Beyond the Hypervisor: A Holistic Approach to Developing Next-Generation Data Center Architecture

As organizations work to evolve their IT infrastructure from a static collection of individual assets into a system capable of rapidly adapting to meet business demands, they have become very interested in the concept of infrastructure virtualization. And for good reason: properly architected and deployed, today's virtualization technologies hold significant promise in helping IT meet the three perpetual business demands of better, faster, and cheaper.

Conceptually, virtualization is not new. Since IBM first offered mainframe virtualization, the idea of getting more value out of existing physical resources through logical abstraction has been appealing. Today, most enterprises thinking about virtualization have in mind technologies such as those available from VMware, Xen, Microsoft, and others that enable server virtualization through the use of a hypervisor. Although this is a good place to start in using virtualization to transform the current-generation data center, the journey toward a more dynamic infrastructure should not end there.

This paper examines the topic of data center virtualization from an architectural perspective and points out where established and emerging network technologies can play a key role, not only in keeping down costs and ensuring operational continuity, but in helping make IT infrastructure an enabler of business agility.

What Is Virtualization?

In computing, virtualization is a broad term that refers to the abstraction of resources. Virtualization technologies can be used to create a logical version of a physical device (such as a server, storage product, or piece of network equipment) or another resource, such as an operating system. Virtualization can be used to divide the resource into one or more execution environments carved out from the original. In this sense, even something as simple as partitioning a hard drive is virtualization because in doing so, you take one drive and logically divide it to create several separate logical drives. Other devices, applications, and human users can then interact with the virtual resources as if they were several real physical devices.

Virtualization is also used to describe a second style of abstraction wherein many different devices are aggregated to present one logical instance. Again, other devices, applications, and human users can regard this virtual resource as a single entity. A RAID array representing a single logical drive composed of several physical spinning disks is a familiar example.

These same basic patterns can be found across many different computing technologies. Let's look at them in the context of server virtualization.

Figure 1. Server Virtualization

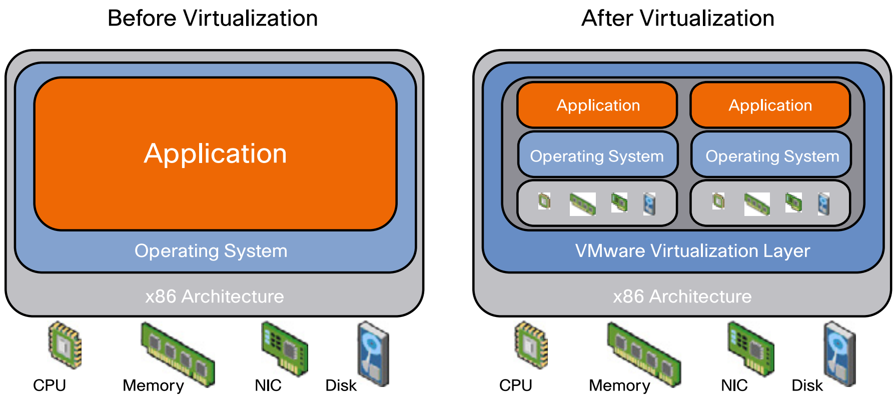

Source: Vmware, Inc.

Pattern #1: Segmentation or Partitioning

This pattern is the one that enables many virtual instances to reside within a single physical device. For example, as in the popular VMware ESX solution depicted in the diagram, the VMware hypervisor acts as a virtualization layer allowing one physical machine to host multiple guest operating systems. The resources of the actual physical host are abstracted and virtually segmented so that each OS instance perceives that it has its own dedicated CPU, memory, NIC, and hard disk.

Instantiation of the segmentation pattern in this case allows for multiple server instances (running either homogeneous or heterogeneous operating systems) to run in the same physical machine. This type of server virtualization can provide application platforms with isolation comparable to that of separate hardware – thereby obviating the need to purchase dedicated equipment for each new project. Applied to compute assets, this pattern allows you to take better advantage of underutilized physical resources on a server, in this case by increasing the number of guest OS images residing on a physical machine.

Pattern #2: Aggregation or Composition

Virtualization also can enable multiple physical devices to be represented by a single virtual instance. We see this pattern manifested in today's utility or grid computing offerings. In utility computing, multiple low-cost PCs may be grouped together to form an "on-demand" metered service as if they were one very scalable machine. Typically in this model, the resource controller for a grid interacts with the underlying physical resources to provide a virtualized environment across which one can run an application instance or task as depicted in the figure below.

Figure 2. Grid Computing

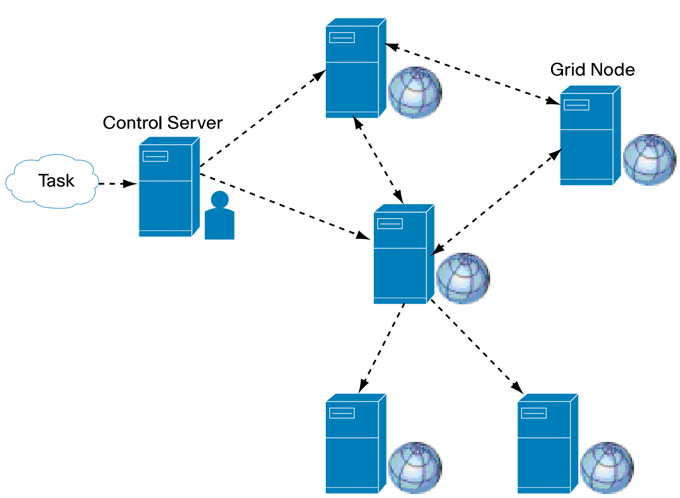

This pattern allows for physical resources to be added to the cluster as needed, ensuring extreme horizontal scalability. Whereas this pattern may do little to reduce the power, cooling, and real estate footprint of a given application, it is an effective way to enable growth without over-provisioning and to provide highly available application services.

