



white paper

Tru64 UNIX
and EV7-based
AlphaServer
systems

april 2003

enabling server consolidation with HP Tru64 UNIX® system partitions and resource management utilities

table of contents

- introduction** 2
- addressing business challenges** 2
 - supporting separate application and test environments 3
 - high availability and workload balancing with clusters 3
- system partitions on EV7-based AlphaServer systems** 5
 - configuring partitions on EV7-based AlphaServer systems 6
 - benefits of system partitions 6
- resource management** 7
 - managing CPU resources 7
 - network and storage resource management 7
- conclusion** 8
- for more information** 8

Abstract: Workload management – spanning system partitions (hardware partitioning) and resource management – is critical to helping IT organizations respond more rapidly and efficiently to changing business demands.

These capabilities are essential to achieving a successful server consolidation initiative and fundamental in enabling flexible workload management.

System partitions and resource management offer extensive benefits, including:

- reduced hardware costs
- increased operational flexibility
- extended scalability
- enhanced security
- greater control of resources
- improved application performance
- simplified management



This white paper defines both system partitions and resource management, and discusses how they are implemented in the current version of HP Tru64 UNIX running on the newest AlphaServer systems. The newest AlphaServer systems are based on the EV7 Alpha chip; to distinguish their implementation details, this paper refers to these systems as EV7-based AlphaServer systems. In addition, the paper will explore how system partitions and resource management can be used independently or in combination to solve real business problems using the advanced capabilities of Tru64 UNIX on AlphaServer systems. Figure 1 shows the family of EV7-based AlphaServer systems.

figure 1



HP AlphaServer ES47 system
up to 2 partitions



HP AlphaServer ES80 system
up to 4 partitions



HP AlphaServer GS1280 system
up to 8 partitions

introduction

These days, it would be nearly impossible to find an IT operation that's not under intense pressure to support a wide range of applications with no downtime and ultimate flexibility, while reducing cost and complexity. To meet these apparently conflicting goals, IT organizations today are consolidating computing resources to keep costs down, and using dynamic resource management and system partitions to keep availability and flexibility up.

addressing business challenges

System partitions and resource management can be applied very effectively to address a variety of business challenges. Among the most common challenges are supporting separate applications for production and test, and allocating appropriate levels of computing resources to multiple applications or user communities. To address these challenges in the simplest, most economical way, many enterprises are turning to server consolidation.

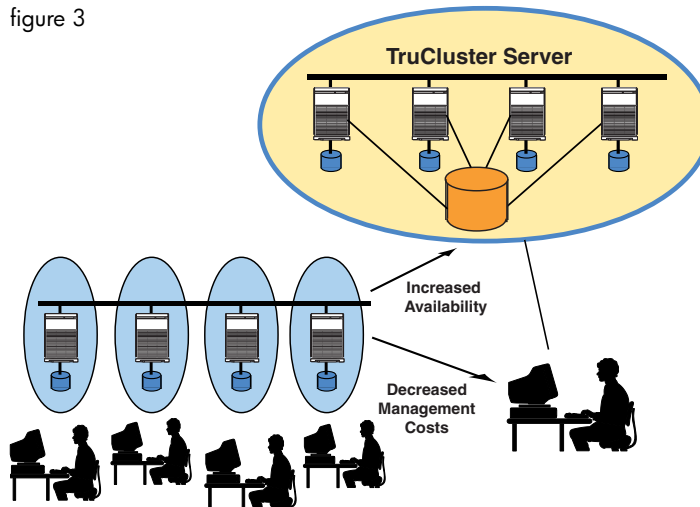
supporting separate application and test environments

One approach to supporting separate application environments is simply to purchase separate physical servers: for example, one for production and one for test. In demanding environments, however, the servers supporting that environment are often large-scale, and the expense of purchasing and managing a second server may not be feasible. In this scenario, using separate, clustered physical machines to achieve the necessary separation of test and production and maintaining high availability, results in two clusters.

