

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

Analyzing the HP-UX11i Roadmap Changes in Terms of Implementing Solutions for System Administration and System Management



Overview	2
1. Cluster Installation and Maintenance	2
1.1 Install O/S and Cluster Software.....	4
1.2 Application Services	5
1.3 Storage Management and Clusters.....	6
1.4 Application Deployment, High Availability and Load Balancing.....	7
1.4.1 Application Deployment.....	7
1.4.2 Application High Availability (HA).....	8
1.4.3 Application Load Balancing	9
1.5 Maintaining Cluster Applications	11
1.6 Periodic Cluster Maintenance	12
1.7 Monitoring System Performance.....	13
1.8 Monitoring Users.....	13
1.9 OS Version Upgrades and Patch Installation.....	13
1.10 Adding a Member Node to the Cluster (Extending a Cluster)	14
1.11 Cluster Troubleshooting.....	15
2. Applications – NFS Servers and TCP Services	16
3. Backup and Restore	17
3.1 Native File System Backup/Restore Command Set.....	17
3.2 Third Party (ISV) and HP Backup/Restore Products	18
Conclusions.....	19

Appendix A: References	20
For More Information	22

Overview

In the related document entitled: “Transitioning Tru64 UNIX V5.x Customers to HP-UX 11i v2: Assessment of Changes to the HP-UX 11i Roadmap in Terms of File Systems, Clusters, and Infrastructure”, we discussed why the HP-UX 11i roadmap was changing, provided an assessment of those changes, and examined how those changes can potentially affect the Tru64 UNIX customer base.

This document is intended to complement and supplement the above referenced document, and as such, is targeted to those customers on Tru64 UNIX V5.x who are planning or designing an HP-UX 11i solution. Its purpose is to provide a look at the HP-UX 11i product roadmap in detail, along with insight and guidance as to how several major system administrative and system management scenarios may be adapted for implementation on the HP-UX 11i-based HP Integrity server platform.

It will also include descriptions of the array of HP-UX 11i products available, how they might be applicable to selected implementation scenarios, and how those products may be deployed as a solution. In some cases, third-party (ISV) software is mentioned in terms of available support. As in the parent document, the phrases “comparable”, “different approaches”, “modified”, and “new functionality” will be used to describe the implementations discussed with regard to the Tru64 UNIX AlphaServer (Source) and HP-UX 11i HP Integrity server (Target) platforms.

This document is organized into the following system administration/management implementation sections:

- Cluster installation and maintenance
- Cluster services
- Backup and restore

1. Cluster Installation and Maintenance

During the course of maintaining clusters, Tru64 UNIX V5.x customers will find that there are comparable approaches to managing individual nodes on HP-UX 11i v2. When managing multiple cluster nodes, there are areas of comparable functionality, as well as areas that may require different approaches and/or task modification, in an HP-UX 11i v2 environment.

If performing system management functions on a cluster (i.e. direct file manipulation such as `/etc/passwd` modification), one will need to connect to a single node, make the changes, and then use manual or scripted methods for reconnecting to a different node. Then the changes may be performed on the next node. Alternatively, one might use the HP SIM product and perform an operation on the Central Management Server (CMS), which is distributed to a defined set of servers, such as the members of a cluster. This would permit the execution of a specific management script across all members of the cluster. In some cases, more steps may be required to achieve similar results.

It is a goal of HP to provide tools designed to allow one-touch management of multiple nodes. Since there will be no shared root Single System Image (SSI) offered in HP-UX 11i, each member of the cluster will have its own copy of the Operating System and various configuration files. While a Cluster File System (CFS) will be available for customer and ISV usage, the HP-UX 11i root file system (/) will not be able to use the CFS. Even so, customers can use existing HP-UX 11i system management tools (and future enhancements) to efficiently manage the members of their cluster.

Existing HP-UX 11i v2 management tools include:

HP Systems Insight Manager (SIM) - delivers inventory, fault, and remote management to all your HP systems. HP SIM can be easily extended to offer additional server, storage, printer, and power management. For the HP-UX 11i environment, it can be extended with software deployment, configuration, and workload management to offer a complete resource lifecycle management solution, allowing the central administration of the HP Virtual Server Environment.

System Administration Manager (SAM) – provides an easy-to-use, out-of-the-box system management tool to offer both a secure, central point of control and effective management for remote systems to perform routine systems management tasks.

Ignite-UX - enables one to build an OS image, which can include applications, and distribute it to multiple servers simultaneously.

Serviceguard Manager - enables one to monitor, administer, and configure multiple Serviceguard clusters and their member systems.

Partition Manager - manages the hard partitions (nPars) of your servers remotely.

LDAP solutions - enables one to ensure consistency in user and group definitions across the enterprise.

HP Process Resource Manager (PRM) – is a workload management tool used to control the amount of resources that processes use during peak system load.

HP Workload Manager (WLM) - a goal-based workload management tool that provides automatic CPU resource allocation and application performance management based on prioritized service-level objectives (SLOs). In addition, real memory and disk bandwidth allocations can be set to fixed levels in the configuration.

