



# Sun StorageTek™ 5800 System Release Notes

---

Version 1.0.1

Sun Microsystems, Inc.  
www.sun.com

Part No. 819-7553-11  
January 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.

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, and 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

---

Upgrading to Version 1.0.1	1
Enhancements in This Release	4
New CLI Option to Display System Firmware Version	4
Enhanced System Status Monitoring	5
New CLI Option to Set Address for Service Node	6
New CLI Option to Power Off All Components	6
Enhanced Schema Definitions	6
Supported Clients	18
Resolved Issues	18
Known Issues	21
Functional Limitations	21
Hot-Swap Limitations	21
Installation and Initial Configuration Issues	22
CLI Issues	23
Metadata Issues	25
Client API and SDK Issues	26
General Issues	28
Release Documentation	33
Service Contact Information	34



# Sun StorageTek 5800 System Release Notes

---

This document contains important information about the Sun StorageTek 5800 system, Version 1.0.1. 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.

- [“Upgrading to Version 1.0.1” on page 1](#)
- [“Enhancements in This Release” on page 4](#)
- [“Supported Clients” on page 18](#)
- [“Resolved Issues” on page 18](#)
- [“Known Issues” on page 21](#)
- [“Release Documentation” on page 33](#)
- [“Service Contact Information” on page 34](#)

---

## Upgrading to Version 1.0.1

Use the upgrade command to upgrade the Sun StorageTek 5800 system cluster to Version 1.0.1.

---

**Note** – Before beginning an upgrade, make sure that the two-gigabit Ethernet network switches on the Sun StorageTek 5800 system are operating correctly. On the back of the system, the first 16 LEDs on the primary switch (the one on the bottom) should be lit. The first 16 LEDs on the secondary switch (the one on the top) should not be lit. Also, issue the `sysstat` CLI command to make sure that the system reports 16 nodes and 64 disks online before beginning the upgrade. See [“Enhanced System Status Monitoring” on page 5](#) for information about the `sysstat` command.

---

To perform an upgrade, follow these steps:

1. **Make sure you have a copy of your current schema definition file.**

Upgrading to Version 1.0.1 deletes your schema definition file, but you can replace it after the upgrade. Your existing data is not lost, but querying will not be available for approximately 12 hours while the system repopulates the database.

If you do not have a copy of your schema definition file, use the CLI command `mdconfig -d` to print the current schema file.

---

**Note** – If you use the `mdconfig -d` command to print your schema definition file, the namespaces and fsViews that come preinstalled on your system are included in the output. You must delete these preinstalled namespaces and fsViews before using the schema definition file with Version 1.0.1, or else the schema will not work correctly. The preinstalled namespaces are the `system` and `filesystem` namespaces. Preinstalled fsViews are as follows: `oidByGidUid`, `oidByGidUidWithTar`, and `oidbyGidUidMimetype`.

---

2. **Edit your schema definition file to define lengths for fields and include table layouts appropriate to your configuration. See [“Enhanced Schema Definitions” on page 6](#) for information about these new features for this version.**

---

**Note** – Your schema definition file will work without updates, but for best performance, edit the file to take advantage of the new features.

---

3. **Insert the DVD into the drive in the service node. Then, log into the CLI and type the following:**

```
hc $ upgrade dvd
```

---

**Note** – You cannot stop the upgrade once it’s started. You also cannot remove or add a node during the upgrade.

---

Once the upgrade begins, the system checks to see that all nodes and disks are available and requires that you answer a few questions confirming the upgrade. After that, the process does not permit any changes or interruptions. However, it does display information about what is going on throughout the process. For example, the system displays messages as it unpacks the install image, upgrades the service node, the switches, and finally each node in the cluster.

If there is a problem with the images during an upgrade, the nodes will boot from the previous version.

If there is a power outage during an upgrade, the system might or might not come back online, depending on where in the process the upgrade was interrupted. If the system does not become operational when the power resumes, call your Sun Service representative.

During the last phase of the upgrade process, the system reboots. First, the service node reboots. Next, each node except one reboots. Then, the switches reboot, and finally, the last node reboots.

---

**Note** – During the reboot, you are logged out of the CLI and are unable to connect again until the system has come back with the new image.

---

**4. Use the CLI command `netcfg clients` to reconfigure your authorized clients.**

You must reconfigure your authorized clients after the upgrade or no data traffic will be allowed on the system. If all clients are to be allowed access, issue the command with the `all` option (`netcfg clients all`).

**5. Use the CLI command `mdconfig` to test your updated schema definition file.**

Make sure the file reports no errors and that it defines the table layout you expect. Refer to the *Sun StorageTek 5800 System Administration Guide* for more information about the syntax and use of the `mdconfig` command.

**6. Use the CLI command `mdconfig -c` to activate your updated schema definition file.**

Refer to the *Sun StorageTek 5800 System Administration Guide* for more information about the syntax and use of the `mdconfig -c` command.

