Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Fujitsu's PRIMEPOWER: Partitioning Capabilities That Enable Consolidation

Sponsored by: Fujitsu

Vernon Turner Stephen L. Josselyn
Jean S. Bozman

April 2003

EXECUTIVE SUMMARY

Server consolidation, the rehosting of workloads supported by a multitude of smaller servers to a few larger servers, promises cost savings for IT managers by reducing the cost of management and increasing utilization of server resources. To support multiple workloads, servers must meet three key criteria:

- System managers must be able to partition server resources and provision workloads with partitions of different sizes.
- System managers must be able to reallocate server resources, both manually and automatically, as workload demands shift.
- Servers that support consolidation must be highly available so that the many, centralized workloads of an enterprise do not go out of service.

Fujitsu's PRIMEPOWER line of servers is designed to host consolidated workloads. PRIMEPOWER models 900, 1500, and 2500 scale to 16, 32, and 128 SPARC64 V RISC processors, respectively, thus providing a wide range of scalable, extensible server configurations. Fujitsu provides physical partitioning (PPAR) at the system-board level and a finer-grained extended partitioning (XPAR) that allows system managers to assign workloads to processors on a system board. Availability is enhanced because physical partitions provide the greatest isolation of software faults.

Fujitsu has developed a cyclical model of server management. Managers must first *visualize* the operations of the server, then *detect* bottlenecks that may occur as workloads shift, and finally *reallocate* resources to better utilize the server's resources. Aligned with the cyclic model is Fujitsu's system management software, WebSysAdmin, which provides a graphical view of server operations to aid in assigning processing resources to workloads. Server resources can also be dynamically reallocated automatically, driven by policies (rules) that identify the priority for different workloads.

To ascertain user requirements for server consolidation, IDC recently surveyed 400 IT managers and found that cost savings was their primary objective in undertaking a server consolidation project. Nearly half of the managers indicated that financial analyses (e.g., return on investment [ROI] estimates) were required to justify the consolidation decision.

IDC believes that Fujitsu's PRIMEPOWER partitioning capability is aligned with industry trends and provides functionality that users will need to succeed in reducing IT spend by consolidating server resources. Fujitsu's focus on availability and manageability (including flexible partitioning) is on track to make PRIMEPOWER a competitive offering.

INTRODUCTION

IDC believes that chief executives have two major concerns about IT systems: high availability and cost efficiency. The need for high levels of availability has emerged over the past decade as businesses have become more dependent on IT systems. When IT systems are down, most businesses are out of business. In short, there is no time for downtime, especially for businesses that are networked to their suppliers and customers.

Making computing systems highly available competes with the goal of cost efficiency. Generally, IT systems achieve greater availability by eliminating single points of failure; for example, by having two power supplies rather than one. Replicated systems with duplicate hardware and software components, however, cost more to deploy and maintain. Moreover, in the traditional primary/backup architecture for datacenters, backup resources are not always fully utilized.

Consolidating server and storage resources is one way to attack the problem of achieving availability and cost efficiency at the same time. Consolidated server resources can be reallocated as necessary to provide flexible processing capacity for different enterprise workloads — an opportunity that is not easily provided by multiple, smaller servers. Consolidated server resources are easier to manage, which adds to the cost savings.

SERVER TECHNOLOGY TRENDS

Server suppliers are designing new platforms to support IT system consolidation. These enterprise servers are designed to be highly available and may contain up to 128 processors. Such scalable systems are ideal for server consolidation — that is, for consolidating computing workloads that had been running on many distributed server machines. The key business benefits of consolidation are unified management and centralized maintenance from a central-site IT staff.

In the server consolidation approach, server resources must be partitioned to provide dedicated resources that are matched to the needs of different applications and workloads. Partitions can be implemented with physical or logical boundaries as follows:

- Physical partitions are defined by a collection of processors dedicated to a workload. Physical partitions are immune to loads that are placed upon the other processors in the server. They are also isolated from faults that may occur elsewhere in the server. Each physical partition runs a single instance of a specific operating system. Multiple partitions within the same server can run different versions of the operating system, if desired, which allows older versions of operating system software to support older workloads and newer versions to run within partitions for test and development purposes.

