White Paper

Cisco Unified Data Center: Redefining the Economics of Business Continuity

What You Will Learn

IT departments tend to focus on trying to improve the economics of their ongoing operations: reducing the cost of maintaining existing infrastructure, improving IT staff productivity, and reducing the cost of implementing new IT projects. The IT event that may have the most significant economic impact on a business, however, is an event that stops ongoing operations—an IT incident or natural disaster that takes down the IT infrastructure so that the organization cannot function. That is why business-continuity strategies are a top priority for IT executives.

Traditionally, business-continuity solutions have relied on expensive replication and duplication of IT resources, further challenging already stretched IT budgets. To avoid this expense, IT departments need to optimize both sides of the business-continuity equation: reducing the likelihood and impact of outages as well as the cost of their overall business-continuity solution.

In this document, you will learn how the Cisco® Unified Data Center redefines the economics of business continuity though a simplified and intelligent architecture and innovative management, and by creating a fundamentally superior and more cost-effective solution than the traditional backup and recovery approach.

What Is at Stake: The Increasing Cost of IT Outages

IT outages and the time required to get processes back to normal have numerous economic repercussions. First, there are the obvious costs of lost revenue, the strain on IT productivity, and wasted employee time. There are also hard-to-quantify but substantial hidden costs such as damage to customer relationships or to an organization's reputation. Two years ago, Aberdeen Group calculated that the average cost of an IT outage to a business was approximately US$110,000 per hour.

The pace of business has continued to escalate since that time. Trends such as online sales, growth of the availability of revenue-producing video and mobile services, and automated business operations have increased business reliance on its IT infrastructure. This caused Aberdeen Group to revise its earlier estimate; the group now believes that these factors have helped raise the cost of an IT outage by an average of 65 percent over the 19 months between June 2010 and February 2012. The middle 50 percent of those organizations surveyed experienced an average cost of over US$400,000 in lost revenue and productivity alone.[1] For enterprises with revenue models that depend solely on the data center's ability to deliver IT and networking services to customers, such as telecommunications service providers and e-commerce companies, downtime can be particularly costly, with the highest cost of a single event topping US$1 million (more than US$11,000 per minute).

But Aberdeen also determined that if 10 percent of these organizations successfully reduce the likelihood of outages and improve recovery times, they would be able to limit the economic impact of annual downtime to just over US$3000.

[1] Aberdeen Group, February 2012: Datacenter Downtime: How Much Does It Really Cost?

However, to understand the full burden of downtime on an organization, the cost of avoiding that downtime must also be considered. In the past, some organizations have invested millions of dollars in redundant IT resources, operations personnel, and remote disaster-recovery site facilities. For organizations with limited IT budgets, the goal must be to reduce both the impact of IT outages and the cost of the business-continuity solution.

Cisco Unified Data Center provides organizations with a way to accomplish both objectives, reducing both the cost of annual downtime and the cost of avoiding downtime—fundamentally redefining the economics of business-continuity solutions.

Optimizing Your Business-Continuity Solution Equation with Cisco Unified Data Center

Less than a decade ago, most IT departments pursued a business-continuity strategy that was based on duplicate servers, redundant resources, dedicated networks, and backup sites. These strategies were inherently expensive and characterized by overcapacity and manual processes. Virtualization and cloud computing have significantly improved both the quality and the economics of business-continuity solutions in many ways: reducing hardware costs by replacing siloed applications with mobile virtual machines, facilitating clustering, and allowing organizations to perform data migration to their backup sites, achieving 99 percent intact data so that all data can be accessed. In addition, with the development of virtualized desktop solutions in which all user applications reside on servers, IT departments can restore business operations to normal more quickly because they can quickly provide employees with access to their data and applications from anywhere.