**7. You can access data stored on the Sun StorageTek 5800 system immediately, but wait approximately 12 hours before issuing queries on the data to allow for full repopulation of the query engine.**

You can use the CLI command `sysstat` to get some feedback about the progress of repopulation. See [“Enhanced System Status Monitoring” on page 5](#) for more information about the `sysstat` command.

---

# Enhancements in This Release

This section describes enhancements included in this version of the Sun StorageTek 5800 system software.

## New CLI Option to Display System Firmware Version

The `version` command has been enhanced to display the system firmware version. To display the firmware version for each node, issue the `version` command followed by `-v` as follows.

---

**Note** – In normal operation, all nodes should be running the same version of SMDC (Service Management Daughter Card) and the same version of BIOS (Basic Input Output System).

---

```
hc $ version -v
Honeycomb 1.0.1 release [1.0-38]
Service Node:
    BIOS Version: 1.1.1
    SMDC Version: 4.13
Switch:
    Overlay Version: 9175
hcb101:
    BIOS Version: 0.0.8
    SMDC Version: 4.18
hcb102:
    BIOS Version: 0.0.8
    SMDC Version: 4.18
hcb103:
    BIOS Version: 0.0.8
    SMDC Version: 4.18
hcb104:
    BIOS Version: 0.0.8
    SMDC Version: 4.18
...
```



# Enhanced System Status Monitoring

The `sysstat` command allows you to monitor the StorageTek 5800 cluster state. For Version 1.0.1, the output of the `sysstat` command has been enhanced to provide more detailed information about the data checks that the system performs.

Issue the `sysstat` command from the CLI:

```
hc $ sysstat
Silo 1, Cell 1: Cluster Online. Estimated free space: 7T.
16 nodes online, 64 disks online, 0 disks unrecovered
Data VIP 10.7.225.222. Admin VIP 10.7.225.223.
Data services Online, Query Available
Data Integrity check not completed since boot
Data Reliability check last completed at Aug 15, 2006 1:31:23 PM
Query Integrity check last completed at Aug 15, 2006 1:20:21 PM
```

The output that the `sysstat` command produces is explained below. Data reported is for all online disks in the entire cluster.

- `Data services Online` indicates that the system is available to read and write to via the API, while `Data services Offline` means that the system is not available to read and write to via the API.
- `Query Available` indicates that it is possible to perform queries and access the system via WebDAV. `Query Unavailable` means it is not possible to perform queries and WebDAV access is disabled.
- `Data Integrity check` indicates when the system last completed checking each fragment on the system for integrity against bit rot. Each cycle of this testing might take up to one week to complete, so the check will be listed as `not complete` for the first week after a system reboot.
- `Data Reliability check` indicates when the system last completed a full cycle of testing to detect and recover any missing fragments, indicating that the system has full reliability. Each cycle of this testing takes approximately 12 hours to complete, so the check will be listed as `not complete` for the first 12 hours after a system reboot.
- `Query Integrity check` indicates when the system last completed indexing all metadata stored by clients, or verified that all metadata is already indexed. When this cycle is complete, you can be sure that any query of metadata will return complete results. Each cycle of this check takes approximately 12 hours to complete, so the check will be listed as `not complete` for the first 12 hours after a system reboot.

## New CLI Option to Set Address for Service Node

The `ifconfig` command has been enhanced to allow you to set the address of the service node. If there is a conflict between the factory configured IP address on the service node and an address on your network, you can now set a new address for the service node using `ifconfig` with the `sp0` option.

Issue the `ifconfig` command followed by `sp0` and the IP address you want to use:

```
hc $ ifconfig sp0 10.9.224.40
```

## New CLI Option to Power Off All Components

The `shutdown` command has been enhanced to allow you to power off the service node as well as the nodes. Issue the `shutdown` command with the `-a` option to power off the nodes and the service node. (The `shutdown` command without the `-a` option powers off just the nodes.)

```
hc $ shutdown -a
```

## Enhanced Schema Definitions

This version of the Sun StorageTek 5800 system includes enhancements to the Document Type Definition (DTD) for the schema definition file.

