

Less Is More

Data Center Automation: Discovering Your True Computing Potential

This article first appeared in the December 2008 issue of *Novell Connection* magazine.

In the November issue of *Novell Connection*, we discussed the goal of data center automation and outlined our three-pronged strategy for achieving it: Discover, Create and Manage/Orchestrate. This month, we'll examine the Discovery phase in more detail.

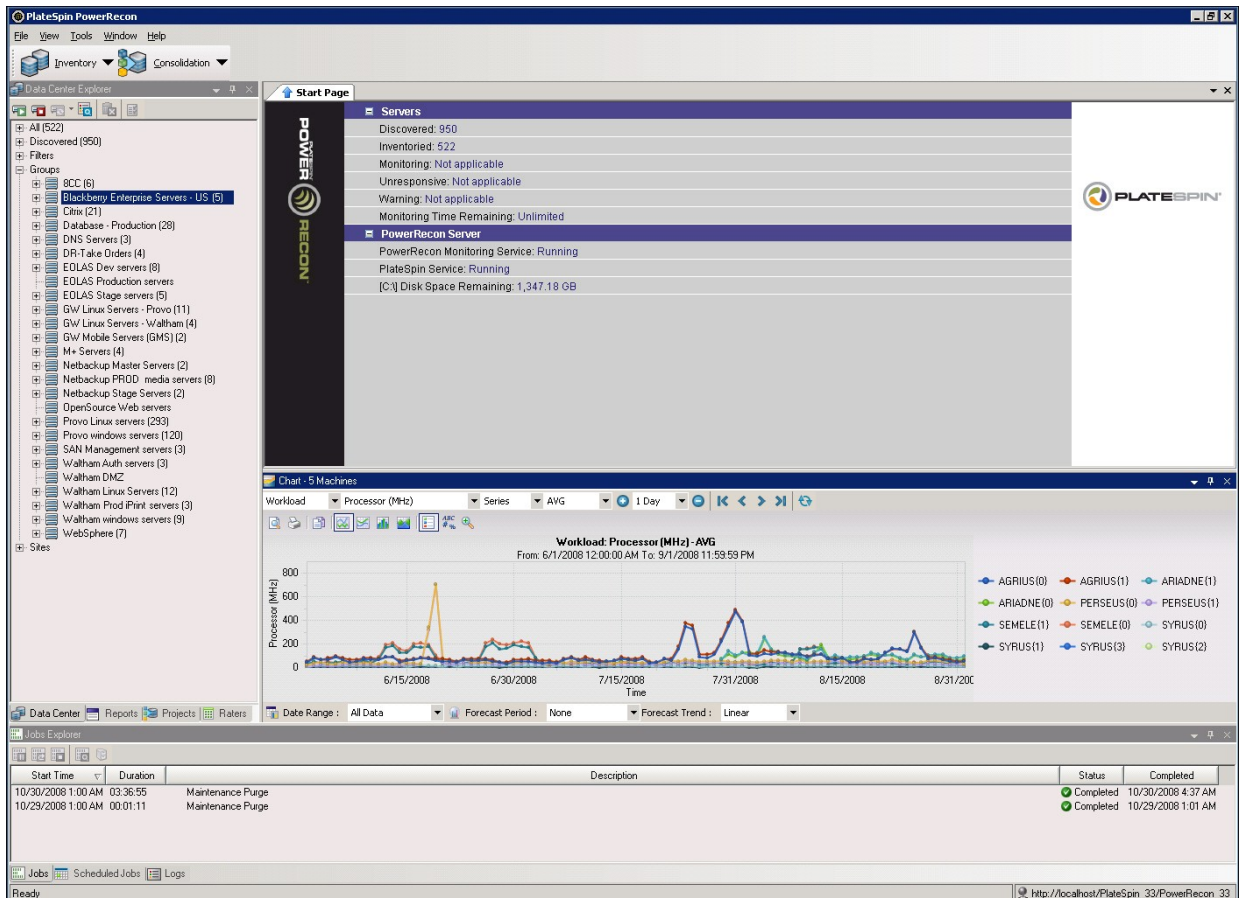
> Knowing What You Have and How It Behaves

Before you can virtualize effectively, you need to know what you have. That includes the operating systems, applications and services that may be candidates for virtualization, as well as the hardware that may host virtualized environments. Most enterprise IT shops already have discovery technologies they use for inventorying assets and managing licenses. However, identifying good candidates for virtualization requires much more fine-grained detail about assets and how they *behave*—as well as the ability to compare those behaviors.

In other words, you need to know not just what CPUs, memory, storage and resources you have available, but also how specific workloads are utilizing these resources over time, as well as when and why utilization spikes occur. You also need the ability to compare these utilization trends between different applications and platforms to determine the best opportunities for virtualizing systems and the best physical machines to run the virtualized systems.

