

Benefits of VMware and Hitachi Data Systems



Table of Contents

Challenge3

Solution.....3

VMware® Infrastructure at Work3

Benefits3

Overview.....3

New Technologies Needed for Consolidating Operations4

Storage Virtualization Enables Centralized Management4

Deciding on VMware Server Virtualization Technology4

Building a Unified Virtualized Environment5

Lessons Learned5

Benefits of Virtualization and Centralized Storage6

Looking to the Future.....6

Challenge

The requirement was to consolidate core IT operations and to model workloads, server farms and storage and backup technology in an efficient way.

Solution

VMware virtualization technology and storage solutions from Hitachi Data Systems were used to meet the need to centralize and consolidate storage resources.

VMware® Infrastructure at Work

- VMware Infrastructure 3 on 4-CPU HP ProLiant 580 G2s and on 4-CPU IBM X460s
- VMware® VirtualCenter running on 2-CPU HP ProLiant ML370s
- VMware® VMotion technology
- Microsoft® Windows® 2003 Server Standard, Windows 2003 Server Enterprise, Windows 2000 Server, Windows 2000 Advanced Server
- Microsoft SQL Server, Microsoft Systems Management Server, SAS, WTS MS.net, Peoplesoft, Solarwinds Orion, Patrol, HP Insight Manager, plus network monitoring, development, testing and staging, and hosting servers
- Hitachi storage systems, including the new Hitachi Universal Storage Platform™ family and legacy Hitachi models such as the Lightning 9900™ V Series and Thunder 9500™ V Series storage systems.

Benefits

- Reduced annual server hardware and associated labor costs
- Reliable business continuity
- Increased performance and flexibility

“Some might say that virtualization is virtualization, but it’s really the tools and packaged deal that VMware offers that led us to our final—and very good—decision to go with VMware® Infrastructure 3.”

Gary Roberge, Chief, Regional Office IT Services

Overview

Statistics Canada is a government agency that collects and analyzes vital data in areas ranging from the country’s economy to its culture. Besides conducting a census every five years, it has about 350 active surveys that shed light on virtually all aspects of Canadian life—and form the foundation for informed decision-making in government, business and the non-profit sector. As you’d expect at a federal agency that lives and breathes data, Statistics Canada requires an IT infrastructure that can handle serious number crunching and data storage. In 2002, the Canadian government issued a directive mandating that all agencies reduce operating budgets. That directive was a tall order for IT Service Delivery which was grappling with ever-increasing processing and storage demands. For instance, the year before, Statistics Canada offered citizens the choice of responding to the census on the Internet—an electronic option that generated an additional seven terabytes of data.

To meet the dual challenge of expanding services while cutting costs, ITSD kicked off a streamlining initiative to centralize and consolidate key IT infrastructure services. The initiative tackled four major areas: centralized storage, centralized backup/restore, file/print and server consolidation.

The streamlining effort also focused on the big picture—Statistics Canada’s overall IT infrastructure. “The infrastructure across the organization’s 65-odd divisions was a mixed bag,” said Gary Roberge, chief, Regional Office IT Services. “Every division had its own unique organically grown hardware infrastructure hosting an assortment of operating systems. Hardware and personnel costs were on the rise, the data center was reaching capacity, and budgets were on the decline. It became apparent that centralization and consolidation of our server infrastructure in Ottawa was of primary importance.”

Solutions from VMware and Hitachi Data Systems are the foundation of the technology infrastructure that Statistics Canada is developing to meet the challenges of centralization and consolidation today—while paving the way for a long-term vision of delivering computing resources based on a 24/7 utility model. VMware Infrastructure 3 has helped the agency cut hardware costs, enhance business continuity and foster centralized server management. And Hitachi Data Systems has provided ITSD with the virtualized storage capacity it needs to safely and economically house its burgeoning data stream, which is stored not only on Hitachi storage systems, but also on other vendors’ gear.

