

Transitioning Tru64 UNIX V4 Customers to HP-UX 11i v2:

Analyzing the HP-UX11i Roadmap changes in file systems, clusters, and infrastructure



Overview	2
1. HP Virtual Server Environment (VSE).....	4
2. Tru64 UNIX V4 File Systems and Volume Management.....	7
2.1 File Systems.....	7
2.2 Volume Management.....	8
2.3 File System and Volume Management bundles	9
3. Clusters and Third Party ISV applications	10
3.1 Applications	10
3.2 Disaster-tolerant Solutions	12
4. Cluster Interconnect.....	12
5. Cluster Configurations	13
6. Custom Code Applications.....	13
7. Storage Implementations.....	14
8. Oracle Database Applications.....	15
8.1 Oracle Single-instance	15
8.2 Oracle High Availability.....	15
8.3 Oracle OPS/RAC.....	15
8.4 Raw Device I/O Options for Databases	17
8.5 File Systems I/O Options for Databases.....	18
9. Conclusions	18
References.....	19
For More Information	20

Overview

Businesses are dynamic – they must respond constantly to market challenges and maximize opportunities as they arise. At the same time, there is the pressure to reduce costs and increase operational efficiency. Today's computing environment is often too complex to be easily managed and responsive to change. In addition, they often require the ability to manage, integrate and utilize different resources from different platforms and operating systems. To deal with this, customers have requested standardized methodologies to help them more readily manage their dynamic computing resources and applications.

The Adaptive Enterprise is HP's vision of an organization where business and IT are synchronized to capitalize on change. An Adaptive Enterprise with integrated and shared computing resources reduces operational costs by simplifying management tasks and optimizing resource utilization. An essential step in becoming an Adaptive Enterprise is to virtualize your IT resources. Virtualization is an approach to IT that pools and shares your IT resources to optimize utilization and match supply to demand automatically.

According to IDC, 49% of companies today are likely to implement a virtualization project in the next 12 months.¹ By 2008, according to Gartner, enterprises that do not leverage virtualization technologies will spend 25% more annually for hardware, software, labor and space for Intel® servers and 15% more for RISC servers.²

Through virtualization, "scale-up" and "scale-out" environments can be used together and be optimized as part of a shared pool that spans servers and storage resources. To address this, HP has re-evaluated its planned product capabilities for its Enterprise class UNIX offering. To provide tighter alignment of clustering, partitioning, multi-platform environments, changes are being made to the HP-UX 11i roadmap. HP has decided not to deliver TruCluster Server and the AdvFS technologies on HP-UX 11i. Tru64 UNIX customers migrating to HP-UX 11i will instead use the VERITAS file system and volume management products and Serviceguard for their clustering solution.

This strengthening of the HP-UX 11i product roadmap is being made to meet the changing needs of HP's UNIX customers. HP is focused on delivering integrated virtualization, clustering, and management tools. HP is confident that this direction is the best long-term-solution for its UNIX customers. It unifies system and storage clustering and virtualization, and simplifies management, resulting in reduced overall TCO. HP plans to deliver this solution first on HP-UX 11i, then on Linux and other operating systems. It is a direction that builds a broad foundation for customers to choose their own path to becoming an Adaptive Enterprise.

This document is targeted at those customers on Tru64 UNIX V4 who are planning or designing an HP-UX 11i solution. Its purpose is to provide a complete assessment of the roadmap changes affecting Tru64 UNIX customers in key areas of potential differences; i.e. the product platform (file systems, clusters, etc) and how the product is deployed or used (Oracle database, application deployment, etc.). In most cases, comparable capabilities are provided that may require different approaches to accomplish the same tasks. In a limited number of cases, a modification of how the cluster or applications are managed may be required to address differences in the tools available. In several key areas, new functionality will be provided to Tru64 UNIX customers.

This document is organized into the following sections:

- HP Virtual Server Environment (VSE)
- File systems and volume management

¹ IDC's 2003 Black Book

² Gartner, "Predicts 2004: Server Virtualization Evolves Rapidly" by Tom Bittman, November 14, 2003

- Clusters and third party ISV applications
- Cluster interconnect
- Cluster configurations
- Custom code applications
- Storage implementations
- Oracle database applications

1. HP Virtual Server Environment (VSE)

HP has always recognized the importance of scale-up and scale-out environments. Both environments address key demands in the IT infrastructure, and the choice is typically driven through the specific application structure and application workload. As such, HP has invested in delivering both environments to our customers. Through virtualization (an approach to IT that pools and shares your IT resources to optimize utilization), scale-up and scale-out environments can be used together, be optimized as part of a shared resource pool that spans servers and storage resources, and managed under common management as part of HP's virtualization offering.

HP Virtual Server Environment (VSE), an integrated solution for both HP 9000 and HP Integrity server platforms, allows enterprises to achieve a greater return on their IT investments by optimizing server resource utilization on a real-time basis according to business priorities. Within a VSE, virtual servers automatically grow and shrink based on the service-level objectives (SLOs) set for each application they host. The design goal for HP VSE is to:

- Double resource utilization by dynamically allocating resources
- Maintain continuous service levels by combining simple policy management and high availability
- Pay only for what you use by integrating with HP's utility pricing portfolio

Today's focus for VSE is to optimize utilization in a scale-up environment, for example, by dynamically moving resources from one partition to another. Today, VSE supports scale-out capability mainly from a high availability perspective – whenever an application is re-directed by Serviceguard to another system or partition, resources are automatically adjusted in accordance with the business priorities and SLOs of all the applications running at that time on that system or partition.