This detailed, comparative trend analysis can help you assign virtual machines to physical machines based on their average and spike resource needs—choosing virtual processes that will “play together well” in a given physical environment. It can also help you decide which processes require maximum resource availability and are best left to run in their own dedicated physical environment.

Figure 1: *PowerRecon lets you discover, compare and report workloads in a variety of ways, such as looking at utilization across all our Blackberry Enterprise Servers.*



> The Right Tool for the Job

At Novell, we use PlateSpin PowerRecon to provide complete and precise details on our available assets and how they are being utilized—and to graph utilization across workloads and machines for intelligent analysis. (See Figure 1.)

PowerRecon automatically discovers hardware, software and services inventory for all workloads, across Windows, Linux and UNIX platforms. And it collects workload utilization statistics to provide a detailed profile of how each server's CPU, memory, storage and I/O resources are impacted over days, weeks or months. PowerRecon can even report on power and cooling usage—resources that are often overlooked although they contribute substantially to overall IT spend.

Workload Management and PowerRecon

The philosophy behind PowerRecon is that server workloads represent the base unit of business value in the data center. A workload is defined as all the data, applications and operating systems that reside on a host—whether the host is physical or virtual. The ability to monitor, assess, manage and move server workloads at will between physical and virtual hosts is the key to data center automation.

Workload profiles can be compared and analyzed in PowerRecon, using scenario modeling, forecasting and planning tools that allow you to create “what-if” scenarios to test the effects of various combinations of hardware and virtual hosts. You can customize the graphical presentation of workload profiles, comparative scenarios and growth forecasts to make analysis practically intuitive. For example, you can:

- Sort workloads by CPU, memory, disk, network, inventory and/or performance to compare power, cooling, space, cost and utilization.

- Identify mismatches between workloads and computing resources at any moment or over time.

- Create customized management reports and automatically deliver them to IT optimization specialists and service providers on a set schedule.

- Calculate workload costs by business unit, based on actual resource usage for both physical and virtual machines, for accurate internal chargeback.

And that's just a sampling of the ways you can use PowerRecon for a more power-, resource- and cost-efficient automated data center. At Novell, we're always thinking about new ways to put PowerRecon to work in our own data centers. Here are just a few of the techniques we've used and the results we've achieved so far.

> Agentless Installation, Discovery and Monitoring

For years, we've been accustomed to installing agents on servers to gather information and perform management

tasks. For example, a typical server might have an antivirus agent, a backup agent, an HP System Insight Manager agent, and a ZENworks Configuration Manager agent. Sound familiar? Agents are often necessary to get a job done. But when you can avoid installing yet another agent, you avoid both the installation work and the extra load that the agent puts on system resources.

The discovery method used by PowerRecon is agentless. At Novell, we simply pointed the PowerRecon discovery tool at our subnets, provided the necessary credentials and then sat back and watched as PowerRecon began discovering and monitoring our servers. The process was amazingly simple; it worked perfectly the first time; and in almost no time we were getting detailed, longitudinal, comparative information we could use to make better operational decisions.

> Server Consolidation and Cost Savings

Once we had it up and running, we immediately began to see PowerRecon provide value in our analysis of our Citrix server farm. This was an older part of our server infrastructure, and we had never before had a clear grasp of resource utilization there. Citrix had been running on 24 physical servers at an annual cost of approximately \$8,700 for power and cooling.

Through our analysis of the utilization data and trends provided by PowerRecon, we concluded that we could consolidate the Citrix workloads on just five servers with no significant performance penalty. By removing 19 unneeded servers, we conservatively calculated the savings in power and cooling to be \$6,700 per year.

