

Fujitsu

ETERNUS CD10000
Software Version V1.0SP2
April 2015

Release Notice

All rights reserved, including intellectual property rights.

Technical data subject to modifications and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner.

1 General	3
1.1 Ordering	3
1.2 Delivery	3
1.3 Documentation	3
2 Brief Product Description	4
3 Software and Hardware Extensions	5
3.1 Software	5
3.2 VSM (Virtual Storage Manager) GUI Enhancements	6
3.3 Hardware	7
4 Technical information	7
4.1 Resource requirements	7
4.2 Software configuration	7
4.3 Product installation	7
4.3.1 Software Upgrade	7
4.4 Product use	8
4.4.1 Rules Regarding the Number and Mixture of Storage Nodes	8
4.4.2 Backup of the Management Node	8
4.4.3 Replacing defect Hard Disks	8
4.4.4 Web Browser	8
4.5 Discontinued functions (and those to be discontinued)	8
4.6 Incompatibilities	9
4.6.1 Internal Administration Switch (Brocade ICX6610) Minimum Firmware	9
4.7 Restrictions	9
4.7.1 Ceph Cluster creation	9
4.8 Procedure in the event of errors	9
5 Hardware requirements	9

1 General

This Release Notice is a summary of the major extensions, dependencies and operating information with respect to FUJITSU Storage ETERNUS CD10000 S1 hardware and the software V1.0

The release level is that of: April 2015.

A collection of hints and information which are necessary for professional and field service can be found in the appropriate and subsequent Support Bulletins, if applicable.

The names used in this Release Notice may be trademarks whose use by third parties for their own purposes may violate the rights of the owner.

1.1 Ordering

ETERNUS CD10000 can be ordered from your local distributors.

ETERNUS CD10000 is supplied subject to a single payment.

1.2 Delivery

The ETERNUS CD10000 files are supplied on DVD.

Delivery of files for ETERNUS CD10000 on DVD as part of hardware delivery:

- ETERNUS CD10000 1.0SP2 Release : 2015-04-14

1.3 Documentation

The following manuals are available:

- ETERNUS CD10000 V1.0 SP2 User Guide

A documentation CD with all the manuals and the release notice are delivered with the product.

The documentation is also available in the form of online manuals at <http://manuals.ts.fujitsu.com>.

2 Brief Product Description

The Fujitsu Storage ETERNUS CD10000 is a hyperscale storage system that addresses next generation storage requirements imposed by current developments in mobile broadband, social networks, big data, and cloud infrastructures.

The appliance is realized on the basis of a storage node cluster and built with standard IT components (network access, server, local hard drives). This grid of storage nodes presents itself as a single large storage repository. From a hardware perspective data access is granted via 10GbE ethernet interfaces, while the storage nodes among each other communicate via an Infiniband Network.

From a software perspective storage access is provided via the major interfaces

- Via a block device interface
- Via an object storage interface (Access by e.g. librados, S3 or OpenStack Swift)

All storage is managed as a single large entity, which greatly eases capacity and workload management.

A central technology element of the appliance is the intelligent data distribution across the storage nodes. In contrast to other systems the ETERNUS CD10000 manages read and write operations with practically no access to a central instance. A special data distribution algorithm called CRUSH (Controlled Replication under Scalable Hashing) enables all storage clients to exactly calculate where to write and read data within the storage cluster. This pseudo-random but still deterministic data distribution method enables an equal utilization of all involved storage nodes in the cluster. In addition all data can be written in a redundant way (via replicas or via a RAID-like algorithm called erasure coding). Using this method the system can tolerate the loss of hard drives as well as the loss of complete storage nodes without losing data. In addition, lost data copies are recreated in the background and the original redundancy level is restored automatically. Because data copies of a single hard disk drive are distributed across multiple nodes rather than copied to just another hard disk the process of restoring the original redundancy level is very fast because more resources are involved in regenerating the old redundancy level.

New nodes can be added easily while the system is running and in order to avoid hot spots data is reshuffled to create free space on all involved storage nodes, i.e. existing as well as new ones. As a result the cluster continues to benefit from the aggregated bandwidth of all storage nodes while writing new data to it. The Infiniband backend offers the necessary bandwidth for

- creating redundant copies
- the regeneration of the redundancy levels after a data loss,
- and for reshuffling of data after new nodes are added.

As a consequence the 10GB Ethernet frontend network is fully available for reading and writing data.

On the basis of this technology the added value is:

- basically unlimited capacity (at the moment 224 nodes are supported)
- no performance hot spots
- parallel access over object, block and file interfaces by consistent storage management at the same time
- full redundancy for all data in the cluster
- automatic error detection and recovery of lost data at the same time
- possibility to transparently phase out old nodes and add new ones
- decouple management effort from storage growth

|

The ETERNUS CD10000 is a storage system that integrates hardware and software in one appliance including an end to end service concept. The hardware was chosen with care to reach a perfect balance between performance and price. The software allows managing all storage and networking components centrally. This includes monitoring, reporting, and the deployment / update of new nodes.

