

Data Center Virtualization and Orchestration: Business and Financial Justification

Overview

Over the past several years, virtualization of server, network, and storage resources has received a lot of IT, media, and vendor attention because of its promise for reducing both capital and operating expenses within the enterprise. For service providers virtualization offers a refreshed business model for outsourcing customer applications and for locating these applications within their regionally managed data centers with well-defined and secured partitions. For enterprises, virtualization offers data center consolidation and reduced power and cooling costs and often defers the need for additional data center sites.

Virtualization decouples the application environment from the constraints of the hosting computing, network, and storage hardware. This decoupling can mean logically partitioning one device into many, consolidating many physical devices into one logical resource pool, or both at once.

It is well understood that virtualization in this sense can produce a short-term capital expenditure slowdown when it equalizes and improves utilization rates of existing equipment. However, siloed virtualization can have unintended downstream consequences that limit its effectiveness as well as create complicated and potentially costly security and regulatory compliance concerns. To gain the full, long-term benefit of virtualization, architects and planners must understand and treat network, storage, and computing devices as interdependent pools of resources within the data center ecosystem; indeed, the network underpins the other two and provides an I/O conduit for them.

Another aspect of holistic virtualization involves next-generation provisioning and orchestration products that offer service-level abstraction on top of the virtualization engines. Security, content load balancing, high availability, the amount of capacity of demand pools, and disaster recovery configurations are all orchestrated as a set of virtualized resources based on application requirements. Moreover these products, with the inclusion of real-time policy-based engines, can quickly reconfigure these virtualized resource pools based on capacity changes, outages, security events, and the need to quickly bring new applications online.

Plan Architecturally, Deploy Incrementally

In many enterprise IT organizations, this new service-level provisioning and orchestration paradigm is viewed as organizationally challenging because these products provision across multiple functional areas, including server, network, and storage operations, thus requiring collaboration. Although conceptually these products hold a lot of promise, many IT managers are so focused on day-to-day challenges that they find it hard to plan for and justify this cross-functional process and operational integration.

However, the user community that IT is servicing is often moving faster and placing more demands on its IT infrastructure than many IT departments are prepared to handle. According to Nemertes Research, approximately 80 percent of collaboration tools (instant messaging, blogs, wikis, etc.) are introduced into an organization by users. This new interaction model, and the associated security and compliance concerns, is unavoidable, and it can be well supported by a network-based, incrementally deployed service-oriented infrastructure (SOI).

Well-orchestrated, virtualized infrastructures with highly available pools of resources offer clear competitive advantages. Business process changes become faster and easier to operationalize, allowing an enterprise to respond to changing requirements more easily. IT becomes an enabling partner instead of a concern in business continuance planning and risk management initiatives. A valuable side benefit of holistically virtualizing the data center is that the data center can largely be managed remotely. This capability facilitates site removal to cheaper locations that may not contain or be attractive to sizable pools of skilled IT workers. The organization can now accomplish more with the same IT resources, and in many cases may attain sufficient capital and operational savings to redeploy some of the 60 to 80 percent of the IT budget that traditionally goes to maintenance for forward-looking projects that can give the organization competitive advantage.

The Cisco Solution

Cisco VFrame Data Center is a cross-functional, rapid services orchestration and provisioning platform. This platform offers a modular approach for configuring and rapidly deploying application services across the data center infrastructure, including the provisioning of server, networks, and storage chassis together as an orchestrated, virtualized set of hosting services. Cisco VFrame Data Center has management and policy-driven provisioning interfaces across these three functional areas and can add server capacity within minutes; reallocate this server capacity based on overall capacity and high-availability demands; and rapidly deploy of a rich set of security, load balancing, and other virtualized services such as virtual storage area networks (VSANs), VLANs, and access control lists (ACLs), based on predefined service templates.

Making the Case for Operational Change

As organizations begin to seriously evaluate virtualization initiatives, they are realizing that such projects entail operational changes that introduce costs that go well beyond the initial purchase price of the technology. Yet these costs are often the component that is hardest to measure within traditional return-on-investment (ROI) models, especially compared to the opportunity costs associated with maintaining existing systems and making reactive, ad hoc infrastructure tweaks. This document describes several of the financial benefits of data center virtualization, in particular with an orchestration offering across many provisioning groups such as Cisco® VFrame Data Center, to help ensure that the operational changes that most companies will incur are worth the near-term costs.