An overview brochure of all HP System Management tools is available at:

<http://h71028.www7.hp.com/ERC/downloads/5982-6788EN.pdf>

If using Tru64 UNIX CFS features for global file names, then one will need to use VERITAS CFS for global files, or modify the naming scheme or use a different approach in HP-UX 11i v2. HP tools can help to maintain consistency of information between nodes. More details on the multi-system management capabilities of the HP Systems Insight Manager are available at:

<http://h18013.www1.hp.com/products/servers/management/hpsim/index.html?jumpid=go/hpsim>

One of the benefits Tru64 UNIX customers will derive from migrating to HP-UX 11i will be the ability to take advantage of the HP Virtual Server Environment (VSE). Within a VSE, virtual servers automatically grow and shrink based on the service-level objectives set for each application they host. It orchestrates the resources in a high availability environment, for all the different forms of partitions (hard, virtual, and resource partitions) and for utility pricing (Instant Capacity, Temporary Instant Capacity, Metered Capacity). For example VSE integrates with Serviceguard and disaster tolerant solutions: when a Serviceguard package fails over from one node to another, the pre-defined SLO will be activated and VSE then allocates resources accordingly. Or VSE can move one processor from one virtual partition to another and can initiate turn-on or turn-off of instant capacity CPUs.

VSE can even move iCAP (Instant Capacity) CPU licenses, which can be dynamically transferred across hard partitions (nPars) to meet service levels by increasing and decreasing available CPU capacity. For a more detailed discussion of how to achieve high availability through system management by integrating HP Serviceguard and HP-UX 11i Workload Manager, see:

<http://h30081.www3.hp.com/products/wlm/docs/wlm.serviceguard.pdf> and
<http://hp.com/products/unix/operating/wlm/info.html#whitepapers>

The following sections will provide more information and insight into the various management tasks performed throughout the cluster lifecycle, from initial installation, to adding applications, establishing application failover, maintaining your cluster, performance, application load balancing, applying patches and troubleshooting.

1.1 Install O/S and Cluster Software

Both Tru64 UNIX and HP-UX 11i users install and configure the O/S and cluster software for a new system installation.

TruCluster Server users create a single initial member cluster and add additional members to the cluster.

HP-UX 11i users must first configure a quorum device, which can be either a cluster lock disk or a quorum service. A quorum service, which is separate from the cluster and is used in place of quorum disks, is not very resource intensive. Therefore, a low-end server is typically used to host the service. HP provides Quorum Server software that runs on both HP-UX 11i and Linux systems. After the storage infrastructure has been configured, the Serviceguard cluster may then be configured to include the specification of the cluster lock disk or quorum service. Once the cluster configuration has been verified, the binary configuration file is distributed automatically to all nodes with the `cmapplyconf` command. Packages and services are configured and verified as the last step. Serviceguard also adds two additional UNIX commands, `cmcp` and `cmexec`, to help with the distribution of files between cluster members and executing programs or commands on cluster members.

HP recommends making the process of rolling out and configuring systems easier by utilizing an Ignite-UX Server to automate the installation. Ignite-UX is an HP-UX 11i administration toolset that helps you install HP-UX 11i on multiple clients on your network, create custom install configurations, or "golden images" to use in multiple installations on clients, recover HP-UX 11i clients remotely,

and manage and monitor multiple client installation sessions. For more information about Ignite-UX, please visit:

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

Ignite-UX can be used in a clustered environment to ensure all members of the cluster have the same “golden image” of the OS, its configurations, and any applications.

To start Serviceguard on a system, you issue the `cmruncl` command or start up the cluster using the Serviceguard Manager GUI. It does not require a booting of the system to make this change. A node can be added to or removed from a Serviceguard cluster without rebooting the system or the cluster.

The HP Serviceguard Manager GUI is a powerful graphics tool, which helps one configure, manage and maintain the Serviceguard cluster and associated packages. HP also offers a for-fee service to help check, manage and monitor cluster configurations called Cluster Consistency Services.

1.2 Application Services

For Tru64 UNIX V5.x, customers use the Cluster Application Availability (CAA) subsystem and for HP-UX 11i, customers use Serviceguard (SG) packages and the Package Manager, to achieve comparable capabilities for application high availability.

For HP-UX 11i with Serviceguard, a package is a collection of resources that are required for the application to run. Examples of resources are: nodes, disk groups, volume groups, file systems, network IP addresses, or services. The package configuration includes the service name, related IP address, related volume groups, policies, run and halt scripts, and application monitors. The Package Manager is responsible for running all application packages on a cluster member.

The Cluster Manager is responsible for initializing the cluster, monitoring the health of the LAN, recognizing node failure, and regulating the reformation of the cluster. Every node in the cluster has a Cluster Manager running on it.

One node in the cluster is the Cluster Coordinator, which is responsible for monitoring the heartbeat from each of the nodes and reforming the cluster if the heartbeat is lost. Reformation of the cluster can include redistribution of packages amongst nodes in the cluster, if necessary.

The Network Monitor looks at the health of networks and recovers from network card adapter failures. Each node can have an active and standby LAN interconnect in each subnet. Cluster membership is maintained by the heartbeat protocol, carried over the LAN via IP addresses.

HP Serviceguard uses a quorum device to prevent a “split brain” situation. In Serviceguard, the quorum device may be a Quorum Service, an Arbitrator, or a cluster lock disk, which is different from how the quorum disk is used in a TruCluster Server environment. Either a cluster lock disk or a quorum service is required to deal with tiebreaker situations (not a vote) for determining which nodes remain in a two-node Serviceguard cluster. One of these mechanisms is also needed to allow a one-node cluster to work. Unlike TruCluster Server, Serviceguard is based on the concept

of dynamic quorum, which allows the cluster to survive more multiple failures and remain running without human intervention to change the quorum votes.

A quorum service, which can be used in place of a cluster lock disk resides outside of the cluster and provides quorum services to the cluster. It can be any HP-UX 11i or Linux server running the HP Quorum Service software. All members in the cluster establish and maintain a connection to the quorum server. An individual quorum server can support up to 50 clusters. For even higher availability, HP recommends setting up the quorum service as a highly available package on a separate highly available Serviceguard cluster.

