

Overview

HP TruCluster Server Version 5.1B-5 for HP Tru64 UNIX Version 5.1B-5 provides highly available and scalable solutions for mission-critical computing environments. TruCluster Server delivers powerful but easy-to-use UNIX clustering capabilities, allowing AlphaServer systems and storage devices to operate as a single virtual system.

By combining the advantages of symmetric multiprocessing (SMP), distributed computing, and fault resilience, a cluster running TruCluster Server offers high availability while providing scalability beyond the limits of a single system. On a single-system server, a hardware or software failure can severely disrupt a client's access to critical services. In a TruCluster Server cluster, a hardware or software failure on one member system results in the other members providing these services to clients.

TruCluster Server reduces the effort and complexity of cluster administration by extending single-system management capabilities to clusters. It provides a cluster-wide namespace for files and directories, including a single root file system that all cluster members share. A common cluster address (cluster alias) for the Internet protocol suite (TCP/IP) makes the cluster appear as a single system to its network clients while load balancing client connections across member systems.

A single system image (SSI) allows a cluster to be managed more easily than distributed systems. TruCluster Server cluster members share a single root file system and common system configuration files. Therefore, most management tasks need to be done only once for the entire cluster rather than repeatedly for each cluster member. The cluster can be managed either locally from any of its members or remotely using Tru64 UNIX Web-based management tools. Tru64 UNIX and TruCluster Server software, and applications, are installed only once. Most network applications, such as the Secure Web server, need to be configured only once in the cluster and can be managed more easily in a cluster than on distributed systems.

A choice of graphical, Web-based, or command-line user interfaces makes management tasks easier for the administrator, flexible for those with large configurations, and streamlined for expert users.

TruCluster Server facilitates deployment of services that remain highly available even though they have no embedded knowledge they are running in a cluster. Applications can access their disk data from any cluster member. TruCluster Server also provides the support for components of distributed applications to run in parallel, providing high availability while taking advantage of cluster-specific synchronization mechanisms and performance optimizations.

TruCluster Server allows the processing components of an application to concurrently access raw devices or files, regardless of where the storage is located in the cluster. Member-private storage and clusterwide shared storage are equally accessible to all cluster members. Using either standard UNIX file locks or the distributed lock manager (DLM), an application can synchronize clusterwide access to shared resources, maintaining data integrity.

TruCluster Server is an efficient and reliable platform for providing services to networked clients. To a client, the cluster appears to be a powerful single-server system; a client is impacted minimally, if at all, by hardware and software failures in the cluster. TruCluster Server simplifies the mechanisms of making applications highly available. A cluster application availability (CAA) facility records the dependencies of, and transparently monitors the state of, registered applications. If a hardware or software failure prevents a system from running a service, the failover mechanism automatically relocates the service to a viable system in the cluster, which maintains the availability of applications and data. Administrators can manually relocate applications for load balancing or hardware maintenance.

TCP-based and UDP-based applications can also take advantage of the cluster alias subsystem. These applications, depending on their specific characteristics, can run on a single cluster member or simultaneously on multiple members. The cluster alias subsystem routes client requests to any member participating in that cluster alias. During normal operations, client connections are dynamically distributed among multiple service instances according to administrator-provided metrics.

TruCluster Server supports a variety of hardware configurations that are cost-effective and meet performance needs and availability requirements. Hardware configurations can include different types of systems and storage units, and can be set up to allow easy maintenance. In addition, administrators can set up hardware configurations that allow the addition of a system or storage unit

Overview

without shutting down the cluster.

For the fastest communication with the lowest latency, use the PCI-based Memory Channel cluster interconnect for communication between cluster members. TruCluster Server also supports the use of 100 Mbps Ethernet or 1000 Mbps Ethernet hardware as a private LAN cluster interconnect. The LAN interconnect is suitable for clusters with low-demand workloads generated by a cluster running failover style, with highly available applications in which there is limited application data being shared between the nodes over the cluster interconnect. Refer to the *Cluster Technical Overview* manual for a discussion of the merits of each cluster interconnect.

New Features

ICSNET Performance Enhancement

This feature optimizes the performance of ics0 pseudo interface in a LAN cluster. The ICSNET pseudo interface is used by cluster aware applications to communicate with peer applications on other member nodes within the cluster. This enhancement makes communication with other members faster using ICSNET driver.

Please refer to the Tru64 UNIX Version 5.1B-5 Release Notes for more information.

Features - TruCluster Server V5.1B-5

Cluster members in a given cluster must all use Memory Channel or must all use LAN. These interconnects cannot be mixed for cluster communication in the same cluster. Using multiple shared buses and redundant Memory Channel or LAN interconnect hardware promotes no-single-point-of-failure (NSPOF) characteristics for mission-critical applications.

A TruCluster Server cluster acts as a single virtual system, even though it is made up of multiple systems. Cluster members can share resources, data, and clusterwide file systems under a single security and management domain, yet they can boot or shut down independently without disrupting the cluster.

Cluster File System

The Cluster File System (CFS) makes all files, including the root (/), /usr, and /var file systems, visible to and accessible by all cluster members. It does not matter whether a file is stored on a device connected to all cluster members or on one that is private to a single member. Each file system is served by a single-cluster member; other members access that file system as CFS clients with significant optimizations for shared access. CFS preserves full X/Open and POSIX semantics for file system access and maintains cache coherency across cluster members. For instance, an application can use standard UNIX file locks to synchronize access to shared files. The member that will serve a given file system can be specified at file system mount time.