Of course, these two basic patterns, segmentation and aggregation, can be used in combination in the context of a given application. Consider the benefits of having a single server virtual machine (VM) (leveraging segmentation) depend upon a network-attached storage array as its “disk” (leveraging aggregation). By combining these two, for a given enterprise application we can achieve both the runtime isolation and storage scalability it requires. At the same time, we concurrently allow for increased server utilization (on the VM host machine) and data recovery capabilities (provided by the storage array) as sought after by those responsible for the underlying infrastructure. From these two basic architectural building blocks, much can be accomplished.

The VM segmentation example just given is also commonly referred to as one type of “platform virtualization,” the virtualization subtype that currently has the most attention from analysts, the press, and IT architects. Unlike many other types of virtualization, this particular set of technologies is perceived to have launched a fundamental disruption in data center architecture. Now the standard “atomic unit” around which most data center activities focus has begun to shift from the server to the server VM.

What Are We Trying To Achieve?

Today’s enterprises are under constant pressure to expand business capabilities, improve real-time information access, and provide richer user interactions. Globalization and new business models are breaking down enterprise boundaries. Increased demand for analytical data driven by industry regulations is introducing new complexities to information processing. And new Internet capabilities are raising customer expectations and demands on user interaction.

Businesses are responding with a new generation of applications taking advantage of the latest architectural trends such as service oriented architecture (SOA) and Web 2.0 in the pursuit of delivering business value. The current generation of data center infrastructure is hard-pressed to keep up. The rate at which new servers can be put into place frequently trails application demand, and at the same time many data centers are running out of floor space. Indeed, many consider the present-generation data center architecture to be at a breaking point.

So what then does the latest wave of virtualization, where hypervisor technology has become the focal point, bring to the table? What opportunities does this trend present? Why are organizations so interested in virtualizing their server infrastructure?

Higher Efficiency

Probably the first reason you’ll hear is the expectation that adopting virtualization technology means better utilization of existing assets. Over the years, due to the past best practice of acquiring and deploying new equipment for each major application to minimize critical dependencies, many data centers have become filled with dedicated equipment for individual applications. Such project-driven procurement practices have led to isolated, over-provisioned hardware environments designed to support peak loads for specific applications. While this made sense from a line-of-business or large application perspective, data center managers soon noticed that many of these isolated hardware environments were significantly underused by the applications they were dedicated to support.

At the same time, data center cooling and power costs and floor space continued to rise linearly with the introduction of each new piece of hardware. Now both rack and floor space are scarce in many data centers, and pressure is building to construct more facilities. Being able to run multiple instances of a virtual resource on fewer physical machines takes off the pressure for new floor space and lowers the projected demand for additional power and cooling.

The first benefit of server virtualization, then, is simply the opportunity to get more out of the gear on hand. Beyond that, virtualization can also ease the current pace at which new hardware is necessary for hosting new applications.

Improved Resiliency

The second major benefit offered by server virtualization is simplified disaster recovery operations. By decoupling the logical functions of a device from its physical container, one important architectural capability that results is portability. With the right configuration information to capture the state of a given execution environment stored in an associated file set, applications can be easily restored after a failure or interruption of service. With the introduction of the virtual machine, an entire logical server and its workload can be restored by retrieving a copy of the VM's files, moving them to a compatible new virtual server host, and importing the VM configuration. This alone can reduce application service recovery time from days to minutes.

These same techniques can also be used to assist in the course of normal operations, such as rolling out or rolling back changes to an application environment or conducting a planned data center migration. The benefit in this case is gained by reducing or eliminating the need for application downtime.

Increased Agility

The final major benefit of a virtualized infrastructure, and perhaps the most important one, is the ability to deploy new execution environments for applications faster than ever before. Traditionally, bringing up a new application environment that required safe partitioning called for provisioning a new server. That could often take months to complete given the associated tasks of capital procurement, racking, stacking and cabling, OS imaging, configuration and testing, properly integrating the new host with the network and storage systems, and so on.

In a virtualized world, multiple VMs share hardware resources without interfering with each other so the effort required to bring up a new isolated instance is significantly reduced. Once a given physical server is in place, new VMs can be created relatively quickly using management and provisioning software. This approach to bringing up a new virtual host significantly reduces the amount of coordination required across teams, saving time in preparing a new applications environment.

Research conducted by the Economist Intelligence Unit in January 2008 suggests why this last benefit may be the most noteworthy for enterprise IT organizations. The survey, sponsored by Cisco and entitled "IT Transformation: Creating a Strategy for Success" , found that in the beginning of 2008, 57 percent of the more than 950 IT professionals surveyed considered their top objective for the year to be improving responsiveness to new business requirements.

Figure 3. What are your IT organisation's top objectives during 2008? (Select up to three)

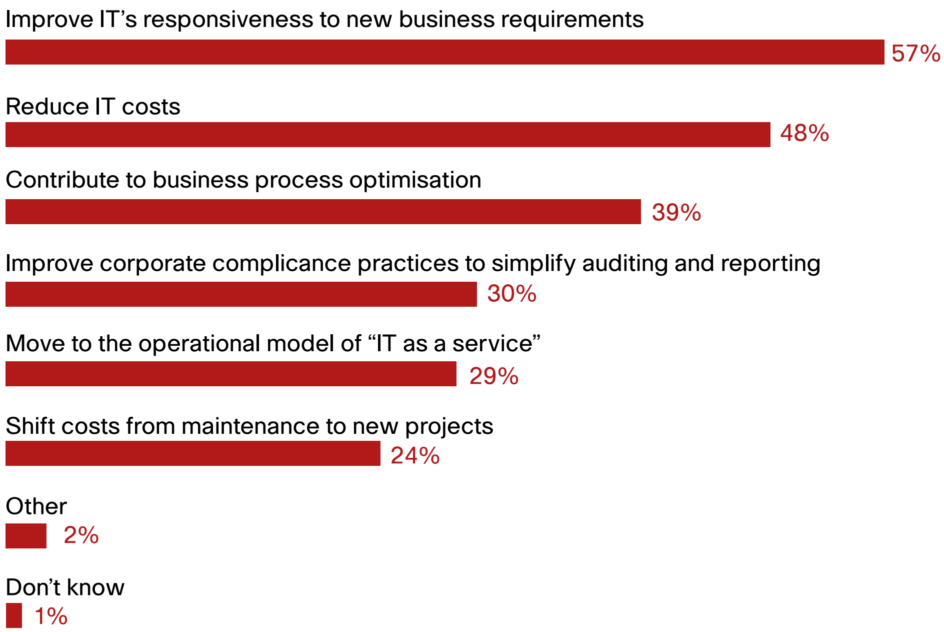

Source: The Economist global survey, Jan./Feb. 2008

When asked what their top projects were for the year associated with meeting these goals, 28 percent of the respondents named virtualization.