But the complexity associated with many virtualized environments can make it difficult to stay compliant with best business-continuity practices and can even lead to outages. A recent Gartner study projected that "Through 2015, 80 percent of outages impacting mission-critical services will be caused by people and process issues, and more than 50 percent of those outages will be caused by change/configuration/release integration and hand-off issues."[2]

In contrast, the Cisco Unified Data Center significantly simplifies virtualized and cloud environments. Consisting of the Cisco Unified Computing System™ (Cisco UCS®), Cisco Unified Fabric, and Cisco Unified Management, Cisco Unified Data Center both facilitates the implementation and management of virtualized and cloud computing environments and dramatically lowers the cost of your business-continuity investment.

The following sections discuss how the Cisco Unified Data Center targets every aspect of the business-continuity solution by reducing the likelihood of an outage, delivering higher availability within the data center, and increasing the speed and reliability of the transfer of workloads between data centers, all while simultaneously reducing both the capital expenditures (CapEx) and operating expenses (OpEx) associated with helping ensure high-availability IT infrastructure.

Reducing the Likelihood of an Outage While Reducing Costs

An effective business-continuity solution requires a high-availability architecture that delivers excellence in disaster recovery, with the capability to address both point and site failures, as well as the capability to quickly restore business processes to normal levels. The Cisco Unified Data Center meets these requirements cost effectively through a simplified architecture designed for greater reliability; intelligence and automation to accelerate problem identification and resolution and reduce the risk of error; and tested and proven architecture designed for stability.

To help ensure continuous business operations, IT departments must help ensure that all workloads, whether virtual or nonvirtual, can be moved quickly and easily. Often, however, they find that nonvirtual workloads require a different set of processes than virtual workloads, which increases management costs and consumes valuable

[2] Ronni J. Colville and George Spafford, Configuration Management for Virtual and Cloud Infrastructures

time. The Cisco Unified Data Center allows IT departments to move all workloads easily and efficiently using the same processes.

Using Cisco UCS service profiles, IT departments can move or rehost a physical or virtualized software environment simply by applying the service profile to a new blade server and rebooting the environment on the new server. The capability to use a similar operational model regardless of whether a workload is virtualized or nonvirtualized improves availability and lowers IT costs.

High-availability infrastructure typically requires rarely used equipment (hot spares) as backup alternatives for point failures or site failures. These hot spares often consume expensive software licenses, facility costs to house and power the equipment, and administrative time to manage those IT resources. Cisco UCS service profiles also address these costs by allowing IT departments to use a cold Cisco UCS spare.

Cisco UCS Manager abstracts server I/O characteristics into service profiles that include server identities, I/O device characteristics, and I/O device identities including MAC addresses and worldwide names (WWNs). Templates can be used to instantiate one or many servers having the same I/O characteristics but unique identifiers, allowing the system to scale on demand. In the event of a point failure, Cisco UCS Manager can apply a service profile template to the cold spare and have it created in minutes so that it aligns 100 percent with all server requirements for that application stack. Cisco UCS service profiles can then interconnect with your deployment software or scripts to automatically install the needed OS and applications.

KORE TELEMATICS

- "Since implementing Cisco UCS solutions, we've experienced zero downtime through the fault-tolerant infrastructure."

Because a software license can be aligned with the down server's service profile, IT departments can move the license to the cold server, often without having to pay for a new software license. This approach can help IT maintain software compliance and cost effectively maintain a pool of Cisco UCS servers that can act as backup spares for multiple application stacks.

IT departments gain significant CapEx and OpEx advantages in this scenario. Saving the cost of an incremental software license can be significant: at least US$75,000 for some mission-critical enterprise applications. In addition, the company benefits from lower ongoing OpEx by using a cold spare instead of a hot spare: the cost of power and cooling can represent 20 percent of the total cost of ownership (TCO) of a server.[3]

Cisco UCS also helps IT departments gain the economic benefits of faster recovery. For instance, if an organization needs to access a hot spare as a replacement for a server with dedicated storage and networking, it can take hours to turn it on and get it running. However, Cisco UCS servers can be up and processing transactions within minutes because the service profile contains all the network and storage information, so you will be restoring processes faster.