For higher performance, applications can use direct I/O through the file system to bypass the buffer cache. CFS also provides a load balancing daemon to monitor and analyze file system usage. The daemon can be configured to automatically relocate file systems based on CFS memory usage, or during changes in the cluster, such as when members join or leave, or when storage connectivity changes. In addition, the member that will serve a given file system can be specified when that file system is mounted.

CFS supports the Advanced File System (AdvFS) for both read and write access and supports AdvFS with BSD-type user and group quotas. NFS client and NFS server are supported for both read and write access, and NFS services are accessible to clients through cluster aliases in addition to the default cluster alias. The UNIX File System (UFS) is supported for read and write access from the CFS server or for read only access on client members. The Memory File System (MFS) is supported for both read only or read and write access by the member on which the file system is mounted. Remote access and failover are not supported. The CD-ROM File System (CDFS) and Digital Video Disc File System (DVDFS) are supported for read access only.

Enhanced Distance Cluster

TruCluster supports an extended distance capability up to 100 KM. These configurations are bounded by parameters that specify the number of separate sites and the distance between sites, the number of members, the means of maintaining cluster quorum, the type of storage used, and the means of replicating data between geographically separated storage. Refer to the Tru64 UNIX Version 5.1B-4 Release Notes for the detailed specification of the cluster configurations supported.

Device Request Dispatcher

The device request dispatcher (DRD) supports cluster-wide access to character and block disk devices, and to tape and tape changer devices. All local and remote cluster disk and tape I/O passes through the device request dispatcher. A member does not need a direct connection to a disk, or tape, or tape changer device to access data on that device. This permits great flexibility in selecting a hardware configuration that is both economical and useful.

Features - TruCluster Server V5.1B-5

Cluster Alias

A cluster alias is an IP address that makes the cluster look like a single system to clients and other hosts on the network. Cluster aliases free clients from having to connect to specific members for services. If the member providing the service goes down, a client reconnects to another member elected by the cluster alias to provide the service. With applications that run concurrently on multiple members, scaling is achieved by permitting multiple clients to connect to instances of the service on multiple cluster members, each using a cluster alias to address the service.

The cluster alias subsystem provides an optional virtual MAC (vMAC) address that can be associated with each cluster alias IP address. When configured, the same MAC address is used in all Address Resolution Protocol (ARP) responses for the cluster alias address, independent of which cluster node is responding to cluster alias ARP requests. This permits faster failover when a new node assumes responsibility for responding to cluster alias ARP requests.

Flexible Network Configuration

TruCluster Server offers flexible network configuration options. Cluster members do not need to have identical routing configurations. An administrator can enable IP forwarding and configure a cluster member as a full-fledged router. Administrators can use routing daemons such as `gated` or `routed`, or they can configure a cluster member to use only static routing. When static routing is used, administrators can configure load balancing between multiple network interface cards (NICs) on the same member. Whether `gated`, `routed`, or static routing is used, in the event of a NIC failure, the cluster alias reroutes network traffic to another member of the cluster. As long as the cluster interconnect is working, cluster alias traffic can get in or out of the cluster.

Cluster Application Availability Facility

The cluster application availability (CAA) facility delivers the ability to deploy highly available single instance applications in a cluster by providing resource monitoring and application relocation, failover, and restart capabilities. CAA is used to define which members can run a service, the criteria under which to relocate a service, and the location of an application-specific action script. Monitored resources include network adapters, tape devices, media changers, and applications. CAA allows services to manage and monitor resources by using entry points within their action scripts. Applications do not need to be modified in any way to utilize CAA.

Administrators can request that CAA reevaluate the placement within the cluster of registered applications either at a regularly scheduled time, or any time at which they desire to manually balance applications by using the `caa_balance` command. Balancing decisions are based on the standard CAA placement mechanisms. Similarly, administrators can request that CAA schedule an automatic failback of a resource for a specific time. This allows an administrator to benefit from CAA automatically moving a resource to the most-favored cluster member without the worry of the relocation occurring at a critical time. The `caa_report` utility can provide a report of availability statistics for application resources. Administrators can redirect the output of CAA resource action scripts to be visible during execution. Lastly, user-defined attributes can be added to a resource profile and they will be available to the action script upon its execution.

Features - TruCluster Server V5.1B-5

Rolling Upgrade

TruCluster Server allows rolling upgrade from the previous version of the base operating system and the TruCluster software to the next subsequent release of the base operating system and TruCluster software. It also allows the rolling of patches into the cluster. Updating the operating system and cluster software does not require a shutdown of the entire cluster. A utility is provided to roll the cluster in a controlled and orderly fashion. The upgrade procedure allows the monitoring of the status of the upgrade while it is in progress. Clients accessing services are not aware that a rolling upgrade is in progress.

To speed the process of upgrading the cluster, the administrator can use the parallel rolling upgrade procedure that upgrades more than one cluster member at a time in qualifying configurations.

Administrators looking for a quicker alternative to a rolling upgrade when installing patches have the option of a patch procedure that favors upgrade speed over cluster high availability. After the first member receives the patch, all remaining members of the cluster receive the patch at the same time followed by rebooting the entire cluster as a single operation.

See the *Cluster Installation* manual for recommended and supported paths to upgrade or roll to the latest version of TruCluster Server.

Cluster Cloning

Tru64 UNIX offers an additional way to install the operating system and TruCluster software upgrades in a process referred to as "cluster cloning". First, you create an exact duplicate of the system on an alternate set of disk drives and then run the dupclone script to apply the update to the alternate disk set. Once completed, the cluster can immediately be rebooted from the alternate set of disks. Refer to the Tru64 UNIX Version 5.1B-4 Release Notes for information on the dupclone command.

