

Realtime
publishers

The Shortcut Guide[™] To



Virtualization and Service Automation

sponsored by



Greg Shields

Chapter 3: Automating Change Management in Virtualized Environments	35
The Sources of Change Requests.....	37
User-Based Changes	37
Team-Based and Project-Based Changes	38
Enterprise-Based Changes	38
Optimizing Resource Use to Meet Business Demand	38
Service Asset & Configuration Management	40
Capacity Management	40
Availability Management.....	41
License Management	42
Short-Term Service Planning.....	42
Long-Term Service Planning.....	43
Automating Resource Provisioning	44
Receiving and Evaluating a Request for Change.....	46
Collecting, Collating, and Prioritizing Change Management Requests	47
Identifying Resources	47
Architecting the Solution	48
Change Management & Release Procedures	48
Technology Architectures for Enabling Managed Self-Service	48
Implementing and Automating Self-Service	49
Governing Self-Service Activities	50
Virtualization Technologies Enable New Automation in Change Management.....	50

Copyright Statement

© 2008 Realtime Publishers, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers, Inc. (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers, Inc or its web site sponsors. In no event shall Realtime Publishers, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library. All leading technology guides from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

Chapter 3: Automating Change Management in Virtualized Environments

Virtualization today is an absolute conversation starter. Organizations both big and small recognize it as a major game changer, forcing IT professionals to shift the ways in which we think about managing our environments. But virtualization arrives as a business enabler as well, realizing that promise because it delivers on value. In an IT ecosystem where technologies arise every day that hint towards solving business problems, virtualization is unique in its ability to rise above the hype cycle and truly ease the processing of business.

What's particularly interesting about virtualization's play is its potential for penetration across the broad spectrum of environment shapes and sizes. Whereas many game-changing technologies don't show their full value until they're incorporated at large scales, virtualization's technologies can assist the small business as much as the enterprise. Although the small business will recognize a different facet of value than the enterprise, part of virtualization's value is in how it touches almost every part of the change management needs of business:

- *Cost savings.* With advancements in technology constantly impacting the speed of new hardware, modern data processing has grown from a focus on demand limits to one of virtually unlimited supply. The average non-virtualized computer system today runs useful work around 5% of the time. That means that 95% of its available process cycles are unnecessarily wasted. Consolidating physical workloads atop a virtual platform enables IT to more efficiently use computer resources. The end result is an overall reduction in need for physical hardware, reducing data center footprints and lowering the cost of doing business.
- *Availability and resiliency.* Virtual workloads are by nature more resilient than those installed directly to physical hardware. This resiliency is a function of their decoupling from physical hardware. When data processing can be abstracted from physical hardware, the loss of any piece of hardware no longer results in a loss of service. Smart businesses see the value of improved availability and its direct impact on the bottom line.
- *Automation.* Virtualization's impact on process completion happens as a function of its capabilities for automation. The discussion in this chapter is written to give you an understanding of what those automation functions are and where they can be implemented to enhance the needs of change management.
- *Process alignment.* Taking a new service from idea to implementation is a difficult task. And getting the hardware resources necessary to make it a reality can add costs to the point where new and needed services may not be fiscally viable. Virtualization's rapid prototyping and rapid deployment capabilities—especially when managed through service solutions as described in Chapter 2—reduce barriers to entry for service development and ease the steps required for effective change control. In the end, virtualization's alignment of technology with process brings levels of agility never before seen in IT.

