

# Open-E Certified Engineer Training

Month 00, 2000

Todd Maxwell  
Director of Technical Sales

OPEN-E RECOMMENDS TO INSTALL DSS V7 ON A SATA DOM, IDE DOM AND OR A SSD DRIVE AS THEY HAVE A HIGHER MTBF.

WE HAVE FOUND OUT THAT MANY OF THE USB FLASH STICKS HAVE LOWER LIFE OF 4 YEARS AND UNDER EVEN WITH SOME THAT HAVE THE WEAR LEVELING SUPPORT THEY CAN STILL FAIL.

## ■ RAID

- Create a 2GB Volume from the RAID
- Create the RAID Array
- Make the 2GB volume the 1st boot order
- Other capacity amount will be available to be used for the DSS V7 Volume manger

- **Virtual Machine**
  - Prepare virtual storage with minimum of 2GB
  - Configure CPU, Memory and NICs for the VM
  - Have DSS V7 ISO image available to be installed as a VM (typically via NFS mounts or others)
  
- **USB Flash Stick**
  - Recommend 2GB or more (please remove all other partitions)
  - Use Wear Leveling (If possible) USB Flash Sticks or DOMs “Disk on Modules”
  - Unpack the DSS V7 ZIP version
  - Format via Fat 16 or 32
  - Use the “bootinst.exe”
  - Make USB the 1st boot order

## **Lab Work**

15 minutes to install Open-E DSS V7  
as a Virtual Machine

## OTHER MEDIUMS THAT CAN BE USED

- Can be used with **Hard Drives**, **FC Targets** (4Gb and 8Gb FC HBA's), **Onboard RAID** controllers with RAID 1, **SATA DOMs**

## ADDITIONAL INFORMATION

- **Defaults to 64 bit mode**
- Can run **Memtest** or other troubleshooting methods
- If carving out 1 partition for the installation of the DSS V7 you will **not be able** to use the rest of the capacity
- **New RAID controllers** that require new drivers that are not in the DSS V7 might **not install or boot** properly

## ADDITIONAL INFORMATION

- Entering the **Product key** in the GUI requires a reboot but **not an Online or Offline Activation**. You can edit the prod.key file to reduce re-boots
- New Volume Group must be created (**5GB is the minimum**)
- Run RAID Array/disk **Read and Write speeds** speed tests before deploying the system along **NIC speeds** as well with small updates – providing during the class for future use
- If you change or add a **CPU, NICs or memory** there is no reactivation only with other hardware components

## ADDITIONAL INFORMATION

- **Existing 32bit** volumes needing to switch to 64bit **must backup the data** and verify before deleting and creating new **64bit Volume Group**
- Once system has been fully configured with Logical Volumes **save your configuration and settings** from the GUI in Maint. > Misc, you can restore them from the settigns.cnf if need be in the future
- The **data on the volumes will always be available** if in the event of a **power hit** or other (except if for RAID issues), just install DSS V7 on another **Boot Media (CD-ROM, USB or other)**
- At the end of the 60 Trial version **performance drops** from 1GbE to 10/100MB but data is still available

## WHAT IS SNAPSHOT?

- Allows you to create a new block device which presents an exact copy of a logical volume, frozen at some **point in time**, so overall it is based on the **Logical Volume Manager**
- This provides access to the data existing on the volume at the snapshot start time
- The original copy of the data continues to be available to the users without interruption, while the snapshot copy is used to perform other functions on the data for **Backup and Data Replication** applications or user access point in time data to access accidentally **deleted or modified files** for FC, iSCSI and FC Volumes
- A Snapshot created in one **Volume Group cannot be used in a different Volume Group**
- Please keep **assigned Snapshots separate for there tasks**, try to keep a designated Snapshot for each task or create additional Snapshots

