

# WHITE PAPER

## FUJITSU ETERNUS DX Quantum StorNext File System Certification

RELIABLE DISK STORAGE SYSTEMS IN STORAGE AREA NETWORKS WITH HIGH PERFORMANCE  
FILE CREATION, STORING AND SHARING



**1.0 SUMMARY**

StorNext is a SAN file system with high performance file creation, storing and sharing. With StorNext data management software, you get high-speed content sharing combined with cost-effective data archiving. StorNext offers data sharing and retention in a single solution, so you don't have to piece together multiple products that may not integrate well. Even in heterogeneous environments, all data is easily accessible to all hosts.

This document describes the interaction from StorNext file system with Fujitsu ETERNUS DX storage systems.

ETERNUS DX disk storage systems are highly reliable and secure Data Safes. For small and medium-sized businesses, ETERNUS DX60, DX80, and DX90 provide fast and highly available data. For mid-range and enterprise customers who need secure and flexible integration into the data center, Fujitsu offers the ETERNUS DX400 series and ETERNUS DX8000 series. ETERNUS DX disk storage systems use the same software stack and hardware functions across the entire range to allow ease-of-use and greater efficiency for users.

For more information's about ETERNUS DX storage system see <http://ts.fujitsu.com/eternus>.

**2.0 TEST ENVIRONMENT**

The test environment was build with ETERNUS DX8400, 2 Brocade switches for SAN connectivity and Primergy Server for high scalable I/O performance.

**StorNext MDC (108): Primergy RX300S5**

- Red Hat Enterprise Linux Server release 5.3 (Tikanga)  
2.6.18-128.el5 #1 SMP Wed Dec 17 11:41:38 EST 2008 x86\_64 GNU/Linux
- 16 x Intel(R) Xeon(R) CPU X5550 @ 2.67GHz, 8192 KB Cache
- 2 x Emulex LPe12002 8Gb 2-port PCIe Fibre Channel Adapter LP,  
FW 1.10A5, Driver 8.2.0.33.3p
- 24 GB MEM
- StorNext file system Version RedHat50AS-3.5.1.10866

**StorNext SAN client (107): Primergy RX300S5**

- Red Hat Enterprise Linux Server release 5.3 (Tikanga)  
2.6.18-128.el5 #1 SMP Wed Dec 17 11:41:38 EST 2008 x86\_64 GNU/Linux
- 16 x Intel(R) Xeon(R) CPU X5550 @ 2.67GHz, 8192 KB Cache
- 2 x Emulex LPe12002 8Gb 2-port PCIe Fibre Channel Adapter LP,  
FW 1.10A5, Driver 8.2.0.33.3p
- 24 GB MEM
- StorNext SAN client RedHat50AS\_2618-3.5.1.10866

**StorNext SAN client (106): Primergy RX300S5**

- Red Hat Enterprise Linux Server release 5.3 (Tikanga)  
2.6.18-128.el5 #1 SMP Wed Dec 17 11:41:38 EST 2008 x86\_64 GNU/Linux
- 16 x Intel(R) Xeon(R) CPU X5550 @ 2.67GHz, 8192 KB Cache
- 2 x Emulex LPe12002 8Gb 2-port PCIe Fibre Channel Adapter LP,  
FW 1.10A5, Driver 8.2.0.33.3p
- 24 GB MEM
- StorNext SAN client RedHat50AS\_2618-3.5.1.10866

**StorNext SAN client (105): Primergy RX300S5**

- Red Hat Enterprise Linux Server release 5.3 (Tikanga)  
2.6.9-78.ELsmp #1 SMP Wed Jul 9 15:46:26 EDT 2008 x86\_64 GNU/Linux
- 16 x Intel(R) Xeon(R) CPU X5550 @ 2.67GHz, 8192 KB Cache
- 2 x Emulex LPe12002 8Gb 2-port PCIe Fibre Channel Adapter LP,  
FW 1.10A5, Driver 8.2.0.33.3p
- 24 GB MEM
- StorNext SAN client RedHat50AS\_2618-3.5.1.10866