1.3 Storage Management and Clusters

While administering storage for TruCluster Server is different from the way storage is administered for HP-UX 11i Serviceguard, there are several steps in the process that are common to both solutions. For the purposes of this discussion, we will assume that both clusters are connected to an HP Enterprise Virtual Array (EVA), which is managed by an OpenView Storage Management Appliance already connected to the network, has discovered its storage and is ready for use.

In both environments, the operations performed from the Storage Management Appliance are identical. Virtual disks must be presented to each host that is to have available access. Once the LUNs are presented, the different operating systems have their own method of configuring the storage.

TruCluster Server

Tru64 UNIX V5 will “autodetect” and “autoname” the LUNs as they become visible to the hosts, and by nature of the Tru64 UNIX Cluster File System (CFS), propagate the device names to all cluster members, assigning each a unique global device id. If logical volumes are to house file systems, Logical Storage Manager (LSM) must first be used to create them from one of the cluster nodes. If not using logical volumes, the LSM step is simply skipped. The remaining steps involve creating the AdvFS domain using the `mkfdmn` command, creating the filesets within the domain using the `mkfset` command, creating the mount points, and mounting the related storage. The TruCluster Server CFS will ensure that the storage is seen by all nodes in the cluster.

Serviceguard clusters

When using HP-UX 11i and Serviceguard, the procedure becomes more involved. In addition, one may achieve different results depending on what solution is chosen. The current release of HP-UX 11i requires the installation of SecurePath 3.0D, along with the respective platform kit for your hardware before connecting to the EVA storage. Once the storage array is connected and the LUNs are presented to the cluster nodes, the HP-UX 11i command `ioscan -fnC disk` should be executed to ensure that the new logical devices are seen by each required host. (HSV110 will be in the description) Running the `insf -e` command from each host will create the required device files on each system.

Next, the LUN storage now visible to HP-UX 11i must be configured using the VERITAS Volume Manager (VxVM). When utilizing VxVM, the disk group device file for a LUN is created using the `mknod` command, and the logical disk is brought under control of the Volume Manager by using the `vxdisksetup` command.

The logical disk group is then created using the `vxdiskadd` command, and subsequent logical volumes are created using the `vxassist` command. An OnlineJFS (a.k.a. VERITAS File System - VxFS) file system is created on the logical volume using the `newfs` command, mount points are created on each node, and the device is then mounted. The disk group needs to be “cluster-enabled” using the `vxchg` command before using them in a Serviceguard cluster. At this point, the storage remains exclusively mounted on one system at a time. The “cluster-enabled” storage needs to be defined within a Serviceguard package to ensure that it follows the application to another server in the event of a failover.

To create a solution today that more closely resembles the TruCluster Server architecture and the Cluster File System (CFS) available under Tru64 UNIX, HP recommends that the VERITAS Cluster Volume Manager (CVM) product be used. It extends the capability of the VxVM volume manager to enable file system accessibility from multiple servers simultaneously, via a cluster-wide logical device naming scheme. Data may be mounted up cluster-wide or locally, depending on your application requirements and all servers have a consistent logical view of the storage volumes.

1.4 Application Deployment, High Availability and Load Balancing

Important clustering features that should be pointed out regarding the implementation of applications are: ease of deployment (and management), effectiveness of load balancing, and the ability to manage recovery from a failure (high availability). In total, clustering capability is measured against how effectively these aspects are implemented. The similarities and differences relating to this capability, as provided for in the Tru64 UNIX AlphaServer and HP-UX 11i Integrity server environments, are discussed in the following sections.

1.4.1 Application Deployment

In the case of either the Tru64 UNIX AlphaServer platform or the HP-UX 11i Integrity server platform, the process of deploying an application on a single node/system is comparable, with any differences being inconsequential.

Deploying an application across a multiple node cluster for the purpose of high availability alone (for recovery from failure) differs from deploying a cluster application for the purpose of both high availability and load balancing.

In the Tru64 UNIX Cluster Application Availability (CAA) environment, both requirements are fulfilled automatically through the features provided by the cluster software. The application may be installed once for the cluster on shared allocated storage, and is “seen” by all members of the cluster by virtue of the Cluster File System (CFS). Rolling upgrades of applications and propagation of OS tuning parameters are simplified due to this capability. Once the application has been configured, a CAA resource profile may be created using the `caa_profile` command or the SYSMAN GUI, the profile edited to include parameters defining the way the application will behave, and a script prepared to launch the application.

In the HP-UX 11i Serviceguard environment (with no extensions), the application and application data may be installed on shared storage that will move with the application package, or on local storage dedicated to each individual node. In some cases, Serviceguard extensions are available for specialized applications like Oracle9i RAC and SAP. Following the installation of the SGeRAC add-on software for Oracle9i, the cluster is started up, and the Oracle9i software is deployed to all the cluster nodes by the Oracle installer program. Either way, the application is copied to a file system on each node, as a Cluster File System (CFS) is not a currently a component of Serviceguard. Therefore, only one cluster node at a time is able to have access to a given application file system. With HP Enterprise Cluster Master Toolkit, an easier, more automated method for deployment of applications across the cluster is provided. Rolling upgrades of applications across the cluster, while maintaining application availability, are also supported in the Serviceguard environment.