New Technologies Needed for Consolidating Operations

Over the past decade, Statistics Canada has added new services, such as online census reporting, and increased the number of statistical surveys and reports it offers. That translates into increased server demand, data storage and backup requirements, and IT operating costs. The task of keeping up with those technology challenges grew even more daunting in 2002, when the Canadian government called for budget cuts at all national agencies. The following year, Statistics Canada decided that the best way to do more with less was to consolidate some of the operations of its distributed network into the datacenter at its Ottawa headquarters. "There are more than 5,000 employees at Statistics Canada, and our services touch almost every one of them," said Sal Carpinone, chief, Managed Storage Services, about ITSD's mission. "We provide enterprise services, centralized storage and centralized virtual machine technology to all our clients."

Although Statistics Canada mandated that operations, like storage and backup, must be handled centrally by ITSD, the agency's business units can still use local solutions for other IT operations. In essence, ITSD has to sell its services to Statistics Canada's various divisions. "In some respects, we compete against client organizations," said Eric Michaud, chief, Managed Server Services. "So we have to offer iron-clad, highly available, highly recoverable services. Our core challenge is delivering those enterprise-class services in a way that's cost-effective to the organization. "Cost is a focus at the agency not only because of the Canadian government's money-saving mandate, but also because ITSD operates on a zero-based budget, which means that all the IT solutions it provides internal clients are charged on a cost-recovery basis."

As the agency moved to consolidate core IT operations, ITSD considered new technologies that would support those efforts. As a result, VMware virtualization technology and storage solutions from Hitachi Data Systems now play major roles in Statistics Canada's revitalized IT infrastructure.

Storage Virtualization Enables Centralized Management

Statistics Canada had about 100 terabytes of data, most of it stored and managed locally on more than 750 servers—with about 49 percent of disk space unutilized on each server. To increase efficiency and deal with anticipated growth in storage and backup needs, the agency mandated a migration to a centralized model. ITSD decided that storage virtualization would be the best way to fulfill that requirement.

The Hitachi Universal Storage Platform family provided the capabilities needed. The Universal Storage Platform's

virtualization and partitioning capabilities via Hitachi Virtual Storage Machines were not only the most robust, but also the least expensive overall when compared to the approaches of other enterprise vendors. "We looked at storage from the industry's leading vendors," said Carpinone. "Given our virtualized partition requirements, Hitachi storage looked like the best bet—and it was also the most cost efficient."

"That's how it started," he added, "but as it's evolved, we've established a good relationship with Hitachi. Hitachi knows our business. As a federal government agency, we don't close the door on other manufacturers, but Hitachi has been consistent through the years and we have a great relationship. I think by understanding the business of Statistics Canada, Hitachi is in a better position to help us make the most of their technology in the environment we have. Our clients' storage requirements are considerable, and Hitachi has helped us meet them."

Carpinone and his team provide three tiers of storage and backup services. First, a high-end fiber channel is used for primary storage of business-critical data. Second, serial ATA disks are used for secondary storage, which his team offers clients at reduced costs. Finally, tape backup is deployed for data that's accessed infrequently.

The Hitachi flexible storage solutions let Statistics Canada grow its storage infrastructure gradually to reflect the agency's needs, which have grown from 100 terabytes of data five years ago to a whopping 400 terabytes today. After the agency's initial deployment of Hitachi storage solutions in 2001, it added Hitachi Thunder 9500 V Series storage systems for departmental use. Hitachi Universal Storage Platforms were first installed in 2005, and new arrays were added in 2007 and again in 2008.

"One of the hardest things in the federal government is the acquisition process, which can take a very long time," Carpinone said. "But we don't want to over buy, since we operate on a cost-recovery basis. We need to be lean and mean—and to buy only what we need."

Deciding on VMware Server Virtualization Technology

Statistics Canada began evaluating server virtualization in February 2005 as a way to centralize server management, consolidate 212 Intel servers at its Ottawa headquarters, and improve business continuity and disaster recovery processes. According to Roberge, the agency evaluated solutions in terms of ease of use, management tools, the number of virtual machines it could get on a physical server, and whether or not virtualization fit into its storage and backup strategies. He and his team considered other virtualization products and non-virtualized blade server implementations, but ultimately decided that VMware offered the best solution.

