Capgemini
ENERGY

Welcome to Capgemini Energy, a business service company that will aim to deliver your immense depth of utility industry knowledge
and business service skills to energy and utility customers across North America.

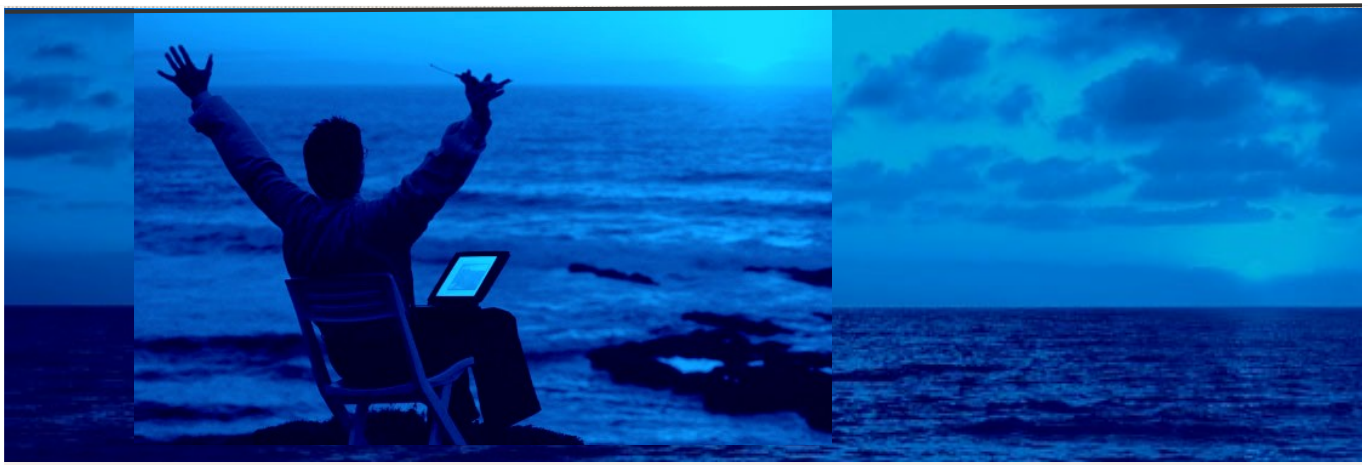

Unix Storage

Introduction Presentation
May 4, 2005

Introduction

- Scope
- Storage Concepts
- Storage Topology
  - Internal SCSI Disks
  - External RAID
  - External SSA
  - NFS
  - SAN
  - NAS
  - Virtual
- Working with HDS command utilities
- Policies, Guidelines and Procedures

Capgemini ENERGY

Scope

- Scope
  - This presentation covers Storage standards, design and concepts in the CGE environment.
  - This presentation is not designed as a tutorial on Storage operations.
  - This presentation will cover necessary command sets to communicate with the CGE storage group.

Capgemini ENERGY

Storage Concepts

- Redundancy
- High Availability
- Disaster Recovery

Capgemini ENERGY

Storage Redundancy

- Purpose of storage redundancy is business continuity. This is used to prevent one or more failures of hardware or software resulting in data loss.

- The common understanding of storage redundancy is to utilize functionalities of the different RAID levels (0-10) or software based technology such as data replication copies.

- Storage Redundancy can be viewed in three levels:
  - Standalone Server
  - Clustered Environment (Spanning multiple servers)
  - Disaster Recovery Environment (Spanning multiple sites)

- Each of the three levels of Storage Redundancy can be observed within the CGE environment and/or its clients.

Capgemini ENERGY

Storage High Availability

- Purpose of storage high availability is business continuity and/or disaster recovery. This is used to prevent one or more failures of hardware or software resulting in data loss.
- The common understanding of storage high availability is redundancy within a cluster of servers when viewed from a server centric aspect and cluster of storage array units from a storage centric view.
- Storage High Availability can be viewed from these aspects:
  - Multiple servers accessing single storage devices
  - Storage Arrays mirroring across multiple frames or within a single frame.
  - Single or multiple servers accessing multiple storage devices.
- Each of these aspects of High Availability can be found within the CGE environment or that of its clients

Capgemini ENERGY

Storage Disaster Recovery

- Purpose of storage disaster recovery is driven by business requirements. This is used to prevent one or more failures of hardware or software resulting in data loss at a site to site level.
- The common understanding of storage disaster recovery is redundancy between sites for servers and storage arrays.
- Storage Disaster Recovery can be viewed from these aspects:
  - Multiple servers replicating single storage devices across multiple sites.
  - Multiple servers replicating transaction changes across multiple sites.
  - Storage Arrays replicating across multiple frames across diverse locations. (Site to Site replication)
- Replication can be synchronous or asynchronous
- Each of these aspects of Disaster Recovery can be found within the CGE environment or that of its clients