In Q3CY2005, the VERITAS Cluster File System (CFS) product is planned to be integrated with Serviceguard to form a new bundle (HP Serviceguard integrated with VERITAS Storage Foundation with CFS). HP recommends this product bundle if multiple-node concurrent file system access is an application requirement. The VERITAS CFS product will provide functionality comparable to that of the Tru64 UNIX CFS environment with regard to application deployment across a cluster file system, with the exception of root file system based applications. Application deployment that relied on a shared root file system in TruCluster Server environments would require modification. The cross-node propagation of application related OS tuning parameters may be handled within HP-UX 11i by using the `kconfig -e <filename>` command to output those parameters to a file which in turn may be used as an input template into the `kconfig -i <filename>` command on another node. A script may be constructed to perform this work in a more automated fashion. Alternatively, a command to adjust a single kernel tunable value (`kctune`) can be applied to all members of the cluster using `cmexec` or HP SIM's `mxexec` feature.

1.4.2 Application High Availability (HA)

In general, application high availability refers to the requirement that an application remain usable even in the event that it experiences a failure on one system, regardless of the reason for that failure. Other phrases that may be used synonymously for this requirement include: "recovery from application failure" and "application failover". The concept of high availability is also directly related the concept of No Single Point of Failure (NSPoF).

The concept of application high availability requires some degree of clustering capability, and thus is handled somewhat differently in the Tru64 UNIX CAA and HP-UX11i Serviceguard environments. Facilities designed to create, manage, and set preferred placement of single instance applications, although functionally comparable, will likely differ in the implementation of that functionality.

From an application high availability standpoint, one may draw an equivalent comparison of functionality between the Cluster Application Availability (CAA) component of Tru64 UNIX TruCluster Server and HP-UX 11i Serviceguard. Both Tru64 UNIX CAA and HP-UX 11i Serviceguard are designed to provide application recovery and protection from a wide variety of hardware and software failures. In the case of either product, automatic detection of various types of failures is provided.

In the Tru64 Unix environment, CAA provides system and application failover from the primary node to a secondary node. In general, Serviceguard provides system failover times of 30-45 seconds. The HP-UX 11i Serviceguard/SGeFF (Serviceguard Extension for Faster Failover) solution shortens that time and provides system failover in 5 seconds. Tru64 UNIX customers should evaluate their current system failover times to determine whether Serviceguard or SGeFF is better suited to their needs. Individual application restart timings will vary and should be considered as well.

HP Serviceguard is a reliable and robust application high availability product. It also affords:

- The ability to survive multiple node failures
- The ability to cluster 16 nodes
- Fast fail-back for applications re-starting on the primary active node
- Configurable, individual, application package failover
- Integration with the HP-UX 11i Workload Manager to ensure that service levels are maintained during planned and unplanned downtime
- Multiple high availability cluster configurations for flexibility; active-active, active-standby, and rotating-standby
- An intuitive graphical user interface, HP Serviceguard Manager, for managing multiple high availability clusters
- Integration with HP Utility Pricing portfolio
- Even faster recovery time with the 2-node Serviceguard Extension for Faster Failover product
- Global application disaster recovery through extended Metrocluster and Continentalclusters (long distance)

Once a base application has been installed and configured, a Serviceguard high availability package may be created for it which will allow for application failover in the event of an application failure detected on the primary node. The package may be created graphically using Serviceguard Manager, or manually by:

- Creating a package template file
- Editing the configuration file
- Editing the package control script (run and halt)
- Applying the configuration file to the cluster
- Starting the package
- Modifying the control file to eliminate any errors

Although comparable functionality is provided by both Tru64 UNIX and HP-UX 11i application high availability solutions, the functionality is achieved using a different approach, and moving from one to another will most likely require some level of modification. For example, the syntax used within HP-UX 11i Serviceguard failover packages and scripts will be different from the syntax used by Tru64 UNIX CAA failover scripts. Additionally, in some cases, syntax will need to be added to handle file system re-mounting and failover.

1.4.3 Application Load Balancing

While examining the application load balancing capabilities of Tru64 UNIX TruCluster Server, HP-UX 11i Serviceguard (and other related components), and VERITAS Cluster File Systems, we again

find that there is much comparable functionality, along with some recognizably different approaches. Even the comparable functionality may be implemented in a different manner, requiring some level of adaptation and/or modification when transitioning from a Tru64 UNIX CAA environment.

In general, application load balancing involves distributing the workload across processor, memory, I/O, and network resources. Such a workload distribution may occur across hard or virtual partitions within the same physical machine (cluster-in-a-box), across physical cluster nodes, or across a combination of both partitions and nodes.

The Tru64 UNIX TruCluster Server environment has a built-in load balancing capability, enabling all the partitions and/or systems in the cluster to share a workload, effectively increasing processing capacity with the number of servers included in the cluster. With the included cluster-alias IP address feature, client access to an application executing on all the cluster nodes may be load-balanced automatically. However, many networks and applications servers are already set up to distribute the application workload and related database server requests across multiple nodes. This being the case, the Tru64 UNIX cluster alias feature is often unnecessary and goes unused. HP-UX 11i Serviceguard, in combination with the HP Auto Port Aggregation (APA) product, offers a comparable capability in that a virtual common IP address may be established for multiple nodes. This is considered comparable functionality implemented using a slightly different approach.

The Tru64 UNIX TruCluster Server implementation also offers a shared cluster-wide file system, cluster-wide ownership of storage devices, and thus a single view of I/O resources devoted to the entire cluster (for more details, please refer to the Cluster section and Storage Integration section found in the related document entitled: "Transitioning Tru64 UNIX V5.x Customers to HP-UX 11i v2: Assessment of Changes to the HP-UX11i Roadmap in Terms of File Systems, Clusters, and Infrastructure"). This allows for a relatively seamless and automatic load balancing capability, especially with regard to I/O heavy applications such as Oracle database management software.

