

Transitioning Tru64 UNIX V5 Customers to HP-UX 11i v2:

Analyzing the HP-UX 11i Roadmap Changes in Terms of File Systems, Clusters, and Infrastructure



Overview	2
1. HP Virtual Server Environment (VSE)	3
2. Tru64 UNIX V5 File Systems, Volume Managers, Cluster File Systems and SSI	6
3. Clusters and Third Party ISV Applications.....	11
3.1 Single-instance Applications	11
3.2 Multi-instance Application, Failover Mode.....	13
3.3 Disaster-Tolerant (DT) Solutions.....	13
4. Cluster Interconnect.....	13
5. Cluster Configurations	14
6. TruCluster Cluster Alias.....	14
7. Custom Code Applications.....	15
8. Storage Implementations.....	16
9. Oracle Database Applications.....	17
9.1 Oracle Single-Instance.....	17
9.2 Oracle High Availability (HA - Failover).....	18
9.3 Oracle RAC.....	18
9.4 Raw Devices I/O Options for Databases	20
9.5 File Systems I/O Options for Databases	20
Conclusions.....	20
Appendix A: References	21
For More Information	23

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. 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.¹

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 our 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 our 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 customers on Tru64 UNIX V5 who are planning or designing an HP-UX 11i solution. The purpose of this paper is to provide a complete assessment of the HP-UX 11i roadmap changes affecting Tru64 UNIX customers and highlight key areas of potential differences; i.e. the product platform (file system, cluster alias, etc.) and how the product is deployed or used (Oracle database, application deployment, patch management, etc.). In most cases, comparable capabilities are provided that may require different approaches to accomplish similar tasks. In a limited number of cases, a modification of how the cluster or applications are

¹ Gartner, "Predicts 2004: Server Virtualization Evolves Rapidly" by Tom Bittman, November 14, 2003

managed may be required to address differences in the tools available. In several key areas, new functionality will be provided, offering new benefits to Tru64 UNIX customers.

This document is organized into the following major sections:

- HP Virtual Server Environment (VSE)
- Features and Functions for file systems, volume managers and clusters
- Effects on Oracle databases and clustering

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 System 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 kcweb (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 V5 File Systems, Volume Managers, Cluster File Systems and SSI

Today's Tru64 UNIX V5 and HP-UX 11i offer feature-rich environments that require additional discussion on how Tru64 UNIX functionality compares to HP-UX 11i, Virtual Server Environment (VSE), and the increasingly sought-after single virtual view environment. When Tru64 UNIX V5 customers migrate, and depending on which functions are being used, they will find cases:

- where new functionality will be available to them with HP-UX 11i v2
- where HP-UX 11i v2 offers comparable capabilities
- where moderately different approaches are required to accomplish similar tasks
- where a modification of how they use their cluster or applications will be required to address the usage of the HP-UX 11i operating system or tools

2.1 File Systems

HP's 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 component must be separately purchased/licensed to enable kernel level Oracle Asynchronous I/O driver support for OnlineJFS (VxFS) file systems under HP-UX 11i. The database accelerator product (QuickI/O or ODM interface) is currently available for purchase from VERITAS for HP-UX 11i on the HP 9000 (PA-

RISC) platform, and is scheduled to be available in Q3CY2005 for HP-UX 11i on the HP Integrity server platform. Currently, kernel level Oracle asynchronous I/O capability is supported for HP-UX 11i on the HP Integrity server only for raw devices or raw device based logical volumes. For more information on the HP and VERITAS bundles for Oracle, please refer to the Oracle section.

Aside from that, the vast majority of AdvFS file system functionality is achievable within HP-UX 11i v2. In some cases, the implementation may require a slightly different methodology in order to achieve that same level of functionality (e.g. multiple filesets per domain). In addition, some features are planned for future releases of HP-UX 11i v2 quarterly updates, such as increased file system size limits. When HP-UX 11i v3 is released, file size limits of up to 16 TB will be supported. Customers requiring a file size limit greater than 2 TB should plan on waiting for the HP-UX 11i v3 release for that capability.