**StorNext SAN client (104): Primergy RX300S5**

- Red Hat Enterprise Linux AS release 4 (Nahant Update 7)  
2.6.18-128.el5 #1 SMP Wed Dec 17 11:41:38 EST 2008 x86\_64 GNU/Linux
- 8 x Intel(R) Xeon(R) CPU X5550 @ 2.67GHz, 8192 KB Cache
- 2 x Emulex LPe12002 8Gb 2-port PCIe Fibre Channel Adapter LP,  
FW 1.10A5, Driver 8.0.16.40
- 24 GB MEM
- StorNext SAN client RedHat40AS\_269-3.5.1.10866

**Storage: ETERNUS DX8400**

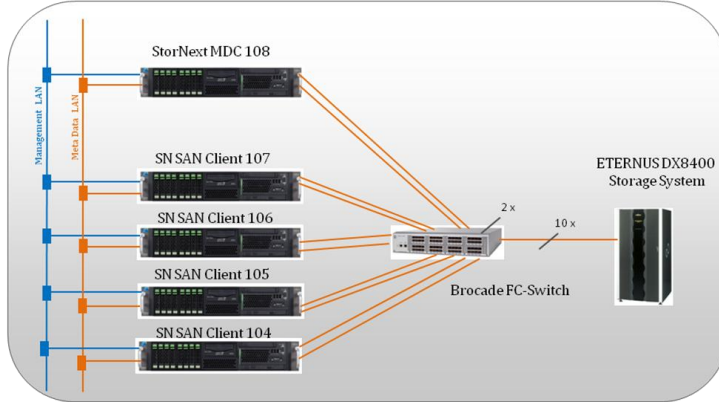
- Firmware: V20L43
  - Controller: 4 x CM, each with 64 GB Memory,
  - Channel Adapter (CA), 4 x 4port FC 8Gb/s (one CA per CM)
- Raid Group:
- RAID1+0 (8 x 146GB FC) - snfs1
  - RAID5 (8 x 146GB FC) - snfs2
  - RAID1+0, 146GB SSD - snfs3
  - 8 x RAID1 (2 x 146GB) for SNFS stripe - snfs3

**Fibre Channel Switch: 2 x Brocade DCX 4S**

- Firmware: 6.3.1a

### 3.0 CONFIGURATION OVERVIEW

For communication of StorNext clients with the Meta Data Controller (MDC), a dedicated network is established.



### 4.0 INSTALLATION

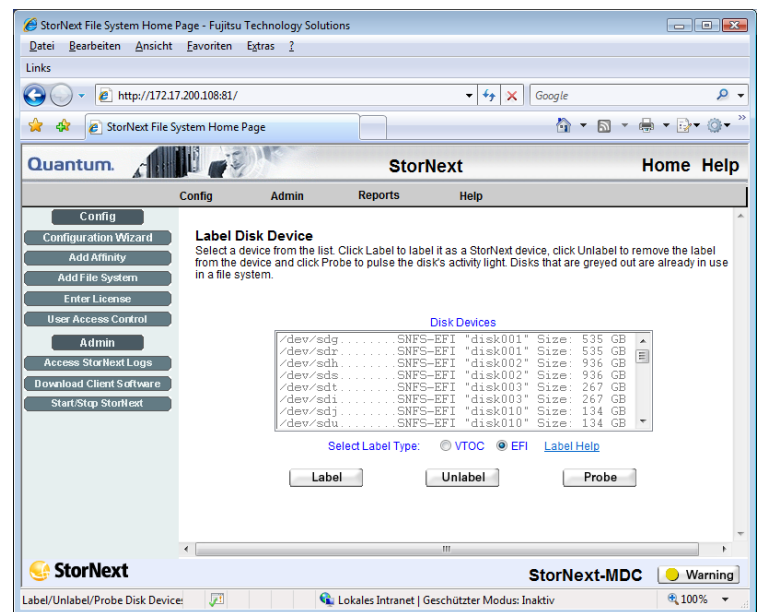
#### 4.1 MDC

Before installation of StorNext files system, the configured LUNs on ETERNUS DX8400 are checked from MDC as follow:

