Mary Johnston Turner
Research Vice President, Enterprise System Management Software

Effective Management of Heterogeneous Datacenters and Multihypervisor Environments

April 2011

Many IT organizations run heterogeneous physical IT infrastructure environments as a result of corporate mergers, acquisitions, and application-specific requirements, as well as to mitigate concerns about single-vendor lock-in. For the same reasons, more and more organizations are testing and deploying multiple hypervisor solutions. Still others are designing and implementing cloud strategies that look to automate and orchestrate workload provisioning and migrate across heterogeneous environments to increase datacenter flexibility and productivity. Operationally, these heterogeneous environments create a number of management challenges for IT organizations. Management tools that can look across diverse physical and virtual platforms and applications are needed to help IT organizations optimize resource performance, seamlessly provision and migrate workloads, plan for future capacity, and quickly troubleshoot problems.

The following questions were posed by Dell to Mary Johnston Turner, research vice president for IDC's Enterprise System Management Software practice, on behalf of Dell's customers.

Q. What are the biggest challenges currently facing IT teams with respect to managing heterogeneous computing environments?

A. Heterogeneous datacenter management and operations are most effective and efficient when processes are standardized and automated and administrators can work with a consistent, accurate, and timely set of information about application and infrastructure performance, availability, and root cause, as well as efficiently coordinate workflows and hand off responsibilities from one group to another as needed. Integration and automation are particularly important for organizations that are considering private cloud strategies to dynamically assign workloads and optimize datacenter resource utilization.

In heterogeneous environments, each vendor's server, storage, network, or hypervisor platform is likely to have its own set of provisioning tools and performance monitoring technologies. Although many will provide open APIs that allow third-party tools to collect and analyze the data generated at the system level, it can often be quite time consuming and expensive to deploy and integrate a large number of adapters and agents to provide IT with an accurate, real-time view of the end-to-end environment. Requests to provision new resources can be delayed because work requests have to be manually routed and approved between different technology groups. We find that coordination and communication across different technology groups can often be inconsistent and ad hoc.

To maintain service-level agreements (SLAs) and keep staffing costs in line, IT teams need to optimize the end-to-end performance of business applications across complex heterogeneous environments. Advanced management and monitoring tools are used to

provide actionable performance and availability insight, support complex root cause analysis, and integrate workflows across teams and technology silos. Automated provisioning, change control, and self-service technologies that permit consistent, rapid resource allocation and ongoing optimization of resource consumption are also required to deploy and manage private cloud environments. They are used to coordinate and speed up end-to-end provisioning of the full set of application, network, storage, and server resources needed to deliver required business services.

Q. How does increasing use of virtualization impact IT's ability to deliver reliable, timely, and cost-effective services to business users?

A. Virtualization adds another layer of complexity to heterogeneous datacenter environments. The promise of virtualization is that it will allow IT organizations to improve server utilization, consolidate workloads into smaller datacenter footprints, and improve the efficiency of IT operations. To date, many organizations have realized measurable benefits from their virtualization initiatives when it comes to optimizing hardware costs and datacenter facilities requirements. However, many have failed to see significant improvements in terms of IT staff efficiency and productivity.

Many organizations report that the current ratio of IT administrators to virtual servers is not much better than their traditional ratio of one IT administrator to 30 or 35 physical servers. Despite the fact that administrators can provision virtual machines in minutes, many organizations find it can still take weeks to fully integrate, test, and deploy the full set of server, storage, network, and application resources required to support the business needs of the end user. In the cloud computing era, where public cloud services can be provisioned almost instantly with the swipe of a credit card, IT teams are experiencing pressure from business groups that want to be able to self-provision some IT resources or at least want IT teams to respond more quickly when the business requests resources.

From the IT perspective, provisioning the resources is only half the battle because it can be very difficult to track the assignment and monitor the use of virtual machines (VMs) once they have been deployed. Many organizations struggle with VM sprawl, unable to reclaim assigned resources. They are frequently unable to validate whether assigned VMs are actually needed by the business or decide if those VM resources should be reassigned to other workloads. This situation results in inefficient use of resources and drives unnecessary purchases of additional physical servers. It also forces IT staff to continue to support VMs and workloads that are no longer being actively used.

Q. Does the use of hypervisors from multiple vendors introduce any additional concerns?

A. In addition to virtualization management challenges, multihypervisor environments experience the same issues that affect heterogeneous physical environments, Each vendor's hypervisor will have its own set of APIs, performance monitors, and VM provisioning and migration technologies that need to be integrated into consistent, standardized workflows and automated provisioning profiles. This added layer of complexity makes it even more challenging for IT teams to operate effectively.

Some organizations might attempt to deal with these challenges by assigning different staff to become experts on each vendor's hypervisor and associated hypervisor-specific management tools and APIs. The danger of this approach is that organizations could end up with even more operational silos and more inefficient ad hoc workflows that can slow their ability to deliver resources to the business. This type of approach can also make it more difficult to use IT staff efficiently as it becomes difficult for single-vendor experts to multitask and address work requests on different platforms.

Alternatively, organizations may opt to train staff on multiple hypervisors and their related management tools. While this approach may be more efficient in the longer term, it can result in extended learning curves and delayed time to value as IT staff switch back and forth across different tools and hypervisor environments.

Q. What types of system management solutions are needed to effectively operate these types of complex virtualized and private cloud environments?

A. The most efficient system management solutions for complex, heterogeneous physical and virtual environments will have many of the following attributes:

- Scalable architecture that allows standards-based integration with a wide array of platforms and systems
- Out-of-the-box integration with leading server, storage, network, and hypervisor environments
- Extensive automation and orchestration abilities including out-of-the-box templates and best practices to enable IT teams to rapidly improve basic operational activities while providing tools to develop and maintain customer-specific workflows and gold images
- Enabled self-service provisioning based on a well-defined, standardized service catalog and SLAs
- Ability to monitor end-to-end performance of applications and systems to provide IT with accurate, real-time information about business impact and root cause
- Ability to monitor and measure resource consumption and support consumption-based capacity planning for real-time workload performance optimization as well as longer-term capital spending

Organizations that implement a complete set of automation, orchestration, and performance management and monitoring solutions not only will be able to significantly improve service levels and IT staff productivity but also will be well positioned to share resources dynamically across many different user groups and workloads, in effect creating a private internal IT cloud.

Q. How can the types of management tools listed above improve the operation of multihypervisor environments?

A. To operate multihypervisor environments efficiently, IT organizations need to use a consistent set of processes, policies, and tools across the environment. If IT organizations can integrate automation, self-service provisioning, end-to-end application monitoring, and consumption-based capacity planning activities using a consistent, integrated set of management solutions, they should be able to streamline operations, reduce complexity, and improve service levels.

By linking information about application service levels to automated workload balancing and VM provisioning tools, IT organizations will be able to better optimize the consumption of IT resources, thereby reducing costs and better matching IT costs to business needs.

As physical and virtual datacenter resources become more diverse and private cloud architectures are introduced, IT organizations will need to take full advantage of these types of integrated and automated management tools to hold down costs while responding more quickly to dynamic business requirements.

ABOUT THIS ANALYST

Mary Johnston Turner, research vice president, Enterprise System Management Software, contributes to IDC's coverage of the system management software industry. Her major areas of interest include systems operations, application service management, datacenter automation, asset management, and management software delivered as a service. Mary's coverage includes market forecasts, competitive assessments, strategic vendor analyses, and delivery of custom consulting projects.

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