The HP AdvFS and VERITAS VxFS products are both log-based journaled file systems with comparable capabilities and are 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 the volume manager, etc.) it is recommended that one consult the technical documentation and consider HP-UX 11i specific training. For more on HP training regarding VERITAS Volume Manager (VxVM) and File System (VxFS), go to:

<http://www.hp.com/education/courses/u4204s.html> and
<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 file system layout or design with 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 and future V4.x releases. The upgrade to VERITAS VxFS V4 will not have an impact on their data as 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 or will be 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 is 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 is not recommended for transactional (OLTP) databases.
- b) OnlineJFS (Online Journaled File System): The enhanced journaled file system provided with the HP-UX 11i Enterprise and Mission Critical Operating Environments; 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) capabilities will be integrated with Serviceguard. There will be several options for customers, which are planned to be available Q3CY2005. The VERITAS CFS provides for file system sharing across clustered nodes and has functionality comparable to that of Tru64 UNIX 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 the Tru64 UNIX AdvFS file system to the HP-UX 11i OnlineJFS (VxFS) file system. 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 and Serviceguard in Q3CY2005.

With the HP-UX 11i v2 quarterly patch update released in October 2004, the file system size limit has grown to 32 TB. For more information on the HP-UX 11i v2 quarterly patch update visit:

http://www.hp.com/products1/unix/operating/hot_topic_unix.html#common_release

2.2 Volume Management

The VERITAS Volume Manager (VxVM) and VERITAS Cluster Volume Manager (CVM) 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 V5 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 product (and the related superset VERITAS CVM when available).

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 11i 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 software with regard to its principles and command set. On HP-UX 11i, VxVM is used in conjunction with the OnlineJFS (VxFS) file system, and has the ability to manage system boot disks as an alternative to LVM. 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 that require 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) when it becomes available on HP Integrity server. For more information on HP's LVM, please refer to:

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

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.

2.4 Cluster File System, Cluster Volume Management and Single System Image

If AdvFS is being used on a single node, the comparable capabilities offered by OnlineJFS (VxFS) may require that a different methodology be employed to achieve the same functionality.

If AdvFS, Tru64 UNIX CFS and LSM are being used in a multi-node environment, there will be comparable capabilities offered through the HP Serviceguard integrated with VERITAS Storage

Foundation with CFS product bundle. There are also features that were available with Tru64 UNIX CFS that will not be available on HP-UX 11i.

The HP Serviceguard integrated with VERITAS Storage Foundation with CFS product bundle is scheduled to be available in Q3CY2005 on HP-UX 11i v2. The VERITAS Cluster File System functionality integrated with Serviceguard will allow globally mounted user files and applications to be accessible by all cluster members. For information that is not globally mounted, users of HP-UX 11i Serviceguard clusters will need to manage data on an individual, per cluster member basis. HP has tools today, like Systems Insight Manager, to help with multi-system management.

Since a shared root (/) capability will not be made available in HP-UX 11i, system directories must be maintained on an individual cluster member node basis. For ease of system management, HP has adopted the strategy of maintaining one system image, sometimes referred to as a "golden image." This image can include the Operating System, aspects of the configuration, and even applications. This one system image can be easily cloned and installed on the respective cluster members. HP-UX 11i software packages are often in Software Distributor format and are managed via the Software Distributor (SD-UX) tool. Golden Images are created and deployed via the Ignite-UX product. In addition, the HP OpenView solution portfolio includes the OpenView OS Manager using the Radia product for software management across a wide range of Operating Systems, including HP-UX 11i. HP-UX 11i provides better system deployment and maintenance technologies than found in Tru64 UNIX through the use of the Ignite-UX, SD-UX, and OpenView OS Manager using the Radia product.

For more information on Ignite-UX, please refer to:

<http://software.hp.com/products/IUX/>

For information on OpenView OS Manager using Radia, please refer to:

http://managementsoftware.hp.com/products/radia_osm/index.html

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

- a) VERITAS CVM (VERITAS Cluster Volume Manager): The VERITAS Cluster Volume Manager product is an extension to the VxVM product described previously. It allows multiple nodes parallel access to disk storage devices, enabling concurrent access to Oracle database datafiles and for other applications that may be written to take advantage of this feature. In addition, it provides the foundation for the flexible mirroring and striping capabilities existing in the base VxVM product. The VERITAS CVM concurrent I/O capability allows multiple connections from any given node to the attached storage array for improved I/O performance. VERITAS CVM must be used in conjunction with VERITAS CFS for customers who wish to implement a cluster file system. The Base VERITAS CVM version currently comes with the HP-UX 11i Mission Critical OE for the HP 9000 PA-RISC platform. The Full VERITAS CVM version can be purchased from HP as an optional product and requires the full version of VxVM. VERITAS CVM product availability for the HP Integrity server platform is planned for Q2CY2005.