Currently HP-UX 11i Serviceguard, in combination with Serviceguard extensions like SGeRAC for Oracle9i RAC, provides a comparable cluster-shared I/O capability for commonly connected storage arrays when using raw devices and raw device based logical volumes, but not when using file systems. A comparable capability is available for file systems on the HP-UX 11i HP 9000 PA-RISC platform in the form of the VERITAS Cluster File System (CFS) product. VERITAS CFS is planned to be integrated with Serviceguard for the HP-UX 11i Integrity server platform in Q3CY2005 (HP Serviceguard integrated with VERITAS Storage Foundation with CFS product bundle).

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.

When HP-UX 11i Workload Manager (WLM) and Process Resource Manager (PRM) are integrated with Serviceguard (SG), load balancing capabilities are provided for:

- Load balancing across hard partitions (nPars), and virtual partitions (vPars) in a cluster-in-a-box scenario

- Re-balancing the workload on the remaining active host following a failover in an active-active scenario

Load balancing across hard partitions (nPars) is currently available using SG/WLM/PRM integration for HP-UX 11i on both the HP 9000 PA-RISC and HP Integrity server platforms. In the future, vPars and SG/WLM/PRM integration for vPars will be supported for HP-UX 11i on the HP Integrity server platform as well.

The capability to dynamically re-balance the application workload on the remaining active host after a failover occurs is currently available for HP-UX 11i on both the HP 9000 PA-RISC and Integrity server platforms. This capability allows load balancing priorities and resource allocation to be dynamically adjusted on the active server picking up the additional workload from a failed node. Thus, the active server resources may be re-allocated to maintain application workload priorities required by established Service Level Objectives (SLOs). Note that, in order to meet the resource demands that may arise as a result of failover workload re-balancing, the HP iCAP (Instant Capacity) feature, which will activate spare CPUs, may need to be deployed.

1.5 Maintaining Cluster Applications

Maintaining applications under Serviceguard is comparable to the way applications are maintained under TruCluster Server. The procedures used to maintain existing cluster applications depends on how you have installed and configured the application, either as a local copy, a shared single-instance configuration, or a cluster wide application using the VERITAS Cluster File System (CFS). One could upgrade a local copy of the application on a server that is not currently hosting the package, and then fail the package over to the newly updated server to grant users immediate access to the new version. If the application was installed on shared storage, both as exclusive storage and also on a CFS volume, a new version would need to be installed into another directory, perhaps residing on the same volume and using links to switch over to the newer version in order to avoid downtime.

Adding storage to applications will also be comparable in a Serviceguard environment as it is in a TruCluster Server environment, especially if Logical Storage Manager (LSM) was used as the volume manager under Tru64 UNIX. The VERITAS File System (VxFS) and Storage Manager (VxVM) offer comparable ways to increase/expand the storage available for a given application.

When the need arises to change how an application fails over from one node to another, failover packages constructed using Serviceguard provide the flexibility of using CLI commands similar to those used under Tru64 UNIX Cluster Application Availability (CAA) while also offering an alternate method of using the Serviceguard Manager GUI. Serviceguard Manager makes use of visual cues like color modification, which enable one to immediately determine if a package has the ability to fail over to another node.

Right-clicking on a specific package defined within Serviceguard will produce check boxes allowing one to re-enable failover capabilities to a specific node, as well as the ability to halt or relocate the package. When integrated with WLM, resource allocations on the member starting the package can be adjusted, and additional resources like idle CPU's can be automatically allocated where they are needed most.

1.6 Periodic Cluster Maintenance

Tru64 UNIX users can manage a single cluster through a CLI or a GUI (SysMan).

HP-UX 11i users can manage a single cluster through a CLI or a GUI (Serviceguard Manager). They can also take advantage of the multiple cluster capabilities of Serviceguard Manager to manage many clusters from a single interface. Using Serviceguard Manager, administrators and operators can see color-coded, graphically intuitive icons in order to get the big picture view of multiple clusters, so that they may proactively manage many clusters, nodes and applications.

The Serviceguard Manager GUI, which can help to automate the management and monitoring tasks of multiple clusters simultaneously, offers Tru64 UNIX users new functionality. Other aspects of the system may be managed using other web-based tools provided by HP, such as kcvweb for kernel tuning, and pdvweb for device and peripheral maintenance.

There also exists integration with HP Systems Insight Manager for cluster discovery and launching Serviceguard Manager, as well as inventory and monitoring across the enterprise. Below is a basic comparison between some of the respective CLI commands:

Tru64 UNIX CAA	Serviceguard	Description
caa_stat	cmviewcl	Provides status on the current state of the cluster members and services
caa_profile	cmmodpkg	manages resource/package attributes
caa_register	cmmakepkg	registers a resource/package
caa_start	cmrunpkg	starts a resource/package
caa_relocate	cmrunpkg& cmhaltpkg	relocates an application
caa_stop	cmhaltpkg	stops a resource/package
caa_unregister	cmmodpkg	removes an application from CAA control
No equivalent	cmcp	copies files to all nodes, or a subset of all nodes in a cluster
No equivalent	cmexec	runs a command or application on designated nodes in a cluster

1.7 Monitoring System Performance

Tru64 UNIX users have typically used the Tru64 UNIX Collect tool, or third party software like BMC Patrol, to monitor system performance.

