



hp TruCluster
technology

october 2002



gold series

The Gold Standard in UNIX® Clusters: TruCluster Technology

Deploying a highly available cluster in a production environment

The Gold Standard in UNIX Clusters: TruCluster Technology

TruCluster Solutions Deliver High Availability and Simplified Cluster Management

The most successful businesses today manage information effectively, ensuring ready access for those who need it. Maintaining and managing a high availability (HA) environment for customers, vendors and employees is no longer optional—it's critical.

Clustering servers is an excellent way to achieve high availability. However, with most traditional cluster configurations, maintenance is complex and time-consuming. HP's TruCluster technology, however, offers an easily managed 24x7 environment with configurations that won't stop working until the last wire fails. TruCluster technology uniquely offers the most comprehensive cluster management environment with scalable performance.

With its unique Single System Image (SSI) and "true" Cluster File System (CFS) technology, TruCluster Server is the leader in providing cluster capability, delivering high value to customers who require continuous uptime, demand unlimited scalability, and are looking for dramatic reductions in the system management costs of their infrastructure.

TruCluster Server with SSI provides reliable redundancy, failover and backup—all key components of a high availability environment. With any IT environment, the administrators are responsible for organizing and operating their systems. Effective systems management requires attention to not only computing resources, but also to procedures such as managing changes and archiving data.

Meeting a "no single point of failure" (NSPOF) requirement can range across a continuum. At one extreme is the "poor man's cluster", a spare server available for use only if needed. At the other extreme is the fully redundant approach—two of everything, including backup sites to avoid the impact of catastrophic failure. Most organizations' high availability needs fall somewhere in between these extremes.

A properly configured and managed TruCluster Server solution offers a cost-effective, highly available environment, with recommended configurations based on redundant topology.

The bottom line--- TruCluster technology can deliver simplified cluster management with 24x7 availability, highly scalable performance, and NSPOF (no single point of failure) so that information can be reliably accessed to ensure continued business success, even in the event of a system failure

Deploying a Highly Available Cluster in a Production Environment

The key to TruCluster technology's simplified manageability is how fully it automates administrative tasks. This comprehensive automation results in efficient and productive system management.

The automation of device discovery, multi-path management, network and storage reconfiguration, and direct i/o, as well as radically simplified failover scripts (available free of charge), combine to provide highly automated and efficient TruCluster solutions.

Additionally, all systems in the cluster share common root, /usr, and /var file systems, as well as application data. All nodes can access disks and files quickly and concurrently, regardless of whether the file is stored on a locally attached disk or on a shared storage bus.

Although Sun with Veritas offers partial cluster file system (CFS) capabilities, TruCluster Server is unique in both its seamless integration with the kernel and its support for common root, /usr, and /var file systems. With CFS, TruCluster Server provides an exceedingly stable, high performance, high availability environment that ensures data integrity.

To fully realize this environment and all of its advantages, systems administrators should follow several basic recommendations. Critical guidelines to achieving a highly available, highly scalable, easy to manage cluster include:

Redundant topology

Having a back up server is critical to high availability, because any single server—no matter how high-end -- represents a single point of failure. The more separate servers a user puts into operation, the greater the degree of redundancy is achieved.

TruCluster technology makes it simple to add additional servers to a cluster environment. Server nodes can be added to a configuration in as little as 15 minutes, after which they can be managed as a single system along with the other clustered nodes. TruCluster technology, often referred to as the “gold standard”, creates simplicity from complexity.

Storage

There are multiple technologies that allow for redundant storage, including hardware such as redundant controllers and software such as Logical Storage Manager (LSM) for managing and mirroring the redundant data.

The operating system and cluster software—through the use of multipath and TruCluster Device Requester Dispatcher (DRD) features-- enable multiple paths to the data as well as distribute storage across the cluster to provide optimum performance as well as ensure NSPOF. Access to the storage will continue uninterrupted until the last wire connecting the storage to another cluster member breaks.

Connection to network

Use reliable and redundant network hardware such as fibre connections, and current generation switches and adapters to create a topology that endures NSPOF.