"Number one, we didn't think other products focused entirely on virtualization, and number two, we knew for sure that the VMware toolset was far superior to anything out there," Roberge said. "Some might say that virtualization is virtualization, but it's really the tools and packaged deal that VMware offers that led us to our final—and very good—decision to go with VMware Infrastructure 3."

As you might expect of a federal agency, Statistics Canada does considerable due diligence on the products and services it employs. "There's always a lot of research and testing before we recommend a solution, especially at an enterprise level," Carpinone said. "If the business case isn't there, it won't make it through all the committees that purchasing decisions have to get through." Since everything ITSD does is on a cost-recovery basis, those business cases tend to focus on the cost of adopting a particular technology versus the revenue it generates. "With VMware, that case is easy to make," he said.

Building a Unified Virtualized Environment

Moving forward, Statistics Canada began researching ways to combine VMware server virtualization benefits with the agency's consolidated storage infrastructure. Although Statistics Canada has storage from a broad array of vendors, it selected Hitachi Data Systems as the vendor of choice in VMware environments. That decision was due to the virtualization potential of the two architectures and early support by Hitachi Data Systems of the VMware Infrastructure 3 environment, not only standalone but also for virtualized external storage environments.

To provide appropriate security, the IT environment of Statistics Canada is divided into three distinct networks. Network A houses business-critical data, Internal Network B contains internal business information, and External Network B contains public information. The two B networks are standard WANs, but security concerns put Network A in a separate category. "It's a lean production network where most of the data processing occurs," Michaud said. "We keep it physically separated from everything else, so there's no routing between it and the other networks.

"That has introduced some challenges," he added, "because we've had to deploy physical infrastructure along all three networks to sustain our client requirements. That spreads our resources thin by making us provision all three networks with hardware that's not utilized as much as we'd like."

In the first phase of its deployment of VMware software, ITSD consolidated 45 physical servers onto virtual machines running on just four physical servers, with approximately 1,000 employees accessing the applications running on virtual machines. The agency began its VMware implementation with VMware ESX 2.5, then migrated to VMware ESX 3.

Carpinone said that migration took his team longer than anticipated because it had to shut down guest images to migrate. "That caused a significant amount of pain because every guest has its own clients and the clients determine the maintenance," he said. "We now know that VMotion technology will help us to do zero-downtime migration next time," he said, looking ahead to a planned migration to VMware ESX 3.5 later this year.

Since then, Statistics Canada's virtual infrastructure has grown gradually. As they are added, all new applications are deployed on virtual machines, and when the time comes to replace older physical servers, they're moved onto virtual machines without any disruption, using VMware VMotion. By Q1 2008, Statistics Canada had VMware licenses for 86 workstations, a number that's expected to rise to about 120 later in the year.

"We have pretty much everything," Michaud said, "database services, application services, web servers and middleware servers. Just about anything out there, we probably host it on VMware."

At this stage, Statistics Canada has approximately 30 hosts running across its three networks, backed up by IBM 3590 tape drives. Roughly 250 virtual guest images run across those hosts, so the ratio is approximately 8 to 1. By mid-2008, the agency plans to increase the number of ESX servers to roughly 350.

Taking full advantage of VMware Virtual Symmetric Multi-Processing (SMP) technology, which lets a single virtual machine use up to four processors, ITSD greatly enhanced the performance of its dual-core IBM x460 processors. "By loading up those boxes with 32 gigabytes of RAM and using pass-through, we are jumping these 4-ways into 8-ways, making them 16-ways because of the dual cores," Carpinone said. "As for operating systems, you name it, we have it: Windows Advanced Server, Sun Solaris, HP-UX, Red Hat Linux, Novell's SUSE Linux, and IBM AIX. We have every major flavor. We're running most of that—Windows, Solaris, and Linux—on our VMware platform."