Figure 4. What are your IT organisation's top project priorities through 2008? (Select up to three)

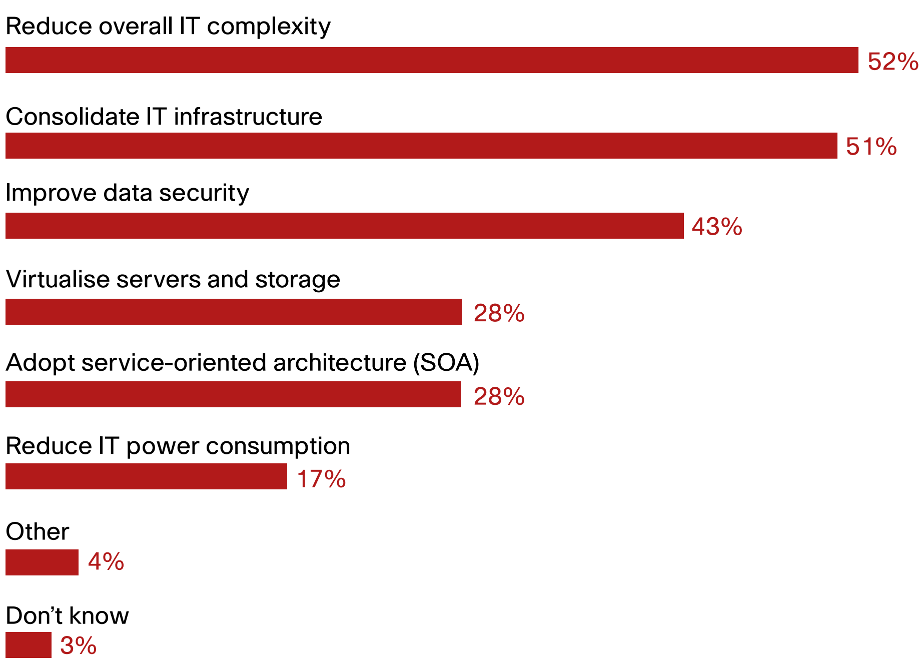

Source: The Economist global survey, Jan./Feb. 2008

Since the time this survey was conducted, the global financial climate has changed significantly. It may well be that the perpetual balance between reducing IT costs and improving IT's ability to be more responsive, flexible, and innovative has shifted toward cost reduction for many organizations. However, the good news is that server virtualization initiatives can often be justified on a cost basis alone while concurrently delivering on the other strategic goal of making IT more responsive to the continual evolution of the business.

Architectural Considerations

Given the foregoing, it appears that server virtualization holds out the promise of helping IT to be better, faster, and cheaper. However, as with the introduction of any new technology, how can we make it most successful within the existing operational context? How can we also ensure that as virtualization takes root beyond the server (in storage, applications, desktop, and so on), over-optimized silos of individual technology virtualization do not degrade the overall data center environment? This is where enterprise architects have a key role to play.

Many organizations have a rather accidental infrastructure architecture. Having at some time selected a few key standards for server, networking, and storage hardware and having built organizations around these individual technologies, IT has since deployed numerous types of appliances, management tools, and one-off hardware solutions at the request of disparate applications teams that have resulted in quite a puzzle. It's no wonder that many organizations spend more than 75 percent of their budgets on maintaining their current assets versus funding innovative projects. The "integration spaghetti" that applications architects often wrestle with commonly has its analog in the world of infrastructure as well.

Even if your organization does not suffer from these symptoms, chances are that the goals and objectives of the various infrastructure teams in your enterprise don't always align. Dependencies between technologies are not always fully understood. The server, storage, and network teams may blame each other when there is an application outage. Against this backdrop, introducing virtualization technology presents both an opportunity and a new set of challenges.

With the introduction of hypervisor technology, the new atomic unit of the data center effectively becomes the VM workload versus the server. Considering that the focus of data center operations has long been on what happens in and around the server, this seemingly small change has begun to have significant impacts on architects, engineers, and operations teams throughout the data center.

Nearly every system in your data center was originally built with a focus on the physical server. From designing network topology, to planning backup operations, to putting together security policies, the likely assumption has been that the server was the center of it all. What happens then when that assumption no longer holds true? Or to put it another way, once you've decided to virtualize the server and focus on the workload, what are the implications for everything surrounding the workload? If you begin to consider advanced virtualization techniques, such as live migration of VMs between physical hosts, perhaps some of the potential challenges are already coming to mind. There is work to be done beyond the hypervisor to achieve a more efficient, resilient, and agile infrastructure.

However, with the disruption of virtualization also comes architectural opportunity. Incorporating server virtualization into your environment is significant enough a disruptor to allow you to refresh (or perhaps establish for the first time) your architectural roadmap for the next-generation data center. The key to achieving success with your plan is to ensure the right scope and perspective, ideally taking a more holistic, systematic approach to the problem, and to make sure that both people and process issues are considered as part of the plan.

Figure 5. IT Infrastructure Needs to Evolve

A Vision for the Next-Generation Data Center

Every successful architectural shift starts with a vision of the desired end state. What should that look like for the data center? For most organizations, it's becoming clear that server virtualization will be crucial. But what else needs to be taken into account? How can other types of infrastructure virtualization assist in completing the picture? Should you even consider budgeting for new capital expenditures when your management team has recently become interested in Amazon's Elastic Compute Cloud (EC2) and similar offerings from Microsoft, IBM, and others? Where to begin?

Let's review what we're aiming for: higher efficiency, improved resilience, and greater agility. Lower cost of ownership, continuous availability, and faster deployment of new solutions are what the business demands. Internally therefore, there is a need for a new way of envisioning, architecting, and ultimately implementing data center designs that will result in a more dynamic infrastructure.