- b) HP SLVM: The HP native volume manager, LVM, is enhanced for clustering when HP Serviceguard/SGeRAC is installed. SGeRAC includes Shared LVM (SLVM) enhancements for concurrent read/write storage device access from multiple nodes of a cluster. This feature is generally used with Oracle's Real Application Cluster (RAC) product, and works only with raw LVM-managed logical volumes. There are no plans to integrate SLVM with the VERITAS Cluster File System when it becomes available on HP Integrity server. For more information on HP LVM, please refer to:

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

3. Clusters and Third Party ISV Applications

For the vast majority of ISV applications, a cluster product, whether it is TruCluster Server or Serviceguard, is used to provide failover of single instance applications. In these cases, the ISV application will work similarly and the cluster will provide comparable functionality with HP-UX 11i v2. There are a few exceptions, not resulting from the ISV software itself, but rather from the way the ISV application is installed and configured on Tru64 UNIX. These exceptions result from the use of cluster alias, or AdvFS features such as global file names, where tasks may need to be done differently or modified.

HP Serviceguard offers a suite of High Availability and Disaster-Tolerant (DT) solutions and allows the configuration of multiple computers to form 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>

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 standard 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, HP-UX 11i v2 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 five second node failover times (as opposed to minutes) 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.

3.1 Single-instance Applications

In a Tru64 UNIX TruCluster Server environment, a single-instance application runs on only one cluster member at a time. If the cluster member executing the application fails, the application is failed over to another member in the cluster using Cluster Application Availability (CAA) scripts.

Tru64 UNIX CAA scripts, in most cases, will have to be re-written, converted, or partially modified to be used in the Serviceguard package manager format for use in the counterpart Serviceguard environment. HP is evaluating the possibility of providing tools to scan scripts and identify areas where modification or rewrite will be needed. ECMT standard scripts will provide, in some cases, a viable alternative to rewriting CAA Scripts. In addition, there are a few other cluster functionality differences, especially when the solution used involved TruCluster SSI capabilities for cluster member management, or cluster alias for load balancing. In these situations, customers will need to modify the affected portions of their implementation.

Tru64 UNIX V5 single-instance applications, requiring application failover to one node and single node exclusive read/write storage access, will work in a comparable fashion using HP-UX 11i v2, Serviceguard, OnlineJFS (VxFS) or raw devices, and VxVM. Instead of using the CAA subsystem and CAA scripting, Serviceguard packages, scripts, package definition files, and the Serviceguard package manager would be used. This solution is currently available on both the HP 9000 (PA-RISC) and HP Integrity platforms.

In Q3CY2005, the HP Serviceguard integrated with VERITAS Storage Foundation with CFS product bundle is scheduled to be available for both HP-UX 11i platforms. With the capabilities of this new product bundle, application code and data can be stored in a cluster file system, providing shared global access concurrently across Serviceguard cluster nodes.

Table 1 -- High Availability options 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 HP 9000 PA-RISC & Integrity	Works today on HP 9000 & HP Integrity. New bundle Q3CY2005 HP 9000 PA-RISC & Integrity	
b) Application and File System data simultaneously shared between nodes			Q3CY2005 HP 9000 PA-RISC & Integrity

3.2 Multi-instance Application, Failover Mode

A multi-instance application can run on multiple cluster member nodes at one time. If one cluster member fails, it does not affect the other instances of the application running on the other cluster members. On Tru64 UNIX, cluster aliases help to provide transparent client access to multi-instance applications. HP-UX 11i achieves comparable results through other means. Please refer to the “Cluster Alias” section (Section 6).