The updated DTD is as follows. The sections that follow the DTD provide an example of a schema definition file that uses the new DTD and also describe the new features in the DTD for this version.

```

<!DOCTYPE HCMETADATACONFIG [

<!ELEMENT metadataConfig (schema, fsViews, tables)>

<!ELEMENT schema (namespace*, field*)>
<!ELEMENT namespace (namespace*, field*)>
<!ELEMENT field EMPTY>

<!ELEMENT fsViews (fsView*)>
<!ELEMENT fsView (attribute+)>
<!ELEMENT attribute EMPTY>

<!ELEMENT tables (table*)>
<!ELEMENT tables (column+)>
<!ELEMENT column EMPTY>

<!ATTLIST namespace name CDATA #REQUIRED>
<!ATTLIST namespace writable (true|false) true>
<!ATTLIST namespace extensible (true|false) true>

<!ATTLIST field name CDATA #REQUIRED>
<!ATTLIST field type (long|double|string) #REQUIRED>
<!ATTLIST field length CDATA>
<!ATTLIST field queryable (true|false) true>
<!ATTLIST field indexable (true|false) true>*

<!ATTLIST fsView name CDATA #REQUIRED>
<!ATTLIST fsView filename CDATA #REQUIRED>
<!ATTLIST fsView namespace CDATA>
<!ATTLIST fsView readonly (true|false) false>

<!ATTLIST attribute name CDATA #REQUIRED>

<!ATTLIST table name CDATA #REQUIRED>

<!ATTLIST column name CDATA #REQUIRED>

]>

```

\* indexable is allowed as a synonym for queryable

## Example of Schema File Using New DTD

The following is an example of a valid schema using the revised DTD.

```
<?xml version="1.0" encoding="UTF-8"?>
<metadataConfig>
  <schema>
    <namespace name="mp3" writable="true" extensible="true" >
      <field name="artist" type="string" length="256" queryable="true"/>
      <field name="album" type="string" length="256" queryable="true"/>
      <field name="title" type="string" length="256" queryable="true"/>
      <field name="year" type="long" queryable="true"/>
    </namespace>
    <namespace name="dates" writable="true" extensible="true" >
      <field name="year" type="long" queryable="true"/>
    </namespace>
  </schema>
  <fsViews>
    <fsView name="byArtist" filename="{mp3.title}.mp3">
      <attribute name="mp3.artist"/>
      <attribute name="mp3.album"/>
    </fsView>
    <fsView name="byAlbum" filename="{mp3.title}.mp3">
      <attribute name="mp3.album"/>
    </fsView>
  </fsViews>
  <tables>
    <table name="reference"/>
      <column name="mp3.artist"/>
      <column name="mp3.album"/>
      <column name="mp3.title"/>
      <column name="dates.year"/>
    </table>
  </tables>
</metadataConfig>
```

## Improved Performance With New Table and Column Elements

Using the new table element in the DTD, you can partition the schema into tables and specify each metadata field as a column within a particular table. You can greatly improve the performance of query and store operations by grouping metadata fields that commonly occur together in the same table and by separating

metadata fields that do not commonly occur together into separate tables. Objects stored in the Sun StorageTek 5800 system become rows in one or more tables, depending on which fields are associated with that data.

For example, suppose you specify columns in a table *reference* in the schema definition file as follows:

```
<table name="reference"/>
  <column name="mp3.artist"/>
  <column name="mp3.album"/>
  <column name="mp3.title"/>
  <column name="dates.year"/>
</table>
```

The reference table you create would have the layout shown in [TABLE 1](#).

**TABLE 1** Example Table (*reference* Table)

OID	mp3.artist	mp3.album	mp3.title	dates.year
Object1	Benny Goodman	The Very Best of Benny Goodman	St. Louis Blues	2000
Object2	Rod Stewart	The Very Best of Rod Stewart	Maggie May	2001
Object3	Bing Crosby	Null	I'm Dreaming of a White Christmas	Null

When an object is stored in the Sun StorageTek 5800 system that has any of the specified metadata attributes associated with it (mp3.artist, mp3.album, mp3.title, or dates.year), that object OID is listed as a row in the reference table and the values of the attributes are listed in the column that corresponds to that attribute. If no value is assigned to an attribute for that object, no value is listed in the corresponding column.

If the object has other metadata associated with it, that object will also be stored in whichever tables include that other metadata as columns.

To take full advantage of the new `tables` element to maximize query and store performance, it is important to understand field length values and to carefully plan which fields are grouped together in which tables. It is also important to plan which fields will be listed together in which indexes. [“Planning Tables” on page 10](#) and [“Planning Indexes” on page 13](#) provide more information on these considerations. For a summary of how to use the new table feature to maximize query performance, see [“Summary of Best Practices for Optimal Query Performance” on page 15](#).

## *New Length Attribute for Fields*

With this version, you can specify a `length` attribute for fields of type `string`. The `length` attribute is important because there are limits to the number of bytes that each table row and each index can store. See [“Planning Tables” on page 10](#) and [“Planning Indexes” on page 13](#) for more information.

---

**Note** – In this version, the Sun StorageTek 5800 emulator now supports the same string length as that supported by the Sun StorageTek 5800 system, within the specified limits.

---

In future releases of the Sun StorageTek 5800 system, trying to store a string that is longer than the specified length will result in an error message. For Version 1.0.1, you can successfully store a string that is longer than the specified length, but the string is truncated to fit the specified length, which might affect query results.