## BASIC EXPLANATION OF SNAPSHOT AND KNOWN CONDITIONS

- Concurrent snapshots can be performed but we recommend to have no more than **10 per Logical Volume** and no more than **20 actively on at the same time per system**
- Deleted data is claimed as free space in a “live” volume mount, but in reality the deleted data is still available in the snapshot mount
- **Starting or stopping a snapshot is very fast**, this only takes a few seconds even for large amount of data
- **Writing speed decreases** with growing number of active snapshots (because of copy-on-write)
- The size of the reserved space for snapshot depends on the amount of changed data while the snapshot is active.
- Daily scheduled snapshots will need less reserved space then weekly scheduled snapshot. The size of the
- Snapshot should be **2 or 3 times the size of the expected changes for the data**

## HOW TO ACCESS SNAPSHOTS FOR NAS VOLUMES

- Go to the **CONFIGURATION -> NAS settings** menu and select the network protocol on which the snapshots will be accessible. You can activate access to snapshots on the following protocols:
  - NFS, SMB(Network neighborhood), FTP, AFP.
- Create a new share that will be assigned to the activated snapshot,
- Go to the **CONFIGURATION -> NAS resources** menu,
- Within the **Create new share** function:
  - Enter share name, use the **Specified path** option and select the snapshot that you want to have access to
- Click **Apply** to create a share, now you can start to explore your share(snapshot) using the specified network protocol

## HOW TO ACCESS SNAPSHOTS FOR ISCSI VOLUMES (SIMILAR TO FC VOLUMES)

- Go to the **CONFIGURATION** -> **iSCSI target manager** -> **Targets** -> **[target\_name]** menu,
- Enter the **Target volume manger** function and click the **Add** button on the right side of the. A new LUN will be added to the target **CONFIGURATION** -> **FC target manager** -> **Groups** -> **[Group\_name]** -> **Function: Add group volumes**