Server management automation enabled by Cisco UCS also helps improve IT productivity and the quality and completeness of your disaster-recovery processes. Cisco UCS Manager allows service profiles that define servers to be exported to XML files and reimported to other locations. Cisco UCS Central Software allows this process to be automated, so that IT departments can systematically and rapidly replicate configurations from a primary site to another disaster-recovery site. Automating this process yields improvements in IT employee productivity, and it also avoids potentially costly errors and delays in restoring business operations. IT administrators typically are focused on their primary environment, dealing with day-to-day workloads. However, if an IT staff member makes a change in the primary environment because of that workload, and if the administrator neglects to make the same

[3] The Linley Group, 2012

change in the disaster-recovery environment, misconfiguration or other problems may occur during a disaster-recovery event and prevent a service from coming back online. By programmatically copying infrastructure configurations and automatically pushing them to the disaster-recovery site, organizations can remove this responsibility from the administrator and help ensure that the disaster-recovery process is successful.

Increasing Efficiency in High-Availability Infrastructure with Network and Management Innovations

Cisco Unified Fabric and Unified Management together help reduce the likelihood of a failure and accelerate the restoration of business operations if a fault should occur.

Delivering Cost-Effective Reliability in the Data Center

In the data center, Cisco Unified Fabric delivers convergence of multiple networks that dramatically reduces the number of potential points of failure. A traditional server may have 8 or 10 cables extending from it (for example, cables for Fibre Channel, Ethernet, application management, application communications, disaster recovery, and more). Each of these cables is a potential source of site failure: any one could fail or accidentally be unplugged. Cisco Unified Fabric, however, unifies multiprotocol and multiple functional fabrics into a single pipe, reducing the average number of server cables to 2. In addition, Cisco Unified Fabric allows IT to consolidate the two storage networks—Ethernet and FibreChannel—into a single physical infrastructure. This consolidation eliminates more potential points of failure, as well as the need to procure, maintain, and manage two sets of switches that also require related CapEx (for example, costs for network interface cards [NICs], host bus adapter [HBA] ports, and networking adapters) and OpEx (for example, costs for power, cooling, and space).

Cisco Unified Fabric not only improves network reliability by reducing points of potential failure, but it acheives this improvement with no reduction in redundancy or security. Consequently, organizations benefit from higher network density so they can buy fewer high performance switches. For example, an organization may decide to buy eight smaller switches at US$40,000 each, spending US$320,000 total. Another organization may decide to buy two higher-capability switches costing US$100,000 each, for a total expenditure of US$200,000. Because of the higher density of the unified fabric, the organization with the two high-performance switches does not waste capacity, and fewer switches are easier and less costly to manage. In addition, fewer points and physical devices mean fewer items that you have to recover or replicate at a disaster-recovery site in the event of a disaster.

In the data center, Cisco Unified Fabric also has a wide range of reliability features designed to prevent faults. For instance, all Cisco Unified Fabric switches, including the Cisco Nexus® and Cisco MDS 9000 families, all run the Cisco NX-OS Software operating system. This Linux-based, modular operating system is designed so that individual portions of the OS or single protocols can fail without taking down the switch. If an application hangs, the OS implements an automatic shutdown and restart process that maintains separation of the control plane and the data plane, so user traffic can continue unaffected without any disruption to the data being processed.

Cost-effective reliability is also designed into the Cisco Unified Fabric hardware. For instance, Cisco Nexus 7000 Series Switches include features that help ensure that they are highly available, so they can tolerate a supervisor failure or power supply failure or a software upgrade without the need to take down the switch.

In addition, Cisco Unified Fabric includes technologies that cost-effectively help ensure business continuity in the event that an Ethernet link goes down. IT departments have traditionally tried to avoid this problem by connecting the active switch to a backup switch in a spanning tree environment. The economic downsides of this strategy are the CapEx and OpEx associated with the backup switch and the time lost moving traffic to the backup switch while the application is not available to users.

