Server Virtualization: Branching Out of the Data Center

Improving Branch Office Infrastructure Efficiency

Overview

Multisite organizations are reducing the number of servers in their branch offices by moving applications to the data center. Yet, they continue to place a few essential applications locally because of performance, survivability, or compliance requirements. By making use of server virtualization, these lean branch offices can increase utilization of local infrastructure, reduce application downtime and recovery time, and shorten the time needed to deploy local applications while lowering infrastructure and operating costs.

Introduction

Some technologies, like some seeds, take longer to germinate than others. And even after they do, market conditions dictate how far and how fast they grow. Virtualization has had such a long period of dormancy, finally taking firm root in the late 1990s when VMware made it available for Intel x86 architecture (x86) servers. By 2009, according to Gartner, 16 percent of all x86 data center server workloads were running on virtualized hardware. Gartner predicts that by the end of 2012, this number will grow to 50 percent[1]

The rapid growth of server virtualization outside the corporate data center is another indicator of the technology's appeal. Cloud-based services, virtual desktops, and virtualized disaster recovery sites are all made possible by server virtualization. Increasingly, the virtualization technology is also finding its way into the branch office.

Server consolidation, energy savings, and rapid server provisioning are the main benefits of bringing virtualization into the branch office. However, to get the most value, the use cases for the technology must reflect the distinct requirements of remote locations. The goal of this document is to (1) compare and contrast data center and branch-office environments in the context of server virtualization, (2) describe how server virtualization addresses specific branch-office challenges, (3) explain how the technology unlocks new capabilities in the branch office, and (4) provide a brief overview of a joint VMware and Cisco® server virtualization solution for the branch office.

From the Data Center to the Branch Office

Economies of scale are achieved by packing a large number of x86 servers into a small number of large data centers. The resulting per-server cost reduction has been the primary driver behind data center consolidation. But large data centers are more complex to manage, consume significant amounts of energy, and require sophisticated capacity planning to account for unexpected loads. Server virtualization addresses all these challenges by inserting a thin layer of software (hypervisor) between the server hardware and the operating system. The hypervisor provides virtual hardware containers for hosting applications and operating systems (virtual servers) as shown in Figure 1. Virtual servers can be provisioned more quickly than their physical counterparts; require less space, power, and cooling than multiple physical servers; and can be cloned, moved, or clustered on demand and without service interruption.

[1] http://www.gartner.com/it/page.jsp?id=1211813

Figure 1. x86 Server Virtualization Architecture

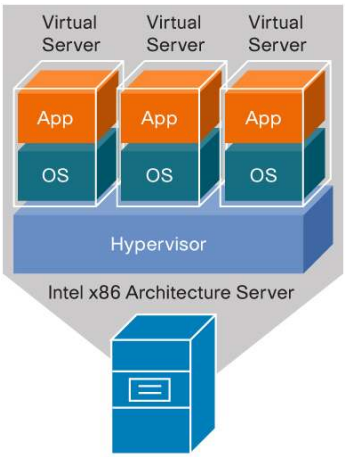

A hypervisor imposes the following set of attributes on virtual servers that traditional x86 servers lack:

- **Isolation:** A virtual server is confined to one container and unaware of other virtual servers.
- **Multiplicity:** A virtual server shares hardware concurrently with other virtual servers.
- **Abstraction:** A virtual server is hardware independent and can run on various platforms.
- **Encapsulation:** A virtual server stores its complete point-in-time execution state in a file.

These fundamental attributes are the building blocks of the higher-level capabilities that server virtualization introduces into the data center (Figure 2).