## THE BUILT-IN BACKUP FUNCTION PROVIDES RELIABLE BACKUPS FOR NAS VOLUMES – THIS IS HOW IT WORKS:

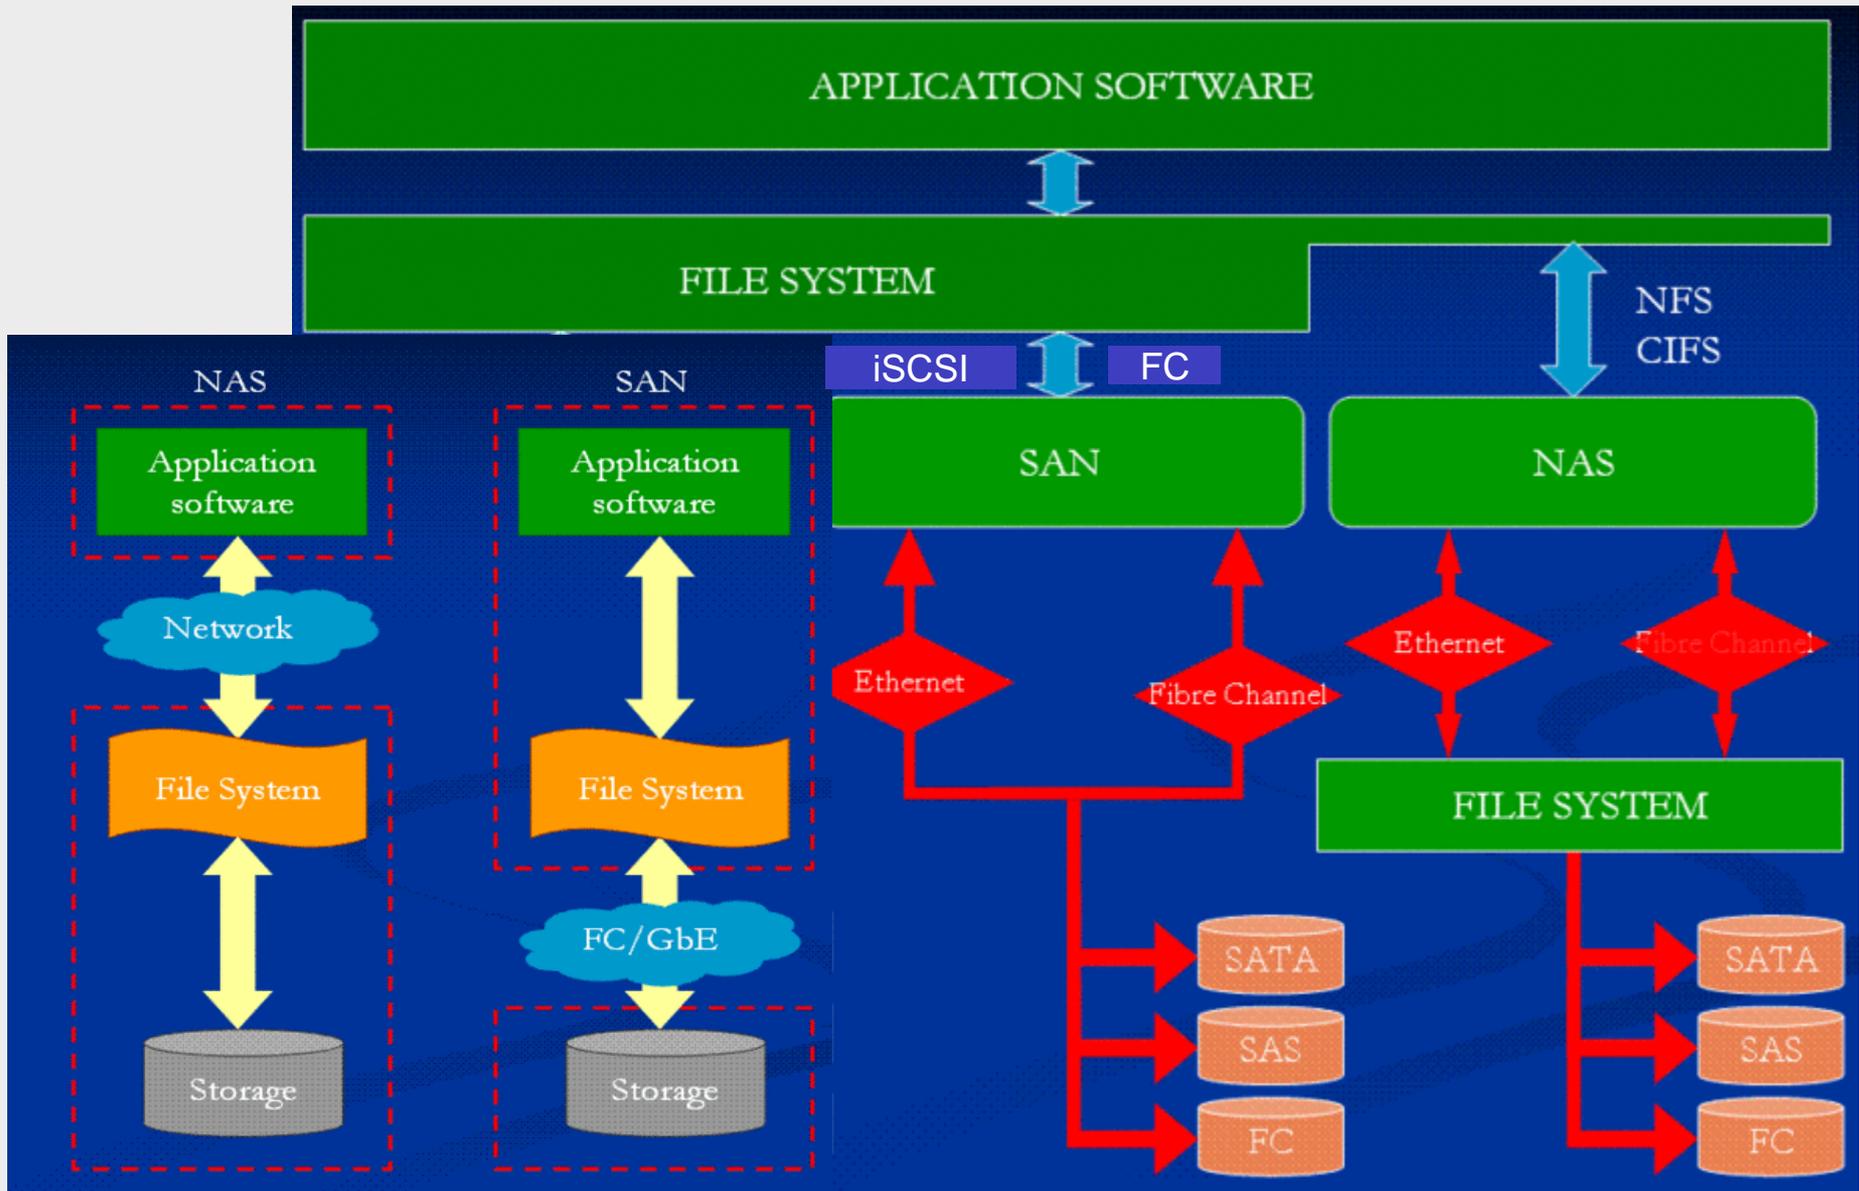
- Backup **only is designed for NAS Logical Volumes** (FC and iSCSI volumes are not supported)
- The Backup function **needs a Snapshot** in order to create the Task
- The Shares are backed up from the NAS Logical Volumes
- **Backup has a higher priority then Data Replication** if scheduled at the same time
- Supported backup types Full (All data on every backup task), **Differential** (new data from last Full Backup) and **Incremental** (only backs up new data)
- Can perform compression (**depending on hardware and file structure**)
- Must have its own **Backup database without existing files in the NAS Share's**
- NAS **WORM volumes cant be backed up**
- Usage of **Pools** used for grouping Tapes or Virtual Taps
- Task scheduled to set time values for your backups
- Restore feature for **backups to restore data** to any NAS Share

