HP Serviceguard for Linux®

Data integrity and application availability for business-critical environments

Executive summary

Enterprises worldwide are embracing the Linux operating system. In a business environment of increasing expectations and decreasing budgets, these companies regard Linux as the way to deliver highly available systems and applications at a lower cost than possible with other operating systems.

HP Serviceguard for Linux enhances Linux capabilities by providing high availability to facilitate managing planned and unplanned downtime. With built-in safeguards against data corruption and the ability to extend protection across geographical boundaries, HP Serviceguard for Linux can be a critical component of the enterprise IT infrastructure. The product leverages the breadth and depth of experience HP has acquired in shipping more than 100,000 server licenses for mission-critical Serviceguard environments. HP Serviceguard for Linux demonstrates HP's leadership in providing enterprise solutions and its commitment to helping its customers maximize return on their information technology investments.

Introduction

This white paper presents the capabilities and benefits of HP Serviceguard for Linux—a solution for high availability clustered computing environments that offers features and functionality similar to HP Serviceguard for the HP-UX platform.

The paper begins with a definition of HP Serviceguard for Linux, makes high availability configuration recommendations, discusses the software architecture, and highlights important solution features. It then discusses typical computing environments in which the solution might be deployed.

What is HP Serviceguard for Linux?

Simply stated, HP Serviceguard for Linux is a high availability solution that leverages HP's experience in designing, deploying, and supporting sophisticated business and technical computing environments. HP Serviceguard for Linux brings robust HP-UX technologies to the Linux operating system environment.

Definition

HP Serviceguard for Linux allows users to create high availability server clusters with shared Fibre Channel or SCSI storage solutions. These server clusters enable application services to remain available despite hardware or software failures or the planned downtime required by normal maintenance or system upgrades. A Serviceguard cluster provides the software and hardware redundancy that eliminates disruptive single points of failure. HP Serviceguard for Linux groups application services (or Linux processes) and their resources into highly available packages. In the event of a server, application service, network, or other resource failure, Serviceguard automatically transfers control of that node's package(s) to another node (server) within the cluster, maintaining service availability with minimal interruption. Through an end-to-end system of alerts, error detection, and dynamic resource re-allocation, HP Serviceguard for Linux safeguards cluster operations to provide a high degree of availability.

Benefits

HP Serviceguard for Linux optimizes a clustered IT infrastructure, offering powerful benefits to high availability enterprise computing environments because it:

- Increases service uptime, preserves data integrity, reduces planned downtime
- Protects against disasters and major site outages with advanced, mission-critical capabilities—at the low total cost of ownership of Linux

- Simplifies the configuration, administration, management, and monitoring of clusters
- Easily integrates clustered Linux solutions into current HP-UX environments while quickly building Linux expertise in IT staff

Linux, the Adaptive Enterprise, and HP Serviceguard for Linux

In an Adaptive Enterprise, business and IT are synchronized to capitalize on change. When business and IT are in alignment, change—though constant, unexpected, and frequently disruptive—presents opportunities that can allow a company to maintain and even strengthen its competitive advantage.

The Linux operating system has penetrated the world of enterprise computing with surprising speed. The open source movement was initially viewed as an interesting facet of IT history, but now Linux has been recognized by enterprise IT managers for its exceptional flexibility and stability—and for its ability to deliver powerful, highly available solutions at a lower cost than proprietary operating systems. Companies invest in high availability solutions for applications that require 24x7 availability and rigorous accountability, because having a system unavailable for as little as a few minutes can mean thousands of dollars in lost revenues and customers who go to the competition. Typical Linux use includes enterprise messaging and groupware, network file systems, shared data through Internet and intranet-based services, and business critical databases—often in environments that run multiple operating systems.

The need for an IT infrastructure that supports business agility drives the selection of Linux for high availability computing environments and underlies the importance of solutions like HP Serviceguard for Linux. HP Serviceguard for Linux is designed to support the data integrity and availability requirements of enterprise-class e-commerce, Internet infrastructure, database, financial, and scientific applications. Further, it delivers the simplicity, agility, and value that are the hallmarks of Adaptive Enterprise solutions.

An overview of HP Serviceguard for Linux operation

This section of the white paper presents an overview of the operation of HP Serviceguard for Linux.

Basic concepts

A Serviceguard cluster consists of networked servers. These servers are known as “nodes” and can be servers, blades, or partitions. Nodes run “packages” that are capable of being started, halted, moved, or failed-over. Packages allow Serviceguard to move applications between nodes easily, and they contain all the resources needed by an application.