Figure 2. Fundamental Attributes of Virtualized Server Platform

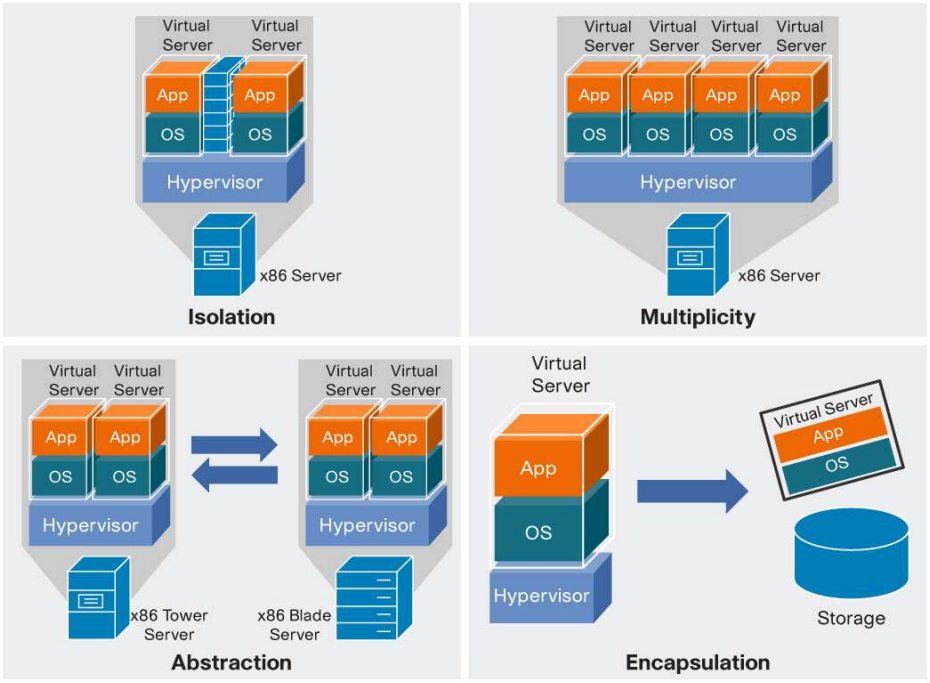

The effects of data center consolidation have spread across the entire IT infrastructure. In the branch office, consolidated data centers and WAN optimization technologies have reduced the need for local servers. However, WAN link limitations (speed and reliability) prevent full server centralization for a majority of the branch offices. Therefore, a small number of servers remain deployed locally in the now leaner branch offices.

This branch-office server infrastructure is dramatically different from a typical data center. Table 1 presents important differences between the data center and the lean branch-office server infrastructures.

Table 1. Typical Data Center and Branch-Office Infrastructure Characteristics

	Consolidated Data Center	Lean Branch Office
Number of x86 servers	Large (tens to thousands)	Small (1 to 2)
Server type	Midrange to high range	Low range
CPU type	Multisocket, many cores	Single and dual socket and a few cores
Memory capacity	Large (tens to hundreds of GB)	Small (4 to 16 GB)
Primary storage type	SAN	DAS
Storage capacity	Large (hundreds to thousands of TB)	Limited (500 GB to 2 TB)
Server load	Dynamic and fluctuating	Static and predictable
Number of applications	Large (tens to hundreds)	Small (6 to 8)
New application rollout	Frequent	Infrequent
IT management	Local	Remote
Server uptime	Critical	Critical
Crash recovery time	Critical	Critical
Application response time	Critical	Critical

A large amount of high-performance, high-capacity server hardware is available in the data center. In contrast, the number of servers in most branch offices is limited, usually to just one or two. Therefore, many server virtualization features such as virtual server mobility, dynamic provisioning, and workload balancing, which require data center-like infrastructure, are not feasible to run in a typical branch office. More importantly, the lack of attractive return on investment (ROI) for deploying these advanced virtualization features in the branch office is deterring investment in richer server configurations.

It may be tempting to think that features such as virtual server mobility or clustering can be delivered from the data center into the branch office over a WAN link. Although this may be possible over a high-bandwidth, low-latency link, the typical branch office today does not have the necessary bandwidth to support such operations. For example, a 4 GB Microsoft Windows Server virtual machine would take up to 6 hours to transfer on a dedicated T1 link with no other traffic. Deploying a WAN optimization solution helps alleviate many of the WAN link issues but cannot completely eliminate them.

Despite the vast differences between data center and branch-office infrastructures, there are compelling reasons and business justifications for using server virtualization in the branch office.

Addressing Current Branch-Office Challenges Through Server Virtualization

Branch-office infrastructure has always been a source of challenge for IT departments. Some of these are specific to the company's line of business, and others cut across all industries. The cost of deploying new applications to remote locations, the impact of insufficient environmental controls on equipment, security and regulatory compliance constraints, the performance of centralized applications, and support for critical uptime requirements are just a few examples. Further, the challenges increase as the number of branch offices grows. According to a 2010 study by the Internet Research Group (IRG), in the past decade the number of branch offices in the United States has grown by 21 percent.[2] Although no one product or technology can cure all of the current branch office headaches, server virtualization does address some of the most difficult issues.

[2] http://www.irg-intl.com/press_releases/2010-branch-office-PR.html

A typical lean branch office has four to six applications running on servers, such as print services, Dynamic Host Configuration Protocol (DHCP) server, Domain Name System (DNS) server, Microsoft Active Directory Domain Services, and performance-critical line-of-business applications. These applications and servers are provisioned in one of the following ways (see Figure 3):

- **Dedicated physical server for each application:** This setup isolates applications from each other for organizational, security, performance, or application availability reasons. However, it increases equipment and operating costs.
- **Multiple applications share one physical server:** This setup reduces the total cost. However it creates access control, operational efficiency, application availability, and time-to-deployment complexity.