## MEDIUMS THAT YOU CAN BACKUP TO

- Backup to a **Tape Library & Drive, RAID controller, RAID array**
- Backup to a Dynamic Volume with - **USB drive, SATA drive, ATA drive**
- Backup to **iSCSI Target Volume, FC Target Volume**

# Installing Open-E DSS V7

*open-e*



## FC – FIBRE CHANNEL

- Fibre Channel Protocol (FCP) is a transport protocol (similar to TCP used in IP networks) which predominantly **transports SCSI commands over Fibre Channel networks**
- iSCSI and FC volumes emulate a raw SCSI drive; they are partitioned and formatted with a regular (non-cluster) file system like NTFS, EXT3, XFS, etc., must be used by a host exclusively. I/O of an FC target is block-based (as opposed to file-based), which means that changes made by one person will not be seen by another person working on the same target/volume
- DSS V7 offer's **both Initiator and Target** (Qlogic only) feature
- **MPIO** in Target mode only is supported

## FC – FIBRE CHANNEL

- FC Volumes can be **initialized** like iSCSI volumes (Writes the volume with all zero's) and blocksize's can be set (VMware is 512, MS 4096)
- **Snapshot and Volume Replication** can be used with FC volumes
- **Public Groups** can be created and is visible to any FC HBA, only 1 Public Group is created
- Privet groups are not publicly seen and can only be visible by the FC HBA's **WWN** that you have assigned from the Host system. This is the FC HBA Initiator from the Host system not DSS V7

## iSCSI – INTERNET SMALL COMPUTER SYSTEM INTERFACE

- **iSCSI encapsulates SCSI frames into a TCP/IP protocol**, so all devices in an iSCSI environment will have addresses
- Every iSCSI address must be unique. Initiators will have IP addresses and targets will have IP Addresses
- **iSCSI and FC volumes emulate a raw SCSI drive**; they are partitioned and formatted with a regular (non-cluster) file system like NTFS, EXT3, XFS, etc . The I/O of an iSCSI target is block-based (as opposed to file-based), which means that changes made by one person will not be seen by another person working on the same target/volume
- Connections can be made with **1GbE, 10GbE and Infiniband**