- Logical partitions are governed by software and allocate a pool of processing resources to a collection of workloads. These partitions are achieved through the use of virtualization software that supports instances of operating systems that run in one or more physical partitions. Logical, software-defined partitions can be reconfigured dynamically to shift server resources among workloads as system demands change.

More sophisticated management tools are needed to support server consolidation. In particular, the IT staff needs tools that manage the partitioning process. Policy-based management tools provide automated management based on rules that IT staff can use to describe how different workloads are to be supported. For example, policies can call for a shifting of server resources overnight to run batch-processing

applications, and they can reassign those resources to support online transaction processing applications during business hours. Policies can also determine priority for applications that compete for limited server resources.

USER REQUIREMENTS

IDC research indicates that server and storage consolidation is a recurring theme for projects under way in IT departments of large and midsize enterprises. Although most IT departments are *centralizing* the location of servers and storage systems, others are taking a *physical consolidation* approach by relocating workloads and data sets on fewer larger servers and storage systems. Three other less frequently used approaches focus on *data integration*, *application integration*, and *storage consolidation*.

SCALE UP VERSUS SCALE OUT

Server consolidation may involve bringing workloads from several smaller servers to a single larger server. This process, which IDC calls *physical consolidation*, is informally called *scaling up*. Server consolidation may refer to the process of colocating servers scattered across an enterprise to a few centralized datacenters. When more processing capacity is needed, then more servers are added. This process is informally called *scaling out*.

Both forms of consolidation attack the high cost of maintenance for multitudes of servers while maintaining an explicit method for scaling to meet increasing capacity demands from the enterprise. Scaling up or out is all about growth, which is one important concern.

Consolidating to fewer larger servers and reallocating (versus simply adding) server resources by partitioning is the focus of this white paper. In principle, multiple smaller servers could be woven into a cluster so that the cluster's resources could be reallocated, but this is not common practice in IT today. The cluster approach is an alternative with its own set of opportunities and challenges.

IDC found that the most common expected benefit of consolidation among the 400 IT managers surveyed was the reduction of IT costs. These managers cited cost savings twice as often as any other benefit. Improved performance and ease of management were other top benefits. Given the importance of cost reduction, nearly half of the IT managers indicated that a formal ROI estimate was required to justify consolidation. Not surprisingly, conducting ROI analyses was more common in larger companies than in midsize companies.

When considering server consolidation, IT planners quickly realize that availability is an increasingly important system requirement. When workloads are consolidated from multiple servers to fewer larger servers, then the failure of a server could lead to widespread IT outages in the enterprise. Selecting servers that are engineered with redundant components (e.g., multiple power supplies and cooling fans) that can be replaced without taking the system out of service is essential.

On average, IDC's sample of 400 IT managers reported more than 200 servers per installation, and servers with Windows and Unix operating environments were most commonly targeted for consolidation. These managers expected support for Unix, Linux, and Microsoft operating environments in their new consolidated architectures.

FUJITSU'S PRIMEPOWER SERVERS

Fujitsu provides three enterprise servers in its PRIMEPOWER product line: the PRIMEPOWER 900, 1500, and 2500. All PRIMEPOWER servers are based on the SPARC64 V RISC processor and the Solaris operating system. As Table 1 shows, PRIMEPOWER models 900, 1500, and 2500 support up to 16, 32, and 128 SPARC processors, respectively.