Figure 3. Application Provisioning Options on Physical Servers

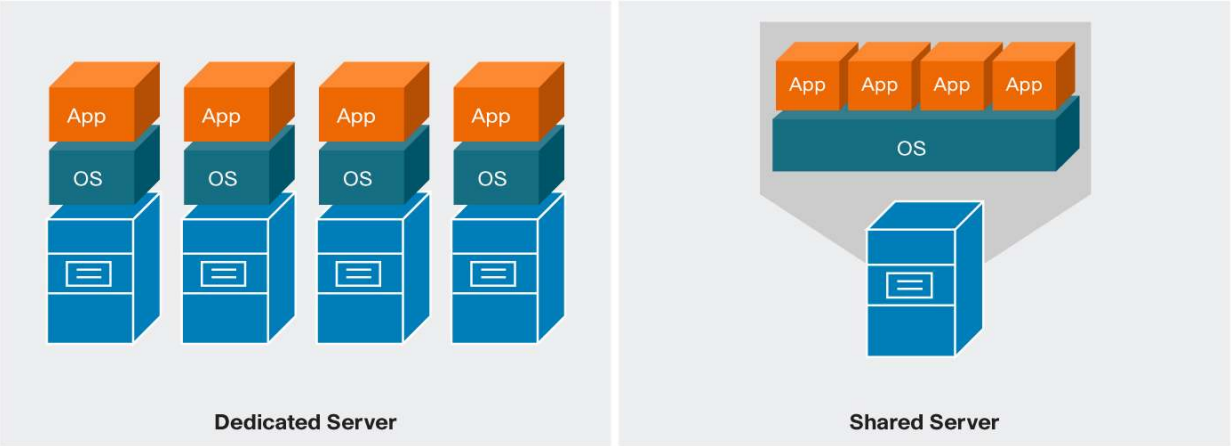

Server virtualization offers the best of both deployment scenarios and eliminates the major drawbacks. The isolation and multiplicity attributes of virtual servers described earlier enable the following provisioning options (see Figure 4):

- **Dedicated virtual server for each application all hosted on one physical server:** This setup isolates applications from each other, while lowering costs, increasing application availability, strengthening access control, and improving time-to-deployment for new branch-office applications.
- **Dedicated virtual server for each application hosted on multiple physical servers:** In addition to the benefits already listed, this setup improves application availability and disaster recovery for a small increase in cost.

Figure 4. Application Provisioning Options on Virtual Servers

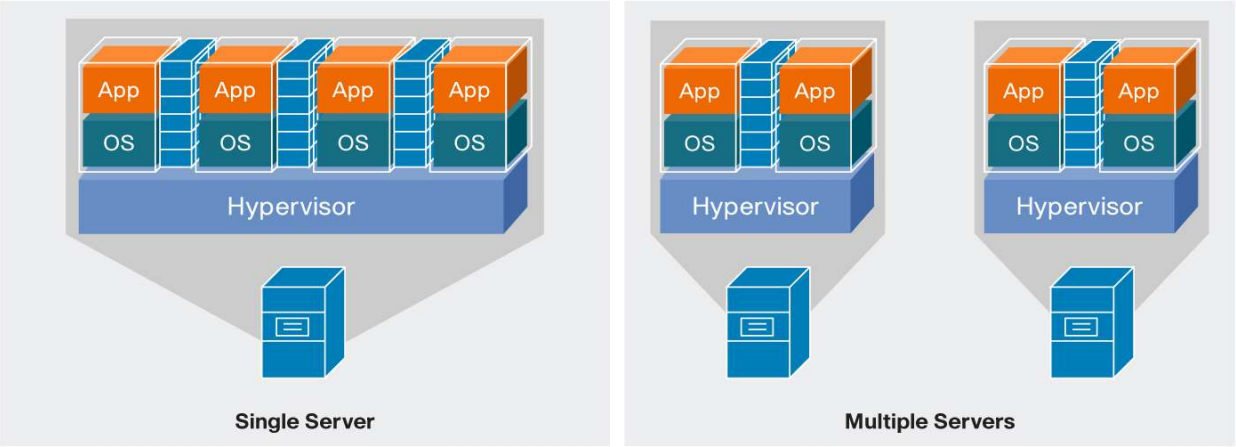

These provisioning options bring the following benefits to the branch office:

- **Lower infrastructure and operating costs:** Most server hardware is underutilized. Consolidating multiple applications to run on a single physical server eliminates unnecessary servers and increases utilization of the remaining servers. The cost savings extend beyond server hardware. Smaller number of servers results in:
  - **Lower equipment and facilities costs:** Fewer Ethernet ports, power supplies, patch panels, and surge protectors and less cabling and rack and physical space are needed.
  - **Lower energy and support costs:** Less energy is needed to power and cool the remaining servers and fewer hardware support contracts, and less time is spent provisioning servers.
- **Less application downtime and quicker response time:** Multiple applications hosted on one operating system share and compete for resources. Isolating applications from each other helps ensure that an action of one application has no impact on the rest of the system. Therefore, dedicating a virtual server to an application results in:
  - **Shorter planned and unplanned downtime:** Patches, updates, or upgrades to one application are transparent to the other applications, one misbehaving application does not affect the other applications, and rebooting of one application does not take the other applications offline.
  - **Explicit allocation of hardware resources:** The hypervisor allocates a specific amount of hardware resources to an application that cannot be taken away by another application, resulting in a guaranteed level of responsiveness for all applications.
- **Faster time to deployment for applications:** Multiple applications hosted on one operating system have various dependencies. Isolating applications from each other helps ensure that a change in one application is transparent to the other applications. Therefore, dedicating a virtual server to an application results in:
  - **Faster implementation of changes for existing applications:** Less coordination, planning, and testing is needed to roll out patches, upgrades, and updates to a single application.
  - **Shorter time to deployment for new applications:** Less coordination, planning, and testing is needed to roll out a new application.

Tables 2 through 4 provide in-depth explanations of each benefit. The solution requires the following changes:

- Install a hypervisor on the original or a new physical server.
- Dedicate a single virtual server for each application.
- Provision each application to its virtual server.
- Deploy all virtual servers on the hypervisor.

Table 2. Lower Infrastructure and Operating Costs Benefits Details

Current setup	Dedicated physical server for each application.
Challenges	• Added capital expenses: additional servers, cabling, power supplies, rack space, switch ports, patch panels, surge protectors, physical space, operating system licenses, and equipment spares. Most of these resources are underutilized. According to one McKinsey & Company study, data center server utilization rarely exceeds 6%.[3] For most branch-office servers, the average utilization may be even lower because they host services with low resource requirements such as DHCP, DNS, or print servers. • Additional operating expenses: increased energy costs to power and cool larger-than-necessary infrastructure, additional per-server hardware support costs, and additional administration costs to provision, configure, and maintain multiple servers.
New setup	Dedicated virtual server for each application, all hosted on one physical server.
Results	• Smaller infrastructure footprint. • Optimized energy consumption. • Faster server provisioning, configuration, and maintenance.

[3] http://www.mckinsey.com/clientservice/bto/pointofview/pdf/Revolutionizing_Data_Center_Efficiency.pdf

Benefits	• Cost savings resulting from infrastructure optimization; combining multiple applications onto a single server and increasing its average utilization provides capital expense savings on the following components: ◦ Fewer CPUs, less memory, and fewer physical network interface cards (NICs) and hard disk drives. ◦ Less cabling and fewer power supplies, patch panels, switch ports, and surge protectors. ◦ Less rack and physical space. ◦ Possibly lower operating system license costs: some operating systems allow multiple running instances per license; for example, Microsoft Windows Server 2008 Enterprise Edition allows up to four running instances per license on one physical server.[4] • Cost savings resulting from smaller infrastructure footprint; combining multiple applications onto a single server and increasing its average utilization provides the following operating expense savings: ◦ Lower energy consumption to power and cool branch-office infrastructure. ◦ Lower annual per-server hardware support services. ◦ Less administration time spent on provisioning, configuring, and maintaining server hardware. • Virtual NICs and switches available at no additional cost.
Costs	Hypervisor and virtual environment management software.

Table 3. Less Application Downtime and Quicker Response Time Benefits Details

Current setup	Multiple applications share one physical server.
Challenges	• One misbehaving application may affect all other applications or take down the entire operating system environment. • One application may consume all available hardware resources and degrade performance or responsiveness of all other applications. In a worst-case scenario, an application can starve other applications of all resources. • Changes to one application, or deployment of a new application, may require you to take the entire server offline for an extended period of time.
New setup	Dedicated virtual server for each application, all hosted on one physical server.
Results	• Applications are isolated from each other, and therefore an operating system failure caused by one application has no effect on the other applications. • Applications can be assigned a guaranteed amount of hardware resources (CPU, memory, and storage) that other applications cannot take away. • A new application or an application change is restricted to one virtual server, which you can take offline without taking the physical server offline.
Benefits	• Improved overall system and individual application uptime. • Reduced system mean time between failures (MTBF) and mean time to repair (MTTR) metrics. • Control over resource sharing based on each application's performance or other business requirements.
Costs	• Hypervisor and virtual environment management software must be added. • Addition of a hypervisor and a new operating system for each application may exceed the hardware resources of the existing physical server. Therefore, this change should be planned along with the next server refresh cycle. • Addition of a new operating system for each application may increase the operating system license cost, depending on the specific operating system being added. Microsoft Windows Server 2008 Enterprise Edition allows up to four running instances per license on one physical server,[4] which is sufficient for most lean branch offices.