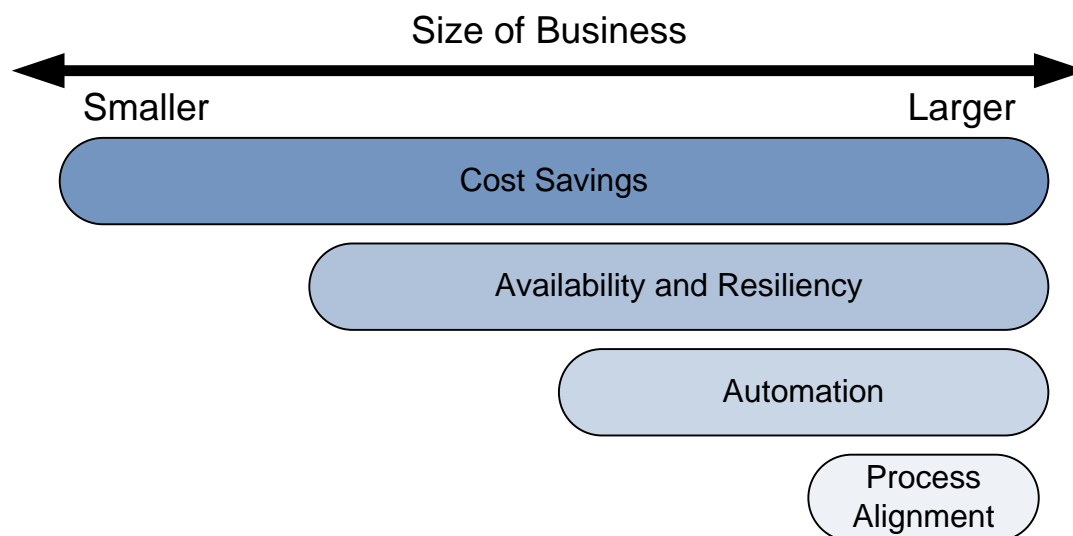



Figure 3.1: Business in the various stages of growth place priority on different parts of virtualization's value proposition. Although all businesses benefit, larger and more mature businesses have the capacity to see more and different areas of value.

Of particular interest within all of these is how they can be realized by essentially any size business. As Figure 3.1 shows, the small business tends to focus exclusively on cost savings and availability, while more mature enterprises look towards process alignment and automation with the intrinsic assumption of the others. No matter your business size, virtualization and the automation capabilities that arise with it have the potential for impacting how you do business.

 Chapter 1 of this guide discussed three tenets of virtualization that are worth repeating here:

Virtualization's tenet of flexibility eliminates the limitations of IT's old and difficult-to-change technology underpinnings, allowing IT to consistently develop service solutions required by the business with a minimum of cost.

Virtualization's tenet of transience means that IT organizations can rapidly and easily reconfigure the IT environment as the needs of business changes over time.

Virtualization's tenet of commonality enables the rapid development of new services as well as the assured reliability of existing ones. By leveraging standardized configurations across technologies, IT reduces its potential for error while speeding the rate in which technology underpinnings can be brought online.

When considering the specifics to be discussed later in regards to virtualization's change management impacts, keep these three points in mind. You'll find that these three concepts of flexibility, transience, and commonality are the backbone of virtualization's impact to IT operations.

With this in mind, we turn towards the specific benefits virtualization provides to the fulfillment of change management needs in an organization. It is obvious to state that different organizations will have different ways in which they fulfill their change management needs. That fulfillment, as will be discussed shortly, tends to be tied to the process maturity of the IT organization itself.

In this chapter, we will explore the details of traditional change management activities with a look towards the impact from virtualization. That discussion starts with a review of the sources of change requests and then moves on to the areas in which those change requests can be fulfilled. You'll discover that the right-sizing of resource use and the automated deployment of assets gains great benefits through virtualization's automation functions.

The Sources of Change Requests

As an IT Service Management discipline, Change Management is concerned with the assurance that correct and approved procedures are used in the handling of change requests. This control over the process by which changes are made ensures that they are researched and completed appropriately, that the proper approvals are obtained, and that metrics are established that validate the change and its efficacy. Changes in an IT environment come from all different directions at once, which is why the formalization of change management procedures is critical for moving an IT environment from a reactive mode to one of proactive management and measured control.

The typical IT environment tends to see changes across a spectrum of inputs. Individual users need updates to configurations and teams and projects require additional resources or the adjustment of existing resources. Change can even come from the enterprise itself, mandating new ways and technologies to accomplish tasks. From the perspective of virtualization, let's look at each of these in turn.

User-Based Changes

The changes often required by individual users are usually small in scope but highly specific. For example, a user may identify that a server setting is incorrect to meet their needs, and as such, requires an update to that configuration. Although virtualization itself deals more with the "chassis" functions of a virtual machine—for example, power off, power on, migrate, snapshot, etc.—an individual user's needs can impact the virtual environment.

As an example of this, a user may discover that he or she requires console access to a server instance. In the physical environment, this need often requires IT to grant data center access to this quasi-trusted individual. With virtualization, all console access is remote, which means the user can gain the necessary console access while preserving the security of the data center.

Team-Based and Project-Based Changes

Teams and projects tend to drive change from a resource provisioning perspective. These groups identify that they require hardware resources and undergo portions of the configuration management process to specify and request the resources they need. Virtualization abstracts and aggregates physical resources such as processing power and memory into what can be called “pools.” From these pools, needed resources can be drawn.

The assignment of resources in and out of pools has a high level of resolution, which means that very discrete quantities can be withdrawn to meet exact team and project demands. This is in comparison with physical resources, which are dependent on their hardware configuration.

