System Must Be Wiped and Rebooted

Bug 6584329 – A restore operation might not work correctly if the system has not been wiped and then rebooted before the restore begins.

Disk Fail or Disable Causes Severe Log Messages

Bug 6585878 – If a disk fails, or if a Sun service technician disables a disk, you might see severe error messages in the external `syslog` host.

Workaround – You can ignore these messages; they simply reflect the fact that the disk has failed and can be ignored.

Some Valid C API Queries Might Return HCERR_BAD_REQUEST

Bug 6588218 – Some valid C API queries to the 5800 system might return the error code `HCERR_BAD_REQUEST`, which seems to indicate that the query is not valid.

Workaround – Use `hc_session_get_status()` to determine if the error string from the query contains the substring `Relalg server involved in current operation failed`. If so, retry the query.

Shutting Down Emulator In Browser Does Not Work for Red Hat 4 Systems

Bug 6589653 – If you are running the 5800 system emulator on a system running Red Hat Version 4, you might not be able to shut down the emulator via the browser.

Workaround – Kill the process manually. The simplest way to do this is to start the emulator from a dedicated command prompt without running it in the background, and use `Ctrl-C` when it is idle to exit the program.

Retrying Restore Operation Might Require 10 Minute Wait

Bug 6595040 – If a restore operation fails, you may have to wait approximately 10 minutes while the system reclaims socket resources.

Some Node and Disk Alerts From Start Up or Shut Down Might Be Missing From the Logs

Bug 6601977 – When a system starts up or shuts down, it might not send a complete set of email and log alerts for all nodes and disks.

reboot --all Might Cause Switch Split Brain

Bug 6603323 – Issuing the `reboot -all` command might cause a switch “split-brain” situation, in which neither switch is fully functioning as the primary switch and both switches are performing some of the primary switch’s duties.

Workaround – Call Sun service to help you troubleshoot and correct the problem.

Delay in Disks Coming Online After shutdown Command

Bug 6604018 – After issuing the `shutdown` command from the CLI, you might have to wait up to two hours when you restart the system before all disks come online.

storeObject Returns Incorrect object_ctime

Bug 6609313 – When you store an object with the `storeObject` API function, the `object_ctime` reported for the object might not match the `object_ctime` that is actually stored with the object.

Workaround – To determine the `object_ctime` that was actually stored with the object, retrieve the metadata for the object after the store operation is completed. The system metadata retrieved will include the `object_ctime` that was actually stored with the object and inserted in the query engine.

Queries of Type Binary Metadata Might Return Incorrect Matches

Bug 6612017 – If you issue a query on a metadata field of type binary, entries in the query engine that include the first bits specified in the query are returned as matches, even if the entry includes more bits than were specified in the query. For example, suppose an entry for binary field, `bfield`, contains the value ABCDEFGHIJ. A query on `bfield = "ABCD"` will return a match to that entry.

Cannot Start Restore on Wiped Cell

Bug 6612146 – You cannot start a restore on a cell that has been wiped in preparation for the restore.

Workaround – After you wipe the cell, reboot it.

Cannot Perform Backup During Cell Expansion

Bug 6612244 – During the process of expanding a cell from 8 to 16 nodes, when you are running the `celladm expand` command, you cannot backup data from the cell.

Schemas with fsViews Cannot Be Defined in GUI

Bug 6613234 – Because the Files Only at Leaf Level checkbox on the Set Up Virtual File System panel in the GUI does not work correctly, it is not possible to use the GUI to define schemas that include fsViews.

Workaround – Use the CLI `mdconfig` command to define schemas that include fsViews.

WebDAV GETs Might Fail for File Names Longer than Directory Names

Bug 6613735 – If the length attribute for one or more directory fields is less than the length attribute for a filename field, WebDAV GETs may fail for file names that are longer than directory names.

