



HP MC/SG 11.17 & Veritas CFS Seminar

김도엽 (do-yup.kim@hp.com)
HP Services Global Solution Center



Serviceguard A.11.17 Delta Training Agenda




- Module 1 SG/CFS introduction
- Module 2 Demo
- Module 3 Requirements / Installation
- Module 4 CFS Concepts
- Module 5 SG A.11.17 new concepts
- Module 6 Use Cases
- Q & A
- Wrap Up



HP Serviceguard CFS Introduction

Module 1





HP Virtual Server Environment integrates all key virtualization techniques



HP Virtual Server Environment (VSE)

Intelligent policy engine



Consolidates & virtualizes server
resources for optimum utilization

Control

Workload Manager
gWLM
Systems Insight
Manager (SIM)

Availability

Serviceguard
Serviceguard
Manager
SGeRAC
SGeFF
SGeSAP
Disaster Tolerance

Partitioning

nPars (Partition
Manager)
vPars
HP Integrity Virtual
Machines with sub
CPU and shared I/O
Secure resource pars
PRM / pSets

Utility Pricing

Instant Capacity
Temporary
Instant Capacity
Pay Per Use

Serviceguard

HP's strategic clustering solution



- **Reputation:** 10 years of delivering and enhancing a high quality, robust product
- **Design Criteria:** Data integrity, availability, performance
- **Investment protection:** 125,000+ HP-UX servers running with Serviceguard today, PA RISC to Integrity failover, up-down version compatibility & support
- **HP-UX Integration:** Tightly coupled to HP-UX to provide O/S specific value added solutions, Infiniband high speed/low latency interconnect, WLM, gWLM, iCAP, Integrity Virtual Machine, nPARS, vPARS, Secure Containers, Secure Resource Partitions
- **VERITAS Integration:** Serviceguard provides a superior integration level with HP-UX to enable advanced management, features and functionality
- **Time to market:** HP Servers, Storage and HP-UX O/S new features, functionality and solutions are supported by Serviceguard at introduction
- **Portfolio:** SGeSAP, SGeFF, SGeRAC, Campus Cluster, Metrocluster, Continentalclusters Disaster Tolerant solutions, ISV toolkits, Serviceguard / VERITAS solution offerings
- **ISV Validation:** tested & validated with 1000's ISV
- **Virtualization:** foundation product for Virtual Server Environment and Integrity Virtual Machine
- **Manageability:** SG Manager integrated with Single Virtual View, Insight Manager, OpenView
- **Multi O/S Investment:** HP-UX, Linux
- **Support:** Single point of contact for the entire solution stack in mission critical environments

HP Serviceguard and the VERITAS Storage Foundation release



- HP will sell a version of VERITAS CFS and VERITAS Storage Foundation Products which are fully integrated with SG and SGeRAC products
 - HP will sell bundles containing SG, SGeRAC, ECMT, VERITAS CFS, and other VERITAS Storage Foundation components
 - These bundles are unique and will only be sold by HP
- In the first release, Serviceguard and VERITAS Storage Foundation products will be integrated with
 - Serviceguard
 - SGeRAC (extension for Oracle RAC)
 - ECMT (application toolkit)
- To be done after SG/SGeRAC testing is completed, not part of EAP program, expect support after SR:
 - SGeFF
 - Extended Clusters
- Within 3 months of SG/VRTS release
 - SGeSAP (toolkit for SAP)
- Under investigation for delivery in 2006 (Disaster Tolerant Solutions)
 - Continentalclusters
 - Metrocluster

Cluster File System overview



- Partnership with VERITAS to include the VERITAS Storage Foundation Cluster File System (CFS) in the HP High Availability solutions portfolio.
- VERITAS Storage Foundation CFS
 - An integrated solution for clustered and SAN computing environments resulting in transparent data sharing.
 - Servers in a cluster can access data at the file or file system level concurrently.
 - Simultaneous access to storage allows for load sharing amongst servers and enables full utilization of enterprise RAID subsystems.

What is the SG Cluster File System?



- Adds an additional disk volume configuration option for an application's data within a SG cluster but retains previous options
 - Within a given SG/CFS Cluster different applications could use any combination of the following mechanisms:
 - Local raw volumes that failover
 - Local file systems that failover
 - Shared raw volumes (e.g., Oracle RAC)
 - Shared file systems (CFS)
 - LVM and VxVM/CVM volumes can coexist in the same cluster but CFS file systems must use CVM

Key Benefits



- HP-UX customers can adopt a Cluster File System free of radical changes to their existing Serviceguard environment.
- Increased ease of management for some applications
- Faster failover expected with use of CFS when there are large numbers of disk groups and file systems
- Customers choice to or not to implement CFS
- The transition from non-CFS to CFS is an easy SG upgrade
- A mixed environment of CFS/non-CFS is supported



Serviceguard bundle product numbers

- T2771BA HP Serviceguard Storage Management
- T2772BA HP Serviceguard Storage Management Premium
- T2773BA HP Serviceguard Storage Management for Oracle
- T2774BA HP Serviceguard Storage Management for Oracle Premium
- T2775BA HP Serviceguard Cluster File System
- T2776BA HP Serviceguard Cluster File System for Oracle
- T2777BA HP Serviceguard Cluster File System for RAC



HP Serviceguard Storage Management

Product Number: T2771BA

Common to all bundles

- Serviceguard A.11.17
- VxVM
- Enterprise Cluster Master Toolkit
- 2TB File System
- VERITAS Enterprise Administrator (GUI)

New with 4.1:

- Dynamic Multi-pathing (enhanced with 4.1)
- Intelligent Storage Provisioning
- Configuration Backup
- Storage Expert
- History log
- Online Intent Log Resize
- Named Data Streams
- Online LUN Resize
- Portable Data Containers
- Hardware Assisted E-Copy
- Multi-Volume File System

From 3.5:

- Online Administration



HP Serviceguard Storage Management Premium

Product Number: T2772BA

Common to all bundles

- Serviceguard A.11.17
- VxVM
- Enterprise Cluster Master Toolkit
- 2TB File System
- VERITAS Enterprise Administrator (GUI)

New with 4.1:

- Dynamic Multi-pathing (enhanced with 4.1)
- Intelligent Storage Provisioning
- Configuration Backup
- Storage Expert
- History log (new with 4.1)
- Online Intent Log Resize
- Named Data Streams
- Online LUN Resize
- Portable Data Containers
- Hardware Assisted E-Copy
- Multi-Volume File System

From 3.5:

- Online Administration

Unique to this bundle

New with 4.1:

- Quality of Storage (QoS)
- Instant Volume snapshots
- Checkpoint rollback

From 3.5:

- Storage CheckPoints
- Disk Group Split and Join (FlashSnap)
- Fast Mirror Resync



HP Serviceguard Storage Management For Oracle

Product Number: T2773BA

Common to all bundles

- Serviceguard A.11.17
- VxVM
- Enterprise Cluster Master Toolkit
- 2TB File System
- VERITAS Enterprise Administrator (GUI)

New with 4.1:

- Dynamic Multi-pathing (enhanced with 4.1)
- Intelligent Storage Provisioning
- Configuration Backup
- Storage Expert
- History log (new with 4.1)
- Online Intent Log Resize
- Named Data Streams
- Online LUN Resize
- Portable Data Containers
- Hardware Assisted E-Copy
- Multi-Volume File System

From 3.5:

- Online Administration

Unique to this bundle

From 3.5, not previously licensed:

- Quick I/O – for pre-9i Oracle
- Oracle Disk Manager – for Oracle 9i and 10g
- VxDBA GUI
- Storage Rollback

Serviceguard Storage Management for Oracle Premium



Product Number: T2774BA

Common to all bundles

- Serviceguard A.11.17
- VxVM
- Enterprise Cluster Master Toolkit
- 2TB File System
- VERITAS Enterprise Administrator (GUI)

New with 4.1:

- Dynamic Multi-pathing (enhanced with 4.1)
- Intelligent Storage Provisioning
- Configuration Backup
- Storage Expert
- History log (new with 4.1)
- Online Intent Log Resize
- Named Data Streams
- Online LUN Resize
- Portable Data Containers
- Hardware Assisted E-Copy
- Multi-Volume File System

From 3.5:

- Online Administration

Unique to this bundle

New with 4.1:

- Quality of Storage
- Instant Volume Snapshots
- Checkpoint Rollback

From 3.5:

- Quick I/O – for pre-9i Oracle
- Oracle Disk Manager – for Oracle 9i & 10g
- VxDBA GUI
- Storage Mapping to LUN Level
- Storage rollback
- Database FlashSnap
- Storage Checkpoints
- Disk Group Split and Join
- Fast Mirror Resync



Serviceguard Cluster File System

Product Number: T2775BA

Common to all bundles

- Serviceguard A.11.17
- VxVM
- Enterprise Cluster Master Toolkit
- 2TB File System
- VERITAS Enterprise Administrator (GUI)

New with 4.1:

- Dynamic Multi-pathing (enhanced with 4.1)
- Intelligent Storage Provisioning
- Configuration Backup
- Storage Expert
- History log (new with 4.1)
- Online Intent Log Resize
- Named Data Streams
- Online LUN Resize
- Portable Data Containers
- Hardware Assisted E-Copy
- Multi-Volume File System

From 3.5:

- Online Administration

Unique to this bundle

- Cluster Files System
- Clustered Volume Manager

New with 4.1:

- Quality of Storage (QoSS)
- Instant Volume snapshots
- Checkpoint rollback

From 3.5:

- Storage CheckPoints
- Disk Group Split and Join (FlashSnap)
- Fast Mirror Resync



HP Serviceguard Cluster File System for Oracle

Product Number: T2776BA

Common to all bundles

- Serviceguard A.11.17
- VxVM
- Enterprise Cluster Master Toolkit
- 2TB File System
- VERITAS Enterprise Administrator (GUI)