Table 4. Faster Time to Deployment for Applications Benefits Details

Current setup	Multiple applications share one physical server.
Challenges	• Applications require periodic security patches, maintenance updates, or functional upgrades. In organizations with multiple teams responsible for different applications that share a physical server, any change to any one application must be coordinated across all teams. The additional coordination and planning slows down the rollout of the change, introduces additional testing, and may trigger unnecessary changes in the other applications. • Application changes may require a security patch, a maintenance update, or even an upgrade of the host operating system. All other applications will be affected whether or not they require the change. The coordination issues highlighted discussed earlier are amplified. • High performance, fast response time, high availability, or other business needs may require a new application to be hosted locally in the branch office. Rollout of the new application must be coordinated to prevent adverse effects on existing applications or the host operating system.
New setup	Dedicated virtual server for each application all hosted on one physical server.

[4] http://download.microsoft.com/download/F/C/A/FCAB58A9-CCAD-4E0A-A673-88A5EE74E2CC/Windows_Server_2008_Virtual_Tech-VL_Brief-Jan_09.docx

Results	• Applications are isolated from each other, and therefore a change in one application or its host operating system has no impact on the other applications. • A new application is provisioned in a new virtual server and deployed on the hypervisor without affecting the other applications.
Benefits	• Faster rollout of application or host operating system patches, updates, or upgrades. • Less time spent on cross-functional coordination required to plan and test application changes. • Improved ability to respond to security risks, application defects, or business requests to add new features or applications.
Costs	• Hypervisor and virtual environment management software. • Adding a hypervisor and a new operating system for each application might exceed the hardware resources of the existing physical server. Therefore, this change should be planned along with the next server refresh cycle. • Adding a new operating system for each application might increase operating system licensing costs. However, this depends on the specific operating system that you are adding. Microsoft Windows Server 2008 Enterprise Edition allows up to four running instances per license on one physical server,[4] which is sufficient for most lean branch offices.

In summary, branch-office server virtualization increases the speed of application deployment, improves application uptime and performance guarantees, and reduces equipment and operating costs. But server virtualization can do much more. Server virtualization offers a rich set of features that can enable new capabilities in the branch office.

Enabling New Capabilities in the Branch Office Through Server Virtualization

Server virtualization brings many new features to the data center that a traditional server cannot provide, such as live virtual server migration, dynamic resource allocation, live CPU and memory increases, and distributed software switches. The small hardware footprint of the lean branch office has limited the use of such features. Nonetheless, a number of new capabilities offered by server virtualization can improve server uptime and failure recovery time and automate server provisioning in the branch office.

A typical lean branch office implements the disaster recovery process either by storing backups on a local storage device with periodic archival to tape or by sending them to a data center for storage and tape archival. A growing number of branch offices rely on continuous data protection by replicating, snapshotting, or mirroring data to a local storage device. Regardless of where the data resides and how frequently it is backed up, two fundamental techniques are used to copy and restore the data (Figure 5):

- **Block-based disk copying with full-system restoration:** This method provides the fastest full-system backup and restore operations, low performance overhead during backups, and full recovery of the server. However, it is an all-or-none approach that requires large storage capacity, individual files cannot be restored, and the restore target must be identical to the original server.
- **File-based file system copying with data and configuration restoration:** This method provides backup and restore operations at the individual file level, platform-independent restoration, and flexibility to adapt the configuration for the restoration target. However, the backup and restore operations take longer, performance overhead is higher during backup, and a server cannot be fully recovered.