For example, suppose you specify the length for the field `mp3.album` as 12 and then store *The Very Best of Rod Stewart* in field `mp3.album` with object ID `010001287969bbdbb511dab28000e08159646c000019970200000000`.

If you issue a query that includes the following specification:

```
"mp3.album = 'The Very Best of Benny Goodman'"
```

The system will return object ID `010001287969bbdbb511dab28000e08159646c000019970200000000` because the first 12 characters match *The Very Best of Rod Stewart*.

---

**Note** – The Sun StorageTek 5800 emulator handles string lengths in queries differently than the system itself. In the emulator, queries extend the length of the shorter string out to the length of the longer string before making a comparison. Therefore, in the previous example, an emulator query for ‘The Very Best of Benny Goodman’ would return neither ‘The Very Best of Rod Stewart’ nor ‘The Very Best of Benny Goodman.’ Instead, you would have to query on ‘The Very Bes’ to retrieve either object, since that is the query that would match the 12 characters that were actually stored.

---

## *Planning Tables*

Both store and query operations are more efficient if you store metadata attributes that always occur together, or even very commonly occur together, in the same table. For best performance, pay close attention to which metadata attributes occur together in your data, especially if those attributes are used in queries, and group those fields together into the same table.

Conversely, avoid putting fields into the same table that do not often occur together, since doing so wastes space and degrades query performance.

When planning tables, be aware that the maximum number of bytes allowed for any single row in the table is 8080.

You might want to specify as small a length value as possible for each field (column) in the table so that you can fit as many columns as possible in the table and any single row will not exceed the 8080 limit.

[TABLE 2](#) lists the number of bytes that each element in a column consumes. The total amount of space consumed by all the columns in a table cannot exceed 8080 bytes.

**TABLE 2** Number of Bytes Used By Each Element in Schema Table

Element	Space Consumed
System overhead	100 bytes per table.
Column (Field)	Each column (or field) in a table uses 2 bytes per column for overhead, plus the number of bytes in the field. For fields of type string, the number of bytes in the field is equal to twice the length of the field. Fields of type long and double use 8 bytes. For example a field of type string with length 80 consumes 2 bytes for overhead plus 160 bytes for the field length, for a total of 162 bytes. A field of type long consumes 2 bytes for overhead plus 8 bytes for field length, for a total of 10 bytes.

Suppose the fields listed in [TABLE 3](#) commonly occur together on data and will be used together in queries. (Note that three of these fields are in the namespace *mp3* and one is in the namespace *dates*.)

**TABLE 3** Examples of Fields to Be Added to a Table

Field	Type	Length Setting	Bytes Required
<i>mp3</i> .artist	string	256	512
<i>mp3</i> .album	string	256	512
<i>mp3</i> .title	string	512	1024
<i>dates</i> .year	long	NA	8

To maximize query and store performance, you could include each of these fields as columns in the same table, called, for example, *reference*. When planning the reference table, calculate the total number of bytes used by all the columns combined, to make sure it is less than 8080, as follows:

100 (for system overhead) +  
8 (2 per column for column overhead) +  
512 (for mp3.artist) +  
512 (for mp3.album) +  
1024 (for mp3.title) +  
8 (for dates.year)

---

= 2164 bytes total

Since 2164 bytes is less than 8080 bytes, the total combined size of all columns is acceptable.

---

**Note** – In this version, the system does not explicitly prevent you from specifying table columns in the schema definition file that would require more than 8080 bytes if full-sized strings of all fields were stored in a row at the same time. If any object in a table requires more than the 8080 bytes in a single row, some metadata might never be stored in a table and therefore might never be queryable. Be very careful to keep the total maximum size of all columns in a table under 8080 bytes.

---

In addition to the limits on the number of bytes consumed by the columns in a table, keep in mind the following considerations when planning the tables:

- Each table can include no more than 254 fields.
- Each metadata field can be a column in *only* one table.
- When specifying a field as a column in a table, be sure to use the field's fully qualified name (`namespace.field`).
- Tables might contain fields from more than one namespace.
- If you do not specify a field as a column in a table, that field is stored in its own separate table, as the single column in that table.
- If you want to assign a field as a column in a table, you must do so when you first add the field to the schema definition.
- Once you have assigned a field to a table, you cannot move that field to another table, or move it to a different column within the table.
- You cannot add or remove columns from a table once you have defined the table.



## *Planning Indexes*

The system creates indexes on metadata fields to allow those fields to be queried more efficiently. Although you may use the `fsViews` element in the schema definition file to create virtual views for applications that have nothing to do with indexes, you also use the `fsViews` elements to specify the content of the indexes the system creates and maximize query performance.

For each `fsView` you create, the system creates an index of up to 15 fields that includes first the attributes in the `fsView`, followed by the fields specified in the `filename` element of the `fsView`, as long as those fields and attributes all come from the same table.