New with 4.1:

- Dynamic Multi-pathing (enhanced with 4.1)
- Intelligent Storage Provisioning
- Configuration Backup
- Storage Expert
- History log (new with 4.1)
- Online Intent Log Resize
- Named Data Streams
- Online LUN Resize
- Portable Data Containers
- Hardware Assisted E-Copy
- Multi-Volume File System

From 3.5:

- Online Administration



HP Serviceguard Cluster File System for RAC

Product Number: T2777BA

Common to all bundles

- Serviceguard A.11.17
- VxVM
- Enterprise Cluster Master Toolkit
- 2TB File System
- VERITAS Enterprise Administrator (GUI)

New with 4.1:

- Dynamic Multi-pathing (enhanced with 4.1)
- Intelligent Storage Provisioning
- Configuration Backup
- Storage Expert
- History log (new with 4.1)
- Online Intent Log Resize
- Named Data Streams
- Online LUN Resize
- Portable Data Containers
- Hardware Assisted E-Copy
- Multi-Volume File System

From 3.5:

- Online Administration

Unique to this bundle

- SGeRAC
- Cluster File System
- Clustered Volume Manager

New with 4.1:

- Quality of Storage
- Instant Volume Snapshots
- Checkpoint Rollback

From 3.5:

- Oracle Disk Manager – for Oracle 9i & 10g
- VxDBA
- Storage Mapping to LUN Level
- Storage rollback
- Storage Checkpoints
- Disk Group Split and Join
- Fast Mirror Resync



HP Serviceguard CFS Demo

Module 2



MC/SG 11.17 & CFS Demo



```
[root@krcipf1:/] swlist |grep T2777BA
T2777BA      A.01.00      HP Serviceguard Cluster File System for RAC
```

```
[root@krcipf1:/]
```

```
[root@krcipf1:/] cmviewcl
```

```
CLUSTER      STATUS
```

```
CFS-Cluster  up
```

```

NODE          STATUS      STATE
krcipf1       up          running
krcipf2       up          running
```

```
MULTI_NODE_PACKAGES
```

```

PACKAGE       STATUS      STATE      AUTO_RUN    SYSTEM
SG-CFS-pkg    up          running    enabled     yes
SG-CFS-DG-1   up          running    enabled     no
SG-CFS-MP-1   up          running    enabled     no
```

MC/SG 11.17 & CFS Demo



```
[root@krcipf1:/] vxdctl -c mode
mode: enabled: cluster active - MASTER
master: krcipf1
[root@krcipf1:/]
```

```
[root@krcipf2:/] vxdctl -c mode
mode: enabled: cluster active - SLAVE
master: krcipf1
[root@krcipf2:/]
```

MC/SG 11.17 & CFS Demo



```
[root@krcipf1:/] bdf
```

Filesystem	kbytes	used	avail	%used	Mounted on
/dev/vg00/lvol3	294912	272488	22424	92%	/
/dev/vg00/lvol1	573440	374592	197352	65%	/stand
/dev/vg00/lvol8	6668288	5538808	1129480	83%	/var
/dev/vg00/lvol7	3440640	2846280	591128	83%	/usr
/dev/vg00/lvol6	6750208	6697832	51976	99%	/opt
/dev/odm	0	0	0	0%	/dev/odm
/dev/vx/dsk/dg01/vol3					
	512000	68019	416239	14%	/cfsmnt1

```
[root@krcipf1:/cfsmnt1] ll
```

```
total 132338
```

-r--r--r--	1	root	sys	1510	Jan 24 22:03	hosts
drwxr-xr-x	2	root	root	96	Jan 24 21:57	lost+found
-rwxr-xr-x	1	root	sys	67754696	Jan 24 22:03	vmunix

```
[root@krcipf1:/cfsmnt1]
```

MC/SG 11.17 & CFS Demo



```
[root@krcipf2:/] bdf
Filesystem      kbytes    used    avail %used Mounted on
/dev/vg00/lvol3 18300928 12213224 6040296 67% /
/dev/vg00/lvol1  311296   150344  159768  48% /stand
/dev/vg00/lvol8  4718592 3499736 1209472 74% /var
/dev/vg00/lvol7  3375104 3033152  339328 90% /usr
/dev/vg00/lvol6  3883008 3747256  134736 97% /opt
/dev/odm          0         0         0      0% /dev/odm
/dev/vx/dsk/dg01/vol3
                    512000   68019   416239  14% /cfsmnt1
```

```
[root@krcipf2:/cfsmnt1] ll
total 132338
-r--r--r--  1 root      sys           1510 Jan 24 22:03 hosts
drwxr-xr-x  2 root      root          96 Jan 24 21:57 lost+found
-rwxr-xr-x  1 root      sys          67754696 Jan 24 22:03 vmunix
[root@krcipf2:/cfsmnt1]
```



HP Serviceguard CFS requirements / installation

Module 3



Serviceguard/SGeRAC CFS configuration requirements – misc.



- SG/SGeRAC/CFS will be supported with HP-UX 11i v2 September 2004 release (0409) and later
 - HP 9000 and Integrity systems - all 11iv2 Update2 09/2004 supported servers will be supported with Serviceguard/SGeRAC A.11.17.
 - Serviceguard/CFS supports mixed clusters (PA & Integrity)
 - SGeRAC does not support mixed clusters
 - Serviceguard/SGeRAC version A.11.17
 - 8-node clusters with SG/CVM/CFS (16-node w/o CFS)
 - VERITAS VxVM/CVM VxFS/CFS version 4.1
 - dedicated cluster interconnect (two are recommended)
 - LVM/SLVM may be used concurrently on the same cluster for non-CFS file systems
 - Oracle 9i will be supported at initial release
 - Oracle 10g will be supported in Q4 2005
- Storage Foundation bundles will be part of the standard HP-UX media with AR1205

Serviceguard/SGeRAC CFS configuration requirements – misc.



- At this time, the SG bundles will NOT be supported with:
 - HP-UX 11i v1
 - HP-UX 11.0 or earlier
- It is not possible to do a rolling upgrade from VERITAS revision 3.5 to 4.1.
- Disk groups created with VERITAS 3.1 and later are compatible with VERITAS 4.1, however some 4.1 features may not be available for disk groups created with these earlier VERITAS revisions.
 - The lab is working on a whitepaper for how to do the upgrade.
- With CVM or CFS, all cluster member nodes must be connected to all shared disks configured for CVM or CFS, even if a node is not actually using CVM or CFS.

Serviceguard/SGeRAC CFS configuration requirements - heartbeat



- Special Heartbeat Configuration Requirements for CVM and CFS
 - When using VERITAS CVM 4.1 with Serviceguard, CVM will use the network interfaces configured as the Serviceguard Heartbeat interfaces for its own inter-node communications.
 - Only Ethernet LAN type is supported for use as Heartbeat subnets if CVM or CFS is used.
 - CVM 4.1 supports multiple network interfaces, and therefore multiple Heartbeat LANs configured in the Serviceguard configuration.
 - Two Heartbeat LANs are required
 - It is recommended, but not required to configure Standby LANs for these Heartbeat LANs, for additional protection.
 - Due to CVM 4.1 restrictions, configurations utilizing APA as the heartbeat network cannot be supported.
 - CVM 3.5 and **4.1** cannot use IPoverIB for their inter-node communication. Therefore if CVM is being used, you must not configure the Serviceguard Heartbeat over InfiniBand interfaces.
 - Since CFS requires CVM, it is subject to the same requirements as for CVM 4.1.

Supported Storage



<u>Bus Type</u>	<u>HP Product #</u>	<u>Comments</u>
HP MSA Disk Arrays:		
	Modular Smart Array 1000 (MSA1000 Active/Active only (not active passive))	Not supported with VxVM 4.1, CVM & CFS, plan to support this soon.
	Modular Smart Array 1500 (MSA1500 Active/Active only (not active passive))	Not supported with VxVM 4.1, CVM & CFS, plan to support this soon.
Enclosures:		
Ultra 160 SCSI	HP SureStore Disk System 2300 (DS2300) A6490A, A6490AE A6490AV, A6490AZ, A6490AD, A6490ED	Supported with VxVM 4.1, CVM, & CFS
Ultra 320 SCSI	Multiple initiator (MSA A30MI) 359645-B21	Supported with VxVm 4.1 CVM or CFS Note: the MSA30MI is a JBOD not a RAID array.