Typical resources include:

- Relocatable IP addresses
- Volume groups
- Logical volumes
- File systems
- Network resources
- Application or system processes
- Monitored services
- Information about where the package will run next

A package may contain one or a combination of these resources. A cluster can have as many as 150 packages, and a package can contain as many as 30 services.

A package also contains shell scripts and functions to enable starting and halting application processes. Essentially, a package is capable of managing application resources that were previously managed by the system. The package also defines a prioritized list of nodes where the application is permitted to run. Packages may be "data-less," in that no disk volume groups (shared storage) are defined. Otherwise, systems that are designed to run packages are physically cabled to the data disks required by that application.

In the event of a single service, node, network, or other resource failure, Serviceguard automatically transfers control of the package to another node within the cluster, allowing applications and services to remain available with minimal interruption.

Figure 1. A typical two-node cluster configuration

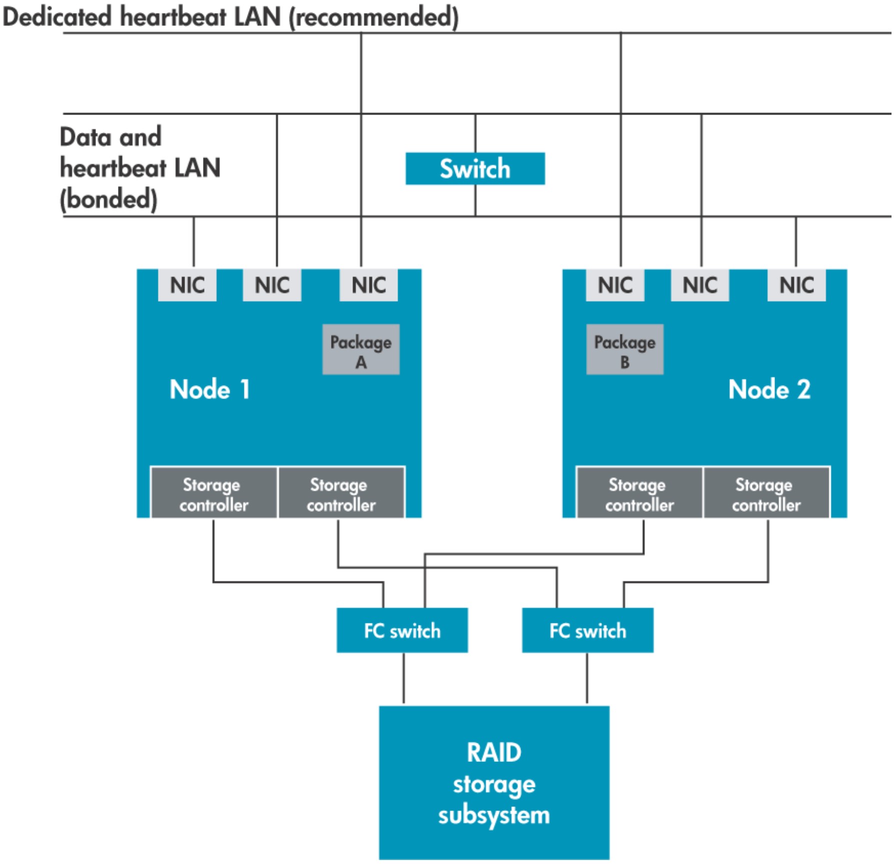

Eliminating single points of failure

To protect critical business applications, the entire environment should be designed to eliminate all single points of failure and to minimize the impact of various component failures. The two major categories addressed in this paper are cluster configuration and software architecture. The cluster configuration section addresses avoiding failures associated with the hardware configuration, and the software architecture section addresses failures associated with the software layer. For a full discussion of high availability theory, please refer to HP's "Providing Open Architecture High Availability Solutions" white paper.

Cluster configuration

In order to provide a high level of availability, a typical cluster uses redundant system components to eliminate single points of failure. In general, as redundancy increases, so does access to applications, data, and services in the event of a failure.

Redundant network components

A number of network components can fail, resulting in lost network connectivity. These include the Network Interface Card (NIC), the network cable, switches, and routers. To survive many of these types of failures, HP recommends the use of redundant networks as shown in Figure 1. Because describing redundancy for an entire network infrastructure is beyond the scope of this paper, please refer to HP's "Highly Available Networks" white paper for more information.