Figure 5. Physical Server State Copy and Restore Techniques

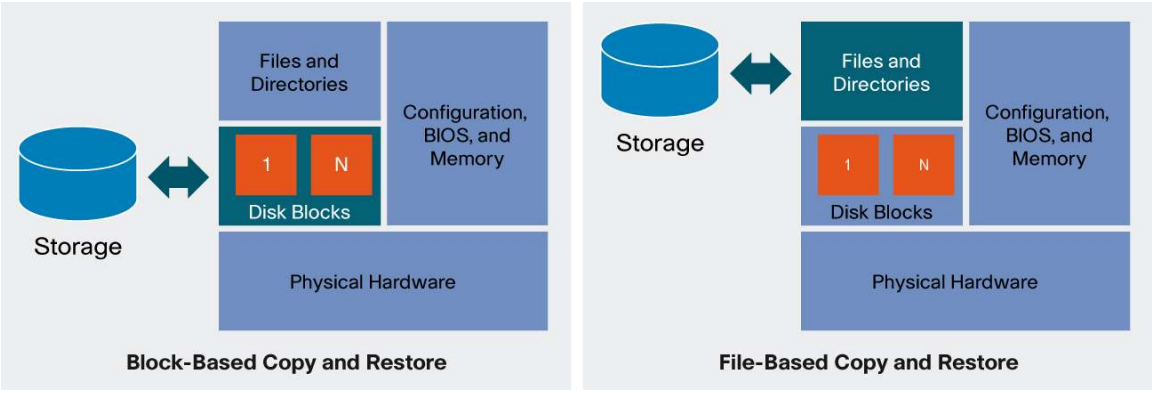

Server virtualization provides new options for capturing a point-in-time server state, which is then stored in several files. The abstraction and encapsulation attributes described previously facilitate the following (Figure 6):

- **Virtual server copy and restore:** This option copies the entire persistent state of the virtual server, which includes the disk, memory, log, configuration, suspended state and snapshot data, NIC teaming, and BIOS. It allows complete restoration of the server. The server must be powered off during the copy operation.
- **Virtual server snapshot:** This option saves the entire persistent and running state of a virtual server at a point in time and allows the server to return to this state. The snapshot is performed on a running server.

Figure 6. Virtual Server State Copy and Restore and Snapshot Techniques

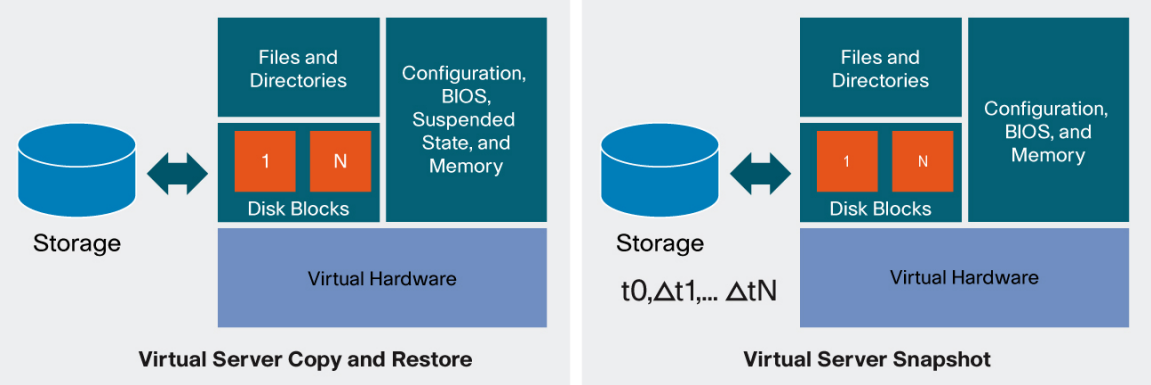

The capability to retain the entire server state in a small number of files provides the following benefits to the branch office:

- **Less application downtime and faster failure recovery time:** Traditional backup and recovery procedures only copy the server state written to the disk. Encapsulating the state of a virtual server in a set of files facilitates preservation of the entire state at a point in time and captures changes over time. Therefore, saving or snapshotting the virtual server state results in:
  - **Faster recovery of the server during hardware migration or after a server failure:** The hypervisor can restart the virtual server from saved files, the files can be exported for backup purposes, and virtual server files can be quickly imported to another hypervisor with the same or a different hardware configuration.
  - **Faster rollback to a stable configuration if an application becomes unstable:** The virtual server can be rolled back to a point in time, reversing any unwanted changes.

Table 5 provides in-depth explanations of each benefit. The solution requires no additional infrastructure changes beyond those discussed in the previous section.

Table 5. Less Application Downtime and Faster Failure Recovery Time Benefits Details

Current setup	Block-based or file-based backup and restore of server state.
Challenges	• Block-based backup and restore mechanisms require nearly identical hardware configurations for disaster recovery or server migration. • File-based backup and restore mechanisms require a long multistep process for disaster recovery or server migration without offering the benefit of full state preservation. • Rollback to a stable server state in case of a misconfigured, defective, or compromised application requires a long multistep restore process and may lead to data loss. Continuous data protection mechanisms with snapshotting provide full rollback capabilities but cost more, consume more processing power, and require more storage capacity than traditional backup operations. Troubleshooting the corrupted application requires restoration of the most recent backup to a spare server. This process is complex and time consuming.
New setup	Dedicated virtual server for each application with the use of snapshotting and copy, restore, import, and export operations.