Time Savings

Increasingly IT organizations are asked to change or roll out new applications more quickly in response to changing market conditions. However, these applications typically require a rich set of infrastructure services, including multi-tier server hosting, well-partitioned and secured centralized storage and client-facing services, and value-added services such as Web caching and intrusion prevention. Current industry averages for provisioning equipment in advance of rolling out a new application are 6 to 8 days for servers and 15 to 20 days for storage. In addition, provisioning is often performed serially, with delays in handoffs between groups. A total deployment cycle can

sometimes last over a month. Because of the amount of coordination involved, very few changes are made over the life of many of these applications.

Cisco VFrame Data Center offers a role-based access control interface that facilitates collaboration across functional organizations. Additionally, Cisco VFrame Data Center provides a service-level design interface in which all the infrastructure services required by an application can be designed through a graphically driven building block approach. Each functional group can log into Cisco VFrame Data Center and add its functional inputs. This graphical approach at a services abstraction level greatly facilitates collaboration and speeds the design operations with secured views across all functional groups. After the design is complete and the defined application services are ready for deployment within the infrastructure (across server, networking, and storage resources), with the push of a button Cisco VFrame Data Center installs up to 500 servers and all their downstream network and storage settings within minutes. One early adopter estimated a reduction in total provisioning time (including design) from 3 weeks to 3 days.

Every enterprise's own costs will vary, but here are some averages:

- Average server deployment time (24 hours) x average US\$75 per person-hour = US\$1800
- Average storage deployment time (130 hours) x average US\$75 per person-hour = US\$9750
- Average application deployment time (40 hours) x average US\$75 per person-hour = US\$3000

Assuming no changes to network devices, and not including the costs of new hardware, service coordination, or system troubleshooting, it costs almost US\$15,000 and takes more than a month to deliver a new application.

With Cisco VFrame Data Center, the calculation is different: Total deployment time = 24 hours x US\$75 per person-hour = US\$1800, or more than 85 percent reduction in personnel costs alone. The person-hours can now be directed to new, higher-value projects.

Capital Savings

Because it can be so time-consuming to coordinate and deploy the resources necessary to support a new application, IT departments commonly overprovision with hardware capacity, typically by 30 to 40 percent. Server utilization rates remain at 35 percent even in the age of server virtualization, and storage utilization rates are often even lower. At the same time, storage purchases are increasing by approximately 60 percent annually. Device virtualization should eventually improve utilization rates, but realistically only when done in the context of system-wide orchestration.

The ease of orchestration and provisioning of the overall infrastructure achieved with Cisco VFrame Data Center allows IT departments to correct for initial overprovisioning, reducing server purchases by as much as one-third. At an average cost of US\$4000 per server (initial hardware costs only), a reduction of only 15 servers completely offsets the initial cost of the Cisco VFrame Data Center appliance.

Other Operational Savings

A significant benefit of reducing the total number of hardware devices in operation is the direct effect on data center power and cooling costs. Gartner has estimated that by 2009, IT's energy costs will grow from an average of 10 percent of the IT budget to more than 30 percent. By throttling hardware proliferation in the data center and by shutting down devices not in use, Cisco VFrame Data Center helps IT gain some control over its energy costs. Cisco VFrame Data Center can address power and cooling concerns in other ways as well, as discussed later in this document.

As previously mentioned, siloed virtualization raises data security and regulatory compliance challenges. Holistic orchestration can prevent these problems, as well as costly fines and lawsuits, which can into the millions of dollars.

Finally, although impossible to calculate here, significant opportunity costs are associated with long and complicated application deployment cycles. These can be calculated in terms of idle user person-hours, in the case of new internal applications; lost productivity, when internal applications cannot be adapted to eliminate manual processes; or lost sales per day, in the case of external-facing applications.

Utility-Based High-Availability Resource Pools

Web applications are typically hosted on x86 servers in multi-tier server architectures. To help ensure high availability, organizations typically provision redundant server, network, and storage services on a per-application basis. In this approach, redundant hardware is not shared between application groups because each piece of hardware has been uniquely provisioned for the failover services it needs to run. Although failover hardware pools are required to ensure uninterrupted operations, statistically application pools are unlikely to all experience failures at the same time on like hardware resources.

One of the primary benefits of a rapid service provisioning and orchestration tool is the capability to detect hardware failures through internal or external monitoring systems and within minutes to dynamically add new hardware from a utility resource pool. This utility pooling of redundant resources across multiple application groups, and the capability to provision resources within minutes with no-touch intervention by an administrator, greatly reduces the amount of over-provisioning within each application group. This approach allows customers to think of high-availability pools differently: as banks of utility servers that can be quickly booted up and added to any application that has experienced failures. These pools include all the network and storage services that need to be associated with this server as it is added to the application group. This approach offers significant cost savings because shared utility pools can reduce the amount of non-operational redundant hardware by 20 to 30 percent.