Channel bonding

Network redundancy is implemented by grouping together two or more NICs in a Linux process known as channel bonding. Channel bonding can be configured in either high-availability or load-balancing mode. In the bonded group (high-availability mode), one interface transmits and receives data while the others are available as backups. If one interface fails, another interface in the bonded group takes over. Load-balancing mode allows all interfaces to transmit data in parallel in an active/active arrangement. In this case, high availability is provided because the bond still continues to function (with less throughput) if one of the component LANs fails. HP highly recommends channel bonding in each critical Internet Protocol (IP) subnet in order to achieve highly available network services. Failover from one NIC to another prevents a package from failing-over to another node and is transparent to the application.

Redundant heartbeats

Serviceguard for Linux also supports redundant heartbeats, which can be located on any combination of dedicated heartbeat or "client" networks. A dedicated heartbeat network is less likely than a client network to lose heartbeats when network traffic is high.

Redundant disk storage

Without redundancy, there are a number of ways failure within the disk subsystem of a cluster can reduce availability. Each node in a cluster has its own "root" disk whose failure can be masked by mirroring. In addition to the root disk, each node may be physically connected to several other disks, allowing other nodes to access the data and programs associated with a package for which it is configured. These disks can be in one or more disk arrays. There are a number of components in the path from a node to the storage array and within the array itself that affect availability. These may include the host bus adapter (HBA), a Fibre Channel switch, cables, storage controllers, disks, and power supplies.

Storage paths

Because storage is shared between nodes, loss of the storage system makes data unavailable and causes most services to fail. Therefore, all storage systems must have redundant controllers and power supplies. If a node has only a single path to shared storage, then any failure in that path may cause all packages relying on shared storage to fail. Multiple paths to shared storage are recommended—or required, depending on the storage system—so that storage path connectivity does not require a failover between nodes but instead fails-over to the redundant path.

RAID

The selection of RAID type and the disks that are part of each RAID set is critical to keeping storage highly available. It is important to configure the RAID such that disks in a set do not share the same SCSI bus. Otherwise, the failure of a SCSI controller can cause the failure of the RAID set. Both RAID-1 and RAID-5 are appropriate for high availability configurations, with RAID-1 typically providing higher availability at a higher price.

Redundant power

Cluster power connections are important. Implementing the recommended power configuration will improve the availability of the cluster. In many cases, failure of a server or storage component from loss of AC power can be prevented. This prevents package failover and restart or storage re-mirroring and the associated interruption of service.

In order to eliminate all single points of failure (SPOFs), HP recommends providing redundant power circuits to prevent a single point of failure for nodes, disks, and disk subsystems. Redundant power supplies should be used whenever possible. Backup of nodes, disks, and disk subsystems with an uninterruptible power source (UPS) is also recommended. These configuration techniques can extend hardware availability. To provide consistent high availability, the power source should be configured in such a way that power loss affects 50% or fewer of the nodes in the cluster.

Multiple power distribution units (PDUs) should be used within a rack and connected to at least two different power circuits. When a component such as a server or switch has two power cords, each should be connected to PDUs that are powered from two or three different circuits. If a component has more than two power cords, each should be connected to different PDUs that are powered from separate power circuits. If this is not possible, it is important to connect all the component's power cords to the same PDU. As in the earlier example, potential power loss on one circuit should affect 50% or fewer of the servers in the cluster.

When dealing with redundant PDUs or power circuits, power sizing calculations for a PDU should be done with the assumption that the PDU might have to handle all the power of the devices attached to it. For example, if a server has two power supplies, each connected to different PDU, assume that all power for the server must be supplied by just one of the PDUs to account for various failures.

HP recommends testing a cluster's power redundancy after it is configured. Part of that test should include resilience to power circuit, PDU, or UPS failure. Each of these components should be "failed" in turn to determine the cluster's ability to recover properly.

Because safety is the highest priority, it is important to follow local codes and regulations that relate to power, even if they differ from these recommendations.

When protection against the loss of power at an entire data center is required, Serviceguard for Linux, with Cluster Extension XP (CLX) for Linux, provides multi-site disaster tolerance at distances up to 100 kilometers.

Software architecture

Hardware redundancy cannot mask all failures. Software is needed to identify and control the transfer of applications after a failure. Serviceguard for Linux is a critical component in creating a complete and robust environment for highly available critical applications.

The architecture of HP Serviceguard for Linux consists of feature-rich "managers," or software components that control vital cluster, package, and network functionality.