Although this renewed interest in virtualization may have begun at the endpoints of the network with the introduction of hypervisor technology, there is a significant role that the network itself can play in helping to complete this vision. After all, the network is the only common single element that connects and enables communications between all components of your IT infrastructure. It may be time to think from the inside out: starting from core network capabilities and reaching out to where the network meets the VM.

The Role of the Network

The network is no stranger to virtualization. Consider our two patterns, segmentation and aggregation: VLANs and VPNs have allowed network administrators to carve out segmented resources allocated to individual applications and users for many years. The introduction of server load balancing devices, conversely, has enabled aggregation of server instances leading up to today's high-performance clusters.

Modern load balancing devices can be partitioned to create virtual contexts on a per-application basis. The same is true of firewall services. Within a given router, the routing table itself can be partitioned using VPN routing and forwarding (VRF) technology. Across Cisco® Catalyst® switches, a virtual switching system (VSS) can be created by aggregating two devices together via multi-chassis etherchannel (MEC) to provide a single logical switch. In latest-generation Cisco Nexus™ switches, Virtual Device Contexts (VDCs) can be provisioned on a per-port basis to provide dedicated Layer 2 and Layer 3 services to a given flow of traffic. And of course, these virtual building blocks can be combined together to create unique solutions such as providing the basis for guest access on a shared corporate network through the combination of VPNs plus VRF associated with particular VLANs.

Indeed, the network has evolved to the point where it has two major roles to play in enabling the next-generation virtualized data center: virtualization of network services themselves and supporting more rapid adoption of server virtualization and some of its more advanced techniques.

Looking more closely at this second role, the network is uniquely positioned to bring together the domains of storage, networking, and server virtualization into a more unified computing experience. Many of the steps down this path were pioneered through innovative technologies that are now open standards or in the process of becoming open standards. The virtualization of storage area networks (VSANs) was introduced by Cisco and became an ANSI standard in 2004, allowing for the consolidation of SAN islands into a more cost-effective redundant virtual fabric made up of a collection of Fibre Channel switch ports. The introduction of lossless Ethernet and Fibre Channel over Ethernet (FCoE) capabilities as part of data center Ethernet offerings now supports the convergence of local and storage area networks into a single unified fabric. Now the network is a key component not only in virtualizing storage, but also in reducing the physical cabling and adapter infrastructure needed to support storage access. Pairing such I/O virtualization technology with the movement toward server virtualization can greatly simplify existing data center infrastructure and deliver increased efficiencies in cooling, power, and space consumption as conceptually depicted in this diagram:

Figure 6. Unified Fabric

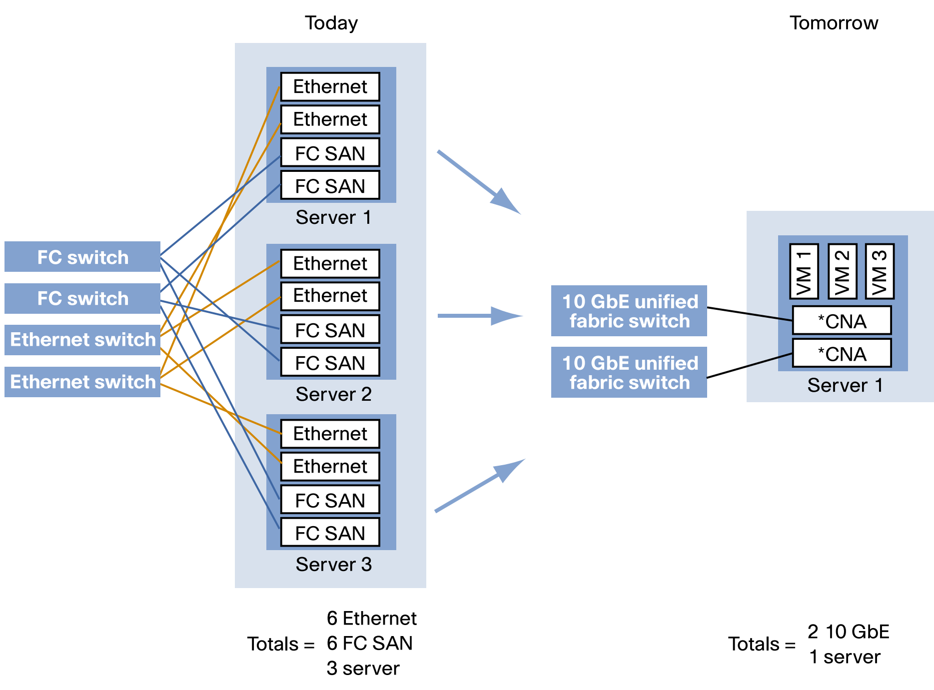

*Converged Network Adapter (CNA)

Source: Forrester Consulting, 2009.

Source: Forrester Research, "How Server and Network Virtualization Make Data Centers More Dynamic" January 2009 .

When this networked storage and compute environment is then tied together with virtualized network services as depicted in the following diagram, the vision of an end-to-end virtualized environment begins to emerge. In this next-generation model, for an individual application you can create a connected set of virtual resources that allows you to consolidate your footprint, provide the end-to-end isolation many applications require, and increase application resiliency by enabling workload portability among and failover between physical devices.