---

**Note** – Because of the index creation process in this version, each `fsView` now consumes more system resources than in previous versions. For best performance, create only `fsViews` that are required for an application or to define an index to speed up queries.

---

For `fsViews` that you create in order to specify indexes that improve query performance, follow these guidelines:

- Include only fields from the same table. All entries in an index must be from the same table.
- Make sure the `fsView` includes 15 or fewer indexable fields. The system first indexes the attributes listed in the `fsView` and then indexes the fields specified in the `fsView` `filename` element. The total of these two items (attributes and `filename` fields) should not exceed 15.

---

**Note** – In Version 1.0.1, all `fsViews`, even those you create for applications (not indexes), must have 15 or fewer fields from the same table or the `fsView` will not work properly.

---

- Make sure that the index requires less than 1024 bytes of space. [TABLE 4](#) lists the number of bytes required by each element of the index.

**TABLE 4** Number of Bytes Used By Each Element in Index

Element	Space Consumed
System overhead	100 bytes per index.
Field	<p>Each field in an index uses 2 bytes for overhead, plus the number of bytes in the field. For fields of type string, the number of bytes in the field is equal to twice the length of the field. Fields of type long and double use 8 bytes.</p> <p>For example a field of type string with length 80 consumes 2 bytes for overhead plus 160 bytes for the field length, for a total of 162 bytes. A field of type long consumes 2 bytes for overhead plus 8 bytes for field length, for a total of 10 bytes.</p>

Suppose you want to have a query on the fields listed in [TABLE 5](#).

**TABLE 5** Example of Fields to Be Added to a Table

Field	Type	Length Setting	Bytes Required
<code>book.author</code>	string	50	100
<code>book.series</code>	string	50	100
<code>book.title</code>	string	50	100
<code>dates.year</code>	long	NA	8

To maximize query performance, you include each of these fields as columns within the same table, called *books*. To maximize performance even further, you create an `fsView` called, for example, *bookview*, that includes these fields and no others so that an index is created on these fields for querying.

Since there are fewer than 15 fields in the `fsView` and all of the fields are from the same table, the system creates an index that includes all of these fields, as long as the total number of bytes required for the index does not exceed 1024.

Calculate the number of bytes required for the index as follows:

100	(for system overhead) +
8	(2 per column for column overhead) +
100	(for book.author) +
100	(for book.series) +
100	(for book.title) +
8	(for dates.year)

---

= 416 bytes total

Since 416 is less than 1024, the system indexes all of the fields, allowing them to be queried at maximum performance.

If you calculate that the fields in a query cannot be indexed because they require too much space, you might want to reduce the length specified for each field. Alternatively, you might want to define an `fsView` with a smaller set of fields. An index of a subset of fields in the query may still help to speed up query performance.

### *Summary of Best Practices for Optimal Query Performance*

To maximize the performance of queries, follow these steps:

- 1. Plan tables that include metadata fields that commonly occur together in queries in the same tables and separate metadata fields that do not occur together into separate tables. Take into account the space limitations on tables, as described in [“Planning Tables” on page 10](#).**
- 2. Determine which fields should be in which indexes. Each query uses at most one index, so the fields in indexes should match the fields in queries as much as possible. Take into account the limitations described in [“Planning Indexes” on page 13](#).**
- 3. Define `fsViews` whose purpose is to create the indexes you want, as determined in Step 2. Limit the number of `fsViews` you configure to those required for these indexes and those required for applications.**

## New Read-Only Option for fsViews

A new `readonly` attribute for `fsViews` allows you to specify that users accessing data through an application (such as WebDAV) cannot upload new objects into the Sun StorageTek 5800 system and also cannot modify metadata attributes for objects.

---

**Note** – All attributes in the `system` namespace are read only. If you include a system attribute in an `fsView`, that entire `fsView` automatically is read only.

---

## New Keyword to Specify Queryable Fields

With this version, you may use `queryable` as a synonym for `indexable`.

By specifying the `indexable` or `queryable` attribute for a field as `false`, you can exclude that field from the metadata that is indexed and available for queries. You might want to specify a field as `indexable = false` or `queryable = false`, for example, if you will access that field only through the `retrieveMetadata` example application, and never through queries.

## New CLI Option to Recover from Timeout While Activating New Schema

When you issue the `mdconfig -c` command to activate a new schema definition file, you may see the following error message if the system is under heavy load:

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

This message indicates that the new schema definition file has been committed to the system, but not all of the tables may have been created.

In this case, reduce the load on the system if possible, and use the new `-r` option for the `mdconfig` command to finish the table creation:

```
hc $ mdconfig -r
```

When you issue the `mdconfig -r` command, the system finishes creating any tables that were not completed during the `mdconfig -c` operation. Tables that had already been created are not affected. You may have to issue the `mdconfig -r` command several times until all tables are created.