## iSCSI – INTERNET SMALL COMPUTER SYSTEM INTERFACE

- DSS V7 offer's both Initiator or Target (**MPIO** for best speeds)
- iSCSI Volumes can be initialized (**Writes the volume with all zero's**)
- **Block IO** - System performs real SEEK operation on disk in order to save/write the data. The XFS table will NOT be updated. Data is stored directly on the hard drive (no memory writes required)
- **File IO with Initialization** - System must do real SEEK in order to read/write the data. XFS table has to be read and updated. Stability is higher for the volume
- **File IO No Initialization** - XFS uses first free block for saving the data – no need to use „SEEK” command

## iSCSI – INTERNET SMALL COMPUTER SYSTEM INTERFACE

- Connections can be made with **1GbE**, **10GbE** and **Infiniband**
- DSS V7 offer's both Initiator or Target (**MPIO** for best speeds)
- iSCSI Volumes can be initialized (**Writes the volume with all zero's**)
- **Block IO** - System performs real SEEK operation on disk in order to save/write the data. The XFS table will NOT be updated. Data is stored directly on the hard drive (no memory writes required)
- **File IO with Initialization** - System must do real SEEK in order to read/write the data. XFS table has to be read and updated. Stability is higher for the volume
- **File IO No Initialization** - XFS uses first free block for saving the data – no need to use „SEEK” command

## iSCSI – INTERNET SMALL COMPUTER SYSTEM INTERFACE

- High performance while „random write”, low performance while sequential read (databases)
- **Snapshot and Volume Replication** can be used with FC volumes
- With the **Write Back** function for the LUN the system memory is automatically caching the iSCSI data
- Use the default of **SCST** if possible as the **IET** will reset all the iSCSI Target connections causing disconnects to the initiators. Look to always switch to SCST (Faster and No resets!)
- Default iSCSI port is **3260** and not changeable in the DSS V7

## BEST iSCSI TARGET VALUES TO BE SET FOR VMWARE

1. From the console, press CTRL+ALT+W
2. Select Tuning options -> iSCSI daemon options -> Target options
3. Select the target in question
4. Change the **MaxRecvDataSegmentLength** and **MaxXmitDataSegmentLength** values to the maximal required data size (check w/ the initiator to match) **262144**.  
Set the **MaxBurstLength=16776192** **maxoutstandingr2t=8**

## BEST iSCSI TARGET VALUES TO BE SET FOR VMWARE

When an initiator issues a Write command, the initiator can include a limited amount of immediate data as part of the command PDU, but the remaining data must be “pulled” from the target. The target “pulls” data from the initiator by sending a series of R2T (Ready to Transfer) PDUs, each specifying an offset and length for a portion of the Write data that it is ready to receive. An individual R2T can request up to the one burst of data, and the R2T is considered to be “outstanding” until the target has received all of the data requested by the R2T.

## NAS – NETWORK ATTACHED STORAGE

- NAS enables **multiple users** to share the same storage space simultaneously
- Connections to a NAS Volume can be accessed with **1GbE, 10GbE, Infiniband and with Bond modes**
- NAS communicates via File Protocols, **NFS** (Network File System), **CIFS/SMB** (Common Internet File System and Server Message Blocks), **FTP** (File Transfer Protocol) or **AFP** (AppleTalk File Protocol), **HTTPS** (Hypertext Transfer Protocol)
- **Antivirus** protection against Viruses, that **Quarantine** the infective file, with automatic updates

## NAS – NETWORK ATTACHED STORAGE

- NAS Shares can be Replicated by use of the **Data Replication** to synchronize the data internally, LAN or WAN (use **Port 873**)
- NAS Volumes can use the **Volume Replication** feature to replicate at the Block level
- **Snapshots** can be used with NAS Volumes
- Provides Authentication methods for **Internal and External LDAP, PDC and ADS** for Users and Groups
- Capable of **Repairing the File system** if in the event there is possible corruption
- Offers **WORM** (Cant be Deleted!!! – only by the RAID controller)

## DSS V7 SYSTEM LOGS AND WHY DO WE NEED THEM?

- The logs contain information of the DSS V7 system and help to trouble shoot issues.
- They provide information on how the system behaves and point to bottlenecks in order to make some tunings.