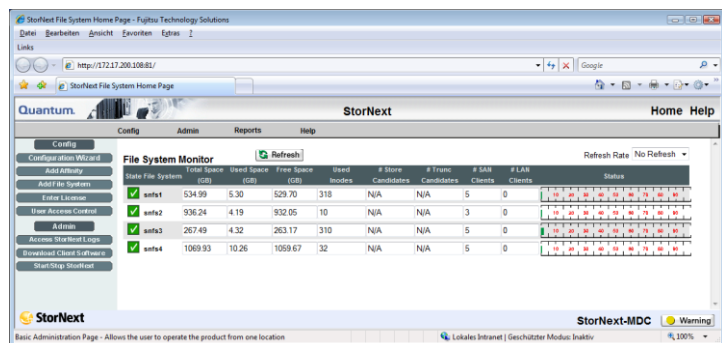
```
[root@StorNext-MDC RedHat50AS_26x86_64]# sginfo -al
/dev/sg0 [scsi0 ch=0 id=10 lun=0]
/dev/sg1 [=/dev/sda scsi0 ch=2 id=0 lun=0]
/dev/sg2 [=/dev/sdb scsi0 ch=2 id=1 lun=0]
/dev/sg3 [=/dev/sdc scsi0 ch=2 id=2 lun=0]
/dev/sg4 [=/dev/sdd scsi0 ch=2 id=3 lun=0]
/dev/sg5 [=/dev/sde scsi0 ch=2 id=4 lun=0]
/dev/sg6 [=/dev/sdf scsi0 ch=2 id=5 lun=0]
/dev/sg7 [=/dev/scd0 scsi1 ch=0 id=0 lun=0]
/dev/sg8 [=/dev/sgd scsi9 ch=0 id=0 lun=0]
/dev/sg9 [=/dev/sdh scsi9 ch=0 id=0 lun=1]
/dev/sg10 [=/dev/sdi scsi9 ch=0 id=0 lun=2]
/dev/sg11 [=/dev/sdj scsi9 ch=0 id=0 lun=10]
/dev/sg12 [=/dev/sdk scsi9 ch=0 id=0 lun=11]
/dev/sg13 [=/dev/sdl scsi9 ch=0 id=0 lun=12]
/dev/sg14 [=/dev/sdm scsi9 ch=0 id=0 lun=13]
/dev/sg15 [=/dev/sdn scsi9 ch=0 id=0 lun=14]
/dev/sg16 [=/dev/sdo scsi9 ch=0 id=0 lun=15]
/dev/sg17 [=/dev/sdp scsi9 ch=0 id=0 lun=16]
/dev/sg18 [=/dev/sdq scsi9 ch=0 id=0 lun=17]
/dev/sg19 [=/dev/sdr scsi11 ch=0 id=0 lun=0]
/dev/sg20 [=/dev/sds scsi11 ch=0 id=0 lun=1]
/dev/sg21 [=/dev/sdt scsi11 ch=0 id=0 lun=2]
/dev/sg22 [=/dev/sdu scsi11 ch=0 id=0 lun=10]
/dev/sg23 [=/dev/sdv scsi11 ch=0 id=0 lun=11]
/dev/sg24 [=/dev/sdw scsi11 ch=0 id=0 lun=12]
/dev/sg25 [=/dev/sdx scsi11 ch=0 id=0 lun=13]
/dev/sg26 [=/dev/sdy scsi11 ch=0 id=0 lun=14]
/dev/sg27 [=/dev/sdz scsi11 ch=0 id=0 lun=15]
/dev/sg28 [=/dev/sdaa scsi11 ch=0 id=0 lun=16]
/dev/sg29 [=/dev/sdab scsi11 ch=0 id=0 lun=17]
```

```
[root@StorNext-MDC RedHat50AS_26x86_64]# sg_map -i
/dev/sg0 FSC SAS Expander 2.01
/dev/sg1 /dev/sda LSI MegaRAID SAS RMB 1.40
/dev/sg2 /dev/sdb LSI MegaRAID SAS RMB 1.40
/dev/sg3 /dev/sdc LSI MegaRAID SAS RMB 1.40
/dev/sg4 /dev/sdd LSI MegaRAID SAS RMB 1.40
/dev/sg5 /dev/sde LSI MegaRAID SAS RMB 1.40
/dev/sg6 /dev/sdf LSI MegaRAID SAS RMB 1.40
/dev/sg7 /dev/scd0 HL-DT-ST DVDROM GT20N CW02
/dev/sg8 /dev/sgd FUJITSU ETERNUS_DX8000 0000
/dev/sg9 /dev/sdh FUJITSU ETERNUS_DX8000 0000
/dev/sg10 /dev/sdi FUJITSU ETERNUS_DX8000 0000
/dev/sg11 /dev/sdj FUJITSU ETERNUS_DX8000 0000
/dev/sg12 /dev/sdk FUJITSU ETERNUS_DX8000 0000
/dev/sg13 /dev/sdl FUJITSU ETERNUS_DX8000 0000
/dev/sg14 /dev/sdm FUJITSU ETERNUS_DX8000 0000
/dev/sg15 /dev/sdn FUJITSU ETERNUS_DX8000 0000
/dev/sg16 /dev/sdo FUJITSU ETERNUS_DX8000 0000
/dev/sg17 /dev/sdp FUJITSU ETERNUS_DX8000 0000
/dev/sg18 /dev/sdq FUJITSU ETERNUS_DX8000 0000
/dev/sg19 /dev/sdr FUJITSU ETERNUS_DX8000 0000
/dev/sg20 /dev/sds FUJITSU ETERNUS_DX8000 0000
/dev/sg21 /dev/sdt FUJITSU ETERNUS_DX8000 0000
/dev/sg22 /dev/sdu FUJITSU ETERNUS_DX8000 0000
/dev/sg23 /dev/sdv FUJITSU ETERNUS_DX8000 0000
/dev/sg24 /dev/sdw FUJITSU ETERNUS_DX8000 0000
/dev/sg25 /dev/sdx FUJITSU ETERNUS_DX8000 0000
/dev/sg26 /dev/sdy FUJITSU ETERNUS_DX8000 0000
/dev/sg27 /dev/sdz FUJITSU ETERNUS_DX8000 0000
/dev/sg28 /dev/sdaa FUJITSU ETERNUS_DX8000 0000
/dev/sg29 /dev/sdab FUJITSU ETERNUS_DX8000 0000
```