To avoid both of these problems, Cisco implements Cisco FabricPath, another member of the Cisco Unified Fabric product line. Cisco FabricPath can connect up to 16 different Cisco Nexus switches that are all forwarding data. This higher number of lanes keeps an organization's data flowing. From a business-continuity perspective, the organization also has 15 potential lanes to which data can be offloaded if one switch fails.

Simplifying and Automating Ethernet and Storage Network Management

Cisco Unified Management plays a critical role in preventing failures and reducing recovery time. One element of Cisco Unified Management is Cisco Prime™ Data Center Network Manager (DCNM). Unlike traditional environments, in which different element managers are used for storage network and other networking products, Cisco Prime DCNM unifies the management of both Ethernet and storage networks with a single dashboard. The dashboard enables network and storage administrators to troubleshoot health and performance concerns across the entire range of Cisco NX-OS platforms, including the Cisco Nexus and Cisco MDS 9000 families, regardless of protocol type, such as Fibre Channel, Fibre Channel over Ethernet (FCoE), Ethernet, IBM Fibre Connection (FICON), and Small Computer System over IP (iSCSI).

Cisco Prime DCNM delivers three main benefits that help organizations implement an effective business-continuity strategy while reducing TCO:

- Better visibility for higher reliability
- Intelligent troubleshooting and provisioning to prevent errors and accelerate service restoration
- Proactive planning to dramatically reduce the chance of outages while improving resource utilization

Single Dashboard for Multiple Products and Multiple Data Centers

In the past, IT departments had to accept the productivity burden entailed in using a different element management environment for every different storage network and network switch in the data center, or they had to pay organizations to integrate these environments. Either way, they paid a significant economic cost to manage their networks. With Cisco Prime DCNM, organizations now have one cost-effective software management tool that provides them with superior, end-to-end network visibility using a single management dashboard.

Cisco Prime DCNM is designed for virtualized environments, extending network visibility to virtual infrastructure by mapping the entire path from the virtual machine and switch to the physical storage and Ethernet networks. The virtual machine–aware dashboard displays performance charts, path attributes, topology, path alerts, and information about the use of virtual machines and virtual hosts. The increased visibility into virtualized infrastructure helps IT administrators locate performance anomalies that may be causing service degradation.

KEMET

- "Before, with our third-party provider, we could only recover a small subset of servers in 72 hours. With UCS, Kemet can recover all servers within eight hours if a disaster were to happen, and many of them in less than an hour."

Through the concept of domain dashboards, Cisco Prime DCNM creates contextual views of the host, switch, and storage infrastructure. These views help IT administrators identify individual components and display their health and performance attributes. In addition, Cisco Prime DCNM provides a search engine that launches these dashboards in the context of the search attributes, so IT administrators can determine the source of a network problem in seconds, thus reducing outage time or preventing an outage from occurring in the first place.

In addition, Cisco Prime DCNM is designed to scale to meet the needs of large and distributed data center deployments. Cisco Prime DCNM maintains a consolidated view of the entire network. As a result, network

operations centers (NOCs) can now have a single dashboard for monitoring network alerts and events across multiple data centers.

Intelligent Troubleshooting and Provisioning for Greater Network Reliability and Faster Recovery

Cisco Prime DCNM is designed to help IT administrators identify a problem before it happens, and to respond more quickly if a problem does occur. Through comprehensive event management, Cisco Prime DCNM consolidates all the events on the network, such as threshold breaches and switch and path alerts, and tracks configuration changes by individual network and storage administrators that may have caused outages or performance degradation. In addition, the event-forwarding mechanism notifies, in real time, the appropriate email address, pager, or third-party management application when predefined thresholds are breached or the health of the network is affected.