## HOW TO DOWNLOAD THE DSS V7 LOGS?

- They are downloadable via WebGUI → Status → Hardware → Logs
- Generation of system logs can take up to few minutes
- They are compressed with gzip and tarball. They can be unpacked using 7zip or WinRAR

## WHAT IS IN THE DSS V7 SYSTEM LOGS?

- Crucial system services, Configuration of system services, Read performance of attached disks and RAID units.
- Information about currently running system tasks with their exit codes (i.e. Data Replication, Snapshots).
- The DSS V7 system logs do not include any kind of information of the stored data on its storage for FC, iSCSI and NAS
- Most viewed logs from the support team: `critical_error`, `dmesg2`, `test`, `Samba`, `iSCSI`, `ha_debug`...

## WHAT IS THE CRITICAL\_ERRORS.LOG FILE?

- It includes all information that could be found in the event viewer (WebGUI).
- It includes timestamps of logged events.
- The DSS V7 uses syslog and smart expressions for filtering messages
- Configure the DSS V7 Email to receive email alerts

## WHAT IS IT THE DMESG?

- The dmesg displays the contents of the system message buffer.
- Each device driver present in the kernel probes the system for the existence of relevant hardware. If the hardware is located, a diagnostic message is produced documenting precisely what was found.
- It can detect a so called I/O errors of attached disks and RAID units.

## WHAT'S IN THE TESTS.LOG FILE?

- It is divided into sections that are named respectively to the command that generated the output from the internal self- test command
- **It includes:**
  1. Information's from sysfs and procfs
  2. Benchmark results of read performance of disks and RAID units
  3. Output of some other commands (apcaces status, net ads info and so on)
  4. Can find RAID controllers firmware (also can be found in the dmesg log)

## WHAT'S IN THE SAMBA LOGS?

- It is placed in „/samba“ directory in logs package.
- It includes log output of 3 main Samba processes (smbd, nmbda and windbind).
- It includes basic level information of each established connection (log's name can be log.ip\_address in case of the connection is established with IP or log.server\_name in case the connection is established using NetBios name).

## EXAMPLE OF RAID UNIT/DISK I/O ERRORS

```
sdb: rw=0, want=4294967052, limit=2429794304  
Buffer I/O error on device sdb, logical block 1073741761  
CCISS controler /dev/cciss/c0d0 reported: Parity/consistency  
initialization complete, logical drive 1 reading directory #81961 offset  
0<3>sd 1:0:0:0: rejecting I/O to offline device
```

## EXAMPLE OF FILE SYSTEM ERRORS

journal commit I/O error

Call Trace:

```
[<ffffffff80395792>] xfs_alloc_ag_vextent_near+0x512/0x980
[<ffffffff80393b87>] xfs_alloc_fixup_trees+0x317/0x3b0
[<ffffffff803ad9bd>] xfs_btree_setbuf+0x2d/0xc0[<ffffffff80395792>]
xfs_alloc_ag_vextent_near+0x512/0x980
[<ffffffff80395cd5>] xfs_alloc_ag_vextent+0xd5/0x160
[<ffffffff80396546>] xfs_alloc_vextent+0x256/0x470
[<ffffffff803a7865>] xfs_bmap_btalloc+0x475/0xa50
[<ffffffff803a827a>] xfs_bmap_i+0x41a/0x12a0
[<ffffffff80271713>] mempool_alloc+0x43/0x120
```

## TEST.LOG/ETHTOOL (BAD PACKETS)

```
eth1  Link encap:Ethernet  HWaddr 00:25:90:21:38:43
      inet addr:192.168.1.220  Bcast:192.168.1.255
      Mask:255.255.255.0
      UP BROADCAST MULTICAST  MTU:1500  Metric:1
      RX packets:0 errors:0 dropped:5689 overruns:0 frame:0
      TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
```

## TEST.LOG /ETHTOOL (NIC SPEED)

Supported link modes: 10baseT/Half 10baseT/Full  
100baseT/Half 100baseT/Full  
1000baseT/Full

Supports auto-negotiation: Yes

Advertised link modes: 10baseT/Half 10baseT/Full  
100baseT/Half 100baseT/Full  
1000baseT/Full

