


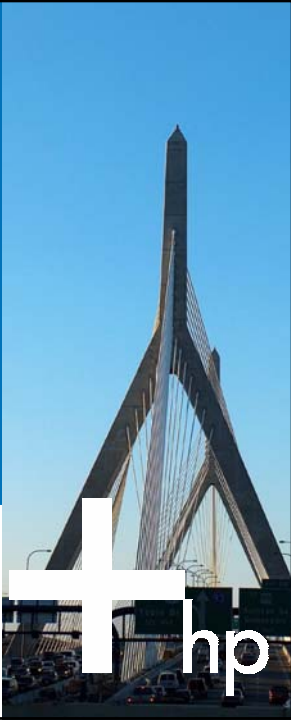
HP UX / Tru64 UNIX® Side by Side

Cluster Comparison & Roadmap Update

Decus Symposium 2005 April 2005

Ludwig Frohnsbeck
HP Services
Ludwig.Frohnsbeck@hp.com


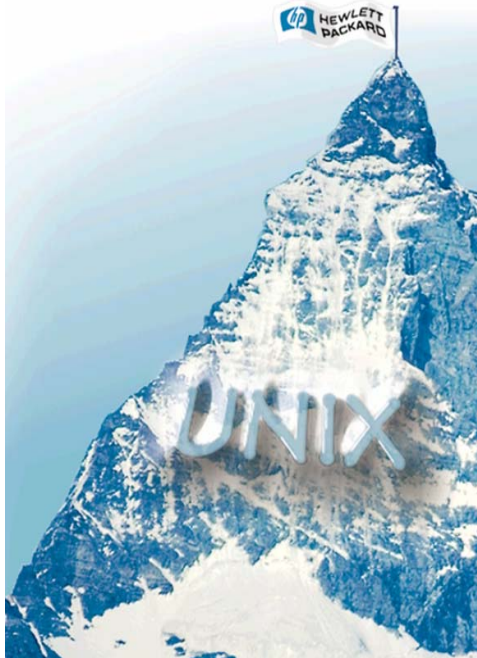
© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice



Agenda


- Product roadmaps and history
- Product design and principles
- Product Portfolio and Solutions
- Cluster Subsystems
- Virtualization
- Clustering roadmap Update

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 2



TruCluster & MC/ServiceGuard release history and roadmaps

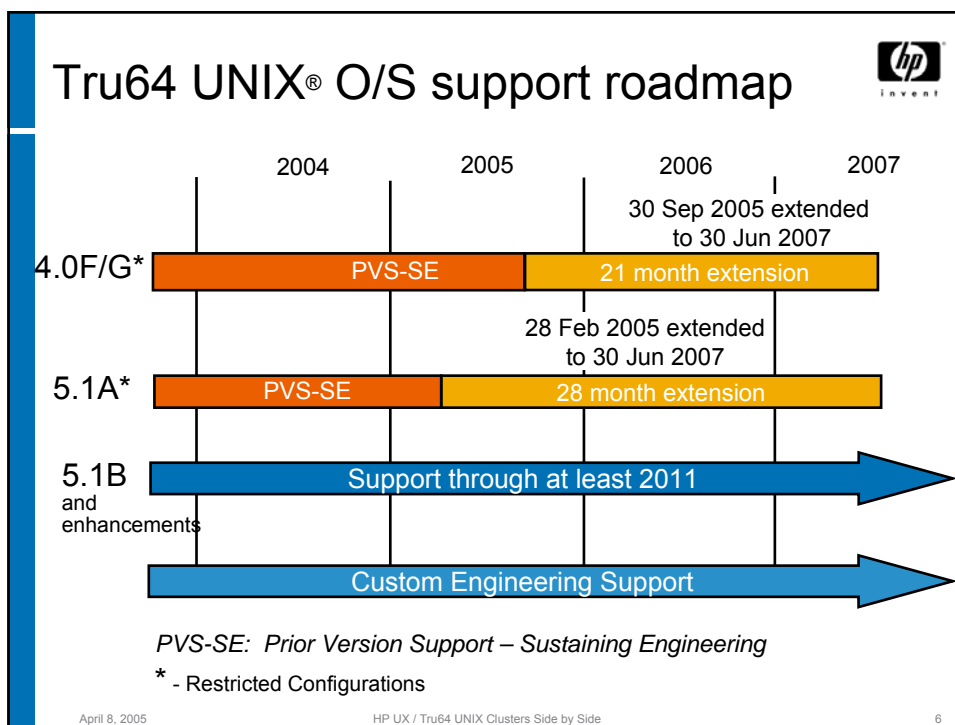
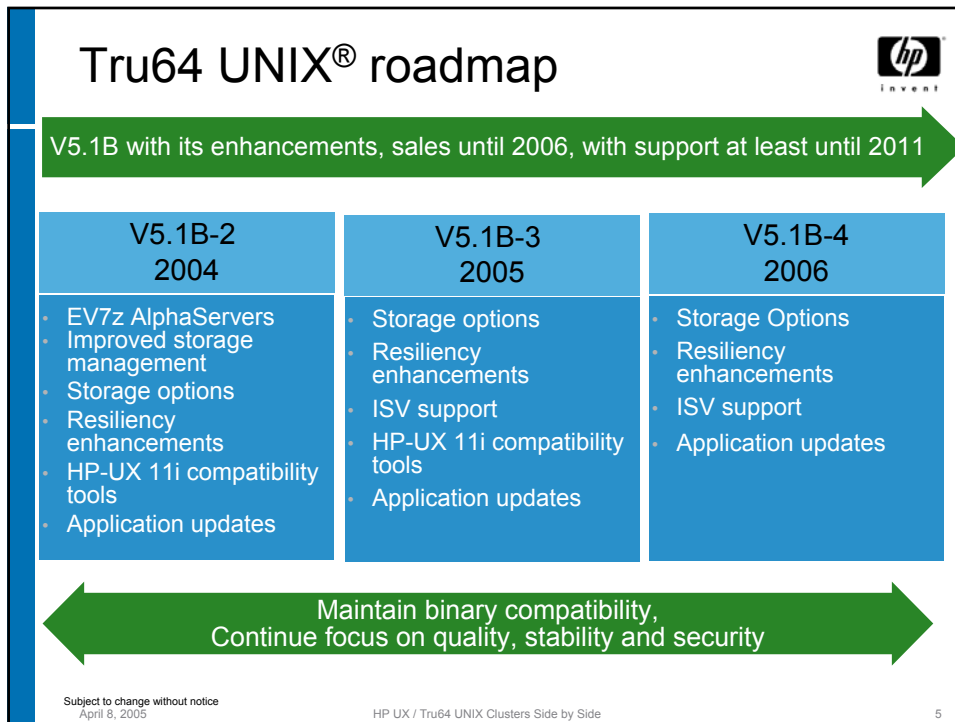
April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 3