Cluster Management

The SysMan system management utilities provide a graphical view of the cluster configuration, and can be used to determine the current state of availability and connectivity in the cluster. The administrator can invoke management tools from SysMan, allowing the cluster to be managed locally or remotely.

Clusterwide signaling allows applications to send UNIX signals to processes operating on other members.

Performance Management

The performance management capability of Tru64 UNIX has been modified from one large performance management tool (*pmgr*) to several smaller and more versatile tools. The performance management tool suite consists of collect, collgui, and two Simple Network Management Protocol (SNMP) agents (*pmgrd* and *clu_mibs*).

The collect tool gathers operating system and process data under Tru64 UNIX Versions 4.x and 5.x. Any subset of the 'subsystems' (Process, Memory, Disk, LSM Volumes, Network, CPU, Filesystems) and Header can be defined for which data is to be collected. Collect is designed for high reliability and low system-resource overhead. Accompanying collect are two highly integrated tools: collgui (a graphical front-end) and cfilt, (which allows completely arbitrary extraction of data from the output of collect to standard output). Collgui is a laborsaving tool that allows a user to quickly analyze collect data. The Performance Manager metrics server (*pmgrd*) is a UNIX daemon process that provides general UNIX performance metrics on request. The *pmgrd* metrics server supports the extensible SNMP agent mechanism (eSNMP).

Features - TruCluster Server V5.1B-5

Cluster MIB	TruCluster Server supports the HP Common Cluster MIB. HP Insight Manager uses this Cluster MIB to discover cluster member relationships, and to provide a coherent view of clustered systems across supported platforms.
--------------------	--

Highly Available NFS Server	<p>When configured as an NFS server, a TruCluster Server cluster can provide highly available access to the file systems it exports. No special cluster management operations are required to configure the cluster as a highly available NFS server. In the event of a system failure, another cluster member will become the NFS server for the file system, transparent to external NFS clients. NFS file locking is supported, as are both NFS V2 and V3 with UDP and TCP.</p> <p>TruCluster Server allows NFS file systems to be served from the cluster through both the default cluster alias and alternate aliases. Alternate cluster aliases can be defined to limit NFS server activity to those members that are actually connected to the storage that contains the exported file systems. NFS clients can use this alternate alias when they mount the file systems served by the cluster.</p> <p>This feature is not supported in an Enhanced Distance Cluster.</p>
------------------------------------	---

Fast File System Recovery	<p>The Advanced File System (AdvFS) log-based file system provides higher availability and greater flexibility than traditional UNIX file systems. AdvFS journaling protects file system integrity. TruCluster Server supports AdvFS for both read and write access.</p> <p>An optional, separately licensed product, the Advanced File System Utilities, performs online file system management functions. See the OPTIONAL SOFTWARE section of this document for more information on the AdvFS utilities.</p>
----------------------------------	---

Increased Data Integrity	<p>Tru64 UNIX Logical Storage Manager (LSM) is a cluster-integrated, host-based solution to data storage management. In a TruCluster Server cluster, LSM operations continue despite the loss of cluster members, as long as the cluster itself continues operation and a physical path to the storage is available. LSM disk groups can be used simultaneously by all cluster members and the LSM configuration can be managed from any cluster member.</p> <p>Basic LSM functionality, including disk spanning and concatenation, is provided with the Tru64 UNIX operating system. Extended functions, such as striping (RAID 0), mirroring (RAID 1), and online management, are available with a separate license. Mirroring of LSM is RAID Advisory Board (RAB) certified for RAID Levels 0 and 1.</p> <p>LSM is supported for use in a TruCluster Server cluster and will support any volume in a cluster, including swap and cluster root and excluding the quorum disk and member boot disks. Hardware mirroring is supported for all volumes in a cluster without exception.</p> <p>LSM RAID 5 volumes are not supported in clusters. See the OPTIONAL SOFTWARE section of this document for more information on LSM.</p>
---------------------------------	--

Global Error Logger & Event Manager	TruCluster Server can log messages about events that occur in the TruCluster environment to one or more systems. Cluster administrators can also receive notification through electronic mail when critical problems occur.
--	---

Features - TruCluster Server V5.1B-5

Cluster Client Network Failover	TruCluster Server supports highly available client network interfaces via the Tru64 UNIX redundant array of independent network adapters (NetRAIN) feature.
--	---

Link Aggregation on the Cluster Interconnect	Link aggregation is now supported on the cluster LAN interconnects. This increases the bandwidth providing increased interconnect throughput. The multiple active interfaces reduces message queue length and delivers lower latency. This results in improved cluster performance.
---	---

Cluster Interconnect Failover	TruCluster Server allows the elimination of the cluster interconnect as a single point of failure by supporting redundant cluster interconnect hardware. You can configure dual-rail Memory Channel, allowing the cluster to survive the failure of a single rail. For LAN interconnect, two or more network adapters on each member are configured as a NetRAIN virtual interface. When properly configured across two or more switches, the cluster will survive any LAN component failure. This not only guards against rare network hardware failures, but also facilitates the upgrade and maintenance of the network without disrupting the cluster.
--------------------------------------	--

Support for Parallelized Database Applications	<p>TruCluster Server provides the software infrastructure to support parallelized database applications, such as Oracle 9i Real Application Clusters (RAC) and Informix Extended Parallel Server (XPS) to achieve high performance and high availability. 9i RAC and XPS are offered and supported separately by Oracle Corporation and Informix Software, Inc., respectively.</p> <p>Parallelized database applications are not recommended to be used across an Enhanced Distance Cluster link.</p>
---	---