When combined with the intelligence in Cisco products, Cisco Prime DCNM troubleshooting can greatly reduce the impact of a fault and its recovery time. For instance, a customer who has installed Cisco UCS, Cisco Nexus Family, and Cisco MDS 9000 Family platforms can rely on these products to send ongoing data to Cisco Prime DCNM, while Cisco Prime DCNM monitors them and their backup platforms. If the customer needs to restore a switch, Cisco Prime DCNM has all the status and event information collected to restore the switch quickly. If a port goes down, Cisco Prime DCNM sends an alert, and intelligence in the network fabric immediately recognizes that the link is down and automatically passes the data to another port to maintain availability until the failed port is fixed. Cisco Prime DCNM also has the intelligence to notify administrators when they no longer have redundancy in a SAN environment, so they can develop a backup plan to create that redundancy and help ensure business continuity.

Cisco Prime DCNM also reduces the opportunities for costly administrator errors by providing wizards and predefined scripts for deploying configuration changes to Ethernet and storage networks. It can guide IT staff through the provisioning of hosts, Inter-Switch Links (ISLs) and trunks, and storage ports, regardless of whether the ports reside on the physical or virtual Cisco switch. In addition, Cisco Prime DCNM provides predefined template scripts for common tasks, which administrators also can customize.

Proactive Planning to Prevent Outages and Service Degradation

By analyzing performance history and projecting future bandwidth consumption, Cisco Prime DCNM helps IT departments understand the baseline traffic pattern for a deployed application and identify any performance anomalies that may affect customer service levels. Cisco Prime DCNM also helps IT design teams plan new deployments and redesign existing deployments based on peak utilization information collected as far back as a year or more. By monitoring historic performance, the design team can optimize switch consolidation activities by choosing modules with higher port densities based on the performance of applications, thus reducing the need to deploy more switches, reducing power use, and leaving more tile space in the data center for computing and storage resources.

Understanding resource consumption and helping ensure effective capacity planning is essential to helping ensure predictable application performance and cost-effective use of resources. To help prevent the Ethernet or storage network from becoming a bottleneck and causing service degradation or an outage, Cisco Prime DCNM helps track port use by tier. Through trending, it can predict when an individual tier will be consumed, providing the capacity management team with enough time to order more ports or switches. Cisco Prime DCNM can detect unused capacity and orphan ports and reclaim those that have not been used, helping IT departments make better use of their existing resources.

Redefining the Economics of Business Continuity Between Data Centers

Traditionally, every business-continuity strategy to ensure ongoing operations in the event of a site failure has either relied on a hot-spare infrastructure located at another site or paid a managed service provider to deliver backup and recovery services. Often, IT departments that use a hot-spare site deploy only the most critical IT assets at the remote site, such as critical servers and, depending on the business, replicated storage or parts of storage, to limit the expense. Even so, the economic costs of a remote backup site are significant, including facilities costs (land, building, lighting, air conditioning, power, etc.), the cost of site personnel, and the CapEx and service contract costs of unused hot spare servers, storage, the software residing on them, and the networking required for data transmission. Most of the time, these resources remain unused—available to the business only in the event of a disaster or site failure.

The virtualization and cloud computing capabilities of the Cisco Unified Data Center give organizations the opportunity to eliminate their unused backup data centers. Instead, with the Cisco Unified Data Center, they can create multiple data centers that can be operated as a single, logical data center. IT departments then can balance workloads across these multiple data centers while using the snapshot feature of VMware vSphere to continually send backup images of the virtual machines to all the data centers, including the disaster-recovery site. Now, instead of maintaining a site with unused resources, organizations can use the disaster-recovery site to process nonproduction workloads, adding the value of that site to the business.

HAMILTON/CLERMONT COOPERATIVE ASSOCIATION

- "We now can provide a business continuity environment with exponentially improved processing power and speed."

Because the production data centers are operated as a single, logical data center, and each is continually sending snapshots to the others, if one part of the logical data center goes down, another can pick up its workload and keep processing without any significant impact on the business. If, however, a site failure occurs and the other data centers are unable to keep up with the workload, the disaster-recovery site can cease processing its nonproduction workload and process the workload from the failed site.