TruCluster V5.x release history

- **V1.x** based products (starting in early 90's)
 - initially simple failover product (DECsafe)
 - evolved to TruCluster V1.x product suite (Production Server, Available Server and MEMORY CHANNEL cluster products)
- **V5.0** (Steel) — Q3CY99
 - limited release
- **V5.0A** (Zinc) — Q2CY00
 - general release, included FC support
- **V5.1** (Zulu) — Q3CY00
 - GS80/GS160/GS320 platform support, CFS enhancements
- **V5.1A** (Yankee) — Q3CY01
 - ES45 platform support, increased low-end cluster support
- **V5.1B** (Wildcat.Alpha) — H2/CY02
 - New Generation HP AlphaServer platform support, CFS load balancing, CAA performance improvements
 - TCR version coincides with base OS version

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 4

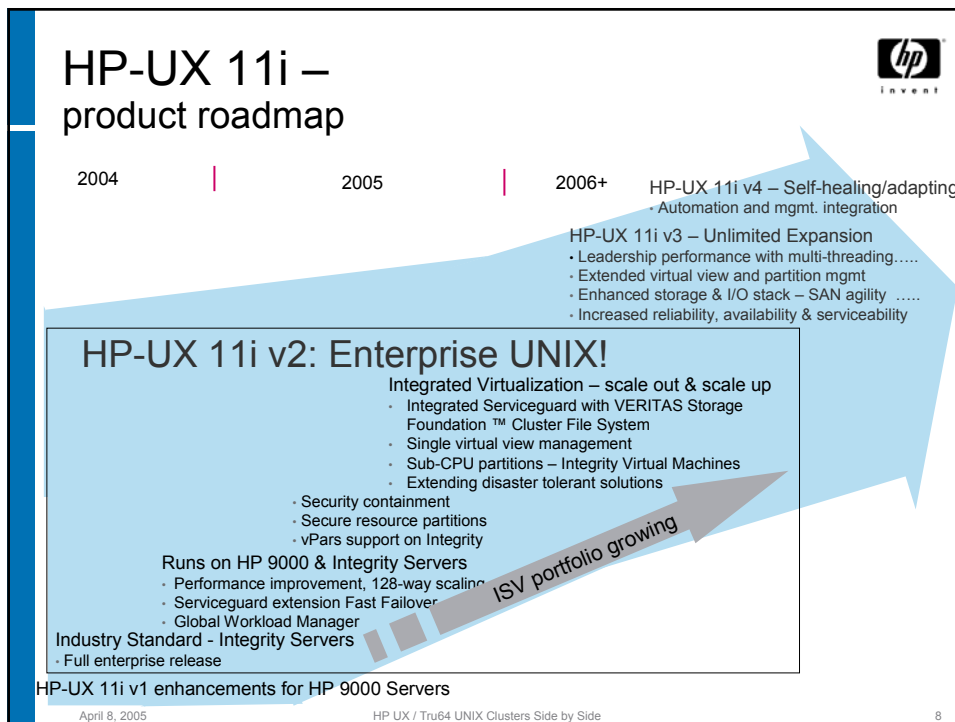





ServiceGuard Release History

OS	SG	Features
10.01	A.10.03 06/95	4 Member, Basic functionality (restart of apps), SAM GUI
10.10	A.10.04 12/95	8 Node Support, RS323 Link for heartbeat, Security Enhancements
	A.10.05 02/96	End of Support March 31 st 2002
10.20	A.10.06 06/96	End of Support December 31 st 2001
	A.10.10 10/97	Online administration, EMS Support
	A.10.11 12/98	Support for EMS V3.0 framework and monitors
10.30	A.10.12 12/99	Platforms Release , misc bug fixes
	A.10.08 N/A	Only for specific customer
11.Xx	A.11.01 02/98	Rotating Standby Feature. New Failover Policies.Up to 200 IP Aliases
	A.11.03 08/98	Up to 16 nodes. New cmquerycl options. EMS Support
	A.11.04 12/98	Advanced Tape Services
	A.11.05 02/99	Support for APA (Auto Port Aggregation)
	A.11.07 06/99	Platform Release. Limited Support for VxVM
	A.11.08 09/99	Not supported on HP-UX 11i. Initial release for SG Manager
	A.11.09 12/99	Platform Release. Full support for VxVM/CVM on 11i. Max pkg 60
	A.11.12 12/00	Platform Release. Quorum Server. Parallel fsck/VG. Max pkgs to 200