HP's goal is to offer common management for scale-up and scale-out capabilities – to manage all virtualized resources in the same way, whether they are partitions, nodes in a cluster, or multiple clusters. These improvements will be integrated with HP Systems Insight Manager and related plug-ins to deliver one unified multi-OS platform management solution.

HP will deliver a single virtual view (SVV) functionality to increase ease of management and visualization. This virtual view is applicable to Serviceguard high availability clusters, Cluster File System or High Performance Technical Computing environments, as well as any group(s) of servers or partitions (hard partitions, virtual partitions, virtual machines, resource partitions). HP's single virtual view includes the following new functionality for cluster management:

- Auto synchronization
- Flexible simultaneous task execution across nodes in a cluster
- Centralized cluster management
- Auto-discovery of Single Points of Failure

These new cluster management capabilities integrate into the System Management Homepage (next generation System Administration Manager (SAM)), HP's out-of-the-box system management solution.

In addition, these new capabilities integrate into HP Systems Insight Manager (SIM) to support multiple clusters through the same management solution, allowing multiple groupings and variable levels of synchronization. This functionality is planned for both HP-UX 11i and Linux.

In addition to the new cluster management capabilities, single virtual view (SVV) will contain enhancements to visualization and configuration. System administrators can see pictures of their physical complexes, as well as how partitioning and clustering technologies are used within and across complexes. System administrators are able to easily view the relationships between workloads

and resources. This management ease is obtained by having a single tool to visualize the complete virtual server environment, configure the various VSE technologies and set SLOs for the applications.

Single virtual view (SVV) functionality is being designed into existing management tools for the HP Virtual Server Environment (VSE). This ensures a common look and feel as well as data sharing for seamless integration. This functionality will be part of the following tools: Serviceguard Manager, System Management Homepage and Systems Insight Manager. It complements server deployment tools such as Ignite-UX and Software Distributor (SD-UX). These tools are all integrated into HP Systems Insight Manager, letting customers take advantage of this functionality and employ HP SIM as the central point of administration for complete resource life-cycle management. In addition, HP will ensure consistent tools across HP-UX 11i and Linux, where applicable. For example, Serviceguard Manager, System Management Homepage (including key management applications like HP's new Disk and File System management utility, fsweb), and HP SIM will be the same tools across HP-UX 11i and Linux, providing consistency and ease of use.

System Management Homepage is the next generation SAM product that supports the out-of-the-box system management applications. SM Homepage will provide a web-based GUI and tighter integration of existing SAM functionality and event management capabilities. Most system management applications that integrate into the SM Homepage support a single OS image, cluster or complex. The SM Homepage framework is the same across HP-UX 11i, Microsoft Windows Server, and Linux, enabling the same system management application to support multiple operating systems. System management applications available from the SM Homepage are also integrated into HP SIM. Some key features of SM Homepage and the new suite of web-based system management applications include:

- Command preview
- Status information for system components (including the cluster)
- Integration between events and system configuration applications
- Greatly improved start-up and screen-to-screen performance

For HP-UX 11i v2 customers will have both SAM (currently available on v2) and System Management Homepage along with web-based system management applications such as pdweb (peripheral and device management) and kcwweb (kernel configuration management).

Single virtual view management enhancements for HP-UX 11i v2:

1. Auto-synchronization: Currently HP SIM provides a basic snapshot comparison capability in which the user can compare basic configuration information among multiple servers. SVV complements HP SIM's inventory capabilities by providing a configuration file management utility with the following features:
 - Compare the content of an HP recommended set of files across nodes and flag differences
 - Execute comparisons at regular intervals via the UNIX scheduler
 - Resynchronize the server configuration by pushing a new file, merging in changed content or executing commands/scripts to modify the configuration automatically
 - Provide extensibility so system administrators can specify additional files for comparison
 - Execute synchronized operations on a down node when the node reboots
 - Report file differences or errors to the system administrator

There are specific instances (like LDAP for user definition/attribute consistency) where auto-synchronization capabilities already exist. This feature will complement existing capabilities by handling additional types of configuration information.

2. Flexible simultaneous task execution across nodes in a cluster: One of the cluster management challenges is ensuring that tasks are executed across all members of the cluster (even when some of the members are down for preventative maintenance). HP SIM provides a

distributed task facility for executing tasks across multiple systems from a central management server. SVV will enhance HP's solution by providing system administrators the ability to:

- Execute tasks from any cluster member to all cluster member nodes
- Execute tasks on a down node when the node reboots
- Use a command line interface (CLI) and/or web-based GUI – available on all systems and from SIM

3. Centralized cluster management:

Cluster log viewer: One of the existing complexities in cluster management involves trying to view and manually “consolidate” the contents of member-specific log files. The cluster log viewer provides both log file consolidation and log file viewing capabilities. Key log files supported include syslog and Serviceguard package logs. Both a web-based GUI and CLI will be provided.

4. Visualization

System administrators can better visualize the physical and virtual resources in their environment:

- Expanded inventory and health information
- Relationships of resources (clusters, partitions, and workloads)

5. Auto-discovery of single points of failure:

Single virtual view includes a utility for analyzing a cluster or group of systems for single points of failure. Key checks supported include fans/power supplies, high availability connectivity of I/O, failover capabilities and other high availability attributes.