Supported Storage



<u>Bus Type</u>	<u>HP Product #</u>	<u>Comments</u>
HP Storage works Disk Array XP:		
	XP48 A5920A & A5925A or A5926A XP512 A5950A & A5955A or A5956A XP128 A7875A & A7876A XP1024 A7905A & *A7906A	Supported w/ VxVM 4.1, CVM, & CFS
	XP12000 AE001A	Supported w/ VxVM 4.1, CVM, & CFS
HP Storageworks Enterprise Virtual Array:	EVA4000, EVA6000, EVA8000	Supported/w VxVM 4.1, CVM, & CFS
HP Storageworks Enterprise Virtual Array:	EVA3000, EVA5000 active/active only	Estimated to be Supported/w VxVM 4.1, CVM, & CFS sometime in 2006 via an upgrade

Supported Storage



<u>Bus Type</u>	<u>HP Product #</u>	<u>Comments</u>
EMC:		
Fibre Channel	Symmetrix 3330, 3430, 3630, 3700, 3830, 3930, 8130, 8230, 8430, 8530, 8730, 8830	Supported/w VxVM 4.1 , CVM &, CFS
Fibre Channel	EMC DMX 800, 1000, 2000, 3000	Supported/w VxVM 4.1, CVM, & CFS

SG/SGeRAC/OS/CFS Compatibility Matrix



HP-UX Version	VERITAS Version	SG/SGeRAC version 11.16 and earlier	SG/SGeRAC version 11.17 and later
11.0	3.5	Yes	No
11i v1 (11.11)	3.5	Yes	No
11i v2 (11.23) (prior to 0509)	3.5	Yes	Yes
11i v2 (11.23 0509 and later)	3.5	Yes	Yes
	4.1 (with or without CFS)	No	Yes
11i v3 (11.31)	4.1 (with or without CFS)	No	Yes

SGeRAC 11.17

RAC support Matrix

- New in 11.17
- Post Cert C activity
- Upgrade to 4.1



		9i RAC				10gR1	10gR2		
		CVM 3.5	SLVM	CVM 4.1	CFS 4.1	CVM 3.5	SLVM	CVM 4.1	CFS 4.1
Cluster size		8	16	8		8	16	8	
Heartbeat limitation		single				single			
SGeFF 2-node, 2HB		no	Y	Y		no	Y	Y	
Extended Cluster	10km	2,4,6,8	2	2,4,6,8		2,4,6,8	2	2,4,6,8	
	100km	2		2				2	
HF2 support (EMS monitoring)	UDP	Yes				Yes		Yes	
	HMP	Yes				PA only		PA only	
Cache Fusion over IB support (SG subnet monitoring, no HB)		No plan				No plan	IPF, (PA support in Sept)	IPF, (PA support in Sept)	
SG Oracle common CI over one IB net								outplan	

- Verify correct version of OS
 - HP-UX 11iv2 Update2
- Verify hardware configuration
 - Supported storage
 - Dual heartbeat
- Install bundle using swinstall
 - will require a reboot of the system
- Verify installation of bundle
- Install VEA (optional)

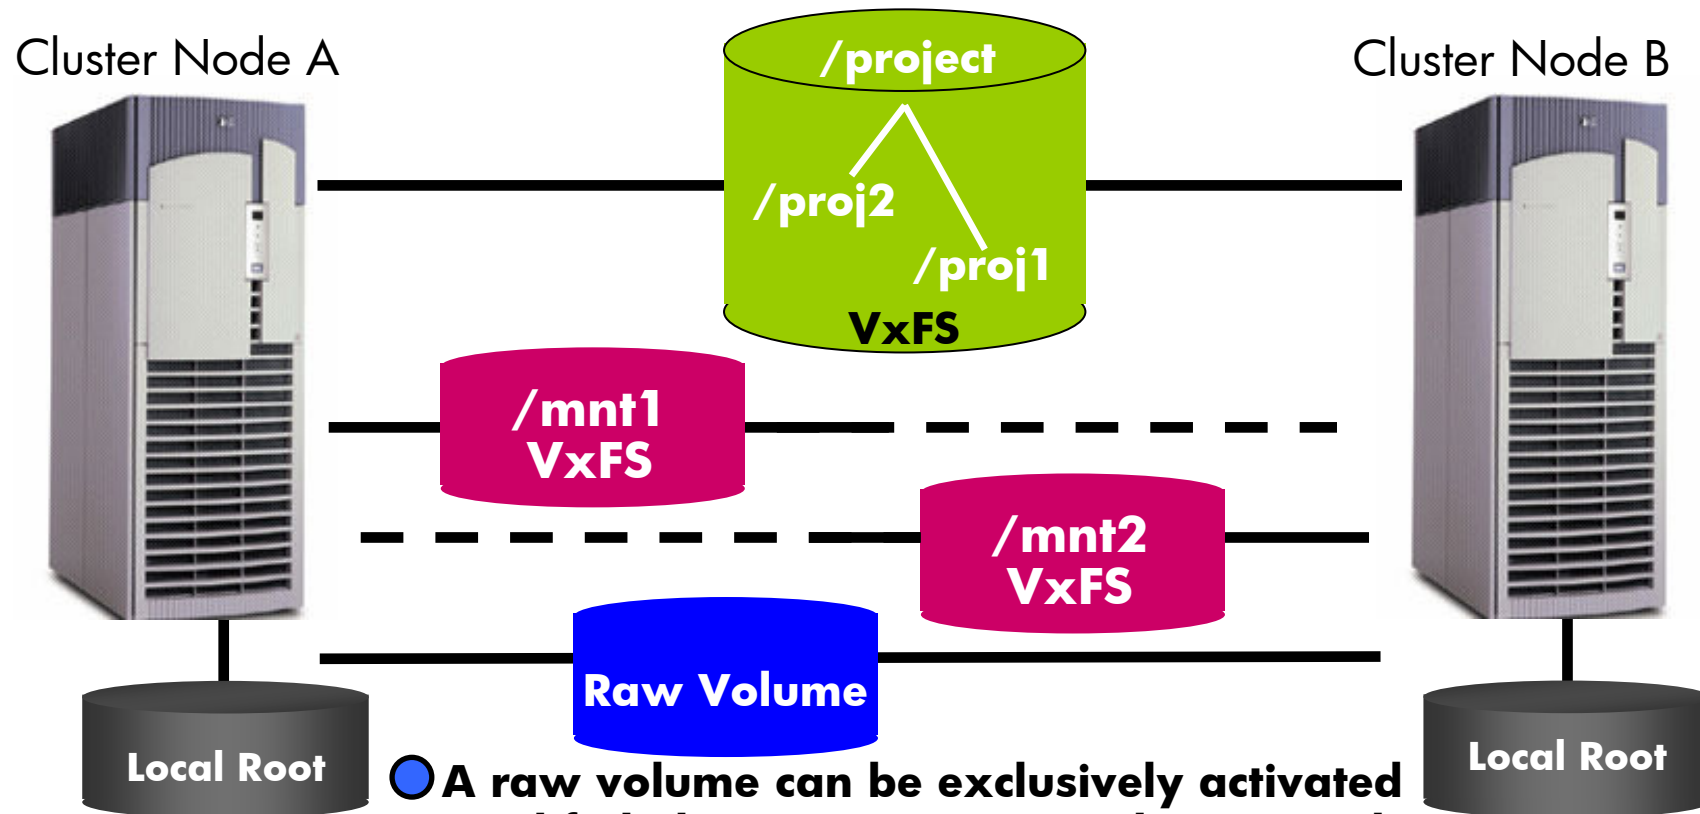


HP Serviceguard CFS concepts

Module 4



Cluster File System versus Failover



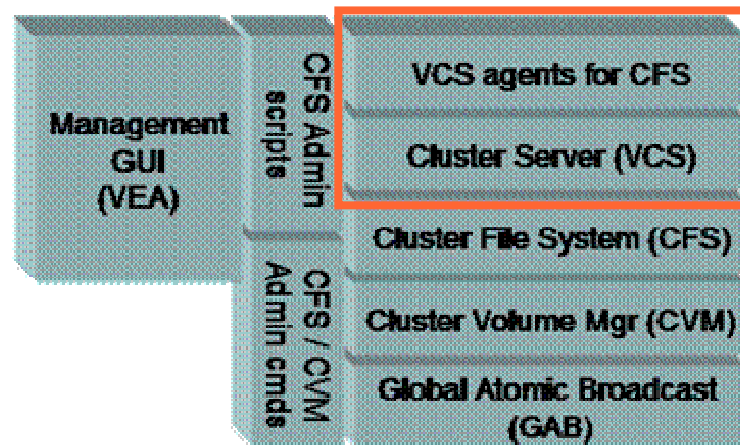
● A raw volume can be exclusively activated and failed over OR concurrently accessed

- A cluster file system is concurrently accessed by all cluster nodes.
- A failover file system is exclusively activated by one node and can transition between nodes.