For example, rather than configuring idle capacity within VMware ESX server pools (and all the downstream related network and storage services), an organization can have a utility pool of x86 servers that Cisco VFrame Data Center can dynamically boot based on change events. When an ESX server failure is detected through the internal monitoring of Cisco VFrame Data Center or through an external monitoring system such as VMware Virtual Center, Cisco VFrame Data Center can dynamically look into the preallocated utility-based server resource pool and find the server that best matches the profile set for hosting ESX server. After a server is located (within seconds), Cisco VFrame Data Center remotely boots this server from a SAN with an ESX image. Cisco VFrame Data Center then communicates to VMware Virtual Center the presence of a new, fully

operational ESX server, including the configuration of all the network attributes, and then allows VMware Virtual Center to load or migrate running virtual machines from the failed or greatly impaired server. Conversely, Cisco VFrame Data Center can release the ESX server back to the utility pool when it is no longer needed. The total operation time to bring up an ESX server is less than 5 minutes, based on the boot time of ESX remotely. This shared utility pool approach greatly reduces the number of idle servers within an ESX pool.

Load Balancing at the Facilities Management Level

There is growing concern about power, cooling, and floor space within data centers. Although much of this concern can be initially addressed through server consolidation, with blade server architectures and hypervisors, power and cooling still need to be load balanced across racks. Through effective load balancing and elimination of rack hotspots, IT departments can significantly reduce their cooling requirements, and therefore their power requirements. Moreover, they can do this with the most commonly installed floor-to-ceiling heating, ventilation, and air conditioning (HVAC) systems and front-to-back cooling, rather than by retrofitting using specialized cooling systems. In most server hosting environments today, applications and the downstream network services are fixed (or hardwired) to the systems to which the services have been assigned. Although it would be nice to shift a bank of servers that are driving temperatures up in one part of the data center to a cooler spot within the facility, these fixed models prohibit this kind of shift of the applications to cooler spots in the data center.

Dynamic provisioning and orchestration tools such as Cisco VFrame Data Center engage with the underlying infrastructure, and through APIs they can be integrated with external temperature-sensing equipment. Through rapid remote server booting and services provisioning capabilities, they can shift the hotspot server workloads to other servers that are available in cooler locations. These shifts can balance server loads based on temperature distribution policies. The core technology behind this load balancing is the capability to dynamically bring up a server in a different location within the data center (through policies with no administration intervention), with remote server rebooting and provisioning of all the downstream services relative to the application definitions of the runtime image being booted to the server. These definitions are part of the stored service templates and configuration workflows.

Cisco VFrame Data Center has a rich discovery mechanism for servers. This discovery includes comprehensive understanding and mapping of the network connections from the server to the adjacent switch ports in Cisco Catalyst[®] switches and Cisco MDS 9000 family storage switches. Through the Simple Object Access Protocol (SOAP)-based Extensible Markup Language (XML) API, Cisco VFrame Data Center can accept external thermal temperature monitoring information, which can indicate to Cisco VFrame Data Center which switches have higher than normal thermals. Using this thermal data and the discovery information, Cisco VFrame Data Center can determine the servers that are adjacent to this switch and can dynamically boot servers in other locations in the data center and launch the same services.

After these other servers are booted, Cisco VFrame Data Center can decommission the servers in the network hotspot by shutting down these nodes. This interaction is based on the available software developers toolkit (SDK) and the integration of the API with power and temperature monitoring systems.

Conclusion

Data center virtualization technologies offer many cost benefits. Device consolidation and improved utilization, enabled by network technologies such as Cisco Application Control Engine (ACE) and tools such as those provided by VMware and others, provide some short-term benefits. These benefits can be extended and enhanced through dynamic, holistic provisioning.

This document provided some introductory methods for calculating business agility benefits, including IT person-hour savings, user productivity, hardware purchase and maintenance reduction, and energy savings. All these virtualization techniques also contribute to greater revenue opportunities because application services can come online more quickly, with better security and responsiveness.

For more detailed ROI analysis, Cisco has developed a downloadable tool that customers can use to model many of their own costs. Many of the costs pre-loaded into this tool are based on inputs from Cisco IT and several other real-life working examples of the current cost of capital, including hardware, software, and labor. For more information, please visit <http://www.cisco.com/go/datacenter> or <http://www.cisco.com/go/vframe>.



Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc. Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc. and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0705R)