TABLE 1

PRIMEPOWER SERVER TECHNICAL DATA

	PRIMEPOWER 900	PRIMEPOWER 1500	PRIMEPOWER 2500
SPARC64 V RISC processors	1 to 8 per system board; 16 maximum per server	1 to 8 per system board; 32 maximum per server	2 to 8 per system board; 128 maximum per server
Memory	2GB to 32GB per system board; 64GB maximum per server	2GB to 32GB per system board; 128GB maximum per server	4GB to 32GB per system board; 512GB maximum per server
PCI slots	8 per system board; 16 maximum per server; additional 20 with PCI/Disk Box	8 per system board; 32 maximum per server; additional 40 with PCI/Disk Box	0 per server; additional 320 with PCI/Disk Box
Footprint	16 U for base cabinet; 8 U for PCI/Disk Box	740 x 1,114 x 1,800mm floor stand for base cabinet	1,066 x 1,144 x 1,800mm floor stand for base cabinet; 1,066 x 1,788 x 1,800mm floor stand for base cabinet and system expansion cabinet
Physical partitions (PPARs) and extended partitions (XPARs)	1 processor minimum; 2 PPAR maximum per server; 8 XPAR maximum per server	1 processor minimum; 4 PPAR maximum per server; 15 XPAR maximum per server	2 processor minimum; 15 PPAR maximum per server; 15 XPAR maximum per server

Source: Fujitsu, 2003

PRIMEPOWER PARTITIONING CAPABILITIES

Fujitsu provides PRIMEPOWER servers with two physical partitioning capabilities — board-level PPAR and processor-level XPAR. WebSysAdmin browser-based system management software manages the PPAR and XPAR capabilities. It enables administrators to assign as many PPARs as there are system boards in the PRIMEPOWER server. System boards have a minimum and maximum number of processors, as shown in Table 1, and the number of processors installed will affect the granularity of PPARs. Hardware and software fault isolation is high — hardware failure of an entire system board will not affect processing on system boards in other PPARs.

A PRIMEPOWER 900 server configured with two system boards and eight processors per system board can support a maximum of two PPARs. As shown in Figure 1, however, XPAR functionality provides a mechanism for establishing four partitions in support of three workloads. Eight processors — four from each system board — are partitioned in XPARs #0 and #1 and then joined with Fujitsu PRIMECLUSTER to form an eight-processor cluster that serves an accounting workload. As Figure 1 also shows, XPAR partitions #2 and #3 are formed with four processors from each system board and provisioned in support of a data mining application and an application development workload, respectively.