	A.11.13 09/01	
	A.11.14 03/02	

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 7






Product design and principles

- Architecture
- Sample Configurations


April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 9



Architecture

- shared nothing
 - independent nodes
 - no data sharing
 - each node must be managed independently
- shared storage connectivity and access
 - sharing of resources such as data and storage
 - locking needs to be used to coordinate shared data access
 - every member can access the same set of data/storage
- SSI (single system image)
 - cluster manages like a single system
 - most management tasks can be issued from any member


April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 10



Architecture

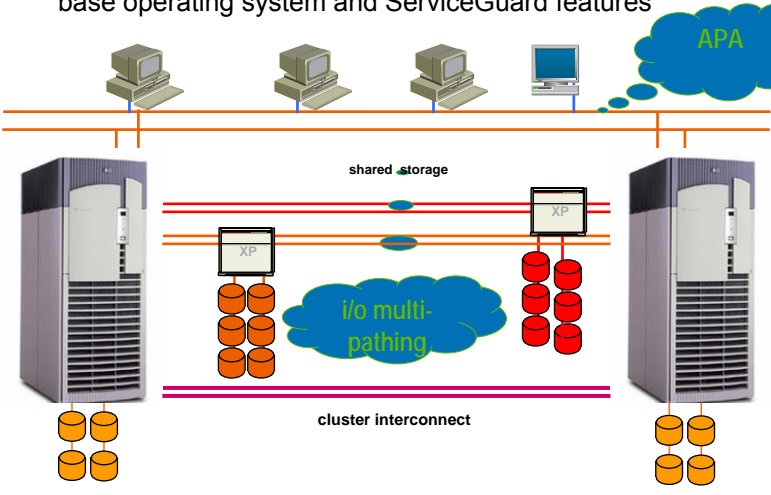
- both products provide highly available clustering solutions
 - applications can be active/standby or active/active
- TruCluster Server implements a shared everything architecture (SSI)
 - single system disk and root filesystem
 - clients and sysadmins see the cluster as one 'big' node
- MC/ServiceGuard implements a shared storage model
 - only one member can access a storage device at a time
 - special rules for Oracle9i RAC and Oracle8i OPS configurations
 - local storage for system disks and "private" storage

April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
11




MC/ServiceGuard configuration

achieve no single point of failure configuration by using base operating system and ServiceGuard features

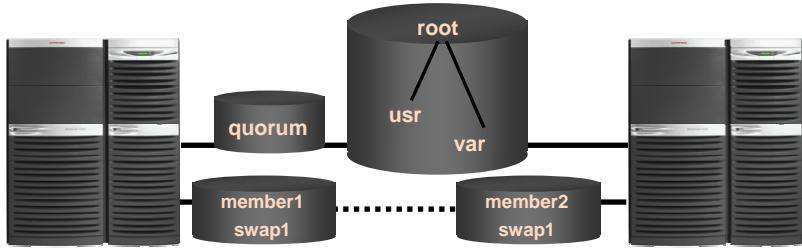


April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
12

TruCluster configuration




cluster file system, shared root, single system image



The cluster file system is shared by all cluster nodes
All nodes see the same data
The cluster manages as a single system, even as you add nodes