Enterprise-Based Changes

Enterprises themselves are often drivers of change. This is particularly the case when businesses change and evolve over time. Enterprise-based changes tend to be high level in nature, requiring the reconfiguration of vast swathes of resources to meet a new business requirement.

Virtualization and the technologies that enable it include automation capabilities such as scripting exposure and automated actions. Virtual machines themselves are also by definition highly similar in nature, meaning that the core of one virtual machine often resembles another. This similarity allows virtualized resources to be reconfigured as a unit rather than one at a time. The end result is that enterprise-based changes can be fulfilled much faster and with higher reliability when necessary.

Understanding where these changes come from, the next sections of this chapter will discuss how each class of change is modified through the introduction of virtualization. To begin, one major result of this addition is through a greater optimization of resource use.

Optimizing Resource Use to Meet Business Demand

In its IT maturity model, Gartner identifies five levels of organizational maturity. As Figure 3.2 shows, that spread starts at the bottom with highly immature organizations, those that Gartner refers to as Chaotic. In these organizations, the processes surrounding IT’s technology investments are ad-hoc and undocumented at best. As organizations recognize the need for process maturity, Gartner sees them elevate through the Reactive, Proactive, and Service stages before ultimately achieving the Value stage. When an IT organization has achieved the Value stage of maturity, it is said to have a high linkage between business metrics and IT capabilities.

As organizations move from one end to the other, a major determinant is in IT’s ability to meet the demands of business. When IT processes are immature, IT is incapable of planning for business and service expansion over time. As IT increases in process maturity, it gains the abilities to more effectively monitor services, analyze trends, and predict problems. IT organizations who effectively meet business demand have quantitative service quality goals in place, are able to measure to those goals, and can plan and budget appropriately.

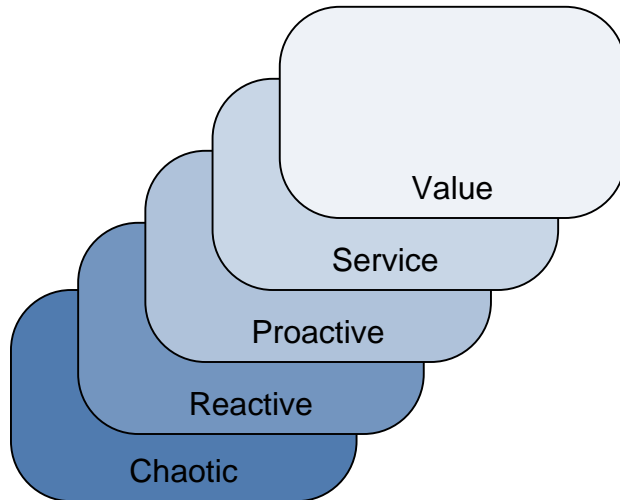


Figure 3.2: Gartner's five-level model of organizational maturity. IT organizations at different levels will recognize different value from their virtualization investment.

One of the major limiting factors that inhibits IT's march towards process maturity is in the device-centric nature of IT itself. Historically, individual IT devices tend to concern themselves with their own processing at the omission of others in the environment. This device-centric approach to management makes it difficult for IT to measure quality and effect change across multiple devices with high levels of accuracy.

Virtualization impacts the maturation process by enabling a greater situational awareness of virtualized workloads across the IT environment. Following the tenet of *commonality*, virtual machines themselves abstract individual differences between devices to gain higher levels of similarity. These common points of administration come as a function of the virtualization platform itself. That virtualization platform—potentially augmented with service-centric management technologies—comes equipped with change implementation functions that work across all virtualized workloads.

With this in mind, let's take a look at some of the activities commonly associated with change management. These activities relate to the logging of assets and changes to those assets, the monitoring of available and needed capacity, the assurance that services stay operational, the implications of licenses and their cost to the organization, as well as short- and long-term planning activities that are regularly required of enterprise IT. Each section will discuss the need and provide a real-world example of how virtualization improves the fulfillment of that activity.

Service Asset & Configuration Management

Service Asset & Configuration Management concerns itself with the need to “Identify, control, record, report, audit and verify service assets and configuration items, including versions, baselines, constituent components, their attributes and relationships” (Source: UK Office of Government Commerce, “ITIL - Service Transition”). What this effectively means is that every component of an IT system must be logged into a CMDB as well as each configuration item associated with that asset.

This task can obviously be a big job as the number of assets and configuration items grows large. The combination of large numbers of assets under management times large quantities of individual configurations of each means that the SACM process has the potential to be very exhaustive. Inventorying each configuration as well as maintaining this log over time through non-automated means becomes increasingly impossible as complexity grows.