FIGURE 1

XPAR PARTITIONING ON THE PRIMEPOWER 900

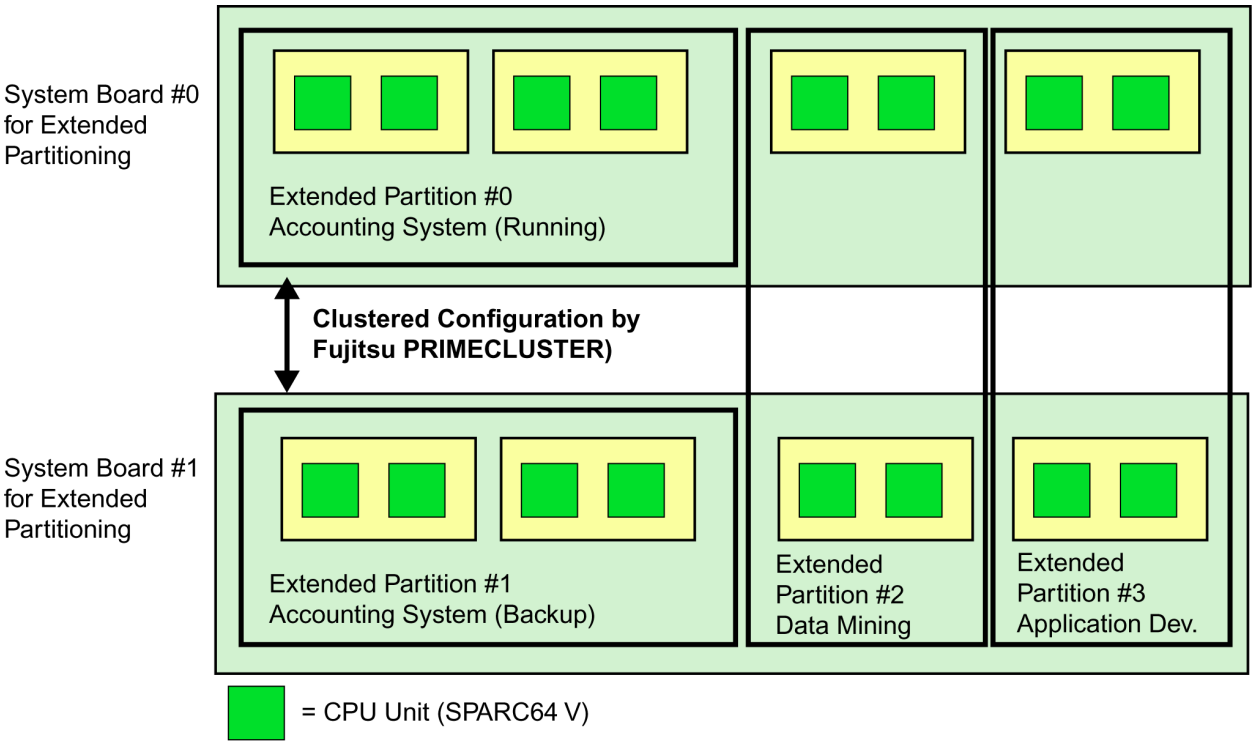

Source: Fujitsu, 2003

Fujitsu's XPAR is capable of providing a partition as small as a single SPARC V processor for PRIMEPOWER 900 and 1500 and as small as two SPARC V processors for the PRIMEPOWER 2500. XPAR provides finer granularity than PPAR and enables partitions to be dynamically reallocated. The advantage of XPAR-based floating operation is the granularity of operated resources. The functionality, offered via WebSysAdmin, permits system administrators to plan for the automatic shifting of server resources, a common challenge in today's datacenter. Using floating operations, for example, the PRIMEPOWER 900 shown in Figure 1 can reassign two processors from XPAR #3 to XPAR #2 if more processing power is needed for the data mining workload and demand for application development processing is low.

MANAGEMENT CAPABILITIES

Fujitsu's management software, which is built in and packaged with all products in the PRIMEPOWER server line, is based on the following three-step cyclical model of system management (see Figure 2):

1. Server administrators investigating system performance first need to see what the server is doing — that is, to *visualize the server's operations and resources*.

2. To improve server performance, administrators next need tools to *detect bottlenecks* — that is, to identify the critical resources that limit the performance of the server. Inadequate main memory and insufficient processors assigned to a workload are examples of resources that may be throttling the server's performance for a particular workload.

3. Performance is improved when *resources can be reallocated*, which is the third step in the cycle. Reallocation allows system managers to *provision* server resources to workloads and workgroups. Provisioning is the key to making sure that the server's tasks are aligned with the IT department's priorities.

FIGURE 2

FUJITSU'S CYCLICAL MODEL OF SERVER MANAGEMENT

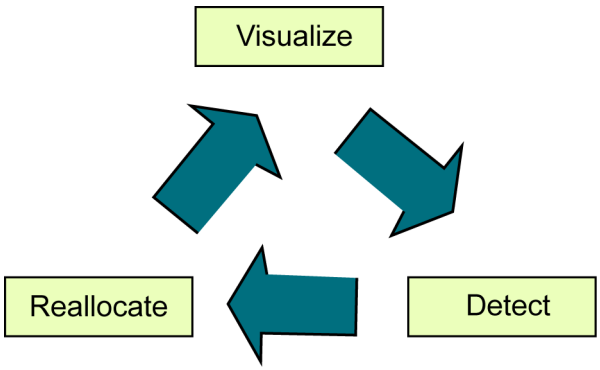

Source: Fujitsu, 2003

When the server's resources have been reallocated, the cycle begins again. Visualization capabilities allow the administrator to see how the server's resources are being deployed. Currently, Fujitsu's WebSysAdmin requires administrator intervention. In the future, Fujitsu intends to use the visualize-detect-reallocate model as the basis for system management that occurs automatically.