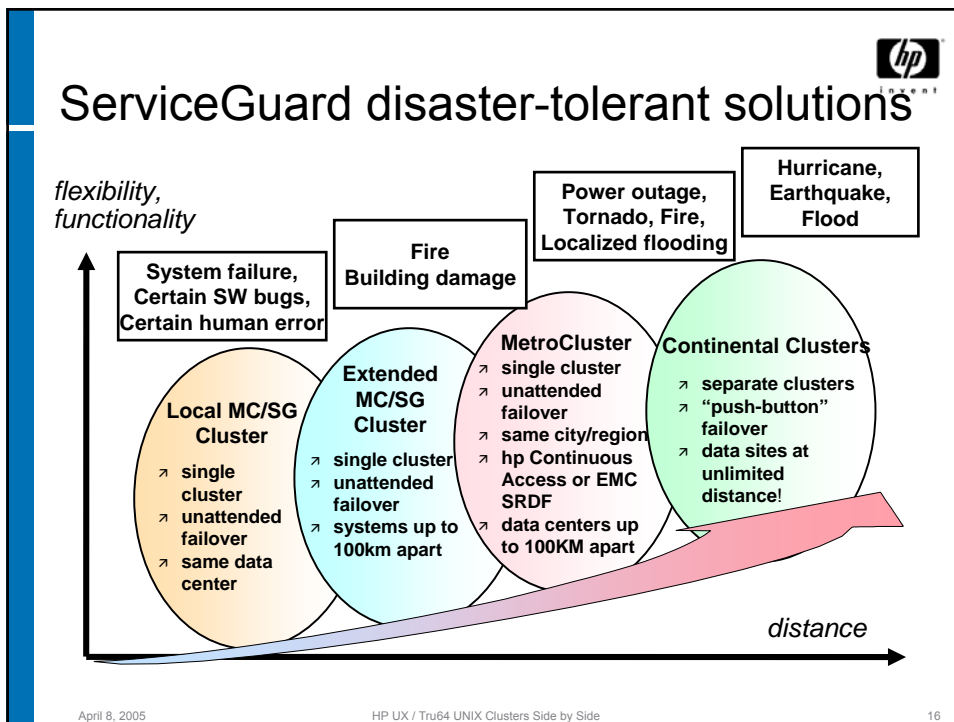
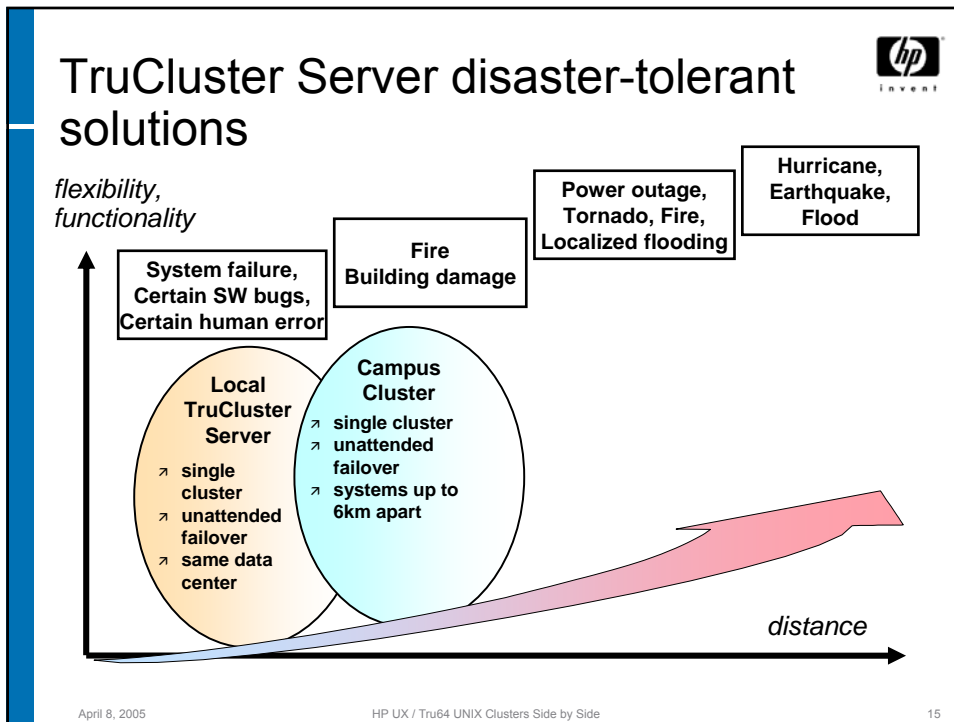
April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 13

Product Portfolio and Solutions



- High Availability Products
- Campuswide Cluster
- Disaster Tolerant Cluster

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 14





Cluster subsystems

- Hardware Requirements
- I/O Infrastructure
- Volume Managers
- Application Availability
- Cluster Administration

April 8, 2005

HP UX / Tru64 UNIX Clusters Side by Side

17



Hardware requirements and configuration

- MC/ServiceGuard supports both the PA RISC and Itanium® platforms
- TruCluster Server supports all AlphaServer models
- you can configure a cluster using standard hardware
 - no need for a specialized cluster interconnect
 - LAN used for heartbeat (MC/ServiceGuard) and cluster interconnect (TruCluster V5.1A)
 - high-speed cluster interconnects are available
 - HyperFabric used for message passing for MC/ServiceGuard Oracle9/ RAC
 - still need LAN for heartbeat
 - Memory Channel for TruCluster
 - SCSI and Fibre Channel storage are supported
 - large number of supported devices and adapters