The vast majority of multi-instance applications are based on Oracle RAC implementations. For recommendations on how to best migrate your Oracle environment, please refer to the “Oracle Database Applications” section (Section 9).

A number of customers have implemented their own multi-instance applications using TruCluster Server technology. For recommendations on how to best migrate your custom developed application environment, please refer to the “Custom Code Application” section (Section 7).

3.3 Disaster-Tolerant (DT) Solutions

HP offers new functionality to Tru64 UNIX TruCluster Server customers in the form of products such as Continentalclusters, Metrocluster, Extended Campus Cluster, or Extended Cluster for Oracle RAC.

For more information on HP Disaster-Tolerant (DT) solutions, please refer to:

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

4. Cluster Interconnect

For Tru64 UNIX V5 cluster connection technology, there are comparable capabilities and new functionality offered with HP-UX 11i v2 and Serviceguard.

Both product clustering technologies support LAN-based interconnects. For low latency, high speed interconnects, TruCluster Server supports Memory Channel, and Serviceguard supports HyperFabric and Infiniband.

Memory Channel support is not and will not be available on HP-UX based platforms. Therefore, customers will need to move to one of the other LAN-based interconnects: for example, GigaBit Ethernet, or to a high speed interconnect such as HyperFabric or Infiniband.

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 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 a TruCluster Server implementation, 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 any given LAN interface goes down, or experiences segment saturation, other heartbeat LANs will stay active and the cluster will remain intact.

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.

By cluster expansion, we mean that Serviceguard supports individual nodes with higher CPU counts (up to 128 CPU's per system), up to 16 nodes in a cluster, and longer distances between cluster data center sites than Tru64 UNIX TruCluster Server environments. Tru64 UNIX V5 TruCluster Server can support individual nodes with only up to 64 CPU's and up to only 8 nodes (for Memory Channel or LAN interconnect) to a cluster.

The maximum distance between cluster nodes will increase from the Tru64 UNIX V5 TruCluster 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. TruCluster Cluster Alias

The TruCluster cluster alias feature provides a single cluster IP address for high availability access, protection from network failures, and load balancing by redirecting traffic to an application which may reside on any cluster member node.

If customers are currently using the TruCluster cluster alias feature for high availability, different alternative mechanisms may have to be employed under HP-UX 11i v2 in order to achieve comparable functionality. If the TruCluster cluster alias feature is being used for load balancing purposes, the approach may need to be modified for an HP-UX 11i implementation. Although the TruCluster cluster alias capability will not be provided in HP-UX 11i, several alternative options do exist.

If TruCluster cluster alias is being used to provide high availability in a light load environment, HP recommends directing traffic to a particular Serviceguard package with a re-locatable IP address. In the event of a node failure, Serviceguard will direct traffic to another node and mount the data to that node. The VERITAS Cluster File System (CFS) functionality, when available, will permit access to globally mounted file system data that is not located on the node where the application is running.

It will be necessary to directly connect to particular nodes or use NFS for access to data that is contained in locations other than VERITAS CFS global files. Other techniques may be used in some environments. These include manual or scripted connections to different nodes, and node selection at the application level for applications that provide this capability.

The TruCluster cluster alias feature is also sometimes used to protect against certain types of network failures. For example, if the network adapter on the member where the cluster alias IP address is currently located fails, then the cluster alias IP address will move to another member. In these cases, additional network hardware may be required to achieve the same level of network availability in the HP-UX 11i Serviceguard environment.

Tru64 UNIX customers also used NetRAIN for link aggregation and higher network availability. HP-UX 11i Serviceguard provides equivalent functionality for load-balancing and network redundancy with its Auto Port Aggregation (APA) product by aggregating multiple network interface cards (NICs) into a single logical network, similar to Network Bonding in Linux. Serviceguard also allows additional LAN interfaces to be designated as failover LANs. Although not configured with an IP address, such interfaces serve as a standby for failover in the event of primary interface failure, and thus aid in eliminating Single Points of Failure conditions. Serviceguard will monitor the status of Network Interface Cards (NICs), and if all NICs associated with a given application subnet fail, Serviceguard will then fail the application over to another node that has maintained network connectivity. For more details on APA, visit:

<http://docs.hp.com/hpux/onlinedocs/J4240-90025/00/00/1-con.html>

Some applications incorporate load-balancing capabilities, which should continue to be used. For other applications that do not provide this functionality, and have instead relied upon the cluster alias capability to achieve load balancing, there are at least four options.