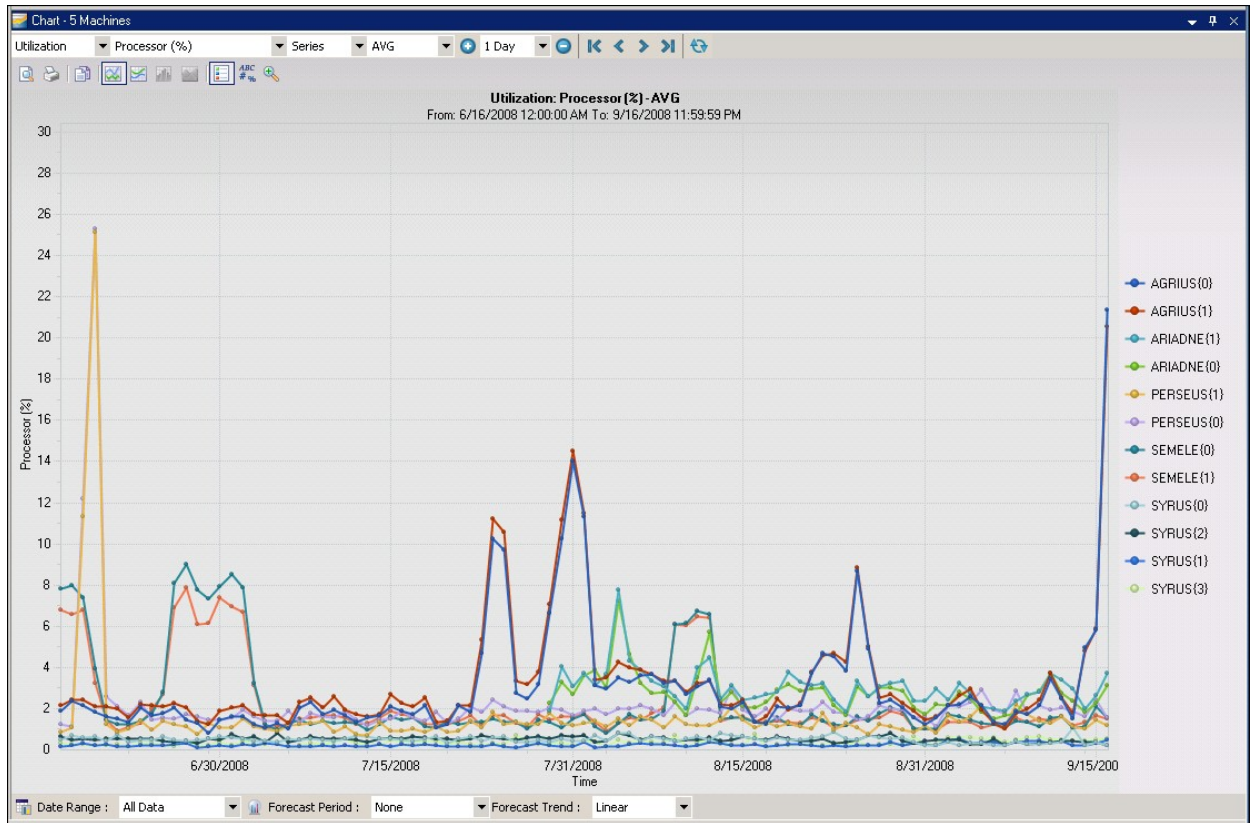
> A More Efficient Development Environment

In developing and staging new software releases, we have traditionally relied on physical servers configured to mirror our production environment. But through analysis of PowerRecon statistics, we learned that these development and staging servers are vastly underutilized compared to our production servers. Since we can replicate the configuration and behavior of physical servers in virtual machines, these development and staging servers became perfect candidates for virtualization—saving on operating expenses and future hardware purchases while freeing up resources for other tasks.

> Doing More with Less

We've had such great success with PowerRecon that we're now evaluating almost every service we run in our data centers, gathering detailed information about how effectively we're using our computing resources to host a wide range of ever-evolving workloads. We're using PowerRecon to discover the timing and cause of utilization spikes, as well as to evaluate utilization trends over time. For most services, we're finding CPU utilization consistently under 20 percent. We're also learning how memory, storage, I/O bandwidth and even power and

Figure 2: This view shows CPU utilization for 12 Blackberry servers over a three-month period. Some servers show obvious spikes, which we can trace to business requirements during the period. But the biggest spike is only about 25 percent, and many servers consistently show very low utilization—making those workloads perfect candidates for virtualization.



cooling resources are being utilized by the wide range of server applications we run in our data centers.

Power and Cooling Costs in Today's Data Center

From 1999 to 2007, according to *Electronics Cooling*, server performance increased 7,500 percent for the same hardware cost. During the same period, performance per watt increased 1,600 percent. That's good news overall, but it means that power and cooling costs are actually beginning to outstrip hardware costs, especially as enterprises rely on more demanding business applications supported by fault-tolerant infrastructure.

Assuming a relatively efficient cooling infrastructure, the three-year power and cooling costs for a typical 1U server in the U.S. essentially equals the hardware cost of the server itself. By 2014, infrastructure and energy will represent 75 percent of costs compared to only 25 percent for IT. In other words, "a paradigm shift has occurred in the data center. The cost of IT equipment is no longer the bulk of the cost, but rather the cost of the power and cooling infrastructure has crept in to be the primary cost driver."

(Source: Christian Belady, Hewlett-Packard, "In the data

center, power and cooling costs more than the IT equipment it supports," *Electronics Cooling*, February 2007.)

We expected to find underutilization. But what's really eye-opening is how utilization patterns vary according to the workload and the user base. From Blackberry Enterprise Servers, to mail servers, to database servers and other types, we're finding that each has its own unique utilization profiles and operating behaviors. (See Figure 2.) As a result of this knowledge, we're planning ways to cluster servers and consolidate workloads in ways we never would have considered—or even dared—to do before. By truly understanding the technical opportunities, we can do more with less to create new business opportunities.

PlateSpin PowerRecon has become an essential tool for us as we move forward with sizing new projects, refreshing old hardware and virtualizing computing environments to provide new efficiencies and cost-savings. Next month, we'll discuss how we create virtual machines based on what we've learned from PowerRecon. See you next month in the (virtual) pages of *Novell Connection*. **N**