Figure 3 shows how Fujitsu's WebSysAdmin system management software detects a hardware fault. In this example, a CPU has failed and WebSysAdmin provides its exact location, including the slot where the CPU is plugged in, the system board, and the server. An illustration, on the right in Figure 3, shows technicians where to access the system board to replace the faulty CPU. This is a simple example that shows WebSysAdmin in action.

FIGURE 3

FUJITSU'S WEBSYSADMIN IN ACTION

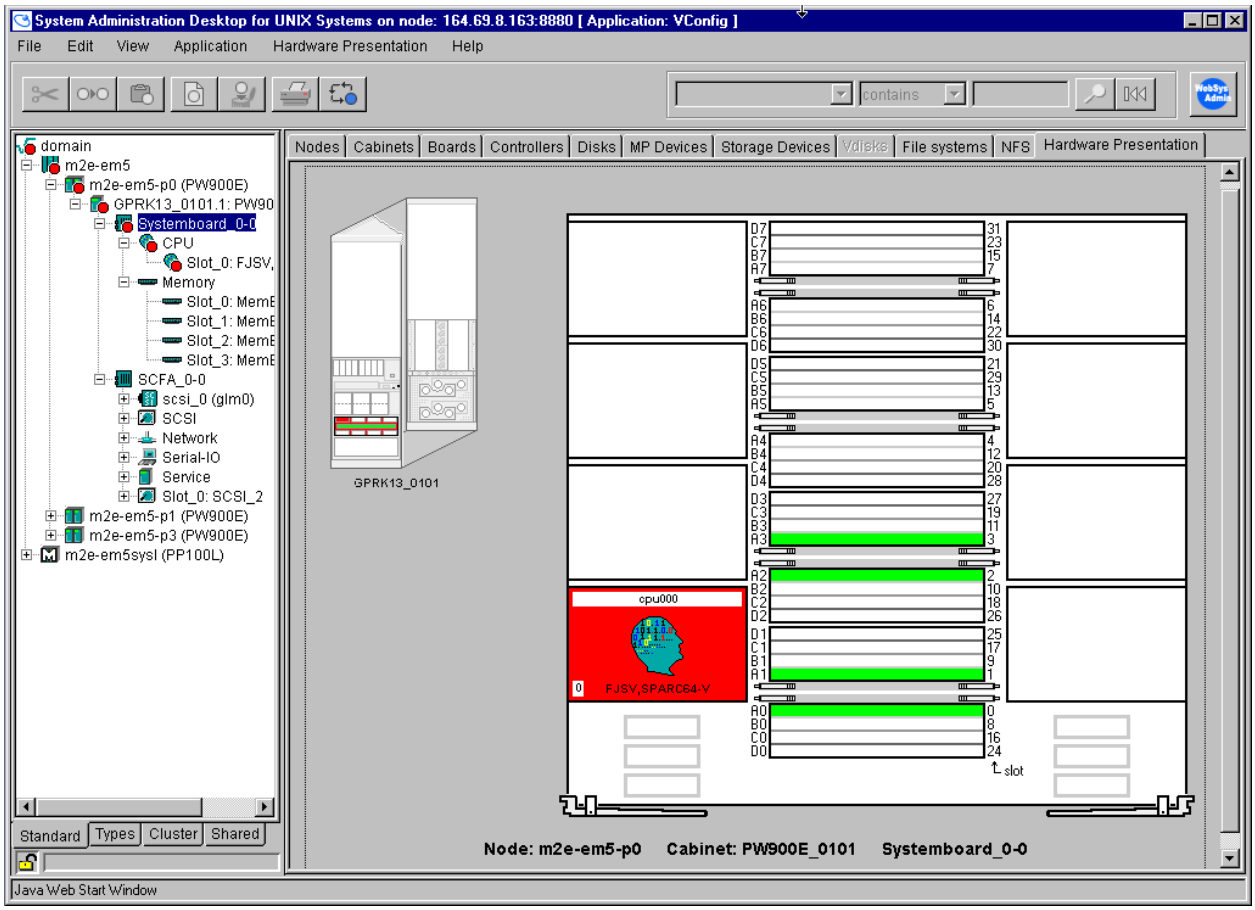

Source: Fujitsu, 2003

PRIMEPOWER AS A PLATFORM FOR CONSOLIDATION

Fujitsu's PRIMEPOWER line of servers are designed to support server consolidation. The PRIMEPOWER design provides four critical features that IT managers demand — scalability, availability, manageability, and the ability to reallocate the server's resources.