As discussed in Chapter 1, virtualization enables the organization to take a step back from the level in which individual configuration items are logged. With virtualization, most virtual machines start their production life as a replica from a template. Once deployed, that replica immediately gains personality traits that separate it from its parent. It also gains added applications and other configurations that allow it to perform its mission. However, in a highly controlled environment, the basis from which every virtual machine starts is the template itself. This basis provides a reference point for change controllers to start from, and reduces the overall scope of asset logging for the environment.




The move to virtualization also tends to bring about a reduction in overall physical equipment in the data center. Notwithstanding their configuration, fewer physical assets under management—the recording of which can still involve manual processes—means less work for change management.

Capacity Management

Capacity Management is the process by which IT determines how well its services are meeting the needs of business. This determination looks at current load levels on servers and services and compares them to perceived quality metrics to validate that services are of good quality. It is also involved with the quantitative monitoring of services with an eye towards identifying remaining supplies of resources. In effect, capacity management provides a way for IT to plan for future needs.

The quantitative parts of capacity management are challenging to accomplish with physical resources. The device-centric nature of physical resources means that the resource load of one server can have little to no bearing on the resource load of another. This can even be the case between servers that work together to support a business service. One server in a service string may be fully loaded in the fulfillment of its mission, while another may be severely underutilized. The goal of capacity management is in identifying where demand exceeds supply as well as the reverse, a process that can involve substantial cross-device monitoring and analysis.

Compare this situation with what can be accomplished through the automation toolsets commonly found with enterprise virtualization solutions. Virtual workloads acquire their resources from available pools of supply, which in and of themselves are abstractions of real, physical resources. These abstractions along with the platform monitoring that measures their use automatically provide the mathematical models that enable quantitative capacity management right out of the box. IT organizations that leverage best-in-class virtualization platforms and/or the service-centric approach immediately and automatically gain an improved vision into the demand for resources in comparison with available supply. As you'll see later in the sections on planning, this data and the long-term analysis that goes with it enables IT to know beforehand when the virtual environment as a whole will require more resources.

 Capacity Management activities are one area in which service-centric management technologies are of particular use. Platform-specific virtualization toolsets do not tend to include the necessary quality monitoring capabilities required for effective measurement. One reason for this is that the measurement of quality comes from many sources and not just the virtual platform itself.

Service-centric management technologies tend to include these necessary capabilities due to their wide-reaching integrations into all levels of the IT environment.


Availability Management

The job of Availability Management is in ensuring that services are available when they are needed. This is done through the assurance and measurement of a service's serviceability, reliability, recoverability, maintainability, resilience, and security.

Virtualization's role in availability management is perhaps best explained through the use of a comparative example. Consider the situation in which an enterprise-based change request comes in the form of a mandate. That mandate requires the Recovery Time Objective (RTO) for all computers with a high level of criticality to be set at 15 minutes or less. This means that should a critical service go down, the business expects the return of that service's functionality to occur within 15 minutes. This policy means that effective backups and extremely fast restoration must be implemented in order to fulfill the requirement.

Accomplishing this with services that lie atop physical servers can be challenging. Servers purchased at one point in time may have different hardware than those from another, requiring the business to keep duplicate spares on hand for each server hardware type. The file-by-file backup of one server can miss files, making a restoration impossible. Traditional backup and restore processes may not be fast enough to meet the deadline.

Virtual servers are another matter entirely. With virtual servers, it is possible to back up a server at the level of the virtual disk itself. This type of backup captures every configuration on that server all at once, ensuring that a restore will occur with a high degree of success. The transience aspect of virtual servers also means that the restoration of that server can happen to any physical hardware that is running the virtualization software. Expensive duplicate spares are no longer necessary to keep on-hand to support emergency operations.

 With virtualization, disaster recovery and business continuity are achieved through the merging of entire-server backups with replication to alternative locations. This process ensures that up-to-date copies of virtual workloads are available and ready for failover whenever necessary and with an extremely low lag time.

License Management

License Management can be considered a subset of asset management, in that it is concerned with the logging and administration of software—and sometimes hardware—licenses across the enterprise. Although license management was discussed in detail in Chapter 2, its importance here is in relation to the automation of workload provisioning.

Later, this chapter will discuss some ways in which requested workloads can be automatically provisioned through virtualization toolsets. In setting up those self-service capabilities, a measure of attention must be paid towards ensuring that license limits are not exceeded. Requestors who have the unrestricted ability to create new virtual workloads will do so, a process that can over time cost the company substantial sums of money in software licensing. Caution to this end in creating automated provisioning systems will pay dividends in license savings over time.