HP-UX 11i users have access to the GlancePlus and OpenView Performance Monitor tools to monitor performance. GlancePlus is a product that displays information about the current CPU load and usage of physical and virtual memory. GlancePlus Pak integrates GlancePlus and HP OpenView Performance Agent for HP-UX 11i (OVPA) products into a single tool to help customers better manage the performance and availability of their servers. GlancePlus Pak is bundled with HP-UX 11i in the Enterprise and Mission Critical Operating Environments. HP-UX 11i Event Monitoring Services (EMS) can be used to monitor individual system resources and trigger package failover, when resource utilization reaches critical user defined values.

1.8 Monitoring Users

Tru64 UNIX customers typically use manual processes or scripts to monitor users on the cluster. They can also use the Event Management (EVM) subsystem.

HP-UX 11i customers use Glance or scripts to monitor users on the cluster.

1.9 OS Version Upgrades and Patch Installation

Tru64 UNIX and HP-UX 11i are comparable in their support for operating system upgrades and patch installations in clustered environments. Both support rolling upgrades of the base operating system, clustering software, and patches.

Tru64 UNIX rolling upgrades are fairly extensive and time-consuming due to the integration of the cluster file system across all cluster members. HP-UX 11i rolling upgrades are more streamlined and performed on an individual node basis. First, node 1 on the cluster is halted and its packages relocated to another node. Next, node 1 is upgraded and patched. After a reboot, node 1 on the cluster is restarted and its native packages are relocated back. The same steps are then repeated across each of the nodes in the cluster.

In some situations (such as upgrades of earlier operating system versions to newer versions) it may be necessary to do a full (Tru64 UNIX) or cold (HP-UX 11i) installation instead of a rolling upgrade. HP-UX 11i cold installations also support the use of "Golden images" that can contain the operating system, patches and implementation specific customizations. The release notes for both operating systems document these situations in detail.

1.10 Adding a Member Node to the Cluster (Extending a Cluster)

TruCluster Server and Serviceguard are comparable in their support for the addition of nodes to a running cluster. Node additions to a halted cluster are not supported by either product. Both support the addition of new nodes through command line utilities. Serviceguard also permits cluster node additions to be performed through the use of the Serviceguard Manager graphical interface.

TruCluster Server supports a maximum of 8 nodes while a Serviceguard cluster can contain up to 16 nodes.

Adding a node to an existing TruCluster Server cluster involves:

- Configuring the new node hardware
- Updating the console firmware
- Disabling of auto boot
- Adding the new node from another node
- Booting the new node
- Creating on-disk copies of configuration files

For a Serviceguard cluster the steps to add a node requires:

- Configuring the new node hardware
- Updating the console firmware
- Installing HP-UX 11i, patches, Serviceguard bundle
- Installing Serviceguard extensions, in specific cases
- Updating the cluster configuration
- Applying the new configuration

Although Serviceguard requires HP-UX 11i patches and Serviceguard to be installed on each new node that is added to the cluster, this process can be automated via Ignite-UX and the use of Golden Images.

TruCluster Server and Serviceguard are also comparable in their support for the temporary addition of nodes to a cluster to perform hardware maintenance or to handle short term increases in workloads. In hardware maintenance scenarios, a temporary node can be added to the cluster and the services and applications of the node to be worked on relocated to it. After the maintenance work has been completed, the node being serviced can be brought back on-line, its services and applications relocated back to it, and the temporary node removed from the cluster. A similar approach may be utilized to handle short-term temporary increases in workloads with the temporary node used to off-set such temporary workloads, and then removed after the increased workload period subsides.

1.11 Cluster Troubleshooting

Troubleshooting problems on both TruCluster Server and Serviceguard clusters are comparable in some aspects yet different in others. Both require the verification of cluster operations and the monitoring of hardware resources. The use of system log files for the reporting of disk, CPU, memory, network and power events is common to both products. However, since the underlying architectures and infrastructures are different, the tools used and the approaches taken to resolve problems that occur differ.

TruCluster Server and Serviceguard both support the monitoring of cluster operations and hardware through command line utilities and graphical interfaces. The Serviceguard Manager graphical interface provides extensive cluster troubleshooting to be performed while the SysMan Menu TruCluster Server interface is primarily intended to provide cluster status.

Cluster event monitoring is comparable for both cluster technologies. TruCluster Server events can be monitored through the Event Management (EVM) subsystem and Serviceguard events through the Event Monitoring System (EMS). Both products log events as they occur and can be configured to trigger other events or actions in response.

Serviceguard clusters may also be proactively monitored through the use of the HP Predictive product. HP Predictive is run on all nodes and gathers status on devices reporting nonfatal errors that have been accumulating over time. These errors can be configured to be automatically escalated to the appropriate support personnel.

For related Tru64 UNIX TruCluster Server documentation, please visit:

http://www-unix.zk3.dec.com:8083/docs/pub_page/cluster51B_list.html

For related HP-UX 11i Serviceguard documentation, please visit:

<http://docs.hp.com/hpux/pdf/B3936-90079.pdf> and
<http://docs.hp.com/hpux/ha/index.html#ServiceGuard>

2. Applications – NFS Servers and TCP Services

If you are managing TCP or NFS services from one node in a cluster, then you may continue to do so in a comparable fashion with HP-UX 11i v2. HP provides additional software called the HA NFS Toolkit, which makes single or multiple NFS servers highly available. Multiple NFS servers can be configured to back each other up in a “cross-mount” fashion. This solution is a high-availability solution and not a load-balancing solution.