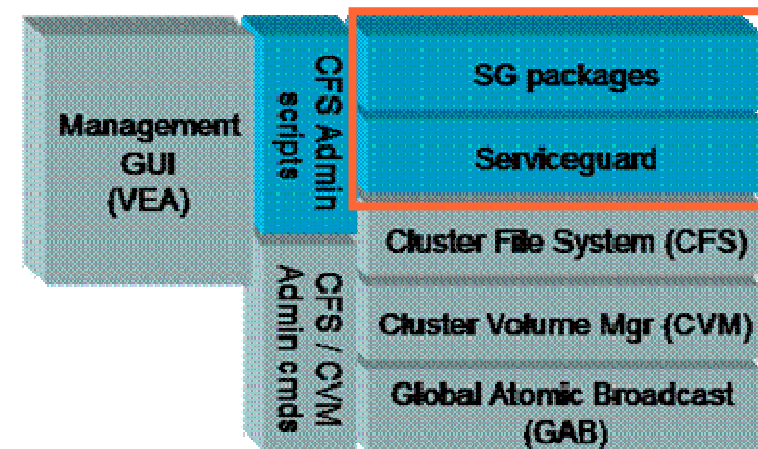
VERITAS and Serviceguard CFS comparison



VERITAS Solution

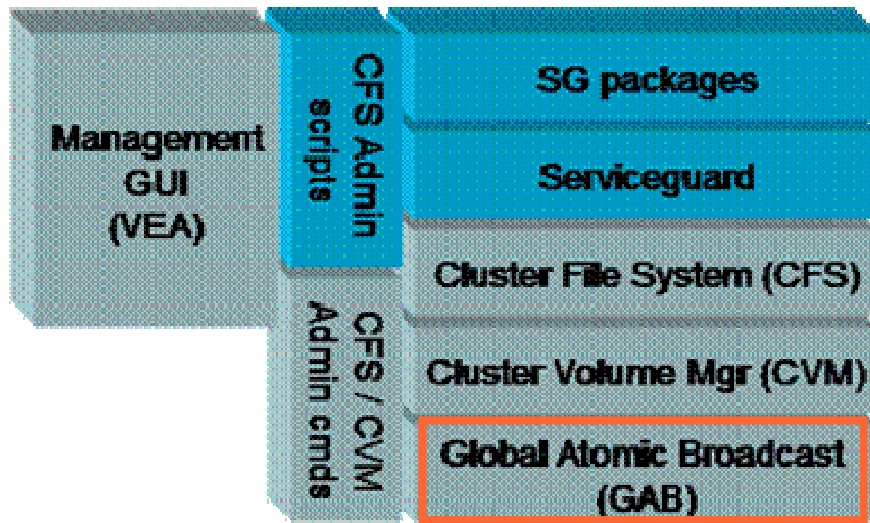


HP Serviceguard Solution



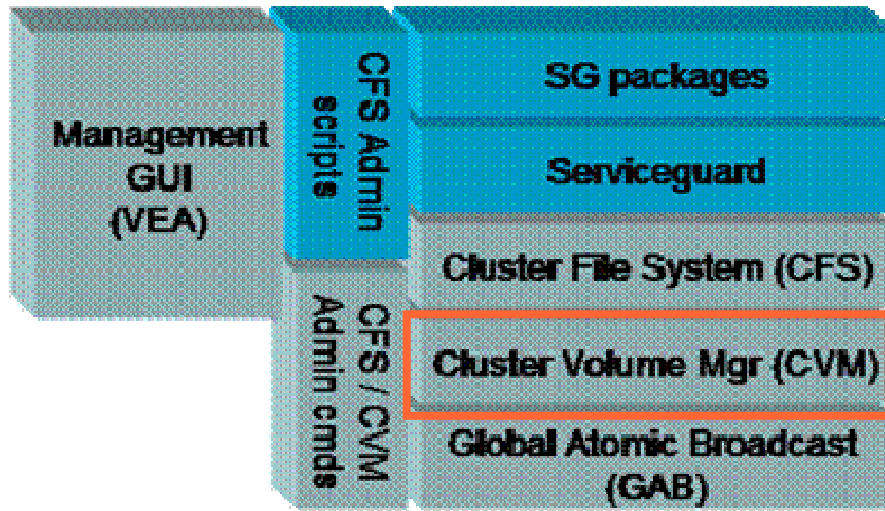
- HP Serviceguard replaces the VERITAS Cluster Server
- HP Serviceguard packages replace the VCS agents
- CFS Admin scripts are utilized by and integrated with Serviceguard

Global Atomic Broadcast (GAB)



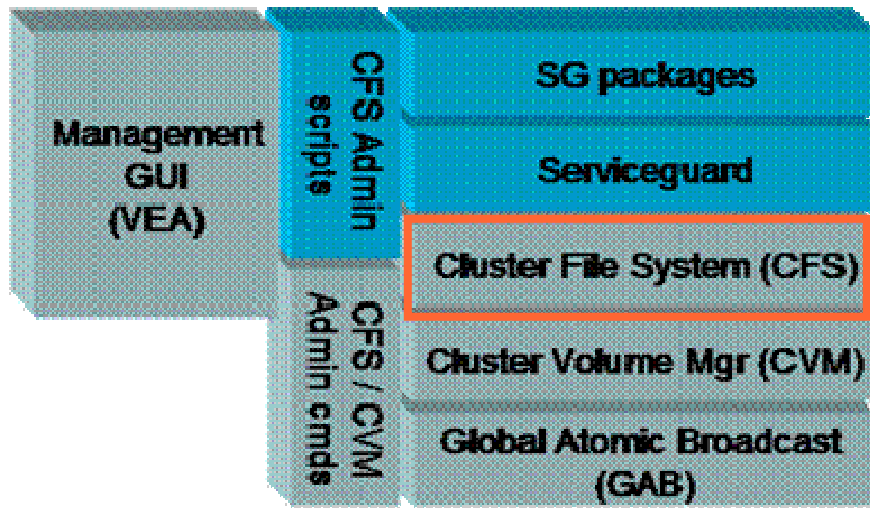
- GAB provides membership and messaging services for the cluster and applications to the file system
 - communication and synchronization layer

Cluster Volume Manager (CVM)



- CVM enables multiple hosts to concurrently access the logical volumes.
 - The cluster monitor component of CVM informs VxVM of changes in cluster membership.
 - When a node joins a cluster it gains access to shared disks.
 - A node joins a cluster when the cluster monitor is started on that node.
 - VxVM objects configured in shared disk groups can be accessed by all nodes that join the cluster.
 - One node acts as the master and all others are slave nodes

Cluster File System (CFS)



- CFS provides shared file system access
 - CFS implements a primary /secondary architecture for management of file system metadata on shared disks.
 - The first node to mount becomes the primary node all subsequent node mounts are secondaries.
 - Secondaries send requests to the primary to perform metadata updates.
 - Only the primary node updates the metadata and maintains the file system's metadata.
 - Data can be updated from any node directly to shared storage.
 - File system operations, such as allocating or deleting files, can originate from any node in the cluster.
 - In the event of a primary node failure an election process begins to define a new primary from the remaining nodes in the cluster.

What is ODM?



- API defined by Oracle to enhance File Management and Disk I/O throughput – Vendors Implement API.
- Designed for Oracle9i or later
- Best suited for Database on File Systems (VxFS).
- Support Oracle's File Management and I/O calls for database storage on VxFS file systems, raw volumes or raw partitions.

- ODM API co-developed by Oracle and VERITAS

Oracle Disk Manager Features



☐ Support for File System I/O

- Raw performance for file systems
- Single system call for all I/O types
- Supported on both raw and file system files

☐ File Management

- Contiguous disk space allocation
- Simplified file creation
- Supports Oracle Autoextend