- Scalability is shown by the top-of-the-line PRIMEPOWER 2500, which can house as many as 128 SPARC64 V processors. Configurations of multiple PRIMEPOWER servers should provide sufficient processing to satisfy large enterprises. PRIMEPOWER 900 and 1500 servers are available to support smaller workloads.

- Availability is designed into all of the PRIMEPOWER servers. Redundant power, storage, cooling, and other system controls eliminate single points of failure. Hot swappable system design allows the PRIMEPOWER series to be maintained without being taken out of service.

- Manageability is supported by Fujitsu's *visualize-detect-reallocate* server management model, which is the basis for functionality built into WebSysAdmin, Fujitsu's software that monitors hardware and workload status across a collection of PRIMEPOWER servers.

- Flexible allocation of system resources is achieved with Fujitsu's XPARs. Server resources can be partitioned at the CPU level to match server allocation to different application needs. Floating resources can be automatically applied to workloads as well.

IDC ANALYSIS: OPPORTUNITIES AND CHALLENGES

OPPORTUNITIES

IDC believes that customers in the market for midsize and enterprise-scale servers will compare the features and functions that suppliers offer in support of server and system consolidation. Thus, Fujitsu's focus on partitioning capabilities for its PRIMEPOWER line of servers is timely. IT buyers are ready to understand how they can improve utilization, decrease the cost of maintenance, and automate the provisioning of centralized server resources.

Clustering and partitioning are the basic mechanisms for allocating consolidated server resources to workloads. Fujitsu's PRIMEPOWER servers integrates both technologies so that demanding workloads can tap the processing power of several processors and light workloads can be matched to just one or two processors. Clustering and partitioning, along with resilient hardware features such as multiple power supplies, address the need for high availability as well.

IDC research indicates that reducing the cost of IT is a top priority for IT managers. Fujitsu is addressing this concern with tools for the PRIMEPOWER servers that provide a graphical view of server performance, detect bottlenecks, and provide manual and automatic mechanisms to reallocate resources. Fujitsu's focus on simplifying and automating server management should be welcome news to customers.

CHALLENGES

IDC notes that system suppliers are pursuing mindshare with respect to partitioning technologies and promises of improved system management. Customers will challenge Fujitsu, along with its competitors, to provide evidence of cost savings based on the experience of early adopters who should be completing server consolidation projects at this time.

Physical partitioning is a strength for Fujitsu, and the XPAR approach takes physical partitioning to its finest level of granularity — the single processor. Logical partitioning provides the ability to share the resources of a single processor with multiple workloads, a useful function for many IT customers. Fujitsu will be challenged by competitors that provide clustering, physical partitioning, and logical partitioning functionality. Fujitsu must be prepared to explain the strengths of its partitioning approach with practical examples of workload balancing using XPAR.

IDC research indicates that users are increasingly interested in a single server platform capable of supporting multiple operating environments. Unix users look to Linux with special interest because Linux is the gateway to Open Source software options. From an IT skills perspective, Unix administrators find Linux techniques familiar. Support for Microsoft server operating environments allows consolidation to include workloads such as Outlook. Readers should note that Fujitsu does offer other server product lines that run the Microsoft and Linux operating environments.

CONCLUSION

IDC studies indicate that IT managers, driven to reduce costs, are moving workloads from multiple smaller servers to fewer larger servers. Consolidation requires highly available servers with resources that can be partitioned to match shifting workloads. Fujitsu's PRIMEPOWER line of servers addresses consolidation requirements. IDC encourages IT managers planning for server consolidation to evaluate Fujitsu's offerings.