Figure 2. HP Serviceguard for Linux

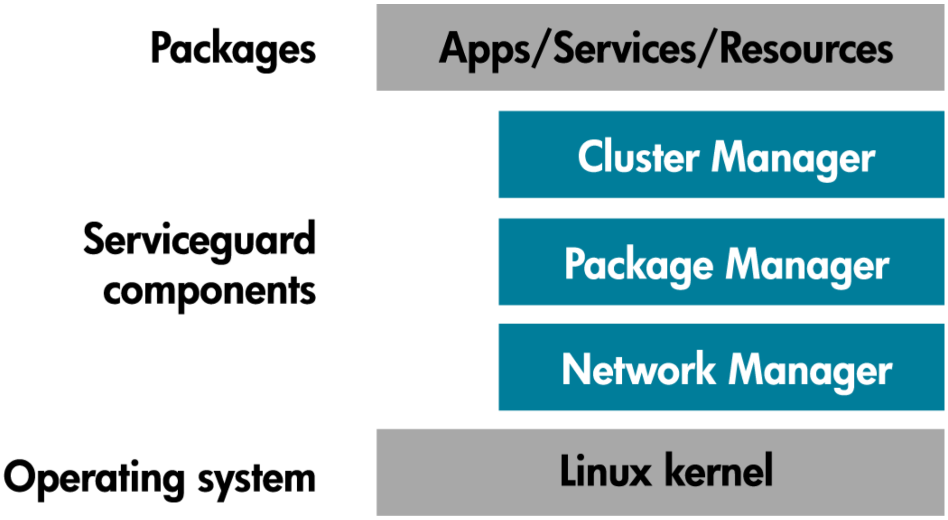

The Cluster Manager

The Cluster Manager component is the key to defining and determining cluster membership.

Cluster Membership

The Cluster Manager initializes a cluster, monitors its health, recognizes node failure, and regulates the re-formation of the cluster when a node joins or leaves it. The Cluster Manager operates as a daemon, or background process, and runs on each node. During cluster startup and re-formation, one node acts as the cluster coordinator. Although all nodes perform some cluster management functions, the cluster coordinator is the central point for inter-node communication.

Split-brain syndrome

Part of the Cluster Manager's responsibilities is to protect against "split-brain syndrome" and protect data integrity. Split-brain syndrome can occur when a two-node cluster, for example, loses all heartbeat connection between the two nodes. Each node attempts to re-form the cluster, with the potential to run two different instances of the same application and corrupt data. Split-brain can also occur in disaster-tolerant clusters, where separate groups of nodes are located in different data centers.

Heartbeat messages

Central to the operation of the Cluster Manager are heartbeat messages. Each node in the cluster exchanges heartbeat messages with the cluster coordinator over each TCP/IP network configured as a heartbeat device. Each node sends its heartbeat message at a rate specified by the cluster heartbeat interval. The rate is established during configuration and contained in the cluster configuration file.

Arbitration

When a configurable number of heartbeats is missed, the Cluster Manager re-forms the cluster. If 50% of the cluster nodes are unaccounted for, a cluster lock is used to arbitrate, or break ties, between nodes during the cluster re-formation process. Essentially, a cluster lock allows the node that acquires it first to re-form the cluster, "locking" other nodes out. Cluster lock can make use of a lock Logical Unit Number (LUN) or a quorum service:

- **Lock LUN**—A lock LUN is created in a 100 KB Linux disk partition. It is not part of a volume group, logical volume, or file system. All nodes in a cluster can write to the cluster lock LUN. When a node obtains the lock, the LUN is marked, indicating to all other nodes that it is "taken." Acquiring the lock follows an arbitration protocol, which adds to failover time. This LUN mark can survive an off-on power cycle of the disk device, unlike some SCSI disk reservation implementations, and is protected via RAID. Once the cluster re-forms, the cluster lock is cleared until needed for future tiebreakers.
- **Quorum service**—A quorum service is an alternative cluster lock that employs a process, running on an external server or PC, as an arbitrator in the event that all network connections between nodes fail and can monitor up to 50 heterogeneous clusters or 100 nodes. The quorum service maintains a special area in memory for each cluster. When a node obtains the cluster lock, all other nodes are locked out. This alternative arbitration method allows implementation of clusters with no shared storage. It also provides mission-critical level protection for disaster-tolerant clusters distributed across sites by allowing a third site to determine which site should have control. This method of arbitration does not add a "single point of failure" because it is used only during the actual arbitration—which is required only when 50% of the cluster nodes are unaccounted for.

  Quorum service systems, when used, should be powered separately from cluster nodes so that when power failures affect some servers in the cluster the cluster can remain available.

After the cluster coordinator and cluster membership are established, information about the new set of nodes within the cluster is passed to the package coordinator. Packages running on nodes that are no longer in the new cluster are transferred to their adoptive nodes.