Advertised auto-negotiation: Yes

Speed: 100Mb/s

Duplex: Full

Port: Twisted Pair

PHYAD: 1

## EXAMPLE OF POOR PERFORMANCE

The system has poor write performance in the **dmesg log** we can find:

```
sd 1:2:0:0: [sdb] Very big device. Trying to use READ CAPACITY(16).
sd 1:2:0:0: [sdb] 8779857920 512-byte hardware sectors (4495287 MB)
sd 1:2:0:0: [sdb] Write Protect is off
sd 1:2:0:0: [sdb] Mode Sense: 1f 00 10 08
sd 1:2:0:0: [sdb] Write cache: disabled, read cache: enabled, supports DPO and FUA
sd 1:2:0:0: [sdb] Very big device. Trying to use READ CAPACITY(16).
sd 1:2:0:0: [sdb] 8779857920 512-byte hardware sectors (4495287 MB)
sd 1:2:0:0: [sdb] Write Protect is off
sd 1:2:0:0: [sdb] Mode Sense: 1f 00 10 08
sd 1:2:0:0: [sdb] Write cache: disabled, read cache: enabled, supports DPO and FUA
```

## EXAMPLE OF CONNECTION ISSUES WITH AN ADS DOMAIN

Cannot renew kerberos ticket

Time skew is greater than 5 minutes

```
[2011/07/04 14:29:31.495026, 0] utils/net_ads.c:285(ads_startup_int)
```

```
ads_connect: No logon servers
```

Didn't find the ldap server!

## PRE-SALES ([PRE-SALES@OPEN-E.COM](mailto:PRE-SALES@OPEN-E.COM))

- Specifically used for inquiring about new features or drivers or hardware certification
- Is not used for technical support
- Can be used for requesting information about the product – “Can the DSS V7 work as a Virtual Machine”
- Please use when you have an opportunity that you want to verify or quantify before your customers purchase

## SUPPORT INFRASTRUCTURE

- Register all DSS V7 products
- All support cases need to be entered from your User Portal

## SUPPORT INFRASTRUCTURE

- Premium vs Standard Support
- Use the Forum and Knowledge Base articles for additional references, use the Open-E “Search” function
- Hardware Compatibility List on what we support, though newer hardware may not be listed it is due to the fact that we do not have them in our labs to test but drivers could be the same or just a small update
- Use the DSS V7 Trial version, it is the full version and allows you to use this for 60 days, after this time period the system will throttle down to a low speed but data is still accessible or you can delete the volume

## SUPPORT INFRASTRUCTURE

- Read the Release notes! Make sure that you enable the “Subscriptions” tab to receive the updates for DSS V7

## BEST PRACTICES AND USEFUL INFORMATION

- Take advantage of the “Solutions, How to resources and the Webcasts and video’s”
- <http://blog.open-e.com/> “Random vs. Sequential explained”, “Just how important is write cache?”, “A Few practical tips about lometer” and “Bonding versus MPIO explained” ....

## OPEN-E DSS V7 UPDATES

- We recommend to have at least one release behind the current version
- Builds older then 5217 should be updated
- Issues with older builds like 4622 “Low memory....” that you have encountered must be updated to prevent this. Even if it has been running perfect for a period of time it is very hard to tell when this will happen and hard for support to determine how to prevent this issue and best to schedule downtime to update.
- If using the iSCSI Auto Failover feature update the Secondary first then restore the connection then Failover and update the Primary then once online then Sync from the Secondary to the Primary and then use the Failback feature.