Figure 7. End-to-End Virtualization

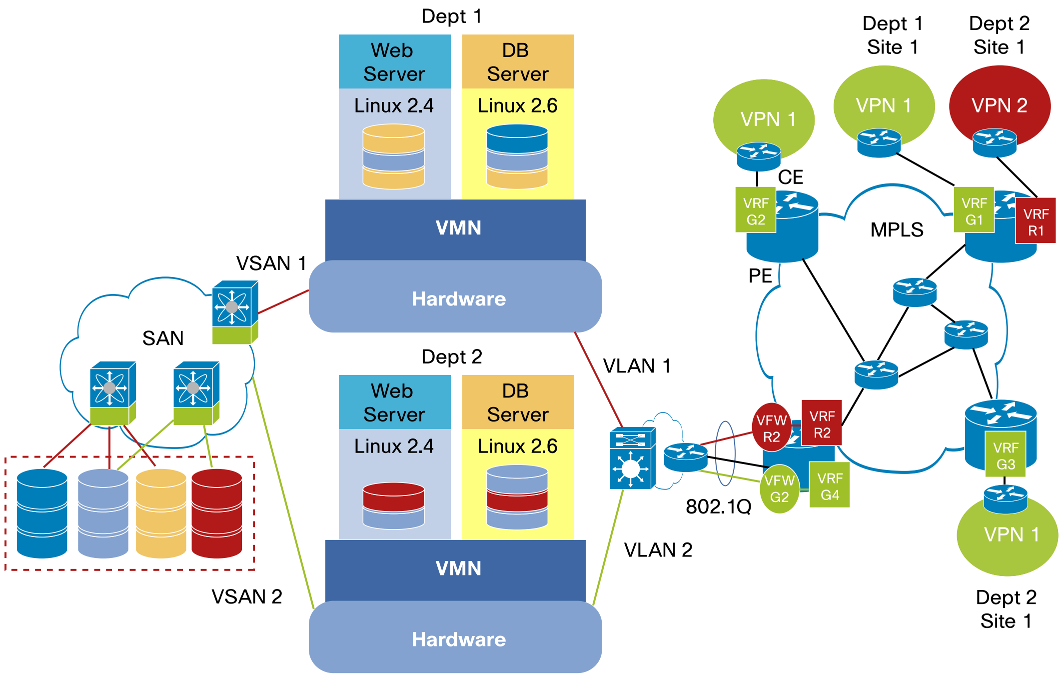

Moving from (physical):	To (virtual):
Server	**VM**
N_Port	**NPIV**
SAN	**VSAN**
Host bus adapters (HBAs) and network interface cards (NICs)	**Converged network adapter (CNA)/unified fabric**
LAN	**VLAN**
Dedicated firewall and load balancers	**Partitioned service contexts**
Per-instance configuration	**Policy-based configuration**

The challenge with operating and maintaining this set of virtualized abstractions then becomes maintaining proper correlation between each virtual resource and its associated workload. In the physical world, the connections were easier to trace. Architecturally, the move toward abstracting the workload from the server has reduced visibility for many tools traditionally used to manage and troubleshoot data center environments. Because the hypervisor abstraction effectively introduces a new access layer between the workload and the physical ports on the server, keeping track of the dependencies between VMs and their surrounding network and storage services becomes opaque from the perspective of the network and storage teams. This in turn requires new tools to consistently apply portable network and security policies to virtual workloads, particularly at enterprise scale.

The needed tools arise from the introduction of technology that can apply uniquely identifiable Layer 2 tags to packets as they leave a given VM. These tags can subsequently be identified downstream to help ensure that the proper network and security policy is matched against each VM instance regardless of current location. With this innovation the network restores the traceability between components and makes virtualization transparent. This VN-Link technology, pioneered in a joint development between Cisco and VMware, effectively reconnects the virtual server, network, and storage elements of the next-generation data center and makes end-to-end virtualization feasible in an enterprise-scale environment.

Figure 8. VN-Link Brings VM Level Granularity

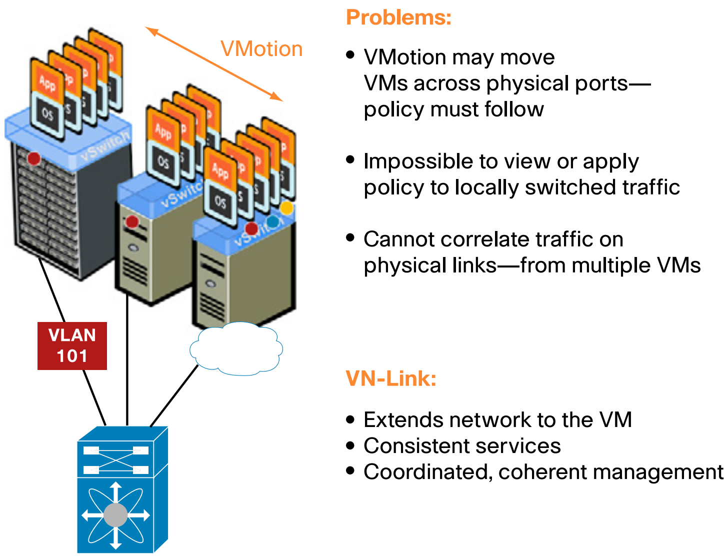

The Need for a Roadmap

Let us assume that the vision is clear, the building blocks identified, and the right mechanisms in place to allow for a comprehensive approach to data center virtualization. How best to get there? The precise answer to that question depends upon your starting point, but broadly the process has three phases as the diagram shows: consolidation, virtualization, and automation. Following the classic formula for change management, the first phase focuses upon preparation for change, followed by the introduction of change during the second phase, and finalizing and mastering the change in the third phase.