Statistics Canada recently decided to phase out its mainframe, so many applications that once ran on big iron are being ported not to physical, but to virtual servers. Even with the efficiencies offered by virtualization, Michaud says ITSD is grappling with server sprawl as the agency addresses processing needs once handled by the mainframe.

Lessons Learned

Focus on operations. Looking back on the rollout of VMware technology at Statistics Canada, Michaud said that operational readiness is just as important as technological savvy. "If you're more focused on technology than operational readiness, it usually catches up with you, because you lack the change-

management disciplines—like process-, problem- and incident-management systems—to properly support an environment that has such density,” he said. “For instance, when a server goes down, it only affects one client. When you put 40 clients on the same box and it goes down, it has a greater impact. When we embarked on the virtualization process, ITSD wasn’t ready to cope with that, but we are working through that as we speak.”

Plan for exceptions, not just the norm. “We try to establish a norm and fit everybody within that norm, but that doesn’t always work,” Michaud said. “I’d say 80 percent of our clients are very satisfied with the services we offer on virtual machines, but at this stage we’re having challenges meeting the needs of the other 20 percent.” Those clients, he explains, don’t fit the norm because of distinct business or performance requirements. “We tried in the past to be everything to everyone,” he said, “but right now we are learning that we need to adapt and be flexible enough to accommodate those exceptions.”

If problems arise, call in the experts. “Our VMware environment has been successful, but in some ways it’s kind of a victim of its own success,” Michaud said. “I call it growing pains. Resources are competing against each other and that’s causing some problems on the platform.”

To fix those glitches, ITSD is working with a VMware partner, Toronto-based Gibraltar Solutions, to tune its network so there are fewer CPU and I/O contentions. For instance, a Gibraltar consultant recently did a VMware Health Check, interviewing the team that supports Statistics Canada’s virtual platform, and offering recommendations on optimizing the agency’s virtual architecture.

Train staff. As befits a division run on a cost-recovery basis, ITSD aims to be as lean and efficient as possible. That’s certainly true of the team that supports Statistics Canada’s virtual infrastructure—just four people. Each is proficient in VMware technology, and training is ongoing.

Provide users with technology they can count on. ITSD’s clients need to know they can rely on the IT services the division provides. “They need be aware of the robustness built into the system that everything is cooled and powered and backed up—that we’re doing much more than deploying a simple storage array on the floor of the data center somewhere,” said Michaud.

Another factor that reassures ITSD’s clients is the knowledge that their IT infrastructure is built on offerings from trusted partners like Hitachi and VMware. “Both companies have huge research and development organizations behind the technologies they offer,” Carpinone said.

Standardizing on virtualization. Virtual platforms have become the platforms of choice at Statistics Canada. All new applications are being run on the agency’s VMware

Infrastructure. And when the time comes to renew the servers for existing applications, those are being ported over to the virtual infrastructure as well. In addition, new storage and disaster recovery demands are being fulfilled from the Hitachi-enabled storage virtualization infrastructure.

Benefits of Virtualization and Centralized Storage

One of the most tangible benefits of virtualization at Statistics Canada has been cost savings. The agency estimates that it has saved a considerable amount in annual server hardware and associated labor costs.

Looking at other numbers, Statistics Canada has achieved an 8:1 server consolidation ratio. As for CPU utilization, the organization has gone from about 4 percent CPU utilization to between 50 percent (average) and 80 percent (at peak times).

In the past, server provisioning was a 2 month process involving the government’s standing-offer process, Statistics Canada’s own paper process, and then resellers’ processes. “With virtual servers, that’s down to days,” Michaud said.

“A huge selling point for our internal clients was business continuity,” said Roberge. “Prior to virtualization, very few divisions had backup systems. If their servers failed, that was the end of it. Today, they can put images on virtual machines and know that, should systems fail, their data will be safe.”

Reliable centralized storage and backup doesn’t just bolster users’ peace of mind about business continuity and disaster recovery, it also saves the agency money. For starters, the Hitachi virtualization architecture, taken as a whole, is often less expensive than alternative, non-virtualized implementations from other storage vendors. And another area of savings comes from the lack of service disruptions.

