TECHNOLOGY BRIEF

January 2002

Compaq Computer Corporation

ISSG Technology Communications

Compaq Resource Partitioning Manager: A Server Consolidation and Optimization Solution for Microsoft Windows 2000

CONTENTS

ABSTRACT

To avoid resource conflicts that may arise when running multiple applications on a server, many system administrators run each application on a dedicated server or servers. However, this approach rapidly expands server networks, which increases management complexity and overhead costs.

Compaq has developed a resource management tool to enable workload consolidation on *Compaq ProLiant* servers running Microsoft Windows 2000 Server, Advanced Server, and DataCenter operating systems. *Compaq Workload Management Pack,* featuring Resource Partitioning Manager (RPM) increases the stability and availability of applications under Windows 2000, thereby allowing customers to confidently deploy multiple applications on a single server. RPM builds upon the embedded job object technology in Windows 2000 to provide a quick and easy way to manage processor and memory resources dynamically.

This technology brief addresses only industry-standard servers running the Windows 2000 operating system. It is assumed that the reader has some familiarity with the Windows 2000 operating system.

COMPAQ

Please direct comments regarding this communication to the ISSG Technology Communications Group at TechCom@compaq.com.

NOTICE

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

Compaq Resource Partitioning Manager: A Server Consolidation Solution For Windows 2000
Second Edition (January 2002)
Document Number TC020102TB

NEED FOR SERVER CONSOLIDATION

Today's corporate computing architectures are often multi-tiered, with each tier of the architecture designed for a different type of application; for example, file-and-print servers, email servers, web proxy servers, web caching servers, and database servers. Because many applications can appropriate all memory and processor resources on a server, enterprises typically run each application on a different server to provide availability, to avoid conflicts over the available resources as the demands of those applications grow, and to provide stability in the case of a runaway process on a single application.

In comparison with mainframes, industry-standard server hardware is relatively inexpensive. It has therefore been expedient and cost-effective for companies deploying multiple database or enterprise resource planning applications to deploy new servers for each application as needed. They just add another server.

As a result of this "just add another server" mentality, companies have to manage increasingly large server networks and deal with the accompanying issues of increased complexity and management overhead. In the current business environment, applications must have access to the resources they need at all times. This often means that processing power goes unused as it sits in reserve for peak loads.

BENEFITS OF SERVER CONSOLIDATION

IT: *Information Technology*

Dual demands to reduce overhead cost and increase resource efficiency make it desirable to consolidate multiple applications onto larger servers. By consolidating applications from two or more under-utilized servers onto a single server, an IT administrator reduces cost and complexity, and uses the computing resources more efficiently.

There are several types of server consolidation. This technology brief only considers workload consolidation[1], in which applications from multiple servers are consolidated onto fewer, more powerful servers. Workload consolidation reduces the total cost of ownership for servers and IT infrastructure by:

- Reducing IT administration costs (personnel, training, travel)
- Reducing the number of software licensing fees companies must pay
- Reducing overall hardware costs

Workload consolidation can also increase performance as newer, more powerful hardware and software are deployed.

MODELS FOR WORKLOAD CONSOLIDATION

There are several ways to increase the number of applications that can run on a single server:

OS: *operating system*

- **No partitioning —** simply loading and running multiple applications on the server.
- **Hardware partitioning —** running multiple OS images on completely isolated hardware.
- **Logical partitioning —** running multiple OS images through a software layer.
- **Resource Management —** isolating specific resources for use by an application.

Partitioning and resource management can be done statically or dynamically. Static partitions can be changed only at a reset of the system; dynamic partitions can be changed at runtime. Dynamic

[1] See the white paper titled "Server Consolidation with Compaq ProLiant Servers," document number 145V-0101A-WWEN for more information.

partitioning requires more sophisticated implementations of the OS, middleware, and applications, so the software can properly de-allocate resources before giving these resources to a new entity.