☐ File Identification

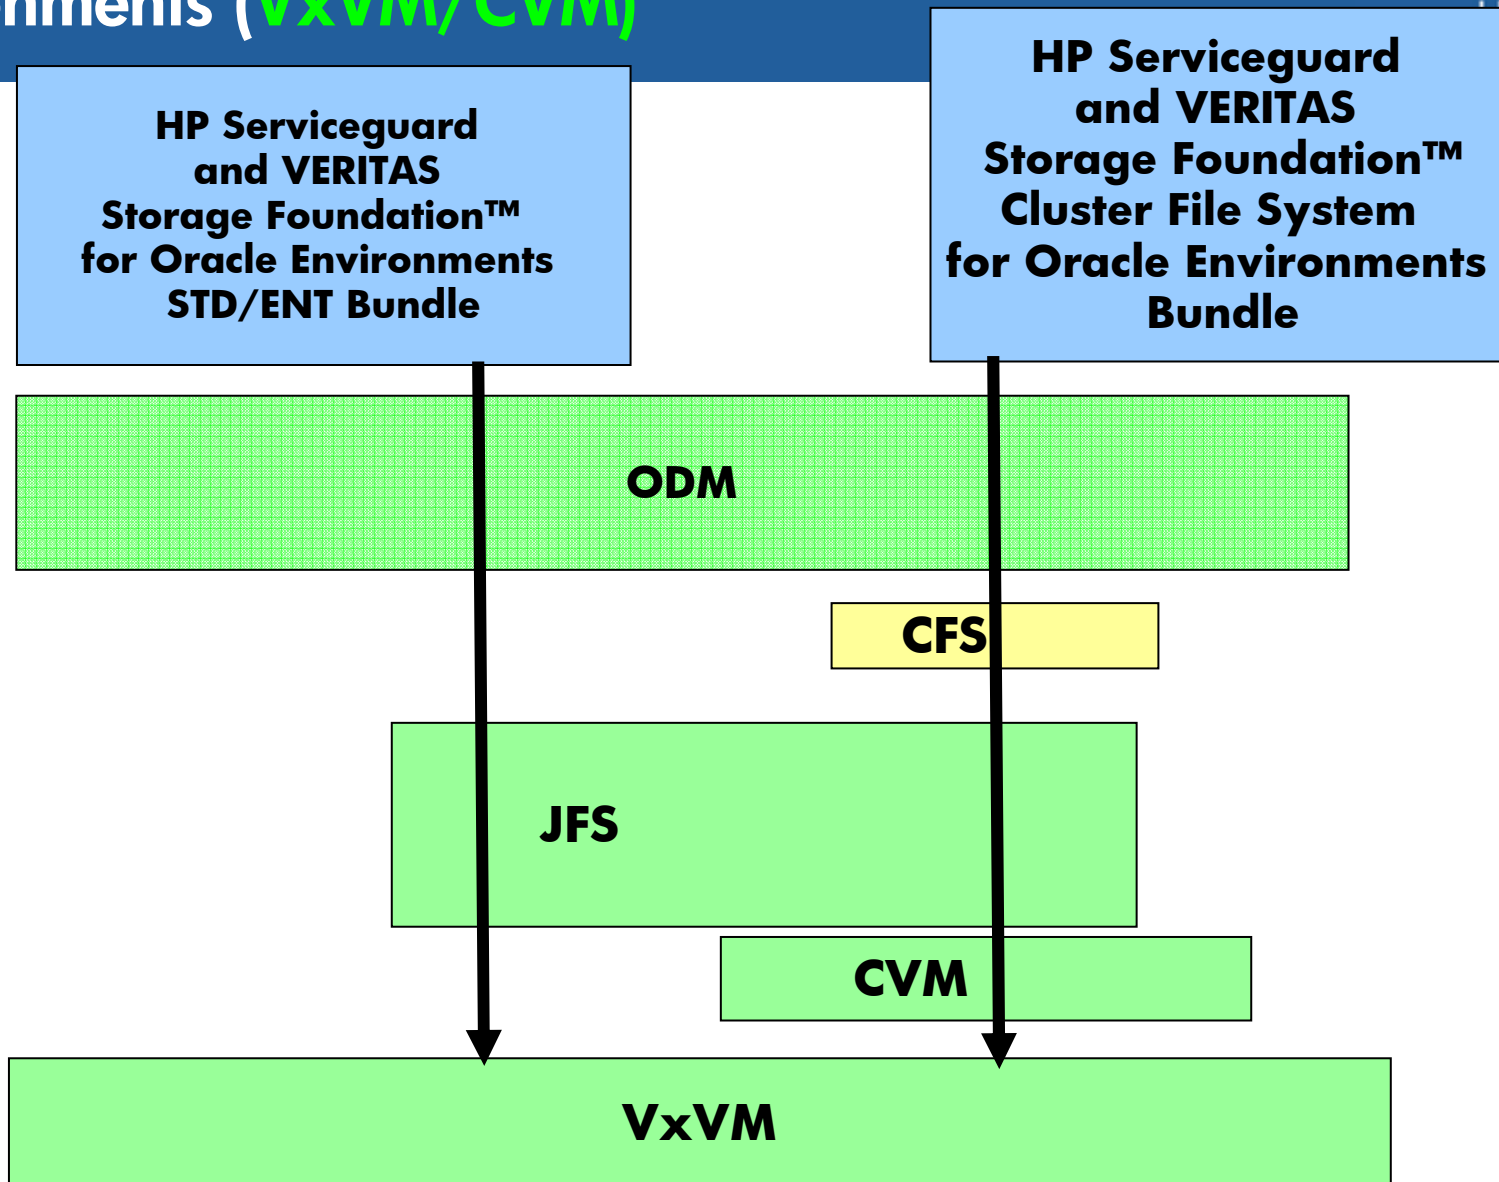
- Reduced CPU overhead

Features/Benefits



- VxFS/CFS With/Without ODM.
- Features
 - File Identification
 - One identifier for all processes – reduced system overhead
 - File Creation and Management
 - Atomic file creation (odm_create) and contiguous file allocation
 - File IO Processing
 - Reduced system call overhead and near raw device performance for file system files.
- Performance and Manageability
 - Kernel Asynchronous IO for files
 - Reduced system call overhead
 - Improved file system layout by preallocating contiguous files on a VxFS file system
 - Performance on filesystems that is near raw device performance
 - Avoid Duplicate file “opens”

ODM in Single instance Oracle and SG environments (VxVM/CVM)



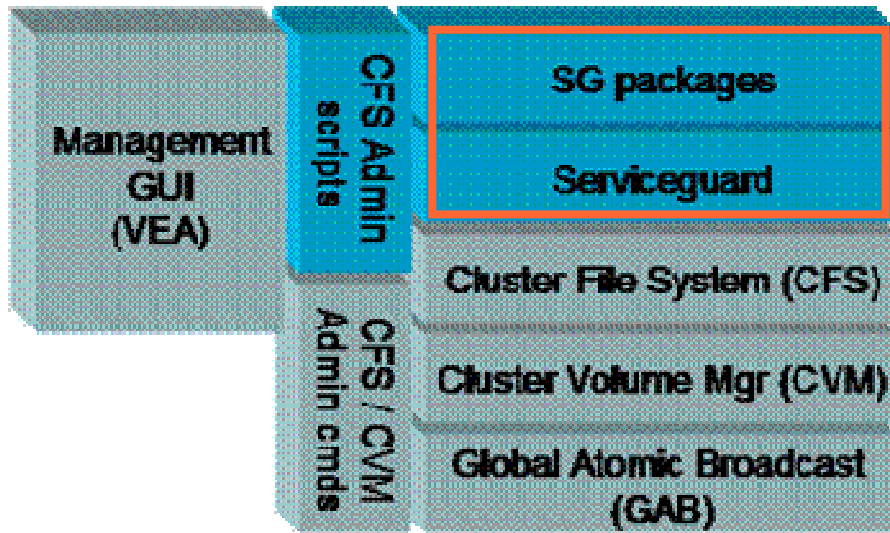


HP Serviceguard concepts

Module 5



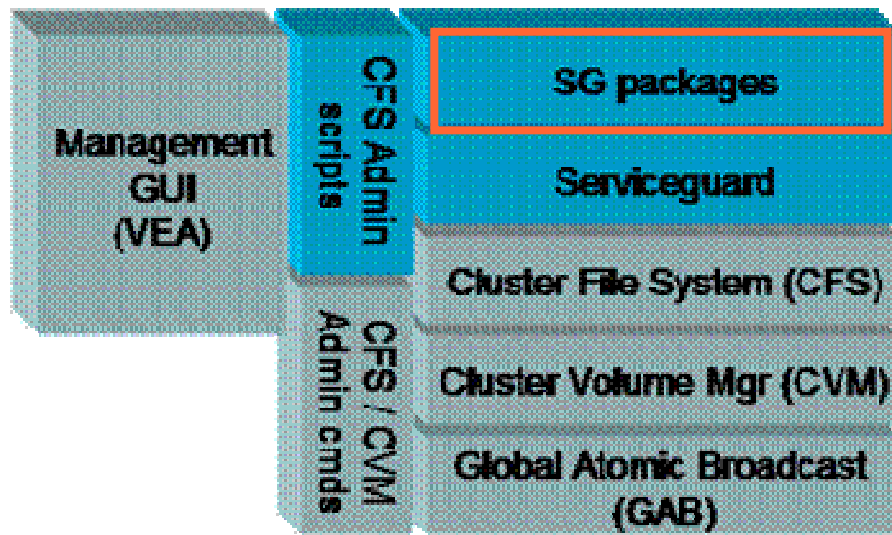
New Serviceguard concepts



- Additions to Serviceguard

- CFS
 - Cluster file system mounted on one or more nodes simultaneously.
- system multi-node package
 - A package that runs on all nodes in a cluster simultaneously.
 - similar to the CVM 3.5 system multi-node package
- multi-node packages
 - A package that can run on multiple nodes in a cluster simultaneously.
- package dependencies
 - A formal dependency where one package is dependent on another.
 - Defined in the package ASCII configuration file
 - DEPENDENCY_NAME *unique_name*
 - DEPENDENCY_CONDITION *package_name =up*
 - DEPENDENCY_LOCATION *SAME_NODE*

New Serviceguard concepts



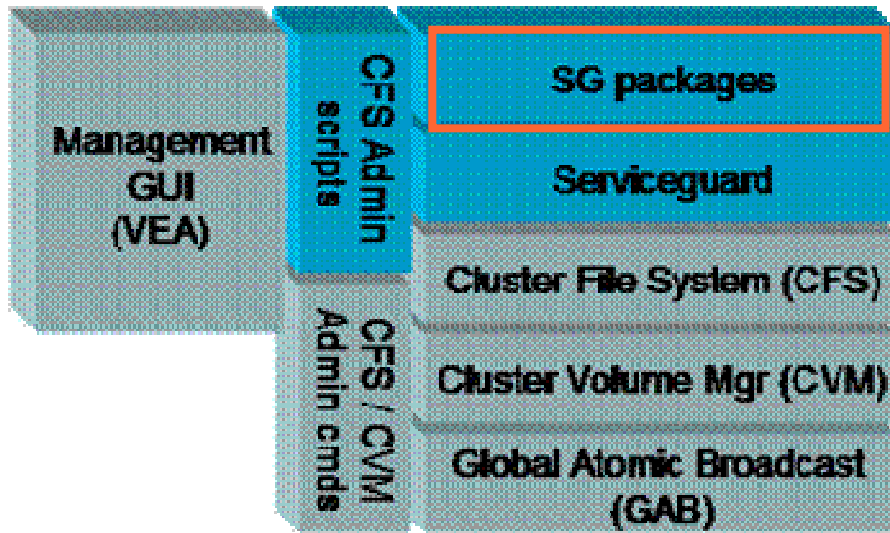
- System multi-node package
 - configures and starts CVM/CFS
 - creates Low Latency Transport (LLT) file
 - /etc/llttab
 - should not be edited
 - uses SG heartbeat network
 - starts LLT
 - creates Global Atomic Broadcast (GAB) file
 - /etc/gabtab
 - should not be edited
 - starts GAB
 - starts CVM and joins the node into the cluster
 - starts CFS daemons as package services
 - is not limited to a single heartbeat IP
 - does not support APA for heartbeat
 - is created using *cfsccluster*

SMNP – SG-CFS-pkg configuration



```
PACKAGE_NAME          SG-CFS-pkg
PACKAGE_TYPE          SYSTEM_MULTI_NODE
NODE_NAME              *
AUTO_RUN              YES
LOCAL_LAN_FAILOVER_ALLOWED  YES
NODE_FAIL_FAST_ENABLED  YES
RUN_SCRIPT             /etc/cmcluster/cfs/SG-CFS-pkg.sh
RUN_SCRIPT_TIMEOUT     300
HALT_SCRIPT            /etc/cmcluster/cfs/SG-CFS-pkg.sh
HALT_SCRIPT_TIMEOUT    60
SCRIPT_LOG_FILE        /etc/cmcluster/cfs/SG-CFS-pkg.log
SERVICE_NAME          SG-CFS-sqcvmd
SERVICE_FAIL_FAST_ENABLED  YES
SERVICE_HALT_TIMEOUT  5
SERVICE_NAME          SG-CFS-vxfscd
SERVICE_FAIL_FAST_ENABLED  YES
SERVICE_HALT_TIMEOUT  30
SERVICE_NAME          SG-CFS-cmvxd
SERVICE_FAIL_FAST_ENABLED  YES
SERVICE_HALT_TIMEOUT  5
```