Short-Term Service Planning

IT must prepare for two time horizons in its service planning activities. The first of these, short-term planning, involves itself with the management of today's needs in comparison with those expected in the short term. This planning activity concerns itself heavily with analyses of performance of services by looking at real-time metrics related to server performance and the ability for servers to perform service needs. When this analysis finds that servers are unable to appropriately fulfill the demand placed upon them, additional resources are brought to bear. These additional resources can take the form of additional hardware on a specific server or additional servers to balance the load.

Yet there is a problem intrinsic to both of these activities in the physical world. Physical servers are notoriously difficult to augment with additional hardware. Smart enterprises tend to purchase servers with a specific load-out, the adjustment of which can be costly or problematic due to limitations of the hardware itself. Towards the addition of new servers to handle the load, some business services are not designed with this capability in mind. Thus, the service may not have the intrinsic capability to horizontally scale. This forces IT to invent “creative” solutions towards meeting load demands that may not be in the best interest of the business.

Virtual workloads by nature are ready-made for horizontal scaling, up to the logical maximums of the virtualization platform or the physical maximums of the hosting hardware. For example, if IT finds that a Web server consistently runs out of memory in the processing of its mission, it is a trivial step to augment that server with additional memory to meet its demand.

Different here is in the mechanism of adding that memory. In the physical world, adding memory requires the physical step of adding that RAM to the server chassis. If IT chooses incorrectly the right amount of memory to add, the process iterates until it is correct. Each iteration requires additional downtime of the server. In the virtual world, the assignment of memory resources can be done through the virtualization platform or service-centric management tool's software interface. Rather than being forced to add memory based on the capacity of memory chips themselves, IT can elect to precisely assign the correct level of memory as needed by the server. This precision conserves memory for use by other servers in the environment.

 Although admittedly server memory is inexpensive, this example illustrates yet another benefit of the move to virtualization: A greater organizational respect for traditional performance management activities. To illustrate this point, consider a short history lesson. Until the advent of pervasive virtualization, the management of performance grew lax across many data centers. IT professionals who saw what amounted to unlimited resources in almost every server found themselves no longer needing to monitor servers for performance. With each server operating at 5% utilization, performance management was seen as an extraneous and/or unnecessary task during nominal operations.

Yet performance management remains a critical skill when operations deviate from nominal. When a business service is not operating to acceptable levels, IT must identify and resolve the source of the problem if they are to bring that service back to full operations. The organizational tribal knowledge associated with good performance management is key towards completing that task. One downstream result of a move to virtualization is that administrative teams must again find themselves looking at server performance if their consolidated virtual environments are to operate as designed. This experience provides them with the knowledge they need to accomplish that task.

Long-Term Service Planning

A major difference between short-term and long-term service planning is in the amount of data used in its analysis. Long-term service planning involves many of the same types of performance monitoring and analysis activities used for short-term planning. Different here is that long-term planning tends to make more use of historical data. This deep analysis of data across a longer amount of time enables IT to understand the levels of use of a business service over the long haul. With this understanding, it is possible to draw trend lines in service utilization that may ultimately point towards the need for expansion, augmentation, or new services entirely.

Also a component of this type of planning is the data gleaned through capacity planning activities. The combination of this data along with information about new initiatives and services that are expected in the future assists IT with determining whether additional resources will be needed.

The metrics associated with virtualization and its use in the environment are of particular help in these activities. Virtualization and its management toolsets tend to include long-term historical performance and utilization data gathering capabilities. This utilization data looks at every virtualized workload and compares its resource use with others in the environment. Unlike with physical systems where performance metrics tend to be unrelated to each other, the types of counters in most best-in-class virtualization solutions today enable enterprises to cross-reference servers against each other and in comparison with total resources in the environment.

The graphical representations enabled by this data make easy the process of analyzing resource use over extremely long-term periods. This data grows even more useful when it is compared with the aggregate level of hardware resources in the virtual “pool.” Environments that analyze this data in relation to the count and load-out of servers that participate in the virtualization environment can determine exactly when additional hardware is needed and what the configuration of that hardware should be.



Although the metrics themselves arrive as a function of the virtualization platform, the visualizations and suggested actions associated with their data may be insufficient with platform tools alone. Here is another area in which external tools such as those that follow the service-centric management approach may better illuminate this data.

Depending on the service-centric management solution chosen, it may be possible to automatically convert raw metrics into actionable information through its interface. This process takes all the guesswork out of long-term planning, deriving actual plans out of raw, quantitative data.