Figure 9. Virtualization is Part of a Journey

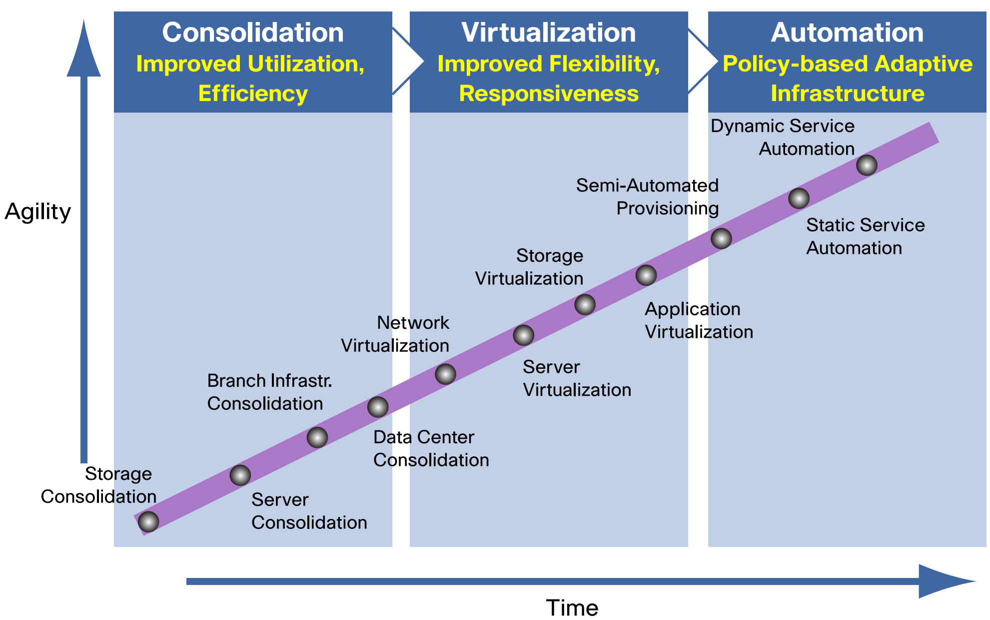

In phase one, the architect's role is to properly discover and capture the current state of the environment, map current dependencies between applications and their infrastructure, and look for immediate as well as longer-term opportunities for infrastructure consolidation. Phase two is where new technologies and techniques for virtualizing the environment can be introduced, piloted, and prepared for scaling in production. The final phase is where operations required to maintain and operate the new virtualized data center are incorporated into existing runbooks and progressively automated to increasingly deliver the sought-after infrastructure agility from end to end.

Planning for such a migration requires the collaboration of multiple teams: not only the server administrators and those responsible for managing the hypervisor environment need to be engaged, but also representatives from the storage, networking, security, and key application teams to ensure a smooth transition that optimizes new capabilities introduced from each domain. Bringing together these teams and getting them to look holistically at the opportunity ahead is where you as an architect have the optimum chance to make a significant impact on the future of your data center.

The Cisco Data Center Advanced Services team, which assists customers who require help moving to a virtualized data center, has assembled a practical set of steps to success along this path as shown below:

Figure 10. Six Steps to Success

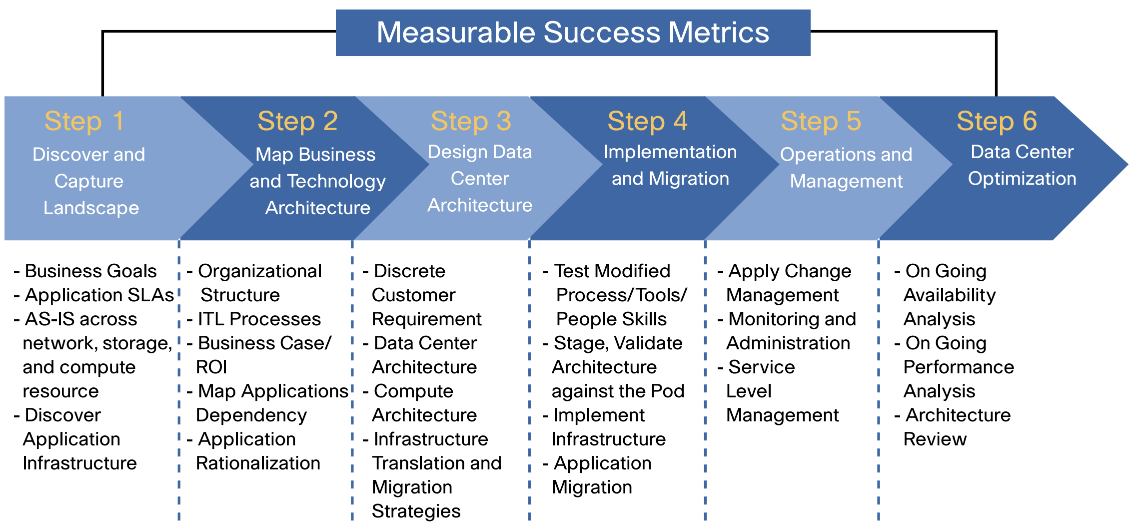

Following such a process should help guide you in assembling a custom plan for implementation. Chances are you're not starting with a greenfield environment. As with any significant change, invest the time up front in analyzing the existing context and understanding the goals of your particular organization. Rely on existing IT service catalogs if you have them. Resist the temptation to quickly move from initial testing to rapid deployment, which could easily result in "virtual server sprawl." Create a business case containing sufficient detail to warrant your proposed investments. Develop a conceptual architecture to help sell your case to key stakeholders. Make sure to consider each technology in turn: server, storage, network, and security, as well as relevant facility and finance aspects of what is needed to make the new data center model a success.