April 8, 2005


HP UX / Tru64 UNIX Clusters Side by Side

18

<h3>MC/ServiceGuard I/O access and infrastructure</h3> <ul style="list-style-type: none">• shared storage model for data disks• I/O multipath through base OS options (licensed)<ul style="list-style-type: none">– AutoPath VA and XP– Veritas VxVM DMP– LVM PVlinks (no license req.)• file systems are mounted on individual members<ul style="list-style-type: none">– cannot be shared amongst member systems– can use NFS to cross mount• each node has its own system disk<ul style="list-style-type: none">– support for multiple concurrent versions– rolling upgrade support	<h3>TruCluster Server I/O access and infrastructure</h3> <ul style="list-style-type: none">• shared storage mode for all data and disks• I/O multipath through base OS and TCR<ul style="list-style-type: none">– integrated into operating system– no third-party or layered product used– use of DRD (device request dispatcher) within the cluster• all file systems are shared amongst all members<ul style="list-style-type: none">– option to enable member-specific access– create member-specific files and directories using CDSL• single-system disk and directory tree<ul style="list-style-type: none">– support for multiple concurrent versions (up to two)– rolling upgrades supported
---	--

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 19

drd (Device Request Dispatcher)



- provides transparent, highly available access to all disks and tapes in the cluster
 - device names are consistent using Tru64 UNIX® V5 naming mechanism
- used for all disk access in a cluster
- LSM, file systems, databases, applications
- supports block and character devices
- provides direct access and served I/O models
- **drdmgr** command for status and management
- no equivalent in MC/ServiceGuard today

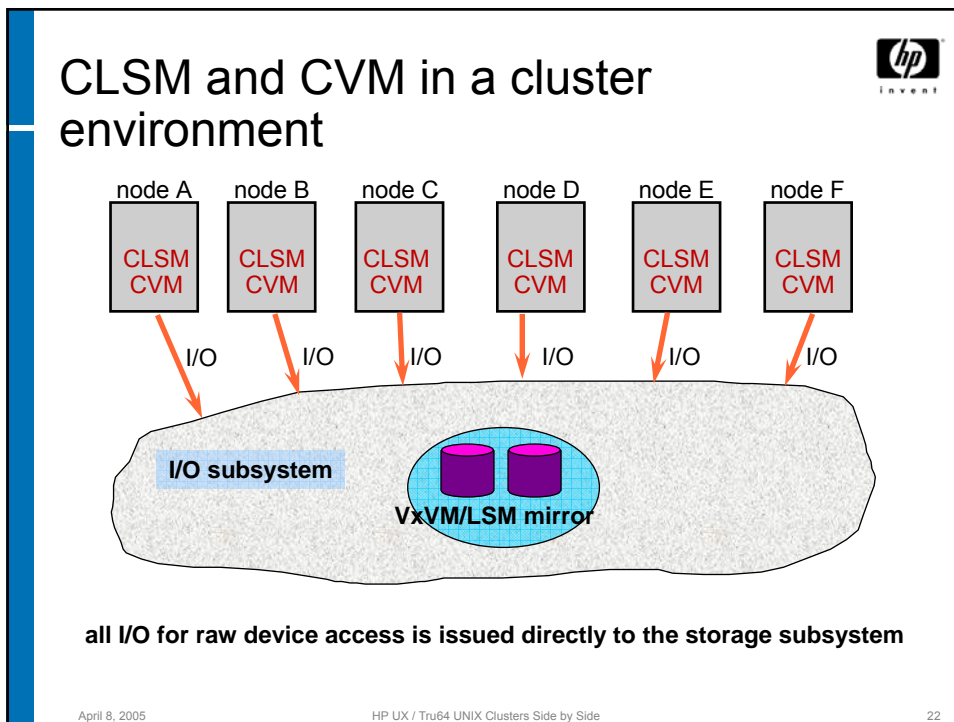
April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 20

logical storage/volume managers


- MC/ServiceGuard
- cluster support is available through CVM or SLVM
 - separately licensed
- shared access to raw volumes
 - used by Oracle9i RAC and Oracle8i OPS
- raw device only support
- nodes must enable a given disk group for clusterwide access
- same management commands as on a standalone system
- Some commands have to be done on every node separately

- TruCluster Server
- cluster support is built into LSM
 - CLSM code enabled if running in a cluster
 - license for mirroring and GUI
- shared access to LSM volumes
- raw device and file systems
- configuration changes can be made from any cluster member
- fully symmetric design
- same management interface as LSM on a single system

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 21




file system access in a cluster



- MC/ServiceGuard
 - an individual file system is mounted by one member in the cluster
 - no clusterwide mounts and file system
 - support for VxFS, HFS, NFS, CDFS and lofs
 - NFS support for client and server access
 - a cluster node can be an NFS client to another member in the same cluster
 - allows sharing of file systems
- TruCluster Server
 - cluster file system (CFS) mounts each file system clusterwide
 - each member can see and (potentially) access every mounted file system
 - additional options to restrict access to the CFS “server”
 - support for AdvFS, UFS, NFS, MFS, and DvDFS
 - cache coherent clusterwide
 - including mmap()'d files
 - NFS support for client and server