Single virtual view management enhancements for HP-UX 11i v3:

Centralized cluster management - fsweb:

This new application replaces the SAM Disk and File Systems functional area. Like the SAM Disk and File Systems area, fsweb provides the ability to view and configure the various disks and file systems including Logical Volume Manager (LVM) configuration. Some key features of fsweb include:

- Improved visualization of file system status
- Improved visualization of LVM components (LV/VG relationships and cluster usage)
- Access to disk array configuration tools
- Improved disk and file system commands, particularly for tracking disk usage
- More device and file system information
- Launch VxVM and VxFS management tools
- Improved LVM cluster configuration capabilities in a Serviceguard environment

For more information on the HP Virtual Server Environment for HP-UX 11i, visit:

<http://www.hp.com/go/vse>

If you would like to learn more about Virtualizing IT in an Adaptive Enterprise, visit:

<http://h71028.www7.hp.com/enterprise/cache/8886-0-0-225-121.aspx>

For information about HP Systems Management, including products like HP SIM, visit:

<http://h71028.www7.hp.com/enterprise/cache/4225-0-0-121.aspx>

2. Tru64 UNIX V4 File Systems and Volume Management

Both Tru64 UNIX V4 and HP-UX 11i offer robust and proven file systems and volume managers with comparable capabilities. The following sections discuss the details and capabilities of both products.

2.1 File Systems

AdvFS, when used on a standalone system or on a single node in a cluster, offers comparable capabilities to OnlineJFS (VxFS - VERITAS File System). The most noteworthy difference is that the OnlineJFS (VxFS) bundled with HP-UX 11i does not support kernel level Oracle asynchronous I/O. A database accelerator, such as VERITAS QuickIO or ODM interface, must be utilized to enable kernel level Oracle asynchronous I/O driver support for OnlineJFS (VxFS) file systems under HP-UX 11i. The VERITAS database accelerator products are currently only available for HP-UX 11i on the HP 9000 (PA-RISC) platform, with a planned availability of Q3CY2005 for HP-UX 11i on the HP Integrity platform. Currently, kernel level Oracle asynchronous I/O is supported for HP-UX 11i on HP Integrity servers only using raw devices and raw device based logical volumes.

The vast majority of AdvFS file system capabilities are available with HP-UX 11i v2, but may require a different methodology in order to achieve the same functionality. In some cases, more steps may be required to achieve comparable results. In addition, some features are planned for future releases of HP-UX 11i v2 quarterly updates, such as increased file system size and file size limits.

The HP AdvFS and VERITAS VxFS products are both log-based journaled file systems having comparable capabilities and managed using similar concepts. With every platform transition, there will be specific tasks, options or commands that will be new or different. In order to better understand these differences (e.g. file system mount options, related "tunables," file system expansion with volume managers, etc.) it is recommended that one consult the technical documentation and consider HP-UX 11i specific training. For HP training information regarding the VERITAS Volume Manager (VxVM) and File System (VxFS), visit:

<http://www.hp.com/education/courses/u4204s.html>

<http://www.hp.com/education/courses/h7085s.html>

Since AdvFS filesets will not be supported on HP-UX 11i, some customers may need to revisit their file system layout and design under HP-UX 11i.

When Tru64 UNIX customers migrate to HP-UX 11i v2 and begin using VERITAS VxFS V3.5, they will have a clear upgrade path to VxFS V4. The upgrade to VxFS V4 will not have an impact on their data as VERITAS VxFS V4 supports backward compatibility for the VxFS V3.5 disk format. If customers wish to convert the format of a VxFS V3.5 file system to a VxFS V4 file system to take advantage of new functionality, VERITAS supplies the vxupgrade utility to do this. The file system may be kept online if there is sufficient free space available to perform the conversion. In cases where there is not enough free space on a file system, an offline upgrade process will be required.

The following list offers a brief description of the file system products that are available for HP-UX 11i v2. Some of the VERITAS products are bundled with HP-UX 11i Operating Environments (OE) while others require additional licenses that can be purchased from HP.

- a) JFS (Journaled File System): The base VxFS journaled file system bundled in the Foundation Operation Environment of HP-UX 11i. It does not support Direct I/O or kernel level Oracle asynchronous I/O options and should not be considered as an option for supporting transactional (OLTP) databases.

- b) OnlineJFS (Online Journalled File System): The enhanced journaled file system provided with the HP-UX 11i Enterprise and Mission Critical Operating Environment; also known as VERITAS VxFS (VERITAS File System). It supports Direct I/O through extended mount options, but not kernel level Oracle asynchronous I/O (planned in the future). It is currently the most advanced file system available under HP-UX 11i v2. It may be used to house database datafiles, but currently supports only synchronous I/O on the HP Integrity server platform. OnlineJFS can also be purchased as a standalone product from HP.
- c) VERITAS CFS (VERITAS Cluster File System): VERITAS CFS provides for file system sharing across clustered nodes and has functionality comparable to that of TruClusters AdvFS/CFS. VERITAS CFS can support shared file system access across cluster members if they are mounted globally. The primary difference between VERITAS CFS capabilities and those of Tru64 UNIX AdvFS/CFS is the lack of shared root access for a single O/S image across cluster members. The impact of this architectural difference is addressed later in this paper.