New Serviceguard concepts



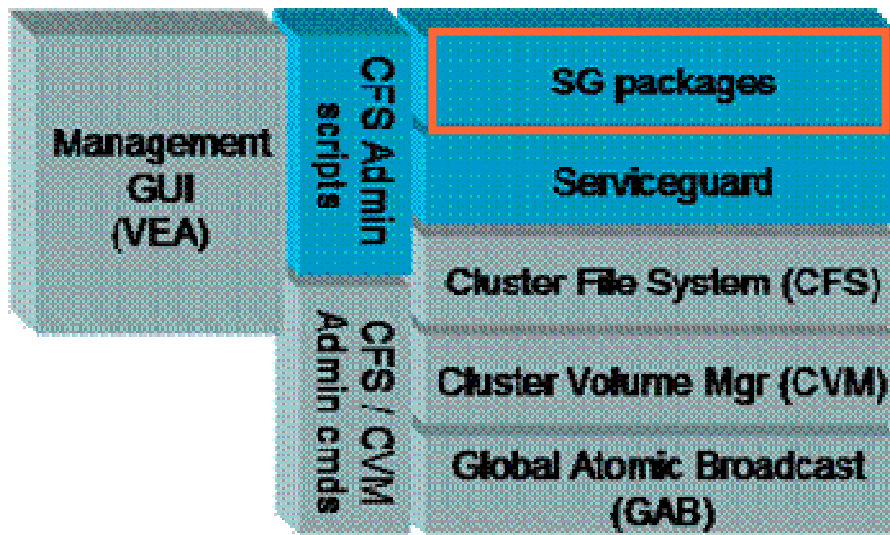
- Multi-node package types
 - Disk group packages
 - sets activation mode of shared disk group
 - controls 1 disk group
 - NODE_FAIL_FAST is not required
 - AUTO_RUN is not required
 - can be started and stopped independently
 - No relocatable IP address support
 - Cannot specify FAILOVER_POLICY or FAILBACK_POLICY
 - has a dependency on SG-CFS-pkg
 - created using *cfsgadm*
 - if name is not specified then it will be automatically generated
 - SG-CFS-DG-#ID (sequence number)

MNP – SG-CFS-DG-# configuration



```
PACKAGE_NAME      SG-CFS-DG-1
PACKAGE_TYPE      MULTI_NODE
NODE_NAME         chives
NODE_NAME         curry
AUTO_RUN          YES
LOCAL_LAN_FAILOVER_ALLOWED  YES
NODE_FAIL_FAST_ENABLED  NO
RUN_SCRIPT        /etc/cmcluster/cfs/sgcfspkg.sh
RUN_SCRIPT_TIMEOUT      NO_TIMEOUT
HALT_SCRIPT       /etc/cmcluster/cfs/sgcfspkg.sh
HALT_SCRIPT_TIMEOUT    NO_TIMEOUT
SCRIPT_LOG_FILE    /etc/cmcluster/cfs/SG-CFS-DG-1.log
DEPENDENCY_NAME    SG-CFS-pkg
DEPENDENCY_CONDITION  SG-CFS-pkg = UP
DEPENDENCY_LOCATION  SAME_NODE
CVM_DISK_GROUP    cfsdgl
CVM_ACTIVATION_MODE  "chives=sw curry=sw"
```


New Serviceguard concepts



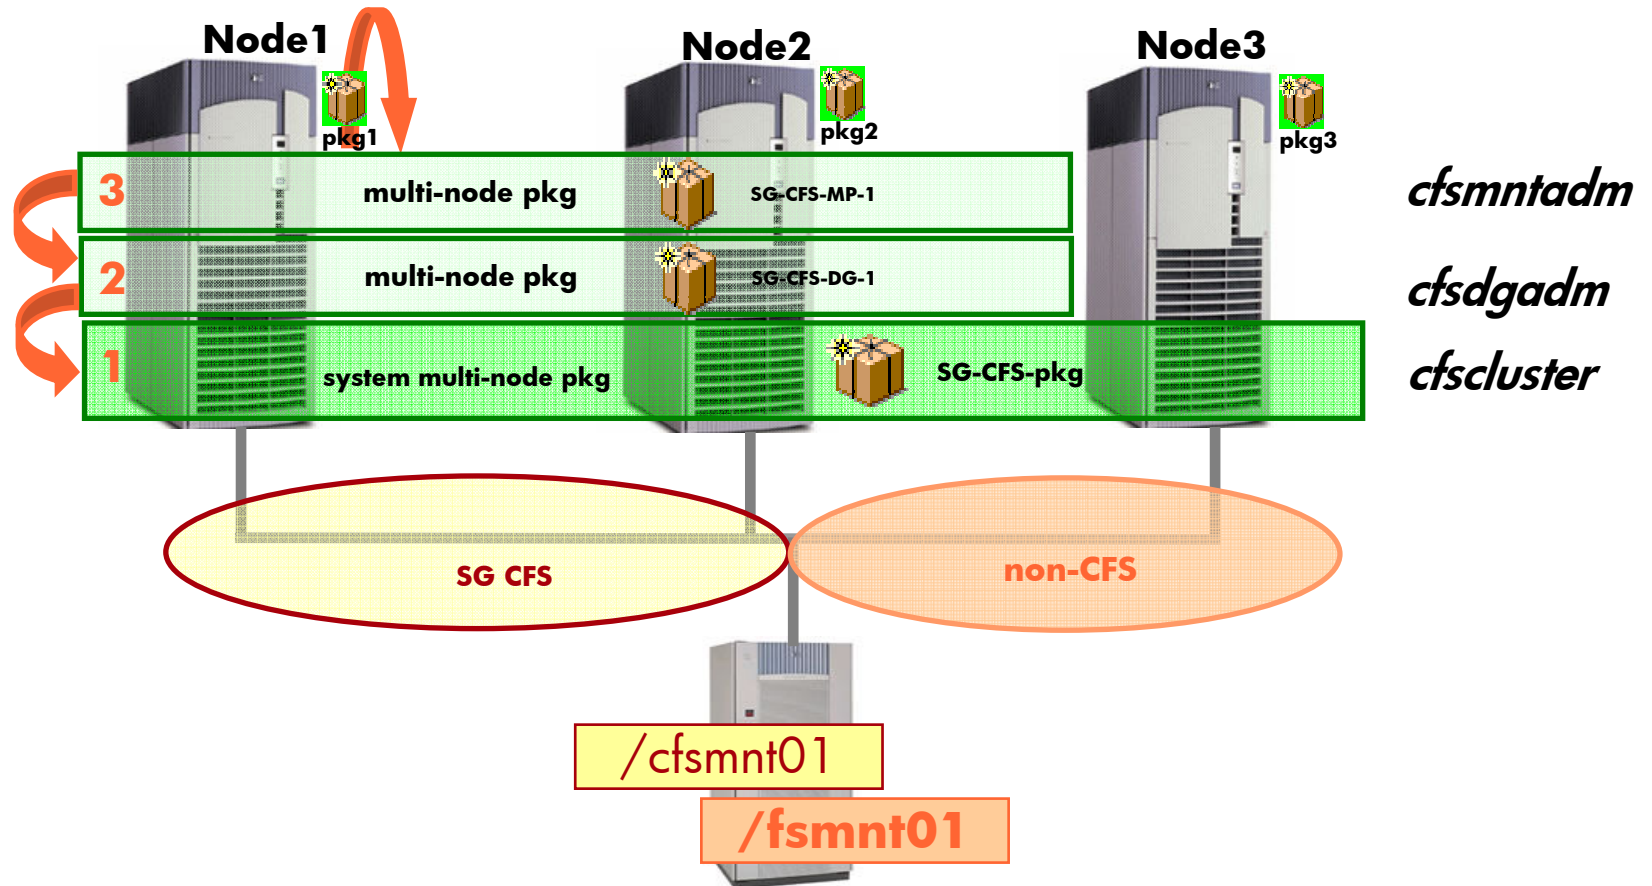
- Multi-node package types
 - Mountpoint packages
 - mounts and unmounts the filesystem
 - controls 1 CFS filesystem
 - NODE_FAIL_FAST is not required
 - AUTO_RUN is not required
 - can be started and stopped independently
 - No relocatable IP address support
 - Cannot specify FAILOVER_POLICY or FAILBACK_POLICY
 - has a dependency on SG-CFS-DG-# pkg
 - created using *cfsmntadm*
 - if name is not specified then it will be automatically generated
 - SG-CFS-MP-#ID (sequence number)