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 23


Application availability



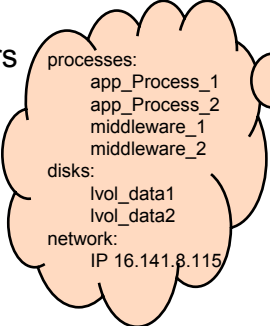
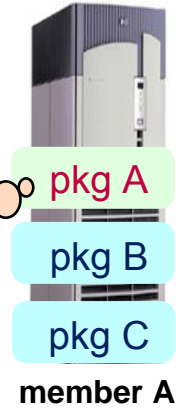
- both clustering solutions provide a comprehensive failover framework
 - MC/ServiceGuard and cm (cluster monitor)
 - TruCluster Server uses CAA (cluster application availability)
- applications can be integrated through scripts and APIs
- extensive collection of example scripts and third-party application integration (MC/ServiceGuard)
- failure recovery either on the local node or by moving the application to another (functional) node in the cluster
- can have dependencies between applications to force ordered startup/shutdown/relocation
 - integrated in TruCluster Server via CAA
 - separate scripting toolkit with MC/ServiceGuard

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 24

MC/ServiceGuard application packages



- easy to configure framework
- flexible and easy to manage application resources
- up to 16 node clusters
- up to 200 application packages




processes:
app_Process_1
app_Process_2
middleware_1
middleware_2

disks:
lvol_data1
lvol_data2

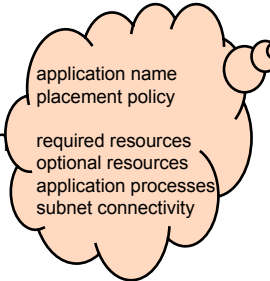
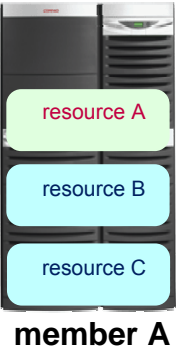
network:
IP 16.141.8.115

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 25

TruCluster CAA application resources



- easy to configure and extensible framework
- central management for applications in a TruCluster
- up to 8 node clusters
- hundreds of application resources



application name
placement policy

required resources
optional resources
application processes
subnet connectivity

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 26

application failover frameworks


- **MC/ServiceGuard**
- packages
- assign IP address to package
- storage configuration kept in ServiceGuard configuration files
- provides a CLI (`/usr/sbin/cm*`) and graphical user interface
- extensive cluster toolkit in ECMT (Enterprise Cluster Master Toolkit)
 - databases and applications
 - NFS
- integrated SAP R/3 solution with SGeSAP Toolkit
 - optionally with Sommersault MC/SGeRAC
- Oracle9i RAC on ServiceGuard
- **TruCluster CAA**
- resources
- can use cluster alias instead of IP alias (no CAA management needed)
- no storage associated with CAA resources
 - storage failover/access part of base OS and TruCluster
- CLI (`/usr/sbin/caa_*`) and GUI through SysMan station
- example scripts for many applications provided as part of the CAA framework
- Oracle9i RAC integration with hp DButility

April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 27

Cluster administration

- cluster installation and addition of cluster members
- both solutions provide a framework to manage the cluster and its services
- can run on multiple platforms
 - Linux, UNIX® (HP-UX, Tru64)
 - Windows®
 - web-based interface available
- TruCluster extends single system commands to cluster
 - all management actions can be launched from SysMan
- administration frameworks
 - ServiceGuard Manager
 - System Management Station


April 8, 2005 HP UX / Tru64 UNIX Clusters Side by Side 28



CAA/ServiceGuard CLI commands

cmviewcl	caa_stat	provides status on the current state of the cluster members and services
cmmodpkg	caa_profile	manages resource/package attributes
cmmakepkg	caa_register	registers a resource/package
cmrunpkg	caa_start	starts a resource/package
	caa_relocate	relocates an application
cmhaltpkg	caa_stop	stops a resource/package
	caa_unregister	removes an application from CAA control

April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
29



Cluster interconnects

- **MC/ServiceGuard**
- LAN for heartbeat and data/client access
 - transparent failover using APA
- high-speed messaging interconnect using HyperFabric
 - up to 2.4Gb/s bandwidth
 - can be shared between clusters
 - transparent failover
 - fiber and copper hardware
 - used in Oracle9i RAC and Oracle8i OPS configurations
 - still require LAN for heartbeat
- HMP for high-speed message passing
 - no adapter failover with HMP in HyperFabric

- **TruCluster Server**
- LAN for cluster communication and client access
 - transparent failover through NetRAIN
- High-speed interconnect using Memory Channel
 - 100MB/s
 - very low latency
 - transparent failover built in
 - specialized API
 - fiber and copper hardware
- RDG for high-speed message passing
 - failover built-in via CI