Not only does this approach fundamentally change the nature of the business-continuity strategy, but it also fundamentally changes the business-continuity economic equation. IT is freed from the expense of maintaining an unused backup data center, and the cost of maintaining a high-availability infrastructure can be dramatically decreased.

To plan a business-continuity strategy based on this capability, IT needs to separate all data center workloads into tiers: production and nonproduction. If disaster procedures are needed, IT must be able to move all production workloads to the designated disaster-recovery site and take down the nonproduction workloads until operations normalize.

To build the technology foundation for this strategy, organizations need to implement three main technologies: LAN extension, SAN extension, and path optimization.

LAN extension allows virtual machines and their policies to move efficiently from one data center to another. SAN extension helps ensure that there is no loss of data and that the data stores are available in the other data centers. Path optimization lets the network resources know that the workload has moved to a new location.

LAN Extension with Cisco Unified Data Center

Normally, when organizations need to move a virtual machine within their own VLAN, they cannot cross Layer 2 boundaries. To move a virtual machine from one data center to another over the Internet—which is mandatory in a business-continuity strategy—you must be able to cross boundaries.

MIAMI CHILDREN'S HOSPITAL

- "These (Cisco Nexus 7000 Series) switches also make executing backups more efficient. We've been able to reduce backup time from two hours to one hour, a 50% time savings."

In the past, mechanisms for the extension of Layer 2 connectivity could not efficiently address connectivity and independence requirements. Some of the challenges include:

- **Fate sharing:** The extension of Layer 2 domains across multiple data centers can cause the data centers to share failures that would normally have been isolated when interconnecting data centers over an IP network. These failures propagate freely over the open Layer 2 flood domain. A solution that provides Layer 2 connectivity yet restricts the reach of the flood domain is needed to contain failures and preserve the resiliency achieved through the use of multiple data centers.
- **Complex operations:** VPNs were not designed to extend Layer 2 domains at the scale needed for data center operations. Although they can provide extended Layer 2 connectivity across data centers, drawbacks include the mix of complex protocols, distributed provisioning, and an operationally intensive hierarchical scaling model. A simple overlay protocol with built-in capabilities and point-to-cloud provisioning is needed to reduce the cost of providing this connectivity.
- **Bandwidth utilization:** When extending Layer 2 domains across data centers, the use of available bandwidth between data centers must be optimized to obtain the highest connectivity at the lowest cost. Balancing the load across all available paths while providing resilient connectivity between the data center and the transport network requires added intelligence above and beyond that available in traditional Ethernet switching and Layer 2 VPNs.

The Cisco Unified Data Center provides customers with the Cisco Overlay Transport Virtualization (OTV) feature to implement LAN extension capabilities. By using Cisco OTV on Cisco Nexus 7000 Series Switches, organizations can implement LAN extension by putting a Layer 2 overlay over an IP network and allowing IT to extend VLANs over Layer 3 networks. Consequently, if an organization is experiencing a problem in one data center that may require it to be shut down, IT can move the virtual machines from that data center to another data center quickly and efficiently because the virtual machines remain in the Layer 2 domain; there is no need to change or reconfigure the existing network design.

SAN Extension

An application server writes data to the local storage system, but IT departments need a SAN extension solution to help ensure that the data also is copied to a different storage array in another data center. Typically, each storage vendor has its own methodologies for copying and replicating data between data centers.

Cisco provides IT departments with a range of options for implementing SAN extension: the use of native Fibre Channel and FCoE protocols and the capability to encapsulate Fibre Channel and FCoE in an IP packet and transport it over the IP network. In addition, Cisco can provide advantages that improve SAN extension performance, compress data to optimize links, and encrypt data to enhance security.