## Clarification of fsView Specification

The *Sun StorageTek 5800 System Administration Guide* might not make it clear that the namespace attribute within an `fsView` is optional. If all of the fields in that `fsView` are from the same namespace, you can specify the namespace attribute so that you do not have to use fully qualified names for the fields in the `fsView`.

You can include fields from different namespaces within the same `fsView`, however. In that case, you should not specify the namespace attribute for the `fsView` and you must specify fully qualified names.

For example, the following `fsView` specification includes fields that are all from the namespace `mp3`. The namespace attribute is specified so the field names are not required to be fully qualified:

```
<fsView name= "byArtist" namespace = "mp3" filename="${title}.${type}">
    <attribute name="artist" />
    <attribute name="album" />
</fsView>
```

In the following `fsView` specification, some of the fields are from the namespace `mp3`, but one field is from the namespace `dates`. In this case, the namespace attribute is not specified and the field names are fully qualified:

```
<fsView name= "byArtist" filename="${mp3.title}.${mp3.type}">
    <attribute name="mp3.artist" />
    <attribute name="mp3.album" />
    <attribute name="dates.year" />
</fsView>
```

---

## Supported Clients

This section lists the clients you can use to access data stored on the Sun StorageTek 5800 system:

- Java™ client application programming interface (API) (Sun Java software developer's kit [SDK] 1.4 or later)
- C client API on the following platforms:
  - Solaris™ 10 Operating system on a SPARC® platform
  - Solaris 10 Operating system on an x86 platform
  - Red Hat Enterprise Linux 4 Update 3
  - Windows XP
- WebDAV protocol

---

## Resolved Issues

This section provides information about functional limitations and bugs that have been resolved in this Version 1.0.1 product release.

### *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.

**Resolution Note** – For this version, the `sysstat` command has been updated to correct this problem. See [“Enhanced System Status Monitoring” on page 5](#) for information about the updated command.

### *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.

## *Shutting Down the Service Node*

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

**Resolution Note** – The shutdown command has a new `-all` option that powers off all components. See [“New CLI Option to Power Off All Components” on page 6](#) for more information.

## *NullPointerException in Alert Code*

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

## *Inconsistent Date Reporting*

**Bug 6408658** – If you attempt to determine when a file or directory was last modified or created, inconsistencies occur 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.

## *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.

**Resolution Note** – You can now set the IP address of the service node. See [“New CLI Option to Set Address for Service Node” on page 6](#) for more information.

## *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 startup. 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 nonoperational.

**Resolution Note** – You can now successfully reboot a node and the cluster when the system is running on the secondary switch. However, you should not perform any networking-related configuration operations on the system, such as those using the CLI commands `ifconfig` and `netcfg`, while operating on the secondary switch.

## *Updating the Schema During a Load*

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

**Resolution Note** – The `mdconfig` command has a new `-r` option that allows you to work around this issue. See [“New CLI Option to Recover from Timeout While Activating New Schema” on page 16](#) for more information.

## *No Node Identifiers With `sysstat --verbose`*

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

## *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.



---

# Known Issues

This section provide 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.

The section contains the following topics:

- [“Functional Limitations” on page 21](#)
- [“Hot-Swap Limitations” on page 21](#)
- [“Installation and Initial Configuration Issues” on page 22](#)
- [“CLI Issues” on page 23](#)
- [“Metadata Issues” on page 25](#)
- [“Client API and SDK Issues” on page 26](#)
- [“General Issues” on page 28](#)

## Functional Limitations

The functional limitations of this version are as follows:

- Scaling to integrated multiple clusters is not enabled.
- There is no backup mechanism currently available for data stored in the Sun StorageTek 5800 system. The only form of disaster protection currently available is for the user to replicate the data using tape backup or WAN replication at the application layer.
- If your primary switch is not functioning and the system is operating on the secondary switch, you should not perform any networking-related configuration operations such as those using the CLI commands `ifconfig` and `netcfg`.

## Hot-Swap Limitations

Do not remove a disk from a live system. When Sun service personnel replace disks, they first power down the system.

In the case of multiple failed disks, Sun service personnel shut down the node for each disk, one at a time, replace the failed disk in that node, and bring the node fully online again before proceeding to the next node.

# 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 two or more missing or faulty disks, that server node might fail during the startup procedure. Contact Sun Support to schedule replacement of the failed drives.

## *Setting Invalid NTP Server IP Address Causes Nodes to Behave Improperly*

**Bug 6466326** – If you set an invalid IP address for an NTP server using the CLI command `netcfg ntp`, the system's nodes might not function properly and might not come up at all.