MNP – SG-CFS-MP-# configuration

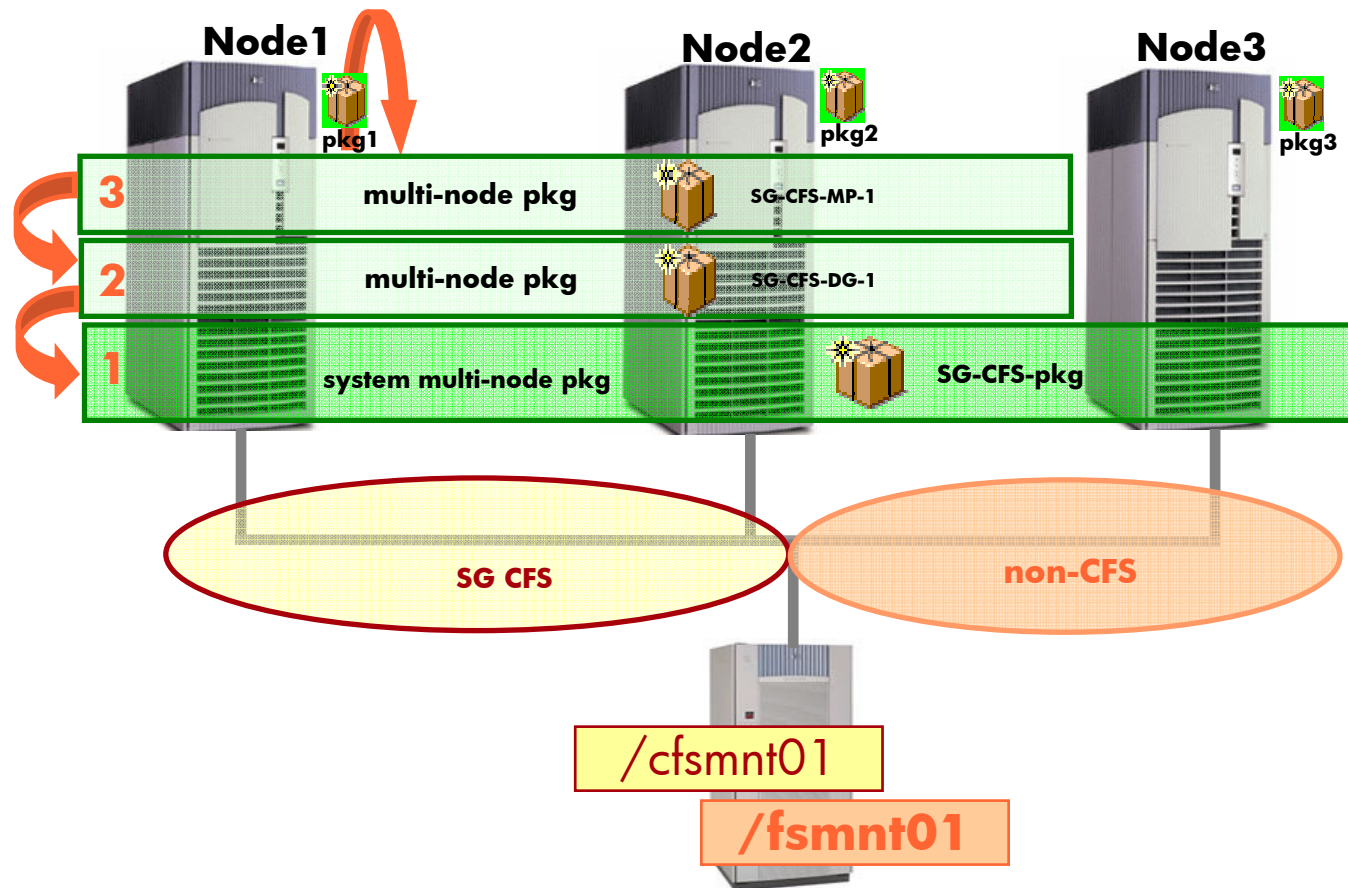


```
PACKAGE_NAME          SG-CFS-MP-1
PACKAGE_TYPE          MULTI_NODE
NODE_NAME             chives
NODE_NAME             curry
AUTO_RUN              YES
LOCAL_LAN_FAILOVER_ALLOWED YES
NODE_FAIL_FAST_ENABLED NO
RUN_SCRIPT             /etc/cmcluster/cfs/sgcfspkg.sh
RUN_SCRIPT_TIMEOUT    NO_TIMEOUT
HALT_SCRIPT           /etc/cmcluster/cfs/sgcfspkg.sh
HALT_SCRIPT_TIMEOUT   NO_TIMEOUT
SCRIPT_LOG_FILE       /etc/cmcluster/cfs/SG-CFS-MP-1.log
DEPENDENCY_NAME       SG-CFS-DG-1
DEPENDENCY_CONDITION  SG-CFS-DG-1 = UP
DEPENDENCY_LOCATION   SAME_NODE
CFS_MOUNT_POINT       /cfs/mnt1
CFS_VOLUME            cfsdgl/voll
CFS_MOUNT_OPTIONS     "chives=rw curry=rw"
CFS_PRIMARY_POLICY    ""
```

SG package concepts



SG package concepts





HP Serviceguard CFS use cases

Module 6



What applications might use SG/CFS?



Oracle RAC (9i or 10g)

- Near-raw volume performance with manageability of a file system
 - Oracle Database Accelerator (ODM - Oracle Disk Manager)

Intra-cluster application file sharing

- Multiple web servers share same data
- Enhance scalability and manageability
- Enhance HA NFS server
 - Improve scalability by enabling all nodes to serve NFS concurrently
 - Enhance HA NFS server by removing "cross mounts"
- Ability to create custom parallel applications

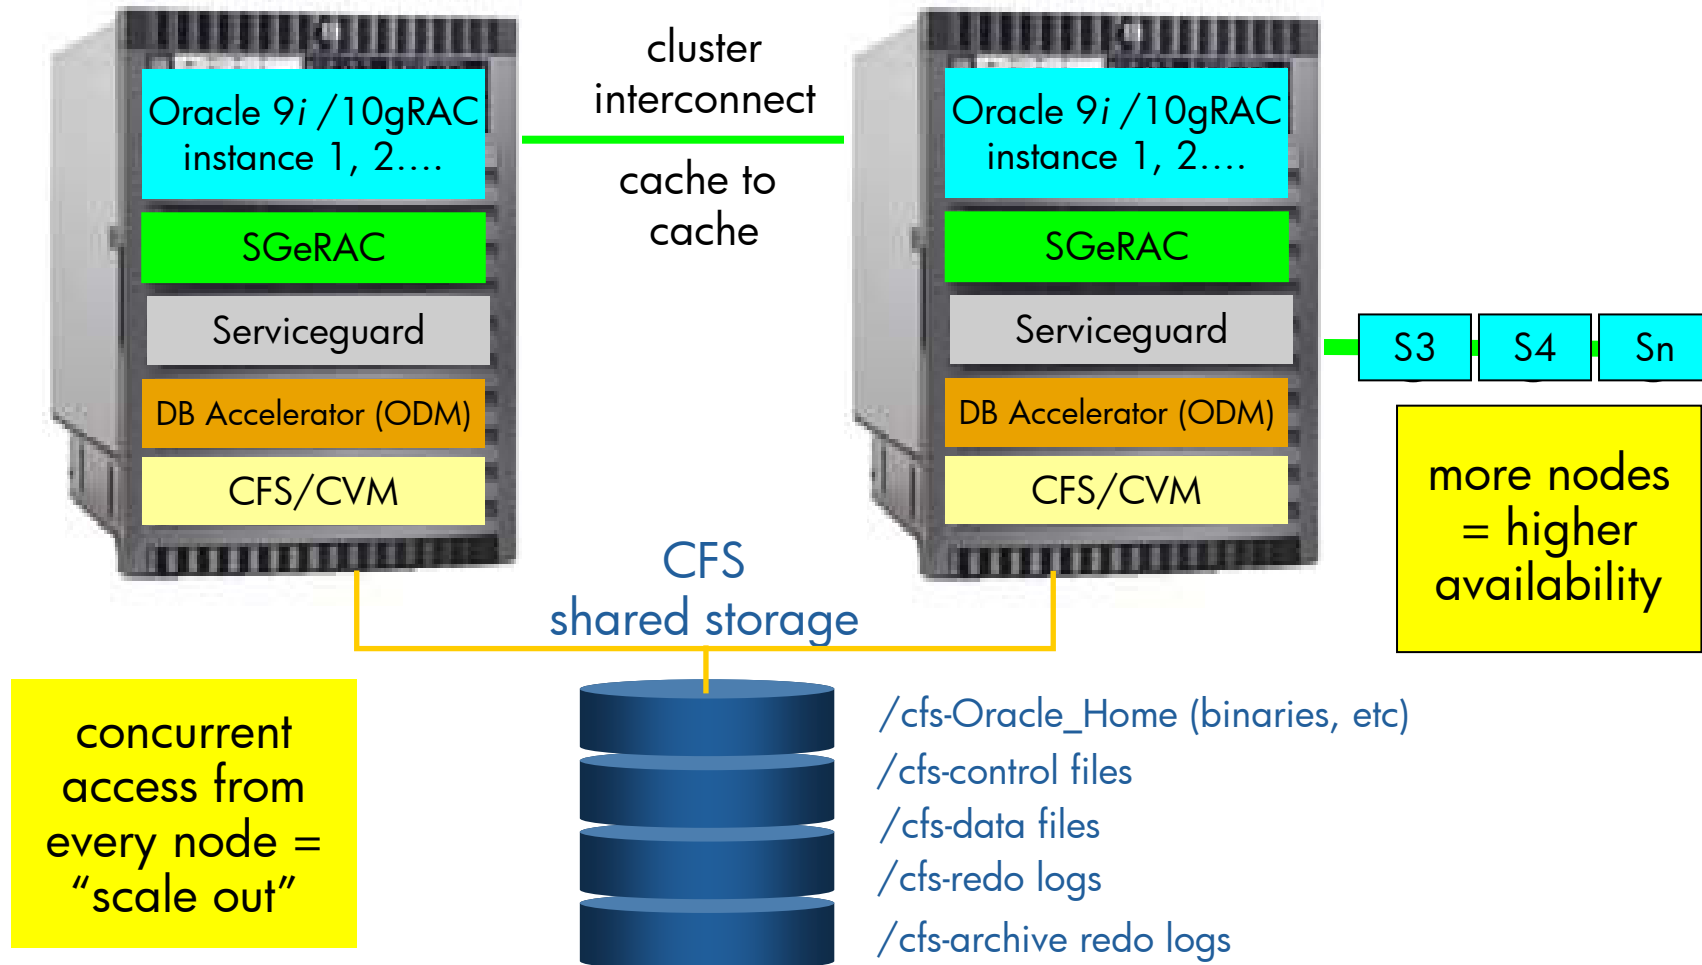
Applications with high number of volumes/filesystems