The installation of StorNext file system was performed on Meta Data Controller (MDC) with the install.snfs script from RedHatAS\_26\_x86\_64 folder on CD. The StorNext application directory was installed on a separate file system on separate local disk. The LUNs are labeled with type EFI.



With the Web GUI Configuration Wizard the license was entered and 4 SN files system were created.



Note: The Linux device-mapper multipath must be disabled. Depending on the storage system (active/active), StorNext may be able to handle multiple paths.

```
[root@StorNext-MDC tools]# mount
/dev/cvfsctl1_snfs1 on /stornext/snfs1 type cvfs (rw,sparse=yes)
/dev/cvfsctl3_snfs3 on /stornext/snfs3 type cvfs (rw,sparse=yes)
/dev/cvfsctl4_snfs4 on /stornext/snfs4 type cvfs (rw,sparse=yes)
/dev/cvfsctl2_snfs2 on /stornext/snfs2 type cvfs (rw,sparse=yes)
```

#### 4.2 STORNEXT SAN CLIENT

On StorNext SAN client make sure that all configured LUNs in ETERNUS are seen from all clients. The client software was installed with rpm command “rpm -ivh snfs-client-RedHat50AS\_2618-3.5.1.10866.x86\_64.rpm”. After installation of this package the cvfs start script was enabled with command “chkconfig --level 3456 cvfs on”. After creation of /usr/cvfs/config/fsnameservers with IP Address from MDC, use the following command to mount the StorNext file system, e.g. # mount -t cvfs snfs2 /stornext/snfs2. This procedure should be performed on all StorNext SAN clients. You can expand the /etc/fstab with the StorNext file systems to be available after reboot as follow:

```
[root@StorNext107 tools]# cat /etc/fstab
/dev/VolGroup00/LogVol00 / ext3 defaults 1 1
LABEL=/boot /boot ext3 defaults 1 2
tmpfs /dev/shm tmpfs defaults 0 0
devpts /dev/pts devpts gid=5,mode=620 0 0
sysfs /sys sysfs defaults 0 0
proc /proc proc defaults 0 0
/dev/VolGroup01/LogVol01 swap swap defaults 0 0
/dev/VolGroup00/LogVol01 swap swap defaults 0 0
#-----
snfs1 /stornext/snfs1 cvfs verbose=yes 0 0
snfs2 /stornext/snfs2 cvfs verbose=yes 0 0
snfs3 /stornext/snfs3 cvfs verbose=yes 0 0
snfs4 /stornext/snfs4 cvfs verbose=yes 0 0
```

