

Virtual Machines You Have to See Them to Manage Them

By Dan Kusnetzky, Principal Analyst

OVERVIEW

Virtualization has tremendous advantages over physical environments, but in the move to virtual, the organization often lost some things it took for granted. One of those things is clarity of what is going on, where, and with whom.

WHAT'S WHAT?

In a physical environment, such as a typical datacenter, all of the systems, network devices, storage devices and even power supplies and air conditioning units are carefully organized and labeled with their IP addresses, with their network names, and often the application it is supporting.



The midrange systems from Sun or Hewlett-Packard are over in the corner running manufacturing or billing applications, the mainframe that sitting in the center of the datacenter is a whiz at processing thousands upon thousands of transactions, and a racks of industry standard (x86) machines are supporting Web-based applications and the organization's collaborative systems.

Finding each of these systems in the datacenter is relatively easy. Traditionally, it has also been easy to learn what each of these systems is doing at any given moment. Quite often datacenter resources are grouped according to the application they're supporting.

As organizations began moving functions from physical systems to virtual systems they ran straight into a new problem. Virtual machines are ephemeral. They can be generated, provisioned with the appropriate software and put into production very rapidly.

They can be halted and deleted when they no longer are needed. No labels define their location or presence.

IT executives are increasingly facing the fact that there really is no good way of telling what virtualized applications are running, where they are located at any given moment, which business unit owns them, when they were created, when they should expire, which physical resources they are using and whole host of other questions. It is no longer easy to determine either what physical systems are doing or, more importantly, not doing.

WHAT'S WHERE?

Technology, such as VMware's Vmotion, XenSource's Xenmotion, and many orchestration products, allow organizations to move virtual systems about at will. Knowing where everything is on a moment by moment basis can be quite challenging.

If an organization needs to maintain an audit trail of where the computing was done, where the data is and other important data points, they increasingly face a very difficult challenge. Complying with some regulations may be impossible unless there is a structured, well-defined way to track everything. Organizations simply must have a trail of their motion, a history, or a “chain of custody” as they transfer from place to place and from test beds to actual production.

W H A T ’ S O L D

Since virtual resources may be created easily, live only as long as needed and then be easily retired, the virtual machine’s “born-on” date and when they are to be terminated is now important information that IT administrators must track. Since this is a new function, traditional products and processes are not likely to be of much help. Will administrators fall back on privately maintained spreadsheets? Will this approach pass muster when an auditor comes? How easily will that scale when there are 5 or even 20 administrators in different facilities are able to create virtual machines?

It is clear that a better way to manage this type of information is needed. Otherwise, organizations will find that they’ve created a virtual machine sprawl that will add additional challenges rather than reducing or removing them

T H E G O L D E N T E M P L A T E C O N C E P T

In order to create some order out of this potential chaos, many organizations have begun creating “golden template” virtual machine files. Each of these templates acts as the foundation for many other virtual machines. Golden templates are created using all of the necessary infrastructure software needed by applications. While this approach has some benefits, managing these elusive golden templates can also be a challenge.

Many organizations opt to use clones, rather than templates, to generate virtual machines, While the solution is different, the same questions remain: those of identification and virtual lineage of each virtual machine in the environment.

Organizations need to know the following information about each and every golden template or clone, every update to that template or clone, and, in turn, all of its children:

- Creation date
- Operating system type
 - Operating system version
 - Operating system patch level
- Type of data management software
 - Data management software version
 - Data management software patch level
- Type of application framework software
 - Application framework version
 - Application framework patch level
- Applications
 - Application versions
 - Application patch levels
- Known incompatibilities
- Planned deletion date
- What children have been created based upon each template?

- Who are the sibling clones in the environment?
- Are there grandchildren as well?

It's exhausting to think about gathering all of this information and keeping it up to date.

W H Y T R A D I T I O N A L A P P R O A C H E S D O N ' T W O R K

The simple answer to this question is that virtual resources are different than physical resources in several ways. These differences mean that traditional approaches often don't work in a virtual world.

Most organizations don't have standard processes to collect and maintain lifecycle information about virtualized resources. Much of this critical information doesn't really fit into their usual tracking systems.

For example, one simply doesn't install a physical system one morning and retire it that afternoon. Virtual systems are often created, used and then destroyed in a day's time.

Furthermore, it is unlikely that a physical system would be picked up and moved from one side of the datacenter to another several times a day. In a virtual world, this is quite possible. There are many vendors offering orchestration tools that can move virtual systems from one physical machine to another or change it from being a virtual resource to a physical resource based upon policies or to achieve service level objectives.

Not only is the fundamental data model of traditional tools flawed in these ways, the architecture is equally inappropriate. Fat agents in each virtual machine could be the end of any performance gains made by virtualization. And log files are often not sufficient for generating real-time insight into the environment. The challenge of monitoring virtual environments has many new dimensions of complexity.

It is not clear who would be responsible for virtual resources, as well. Would the system administrators who operate physical systems be also be responsible for the virtual resources that machine might be supporting? Would that job fall to the person in one of the business units who created and is using that virtual system? In the end, without help, administrators really have no reliable way to know what's running, where it's running, and what compatibility issues it may generate.

S U M M A R Y

As organizations increasingly turn to virtual machine technology for both client and server environments, it would be wise to consider the challenges this technology imposes as well as the benefits they may bring.

The Kusnetzky Group recommends that organizations develop a sound, well planned set of processes and procedures to manage these resources. It would also be wise to seek out tools that can automate this process.