Path Optimization

In some cases, organizations have advance notice of an impending disaster (for example, before a hurricane or blizzard). To avoid the negative ramifications of a potential site failure, IT needs to be able to move the most important production workloads from one data center to another quickly and without errors. In the past, VMware vSphere virtualization software provided organizations with only two options for moving a workload, neither of which provided the necessary mobility to help ensure successful transport of mission-critical workloads in the event of a site failure. Using one option, VMware vMotion movement could be accomplished only in data centers that were in close proximity. To avoid problems related to a natural disaster, however, it is obviously unwise to rely solely on a nearby site in a backup strategy. Using the second option, for data centers that are far apart, IT must shut down VMware vSphere Site Recovery Manager, move it to a new VMware vSphere cluster, and then reconfigure the application with a new IP address so it can work in the new network. However, the more applications and workloads that a company needs to move, the more time is required for IT staff to perform this reconfiguration and the greater the risk of errors—just when time and the restoration of business operations matter most.

The Cisco Location Identifier Separation Protocol (LISP) enables the multihoming capability required for the portability of IP networks so that workloads can be moved between data centers. It allows organizations to move a workload and keep the workload's IP address the same, making workload movement easier, safer, and faster. By using Cisco LISP Virtual Machine Mobility (VM-Mobility), organizations can deploy virtual machines and workloads anywhere, regardless of their IP addresses, and can help ensure their quick movement across data center racks and rows, to separate locations, and globally across organizations.

Cisco LISP VM-Mobility helps guarantee optimal routing to the desired IP endpoint, regardless of its location. The location flexibility provided by Cisco LISP VM-Mobility can help IT departments streamline disaster-recovery procedures and save precious time when it matters most.

Implementing a Tested and Proven Architecture for Business Continuity

If your business-continuity strategy relies on multiple data centers that are acting as shared infrastructures, it is essential to design secure separation into the infrastructure. If a problem should occur in one application, secure separation built into the architecture will prevent the problem from affecting other applications. Secure separation allows you to segment separate tiers of applications, different tenants, and different use cases on the network, helping ensure, for instance, that development environments are separated from back-office applications or a pool of virtual desktops. The Cisco Unified Fabric's Cisco Nexus switches and Cisco Unified Network Services enable virtualization that allows secure separation of users, containers, and applications in your data center infrastructure.

> "Organizations struggle to afford a separate infrastructure for disaster recovery. Cisco IT realized we could take advantage of virtualization and automation to use the same infrastructure for disaster recovery and nonproduction applications."
>
> —Cisco IT

To help ensure that your data center architecture is designed for optimal availability and stability, enabling secure separation and using the cost-effective business-continuity strategy discussed here, organizations can use a tested and proven detailed design known as the Cisco Virtualized Multiservice Data Center (VMDC). Cisco VMDC is tested and validated reference architecture for the Cisco Unified Data Center. It provides a set of guidelines and best practices for the creation and deployment of a scalable, secure, and resilient infrastructure in the data center. The Cisco VMDC architecture demonstrates how to bring the latest Cisco routing and switching technologies, network services, data center and cloud security, automation, and integrated solutions together with those of Cisco's ecosystem of partners to develop a trusted approach to data center transformation.

Conclusion: Optimizing the Business-Continuity Equation

Cisco Unified Data Center innovations give IT departments the power to rewrite the economics of business continuity. In the past, the main way to increase infrastructure availability was to increase the amount of backup resources: higher availability meant higher costs.

Today, the Cisco Unified Data Center allows IT to get the best of both worlds: higher availability and dramatically lower technology, management, and facility costs while ensuring continuous operations. Organizations gain increased business agility and customer satisfaction while increasing their technology return on investment (ROI).

In a time when loss of availability has a direct impact on business profitability, organizations that can optimize both sides of the business-continuity equation gain a tremendous competitive advantage. They obtain benefits that increase profitability—such as greater business agility, customer satisfaction, and employee productivity—and they also increase their technology ROI.

For More Information

- http://www.cisco.com/go/unifieddatacenter
- http://www.cisco.com/go/unifiedmanagement

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 C11-726816-00 03/13