If using cluster alias to enable one common IP address for your cluster, then the approach will need to be modified or switched to alternative mechanisms under HP-UX 11i v2. The HP-UX 11i Serviceguard package with a re-locatable IP address provides an IP address to one package of one or more applications, and not necessarily for the entire cluster. (Please refer to the Cluster Alias section in the related document entitled: “Transitioning Tru64 UNIX V5.x Customers to HP-UX 11i v2: Assessment of Changes to the HP-UX 11i Roadmap in Terms of File Systems, Clusters, and Infrastructure”).

If you are using TruCluster Server for dynamic load balancing, one option would be to split your TCP or NFS workloads across multiple members in the cluster. Some larger TruCluster Server customers, which had exceeded their Cluster Interconnect bandwidth, have already used this approach. HP-UX 11i Serviceguard failover capabilities allow the affiliation of an IP address with a Serviceguard package. However, this Serviceguard package can only run on one node at a time. (It could be used for export connections, FTP, telnet, etc.) Alternatively, something similar to a smart load-balancing hardware device for FTP/telnet connections could be used. Another option would be to match the capacity of the individual node server to the size of the workload.

In order to implement an effective scalable solution, the workload will need to be partitioned as Serviceguard packages. For more details on how to design and manage NFS in a Serviceguard environment, please refer to:

<http://docs.hp.com/hpux/pdf/B5140-90020.pdf>

For an overview on how to manage Highly Available NFS, please refer to:

<http://docs.hp.com/hpux/pdf/B5125-90001.pdf>

Connections between Serviceguard nodes can be done manually, via scripts or through the use of a Serviceguard package with a re-locatable IP address.

If near instantaneous node failover times in the event of a catastrophic loss situation are required, consider using Serviceguard Extensions for Faster Failover (SGeFF), a new Serviceguard add-on available from HP. SGeFF functionality is available only for two-node Serviceguard clusters and requires the use of a separate Quorum Service and an additional dedicated heartbeat LAN. SGeFF is designed to reduce system/node failover times to approximately 5 seconds (this does not include application restart time on secondary node). For more details on SGeFF, please refer to:

<http://docs.hp.com/hpux/onlinedocs/T2389-90001/T2389-90001.html>

Other ways that Tru64 UNIX TruCluster Server systems are currently being used must be evaluated as well, such as cluster-wide “Login Farm”, Print, and NFS servers.

If you are using a TruCluster Server system as a “Login Farm” for the majority of your existing users,

you may need to look at modifying your approach, especially in the areas of how users gain access to all cluster members, as well as security access to project data and files. If NIS or LDAP are not being used, you should consider implementing it to ensure near functional equivalence.

If you are using a TruCluster Server system as a Cluster-wide Print Server, you may need to modify your approach such that print queues are managed on a per member basis or consider using Samba with your Serviceguard cluster.

Customers will be able to use the VERITAS CFS product to access files from any node in the cluster. Files that are not mounted with the `-g` option will have to be accessed from the node on which they reside.

For NFS servers, note that only file systems that are shared by all nodes in the cluster can be exported.

For managing a single security domain, use LDAP as a solution with HP-UX 11i V2. With the use of LDAP, one will need to look at the security requirements, scalability of the solution, and capabilities of LDAP. Alternatively, security files will need to be managed on a per cluster member basis. HP customers can also use HP SIM's role-based access control capability where users can create rules on who can do what and apply these security rules to all members of the cluster.

With HP-UX 11i, a single shared copy of the `/root` file system for all Serviceguard cluster member nodes is not available. Users will need to connect to the correct node to access the local cluster member file. Similar cluster management functions can be scripted and automated via HP SIM, and reused on any cluster member. Users will need to be aware of the node they are pointing to for system administration tasks and locating cluster member specific files.

If using an NFS Server to distribute OS patches or configuration information, Ignite-UX and other HP configuration management tools may be used instead.

3. Backup and Restore

Overall, the backup and recovery product features available for HP-UX 11i on the HP 9000 and Integrity server platforms are comparable to those found in the Tru64 UNIX environment. There are differences and timing issues with regard to some of the more common backup and recovery software products which will be discussed below. For differences in backup/recovery hardware device compatibility across platforms, please refer to the Tru64 UNIX AlphaServer to HP-UX 11i Integrity server Platform Transition Modules.

3.1 Native File System Backup/Restore Command Set

While the syntax/options of the commands may differ, the essential backup/restore functionality of the Tru64 UNIX and HP-UX 11i native commands is comparable.

The tables below illustrate the majority of backup and restore native commands that exist for Tru64 UNIX and HP-UX 11i.

OS/FS	Tru64 UNIX AdvFS	HP-UX 11i OnlineJFS VxFS
Backup Command	<code>vdump</code> or <code>rvdump</code>	<code>vxdump</code>

Restore Command	vrestore or rvrestore	vxrestore
-----------------	-----------------------	-----------

OS/FS	Tru64 UNIX UFS/NFS	HP-UX 11i HFS/NFS
Backup Command	dump	dump
Restore Command	restore	restore

3.2 Third Party (ISV) and HP Backup/Restore Products

In the case of many customer IT environments, functional requirements of the backup and recovery strategy are implemented through the use of third-party (ISV) software products. Such products provide varying degrees of success in automating, facilitating and enhancing backup/recovery functionality. Availability and compatibility of some of the more popular backup/recovery software products will be addressed herein.

3.2.1 Legato NetWorker (Power Edition, etc.)

The Legato family of backup/recovery products is fully supported on the Tru64 UNIX AlphaServer, HP-UX 11i HP 9000 PA-RISC, and HP-UX 11i Integrity server platforms. This being the case, functionality is comparable and relatively uniform in external implementation across the platforms supported, making the Legato NetWorker product an excellent product choice through transition.