For migration purposes, customer data will need to be moved from Tru64 UNIX AdvFS file systems to HP-UX 11i OnlineJFS (VxFS) file systems. For any Tru64 UNIX migration, data movement between file systems on different operating systems (e.g., any file system on Tru64 UNIX to any file system on HP-UX 11i) has always been considered a requirement and part of the plans for the migration process to HP-UX 11i. For further information on the migration of data and file systems from Tru64 UNIX to HP-UX 11i, please see the HP Transition Modules at:

<http://h30097.www3.hp.com/transition/modules.html>

Currently VERITAS V3.5 base Storage Foundation products are integrated with HP-UX 11i v2. HP plans to offer VERITAS V4.x products integrated with HP-UX 11i v2 in Q3CY2005.

The initial release of HP-UX 11i v2 supports file system sizes up to 8 TB and file sizes up to 2 TB. With the HP-UX 11i v2 quarterly patch update released in October 2004, file system sizes are able to grow to 32 TB. This is a significant increase over the current 16 TB file system size limit on Tru64 UNIX V4. For more information on the HP-UX 11i v2 quarterly patch update visit:

<http://www.hp.com/products1/unix/operating/index.html>

2.2 Volume Management

The VERITAS VxVM (VERITAS Volume Manager) and CVM (VERITAS Cluster Volume Manager) products currently offer volume management capabilities comparable to those provided by the Tru64 UNIX AdvFS and Logical Storage Manager (LSM) components. Tru64 UNIX LSM was based on an OEM version of VxVM. Therefore, Tru64 UNIX V4 LSM users will already be familiar with the concepts of these products. In light of this, HP recommends that Tru64 UNIX customers use the VERITAS VxVM (and the related superset CVM) product.

The following sub-section lists volume management products that are available on HP-UX 11i:

- a) VERITAS VxVM (VERITAS Volume Manager): A storage management software sub-system that provides capabilities for configuring storage presented to HP-UX as logical volumes for use on an individual node. The VERITAS volume manager provides additional new functionality, especially in the area of online configuration changes, over and above the HP LVM product, which is bundled with HP-UX 11i. VxVM is very similar to the Tru64 UNIX LSM (Logical Storage Manager) software with regard to its principles and command set. On

HP-UX 11i, VxVM can be used in conjunction with the OnlineJFS (VxFS) file system, and has the ability to manage system boot disks as an alternative to LVM. This is an attractive option for customer environments that have standardized on VERITAS products. For information on the VERITAS Volume Manager on HP-UX 11i, please refer to the release notes:

<http://docs.hp.com/hpux/onlinedocs/5187-1373/5187-1373.html>

- b) HP LVM (Logical Volume Manager): HP includes its own Logical Volume Manager (LVM) product with HP-UX 11i. In general, HP recommends that Tru64 UNIX customers requiring the use of a volume manager with HP-UX 11i migrate to the VERITAS Volume Manager (VxVM). There are no plans at this time to integrate LVM with the VERITAS Cluster File System (CFS). For more information on HP's LVM, please refer to:

<http://docs.hp.com/hpux/onlinedocs/B2355-60103/00/42/4255-con.html>

- c) HP SLVM (Shared Logical Volume Manager): SLVM is enabled during the installation of the Serviceguard Extension for RAC (SGeRAC) product and provides concurrent read/write storage device access from multiple nodes in the cluster. This feature is used only for Oracle's Real Application Cluster (RAC) product, and works on raw LVM-managed logical volumes only. There are no plans to integrate SLVM with the VERITAS Cluster File System.
- d) VERITAS Cluster Volume Manager (CVM): Veritas CVM is an extension to VERITAS Volume Manager (VxVM) and provides concurrent read/write storage device access from multiple nodes in the cluster. This feature is used with Oracle's RAC product and user-written applications that are written to coordinate the concurrent reads and writes. VERITAS CVM is typically used with SGeRAC for Oracle RAC environments on HP 9000 servers. Those customers who wish to implement a cluster file system instead of using raw logical volumes must use VERITAS CVM together with VERITAS CFS (see the next section). The Base VERITAS CVM version comes with the Mission Critical OE. The Full VERITAS CVM version can be purchased from HP as an optional product. The full version of VERITAS CVM requires the full version of VxVM. The VERITAS CVM product availability for the HP Integrity server platform is planned for Q2CY2005.

2.3 File System and Volume Management bundles

As an aid to Tru64 UNIX customers migrating to HP-UX 11i v2, HP plans to provide trade-in bundle options for specific Tru64 UNIX AdvFS file system and LSM volume manager software. The three bundles are planned to be made available in Q2CY2005 and are the:

- VERITAS VxFS and VxVM integrated with the HP File Mover Tool
- VERITAS VxFS integrated with the HP File Mover Tool
- VERITAS VxVM integrated with the HP File Mover Tool

The HP File Mover Tool (FMT) is designed to simplify the process of migrating data from a Tru64 UNIX AdvFS file system to an HP-UX OnlineJFS file system. It automates the analysis, transfer, and verification of files when migrating to the new HP platform. It is intended to move data files and file types you specify, and it will not move application executables or ISV binary files.

For the HP File Mover Tool (FMT), HP is looking at providing an endian conversion tool for binary files from in-house applications. The customer will have to provide the information on the format of the data to make this possible. For third party applications that store data in proprietary binary format (like Oracle), customers will have to use migration methods provided by that application vendor when moving between different endian platforms.