While this approach is effective, consolidating servers and taking advantage of system partitions presents a much more cost-effective and more easily managed alternative. With the EV7-based AlphaServer system and Tru64 UNIX, HP has set a high watermark for system partitioning. The EV7-based AlphaServer system supports up to eight system partitions, introducing unprecedented flexibility to support diverse business requirements.

Instead of separate physical servers, system partitions offer hardware isolation between applications on a single EV7-based AlphaServer system. These system partitions prevent applications running in one partition from interfering with applications running in another, or from accessing hardware resources owned by another partition. In addition, individual system partitions are capable of not only running different versions of the same operating system, but also running different operating systems entirely. That means, production applications can be run safely while testing is performed on new releases concurrently—with lower hardware expenses thanks to the shared platform infrastructure. Figure 3 provides an illustration of this capability.

figure 3



System partitions on the EV7-based AlphaServer system may also be used to support online maintenance and repair. That is, if certain system components require attention, a new partition can be created for just those components. Then, while the system continues to run applications, the problem can be diagnosed on the isolated partition and action taken without affecting the operation of other production applications.

high availability and workload balancing with clusters

High availability can be achieved for separate physical servers, as well as system partitions. This is where the TruCluster Server software comes into play. TruCluster Server software enhances Tru64 UNIX in configurations of both multiple systems and partitioned single systems by providing automated load balancing and featuring a single, cluster-wide file system that is tightly integrated with the Tru64 UNIX operating system.

With automated load balancing, clients that are connecting to the cluster are assigned to a cluster member on a round-robin basis—whether a standalone server or a partition. The system administrator determines which system will handle the most connections and can adjust the load sharing online as needed. By continuously rebalancing client connections, TruCluster Server v5 ensures optimum performance and resource utilization. Even as workload intensifies, systems and partitions can be easily added or resized without affecting user service.

By managing a system with clustered partitions—or multiple systems (partitioned or not)—with a single file system, system administration is greatly simplified, lowering management costs while enjoying the added flexibility. The single system image technology also enhances storage management cluster-wide, providing improved data management, higher availability, and lower costs. Figure 4 indicates how TruCluster Server software enables an administrator to manage multiple systems as a single system.

figure 4

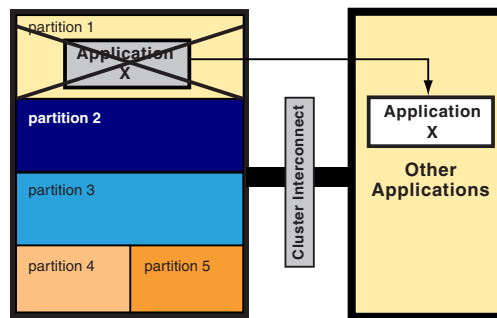



figure 5



In addition to supporting multiple clustered servers, TruCluster Server software also enables clustering of system partitions within a single server, as well as system partitions across multiple servers, with the same ease of management. Figure 5 illustrates how two partitions within a single server can be clustered.

Figure 6 below shows clustering of two partitions across two separate servers.




figure 6

Should application requirements change, additional servers or system partitions can easily be added to the cluster dynamically. In fact, combined with the system partitioning capability of the EV7-based AlphaServer system, HP provides an unprecedented level of flexibility to expand or reconfigure systems to meet changing needs: 1) by creating additional partitions within a single server, 2) by expanding the number of servers and system partitions across those servers, and 3) by expanding clusters across both physical servers and system partitions.

system partitions on EV7-based AlphaServer systems

Today, HP provides system partitioning on its EV7-based AlphaServer systems. System partitioning is the process of dividing system resources – processors, memory and I/O bandwidth – into subsets where each subset is capable of running an instance of an operating system. An undivided system can be considered a single hard partition. Each partition must contain some quantity of each resource type to allow booting an operating system instance and running applications.