Features such as link aggregation groups (LAG), redundant array of network adapters (netRAIN), and cluster alias (which manages network connections, rerouting traffic transparently if a failure occurs) provide the technology to ensure NSPOF to the network.

Good System Management Practices

Regardless of whether a system administrator is managing a single non-clustered node, a “traditional” failover cluster, or a single system image cluster such as TruCluster Server, good system management, today and in the future, will be based on proper training and procedures tempered with common sense.

TruCluster technology simplifies system management by eliminating time-consuming tasks, allowing multiple systems to be managed as one system. The storage stack, single security protocol and authentication domains are a just a few examples of how the Tru64 UNIX[®] TruCluster environment implements simplified management.

System management practices should include the following:

- Give administration privileges to only those who need it to protect against human error. Use appropriate security and auditing.
- Periodically backup system disks and member boot disks.
- Develop an availability plan with operational guidelines. Follow the plan, and periodically test and modify if needed.
- Extend procedures to include Disaster Recovery strategy---from regular backup operations to fully redundant disaster tolerant solutions.

On a more technical level:

- On large deployments, proper usage of the Division of Privileges (dop) tool-- or similar tools such as sudo-- ensures that system managers with responsibilities for certain components of the overall environment are given access only to the components they manage. This practice will limit the usage of the super-user and thus limit the possibility of accidental deletions of critical files or file-systems.
- Smart management practices, even in non-clustered production environments, include mirroring the root /usr and /var file systems, plus any file systems with important data.
- Periodically restore system backups to alternate disks to ensure the integrity of the system backup

Additionally, if you are using hardware RAID, you should consider:

- Using hardware RAID to mirror the clusterwide root (/), /usr, and /var file systems, the member boot disks, quorum disk (if present), and data disks.
- Using at least two shared storage buses to ensure continuous access to the dual-redundant RAID array controllers (HSZ70, HSZ80, HSG60, and HSG80).

The Gold Standard in UNIX Clusters: TruCluster Technology

Computer Center
Considerations for
NSPOF

Good system management and practices/policies extend beyond the cluster configuration. Despite redundant fabric, switches, hardware and software, a single environmental event can still result in a SPOF. The physical configuration of the data center is critical to achieving high availability.

For example, let's say a customer places dual servers in ONE cabinet in the computer room. This cabinet is located under the room's cooling unit, and water leaks into the cabinet causing a failure. The redundant configuration shorts out and stops short of achieving NSPOF. The solution? Use more than one cabinet to avoid a SPOF.

Another example is as simple as the power circuit. Are all servers connected to the same circuit breaker? If the circuit breaker fails, will the cluster fail?

To achieve even higher levels of availability, customers need to look beyond cluster configurations that are housed in a single computer room to protect against a catastrophe in the building. A high availability NSPOF configuration that includes disaster tolerance or disaster recovery requires remote backup and failover sites.

The Gold
Standard in
UNIX Clustering

Proponents of traditional failover models frequently attempt to minimize the value of TruCluster technology to customers that require a highly available environment. The reality is that a properly configured TruCluster environment is highly available, extremely reliable, flexibly scalable and much easier to manage than traditional clusters.

TruCluster technology preserves transparency, high performance, and predictable behavior across the cluster, while protecting data integrity. In the event that a system fails, data can be recovered quickly and efficiently, so that users are not even aware that a failure ever occurred.

TruCluster technology is planned for incorporation into HP-UX 11i version 3 (expected availability 2004) as stated in our public roadmap.

With TruCluster Server you can achieve the Gold Standard in UNIX clustering. High availability no longer needs to be weighed against the burden of typical cluster approaches. HP is truly delivering clustering for all, enabling enterprises of all sizes to protect their critical information assets and grow seamlessly while maintaining the ease of management of a single system.

The Gold Standard in UNIX Clusters: TruCluster Technology prepared by UNIX Software Division
First Edition (October 2002)

©2002 Hewlett-Packard Company. Printed in the U.S.A.

UNIX is a registered trademark of The Open Group.

All other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for HP products are set forth in the express limited warranty statements accompanying such products. Nothing herein shall be construed as constituting an additional warranty.