- Availability of volumes and file systems to all nodes

Oracle9i/10g Real Application Cluster (RAC) – using CFS shared storage



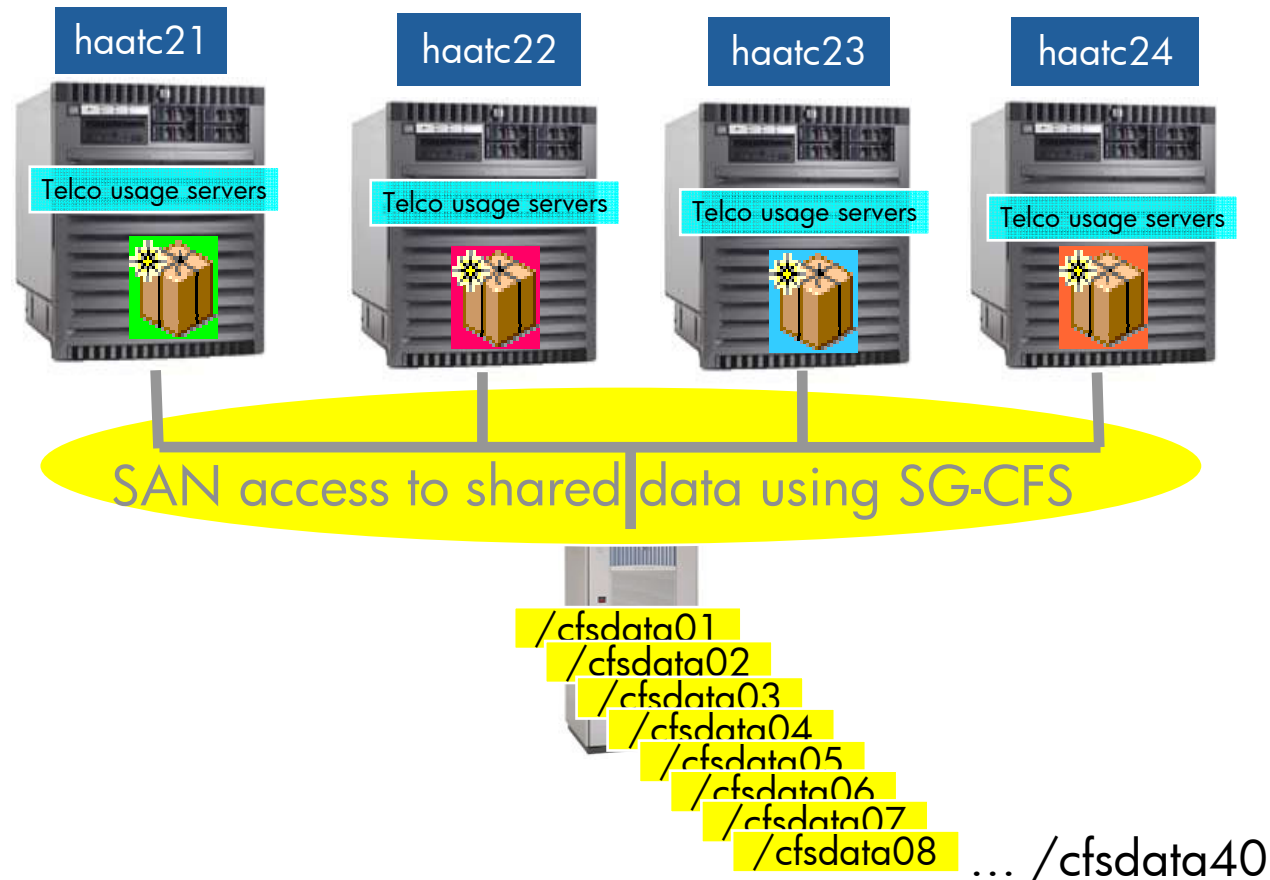
server 1



Applications with high number of volumes/file systems



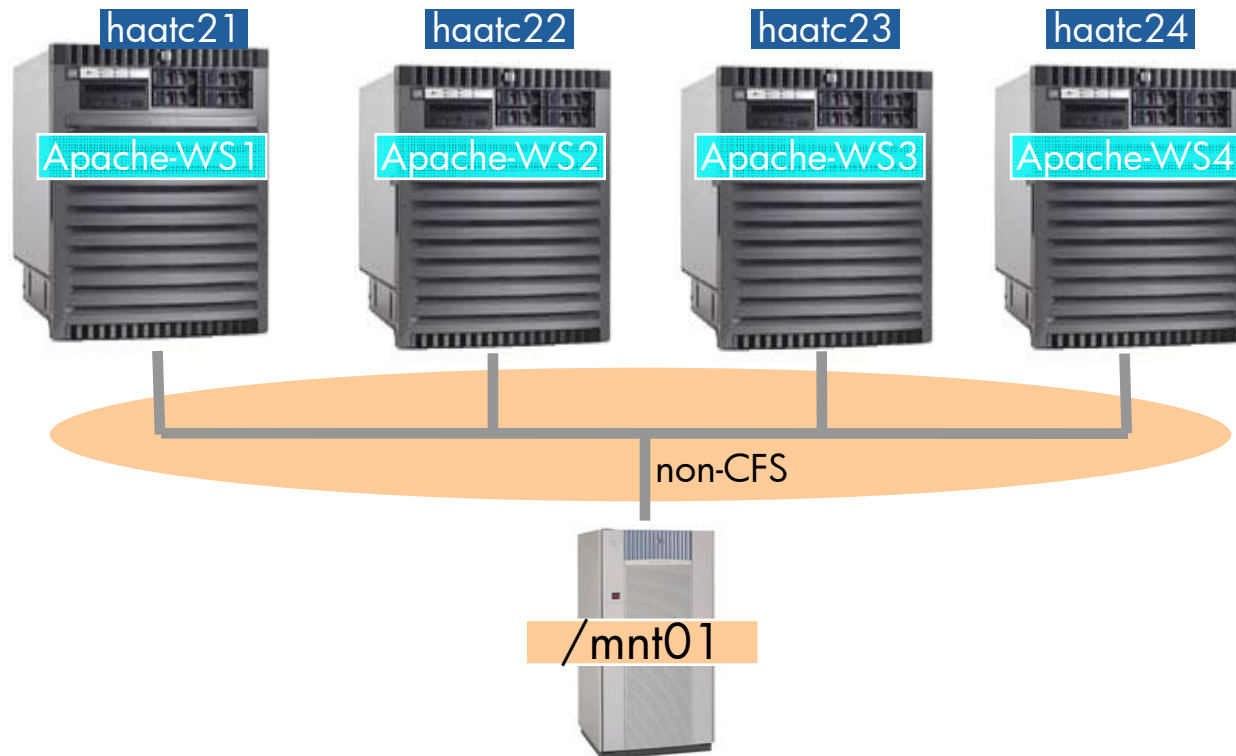
- Four node cluster
 - HA telco usage application failover package
 - CFS
 - faster failover time



Serviceguard non-CFS Web Server Farm



Intra-cluster application file sharing



Slide 57

REC1

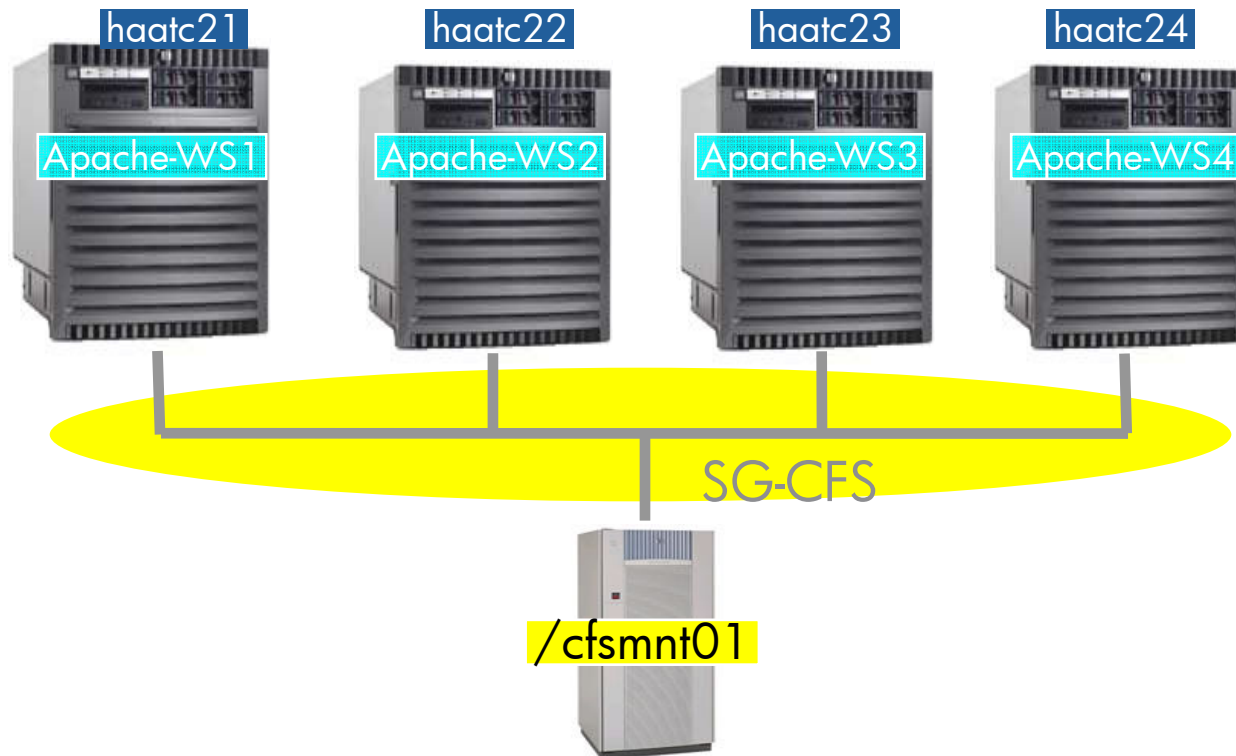
Need help here, this is to show HA NFS without CFS

Randi E. Constable, 04/29/05

Serviceguard CFS Web Server Farm



Intra-cluster application file sharing




Slide 58

REC2

Need help here, this is to show HA NFS with SG CFS

Randi E. Constable, 04/29/05





i n v e n t