Memory Leak: *Memory that is allocated by an application but never de-allocated. This can cause system degradation and OS crashes over time as multiple connections are made (and more memory allocated) to the application.*

Clearly, the first option of running multiple applications without any type of partitioning or resource management is to be avoided for reasons already explained. An OS failure due to resource conflicts in a consolidated environment can affect a great number of individuals and have an adverse effect on the productivity of an entire organization. Specifically, with multiple applications loaded onto a single OS image, the potential for *memory leaks* within those applications increases, potentially leading to fatal OS crashes. While an OS crash on a server running a single application is costly and time-consuming, an OS crash on a server that holds a procurement database, website application, and email application can bring down an entire small business.

ccNUMA*: cache-coherent non-uniform memory access*

Hardware partitioning has been available for years in the mainframe world, for example, in server architectures such as *ccNUMA*. The hardware is isolated so that the OS has no indication that other resources are available and multiple images of the OS can run on a single server. The amount of fault tolerance and reliability available from the server depends upon how well the server hardware is truly isolated. The cost of the server can increase tremendously as the amount of redundant hardware increases for fault tolerance. If a hardware partition allows dynamic changes, a vast amount of engineering development time goes into ensuring that the partitions work flawlessly and reliably.

With logical partitions, the amount of fault tolerance and reliability depends upon how well the software truly isolates the OS. Developing the software to isolate multiple instances of the OS within a server demands a large investment in engineering.

Resource management brings many of the benefits associated with a logical partition. It provides the stability needed to run multiple applications on a single server without the added capital cost of a server using hardware partitioning. Resource management also reduces the complexity of a logical partition by managing only specific resources such as memory and processors, allocating these resources to specific applications.

OVERVIEW OF COMPAQ RESOURCE PARTITIONING MANAGER

Compaq Resource Partitioning Manager (RPM) is a resource management tool that enables IT administrators to dynamically allocate, or partition, a server's memory and processor resources for specific applications. RPM is a specific solution to enable workload consolidation on *Compaq ProLiant* servers running Windows 2000 Server, Advanced Server, and DataCenter operating systems.

RPM and Microsoft Job Objects

Process*: In this context, a process can be a single executable file, an application, a driver, or a service.*

RPM makes use of the job objects feature in the Microsoft Windows 2000 operating system. The job objects feature allows bundling *processes* together and applying boundary conditions so that the entire group can be managed as a single unit. While job objects can consist of a set of processes, it is also useful to create job objects containing a single process so that stringent resource restrictions can be applied. For example, a job object could include an entire application, such as visio.exe with all its associated drivers and services, or it could include only a specific driver within that application, such as CallPOR.dll. If, for example, an administrator knows that two executable files are particularly processor-intensive, the administrator can potentially use the job objects functionality to separate, or partition, the resources needed for these two files to satisfy the compute demands of both.

In the Compaq RPM software, Resource Partitions incorporate the job object technology and enhance the job object functionality. The RPM application automates the task of creating, defining,

and activating these resource partitions, allowing them to be developed in a manner normally unavailable within the operating system.

RPM Provides Dynamic Control

Most importantly, RPM provides the ability to dynamically alter the assignment of memory and processors to a resource partition. As the needs of an application change, RPM can increase or decrease the resources allocated to that application.

For example, when Java web servers are consolidated, RPM typically improves processor utilization. Suppose a company called Acme Web Hosting typically uses 25 dedicated Java web servers to host 25 individual client websites. The load on each server varies according to time of day, seasonal variations, and other factors. Consequently, less than 50 percent of the resources are being used much of the time. By using RPM and taking advantage of its scheduling and optimization rules, Acme can achieve significant server consolidation, reduce costs, and increase performance.

On the other hand, if an administrator loads two applications onto the same 8-way server, one or both applications may try to maximize its memory resources. Under normal operating conditions, this could cause one application to become starved for memory. Using RPM to control memory resources can prevent one application from using memory resources needed by another application. Likewise, RPM's dynamic reallocation of memory resources can reduce memory leaks, which can ultimately cause a server to crash.

RPM ARCHITECTURE

As shown schematically in Figure 1, RPM consists of a graphical user interface (GUI) and the RPM Service. The RPM Service contains a rules engine that provides the rule-based parameters for dynamically altering the resource allocation. Both the GUI and the RPM Service interface directly with the Windows 2000 job objects API, the registry, and the performance counter, using standard Windows API function calls.