April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
30



Networking client access

- both products provide network failover for clients
- they use IP aliases assigned to interfaces
- IP alias moves with application
- can monitor for connectivity to IP address/subnet
- in **ServiceGuard**, part of the package configuration process
- in **TruCluster Server**, use cluster alias or IP alias as part of the resource definition
 - cluster alias can be configured:
 - as virtual IP address
 - to provide support for single and multi-instance services

April 8, 2005

HP UX / Tru64 UNIX Clusters Side by Side

31




Clustering Update

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice



HP-UX 11i – Proven Foundation for the Adaptive Enterprise



2004
2005
2006+

HP-UX 11i v4 – Self-healing/adapting

- Automation and mgmt. integration

HP-UX 11i v3 – Unlimited Expansion

- Leadership performance with multi-threading.....
- Extended virtual view and partition mgmt
- Enhanced storage & I/O stack – SAN agility
- Increased reliability, availability & serviceability

HP-UX 11i v2: Your Enterprise UNIX!

Integrated Virtualization – scale out & scale up

- Integrated Serviceguard with VERITAS Storage Foundation™ Cluster File System
- Single virtual view management
- Sub-CPU partitions – Integrity Virtual Machines
- Extending disaster tolerant solutions

- Security containment
- Secure resource partitions
- vPars support on Integrity

Runs on HP 9000 & Integrity Servers

- Performance improvement, 128-way scaling
- Serviceguard extension Fast Failover
- Global Workload Manager

Industry Standard - Integrity Servers


- Full enterprise release

HP-UX 11i v1 enhancements for HP 9000 Servers

ISV portfolio growing

April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
33

VxFS v4 compared to no longer planned AdvFS on HP-UX 11i (1 of 4)



Feature	VxFS v4	AdvFS**
Log Based File System - Fast file system recovery	✓	✓
Max file size – tested / arch	16TB / 8EB (2TB on HP-UX today) 32TB / 8EB	20TB / 8EB (16TB on Tru64 Today)
Max file system size – tested / arch	(32TB on HP-UX today)	512TB / 8EB (16TB on Tru64 Today)
Extent-based allocation	✓	✓
Multiple Volume Aware File Systems	✓ VxVM volume sets (4.0)	✓ domains
On-line balance and migrate	✓	✓
Automated placement policies and quality of storage service (QoS)	✓	✗


Legend comparison table

✓	Has or is/was a planned capability for VxFS v4 or AdvFS on HP-UX respectively
(4.0)	release introduced
✗	Does not have or is/was not a planned capability for VxFS v4 or AdvFS on HP-UX respectively

*as planned for HP-UX in Q2CY05 as an update on the 11i v2 release

** as no longer planned for HP-UX as part of the EndCY05 11i v3 release


April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
34



VxFS v4 compared to no longer planned AdvFS on HP-UX 11i (2 of 4)

Feature	VxFS v4*	AdvFS**
Online Point-in-time File System Copies	✓ snapshots/ checkpoints (3.5)	✓ clones (was to be renamed snapshots in HP-UX version)
Persistent copies	✓ checkpoints	✓
Writeable copies	✓	✗
Metadata only copies to support block level incremental backup (BLIB)	✓	✗
Quotas Clean-up policies on copies	✓ (4.0)	✗
Integrated administrative interface for previous file or file system version recovery from copies	✓ (4.0)	✗
Hw copy support – Freeze/Thaw interface	✓	✓
Heterogeneous Data Portability and Access	✓ cross platform data sharing (cds) (4.0)	✗
On-line growth	✓	✓
Underlying volume	✓ expandfs, fsadm -b	✓ mount -o
Volume addition	✓ vxvset - addvol,rmvol (VxVM)	✓ addvol, rmvol


April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
35



VxFS v4 compared to no longer planned AdvFS on HP-UX 11i (3 of 4)

Feature	VxFS v4*	AdvFS**
I/O Performance Enhancements and Optimizations	✓	✓
Log file isolation	✓ quicklog (3.5) / volume sets (4.0)	✓
Small File Support; space efficient and high performance	✓	✓
Fast directory operations	✓ hashed	✓ b-tree indexes-fast lookup algorithm
Asynchronous I/O option	✓	✓
Local Filesystem Cache bypass option (Direct IO)	✓	✓
CFS Direct Path	✓	✓


April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
36




VxFS v4 compared to no longer planned AdvFS on HP-UX 11i (4 of 4)

Feature	VxFS v4*	AdvFS**
Online administration and tuning	✓ GUI: fsweb or VEA from VERITAS	✓ GUI: FS Web
Online defragmentation	✓ fsadm	✓ defragment
Continuous automated defragmentation	✗	✓ vfast
File Recovery	✓ checkpoint based	✓ trashcan
Standard and Specialized Programming Interfaces	✓	✓
POSIX, X/OPEN application programming interface	✓	✓
Named data streams	✓(4.0)	✗
File change log API	✓(4.0)	✗
Miscellaneous		
Access control list (ACL) support	✓	✓(property lists)
Log file on-line resize	✓	✗
DMAPI	✓	✗