**Workaround** – Be very careful to enter the correct IP address for an NTP server.

## *Authorized Clients Must Be Reconfigured After Upgrade*

**Bug 6470857** – When you upgrade the Sun StorageTek 5800 system to a new version of software, any authorized clients you have configured are lost and no data traffic can occur until you reconfigure authorized clients, even if previously all clients had been allowed.

**Workaround** – After the upgrade is complete, you must use the `netcfg clients` command to re-enter your authorized clients. If all clients are allowed, issue the command with the `all` option. (`netcfg clients all`).

## *Upgrade Server Must Be Authorized Client*

**Bug 6471001** – If you want to upgrade your software version over the network from an http server, that server must be configured as an authorized client on the Sun StorageTek 5800 system.

**Workaround** – Use the `netcfg clients` command to configure the HTTP server as an authorized client.

## 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 might 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 Causes an Error*

**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 the `sysstat` command before using the `mdconfig` command. The output of the `sysstat` command should specify `Query Available`. If the `sysstat` command reports `Query Unavailable`, do not issue 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 SMTP configuration does not honor the port setting. Regardless of the port value you specify with the `netcfg` command, the system sends mail to your SMTP server using the default port 25.

## *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 `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 Nonoperational*

**Bug 6241900** – If you log in to the system when it is not operational, 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** – If you inadvertently add the `-v` or `--verbose` option to certain CLI commands, such as `hwcfg`, where its use is not supported, the following error is displayed:

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

## *Clarification Needed for Reported Storage Utilization*

**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.

### *System Hangs After reboot --all*

**Bug 6450643** – The system might hang after you issue the `reboot --all` command. This issue is most commonly seen during initial configuration.

**Workaround** – If the system does not reboot after ten minutes, turn off the power to the Sun StorageTek 5800 cell and then turn the power back on to reboot all components.

### *passwd Command Syntax Documented Incorrectly*

**Bug 6466323** – The syntax for the `passwd` command is documented incorrectly in the command help and also in the *Sun StorageTek 5800 System Administration Guide*. The correct syntax is as follows:

```
passwd [-K | --pubkey]
```

Use either the `-K` or the `--pubkey` options to provide a public key to be used for login authorization.

## Metadata Issues

### *Supported Metadata Type*

**Bug 6194366** – The supported metadata type field `double` does not work in views.

### *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.

### *System Allows a Table Beyond Maximum Size*

**Bug 6464058** – The total number of bytes allowed in a schema table is 8080. However, the system does not prevent you from specifying a table in the schema definition file that would require more than 8080 bytes if full-sized strings of all

fields were stored in the table at the same time. If you specify a table that requires more than 8080 bytes, the metadata for some objects might never be stored in a table and therefore might never be queryable.

**Workaround** – Be careful to keep the maximum size of any table that you specify in the schema definition file under 8080 bytes. See [“Planning Tables” on page 10](#) for information about calculating the number of bytes that a table requires.

### *System Does Not Create Index for fsViews With More than 15 Fields*

**Bug 6464061** – If an fsView includes more than 15 fields, the system does not create an index for any of those fields.

**Workaround** – Do not create fsViews with more than 15 fields.

### *Complex Queries Might Time Out*

**Bug 6467475** – During a complex query operation that requires the query to access more than one table, the system might generate the error message “Table memory space exhausted” and the query might time out.

**Workaround** – Follow the guidelines in [“Planning Tables” on page 10](#) and [“Planning Indexes” on page 13](#) to maximize query performance.

## Client API and SDK Issues

This section describes known issues and bugs related to using the client API (application program interface) and SDK (software developer’s kit).

### *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.

### *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). This limitation does not apply to data stored using the Java API.

### *Limitations on Metadata Values Containing Nonprintable Characters*

**Bug 6427141** – The behavior of metadata values that contain nonprintable 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.

### *Performance Might Degrade With Too Many Authorized Clients*

**Bug 6466803** – System performance might degrade significantly if you have a large number of authorized clients configured.

**Workaround** – Limit the number of authorized clients to 4 or fewer.

### *Configuring Authorized Clients Seems to Hang CLI*

**Bug 6468507** – When you configure authorized clients using the `netcfg clients` command, the CLI seems to hang. After a few minutes, however, the configuration is updated and you can enter new CLI commands.

**Workaround** – Wait a few minutes after issuing the `netcfg clients` command to allow the system to update the configuration.

### *SDK Example Scripts Fail on csh*

**Bug 6471588** – You may encounter problems running the SDK example scripts for Unix systems from the C shell (`csh`).

**Workaround** – Run the scripts from a `bash` or `sh` shell only.

### *RetrieveMetadata SDK Example Application Fails With FILE Argument*

**Bug 6472509** – If you specify the `FILE` argument in the `RetrieveMetadata` SDK example application, the application fails.