“Hitachi Data Systems has improved our IT service levels,” Carpinone said. “Has it improved uptime? Absolutely! We have been able to maintain services to our clients even with hardware failures.” In fact, a planned December 2007 shutdown was the only time ITSD has had any downtime in its storage and backup services.

Looking to the Future

ITSD will continue to extend the benefits of virtualization across Statistics Canada’s organization. It plans to upgrade to VMware ESX 3.5 in the coming months and buy licenses for more than 30 additional VMware ESX hosts. “We’re moving virtualization into different areas and divisions,” Roberge said. “Attitudes have changed a lot. People who were at first hesitant because they couldn’t physically touch the virtual machines are now starting to see the true value of virtualization. We’re probably putting up a server a week now.”

"Right now our drive to consolidate is being driven by our desire for cost-reduction and greater manageability," Michaud said. "And down the line, we hope consolidation will make us a greener operation."

Another benefit Michaud hopes to reap from virtualization is better trending. "I'd like to see us doing performance mapping so we can move workloads around our infrastructure based on server performance and clients' needs at the time," he said. With VMware® Distributed Resource Scheduler (DRS), organizations like Statistics Canada can align their IT infrastructure with their business goals by dynamically allocating and balancing computing resources. DRS continuously monitors utilization across resource pools and intelligently allocates available resources among virtual machines based on pre-defined rules that reflect business needs and changing priorities. VMware® vMotion™ leverages the complete virtualization of servers, storage and networking to move an entire running virtual machine instantaneously from one server to another. This keeps the IT environment up and running, providing unprecedented flexibility and availability to meet the demands of business and end users.

On the storage front, Carpinone's group is currently looking at implementing a faster, more reliable method for backing up data that accommodates the dramatic increase in volume. The agency is planning to implement the latest Hitachi storage product—the Hitachi Universal Storage Platform VM—to virtualize and manage external storage (from both Hitachi and other vendors) as part of a disk-to-disk backup option.

Carpinone and his team are continuing their quest to make their storage and backup infrastructure as efficient as possible. "Because we are a central service, if we allocate 300 terabytes and the usage is only 200 terabytes, that's 100 terabytes of unused storage space that's costing the agency money," he said. His goal is to make that unused portion smaller, while still ensuring that Statistics Canada's clients have all the storage capabilities they need. He's investigating Hitachi's latest storage improvement—"thin provisioning"—which is a way to automatically allocate storage as needed. "Hitachi allows us to buy an enterprise storage system that's scaled to our current requirements, but can be scaled up as our capacity requirements increase," he said. Solutions from Hitachi Data Systems and VMware are helping Statistics Canada give its statisticians and other processing- and storage-hungry users all the IT capabilities they need—and to deliver that processing and storage muscle as efficiently as possible. The model is a utility that's available 24/7. "When people use power, they don't necessarily need to know if it's coming from a hydroelectric turbine or a nuclear power plant," Michaud said. "They just need to know that their power works. If they want to turn on the air

conditioner, the power's there. If they want heat, it's there. That's the model we need to get to. And VMware and Hitachi are helping us do that in a cost-effective manner. Their technology allows us to model workloads, server farms and storage and backup technology in an efficient way."



VMware, Inc. 3401 Hillview Ave Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Copyright © 2008 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,961,806, 6,961,941, 6,880,022, 6,397,242, 6,496,847, 6,704,925, 6,496,847, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,944,699, 7,069,413, 7,082,598, 7,089,377, 7,111,086, 7,111,145, 7,117,481, 7,149,843, 7,155,558, 7,222,221, 7,260,815, 7,260,820, 7,268,683, 7,275,136, 7,277,998, 7,277,999, 7,278,030, 7,281,102, 7,290,253; patents pending. VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies

Hitachi is a registered trademark of Hitachi, Ltd., and/or its affiliates in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries. Universal Storage Platform, Lightning 9900 and Thunder 9500 are trademarks of Hitachi Data Systems Corporation.