Results	• A “golden” image of the virtual server is stored alongside the running instance to serve as a fail-safe copy. • A snapshot of the virtual server is taken when application patches, updates, or upgrades are deployed. • The production instance of a virtual server is exported from an existing hypervisor and imported to a different hypervisor when migrating to new server hardware. • The most recent copy of a virtual server is exported from an existing hypervisor, stored on a backup device, and imported to a different hypervisor when a server is restored after a catastrophic failure. (Note: A number of backup and restore solutions automate the disaster recovery process for virtualized infrastructure.)
Benefits	• Faster recovery time in case of a catastrophic hardware failure. • Faster migration time when moving to new server hardware. • Faster rollback to a well-known state in the case of an unstable application. • Capability to troubleshoot an unstable application on a backup replica of the virtual server while the primary instance is rolled back to an earlier state.
Costs	• Hypervisor and virtual environment management software. • Optional backup and restore solution for virtualized infrastructure.

In summary, the capability to copy and restore the entire state of a virtual server to a specific point in time improves the speed of recovery after a server failure, speeds up the migration process to new server hardware, and helps enable faster rollback to a fail-safe state in case of an unstable application.

VMware and Cisco Branch-Office Server Virtualization Solution

Cisco Unified Computing System Express (UCS Express) combines the on-demand application provisioning capabilities of the Cisco Service Ready Engine (SRE) x86 blade and the hardwarelike reliability and performance of the VMware vSphere Hypervisor™ (ESXi) into a server virtualization platform for the lean branch office. Cisco UCS Express facilitates the consolidation of all branch-office network and application services into the second generation of Cisco Integrated Services Routers (ISR G2). This solution is best suited for multisite organizations with centralized IT infrastructure that need to host a small number of essential applications locally in the branch office. Unlike standalone x86 servers with virtualization, Cisco UCS Express combined with the multiservice Cisco ISR G2, provides more agile, simpler, and lower-cost branch-office infrastructure integrated into a single device. Unlike appliance-based virtualization products, Cisco UCS Express combined with VMware vSphere Hypervisor provides enterprise-class virtualization hosted on high-performance, feature-rich hardware.

The Cisco ISR G2 acts as a blade server chassis capable of hosting one or more Cisco SRE blades, depending on the Cisco ISR model. Therefore, with just a small increase in power consumption (50W per blade) and no increase in rack space or cabling, one or more servers can be deployed in the branch office. The Cisco SRE blades communicate with each other and the router over a multigigabit fabric backplane, eliminating the need for external cables. IP addresses, switch-port assignment, and routing services including Network Address Translation (NAT), firewall, and intrusion prevention system (IPS) can be added dynamically as required without the need to physically recable the blades.

Data center consolidation has triggered centralization of branch-office applications. However, some applications continue to be deployed in the branch office for the following reasons:

- **Performance:** Applications that cannot tolerate the latency and bandwidth limitations of a WAN link because they either require fast response times or generate large amounts of data (for example, software distribution depots, performance monitoring services, and interactive applications)
- **Availability:** Applications that cannot tolerate the unreliability of a WAN link because they require continuous availability (for example, DNS server, DHCP server, Microsoft Active Directory Domain Services, print service, and point-of-sale applications)
- **Compliance:** Applications that cannot be hosted outside the branch office or that require a local copy of data because they must meet security, regulatory, or internal policy requirements (for example, Payment Card Industry Data Security Standard [PCI DSS], Health Insurance Portability and Accountability Act [HIPPA] compliance, and Sarbanes-Oxley regulations)

The Cisco SRE blades have been designed to help organizations provide services to lean branch offices. On-demand application provisioning helps enable organizations to remotely deploy networking and computing services to the branch office at any time. It eliminates costly onsite visits and infrastructure modifications when performance, availability, or compliance considerations require deployment of a local application. Deploying Cisco SRE with Cisco ISR G2 removes the burden of having to decide at the beginning which services have to be provisioned in the branch office, and it provides the flexibility to change service placement decisions in the future, eliminating costly infrastructure lock-in.