System partitions allow a server to be segmented into several smaller systems and recombined back into a single server as necessary. This capability is ideal whenever multiple instances of an operating system are needed, such as for testing new releases, performing upgrades, or segregating applications. With multiple system partitions, a single server can support evolving application needs and requirements.

For example, each partition can be configured to run in a variety of combinations, including:

- the same version of the same operating system—this would enable different applications to be used in a mixed environment of production, development, and test
- different versions of the same operating system—this would support legacy applications or testing new versions of operating systems, databases, or applications
- different operating systems—this provides the foundation for server consolidation or workload management, and also would support existing applications on one operating system, while running new applications on a different operating system

Applications and operating systems running in different system partitions are completely isolated and protected from one another by hardware-enforced firewalls. An application or operating system failure in one independent system partition will have no effect on applications running in other partitions. This can help especially in environments where individual business resources must be kept separate, or where certain applications that are not particularly well behaved must reside on separate systems. In addition, clustering may be desirable to achieve necessary levels of availability for individual partitions or entire systems.

System partitions are also very useful when an organization is running multiple applications, each requiring different versions of various operating systems or databases. This is a powerful way to consolidate single or multiple applications from many distributed servers onto a single system or onto a few systems while maintaining necessary security and data protection. The key to optimizing system utilization in these scenarios is to manage system partitions flexibly and efficiently. HP provides the HP AlphaServer Partition Manager, an intuitive management application that enables an administrator to respond to changing workload requirements by creating, adding, or deleting system partitions quickly and easily. The AlphaServer Partition Manager also enables an administrator to detect if any part of the system is idle, and to rectify the situation by creating a partition that allows that part of the system to become active.

The AlphaServer Partition Manager provides a user-friendly graphical interface that can be launched as an application from the AlphaServer System Management Console, without the need to know console commands. It can also be run through a web browser, making it possible to manage system partitions from anywhere at any time. The AlphaServer Partition Manager is an important advance toward enabling the dynamic configuration of partitions, which will allow administrators to create, resize, or delete partitions online without rebooting the system.

configuring partitions on EV7-based AlphaServer systems

All EV7-based AlphaServer systems support hard partitions (except for the HP AlphaServer ES47 tower system). Partitions must contain at least two processors but can contain a larger number.

If the system is partitioned at building block boundaries – eight-processor (8P) drawers in the HP AlphaServer GS1280 series, and the 2P drawers in the HP AlphaServer ES47 and ES80 systems – then the hardware partitions are completely electrically isolated. This means that any error, whether hardware or software, in one partition will have no effect on other running partitions. Partitioned systems configured this way have no single points of failure.

Shared system elements provide for modular growth and investment protection in the EV7-based AlphaServer SMP systems. The major components of the AlphaServer ES47, ES80, and GS1280 systems are the dual-processor module, system building block drawers and I/O building block drawers.

The dual-processor modular building block is built with two EV7 system chips, memory cards, power regulators and a CPU management card.

There are two types of system drawers: the 2P drawer contains one dual-processor module, the 8P drawer can contain up to four dual-processor modules. Up to four 2P drawers can be connected to create an eight-processor system; eight 8P drawers can be connected to create a 64-processor system.

benefits of system partitions

System partitions offer a number of distinct benefits:

- increased system utilization, achieved by sizing partitions according to the requirements of individual applications
- higher efficiency, by running different operating systems and applications for production or testing in a single system
- improved system reliability through the isolation of most software faults
- dynamic configuration to expand or contract the system environment as needed
- enhanced security for specific applications by isolating security requirements within a partition
- reduced hardware and operations costs through consolidating multiple servers into a single partitioned SMP system
- greater operational flexibility, enabling resources to expand without moving the physical hardware

System partitions provide a solid foundation for meeting the larger objectives of an IT organization to reduce costs, simplify system management, improve security, enhance system and application availability, and increase overall business flexibility. These benefits are further enhanced by the extensive resource management capabilities supported by Tru64 UNIX.

resource management

Resource management complements system partitions by providing more fine-grained system control. Tru64 UNIX offers a variety of resource management capabilities within the operating system that give system administrators the flexibility to tailor a system to fit specific business requirements.