Deadman timer

Another potential failure can occur when a server does not respond to heartbeats because it is temporarily "hung." When this occurs, the cluster can re-form without the server. However, the server may recover and continue its previous operations as if it were still part of the cluster, potentially corrupting data. The solution to this situation is the "deadman timer" process, which runs at the highest priority and initiates a node reset when the server recovers.

Package Manager

Each node in a cluster runs an instance of the Package Manager. The Package Manager that resides on the cluster coordinator is known as the package coordinator. The package coordinator decides when and where to run, halt, or move packages. The Package Manager on each node executes a user-defined control script that runs and halts packages and package services as well as reacting to changes in the status of monitored resources.

Package configuration

The initial package configuration process defines a set of application services that are run by the Package Manager when a package starts up on a node. The configuration also includes a prioritized list of cluster nodes on which the package can run and the types of failover allowed for the package. Serviceguard A.11.16 allows the configuration process to be done through the Serviceguard Manager graphical user interface (GUI).

A package starts up on an appropriate node when the cluster starts. Package failover occurs when the package coordinator initiates the startup of the package on a new node. A package failover involves both halting the existing package—in the case of a service, network, or resource failure—and starting a new instance of the package.

Package parameters

Serviceguard offers important package flexibility options. Packages can be defined, for example, to create dependencies on other packages, enabling load balancing of system resources or sequencing of package startup/shutdown. After a failover transfer, a package typically remains on the adoptive node and the primary node becomes redundant—a feature referred to as "rotating standby." Because all packages are capable of running on all nodes, if a failure occurs, a package fails-over to the node with the fewest running packages. However, the package can be configured to return to its primary node as soon as the primary node comes becomes available or at another appropriate time—a feature referred to as "automatic failback."

Alternative configurations include traditional active/passive, with one node serving as the backup for other node(s) running packages—or active/active, where all nodes are running packages that access separate data and where nodes have additional capacity to handle failed nodes' packages.

Package control script

Essentially, the behavior of a package—startup, halt, and failover—is under the control of a package control script. This script contains the information necessary to run all the services in a package, monitor them during operation, react to a failure, and halt the package when necessary. Whereas the Cluster Manager monitors node availability, the Package Manager detects and recovers from package-level interruption. Each package must have a separate, executable control script. Separate scripts can be created for run and halt operations. When the package control script is complete, the user propagates it to all nodes of the cluster to provide uniformity of operation. If configured through the Serviceguard Manager GUI, package distribution is automated and a consistency check of the packages is performed. In addition, HP provides toolkits for common applications to simplify their deployment in a Serviceguard for Linux cluster and provide ongoing support and maintenance.

Network Manager

Networks remain highly available through the standard Linux kernel bonding capability. The Network Manager detects network cards and cable failures and migrates packages to a designated alternative node.

Each node, or host system, has an IP address for each active network interface. This address is known as a stationary IP address, is not associated with packages, and is not transferable to another node. Stationary IP addresses transmit heartbeat messages, as described in the "Cluster Manager" section of this white paper.

In addition to the stationary IP address, each package is normally assigned one (or more) unique IP address, called a "relocatable IP address" or "package IP address" because it can move from one node to another.

The Network Manager makes provisions for load balancing. Individual services can be placed in separate packages with unique IP addresses, making it possible to shift them quickly and easily to less burdened systems.

Failover scenarios

HP Serviceguard for Linux clusters takes action based on the failure of:

- The system (node)
- The network (heartbeat and/or user)

- The application (services)
- The storage connectivity (if used)
- Other resources (as configured)

Examples of the support HP Serviceguard for Linux provides to clusters are as follows:

- If a local LAN fails, the supported Linux bonding facility provides failover to a local standby LAN, or Serviceguard moves packages and their associated IP address(es) to another node.
- If the node fails, Serviceguard quickly and automatically transfers an application to a functioning node.
- If the storage path fails, Serviceguard uses path failover software—e.g. the Fibre Channel HBA driver or, for SCSI, the Linux multi-device (MD) driver—to use the redundant path or, if that is not possible, moves packages and their associated IP address(es) to another node.
- If a disk drive in shared storage fails, RAID provides redundancy, preventing loss of data.
- If software fails, Serviceguard restarts an application on the same or another node with very little disruption.

Serviceguard also allows users to easily transfer control of applications to another server for system administration, maintenance, or version upgrades without bringing down the cluster.

Under normal conditions, a Serviceguard cluster monitors the health of cluster components while the packages run on individual nodes. When creating a package, a primary node is specified as well as one or more "adoptive" nodes where the package runs if the primary node fails. Several nodes can be defined as adoptive to maintain application performance.