Capgemini ENERGY

Storage Topology

- Internal SCSI Disks
- External RAID
- External SSA
- NFS
- SAN
- NAS
- Virtual

Capgemini ENERGY

Storage – Internal SCSI Disks

- Internal SCSI disks are used in the CGE AIX environment for storage of the AIX operating system.
- One or more internal SCSI disks may exist in an enclosure
- In systems where more than one internal SCSI disk exists, software mirroring is used to replicate the AIX operating system.
- In standalone database systems not connected to the SAN, striping may be used for enhanced performance
- Quorum requirements is generally not a consideration due to the CGE environment strategy
- Current architectural direction is to utilize external disk technology for both Application and Operating System volumes.

Capgemini ENERGY

Storage – External RAID

- This segment refers to raid configurations outside of the SAN environment
- Raid is used on IBM SCSI external units. (T42)
- Current Raid levels utilized within these constraints is 0, 1 and 5.
- This configuration is currently used on legacy environments and on non-production systems in a standalone mode. (not highly available)

Capgemini ENERGY

Storage – External SSA

- SSA is direct attached storage and is currently utilized only in Non-HA environments.
- SSA is utilized in development as well as production environments.
- SSA provides multiple paths to an SSA disk and provides redundant loops to an SSA drawer to provide a Highly Available Disk subsystem.
- SSA is currently not the desired direction of CGE for its storage solutions and will be phased out as time passes.
- There are currently only about 6 SSA drawers utilized within the CGE environments.

Storage – NFS

- NFS – Network File System. Storage provided by other servers across the network. (Non-NAS storage)
- NFS storage requirements are based on application requirements.
- NFS storage presentation happen from Unix to Unix, Windows to Unix and Unix to Windows.
- NFS version 2 and 3 are utilized based upon Operating System limitations
- NFS is utilized by NIM to provide installation, patching and CD image presentation.

Storage – SAN

- Current SAN environment utilizes Hitachi 9980s and 9585s within the midrange Unix environment. There is currently over 20 Terabytes within this environment.
- The switch connectivity between storage arrays and systems are Cisco and Brocade.
- The storage subsystems are managed and maintained by the CGE Storage Management Group.
- Each system utilizes redundant paths through the SAN fabric to provide high availability in the event of a hardware failure.
- Disk mirroring is generally not done within the SAN environment.
- Server based striping is not performed because the SAN environment is already performing this action.
- The current direction is for all storage to reside on the SAN environment, including boot volumes.

Capgemini ENERGY

Storage – NAS

- Network Attached Storage (NAS) – Provides storage across the ethernet network utilizing NFS and SMB protocols.
  - Network File System (NFS)
  - Server Message Block (SMB)
- NAS Storage provides a mechanism for multiple machines to have simultaneous access to programs, applications, and data
- NAS storage is managed and maintained by the CGE Storage Management Group
- NAS Storage is utilized by servers running AIX, Solaris, Linux, Windows
- Access to NAS Storage is dependent upon business and application requirements
- The architectural direction for NAS is non-production environments
- CGE currently utilizes NAS for both application and database storage.

Capgemini ENERGY

Storage – Virtual

- Virtual Storage is storage presented to virtual I/O servers which in turn present this storage wholly or partially to a Virtual LPAR.
- In this scenario Virtual I/O Servers act as the SAN or storage management subsystem in storage presentation to LPARs.
- Virtual I/O Server is an AIX based platform that runs within an LPAR and is restricted to the Power 5 technology at this time.
- Physical storage is presented to the Virtual I/O Server which in turn reallocates the storage to the virtual LPARs.
- The Physical storage can be SAN, internal or external disks.
- Virtual storage presentation is viewed by the virtual LPAR as a standard (virtual) SCSI drive.
- The Virtual I/O Server can present entire disks or logical volumes as storage to a virtual LPAR.

Capgemini ENERGY

Working with HDS Command Utilities

- Each disk in the AIX environment has a variety of identities. These identities are used for different purposes. The identities include:
  - Location Code
  - PVID
  - Hdisk Name
  - DLM Name
  - World Wide Name (WWN)
  - LDEV
- To determine each of the disk identities requires knowledge of multiple tools, both AIX and Hitachi utilities.
- Communication with the Storage Management team will require WWN's and LDEV information for each disk.