3. Clusters and Third Party ISV applications

For the vast majority of ISV applications, a cluster product, whether it is Tru64 UNIX ASE or Serviceguard, is used to provide failover of applications. In most cases, ISV applications will work similarly and the cluster will provide comparable functionality under HP-UX 11i v2.

HP Serviceguard offers a suite of High Availability (HA) and Disaster-tolerant (DT) solutions. Serviceguard allows the configuration of multiple computers into a highly available cluster. Within such a cluster, applications are grouped into packages, which describe how to start and stop an application, as well as all the resources required. Serviceguard monitors nodes, networks, and processes. If any fail, the package can be configured to automatically be stopped and moved to a different node, along with all the resources it needs. For more information on the features and capabilities of Serviceguard, please refer to:

<http://h71028.www7.hp.com/enterprise/cache/4174-0-0-121.aspx>

For Tru64 UNIX ASE customers that require high application availability, HP recommends the use of HP-UX 11i v2, OnlineJFS (VxFS), VxVM, and Serviceguard. This solution is available on both the HP 9000 (PA-RISC) and HP Integrity platforms today.

HP-UX 11i and Serviceguard will provide comparable or new functionality with various toolkit products such as the Enterprise Cluster Master Toolkit (ECMT) and Serviceguard Extension for SAP (SGeSAP). The ECMT provides a set of pre-tested, fully supported integration scripts for databases such as Oracle and Sybase, as well as applications such as Web services for Apache. Serviceguard and Metrocluster integration for SAP is included in the SGeSAP toolkit. In addition, the HP-UX Mission Critical Operating Environment provides customers with the Workload Manager (WLM) product, which can be used to adjust the resource allocations on the affected servers receiving the new workloads. Serviceguard Extension for Faster Failover (SGeFF) can be used to provide 5 second node failover times with certain restrictions, such as in a 2-node cluster configuration. The restart times of applications after a node failover will vary, depending on the characteristics of each application.

Tru64 UNIX ASE action scripts, in most cases, will have to be re-written, converted, or partially modified for use in a Serviceguard environment. HP is evaluating the possibility of providing tools to scan such scripts and identify areas where modification or rewrite will be needed.

HP provides standard scripts as part of ECMT (Enterprise Cluster Master Toolkit) for applications like Oracle, Apache, and others. These standard scripts will, in some cases, provide an attractive alternative to rewriting ASE action scripts.

3.1 Applications

In a Tru64 UNIX ASE environment, applications run on only one cluster member at a time and are known as single-instance applications. If the cluster member running the application fails, then the application is failed over to another member in the cluster via ASE action scripts.

From the client side, connections to the environment are done through an ASE service. A service can be of type NFS, disk, user defined or shared tape. Tru64 UNIX Production Server also provides a distributed raw disk (DRD) service for Oracle8i Parallel Server (OPS) environments. For more details on OPS environments, please see section 8.3 of this paper.

Tru64 UNIX ASE applications will work in a comparable fashion under HP-UX 11i v2 and Serviceguard. Specifically, when the file system is used on only one node at a time, then this

approach will satisfy the majority of system availability requirements. HP-UX 11i, OnlineJFS (VxFS) or raw devices, VxVM and Serviceguard provide comparable capabilities. Instead of using the Tru64 UNIX ASE services subsystem, customers will use Serviceguard packages, package definition files, and the package manager.

For Tru64 UNIX ASE customers migrating to HP-UX 11i, HP recommends the use of HP-UX 11i v2, OnlineJFS (VxFS), VxVM, and Serviceguard together as a solution. It is available on both the HP 9000 (PA-RISC) and HP Integrity platforms today.

In the future (Q3CY2005), HP plans to integrate Serviceguard and the VERITAS CFS product in the form of a new product bundle, HP Serviceguard integrated with VERITAS Storage Foundation with CFS, which will allow sharing of file system data simultaneously across cluster nodes. VERITAS CFS can be used for global file access across cluster members for data files and applications. Table 1 lists the current and planned cluster and file system options:

Table 1 – High Availability option plans for HP-UX 11i

Cluster Software	HP Serviceguard	HP Serviceguard	HP Serviceguard
Bundle*		Integrated with VERITAS Storage Foundation	Integrated with VERITAS Storage Foundation with CFS
File System	Raw or OnlineJFS (VxFS)	VxFS	VERITAS CFS
Volume Manager	LVM or VxVM	VxVM	VERITAS CVM
High Availability - Failover (Single instance application on one node, with failover capability)			
a) Physically shared connection to storage device with no concurrent shared I/O access	Available today for HP 9000 PA-RISC & HP Integrity	Works today on HP9000 & HP Integrity. New bundle Q3CY2005 for HP 9000 PA-RISC & HP Integrity	
b) Application and File System data simultaneously shared between nodes			Q3CY2005 HP 9000 PA-RISC & HP Integrity

3.2 Disaster-tolerant Solutions

HP offers new functionality to Tru64 UNIX ASE customers in the form of products such as Continentalclusters, Metrocluster, Extended Campus Cluster and Extended Cluster for RAC. In addition to supporting larger distance failover for more servers, these products also provide better resource utilization and data replication capabilities.

For more information on HP’s Disaster-tolerant solutions, please refer to:

<http://h71028.www7.hp.com/enterprise/cache/6469-0-0-225-121.aspx>

4. Cluster Interconnect

For Tru64 UNIX ASE cluster interconnect technology, there are comparable capabilities and new functionality offered with HP-UX 11i v2 and Serviceguard.