3 Software and Hardware Extensions

Only the new features and functional enhancements available with ETERNUS CD10000 V1.0 SP02 are described in the following section. The ETERNUS CD10000 User Guide provides detailed information about the features.

3.1 Software

- Web based GUI for management and monitoring of HW components
The web based GUI (basking) is available for managing networks, SNMP and some of the nodes and switch parameters, such as serial numbers. Furthermore information about the nodes and switches and operations such as power off and on are available.
To access the Basking GUI enter the following URL in the address field of your web browser:
`https://hostname/basking`
Login with user `admin` and password `admin`. This default password must be changed at first login.
- Simplification of Ceph cluster creation
A new command
`cd10000 setup cluster`
is available which simplifies the initial cluster creation. Instead of preparing the necessary VSM configuration files and starting the processes on each storage node manually, the new command can be used. The command creates the configuration files with default values. If these values are not sufficient for your needs the manual creation process can be used.
- Additional values stored for each component
Additional parameters are stored or can be updated by the customer for each HW component of the cluster. It is possible to add customer specific location values such as Datacenter, Rack etc.
- Automatic backup of the management node
ETERNUS CD10000 is equipped with one management node. If this management node must be replaced, a backup of the system should be available. For this purpose the tool `fbart` is available. It provides backup and restoration of the management node. Multiple backups according to a predefined schedule are created automatically.
During the "Getting started" process the backup has to be activated manually.
- Collecting log files from switches
The log files generated on the 1Gb and Infiniband switches are automatically transferred to the management node and inserted in the Elasticsearch database where all log files are stored.
The switches are pre-configured accordingly.
- Enhancement to configure SNMP
The command `fsnmpconfig` has been renamed to
`cd10000 setup snmp`
- SNMP monitoring of system disks
SNMP monitoring of the system disks of all node, including the management node, has been implemented. If a defined threshold about the used capacity is exceeded an trap is generated.

- SNMP monitoring of Placement Group (PG) status
SNMP traps are sent when PGs go into one or more of „fatal“ / „warning“ states (degraded, down, inconsistent, incomplete, stale).
- SNMP Trap interval is now configurable.
To change it, edit the file `/etc/snmp/snmp.conf` and add the following line:

```
cd10000TrapInterval 20
```

where “20” is the interval in seconds. Default is 10 seconds. Restart agents afterwards by issuing

```
cd10000agents reload
```

- CD10000 CLI and `vsm_cli` operations in different formats
All CD10000 and `vsm_cli` commands displaying a status of components or executed checks are enhanced to print the output in one of the formats `xml`, `csv`, `json`.
If the option
`-f|--format [json,xml,csv,text]`
is specified the output is printed in selected format. If the option is not specified the output is in ASCII format.
- Automated firmware update of BIOS and iRMC
Updating the firmware of BIOS and iRMC of the storage nodes and the management node is possible with the command `firmware-manager`.
The command is executed on the management node and the firmware is update on the specified nodes one by one.

3.2 VSM (Virtual Storage Manager) GUI Enhancements

- New version V1.0.1
The new version 1.0.1 is integrated and the version is displayed on the Cluster Status page.
- HTTPS for VSM web server
For security reasons HTTPS must be used to access the VSM web server, e.g. <https://managementnode1/dashboard>
- Support erasure coded pools
Under the “Manage Pool” section you can create erasure coded pools. The profile name used for the creation is displayed for such pools.
See the ETERNUS CD10000 User Guide for a detailed description.
- Support for cache tiering
Under the “Manage Pool” section you can create cache tier pools with the needed roles. The pools tier role (cache or storage) and the storage group to which it is paired is displayed for existing cache tier pools.
See the ETERNUS CD10000 User Guide for a detailed description.
- Replicated pools
A pool can be configured in such a way that the primary data is going to one pool and the replicated data to another pool.
- Setting pool quotas
For all types of pools quotas can be specified during pool creation.

3.3 Hardware

none

4 Technical information

4.1 Resource requirements

ETERNUS CD10000 provides all required resources on its own hardware. Connection cables for the connection to the customer network have to be present or have to be ordered separately by the customer.

4.2 Software configuration

The ETERNUS CD10000 V1.0 User Guide provides basic configuration instructions for the ETERNUS CD10000 system.

4.3 Product installation

4.3.1 Software Upgrade

All ETERNUS CD10000 running the software version V1.0SP01 can be upgraded to the software V1.0SP02.