3.2.2 VERITAS NetBackup Data Center Edition

Currently, VERITAS NetBackup Data Center Edition is fully supported for the Tru64 UNIX AlphaServer and HP-UX 11i HP 9000 (PA-RISC) platforms. It can be tightly integrated with other VERITAS software components (VxFS, VxVM) as well (for HP-UX 11i HP 9000 PA-RISC).

Only the client component of this software is currently supported for HP-UX 11i on the Integrity server platform. However, it is planned and expected that VERITAS NetBackup Data Center Edition (client/server) will be fully supported on the HP-UX 11i Integrity server platform in Q2CY2005.

3.2.3 VERITAS NetBackup Cluster Edition

Currently, VERITAS NetBackup Cluster Edition is fully supported for the Tru64 UNIX AlphaServer TruCluster Server and HP-UX 11i HP 9000 PA-RISC VERITAS Cluster Server environments. In the future, it may be tightly integrated with other VERITAS software components (CFS, CVM) as well for HP-UX 11i HP 9000 PA-RISC and HP Integrity server platforms.

The VERITAS NetBackup solution may be configured as a package in a Serviceguard cluster to provide a highly available backup solution.

This software product is not currently supported for the HP-UX 11i V2 Integrity server platform. However, full support for the HP-UX 11i Integrity server platform is planned for Q3CY2005.

3.2.4 Syncsort Backup Express

The Syncsort Backup Express product is supported on Tru64 UNIX AlphaServer, HP-UX 11i HP 9000 PA-RISC, and HP-UX 11i Integrity server platforms. It supports raw device and file system backup/restore functions and supports database backup/recovery for Oracle, Sybase, Informix and MS SQL Server. It currently supports cluster backup/recovery support for the Tru64 UNIX AlphaServer platform, but not for HP-UX 11i platforms.

Hardware backup/restore device compatibility for Syncsort Backup Express may be found at the URL:

<http://www.syncsort.com/bex/bexlib.htm>

3.2.5 HP OpenView Data Protector 5.1

Data Protector 5.1 supports comparable feature/functionality for the Tru64 UNIX AlphaServer, HP-UX 11i HP 9000 PA-RISC, and HP-UX 11i Integrity server platforms. For more compatibility details, please refer to URL:

http://www.openview.hp.com/products/storage_data_protector/device_matrices/Platform_Integrtn_SptMtx_DP51.pdf

Conclusions

As exhibited by this document, there are a number of ways in which a desired end-solution may be implemented on the HP Integrity server platform using the HP-UX 11i family of products, and/or available third-party (ISV) software products. It is true that some desired system administration and system management solutions may be implemented in a different manner, or by using a slightly different approach. However, it is important to take note of the benefits of the breadth of new functionality, flexibility, and reliability brought to bear through the use of the full arsenal of HP-UX 11i software products, especially in the areas of High Availability, Workload Management, Computing Resource Management, and Data Center Management. As HP-UX 11i products continue to evolve with the concept of shared virtualization as a goal, so will the benefits and advantages those products bring to the efficient, cost-conscious, well-managed data center.

Appendix A: References

The following URLs are those already listed in the body and Appendix A of the partner paper, “Transitioning Tru64 UNIX V5.x Customers to HP-UX 11i v2: Assessment of Changes to the HP-UX 11i Roadmap in Terms of File Systems, Clusters, and Infrastructure.” 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>

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

o HP Systems Management Tools

- HP System Management Tools Overview
<http://h71028.www7.hp.com/ERC/downloads/5982-6788EN.pdf>
- HP Systems Insight Manager
<http://h18013.www1.hp.com/products/servers/management/hpsim/index.html?jumpid=go/hpsim>
- HP Serviceguard and HP-UX 11i Workload Manager
<http://h30081.www3.hp.com/products/wlm/docs/wlm.serviceguard.pdf>
<http://hp.com/products/unix/operating/wlm/info.html#whitepapers>
- HP Ignite-UX
<http://software.hp.com/products/IUX/>

o HP Cluster Solutions

- Tru64 UNIX TruCluster Server
http://www-unix.zk3.dec.com:8083/docs/pub_page/cluster51B_list.html
- HP-UX 11i Serviceguard
<http://docs.hp.com/hpux/pdf/B3936-90079.pdf>
<http://docs.hp.com/hpux/ha/index.html#ServiceGuard>

o Cluster Applications

- NFS and the Serviceguard Environment
<http://docs.hp.com/hpux/pdf/B5140-90020.pdf>
- Managing Highly Available NFS
<http://docs.hp.com/hpux/pdf/B5125-90001.pdf>
- HP Serviceguard Extensions for Faster Failover (SGeFF)
<http://docs.hp.com/hpux/onlinedocs/T2389-90001/T2389-90001.html>

o Third Party (ISV) and HP Backup/Restore Product Hardware Compatibility

- Syncsort Backup Express
<http://www.syncsort.com/bex/bexlib.htm>
- HP OpenView Data Protector 5.1
http://www.openview.hp.com/products/storage_data_protector/device_matrices/Platform_Integrtn_SptMtx_DP51.pdf

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.

Intel and Itanium are trademarks or registered trademark of Intel Corporation in the United States and other countries and are used under license.

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.

LEGATO and LEGATO NetWorker are trademarks or registered trademarks of LEGATO Systems, Inc.

Syncsort and Syncsort Backup Express are trademarks or registered trademarks of Syncsort, Inc.

11/2004