API: *application programming interface*

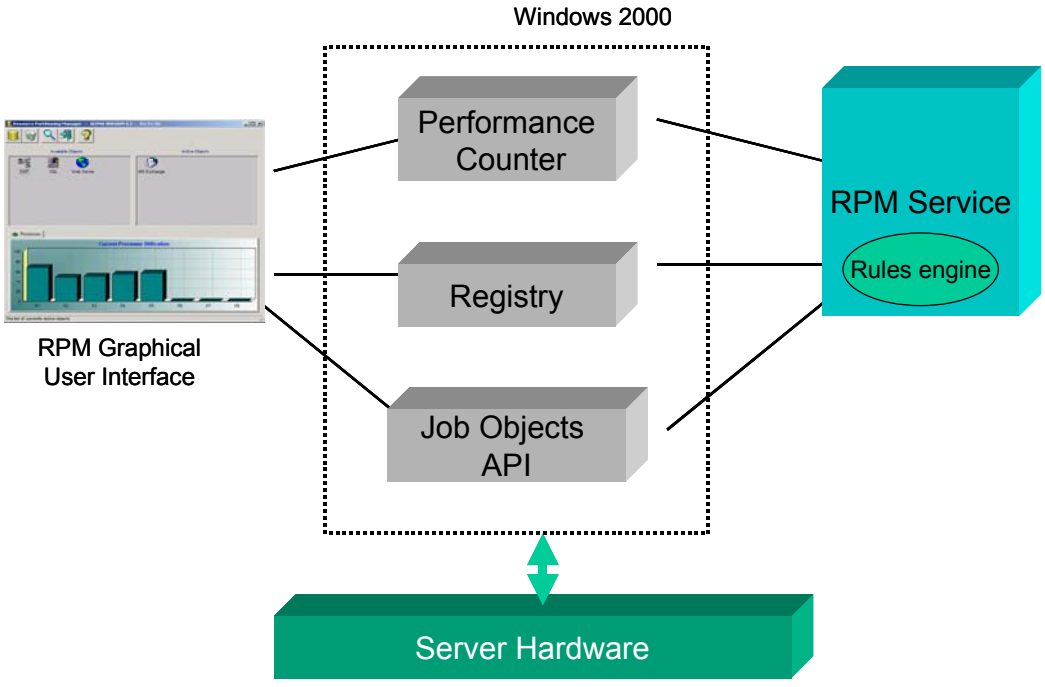

Figure 1. RPM architectural diagram

User Interface

The RPM GUI allows an administrator to easily create and activate a resource partition by defining the required processes, their associated properties, and the parameter rules in a resource partition.

The GUI consists of three major sections: an icon list of available resource partitions, an icon list of active resource partitions, and a performance chart showing current processor and memory utilization on the server (Figure 2).