If the cluster coordinator node does not receive heartbeat messages from all other cluster nodes within a prescribed time, a cluster begins to re-form. When the process is complete, information about the re-formed cluster is passed to the package coordinator (described in the "Package Manager" section of this white paper). Packages running on nodes that are no longer in the re-formed cluster are transferred to their adoptive nodes. In the case of a very brief loss of heartbeat, the cluster may re-form with the same nodes as before, so packages do not halt or switch.

Following are example scenarios where Serviceguard takes control of the situation:

- If a node leaves the cluster, the cluster re-forms with new nodes as members.
- If a node hangs or goes into an infinite loop, it is re-set.
- Network failure causes local network switching (if designed into the cluster).
- Service or application failure causes re-start (if configured to do so).
- Resource failure causes application failover to the node with available resources.
- Node failure causes application failover to a healthy node in the application's node list.

When a node or its network communications fails, Serviceguard transfers control of the package to the next available adoptive node.

Serviceguard provides protection beyond single-node failures within a cluster. Using a dynamic quorum capability, Serviceguard can survive multiple and cascading node failures and still provide protection for mission-critical applications. In fact, it is even possible for a single node to survive with appropriate connectivity to data and networks.

Managing Serviceguard clusters

Serviceguard Manager, which can run on HP-UX, Linux, and Windows® systems, allows users to view the status of clusters, nodes, and packages on HP-UX or Linux nodes through a single console. Serviceguard Manager can save snapshots of cluster status to document configurations and to provide a basis for comparison at a later date. Figure 3 shows a sample display of a cluster map and hierarchical view of the objects in one cluster.

Serviceguard Manager gets cluster data from an Object Manager daemon, or background process, which establishes connections with all managed cluster nodes. The Object Manager daemon may be located on a different system from Serviceguard Manager.

Basic control of the cluster is possible with Serviceguard Manager. This includes starting, stopping, and moving packages as well as starting and stopping nodes and the cluster itself.

With Serviceguard A.11.16 and later versions, Serviceguard Manager can create or modify the configuration of clusters and packages. Many of the configuration options are discovered by Serviceguard and displayed in user-selectable lists. Configuration files and control scripts can be automatically created with the click of a button.

Figure 3. Serviceguard Manager cluster configuration and status

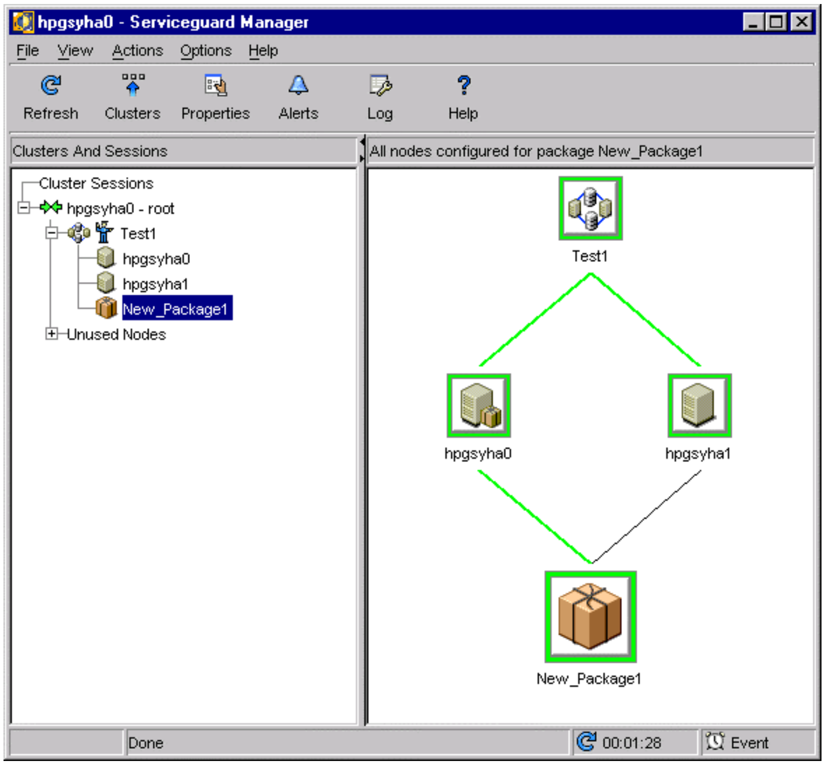

HP Serviceguard for Linux in action: application highlights

HP provides tools and products to simplify the deployment of applications from leading independent software vendors (ISVs). HP has built on its long-term partnerships with Oracle® and SAP to develop HP Serviceguard for Linux extensions that improve the availability and disaster tolerance of these environments as well as ease the integration of Serviceguard for Linux into other leading Linux application environments.