*IMPORTANT NOTE: In case you experience a problem with the availability of iSCSI / FC volumes after the upgrade from version <= 4622, please change the Identification Device (VPD) compatibility to SCST 1.0.*

*Run this from the Console:*

*(CTRL+ALT+W -> Tuning options -> SCST subsystem options -> Device Identification (VPD) compatibility -> SCST).*

*Important notes regarding updating a system with failover in use:*

*\* In case of using VMware ESX(i) or MS Hyper-V Initiator system, you need to change the Identification Device (VPD) compatibility to SCST 1.0 on the secondary node. This is located in the Console tools (CTRL+ALT+W -> Tuning options -> SCST subsystem options -> Device Identification (VPD) compatibility -> SCST VPD)*

- Then once the Secondary is running click on the start button in the Failover manager.*
- Now update the Primary system using the software update functionality and reboot.*

*\* In case of using VMware ESX(i) or MS Hyper-V as Initiator system, change the Identification Device (VPD) compatibility to SCST 1.0 on the Primary node. This is located in the Console tools (CTRL+ALT+W -> Tuning options -> SCST subsystem options -> Device Identification (VPD) compatibility -> SCST VPD)*

- Once the Primary is running go to the Secondary and click on the Sync volumes button in the Failover manager.*
- Then click on the Failback button in the Secondary system.*
- The Primary system now be go back to the active mode and ready for another failover.*

## INTRODUCTION TO HARDWARE CERTIFICATION PROGRAM

### What does the HCP provide for your company?

- Ensures your systems are fully compatible with the Open-E DSS V7
- 2 types of certification tests offered by the HCP program:
  - **OPEN-E HARDWARE CERTIFICATION FULL**  
You send Storage server to our facility to be certified by our engineers 3000€. This process take approximately 2 weeks to be completed
  - **OPEN-E HARDWARE CERTIFICATION PROOF**  
You test your own Storage server with our testing procedures, logs and other supporting tools 1000€.



## INTRODUCTION TO HARDWARE CERTIFICATION PROGRAM

- Once that tests have been completed your Storage servers will be listed on our website
- Use the Open-E Certified System Logo for your hardware

The collage illustrates the hardware certification process. It features a document titled "Intel SR2612UR storage system" with a "CERTIFIED" logo. Another document shows "Test description and environment" with a diagram of a server rack. A third document, labeled "1.1 Intel SR2612UR pictures", shows three views of the server rack. A fourth document, labeled "3. Writing and read test on SMB shares connected to Workstation via iSCSI controller", includes a bar chart showing performance metrics. The right side of the collage shows a screenshot of the Open-E website, displaying the "Storage Certified Systems by Open-E" page. The website shows the "Where to buy" section with a "Where to buy" button and a "Where to buy" dropdown menu. Below this, there are two entries for the Intel SR2612UR server, each with a "Full report" button and a "PDF" button. The first entry shows the "Specification" and "Recommendation" for the Intel SR2612UR server, including details like "Motherboard: Intel S5520UR", "CPU: 2x Intel Xeon Quad-core E5540 2.53GHz", "Memory: 8GB DDR3 ECC-REG", "Storage: LSI SAS1078 RAID, 8Gb Qlogic FC HBA", "HDD slots: 12 x SATA/SATAII", and "Network: 4x1GbE Intel, 1x10GbE Intel". The "Recommendation" section lists "Backup", "HA cluster", "iSCSI storage", and "NAS filer". The "Certification Date" is 21.06.2011.

## COMPARE THE OPTIONS OF FULL AND PROOF FOR THE OPEN-E HARDWARE CERTIFICATION PROGRAM

Step	Open-E Hardware Certification Full	Full	Proof	Workdays
1	Getting certification request	Open-E		0
2	Verification of certification request	Open-E		0.5
3	Selection of proper set of tests (depends on hardware configuration)	Open-E	Partner	0.5 – 1
4	Preparation of test environment (depends on set of tests)	Open-E	Partner	0.5 – 1
5	Performing the tests and report the test results (depends on set of tests)	Open-E	Partner	5 – 10
6	Verification of test results (depended on set of tests)	Open-E	Partner	0.5 – 1
7	Taking photos of certificated system	Open-E	Partner	0.5
8	Preparation of Certification Report (PDF document)	Open-E	Open-E	1
9	Preparation of Certification Document	Open-E	Open-E	0.5
10	Upload information to the Open-E Website - Subpage	Open-E	Open-E	0.5
11	Upload information to the Open-E Website - Hardware Compatibility List	Open-E	Open-E	0.5
<b>TOTAL - max.</b>				<b>16.5</b>

# Thank you!