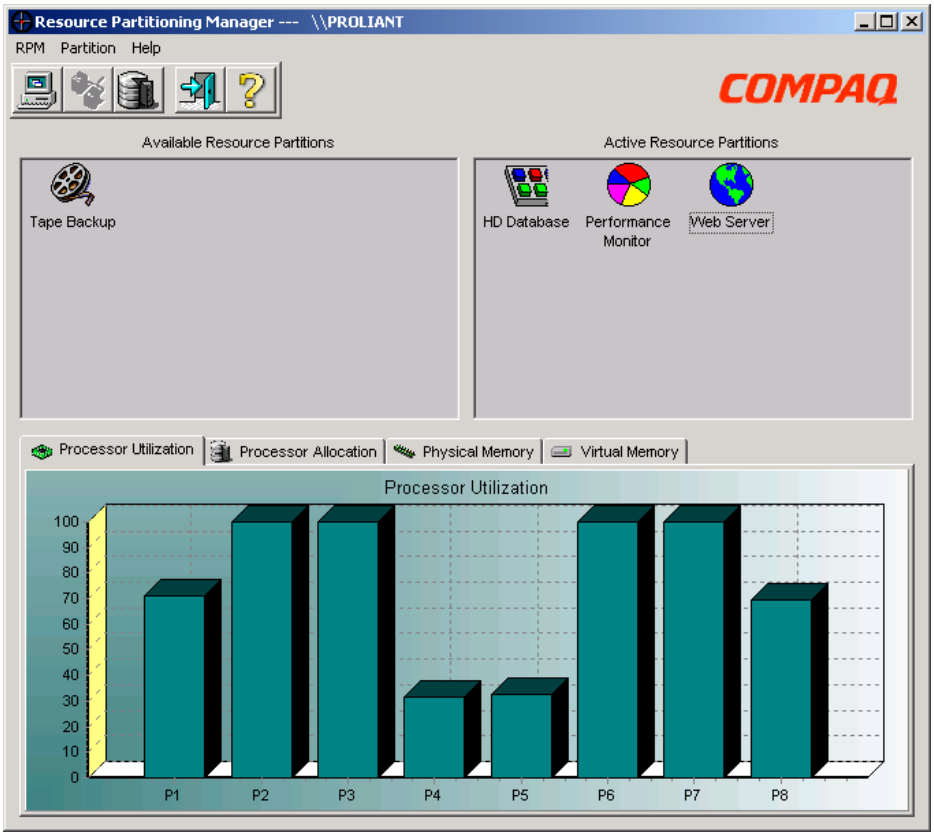

Figure 2. The RPM GUI includes windows for available resource partitions, active resource partitions, and real-time processor and memory utilization.

When creating a new resource partition, the administrator simply adds the desired processes and then assigns the properties for the desired number of processors and amount of memory. The administrator can also set a third property, Autostart. With Autostart selected, when the server is restarted, the administrator does not have to drag and drop each resource partition to activate it.

With data acquired from Windows 2000, the GUI can show the processor and memory utilization of the *currently active* resource partitions; or the GUI can show the processors and memory allocated to *all* resource partitions on the server. This enables an administrator to easily verify how resources are assigned across multiple applications.

The GUI builds and stores the description of the resource partition: the name of the resource partition, the amount of allocated memory, and the number of allocated processors.

RPM Service

The RPM Service is the heart of the RPM tool. Once a resource partition has been activated through the GUI, the RPM Service will retrieve all stored information about that resource partition. If the application is launched using RPM, the service will monitor all activity and ensure that all resources are properly assigned. If an application is launched without RPM but a resource partition containing the application has been activated, RPM will monitor the activity and assign resources as required. If the resource partition containing the application is not active when the application is launched, then the application will run normally but RPM will be ignored.

The RPM Service is active on the server continuously. Even so, the load on the server from running the RPM service is minimal and the performance overhead will typically be minimal, especially when compared to the potential benefits.

Rules Engine

The rules engine within the service provides the mechanism for dynamically altering allocation of processor and memory resources. If external server conditions change (such as increasing hits to a web server or a large and complex data query on a database) and a process requires more resources, the rules engine will automatically execute the rules defined when the resource partition was created. For example, an administrator could set up a rule to increase the number of processors allocated to a resource partition up to a maximum of eight, when all currently assigned processors reach 70 percent utilization for more than 30 seconds (Figure 3). Likewise, the administrator could set up a rule to make processors assigned to a resource partition available to other partitions during periods of low usage. These rules automate the way a server is able to dynamically manage changing needs at critical moments, thereby enhancing availability and performance for the server.