**Workaround** – Use `RetrieveMetadata` example application without the `FILE` argument.

## General Issues

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

### *Cryptic Error Message When Deleting a Deleted Object*

**Bug 6187582** – If you try to delete the same object twice, the system 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
```

### *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.

### *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 might 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. The source address will include the Administrative VIP configured for the cluster.

### *Unable to Start or Shut Down the Database*

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

### *Node BIOS Issue*

**Bug 6413587** – When the Sun StorageTek 5800 node BIOS is booting up, very rarely, the following two conditions might occur:

- The BIOS fails to boot successfully and becomes nonoperational.
- 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 Email Alert*

**Bug 6402478** – When the switch fails over, it does not send an email alert. The following alert is displayed 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.

**Workaround** – Visually inspect the cluster 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, the node might 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 can result in data loss for recently stored files. Testing of the failure scenarios has not yet resulted in any data loss.

### *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.

### *Full System Cannot Heal Itself After Six Disk Failures*

**Bug 6465815** – If six disks fail on a full cluster (greater than 85% utilized), the system cannot heal itself in a reasonable time.

**Workaround** – In the unlikely event that five disks fail within the quarterly maintenance window, schedule immediate replacement of those disks; do not wait for the next quarterly maintenance.

### *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 shut down or reboot process, for best performance, you might want to wait until the query engine is fully initialized before proceeding.

## *Query Engine Not Recreated Under Heavy Recovery Load*

**Bug 6422739** – If the Sun StorageTek 5800 system experiences an unexpected power loss, the query engine must be recreated and repopulated with user metadata. If, after such a power loss, the query engine remains unavailable for more than four hours following system reboot, it may indicate that the internal recovery load on the system is too high.

**Workaround** – Contact your Sun service representative. A system technician will tune internal recovery to allow the query engine to be recreated. Once the query engine has been recreated, internal healing processes will repopulate the query engine's indexes with user metadata. This healing process may take 12 or more hours to complete. During the healing process, query and WebDAV access will be operational. Not all objects will be accessible, however, until the query engine's indexes are fully repopulated.

## *Intermittent Failures to Store or Retrieve After Reboot*

**Bug 6481942** – In rare cases, when some (but not all) nodes are rebooted after a temporary loss of internal network connectivity in the Sun StorageTek 5800 system, the nodes that were not rebooted may not be able to utilize the disks on the rebooted nodes. In this situation, Sun StorageTek 5800 clients experience intermittent failures to store or retrieve data. For example, a client cannot retrieve an object, but then a few seconds later, the client does successfully retrieve the same object.

**Workaround** – Clear this state using the `reboot` command. (It is not necessary to use the `reboot - - all` option; it is sufficient to just reboot the nodes):

```
hc $ reboot
```

## *Messages Indicating Nodes Leaving Cluster and Rejoining*

**Bug 6473958** – If the Sun StorageTek 5800 system is configured to send email alerts, the system may send a message indicating *The node XXX has left the cluster* and then several seconds later another message indicating *The node XXX has joined the cluster* several times a day during normal operation.

**Workaround** – Messages indicating that a node has rejoined the cluster a few seconds after it left are innocuous. You can ignore them.

## *Reboot or Shutdown During Client Access May Corrupt Query Engine*

**Bug 6481952** – If a reboot or shutdown operation occurs while clients are accessing the Sun StorageTek 5800 system, the query engine database may be corrupted. If the query engine is corrupted, the system may require 12 or more hours after the reboot or shutdown to fully repopulate the database with all user metadata so that query and WebDAV can access all objects.

**Workaround** – If you need to reboot or shut down the system for administrative reasons, stop all clients first and then issue the `reboot` or `shutdown` commands.

## *Reboot or Shutdown with Nonoperational Mail Server May Corrupt Query Engine*

**Bug 6483145** – If the Sun StorageTek 5800 system is configured to use a mail server that is nonoperational, the query engine database may be corrupted during a reboot or shutdown operation. If the query engine is corrupted, the system may require 12 or more hours after the reboot or shutdown to fully repopulate the database with all user metadata so that query and WebDAV can access all objects.

**Workaround** – Make sure your SMTP host setting is for a valid and operational mail server. If there is no valid SMTP host available, before shutting down or rebooting, configure an invalid SMTP host (for example, 0.0.0.1).

## *CLI May Appear to Hang if Mail Server is Nonoperational*

**Bug 6450643** – If the Sun StorageTek 5800 system is configured to use a mail server that is nonoperational, some CLI commands may appear to hang because they are trying to send email alerts and timing out.

**Workaround** – Make sure your SMTP host setting is for a valid and operational mail server. If the mail server is not operational and the CLI commands appear to hang, try waiting three to five minutes for the commands to take effect.

## *Queries May 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 the `maxResults` parameter to a value less than 10000.

## *Disks Unrecovered Count May 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.

---

# Release Documentation

<b>Subject</b>	<b>Title</b>	<b>Part Number</b>
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, two programmer guides are 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>