Tru64 UNIX Available Server and Serviceguard both support LAN-based interconnects. For low latency, high speed interconnects, Tru64 UNIX Production Server supports Memory Channel, and Serviceguard supports HyperFabric and Infiniband.

Memory Channel is not supported on HP-UX based platforms and customers will need to move to one of the supported interconnects.

Both clustering technologies support LAN-based interconnects for heartbeat communications, and low latency, high speed interconnect options for application traffic and/or message passing. A HyperFabric or Infiniband solution should be considered when applications require high speed data transfer between nodes or support for high speed traffic and/or message passing.

In Serviceguard, the cluster interconnect is used only for heartbeat, and thus LAN connections are typically sufficient. Unlike Tru64 UNIX ASE, Serviceguard requires redundant paths for the heartbeat, which eliminates the network itself as a Single Point of Failure (SPoF). This redundancy can be accomplished in a variety of ways such as:

- Auto-port Aggregation (APA)
- A standby LAN on the same subnet
- Multiple NICs on different subnets
- A cross-over cable directly between NICs

With Serviceguard, if a given LAN interface goes down, or experiences segment saturation, then other heartbeat LANs will remain active and the cluster will stay in formation.

Although LANs may be shared for use with a heartbeat signal and application data, it is recommended that there be at least one cluster interconnect dedicated to heartbeat.

5. Cluster Configurations

For cluster expansion capabilities, there are comparable capabilities and new functionality offered with HP-UX 11i v2.

Serviceguard supports individual nodes with higher CPU counts (up to 128 CPU's per node), up to 16 nodes in a cluster, and longer distances between cluster computer center sites than Tru64 UNIX TruCluster.

The number of nodes supported increases from 8 (Memory Channel) using Tru64 UNIX Production Server, to 16 under HP-UX 11i and Serviceguard. Tru64 UNIX Available Server supports a maximum of 4 nodes (LAN interconnect only).

The maximum distance between cluster nodes will increase from Tru64 UNIX V4 TruCluster's Memory Channel based limit today of 6 KM, to 100 KM in a Serviceguard Extended Cluster configuration or with Metrocluster. For greater distances, applications installed in Serviceguard clusters can failover across a virtually unlimited distance in a Continentalclusters configuration.

6. Custom Code Applications

In general, the comparison findings that apply to Third Party ISV applications may also be applied to in-house developed software. Please refer to section 3. Clusters and Third Party ISV Applications for more detail.

When using Tru64 UNIX ASE, a distributed application may have been designed to allow different cluster members to perform specific tasks or portions of the total application. To implement this, Tru64 UNIX Production Server Memory Channel, DLM, or cluster information service API's may have been used to integrate applications with cluster resources.

There are relatively few Tru64 UNIX ASE applications that have been implemented that make use of these API's. In such cases, the customer will need to modify and rewrite those portions of their application code that utilize the Memory Channel, DLM or cluster information service API's.

Applications that depend on Tru64 UNIX ASE action scripts may need to be modified depending on their interaction with the scripts. The scripts themselves will need to be re-written, based on the appropriate Serviceguard package syntax.

7. Storage Implementations

There are both comparable functionality and differences between the Tru64 UNIX V4 and HP-UX 11i v2 storage implementations. These relative similarities and differences may vary, depending on the implementation chosen.

HP StorageWorks XP and EMC storage currently support HP-UX 11i v2 with VxVM, VxFS, and VERITAS CFS, and will support the VERITAS CVM and HP Serviceguard integrated with VERITAS Storage Foundation with CFS product bundle in the future (planned for Q3CY2005). HP StorageWorks XP and EMC storage arrays support parallel multi-pathing capability with HP-UX 11i v2 using AutoPath in an active-active mode on both the HP 9000 PA-RISC and HP Integrity server platforms.

HP StorageWorks EVA, EMA, RA or MA storage arrays (HSG80 controller based) are supported by HP-UX 11i v2 LVM, SLVM and OnlineJFS (VxFS) file system. SecurePath is required to support the active-passive path failover provided on these controllers. LVM may be employed to provide emulated multi-pathing by striping logical volumes across a number of single-path LUNs. VxVM is also supported, but without multi-pathing. Although EVA storage support will continue, support for EMA, RA, and MA storage arrays is not planned to be incorporated into HP-UX 11i V3. Therefore, customers should consider moving to either XP or EVA based storage arrays when they transition to HP-UX 11i v2 or v3.

In 2HCY2005, support for EVA storage is planned for HP-UX 11i v2 with VxVM, VxFS, VERITAS Cluster File System and VERITAS Cluster Volume Manager. For EVA storage arrays, HP plans to provide full multi-pathing in Active-Active mode on both the HP 9000 PA-RISC and Integrity Server platforms.

Two storage features familiar to Tru64 UNIX users that are not currently included in HP-UX 11i v2 are:

- Global Device Names
- Single Device Naming for LUNs accessed via multi-pathing

This will require a different approach to device naming under HP-UX 11i v2. The above storage features are still being considered for HP-UX 11i v3.