Automating Resource Provisioning

Yet another component of change management that is impacted by virtualization is the need for resource provisioning. This component represents the need for IT to deploy additional servers or change server configurations based on an internal or external request. The process of provisioning servers or changing their configuration in the physical world is exceptionally dependant on manual labor and non-automated tasks. Its process involves making and receiving the request, locating the physical asset, mounting and connecting that asset within the data center, provisioning operating systems (OSs) and applications, and testing for accuracy. With many of these tasks requiring the physical relocation of server assets, the time to complete can be an extended period.

Contrast this with the click-and-go nature of virtual machine provisioning. With essentially all enterprise-class virtualization platforms today, the deployment of a new virtual server involves little more than a few clicks to complete. That process includes making and receiving the request, identifying the necessary resources from the virtualization environment's available pool, file-copying the server template from its storage location, re-personalizing the new server asset, and installing any delta applications or configurations as necessary. The longest part of this process in many cases is the file copy itself.

Yet even with this level of automation, enterprise organizations can still experience problems with the correct distribution of physical resources to virtual assets. Consider a common situation in which multiple business units have collocated their assets into a single data center. Often, the separation between assets of one business unit is done by physical location, with individual server racks owned by one unit or the other. The provisioning of an additional server by one business unit typically means the addition of that server to the rack owned by the business.

With virtualization, there are reliability gains to be recognized by the consolidation of server resources into centralized pools. By eliminating the hard lines between physical server instances, the hardware resources of each can be aggregated into a single pool for distribution to requestors. However, in the situation where multiple business units may have joined into the pool, each unit likely wants to get out of the pool what they paid into it.




Figure 3.3: Enterprise virtualization solutions include multiple tiers of access. This tiering provides a way to distribute the tasks associated with provisioning to those responsible for hardware resources.

As Figure 3.3 illustrates, best-in-class virtualization products include the ability to subdivide resource pools into ever smaller groupings. For example, if 10 physical servers, each with 32GB of RAM and eight 3GHz processors, are aggregated into a virtualization environment, the sum total of resources the environment is said to have available for distribution is 320GB of RAM and 240GHz of processing power. In this example, if three business units joined equally in the purchase of the environment, each unit would expect to receive one third of that total.

Effective virtualization solutions include the ability to subdivide that 320GB of RAM and 240GHz of processing power into three separate resource pools. At the same time, these virtualization solutions also provide the permissions model necessary to assign rights to pool administrators. Those rights give the administrator the ability to assign resources out of their pool, but not out of the pools of others. Above the level of the pool administrator is the virtual environment administrator. This individual carries the privileges to administer the environment and adjust pool settings. Below the pool administrator may be individual server administrators. These individuals do not require the privileges to provision new servers but must be able to work with those that are already created.

This segregation of duties provides a convenient way for multiple environment stakeholders to aggregate their environment for the benefits associated with economies of scale. At the same time, it provides secure virtual breakpoints that ensure each purchaser gets out of the system what they put in.

Receiving and Evaluating a Request for Change

And yet the technology that allows the previous example to happen works only when it wraps around existing change management processes. In effect, mature IT organizations are not likely to recognize value in the technological benefits described earlier if those benefits are not part of a defined process flow.

Consider what a typical organization's Change Advisory Board (CAB) requires in order to fulfill this request. The CAB requires tools that evaluate the potential impacts associated with the provisioning and align its timing with other requests. Also key here are discovery and dependency mapping tools that look across the entire environment to understand infrastructure components and their relationship with hosted applications and data as well as the requested asset to be brought into production. ITIL refers to the CMDB construct as the storage location for this data, but the monitoring and analysis components are necessary parts of collecting it as well.

Traditional platform-specific tools for virtualization are not likely to include these needed add-on components. These tools tend to focus on enacting change to the virtual environment itself, and do not often concern themselves with the process structures that wrap around that change. Organizations with high levels of process maturity are likely to need these extra bits if they are to be able to manage changes over a large scale.

It is in these scenarios that technologies such as those that embrace the service-centric approach make sense. Like what is shown in Figure 3.4, the service-centric approach wraps around an organization's virtualization environment to link in needed components such as asset discovery and inventory, dependency mapping, CMDB, the service desk, and the business calendar. Also integrated are the administrative processes of the CAB itself, streamlining the request-and-approval process.