HP Cluster Extension XP

HP Serviceguard for Linux is tightly integrated with the HP Cluster Extension XP product, an online, mirrored data storage solution that uses real-time data replication of a production (primary) site to a secondary site. If the primary site becomes unavailable, HP Cluster Extension XP enables rapid storage data failover to the secondary site, allowing critical data to be available to business operations in minutes.

HP Serviceguard for Linux and HP Cluster Extension XP, together with HP StorageWorks XP disk array storage resources, have the capacity to protect mission-critical computing environments against downtime from fault, failure, or disaster. In particular, HP Cluster Extension XP with HP Serviceguard for Linux extends a single cluster over metropolitan-wide distances, supporting:

- Automatic failover for data centers located up to 100 kilometers apart
- High availability of computing resources
- Disaster-tolerant configurations that optimize server resources
- Significant reduction in the risk of operator error

HP Serviceguard for Linux Oracle Toolkit

This specialized toolkit enables Serviceguard clusters to support high availability for Oracle Database server applications. Working as a subsystem of Serviceguard, it provides functionality for starting, stopping, restarting, and monitoring Oracle database services in a cluster environment. Not only does the toolkit simplify integrating a complex database environment into the protection of a Serviceguard cluster; it provides full support from HP for the implementation when support is purchased.

HP Serviceguard Extension for SAP for Linux

Serviceguard Extension for SAP delivers instance health monitoring, automated failure detection, and robust application failover for mission-critical mySAP components based on SAP Netweaver technology. The product also provides failover capabilities for underlying Oracle or SAP DB database technologies.

Proven through thousands of successful implementations worldwide on HP-UX, HP Serviceguard Extension for SAP for Linux simplifies the integration of complex, high-availability SAP environments, offering:

- mySAP and Netweaver instance virtualization using simple Serviceguard commands
- SAP application availability during failures as well as hardware and software maintenance
- Full utilization of failover nodes during normal operation to maximize return on investment
- A fully tested, low-risk implementation that is approved by SAP
- The ability to integrate SAP applications with HP Cluster Extension XP
- A complete solution support portfolio and planning support from HP

Contributed toolkits for HP Serviceguard for Linux

In addition, HP has developed toolkits designed to speed deployment of common Linux applications. Simple script templates for HP ProLiant and Integrity servers and widely used applications like NFS, Samba, Apache, MySQL, Postgres, and SendMail are available at no extra charge. Users can easily customize configurations based on preferred cluster parameters and quickly install applications.

The list of extensions and toolkits available for HP Serviceguard for Linux is continually expanding. For the most updated list available, please see www.software.hp.com. HP also provides a free application validation service for ISVs through the Partner Technology Access Center, and additional applications can be supported through simple custom shell scripts developed in-house or through HP Consulting.

Complete solutions from HP

The world of Linux is evolving rapidly, with companies that offer Linux solutions entering and leaving the market relatively frequently. HP offers enterprise customers the peace of mind of working with an industry leader committed to a complete portfolio of solutions and full accountability.

HP Support and Services

HP understands that in the real world, Linux deployments generally mean application migration. In this environment, HP has Linux master builder status. As a company publicly committed to multiple operating systems, HP is uniquely qualified to manage successful application migrations to Linux and has the know-how to integrate Linux into heterogeneous OS environments. In support of these efforts:

- HP fields more than 5,000 trained Linux service professionals who provide migration, enterprise support, and deployment services.
- HP has created a comprehensive portfolio of services and support, including consulting, education, and training services, and critical support for Linux operating systems as well as HP hardware and software products. As the single point of contact, HP is solely accountable for all aspects of the Linux solutions we deploy, eliminating wasted time and the need to coordinate the work of multiple suppliers.
- HP delivers mission-critical competency 24x7 in 160 countries around the globe.

Industry-standard HP ProLiant and Integrity servers

HP ProLiant and Integrity servers deliver simplicity, agility, and value through innovation based on industry standards. This unique approach saves money for businesses with solutions that aim to reduce cost and complexity and improve integration across the infrastructure.

HP has the experience, vision, and commitment to help businesses leverage the value of industry-standard solutions. To provide the right technologies at the right price for Linux environments, HP markets a wide range of servers designed to meet user requirements consistently and without compromise.