Device naming conventions for logical unit numbers (LUNs) currently differ between Tru64 UNIX V4 and HP-UX 11i. Consistent logical volume device names may be provided under HP-UX 11i using VERITAS Volume Manager software products. Cluster-wide management of logical volume devices may be achieved through the use of the cluster-extended software, VERITAS Cluster Volume Manager (CVM), currently on the HP 9000 PA-RISC platform and, in the future, on the HP Integrity server platform. Although storage array configuration options serve to protect against physical device failure, the VERITAS volume manager products are able to support logical volume striping/mirroring and RAID5 protection as well.

Multi-bus support under HP-UX 11i is provided through the HP Business Copy EVA product (also known as EVM – Enterprise Volume Manager). The EVM software also provides clone, snap-clone and snapshot capabilities for the HP-UX 11i environment. For more information, please refer to:

<http://h18006.www1.hp.com/products/storage/software/bizcopyeva/>

8. Oracle Database Applications

Currently, Oracle database implementations will work in a comparable fashion and provide comparable functionality in the HP-UX 11i v2 environment. In addition, new functionality becomes available for customers who choose to upgrade from Oracle8i on Tru64 UNIX V4 to Oracle9i (or above) on HP-UX 11i v2.

For those customers migrating from an Oracle8i, non-OPS implementation under Tru64 UNIX ASE to Oracle9i on HP-UX 11i that requires data access to only one node at a time, HP recommends that customers use HP-UX 11i v2, Serviceguard, OnlineJFS (VxFS), and VERITAS VxVM.

There are two options for migrating Tru64 UNIX V4 Oracle8i OPS implementations to HP-UX 11i Oracle9i RAC (or 10g RAC). Customers may use the Serviceguard Extension for RAC (SGeRAC) product and house their data on raw devices or on logical volumes created over raw devices. As a future option, HP plans to integrate Serviceguard Extension for RAC (SGeRAC) capabilities with the VERITAS Cluster File System (CFS) product for customers who prefer that their database data files reside within file systems. The new product bundle, HP Serviceguard (SGeRAC) integrated with Storage Foundation with CFS for Oracle RAC, is planned for release in Q3CY2005. For further details, please see the Tru64 UNIX to HP-UX 11i transition module titled "Database for Oracle Planning module V1.2":

<http://h30097.www3.hp.com/transition/modules.html>

8.1 Oracle Single-instance

Oracle executing as a single-instance application, on a single node, will work in a comparable fashion on HP-UX 11i v2. HP-UX 11i v2, OnlineJFS (VxFS) or raw devices, and VERITAS VxVM provide comparable platform capabilities for Oracle 9i.

In Q3CY2005, an integrated HP and VERITAS bundle is planned to be made available. This bundle, called HP Serviceguard integrated with VERITAS Storage Foundation for Oracle Environments, will be available on both HP 9000 (PA-RISC) and HP Integrity servers.

8.2 Oracle High Availability

Oracle running in failover mode on Tru64 UNIX Production Server will also work in a comparable fashion on HP-UX 11i v2. HP-UX 11i v2, OnlineJFS (VxFS) or raw devices, VERITAS VxVM, and Serviceguard, operating in combination, can provide comparable capabilities for failover with few exceptions. Tru64 UNIX ASE action scripts will need to be revised for use under Serviceguard.

Tru64 UNIX customers migrating to HP-UX 11i, who are implementing an Oracle High Availability solution requiring single node exclusive read/write storage access, will find a comparable solution in the combination of HP-UX 11i v2, OnlineJFS (VxFS), VxVM and Serviceguard products. This solution is available on both the HP 9000 (PA-RISC) and HP Integrity platforms today. In Q3CY2005, HP plans to offer HP Serviceguard integrated with VERITAS Storage Foundation with CFS for Oracle, forming a new and enhanced high availability product bundle that provides customers with new functionality not found in the Tru64 UNIX V4 environment. Please see Table 2 in the Oracle RAC section (section 8.3)

8.3 Oracle OPS/RAC

A Tru64 UNIX TruCluster Oracle8i OPS raw device based implementation will work in a comparable fashion on HP-UX 11i with Serviceguard on the HP 9000 PA-RISC platform. As Oracle8i OPS is not supported on the HP Integrity platform, Oracle8i OPS databases must be upgraded to Oracle9i RAC for deployment on Integrity servers. Moving to Oracle9i RAC (or 10g RAC) on HP-UX11i v2 will offer improved multi-instance Oracle clustering functionality and increased performance when used in combination with Serviceguard Extension for RAC (SGeRAC) and SLVM logical volumes created over raw devices.

Tru64 UNIX V4 customers currently using Oracle8i OPS datafiles residing on raw devices, and wanting to evolve to using file systems to house Oracle datafiles, will be able to do so with the future integration of Serviceguard SGeRAC, Oracle 9i RAC and the VERITAS Cluster File System (CFS) products. Such a future integration is planned for Q3CY2005 in the form of a new product bundle, HP Serviceguard (SGeRAC) integrated with VERITAS Storage Foundation with CFS for Oracle RAC. The available options are summarized in Table 2 below:

Table 2 – Database and Clustering option plans for HP-UX

Cluster Software	HP Serviceguard	HP Serviceguard	HP Serviceguard Extensions for RAC (SGeRAC)	HP Serviceguard Extensions for RAC (SGeRAC)
Bundle*		Integrated with VERITAS Storage Foundation with CFS for Oracle		Integrated with VERITAS Storage Foundation with CFS for Oracle RAC
File System	Raw or OnlineJFS (VxFS)	VxFS	Raw Device	VERITAS CFS
Volume Manager	LVM or VxVM	VxVM	SLVM or VxVM	VERITAS CVM
High Availability – Failover (Single instance application on one node, with failover capability)				
Physically shared connection to storage device with no concurrent shared I/O access	Available today for HP 9000 PA-RISC & HP Integrity	New bundle including I/O Accelerators and DB tools Q3CY2005 HP 9000 PA-RISC & HP Integrity		
Applications executing concurrently on multiple nodes, including failover capability (Multi-instance application)				
a) Application data resides on raw devices and is simultaneously shared between nodes			Available today for HP 9000 PA-RISC & HP Integrity	
b) Application and File System data simultaneously shared between nodes				Q3CY2005 HP 9000 PA-RISC & HP Integrity

8.4 Raw Device I/O Options for Databases

Raw devices are an option for housing Oracle database datafiles for HP-UX 11i platforms. The features and behavior of raw device I/O access is basically the same on Tru64 UNIX and HP-UX 11i. Usage across platforms remains essentially the same, with the exception of raw device (LUN) identification. Under Tru64 UNIX V4, raw device LUN's are identified as "rz/rrz" numbers, and under HP-UX 11i v2, they are represented in "controller-target-device" address format (cmtndn). On both platforms, volume management software allows for custom identification/naming of logical volumes created on top of LUN raw devices by the system administrator (please refer to Table 3).

Table 3 - Raw device options

OS	Synchronous I/O	Asynchronous I/O *	Volume Management	LUN Identification
Tru64 UNIX V4	Yes	Yes	LSM	/dev/rrzn
HP-UX 11i v2	Yes	Yes	VxVM, LVM, SLVM	/dev/rdisk/cmtndn

* Recommended for database usage

Raw database I/O is efficient in that no file system buffering is required. Both synchronous and asynchronous I/O is supported on the Tru64 UNIX Alpha Server, HP-UX HP 9000 and HP-UX 11i HP Integrity platforms. With raw devices, HP-UX 11i offers a kernel level asynchronous I/O driver which is used in the access of databases such as Oracle (and Oracle RAC) to optimize I/O performance.

8.5 File Systems I/O Options for Databases

Please refer to the File Systems and Volume Management section for more information on the various options.

9. Conclusions

HP is committed to providing IT solutions that meet the dynamic needs of businesses in today's challenging market climates. The HP Adaptive Enterprise is HP's approach to providing businesses with computing environments that evolve and adapt to meet those needs. By providing a tighter alignment across clustering, partitioning, multi-platform environments, system management and virtualization technologies, the HP Adaptive Enterprise can best meet those requirements while maximizing IT investments. Although TruCluster Server and AdvFS will not be delivered in HP-UX 11i v3, the VERITAS Storage Foundation technologies will be integrated with Serviceguard and made available earlier in HP-UX 11i v2. HP is confident that these changes will provide customers with the best IT solutions for meeting the business/IT challenges of today and tomorrow.

References

HP Virtual Server Environment and the HP Adaptive Enterprise

- HP Virtual Server environment for HP-UX
<http://www.hp.com/go/vse>
- Virtualizing IT in an Adaptive Enterprise
<http://h71028.www7.hp.com/enterprise/cache/8886-0-0-225-121.aspx>
- HP System Management for HP-UX 11i
<http://h71028.www7.hp.com/enterprise/cache/4225-0-0-0-121.aspx>

HP-UX File Systems and Volume Management

- JFS Tuning and Performance
http://docs.hp.com/hpux/onlinedocs/5576/JFS_Tuning.pdf
- OnlineJFS information
<http://www.hp.com/products1/unix/operating/onlinejfs.pdf>
- HP-UX 11i V2 Update 2 Release
http://www.hp.com/products1/unix/operating/hot_topic_unix.html#common_release
- VERITAS Volume Manager release notes
<http://docs.hp.com/hpux/onlinedocs/5187-1373/5187-1373.html>
- HP's LVM
<http://docs.hp.com/hpux/onlinedocs/B2355-60103/00/42/4255-con.html>

Oracle Database Applications and Serviceguard Clusters

- Serviceguard HA and DTS solutions for HP-UX 11i and Linux
<http://h71028.www7.hp.com/enterprise/cache/6469-0-0-225-121.aspx>

HP-UX and VERITAS Volume Manager and File System Training

- HP-UX Education Program
<http://www.hp.com/education/sections/hpux.html>
- VERITAS Volume Manager and File System Administration
<http://www.hp.com/education/courses/u4204s.html>
- VERITAS Volume Manager for HP-UX
<http://www.hp.com/education/courses/h7085s.html>

For More Information

For information on HP-UX 11i, the proven foundation for the Adaptive Enterprise

- HP-UX 11i Operating Environment
<http://www.hp.com/products1/unix/operating/index.html>

For training information on HP-UX 11i and VERITAS Volume Manager and File System

- HP-UX 11i Education Program
<http://www.hp.com/education/sections/hpux.html>
- VERITAS Volume Manager and File System Administration
<http://www.hp.com/education/courses/u4204s.html>
- VERITAS Volume Manager for HP-UX 11i
<http://www.hp.com/education/courses/h7085s.htm>

© 2004 Hewlett-Packard Development Company, L.P. The information and roadmap information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

VERITAS, VERITAS Software and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation in the US and/or other countries.

Oracle is a registered trademark of Oracle Corporation. Various product and service

12/2004



i n v e n t