#### 5.0 TESTS

##### 5.1 SWIFT-AD

Swift-AD is a test program that creates test files from 4K byte to 512M byte under specified directory and executing write/read/compare of files. The following script was used to test the file system on each StorNext SAN client. This script launches 400 processes which creates the test files. Each client used an own directory on the shared file system.

```
#!/bin/bash
i=0
for i in `seq 0 99`
do
    ./rdflwr -d /stornext/snfs1/107/"$i" -r wr -c 100 -o
log_107/snfs1_swift.log"$i" -f rnd -e3 -l 100 &
    ./rdflwr -d /stornext/snfs2/107/"$i" -r wr -c 100 -o
log_107/snfs2_swift.log"$i" -f rnd -e3 -l 100 &
    ./rdflwr -d /stornext/snfs3/107/"$i" -r wr -c 100 -o
log_107/snfs3_swift.log"$i" -f rnd -e3 -l 100 &
    ./rdflwr -d /stornext/snfs4/107/"$i" -r wr -c 100 -o
log_107/snfs4_swift.log"$i" -f rnd -e3 -l 100 &
done
```

##### 5.2 IOZONE

The iozone performance test program was primary used to generate IO stress to the file system. On one StorNext SAN client connected with 2 paths to ETERNUS storage a write performance of about 880 MB/s and read Performance of 1,23 GB/s was reached. The values are oriented performance results and can be tuned in StorNext files system (amount of stripe LUNs, stripe size, special Meta Data devices etc.) as well as in LINUX Operating system and the amount of data paths to ETERNUS Storage system. The tuning of this parameter mainly depends on the customer application facts, e.g. streaming or random file system access or ratio of read and write IO.

```
ALLFILES=""
rm iozone.out
for i in `seq 1 20`
do
    ALLFILES="$ALLFILES /stornext/snfs4/iozone/file_$i"
    echo $i
    /opt/iozone/bin/iozone -R -i0 -i1 -t"$i" -s 1000m -r
1024k -F $ALLFILES >>iozone.out 2>&1
done
```

##### 5.3 RAS

For reliability tests the swift-AD test program was used for write/read/compare data. Swift-AD was performed on all SAN clients. Parallel a shell script was launched from a control host which performs a “portdisable / portenable” on the fibre channel switch for all ports within the StorNext configuration sequentially.

```
while :
do
    cmd "./switch_admin -s 172.17.206.161 -c switchshow;
sleep 10"
    cmd "./switch_admin -s 172.17.206.161 -c
\"portdisable 7/0\"; sleep 180"
    cmd "./switch_admin -s 172.17.206.161 -c \"portenable
7/0\"; sleep 180"
    cmd "./switch_admin -s 172.17.206.161 -c
\"portdisable 7/1\"; sleep 180"
    cmd "./switch_admin -s 172.17.206.161 -c \"portenable
7/1\" ; sleep 180"
    :
    :
    cmd "./switch_admin -s 172.17.206.164 -c switchshow;
sleep 10"
    cmd "./switch_admin -s 172.17.206.164 -c
\"portdisable 7/20\"; sleep 180"
    cmd "./switch_admin -s 172.17.206.164 -c \"portenable
7/21\"; sleep 180"
    cmd "./switch_admin -s 172.17.206.164 -c
\"portdisable 7/22\"; sleep 180"
    cmd "./switch_admin -s 172.17.206.164 -c \"portenable
7/23\"; sleep 180"
done
```

After 16 hours tests with sequential portisable / portenable, the whole fibre channel switch was disabled with “switchdisable” command. Finally the log files of swift-AD were investigated for file corruption. All tests passed successful.

#### 6.0 FINAL RESULT

The result of the tests shows that ETERNUS DX storage systems in conjunction with Quantum’s StorNext file system are a well suitable team for best performance and scalability. The tests described above in the document passed without any problems.

---

#### CONTACT

FUJITSU TECHNOLOGY SOLUTIONS  
Address: Mies-van-der-Rohe-Strasse 8,  
D-80807 Munich, Germany  
E-mail: [storage-pm@ts.fujitsu.com](mailto:storage-pm@ts.fujitsu.com)  
Website: [ts.fujitsu.com/eternus](http://ts.fujitsu.com/eternus)  
2010-05-12 EMEA EN

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. For further information see [ts.fujitsu.com/terms\\_of\\_use.html](http://ts.fujitsu.com/terms_of_use.html)  
Copyright © Fujitsu Technology Solutions GmbH 2010