Distributed Lock Manager	The distributed lock manager (DLM) synchronizes access to resources that are shared among cooperating processes throughout the cluster. DLM provides a software library with an expansive set of lock modes that applications use to implement complex resource-sharing policies. DLM provides services to notify a process owning a resource that it is blocking another process requesting the resource. An application can also use DLM routines to efficiently coordinate the application's activities within the cluster.
---------------------------------	--

Support for Memory Channel API	<p>TruCluster Server provides a special application programming interface (API) library for high-performance data delivery over Memory Channel by giving access to Memory Channel data transfer and locking functions. This Memory Channel API library enables highly optimized applications that require high-performance data delivery over the Memory Channel interconnect. This library is supported solely for use with Memory Channel.</p> <p>High performance within the cluster is achieved by providing user applications with direct access to the capabilities of the Memory Channel. For example, a single store instruction on the sending host is sufficient for the data to become available for reading in the memory of another host.</p> <p>The Memory Channel API library allows a programmer to create and control access to regions of the clusterwide address space by specifying UNIX style protections. Access to shared data can be synchronized using Memory Channel spin locks for clusterwide locking.</p> <p>The Memory Channel API library facilitates highly optimized implementations of Parallel Virtual Machine (PVM), Message Passing Interface (MPI), and High Performance Fortran (HPF), providing seamless scalability from SMP systems to clusters of SMP machines. This provides the programmer with</p>
---------------------------------------	--

Features - TruCluster Server V5.1B-5

comprehensive access to the current and emerging de facto standard software development tools for parallel applications while supporting portability of existing applications without source code changes.

NOTE: To users of the Memory Channel API V1.6 product on Tru64 UNIX Version 4.0*:

On Tru64 UNIX V5.*, the TruCluster Memory Channel Software V1.6 product is bundled as a feature of TruCluster Server V5.*. To run the Memory Channel API library on Tru64 UNIX Version 5.*, you must install a TruCluster Server license to configure a valid TruCluster Server cluster.

NOTE: To users of the Memory Channel application programming interface with Memory Channel virtual hub (vhub) configuration:

The Memory Channel API is not supported for data transfers larger than 8K bytes when loopback mode is enabled in two member clusters configured with MC virtual hub. For more information on loopback mode go to http://h30097.www3.hp.com/docs/pub_page/cluster_list.html and refer to the 7.Aug.2003 Technical Update titled "TruCluster Server Version 5.1A Technical Update" and refer to the article titled "MC API Applications May Not Use Transfers Larger Than 8 KB with Loopback Mode Enabled on Clusters Utilizing Virtual Hubs."

Connection Manager

The connection manager is a distributed kernel component that ensures that cluster members communicate with each other and enforces the rules of cluster membership. The connection manager forms a cluster, and adds and removes cluster members. It tracks whether members in a cluster are active and maintains a cluster membership list that is consistent on all cluster members.

Cluster Storage I/O Failover

TruCluster Server provides two levels of protection in the event of storage interconnect failure. When configured with redundant storage adapters, the storage interconnect will be highly available. Should one interconnect fail, traffic will transparently fail over to the surviving adapter. When a member system is connected to shared storage with a single storage interconnect and it fails, transactions are transparently performed via the cluster interconnect to another cluster member with a working storage interconnect.

Support for Fibre Channel Solutions

TruCluster Server supports the use of switched Fibre Channel storage and Fibre Channel arbitrated loop. Compared to parallel SCSI storage, Fibre Channel provides superior performance, greater scalability, higher reliability and availability, and better serviceability. Compared to parallel SCSI storage, Fibre Channel is easier to configure and its long distance permits greater flexibility in configurations. Fibre Channel can be used for clusterwide shared storage, cluster file systems, swap partitions, and boot disks.

Compared with a switched Fibre Channel topology, arbitrated loop offers a lower cost solution by trading off bandwidth, and therefore some performance. Arbitrated loop is supported for two-member configurations only.

For more information on supported TruCluster Server configurations and specific cabling restrictions using Fibre Channel, see the Cluster Hardware Configuration manual at the following URL: http://h30097.www3.hp.com/docs/pub_page/cluster51B_list.html

Enhanced Security with Distributed Authentication

TruCluster Server supports the Enhanced Security option on all cluster members. This includes support for features for enhanced login checks and password management. Audit and access control list (ACL) support can also be enabled independently of the Enhanced Security option on cluster members.

Configuration

Software Requirements TruCluster Server Version 5.1B-5 requires the Tru64 UNIX Version 5.1B-5 operating system. The Tru64 UNIX operating system is a separately licensed product. See the Tru64 UNIX operating system *QuickSpecs* for more information.

TruCluster Server requires that additional software subsets be installed. See the *TruCluster Server Cluster Installation* manual for more information.

Software Configuration Requirements When configuring TruCluster Server software, an additional 64 MB of memory is required on each member system.

Each system requires a disk for a member boot disk and the cluster requires a minimum of one disk for clusterwide root, /usr, and /var file systems. A quorum disk is optional. See the *TruCluster Server Cluster Hardware Configuration* manual for details. Free disk space required for use (permanent): 62 MB to load TruCluster Server software onto a Tru64 UNIX system disk.

These requirements refer to the disk space required on the system disk. The sizes are approximate; actual sizes may vary depending on the system environment, configuration, and software options.

Growth Considerations The minimum hardware and software requirements for any future version of this product may be different from the requirements for the current version.