Proceed as follows to upgrade to the new software version:

- Upload the provided RPM packages to the repository:
`update-manager upload -d ABSOLUTE_PATH_TO_RPM_DIR`
- Update the cluster manifest partials using Update Manager (manifest partials contain defined cluster state)
`update-manager update_partials`
- Publish repository
`update-manager publish`
- Merge the custom/changed manifests with Updated (this step can be omitted when there are no custom changes or no custom manifests).
- Upgrade the management node only using Update Manager. By applied manifests, management node will be rebooted. After rebooting wait about five minutes before going to next step.
`update-manager upgrade -n pmgmt`
- To update the VSM related files and database execute
`vsm-update`
- Upgrade the cluster using Update Manager (this function upgrades only storage nodes).
`update-manager upgrade`
- Upload the provided iso installation image for future installation of storage nodes:
`update-manager upload -d ABSOLUTE_PATH_TO_MOUNTED_ISO -iso`
`mount -o loop path_to_iso ABSOLUTE_PATH_TO_MOUNTED_ISO`
 where
`ABSOLUTE_PATH_TO_MOUNTED_ISO=/tmp/iso`
`ABSOLUTE_PATH_TO_RPM_DIR=/tmp/iso/Packages`

All nodes are upgrade one by one. If the upgrade of a node fails the process exits with an error message.

For a more detailed description of the upgrade process see ETERNUS CD10000 User Guide, chapter 9.

4.4 Product use

4.4.1 Rules Regarding the Number and Mixture of Storage Nodes

Regarding the number of storage nodes the following rules apply

- the minimum number of storage nodes is four
- the maximum number of storage nodes is 224

Regarding the number and mixture of storage node types the following rules apply:

- all storage nodes must be of the same type (basic, performance, or capacity storage nodes)
- Or in case capacity storage nodes are mixed with basic or performance nodes, the number of storage nodes per chosen type must be equal to the number of replicas
 - for example a cluster consisting of storage nodes of type basic and capacity must at least consist of 2 x basic and 3 x capacity storage nodes
 - for example a cluster consisting of storage nodes of all three types must at least consist of 2 x basic, 2 x performance, and 3 x capacity storage nodes
- The number of replicas has influence on the mixture of storage node types. For example a cluster with 2 basic storage nodes and 2 capacity storage nodes, although it is an allowed configuration, doesn't make sense in case three replicas are planned.

In general smaller configurations below 8 storage nodes are not recommended and if less than 8 storage nodes are necessary the usage scenario and potential implications on operating the cluster should be considered very carefully. In any case please note that configurations with less than 4 storage nodes are not supported under any condition.

4.4.2 Backup of the Management Node

ETERNUS CD10000 is equipped with one management node. If this management node must be replaced, a backup of the system should be available. For this purpose the Fujitsu Backup and Restoration Tool (**fbart**) is created. It provides backup and restoration of the management node of the Fujitsu ETERNUS CD10000. Multiple backups according to a predefined schedule are created automatically.

The backups include all the configuration data of the management node which are needed for a complete restoration.

During the "Getting started" process the backup has to be activated manually.

4.4.3 Replacing defect Hard Disks

If you have to substitute a new disk for a failed/retired disk keeping the overall storage cluster configuration intact this must be done with the `vsm_cli replace` disk commands (see User Guide for detailed information).

The web-based GUI VSM does not support this process!

4.4.4 Web Browser

VSM requires Google Chrome or Internet Explorer 10, or another modern Web-browser (tested with Google Chrome and Internet Explorer 10 with compatibility mode turned off).

4.5 Discontinued functions (and those to be discontinued)

None

4.6 Incompatibilities

4.6.1 Internal Administration Switch (Brocade ICX6610) Minimum Firmware

For proper working the Internal Administration Switch (Brocade ICX6610) requires to have at least the firmware version 08.0.10d.

In addition, any firmware from the series 07.#.### and 09.#.## is not supported. In this case the switch's firmware has to be upgrade or downgrade to fulfil the requirements before upgrading the software.

4.7 Restrictions

4.7.1 Ceph Cluster creation

The ETERNUS CD10000 storage cluster provides the web-based user interface Virtual Storage Manager for Ceph (VSM) and the command line `vsm_cli` to manage and monitor the storage cluster. You can perform many configuration, administration, and monitoring functions with VSM.

Ceph commands are installed on the system and can be used for some purposes.

It is highly recommended to create the Ceph cluster with VSM or `vsm_cli` and not with Ceph commands. If you are using Ceph to create the cluster it is not possible to manage or monitor the cluster with VSM afterwards.

4.8 Procedure in the event of errors

If an error occurs, first consult the troubleshooting chapter in the ETERNUS CD10000 User Guide.

The following error documents are needed for diagnostics:

A detailed description of the error condition, indicating whether and how the error can be reproduced.

The ETERNUS CD10000 serial number can be obtained with the following command on the management node:

```
# cd10000 ip cluster show
```

Cluster SN	Management node SN	Administration net
YM2D0010034	YLNT004583	192.168.20.0/23

The entry in column "Cluster SN" is the relevant serial number.

The error must be reported to the appropriate service provider.

An AIS Connect access is essential for diagnostics.

5 Hardware requirements

None