In order to implement a highly available service providing a virtual IP address to the client network, one can use a Serviceguard package with a re-locatable IP address to transfer to a different node in the cluster. Another approach is to connect directly to a single node and use scripts to reconnect to a different node. Serviceguard, in conjunction with Auto-Port Aggregation (APA), can provide load balancing and network redundancy.

If TruCluster cluster alias is being used to provide high availability and load balancing in a heavy load environment, HP recommends considering the use of a smart load-balancing hardware device to present a single IP address for the service, distribute traffic to the various cluster members, and provide high availability. The HP Serviceguard integrated with VERITAS Storage Foundation with CFS product bundle will permit applications on various cluster member nodes to access the same data residing in globally mounted files. It will be necessary to directly connect to individual member nodes to access data residing in locations other than VERITAS CFS global files. This will provide a high degree of equivalence to TruCluster cluster alias functionality.

7. 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 the "Clusters and Third Party ISV Applications" section (Section 3).

When using TruCluster Server, 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 type of distributed functionality design, Memory Channel, Distributed Lock Manager (DLM), or cluster alias API's may be used to integrate applications with cluster resources.

Although relatively few in number, such Tru64 UNIX customer applications may have been developed to incorporate various TruCluster Server API's, related to the utilization of DLM, Memory Channel, Cluster Alias, and shared root (/) capabilities. Since no equivalent to the TruCluster Server API exists within HP-UX 11i v2, such cases will require that portions of the application code be modified and rewritten.

Applications that depend on Tru64 UNIX Cluster Application Availability (CAA) scripts may need to be modified depending on their interaction with the scripts. The scripts themselves will need to be converted/re-written, based on the appropriate Serviceguard package syntax.

8. Storage Implementations

There are both comparable functionality and differences between the Tru64 UNIXv5 and HP-UX 11i storage implementations. These relative similarities and differences may vary, depending on current versus future supported configurations.

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

HP StorageWorks EVA, EMA, RA or MA storage arrays are supported by HP-UX 11i v2 with LVM, SLVM and OnlineJFS (VxFS) file system. SecurePath may be used in active-passive mode to support single-pathing with automatic single-path failover. 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.

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.

VxVM and VxFS will not be supported with multi-pathing for EMA, RA, or MA storage arrays. The HP Serviceguard integrated with VERITAS Storage Foundation with CFS product bundle will not be supported for EMA, RA, or MA storage arrays. Support for these storage arrays is not planned to be incorporated into HP-UX 11i v3. Therefore, customers should consider moving to either EVA or XP storage when they transition to HP-UX 11i v2 or v3.

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 HPUX-11i. 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 V5 and HP-UX 11i. Consistent logical volume device names may be provided under HP-UX11i 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 available 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/>

9. Oracle Database Applications

Currently, the majority of Oracle database implementations will work in a comparable fashion and provide comparable functionality in the HP-UX 11i v2 environment. There are a few exceptions: when an application is installed or configured to use cluster alias, or AdvFS features such as global file names. In these cases, the method of implementation may need to be done differently or modified.

In general, the Oracle database migration information published 31 March 2004, in the “Tru64 UNIX to HP-UX 11i Transition Database Planning Module V1.0 for Oracle” still applies. This database transition module for Oracle can be accessed at:

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

Three possible Oracle implementation scenarios will be discussed in this section:

- Oracle Single-Instance
- Oracle High Availability (HA - failover)
- Oracle Real Application Clusters (RAC)

9.1 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, OnlineJFS (VxFS) or raw devices, and VERITAS VxVM provide comparable platform capabilities.

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.

9.2 Oracle High Availability (HA - Failover)

Oracle running in failover mode on Tru64 UNIX using CAA will work in a comparable fashion on HP-UX 11i v2. Currently, HP-UX 11i v2, OnlineJFS (VxFS) or raw devices, VERITAS VxVM, and Serviceguard, operating in combination, can provide comparable capabilities for failover with a few exceptions. When TruCluster SSI capabilities are used for cluster member management or file sharing across the cluster, affected customers will need to modify those portions of their implementation under HP-UX 11i v2.