A rolling upgrade to the next version of the cluster software requires the following:

- At least 50 percent free space in root (/), cluster_root#root
- At least 50 percent free space in /usr, cluster_usr#usr
- At least 50 percent free space in /var, cluster_var#var, plus an additional 425 MB to hold the subsets for the new version of the Tru64 UNIX operating system
- At least 50 percent free space in /usr/i18n, cluster_i18n#i18n when used.

Ordering Information

TruCluster Server V5.1B-5 licenses include:

Systems	TruCluster Plus packages*	TruCluster Server License	TruCluster Server Migration License**
AlphaServer 800, 1000A, 1200, DS10, DS15, DS15A, DS10L, DS20, DS20E, DS25, ES47 server tower	QP-6R9AC-AA	QL-6BRAC-AA	QL-6J1AC-AA
AlphaServer ES40, ES45, ES47, TS20	QP-6R9AE-AA	QL-6BRAE-AA	QL-6J1AE-AA
AlphaServer 2000, 2100, 2100A, 4000, 4100, GS60E, GS80, GS80, ES80	QP-6R9AG-AA	QL-6BRAG-AA	QL-6J1AG-AA
AlphaServer 8200, 8400, GS60, GS140, GS160, GS320, GS1280	QP-6R9AQ-AA	QL-6BRAQ-AA	QL-6J1AQ-AA

Configuration

Software Documentation: QA-6BRAA-GZ

* TruCluster Plus Software packages include licenses for TruCluster Server, Logical Storage Manager, and AdvFS Utilities.

** If you currently have TruCluster Available Server or TruCluster Production Server and want to convert to TruCluster Server, use the QL-6J1A*-AA migration license.

Software Licensing

The HP TruCluster Server license provides the right to use the software as described in this QuickSpecs, and is furnished under the licensing of HP Computer Corporation's Standard Terms and Conditions. The version of HP TruCluster Server described in this QuickSpecs qualifies as a minor version release. Licenses for prior versions must be updated to this version either through the purchase of a Service Agreement that includes the rights-to-use new versions, or through the purchase of Update Licenses. Each system in the TruCluster Server environment requires separate Tru64 UNIX and TruCluster Server licenses. Hard partitions of ES80, GS80, GS160, GS320, and GS1280 AlphaServers can be clustered together either across separate systems or within systems, and only one TruCluster license is required per system. For more information about the HP licensing terms and policies, contact your local HP representative or reseller.

This product supports the Tru64 UNIX License Management Facility (LMF). License units for the TruCluster Server product are allocated on an unlimited-system-use basis.

For more information on the License Management Facility, see the Tru64 UNIX Operating System QuickSpecs or the Tru64 UNIX operating system documentation.

Distribution Media

TruCluster Server is a separately licensed product and is distributed on the Tru64 UNIX Associated Products Volume 2 CD-ROM. The TruCluster Server documentation is available on-line at http://h30097.www3.hp.com/docs/pub_page/cluster51B_list.html

Software Product Services A variety of service options are available from HP. For more information, contact your local HP service representative.

Optional Software

HP Advanced Server for UNIX (ASU)

HP Advanced Server for UNIX (ASU) provides Windows networking services, such as file sharing, print sharing, and security for Tru64 UNIX. In addition to basic file and print services, ASU provides full Windows domain controller support, support for enterprise-wide trust relationships, and support for Windows security - including file permissions and Windows local and global groups.

Additionally, you can manage users, file shares, and printers using native Windows administrative tools. When combined with TruCluster Server software, ASU provides highly available and highly scalable file shares, print shares, and even Primary Domain Controller resources to Windows clients. For more information on ASU, visit the HP Advanced Server for UNIX Web site at:

http://h30097.www3.hp.com/products/adv_server/.

Continuous Access (CA) by HP

Continuous Access is controller-based data replication software for disaster tolerance solutions. CA is supported in a TruCluster with the XP and EVA product families.

The SANworks Data Replication Manager (DRM) is a legacy controller-based data replication software product for disaster tolerance and data movement solutions. DRM works with StorageWorks Fibre Channel MA8000 /EMA12000 and the RAID Array 8000 (RA8000) and Enterprise Storage Array 12000 (ESA12000) solutions from HP.

Multiple clusters or standalone systems can be connected to replicate application data. Use of disaster tolerant data replication within a single cluster is supported only through the "Campus-Wide Disaster Tolerant Cluster" product offering. For more information about the Campus-Wide Disaster Tolerant Product, see http://h30097.www3.hp.com/cluster/tru64_cluster_clusters.html.

Advanced File System (AdvFS) Utilities

The Advanced File System (AdvFS) log-based file system provides flexibility, compatibility, high availability, and high performance for files and filesets, up to 16 terabytes (TB). Administrators can add, remove, reconfigure, tune, and defragment files – and back up storage – without unmounting the file system or halting the operating system. By supporting multivolume file systems, AdvFS enables file-level striping to improve file transfer rates, and integrates with the functionality provided by the Logical Storage Manager (LSM).

A graphical user interface simplifies management tasks and utilities to dynamically resize file systems, load balance, undelete files, and clone files for hot backup.

The AdvFS Utilities is a separately licensed software product for Tru64 UNIX. See the AdvFS Utilities *QuickSpecs* for more information.

Logical Storage Manager (LSM)

The Tru64 UNIX Logical Storage Manager (LSM) is an integrated, host-based solution to data storage management. Concatenation, striping, mirroring, hot-sparing, and a graphical user interface allow data storage management functions to be done online, without disrupting users or applications. LSM manages storage as a single entity in both cluster and single node environments. LSM is a separately licensed software product for Tru64 UNIX. For more information, see the Logical Storage Manager *QuickSpecs*.