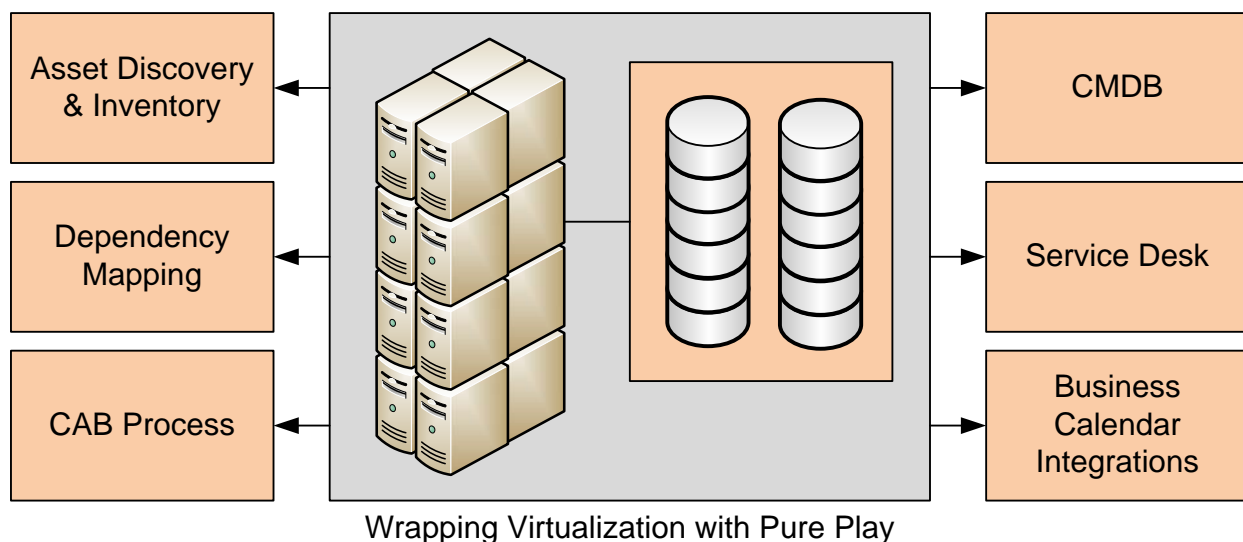


Figure 3.4: Enterprises with mature processes need virtualization that integrates into existing process elements such as those shown. Service-centric management's agnostic approach to integrations enables this to occur.

Considering this, let's explore an example of how a typical provisioning activity can be augmented through the incorporation of virtualization and enterprise framework enablers such as service-centric management. This example will take you through a typical request from its analysis and architecting phases through the release of resources.

Collecting, Collating, and Prioritizing Change Management Requests

When a provisioning request arrives, that request typically arrives through the organization's service desk. The service desk often has the internal logic necessary to submit the request to the proper channels for provisioning and approval. The collection and collation of requests involves understanding the environment and its current capacity, recognizing the validity of the request, approving it, and ultimately sending that request on to implementation teams for the actual provisioning activity. With the right integrations in place, this process works mostly within the organization's service desk application.

Identifying Resources

Once the provisioning request has been identified and submitted to the proper implementer, that individual must further determine whether the available set of resources can support the addition. In creating virtual machines, a number of elements must be configured at their moment of build:

- Assigned *processing power*, usually measured in MHz or GHz
- Assigned *memory*, usually measured in MB or GB
- Assigned *storage*, usually measured in GB
- Assigned *network*, to include the number of network cards, their location on the network, their assigned speed, and how they interconnect with virtual switches within the host
- Attached *peripherals*, which can be CD/DVD drives, storage LUNs, USB connections, or other external connections required by the virtual machine
- *Storage location*, which is dependent on the virtual machine's requirements for hot migration and collocation with other servers as well as the storage available in the environment
- *Collocation rules*, which ensure that the provisioned virtual machine exists on the same host as other defined virtual machines, which is often done to enhance inter-machine processing; conversely, some virtual machines must never collocate with others on the same host due to business rules or security reasons
- *High-availability and load-balancing rules*, which determine the machine's behavior in the case of a host failure; load-balancing rules additionally ensure that resource overuse by the provisioned virtual machine does not impact the performance of other collocated servers.

This information can be included as part of the workflow surrounding the provisioning request or can be an action to be determined by the pool administrator. Each of these elements must be identified at the time the virtual machine's resources are carved out of the available pool.

Architecting the Solution

Once resources have been identified and analyzed in comparison with established supplies and short-term expectations, the pool administrator must then design the solution. This design is involved with actually creating the virtual machine from its template, assigning its personality, and installing any additional applications or configurations as required by the requestor.

The follow-on activities associated with re-personalization and application/configuration installation are often activities that are not actionable through platform-specific virtualization tools. Lacking the additional toolset support, this process can be a substantially manual process, requiring manual steps on the part of the implementer. Effective toolsets—especially those that support systems management functions—enable the implementer to instead initiate scripted actions for the completion of these follow-on activities as required by the requestor.