Cisco UCS Express powered by the VMware ESXi enhances the Cisco SRE on-demand application-provisioning model by providing the capability to host any application in the Cisco ISR G2 while taking advantage of all the benefits provided by virtualization as described in this document:

- **Lower infrastructure and operating costs**
- **Less application downtime and faster response time and failure recovery time**
- **Faster time to deployment for applications**

Cisco UCS Express allows one or multiple instances of Microsoft Windows Server to run directly on the Cisco ISR G2. This capability helps organizations with Cisco ISR G2 consolidate Microsoft Windows applications and core Microsoft Windows services such as Microsoft DHCP server, DNS server, Active Directory Domain Services, and print services on a single branch-office device. The Microsoft Windows Server on Cisco UCS Express combined with Cisco ISR G2 creates the industry's first and only solution for hosting all branch-office applications and services such as routing, switching, security, voice, video, wireless, computing, storage access, server virtualization, Microsoft Windows core services, and line-of-business applications. This "branch office in a box" provides the following benefits:

- **Cost savings:** The multiservice Cisco ISR G2 is the lowest total cost solution for branch-office services.[5] The Cisco SRE deployed on the Cisco ISR G2 adds to these cost savings by eliminating onsite visits and infrastructure changes, lowering energy consumption, and attaching Cisco SRE hardware support services to the Cisco ISR G2 at no additional cost. Microsoft Windows Server on Cisco UCS Express adds cost savings from the benefits of server virtualization such as infrastructure footprint reduction, energy savings, lower hardware support services, and increased IT productivity.
- **Consolidation:** Providing all branch-office services and Microsoft Windows applications in a single device simplifies and reduces the footprint of branch-office infrastructure. With fewer servers and appliances, less cabling, fewer Ethernet ports, fewer power supplies, and less rack and physical space, a hard-wired physical infrastructure that needs frequent onsite support is replaced by a soft-wired virtual infrastructure that is easier to manage remotely.
- **Integration:** The Cisco UCS Express platform is connected to the rest of the branch-office network and other Cisco ISR G2 service modules through a multigigabit backplane switch. Virtual servers running Microsoft Windows Server can now directly take advantage of router and switch features such as VLANs, security zones, access control lists (ACLs), and firewall with no impact on performance. By unifying routing, switching, computing, and storage access, the Cisco ISR G2 provides more flexibility, more granular control over security, and higher performance.

[5] http://cio.cisco.com/en/US/prod/collateral/routers/ps5855/prod_white_paper0900aecd805898e5.html

Cisco UCS Express Use Case 1: Core Windows WAN Edge Services

An organization decides to implement Microsoft Active Directory Domain Services, DHCP server, and DNS server locally in the branch office for these reasons:

- **WAN reliability:** The WAN is unavailable from time to time, and the cost of a backup link is prohibitive; or a second service provider is not available in the specific geographic area.
- **WAN performance:** Latency between the branch office and the nearest data center is significant, and branch-office users constantly generate queries to Microsoft Active Directory Domain Services, DHCP server, and DNS server.

The branch office experiences productivity slowdown when these services are either not available or perform slowly. With Cisco UCS Express, core Windows services can be implemented in the following ways:

- Each service can be run as a dedicated virtual server.
- A group of services can be co-located and run on a virtual server.

(Actual implementation depends on the needs and requirements of the individual organization.)

This solution decreases the infrastructure footprint at the branch office.

Cisco UCS Express Use Case 2: Bank Teller Line-of-Business Applications

A retail bank cannot tolerate any downtime for its branch-office teller application during business hours. The organization implements a thin in-office control point (IOCP) that serves as a standby system for a centrally hosted teller server. Under normal conditions, the branch-office teller client interacts with the central teller system. However, if the WAN connection is disrupted or degraded, the local IOCP takes over serving requests from the teller client. After the WAN connection is restored, the IOCP synchronizes the in-office transactions and returns control to the central teller system. With Cisco UCS Express, the IOCP can run on a dedicated virtual server hosted in the Cisco ISR G2. This solution provides higher availability for a business critical application.

Cisco UCS Express Use Case 3: Virtualized Network Appliances

A number of appliances providing network-level functions, including firewall, wireless LAN controller, private branch exchange, and WAN optimization controller functions, continue to be deployed in the branch office. All such devices can be virtualized and hosted on the Cisco UCS Express. This solution reduces complexity of the remote infrastructure.

Summary

A typical lean branch office has a significantly different server infrastructure than a data center, but despite the differences, multisite organizations have much to gain by deploying server virtualization in their branch offices. The virtualization technology helps meet common branch-office challenges, and it introduces new capabilities not available with traditional servers. The main benefits of virtualization in the branch office include:

- Lower infrastructure and operating costs
- Less application downtime and faster response time and failure recovery time
- Faster time to deployment for applications

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

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 C11-611291-00 03/11