Optional Software

StorageWorks Software The StorageWorks Software package includes the licenses for Tru64 UNIX Logical Storage Manager and the Advanced File System Utilities. The part number for the StorageWorks software package is QB-5RXA*-AA.

Service Tools - WEBES *WEB-Based Enterprise Service (WEBES)* tools integrate a high availability system fault management architecture, *Distributed Enterprise Service Tools Architecture (DESTA)*, with HP's architecture for distributed, Web-Based System Management. The tool functionality contained in the WEBES kit includes the following: *HP Analyze* (symptom-directed hardware diagnosis tool), *HP Crash Analysis Tool [CCAT]* (symptom-directed operating system software diagnosis tool), and *Revision and Configuration Management [RCM]* tool (system configuration and revision data collection tool).

HP Analyze is a hardware diagnosis software tool that provides analysis for single errors or fault events at a rudimentary level, as well as multiple event and complex analysis. *HP Analyze* provides automatic notification and isolation of hardware components to quickly identify areas of the system that may be having problems. *HP Analyze* is the successor to *DECEvent* and supports the newer EV6-based systems. Refer to the release notes for the products that are supported.

CCAT is a software application tool that helps service engineers and system managers to analyze operating system crashes. This tool collects data that describes system crashes and matches that data against a set of operating system specific rules.

The RCM tool collects system configuration and revision information. The data is stored in the RCM Server at HP Services and the server is then used to create detailed revision and configuration reports.

Hardware Information

Supported Systems and Cluster Interconnect Hardware Requirements

TruCluster Server Version 5.1B-5 supports the systems listed in the following table with up to eight systems in a configuration.

TruCluster Server supports the KZPSA-BB, KZPBA-CB, 3X-KZPBA-CC, 3X-KZPEA-DB, KGPSA-BC, KGPSA-CA, KGPSA-DA, KGPSA-EA, FCA2684, and FCA2684DC as shared storage bus adapters, subject to the current maximum number of adapters and any other restrictions for a given system. TruCluster Server supports a maximum of 62 shared buses per system in any combination.

Information on firmware release can be found at: <http://www.hp.com/support/> or from the current Alpha systems firmware update CD-ROM.

For additional information on supported hardware refer to the AlphaServer Supported Options list at the following web site: <http://h18002.www1.hp.com/alphaserver/products/options.html>.

Supported Systems

System	Shared Storage IO Adapter	Cluster Interconnect ³
AlphaServer 800	KZPSA, KZPBA-CB/CC, KGPSA-BC/CA	MC1, 1.5 & 2, 100 Mbps & 1000 Mbps LAN
AlphaServer 1000A	KZPSA	MC1.5, 100 Mbps & 1000 Mbps LAN
AlphaServer 1200	KZPSA, KZPBA-CB/CC, KGPSA-BC/CA	MC1, 1.5 & 2, 100 Mbps & 1000 Mbps LAN
AlphaServer 2000	KZPSA	MC1, 1.5, 100 Mbps & 1000 Mbps LAN
AlphaServer 2100	KZPSA	MC1, 1.5, 100 Mbps & 1000 Mbps LAN
AlphaServer 2100A	KZPSA	MC1, 1.5, 100 Mbps & 1000 Mbps LAN
AlphaServer 4000, 4100	KZPSA, KZPBA-CB/CC, KGPSA-BC/CA	MC1, 1.5 & 2, 100 Mbps & 1000 Mbps LAN
AlphaServer 8200	KZPSA, KZPBA-CB/CC, KGPSA-BC/CA	MC1, 1.5 & 2, 100 Mbps & 1000 Mbps LAN
AlphaServer 8400	KZPSA, KZPBA-CB/CC, KGPSA-BC/CA	MC1, 1.5 & 2, 100 Mbps & 1000 Mbps LAN
AlphaServer DS10	KZPBA-CB/CC, KGPSA -BC/CA/DA/EA, KZPEA-DB ¹	MC2, 100 Mbps & 1000 Mbps LAN, DEGXA-S*/T*
AlphaServer DS10L	KZPBA-CB/CC, KGPSA-DA/EA	100 Mbps LAN, DEGXA-S*/T*
AlphaServer DS15, DS15A	KZPBA-CC, KGPSA-DA/EA, FCA2684, FCA2684DC, KZPEA-DB ¹	MC2, 100 Mbps LAN, DEGXA-S*/T*
AlphaServer DS20	KZPBA-CB/CC, KGPSA-BC/CA	MC2, 100 Mbps & 1000 Mbps LAN
AlphaServer DS20E	KZPBA-CB/CC, KGPSA-BC/CA/DA/EA, KZPEA-DB ¹	MC2, 100 Mbps & 1000 Mbps LAN, DEGXA-S*/T*
AlphaServer DS20L	KZPBA-CB/CC, KGPSA-CA/DA/EA	100 Mbps & 1000 Mbps LAN, DEGXA-SA/TA
AlphaServer DS25	3X-KZPBA-CC, KGPSA-CA/DA/EA, FCA2684, FCA2684DC, KZPEA-DB ¹	MC2, 100 Mbps & 1000 Mbps LAN2, DEGXA-S*/T*
TS20 ⁴	KZPBA-CB/CC, KGPSA-CA/DA	100 Mbps & 1000 Mbps LAN
AlphaServer ES40	KZPBA-CB/CC, KGPSA -BC/CA/DA/EA, FCA2684, KZPEA-DB ¹	MC2, 100 Mbps & 1000 Mbps LAN, DEGXA-S*/T*
AlphaServer ES45	KZPBA-CB/CC, KGPSA-CA/DA/EA, FCA2684, FCA2684DC, KZPEA-DB ¹	MC2 (see notes), 100 Mbps & 1000 Mbps LAN, DEGXA-S*/T*