Workaround – Specify equal lengths for directory and filename fields. For example, instead of specifying:

```
<namespace name="space1" writable="true" extensible="true">
  <field name="dir1" type="string" length="2" />
  <field name="dir2" type="string" length="2" />
  <field name="fname" type="string" length="128" />
</namespace>

<fsView name="HashDirs" filename="${ofoto.fname}" filesonlyatleaflevel="false">
  <attribute name="space1.dir1" unset="unk" />
  <attribute name="space1.dir2" unset="unk" />
</fsView>
```

specify the following:

```
<namespace name="space1" writable="true" extensible="true">
  <field name="dir1" type="string" length="128" />
  <field name="dir2" type="string" length="128" />
  <field name="fname" type="string" length="128" />
</namespace>

<fsView name="HashDirs" filename="${ofoto.fname}" filesonlyatleaflevel="false">
  <attribute name="space1.dir1" unset="unk" />
  <attribute name="space1.dir2" unset="unk" />
</fsView>
```

Stores of Very Large Files Might Fail

Bug 6615347 – Storing a very large file (greater than 400 GB) might fail with an error similar to the following:

```
com.sun.honeycomb.common.ArchiveException: Failed to get system
metadata from the fragments
```

Workaround – Break the file into smaller pieces and then retry the store operation.

After Failed Restore Operation, New Restore Operation Might Fail

Bug 6616306 – If a restore operation fails and you initiate a new restore operation, the new restore operation might fail if it is started too soon.

Workaround – Wait at least 20 minutes before initiating a new restore operation.

Restrictions on Dates in C API

Bug 6619221 – In the C API, time stamps supplied to `hc_nvr_add_timestamp` and dates supplied to `hc_nvr_add_date` should be in the range 00:00:00 January 1, 1970 UTC to 00:00:00 January 1, 2038 UTC. This is due to limitations in converting “seconds since the epoch” to the human-readable dates used for metadata storage. The Java interface is not as restrictive; however, dates outside of these limits stored in Java might not be retrieved properly by the C API.

JAVA API and SDK RetrieveMetadata Program Always Return False Indexed Status

Bug 6621320 – In the JAVA API and SDK `RetrieveMetadata` program, the `SystemRecord.isIndexed()` method always returns `False`.

Workaround – Ignore the `SystemRecord.isIndexed()` value.

After Restoration of Data, Backup Might Remain Enabled

Bug 6624848 - After you have restored data to a system following a disaster, you might not be able to resume backup of the system.

Workaround - Contact Sun service for assistance.

Usage Message for SDK Example Application CheckIndexed is Incorrect

Bug 6625515 – The usage message displayed by the system for the SDK Java example application CheckIndexed is actually the usage message for the RetrieveMetadata example application.

Workaround – Refer to the *Sun StorageTek 5800 System SDK Developer's Guide* for the correct usage for CheckIndexed.

max_records Variable Should Be Renamed results_per_fetch

Bug 6627590 – The signature for `hc_query_ez()` in the C API `hcclient.h` has an int variable named `max_records`. This int variable should be renamed `results_per_fetch` to make its function more apparent. The *Sun StorageTek 5800 System Client API Reference Guide* refers to the variable as `results_per_fetch`, but in the code it is named `max_records`. Functionality is not affected.

CLI Reboot Generates Connection Refused Message

Bug 6628840 – Occasionally, when you attempt to reboot the 5800 system using the CLI command `reboot`, the system returns the error message `Connection Refused`.

Workaround – The error message indicates that one or more storage nodes did not reboot. Wait at least 10 minutes and then issue the CLI command `reboot --all`.

Release Documentation

The following table lists the documents for this product. The online documentation is available at:

<http://docs.sun.com/app/docs/prod/stortek.5800#hic>

Title	Part Number	Format	Location
<i>Sun StorageTek 5800 System Regulatory and Safety Compliance Manual</i>	819-3809-xx	HTML PDF	Online
<i>Sun StorageTek 5800 System Administration Guide</i>	819-7555-xx	HTML PDF	Online
<i>Sun StorageTek 5800 System Site Preparation Guide</i>	820-1635-xx	HTML PDF	Online
<i>Sun StorageTek 5800 System Client API Reference Guide</i>	819-7557-xx	PDF	Online
<i>Sun StorageTek 5800 System SDK Developer's Guide</i>	819-7558-xx	PDF	Online

Service Contact Information

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

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