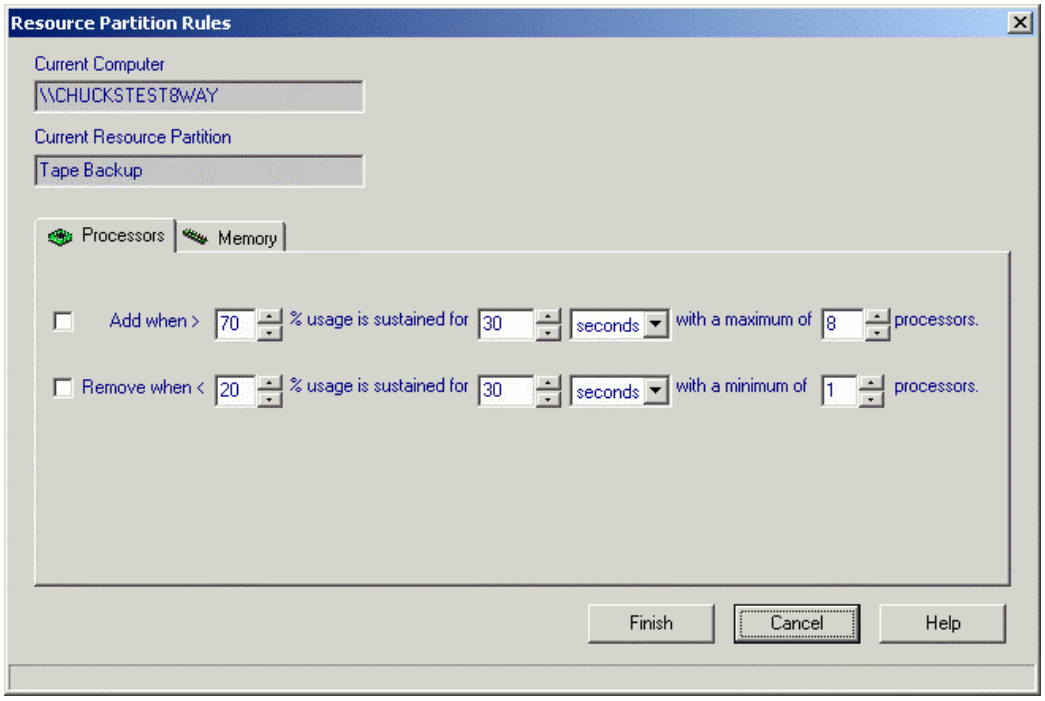

Figure 3. The RPM GUI leads the administrator through choices to establish rules for the processor and memory allocation.

Because of the dynamic rules engine, the RPM Service automatically begins to apply the properties and resource rules to the application, modifying resource allocation as needed. The RPM Service works with the application to gracefully release resources, if that is required. The newly created or edited resource partition information is saved so that the information can be retrieved anytime the application uses these same resource partition rules.

Memory Allocation

Using the job object functionality within Windows 2000, RPM enables IT administrators to control the amount of physical and virtual memory allocated to a resource partition and also to each process within the resource partition. The total amount of virtual memory that can be allocated to each resource partition is determined by the maximum amount of virtual memory configured on the system.

At the resource partition level, RPM limits the maximum amount of virtual memory that can be allocated to a resource partition. At the process level, RPM enables an administrator to set allocation limits on the working set per process (the physical memory) and on the virtual memory per process within the partition. These controls allow an administrator to tailor memory use to match specific application needs.

Because both physical and virtual memory can be allocated *per individual process* within the resource partition, an administrator must understand the amount of resources each process typically requires before setting the limit for virtual memory. The administrator must also know how many total processes the resource partition contains to allocate sufficient resources for each.

Processor Allocation

RPM automatically recognizes how many processors are configured on the server. For example, if a customer installs RPM on a four-processor system, the GUI will automatically show only four processors available. RPM is designed for expandability and can support more than eight processors if future server designs incorporate more processors.

APPLICATION LIMITATIONS

Some applications, such as Microsoft SQL Server, already contain resource management functionality similar to that found in Compaq Resource Partitioning Manager. With such an application, the built-in resource management functionality should be used to control resources for that application. However, users can deploy other applications on the same server using RPM.

If an application has limitations on the number of instances that can be run on a single OS instance, RPM cannot overcome those limits imposed by the application vendor. For instance, Microsoft IIS, Terminal Server, and Exchange are three applications that cannot run multiple instances on top of a single OS instance. However, using Compaq RPM, a user can run a single instance each of Microsoft IIS, Exchange, and Terminal Server simultaneously on a server and ensure that each application gets the resources needed to execute efficiently.

POTENTIAL FUTURE ENHANCEMENTS

Compaq RPM is an excellent tool for increasing the robustness and fault-tolerance of Microsoft Windows 2000 platforms. In the foreseeable future, Compaq plans to enhance this industry-leading application by integrating it with complementary tools. Tools with which RPM could potentially be integrated include remote access and management software, performance monitoring, and multiple platform (multi-node) support for RPM.

CONCLUSION

Compaq Resource Partitioning Manager simplifies management of complex multiprocessor server environments by simplifying workload consolidation, increasing application stability, and improving resource utilization. RPM provides end-user control of server resources in a Microsoft Windows 2000 environment. This product complements such management building blocks as Compaq Insight Manager XE and Compaq services that target capacity planning and performance monitoring. RPM can be an essential component in optimizing application performance and consolidating workloads onto 8-way servers, particularly those running middle-tier applications, databases and Exchange/Domino and e-commerce environments.

For further information about *Compaq Workload Management Pack*, featuring Resource Partitioning Manager, visit http://www.compaq.com/manage/rpm.html.