Hardware Information

AlphaServer ES47 server tower	KZPBA-CC ⁵ , KGPSA-DA/EA, FCA2684, FCA2684DC, KZPEA-DB ¹	MC2, 100 Mbps LAN, DEGXA-S*/T*
AlphaServer ES47	KZPBA-CC ⁵ , KGPSA-DA/EA, FCA2684, FCA2684DC, KZPEA-DB ¹	MC2, 100 Mbps LAN, DEGXA-S*/T*
AlphaServer ES80	KZPBA-CC ⁵ , KGPSA-DA/EA, FCA2684, FCA2684DC, KZPEA-DB ¹	MC2, 100 Mbps LAN, DEGXA-S*/T*
AlphaServer GS60, GS60E	KZPSA, KZPBA-CB/CC, KGPSA-BC/CA	MC1, 1.5 & 2, 100 Mbps & 1000 Mbps LAN
AlphaServer GS80	KZPBA-CB/CC, KGPSA-CA/DA/EA, FCA2684	MC2, 100 Mbps & 1000 Mbps LAN, DEGXA-S*/T*
AlphaServer GS140	KZPSA, KZPBA-CB/CC, KGPSA-BC/CA	MC1, 1.5 & 2, 100 Mbps & 1000 Mbps LAN
AlphaServer GS160	KZPBA-CB/CC, KGPSA-CA/DA/EA, FCA2684	MC2, 100 Mbps & 1000 Mbps LAN, DEGXA-S*/T*
AlphaServer GS320	KZPBA-CB/CC, KGPSA-CA/DA/EA, FCA2684	MC2, 100 Mbps & 1000 Mbps LAN, DEGXA-S*/T*
AlphaServer GS1280	KZPBA-CC ⁵ , KGPSA-DA/EA, FCA2684, FCA2684DC, KZPEA-DB ¹	MC2, 100 Mbps LAN, DEGXA-S*/T*

NOTES:

- Hard partitions of GS80, GS160, and GS320 AlphaServers can be clustered together either across separate systems or within systems, but each hard partition must have at least one cluster interconnect connection.
- TruCluster Server requires all members to be connected with either all members using Memory Channel hardware or all members using a private 100 Mbps or 1000 Mbps full-duplex LAN. Note that there are two variants of Memory Channel (MC).
- ES45 models 1, 1B, 2, and 2B support single rail Memory Channel (MC) configured for either 512 MB or 128 MB. Dual rail MC is supported with both Memory Channel adapters placed on the same PCI bus and jumpered to run at 128 MB.
- ES45 models 3 and 3B can support dual rail Memory Channel (MC) with both rails configured either for 512 MB or 128 MB. Memory Channel adapters must be placed on separate PCI busses when jumpered to run at 512 MB.
- For further information on deploying an ES45 in a TruCluster, refer to the *Cluster Hardware Configuration* manual at: http://h30097.www3.hp.com/docs/pub_page/cluster51B_list.html.

1. KZPEA-DB is supported for configuring a shared bus in a TruCluster with a maximum of two members per shared bus. Patch Kit 1 is required. Please check the specific platform Supported Options Lists to see if a particular platform is supported.

<http://h18002.www1.hp.com/alphaserver/products/options.html>

2. The embedded Gigabit Ethernet adapter (10/100/1000Mbps) on the DS25 is supported by TruCluster Server as a LAN cluster interconnect if patch kit 2 is installed.

3. 3X-DEGXA-S*/T* Gigabit Ethernet adapter is supported by TruCluster Server as a LAN cluster interconnect if patch kit 2 is installed. Please check the specific platform Supported Options Lists to see if a particular platform is supported.

<http://h18002.www1.hp.com/alphaserver/products/options.html>

4. TS20 is supported for two member LAN interconnected cluster only.

5. When installed in ES47s, ES80s, & ES1280s, the KZPBA-CC supports ONLY the HSZ80 storage controller.

Note that there are two variants of Memory Channel (MC):

Supported Memory Channel Hardware

Hardware Information

MC 1 & 1.5	MC2	Description
	CCMAB-BA	5.0 Volt/3.3 Volt Compatible
CCMAA-AA or CCMAA-BA	CCMAB-AA	PCI adapter
CCMHA-AA	CCMHB-AA	Hub
CCMLA-AA	CCMLB-AA	Line card
N/A	BN39B-04	4-meter cable
BC12N-10		10-foot copper cable
	BN39B-10	10-meter cable
	BN39B-01 (one meter)	Connects MC adapter to CCMFB optical converter
	CCMFB-AA	Fiber Optics converter
	CCMFB-BA	5.0 Volt/3.3 Volt Compatible Fiber Optics converter
	BN34R-10 (10 meter)	Fiber-optic cable: Connect one optical converter to another
	BN34R-31 (31 meter)	Fiber-optic cable: Connect one optical converter to another

Memory Channel configuration notes:

- At least one Memory Channel adapter must be installed in a PCI slot in each member system. One or more link cables are required to connect systems to each other or to a hub. A cluster environment with two nodes does not require a hub. A configuration of more than two members requires a Memory Channel hub.
- MC1 and MC1.5 are no longer saleable items and cannot be mixed on the same rail with MC2.

There are special rules about circumstances where Memory Channel 1 and Memory Channel 2 can be used together in the same cluster. The *TruCluster Server Hardware Configuration* manual provides information regarding supported Memory Channel configurations.