Figure 11. Elements to Consider in Your Plan

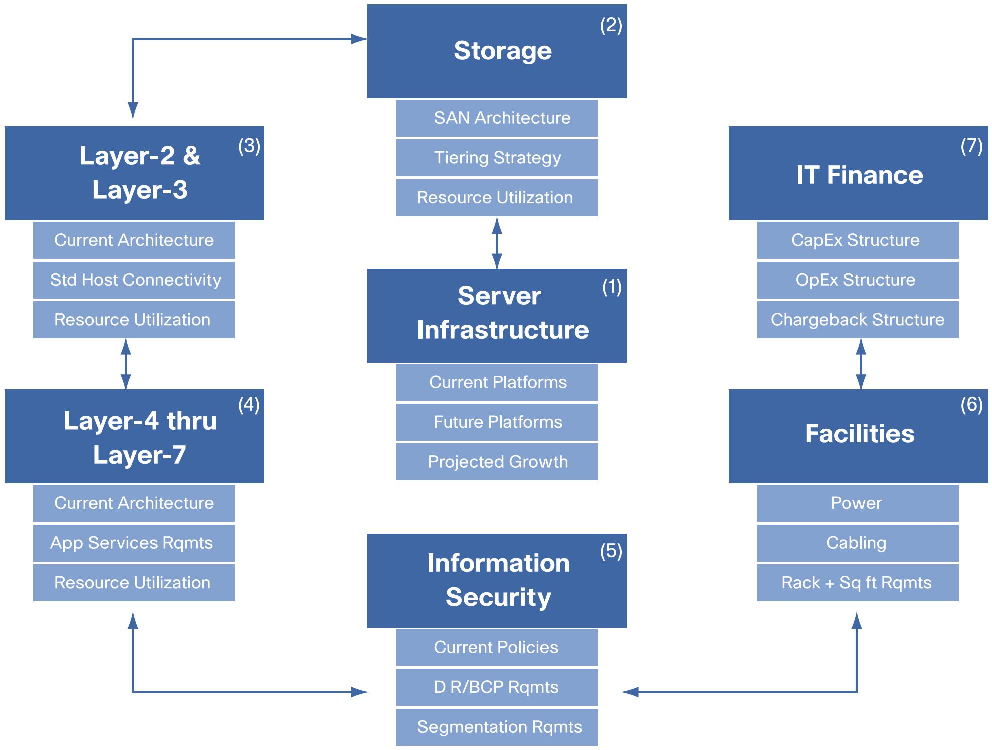

Establish metrics that will ensure you're able to measure your progress and communicate your results. From there, begin your more detailed architectural analysis and develop plans for operationally validating your proposed architecture. Next comes the moment you've been building up to: putting your architecture into action with a phased rollout from staging to production. Finally, make sure to periodically revisit what you have achieved to look for opportunities for further optimization. This strategy has been proven to accelerate the adoption of virtualization by focusing on three key areas: the development of an end-to-end data center architecture, quantification of the value such an approach brings, and the establishment of operational procedures necessary to sustain such a transformation.

The Cisco SONA Framework

One additional artifact that can be helpful in determining which network-based services (Layer 2 through 7) to incorporate into your next-generation plan is Cisco Service-Oriented Network Architecture or SONA. Cisco SONA is a multilayered framework of infrastructure elements and core services that provide reusable capabilities to applications. The framework was constructed as a way to catalog and illustrate how network services can be employed by enterprise applications to achieve desired business outcomes.

Figure 12. Cisco SONA Framework

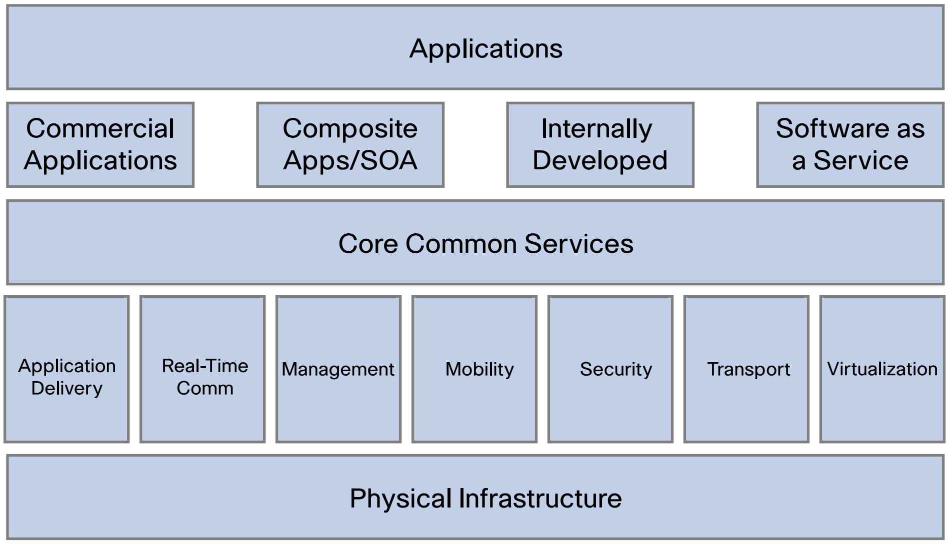

As the illustration shows, Cisco SONA's Physical Infrastructure Layer, which consists of modular, connected infrastructure elements, instantiates the capabilities provided by the Core Common Services Layer.

The Core Common Services Layer comprises a library of network-based service categories working together to create functionality that can be used by the Applications Layer. These service categories include:

- **Real-Time Communication Services** that offer session and media management capabilities, contact center services, and presence functions
- **Mobility Services** that provide location information as well as other context-aware functionality
- **Application Delivery Services** that use application awareness to optimize performance
- **Security Services** that help protect the infrastructure, data, and application layers from threats and also offer access-control and identify functions
- **Management Services** that offer configuration and reporting capabilities
- **Virtualization Services** that deliver abstraction between physical and functional elements in the infrastructure
- **Transport Services** that help with resource allocation and deliver on the overall QoS requirements of the application, as well as providing routing and topology functions

The VPN, VLAN, VSAN, VSS, I/O, Load Balancing, and Service Partitioning capabilities of the network described earlier in this paper all fall within the Virtualization Services category of the Cisco SONA framework. Certain capabilities from other categories (many of which can be segmented into virtualized service contexts) are also important to consider when planning the full list of network-based services of which a virtualized environment can take advantage. Security, transport, and management services are the foremost of these. As one example, the recently announced Cisco EnergyWise power management service that allows administrators to monitor and manage energy consumption of IP devices can help further the efficiency objectives of this next-generation environment. More information about Cisco SONA and its related set of Design Zone infrastructure and solution building blocks is available at http://www.cisco.com/go/architecture.

The Cloud Connection

What about clouds? Is cloud computing a threat to your plans to transform your own data center architecture? There are a number of reasons why this is not likely. The appeal of cloud services (whether of a software-as-a-service, platform-as-a-service, or infrastructure-as-a-service type) is that they are in theory accessible from anywhere, they can scale dynamically, and they can be paid for incrementally with no additional upfront capital expense. In a sense, with the next-generation virtualized data center, you are laying the groundwork for an internal or "private" cloud with the same advantages. Providing an environment for segmented, rapid application provisioning is one of your fundamental goals. With the right network and security architecture in place, the application services you provide internally should scale and be accessible from any location, enabling a borderless enterprise.