Change Management & Release Procedures

The change management process at this point requires a number of additional steps. First, prior to turning over the requested resource to its requestor, the implementer typically requires a final approval step. This step is often a function of the CAB. The CAB in this instance looks at the request in comparison with the rest of the environment with an eye towards conflicts between this request and others. Those conflicts can occur based on future demand for resources, conflicts with applications themselves, or timing issues in the implementation.

A further activity required by change management is the logging of the asset itself as well as the individual configurations that make up that item. Configuration management can accomplish this process through automated discovery and inventory tools, or it can be a manual activity. The logging of the asset is often considered the last step prior to turning it over for production use.

Once approved by change management and logged into the organization's CMDB, the virtual asset can then be turned over to its requestor for service. Virtualization improves this activity through a universal capability to remote control the console of each virtualized server. This increases the overall security of the environment while at the same time provides easy access for administration.

Technology Architectures for Enabling Managed Self-Service

Taking this level of automation one step further, many virtualization platforms today include the native or add-on capability for self-service provisioning. The concept here of self-service provisioning entails an interface whereby a requester can submit a request and have that request be automatically provisioned. The interface determines the validity of the request in terms of permissions, requested resources, and available supply. If accepted, the interface then automatically provisions the asset complete with needed personalization. Once complete, the user is granted access to manage the server through the virtualization platform's remote console.

Implementing and Automating Self-Service

Self-service goes a step beyond the automation processes discussed to this point because it fully automates every step in the provisioning process. This implementation obviously requires a substantial level of up-front work to set up the workflow and provisioning steps within the interface. However, once set up, provisioning teams can take a completely hands-off approach to fulfilling new server requests.

A few elements must be put into place for self-service technologies to work in this manner:

- *Workflow.* Any resource requesting process must have the right kinds of workflow elements enabled within the system. This means that the interface used for automatic provisioning must likely integrate with the service desk application in some way. This enables a provisioning request to work within the workflow rules as assigned for manually fulfilled requests.
- *Resource pools.* Pools of available resources must be supply-limited. This forces requesters to work within their constrained limits when determining what servers and server configurations they require.
- *Templates.* Individual virtual machine templates must be created and de-personalized so that they are available for provisioning and re-personalization upon request. This process can involve a substantial amount of time if the level of individual template configuration is high.
- *The interface itself.* Although the virtualization platforms that include self-service capabilities tend to have off-the-shelf interfaces built-in to the product, your level of workflow integration may require additional customization. Understanding what capabilities are available through off-the-shelf interfaces in comparison with your needs will determine whether custom development is necessary.

As you can see, self-service provisioning tools have particular implementations where they tend to work better than others. Environments that have a high turnover of server assets, such as hosted desktop or development and test environments, tend to function very well with self-service due to the high rate of change in these environments.




In these environments that see high rates of asset churn, the up-front cost of building the self-service interface is paid by the time-savings gained through each individual request.

Governing Self-Service Activities

As should be obvious, an effective plan of governance must be set into place with self-service provisioning tools. Otherwise, putting these technologies in place can lead to a massive oversubscription of resources as virtual machines are self-generated over time. The concept of resource pools is one mechanism that places the onus of responsibility for smart virtual machine creation onto the requester.

In this example, an organization can decide that a particular project should be given a certain level of resources. Once that level has been identified and assigned to the project, the project personnel can determine on their own how to “spend” their available resources. This process gives flexibility to virtual administrators, eliminating their involvement in the day-to-day decisions of each project that makes use of the virtualization environment. At the same time, it gives flexibility to project teams to allow them to distribute their assigned resources as they see fit without the need to iteratively return back to the CAB for each individual change. In the end, this flexibility enables greater agility to the organization as a whole while maintaining the control as required by change management.

 Some technologies use a type of quota point system rather than resource pools. In this example, each quota point refers to a server of some predefined level of capability. Projects can be assigned quota points based on their need, and can assign multiple quota points to a single server instance to realize a more powerful server. The end result is effectively the same as with resource pools but with slightly less granularity of resource assignment.

Virtualization Technologies Enable New Automation in Change Management

This chapter has but touched on the ways in which virtualization and its accompanying automation capabilities can enhance the fulfillment of change management tasks for an enterprise organization. Enterprises that leverage virtualization stand to gain substantial improvements to the fulfillment of processes while easing the process of monitoring and managing change over time.

But change management is only one area where virtualization can impact your operations. Chapter 4, the last of this guide, will continue this discussion with a look at problem resolution and the areas in which virtualization and its service-centric management augmentations provide an assist.

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.