April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
37

- 
- ### What has changed ?
- Change from TCR/CFS to VERITAS Cluster Filesystem (CFS)
 - Serviceguard continues to be the strategic clustering solution and will be integrated with VERITAS CFS
 - Enhanced management tools for multi system and multi cluster management
 - Cross platform management for config files, kernel tuning, merged logfile view etc...
 - Accelerated virtualization features and products
 - HP VM, Secure Containments for HP-UX 11i, VSEmgr, gWLM
 - “VSE features replace and extend SSI components”
 - Except for shared root and cluster alias
 - New and enhanced tools to facilitate management of individual cluster members and a cluster as a whole
 - Unified processes and tools to manage VSE components and systems
- April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
38


BaseOS Component Review



Tru64 UNIX® Component	HP-UX 11i Component
Event Management (EVM)	Event Management (EVM)
LAG (Link Aggregation)	APA (Auto Port Aggregation)
NetRAIN	APA (Auto Port Aggregation)
NIFF	LAN Monitor
I/O Multi Pathing	HP LVM plinks, VERITAS DMP, HP Securepath
Location Independent Device Names	Location dependent device names, Volume Manager for abstraction
Volume Manager (LSM)	VERITAS VxVM, HP LVM

April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
39


Cluster Component Review



TruCluster Component	HP-UX 11i Serviceguard Component
Cluster Application Availability (CAA)	Serviceguard Package Manager
Cluster Filesystem (CFS)	VERITAS Cluster Filesystem
Single System Image/Shared Root (SSI)	single virtual view (SVV) and System Management Homepage (SMH) / no shared root
Distributed Lock Manager (DLM) Distributed Lock Manager API	Need to redesign locking portion of application
MEMORY CHANNEL API (MC/API)	Need to redesign MC/API specific code
Cluster Alias (CLUA)	IP Alias / no Cluster Alias
Device Request Dispatcher (DRD)	VERITAS CVM / HP SLVM and SAN storage
EVM (Cluster Wide Events)	EVM (cluster wide view via SIM)
Connection Manager (CNX)	Serviceguard Membership Manager
Clusterized Volume Manager (CLSM)	VERITAS CVM, HP SLVM

April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
40

Extending the HP Serviceguard Portfolio



Integrated bundles including VERITAS software

Available on HP-UX 11i v2 for HP Integrity & 9000 Servers, planned for Q3 2005

HP Serviceguard
and
VERITAS Storage Foundation™
(volume manager, file system, management tools)

HP Serviceguard
and
VERITAS Storage Foundation™
Cluster File System

HP Serviceguard
and
VERITAS Storage Foundation™
for Oracle Environments


HP Serviceguard
and
VERITAS Storage Foundation™
Cluster File System for Oracle

HP Serviceguard
and
VERITAS Storage Foundation™
Cluster File System for Oracle RAC

Delivered and supported by HP

April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
41

Tru64 to HP-UX Migration Solutions*




<i>Existing Tru64 Solutions</i>	<i>Equivalent HP-UX Solutions</i>
<i>LSM</i>	<i>VxVM + HP File Mover Tool</i>
<i>AdvFS</i>	<i>VxFS + HP File Mover Tool</i>
<i>AdvFS & LSM</i>	<i>Storage Foundation + HP FMT</i>
<i>TruCluster (v4 and v5)</i>	<i>HP Serviceguard + Storage Foundation (v4) or Storage Foundation CFS (v5)</i>
<i>TruCluster for Oracle customers (v4 and v5)</i>	<i>HP Serviceguard + Storage Foundation (v4) for Oracle or Storage Foundation CFS for Oracle (v5)</i>
<i>TruCluster for Oracle RAC customers (v5)</i>	<i>HP Serviceguard + Storage Foundation CFS for Oracle RAC</i>

*Customers under service support contract on 11/19/04 will receive a comparable set of licenses

April 8, 2005
HP UX / Tru64 UNIX Clusters Side by Side
42

HP-UX 11i v2 environments

What Tru64 UNIX® customers can expect ---- key areas



	Tru64 UNIX V4	Tru64 UNIX V5	Tru64 UNIX V5 using CFS
HA and DT clusters	✓ PLUS : Today Serviceguard and Metrocluster		≈ Comparable in Q3 2005 Serviceguard integrated with VERITAS Storage Foundation™ Cluster File System <small>May require a different approach</small>
Supported storage	≈ Fibre channel devices are the same: HSG80*, MSA, EVA, XP <small>Different approaches may be required to align preferred volume manager, multipathing and storage array.</small>		
File systems and volume management	✓ PLUS File System: OnlineJFS (VERITAS VxFS) Volume Manager: VxVM		
Overall functionality	✓ PLUS virtualization, security, management, ISV applications		

* HSG80s are supported on HP-UX 11i v2 Integrity servers and in a Serviceguard cluster environment, with restrictions.
HP-UX / Tru64 UNIX Clusters Side by Side

April 8, 2005 43