The key to effective system control is the management of critical system resources, including CPU, disk, and network bandwidth. Through a tightly integrated set of tools and utilities, Tru64 UNIX provides various levels of control of these resources within single servers, across system partitions, and among clustered servers.

managing CPU resources

HP has a long history of innovation in the realm of resource management, and Tru64 UNIX supports a variety of tools to manage CPU resources.

Processor sets and class scheduling have been provided with the Tru64 UNIX operating system since 1997. With processor sets, individual CPUs can be grouped and dedicated to specific applications or users. CPUs can then be moved between processor sets administratively or, should one set require additional resources during peak load periods, through the resource usage monitoring capabilities provided within the Tru64 UNIX operating system.

Class scheduling adds further definition to how CPU resources can be managed by assigning users and applications to a class, which is allotted a percentage of the CPU time available on the system—as a whole or within processor sets. Up to 100 different classes may be defined for a given system or system partition, enabling an administrator to apportion the necessary CPU cycles among all the jobs running on the system. As a result, the risk that a highly compute-intensive job will dominate the system and degrade services for other users can be effectively mitigated.

network and storage resource management

A network bandwidth tool is also provided in Tru64 UNIX. This tool allows administrators to limit network bandwidth at the port level—an important capability that will continue to be enhanced in future releases.

Figure 7 below shows a storage area network with the I/O resources that must be managed.




figure 7

Finally, disk quota tools are provided that restrict the amount of file system space available for users or groups by establishing quotas for the file system. Both soft and hard limits can be specified, with soft limits issuing a warning and hard limits preventing writes to the file system.

Today, it is possible to control disk I/O for applications by segregating application data to different file systems on different physical disks and different buses. With HP Storage Area Networks (SANs), the utilization of storage bandwidth across the available SAN connections is an important resource to be managed. Future releases of Tru64 UNIX will support allocation of storage bandwidth in a SAN configuration.

conclusion

System partitions and resource management can solve numerous business challenges, providing extensive flexibility to support various application and user demands without the need for purchasing separate physical servers. System partitions make it possible to consolidate computing resources for cost advantages. Finally, the addition of clustering balances workload while providing the highest levels of availability, scalability, and manageability for two or more system partitions.

With the EV7-based AlphaServer systems and Tru64 UNIX v5.1B, HP has reinforced its leadership position in high-performance, high-availability solutions for enterprise IT, offering an extensive array of system partitioning and resource management capabilities.

Workload management with Tru64 UNIX and EV7-based AlphaServer Systems: functions at a glance

the capability	its function	granularity
hardware partitioning	1. Expand or contract the system environment as resource requirements vary, and 2. Enhance security for specific applications by isolating applications within a partition	Minimum of 2 processors per partition; maximum of eight 8-processor drawers on one partition
class scheduler	Dynamic management of system resources – allows the system administrator to define a set of classes, where each class has a maximum allowable CPU usage percentage	A “class” is a collection of: Group IDeNtifiers (gid), User IDeNtifiers (uid), Sessions (sess), Process GRouPs (pgrp); and Process IDeNtifiers (pid), any combination
processor sets	Dynamic management of system resources – dedicating CPUs to critical tasks	By CPU, restricts which CPUs are available to a process, or allows exclusive use of a processor or set of processors by a process
clustering with TruCluster Server software	Ensures application and data availability; simplifies management of clustered nodes; balances incoming requests for CPU resources and data access across the nodes in the cluster	Systems or partitions defined as part of the cluster

for more information

For more information, contact the Tru64 UNIX software division at www.hp.com/go/tru64unix and the AlphaServer division at www.hp.com/servers/alphaserver

The information in this document is subject to change without notice.

UNIX is a registered trademark of The Open Group in the U.S. and other countries.

©2003 Hewlett-Packard Company 04/2003

P/N 5981-6396EN