Ideal for scale-out architectures and a variety of application types, HP ProLiant servers are the most trusted and best-selling industry-standard servers in the world[1]. Businesses worldwide have purchased more than eight million ProLiant servers in the past 10 years. And ProLiant is the development platform of choice for leading business environments, including Oracle, PeopleSoft, SAP, and Siebel.

[1] IDC Quarterly PC Tracker, August 2004

HP Integrity servers, available in entry-level, midrange, and high-end models, are architected to handle large databases, advanced application development, and critical business processes quickly and efficiently. HP Integrity servers combine industry-standard Intel® Itanium® 2 architecture with specialized HP chipsets to increase memory and I/O subsystem scalability. Today, HP produces record-breaking performance on Integrity servers running Linux, achieving industry-leading TPC-C performance with 1,184,893.38 tpmC @ $5.52/tpmC[2].

Through innovation based on standards and a broad portfolio of flexible choices, HP servers provide the essential building blocks of the Adaptive Enterprise. HP software solutions, global partnerships, and expert professional services and support offer a powerful approach to transforming change into opportunity.

HP storage

From the market leader in network storage, HP StorageWorks storage area networks (SANs) are characterized by their outstanding interoperability, scalability, and availability:

- **Interoperability**—interoperate with a broad spectrum of operating systems, HP clustering environments, and HP servers, enhanced by HP's after-sales support for significant business value
- **Scalability**—support SAN architectures that scale in connectivity, performance, and distance (LAN, MAN, WAN)
- **Availability**—meet the requirements of mission-critical applications and environments, including directors and core switches, dual fabrics, and multi-pathing software

The HP StorageWorks MSA family is ideal for small and medium business (SMB) customers and enterprise departments making the first step toward external storage consolidation and SAN. Typical customers tend to believe investment protection is critical and have moderate concerns about scalability. The MSA family is a perfect solution for customers with deployed Smart Array controllers and/or ProLiant servers, because it is built with HP Smart Array technology and features. With the MSA family and its unique DtS (DAS to SAN) migration technology, customers can easily move from internal ProLiant server storage to external storage and even SAN, leveraging common disk drives, management tools, data, and technology. The MSA family delivers the simplicity and cost saving of consolidated storage.

The HP StorageWorks EVA family is ideal for midrange customers who require both high availability and simplicity of management. When scalable growth and ongoing management costs are very important, EVA family products offer a simple-to-use and highly automated solution. With powerfully simple virtualization technology and tools such as Vsnaps, EVA products enable customers to consolidate storage and manage more storage per administrator to attain a desirable total cost of ownership (TCO) and the capacity to grow with the demands of their business.

The HP StorageWorks XP family is ideal for enterprise customers requiring the highest level of scalability, availability, disaster recovery, and business continuity, as well as very high service levels. In addition, the XP series provides an effective platform for data-center consolidation.

[2] Powered by 16 HP Integrity rx6570 servers, each with 4 Intel Itanium 2 processors

Conclusion

HP believes that the Linux operating system broadens the opportunities for enterprise computing. And enterprises that have chosen—or are considering choosing—the Linux operating system are clearly concerned about infrastructure stability and application availability in business-critical computing environments.

In the high-stakes world of enterprise IT, the choice of any solution must be influenced by the availability of a partner who can provide products and support that extend its capabilities. HP is well known for its support of open source computing solutions and industry standards, for its support of Linux, and for its ongoing commitment to developing robust solutions based on the Linux platform.

HP has sold more than 100,000 Serviceguard server licenses for HP-UX environments and has leveraged its UNIX® expertise for the benefit of customers running Linux in single- and multi-OS environments. HP will continue to develop the solutions, services, and support that customers expect for this powerful, flexible platform.

For more information

For more information about HP Serviceguard for Linux, please visit www.hp.com/go/sglx. For more information about Linux and HP ProLiant and Integrity servers, please visit www.hp.com/go/linux. For additional Serviceguard product family information, please visit www.hp.com/go/serviceguard.

The HP documents listed below may also be of interest:

- "Arbitration for Data Integrity in MC/Serviceguard Clusters" (July 2002)
  www.docs.hp.com/linux/onlinedocs/B3936-90078/B3936-90078.html
- "Highly Available Networks" (1996)
  www.docs.hp.com/hpux/onlinedocs/ha/wpnetwork.pdf
- "Providing Open Architecture High Availability Solutions" (February 2001)
  www.docs.hp.com/hpux/onlinedocs/ha/HA_Solutions.pdf
- "Understanding High Availability" (1996)
  www.docs.hp.com/hpux/onlinedocs/ha/wpunderha.pdf

5982-8326EN, 09/2004