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. Please refer to Table 2 in Section 9.3.

9.3 Oracle RAC

There are two options for Tru64 UNIX Oracle RAC implementations being migrated to HP-UX 11i; one currently available and one available in the near future.

A Tru64 UNIX Oracle RAC raw device-based implementation will work in a comparable fashion on HP-UX 11i v2 with Serviceguard. The combination of HP-UX 11i v2, VERITAS Shared Logical Volume Manager (SLVM) over raw devices, and Serviceguard Extension for RAC (SGeRAC), provides comparable capabilities for Oracle RAC with a few exceptions.

When transitioning to HP-UX 11i v2, customers that have been using TruCluster SSI for cluster member management, or cluster alias for load balancing, will need to modify those portions of their implementation that currently rely on those capabilities.

Tru64 UNIX TruCluster Server customers currently using Oracle RAC datafiles residing within file systems, and wanting to continue using file systems to house Oracle datafiles, will need to wait for the scheduled availability in Q3CY2005 of the HP Serviceguard (SGeRAC) integrated with VERITAS Storage Foundation with CFS for Oracle RAC bundle. Please refer to Table 2.

Table 2 -- Database and Clustering options for HP-UX platforms:

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 Devices	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 HP 9000 PA-RISC & Integrity	New bundle including I/O Accelerators and DB tools Q3CY2005 HP 9000 PA-RISC & 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 HP 9000 PA-RISC & Integrity	
b) Application and File System data simultaneously shared between nodes				Q3CY2005 HP 9000 PA-RISC & Integrity

9.4 Raw Devices 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. Behavior and usage across platforms essentially remains the same, with the exception of raw device logical unit number (LUN) identification. Under Tru64 UNIX, raw device LUN's are identified as "dsk" numbers, and under HP-UX 11i v2, they are represented in "channel-target-device" address format (cntndn). On either platform, 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 below.

Table 3 - Raw device I/O options

OS	synchronous I/O	asynchronous I/O - Oracle*	Volume Management	LUN Identification
Tru64 UNIX V5	Yes	Yes	LSM	/dev/rdisk/dsknn
HP-UX 11i v2	Yes	Yes	VxVM, SLVM, CVM	/dev/rdisk/cntndn
HP-UX 11i v3	Yes	Yes	VxVM, CVM	/dev/rdisk/dsknn

* Recommended for database usage

Raw database I/O is efficient in that no file system buffering is required. Both synchronous and kernel Oracle asynchronous I/O are supported on the Tru64 UNIX AlphaServer, HP-UX HP 9000 and HP-UX 11i Integrity Server platforms. With raw devices, HP-UX 11i provides a kernel asynchronous I/O driver for Oracle and Oracle RAC in order to optimize database I/O performance.

9.5 File Systems I/O Options for Databases

Please refer to the "File Systems and Volume Management" section (Section 2) for more information on the various options.

Conclusions

Within this whitepaper we have addressed the reasons for the change in the HP-UX 11i roadmap, laid out the framework for the revised roadmap, and assessed the changes made to that roadmap relative to making the transition from Tru64 UNIX AlphaServers to HP-UX11i Integrity servers. These changes obviously affect the nature of the similarities and the differences between the two UNIX platforms, now and in the future. We have also noted the potential benefits these changes will bring to Tru64 UNIX environments by transitioning to HP-UX 11i.

To a great degree, the similarities and differences discussed will have an impact on how system administration, system management, and application solutions are implemented on the new HP-UX 11i Integrity server platform. A follow-on related document, entitled, "Transitioning Tru64 UNIX V5 customer to HP-UX 11i v2: The HP-UX11i Roadmap Changes in Terms of Implementation Solutions for System Administration and System Management" is also available, and attempts to offer some guidance on how such implementations may be affected.

Appendix A: References

The following URLs are those already listed in the body of this paper. They have been included here for the convenience of the reader.

- Strengthening the HP Virtual Server Environment:
 - 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>
- Training information on HP-UX and VERITAS Volume Manager and File System:
 - 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 names referenced herein may be trademarks of Oracle Corporation

11/2004