Capgemini ENERGY

Working with HDS Command Utilities (cont)

- Bare metal installations and maintenance booting will require knowledge of location codes.
- Volume group management will require knowledge of PVID's, Hdisk names, and DLM names.
- The standard AIX utilities “lspv”, “lsdev”, “lsattr”, and “lscfg” provide some of this information
- The Hitachi HDLM utilities exist in the directory “/usr/DynamicLinkManager/bin” on each SAN attached system.
- The primary HDLM tool for obtaining disk information is “dlnkmgr”.
- Hitachi DLM utilities is an older HDS Command utility that still exists within the CGE environment and has a separate set of command utilities
- Example scripts to generate this disk identification information exist at the following URL:
  - http://eperf.tu.com/GlobalSysAdmin/AIX/dlmfdrvmap.shtml

dlnkmgr Help

- Obtaining help for the “dlnkmgr” command is not intuitively obvious:

```
$ dlnkmgr view -help

view:
  Format
    dlnkmgr view -sys [ -sfunc | -msrv | -adrv | -pdrv | -lic ] [-t]
    dlnkmgr view -path -c [-stname] [-srt {lu | cp}] [-t]
    dlnkmgr view -path [ -hdev HostDeviceName ] [-stname]
                                                  [-srt {pn | lu | cp}] [-t]
    dlnkmgr view -path -item [pn] [dn] [lu] [cp] [type] [ic] [ie] [dnu] [hd]
                  [ -hdev HostDeviceName ] [-stname] [-srt {pn | lu | cp}] [-t]
    dlnkmgr view -drv [-t]
```

Capgemini ENERGY

dlnkmgr Output

- The following is example output from the “dlnkmgr” command presenting DLM Fibre Drive to AIX hdisk relationships with associated LDEV numbers.

```
$ dlnkmgr view -drv
PathID HDevName  Device  LDEV
000000 dlmfdrv1  hdisk1  9970/9980.10066.01C5
000001 dlmfdrv1  hdisk16 9970/9980.10066.01C5
000002 dlmfdrv2  hdisk17 9970/9980.13127.00A0
000003 dlmfdrv2  hdisk2  9970/9980.13127.00A0
000004 dlmfdrv3  hdisk18 9970/9980.13127.00A1
000005 dlmfdrv3  hdisk3  9970/9980.13127.00A1
000006 dlmfdrv4  hdisk19 9970/9980.13127.00A2
000007 dlmfdrv4  hdisk4  9970/9980.13127.00A2
000008 dlmfdrv5  hdisk20 9970/9980.13127.00A3
000009 dlmfdrv5  hdisk5  9970/9980.13127.00A3
000010 dlmfdrv6  hdisk21 9970/9980.13127.00A4
000011 dlmfdrv6  hdisk6  9970/9980.13127.00A4
```

Capgemini ENERGY

DLM Command Utilities

- The commands utilized in the DLM utility set are:
  - cfallvpath
  - extendvg4vp
  - lsvpcfg
  - mkvg4vp
  - restvg4vp
  - savevg4vp
  - vp2hd
  - hd2vp

- The utility “lsvpcfg” is most commonly used to see if redundant paths are available and LUN identifier information.

- A sample output of the command:

vpath0 (Avail pv egate1vg) HI_R450294B0E67 = hdisk5 (Avail) hdisk50 (Avail)
vpath1 (Avail pv egate1vg) HI_R450294B0E68 = hdisk6 (Avail) hdisk51 (Avail)
vpath2 (Avail pv egate6vg) HI_R450294B0E69 = hdisk7 (Avail) hdisk52 (Avail)
vpath3 (Avail pv egate104vg) HI_R450294B0E6E = hdisk8 (Avail) hdisk53 (Avail)
vpath4 (Avail pv egate104vg) HI_R450294B0E6F = hdisk9 (Avail) hdisk54 (Avail)

Capgemini ENERGY

Storage Policies

- Storage policies for this presentation denote policies related to the AIX environment.
- All systems boot from SAN and may or may not have an internal disk mirror.
- All storage required by application and data storage is external to the server.
- All Servers have redundant paths to the storage.
- Root Volume Group mirrors may be accomplished by AIX or by the Storage Array.
- Application Volume groups will rely on Storage Array internal redundancy in place of mirroring.
- All NFS mounts should be presented by NAS devices

Capgemini ENERGY

Storage Guidelines

- SAN-HDS 9980 Storage Arrays are utilized for Production environments only.
- SAN-HDS 9585 (SATA) Arrays are utilized for Development and Test environments.
- NAS Storage is presented whenever a file share is needed between Unix systems and also provides non-production storage for servers.
- If an internal disk exists on a server, the preferred boot device is the external attached storage. (exception is NAS)

Capgemini ENERGY

Storage Procedures

- Storage Procedures
  - Storage Requests
  - Lead Time
  - Change Control
  - Implementation
  - Documentation
- Storage Reclamation
- For Further storage procedures and examples:
  - http://eperf.tu.com/GlobalSysAdmin/storage/procedures.shtml

Capgemini ENERGY

Helpful Links

- HDLM for AIX Reference
  - http://eperf.tu.com/GlobalSysAdmin/Storage

Capgemini ENERGY

Q&A

Capgemini ENERGY