As long as you incorporate an internal chargeback model into your planning so that application developers and owners can be appropriately billed for usage of a common, shared infrastructure, it becomes a sourcing decision whether to host a particular application internally in your virtualized data center or in a public cloud provider environment. If you already have such a financial model in place, you need to determine how to appropriately meter and bill for usage of internal virtualized resources. If not, constructing your own internal cloud and implementing a chargeback model will in many ways prepare the organization to become a better external cloud consumer in the future, because that is how they will need to plan and budget for projects destined for the external cloud

As an enterprise architect, you can then take on the important role of helping your organization decide whether to insource or outsource a given solution to meet business objectives. Your particular enterprise will have its own unique set of core competencies (both in IT and business), industry regulations to follow, level of desire to have complete control over the operating environment, and price point at which you are able to deliver services. These should go into the creation of a decision matrix that will allow you to guide your business counterparts in making informed investment decisions. If rapid results are what matter most to them, then turning your data center into an internal cloud based upon a virtualized infrastructure can change your discussions about IT service delivery into business-oriented conversations about issues such as risk tolerance or core versus context, rather than the same old conversation about how technical interdependencies will add extra time to their desired schedule. Isn't that what they've wanted all along? Looking forward, as proposed standards such as the Open Virtualization Format (OVF) mature, the opportunity for portability between private and open public clouds will increase and your internal experience will continue to pay dividends in prudently guiding IT solution investments.

You should be able to claim victory on another front as well: in moving toward a data center architecture that can rapidly provide homes for new applications, you complement the path taken by SOA and Web 2.0 applications initiatives in the search for greater business agility. Rapid application development solves only one-half of this equation: agile infrastructure delivery is also required. With the proper planning, the next-generation data center can ensure that each application service has the level of security, availability, and scalability it requires with the help of network-based services. By pairing the loosely coupled, abstracted application services offered by SOA with the virtualized abstracted infrastructure services of the next-generation data center, your ability to rapidly deliver solutions to the business can take an important leap forward.

Organizational implications

Apart from the technology, the roles and processes of the people supporting the data center are likely to change in a virtualized environment. While as an architect you may have the power to bring the server, storage, network, security, and applications teams together to draft an initial plan, the consideration of people issues shouldn't stop there. Continued collaboration across IT departmental boundaries will be necessary to support the virtualized data center as new operational best practices are developed. Certain political boundaries may need to be redrawn. Make sure you prepare your organization for change on this front as well.

You can position this transition as an opportunity to move more rapidly to an ITIL-like model, transitioning from an organizational structure optimized to support individual technologies to one optimized for integrated service delivery. You can shift the conversations with application stakeholders to be about delivering capabilities (such as isolation, portability, and rapid instance provisioning) rather than hardware instances. Infuse the same way of thinking into conversations with your various infrastructure teams.

- Point out how such a mindset will ultimately position IT to move into better alignment with the way business clients view the value provided by IT.
- Explain how adopting a holistic plan for data center virtualization expedites such a transition.
- Finally, emphasize how you envision such changes taking place incrementally so that natural human resistance to change doesn't become your biggest impediment.

One such incremental approach is suggested in the following figure. As an architect, use your persuasive talents to work through the issues that will naturally arise as you bring fundamental changes to the conceptual models people have grown comfortable with in the data center. And be prepared to allot at least as much time on this as on technology planning and implementation.

Figure 13. Optimizing People, Processes, Technology

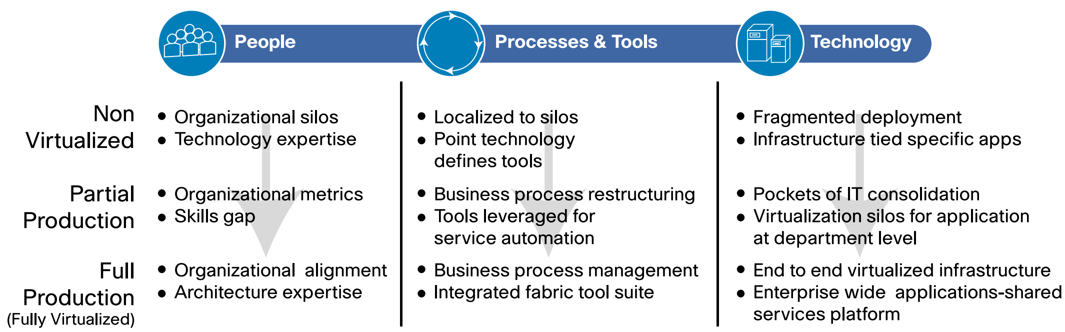

Conclusion

Virtualization has been with us in IT for many years, but the introduction of server virtualization has caused a disruption in the data center as fundamental as the introduction of client-server technology. This should be considered an opportunity to change your current data center architecture to achieve the improved efficiency, resiliency, and agility that server virtualization introduces. To optimize the delivery of the benefits that virtualization promises, as an architect you should approach the problem holistically: look beyond the hypervisor and consider the entire system of services connected to the virtual workload via the network. As the only common element that connects and enables all components of the IT infrastructure, the network is uniquely qualified to enable virtualization at scale. As Stephen Herrod, Chief Technology Officer of VMware, has suggested, "I think of the network as the glue that holds together the new virtual infrastructure of data centers."

Take time to consider how to incorporate network innovations based on open standards into your planning to allow for the virtualization of storage access and the virtualization of network services themselves. Understand how the network provides a critical mechanism for linking these through policy to the virtual machines in which your applications will increasingly reside.

What will this newly unified computing environment mean to your business with its ability to rapidly deliver new execution environments, while lowering operational costs and providing greater redundancy? Anticipate how this will change your solution planning and budgeting cycles for the future, and consider what this change will mean to the stakeholders involved.

Be active in bringing together the appropriate teams in your organization to plan the transition to this next-generation environment, leveraging external expertise where required, and help new operational procedures and organizational boundaries evolve as needed to ensure operational integrity of this new architecture.

Virtualization will be at the heart of next-generation data centers. Architects need to take an active role in planning for this transformation. As with all transformations, consider it a journey that will surely evolve further as more innovations are introduced by different vendors, both for on-premises solutions and for the cloud.

Ensure that along the way your focus remains on the objective of improving IT's responsiveness to new business requirements, part of your ongoing pursuit of making sure that IT is a strategic enabler to your business.

For more information, visit http://www.cisco.com/go/architecture.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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.

Printed in USA C11-525495-00 04/09