Supported LAN Interconnect Hardware

Use of dedicated LAN is supported for use as a cluster interconnect. A LAN interconnect must be private to cluster members. As long as any packet that is transmitted by a cluster member's interconnect adapter can only be received by interconnect adapters of other members of the same cluster, the interconnect meets the privacy requirement.

A LAN interconnect can be a direct connection between two cluster members or can employ hubs or switches. In general, any Ethernet adapter, switch, or hub that works in a standard LAN at 100 Mbps or 1000 Mbps should work within a LAN interconnect. (Adapters on combo cards such as the KZPCM, DEPVZ, and the DEPVZ are not supported.) Check the supported options list (at <http://h18002.www1.hp.com/alphaserver/products/options.html>) for the hardware platform in question to verify if the DEGXA-S*/T* is supported for LAN interconnect. Fiber Distributed Data Interface (FDDI), ATM LAN Emulation (LANE), and 10 Mbps Ethernet are not supported.

Although hubs and switches are interchangeable in most LAN interconnect configurations, switches are recommended for performance and scalability. Most hubs run in half-duplex mode and do not detect network collisions, so their use in a LAN interconnect may limit cluster performance. Overall, using a switch, rather than a hub, provides greater scalability for clusters with three or more members.

Adapters and switch ports must be configured compatibly with respect to speed (100 Mbps or 1000 Mbps) and operational mode (full-duplex). A maximum of three hops is allowed between cluster members, where a hop means passing from a system, switch, hub, or router, to another system, switch,

Hardware Information

hub, or router. That is, any combination of up to two hubs, switches, or routers is supported between two cluster members. You must not introduce unacceptable latencies by using, for example, a satellite uplink or a wide area network (WAN) in the path between two components of a LAN interconnect.

A fully redundant LAN interconnect configuration employs two or more Ethernet adapters in a NetRAIN set on each member, with redundant wiring to two or more switches interlinked by two crossover cables. These Ethernet switches must be capable of one of the follow mechanisms for managing traffic across parallel inter-switch links: link aggregation (also known as port trunking), resilient links, or per-port-enabled spanning tree algorithm.

Link aggregation of Ethernet adapters using Tru64 UNIX features (including the lagconfig command) is not supported for a LAN interconnect.

Fibre Channel Arbitrated Loop Support TruCluster Server supports Fibre Channel arbitrated loop for clusters with a maximum of two members only. The DS-SWXHB-07 seven-port Fibre Channel hub is required to build an arbitrated loop configuration and is restricted to use with DS10, DS20, DS20E, and ES40 systems. The other AlphaServers including the DS10L and ES45 systems are not supported for arbitrated loop.

Supported SCSI Controllers The 3X-KZPEA-DB is the currently saleable SCSI controller supported in a TruCluster environment. More information on the 3X-KZPEA-DB is available from the following web site:
<http://h18002.www1.hp.com/alphaserver/products/storage/kzpea/index.html>.

The TruCluster Server Cluster Hardware Configuration manual and Release Notes also provide information regarding SCSI controller configuration.

Supported Tape Devices and Media Changers TruCluster Server supports the configuration of specific tape devices on a shared SCSI bus and on Fibre Channel. These devices will function properly in a multi-initiator environment. Refer to the AlphaServer Supported Options list at the following web site for more information on tape devices:
<http://h18002.www1.hp.com/alphaserver/products/options.html>.

Backup software must be explicitly capable of handling and recovering from such events and must utilize the cluster application availability (CAA) facility to facilitate highly available backup. For more information on supported backup software refer to the EBS Compatibility Matrix available from the following web site: <http://h18000.www1.hp.com/products/storageworks/ebs/>.

Hardware Restrictions TruCluster Server has the following hardware restrictions. The TruCluster Server Cluster *Hardware Configuration* manual provides additional information regarding hardware restrictions.

- Prestoserve NVRAM failover is not supported on shared disk devices.

TruCluster Server supports up to eight member cluster configurations as follows:

- **Switched Fibre Channel:** Eight-member systems may be connected to common storage over Fibre Channel in a fabric (switch) configuration.
- **Parallel SCSI:** Only four of the member systems may be connected to any one SCSI bus. Multiple SCSI buses may be connected to different sets of members and the sets of members may overlap. Use of a DS-DWZZH-05 UltraSCSI hub with fair arbitration enabled is recommended when connecting four member systems to a common SCSI bus. Use of fair arbitration drives are recommended when the KZPEA shared bus is used.

Hardware Information

Hardware Configuration Examples The TruCluster Server *Hardware Configuration* manual provides hardware configuration examples.

Updates to Supported Hardware The TruCluster Server QuickSpecs is updated and corrected periodically to reflect new hardware options and platforms support. Please check online for the latest revision. Go to <http://h18000.www1.hp.com/products/quickspecs/productbulletin.html> and select either "Worldwide QuickSpecs" or "U.S. QuickSpecs" and then navigate through "High Availability and Clustering" to "Tru64 UNIX Clustering."

© Copyright 2003-2008 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice.

Windows is a trademark of Microsoft Corporation in the U.S. and/or other countries. UNIX and X/Open are trademarks of The Open Group in the U.S. and/or other countries. Informix and Informix Extended Parallel Server are trademarks of Informix Software, Inc. NetWorker and Prestoserve are trademarks of Legato Systems, Inc. Oracle is a trademark of the Oracle Corporation in the U.S. and/or other countries. NFS is a trademark of Sun Microsystems, Inc. All other product names mentioned herein may be the trademarks of their respective companies.

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.