IDC
Analyze the Future

I D C A N A L Y S T C O N N E C T I O N

Gary Chen
Research Manager, Cloud and Virtualization System Software

The Value of Virtualizing Remote and Branch Offices

August 2014

Virtualization began in datacenters because of the massive savings and measurable paybacks from consolidation, but today the agility benefits from virtualization have become equally important. These benefits are driving more organizations to virtualize their remote-office/branch-office (ROBO) sites. Virtualized servers allow much more flexibility and standardization as well as enable advanced features like high availability (HA) to be implemented more easily and cost-effectively. This makes virtualization valuable for smaller environments like ROBOs, and makes managing distributed and remote sites more cost-effective, as infrastructure can be manipulated primarily through software rather than physical access to devices being required for incremental changes.

The following questions were posed by VMware to Gary Chen, a research manager in IDC's Cloud and Virtualization System Software program, on behalf of VMware's customers.

Q. What are customers' top challenges when it comes to their ROBO infrastructure?

A. IDC research shows that availability, security, regulatory compliance, and data backup/recovery are top priorities in enterprise IT infrastructure, including organizations with remote offices and branch offices. When you look at ROBO sites, they're generally the "endpoints" for a business and include important activities such as interfacing with customers in locations such as retail stores or manufacturing sites producing the company's core products. A big challenge with these endpoints is that they're physically distributed and remote, with possibly a large number of sites, sometimes hundreds or thousands. To further complicate matters, most ROBO sites have little or no IT staff readily available and some sites present harsh operating conditions for computer equipment.

The availability challenge for ROBOs has a few unique aspects. Workloads are run locally in ROBOs so that sites can (and often need to) function independently and not be reliant on remote infrastructure. While that lends a certain kind of sustainability, ensuring high availability locally is not trivial. The physical compute facilities in a ROBO are usually not as inherently resilient as in a corporate datacenter. In addition, implementing classic high-availability solutions has traditionally been complex and costly, making it difficult to apply to a wide scope of workloads.

The other major concerns are regulatory and corporate policy compliance, as well as security, in which configuration management is an underlying problem. With lots of remote infrastructure, it can be tough for companies to push out consistent, compliant, and secure configurations and, more importantly, maintain these configurations over time — and evolve and update these systems in a normal cadence.

Last, data backup is the other big worry. Data is unique to a company and something that can never be fixed or recreated once lost, thus companies worry about data protection obsessively. With ROBO, you often have a lot of crucial data being generated in the remote locations and you have to store and replicate this data locally as well as centrally.

Q. How can virtualization help with some of these ROBO challenges?

A. Virtualization significantly improves availability, compliance, and networking. In fact, IDC's *Virtualization and Cloud Study* shows that 44% of all virtual machines (VMs) are configured for high availability. Even without automated HA, the flexibility and agility of VMs allow for both planned and unplanned downtime as VMs can be evacuated from the physical server if needed and temporarily relocated very easily. This allows admins to dynamically reconfigure the infrastructure as needed in case of a hardware failure or simply for upgrades and maintenance — and to do so remotely.

Virtualization can also help with the "configuration drift" problem for better compliance and security. With advanced virtualization products, everything is controlled from a central management console. Using this console, an admin can push down a standard hypervisor configuration and then use templates to ensure consistent guest virtual machine images. Monitoring tools in the console can monitor configurations over time and alert admins if something gets changed. A VM snapshot is another great tool that has a variety of uses, like easily rolling back VMs to a known configuration. It can also allow for non-disruptive and consistent backups.

Virtual network switching is another aspect of virtualization that can improve compliance and security. Virtual infrastructures have software-based virtual network switches embedded in the hypervisor, which allows a lot of flexibility for admins to easily create and isolate networks, often a crucial requirement in many regulations.

Q. What is the typical hardware configuration in the ROBO site, and how might these environments be impacted by virtualization?

A. Since ROBO sites have modest infrastructure resources, consolidation is not the primary driver for virtualization as it would be in a datacenter. ROBO sites are instead virtualized for operational efficiencies and improvements, particularly in the areas of availability, standardization, and management. Nevertheless, companies with a large number of ROBO sites could still see significant savings through consolidation onto a smaller number of physical servers at each ROBO location.

Many ROBO sites have only a couple of lower-end tower form factor servers with one or two CPUs and internal direct-attached disks, with more advanced sites having a small networked storage appliance or RAID setup and a relatively narrow WAN connection. These servers often run in conditions much different than those in a datacenter with optimized environments, such as a backroom IT closet. In some cases, particularly with manufacturing companies, servers could potentially be located out in the manufacturing area, with systems residing in harsh conditions that may lead to premature server failures.

According to the IDC's Server Virtualization Tracker, the overall industry average virtual machine density today is about 10 VMs per server. At the high end, there are datacenter servers running 20–30 VMs or more per server. On the low end of the average, such as small and medium-sized business (SMB) and ROBO, it's common to only see 5–7 VMs per server.

©2014 IDC

Q. What are the three biggest considerations for pricing and licensing ROBO sites?

A. ROBO sites are very different from corporate datacenters, thus customers tend to look for a different set of pricing and licensing requirements:

- **Feature set.** The product needs to have a feature set that takes into account the specific needs of ROBOs without delivering (and charging for) unneeded functionality. While ROBO sites are smaller and less complex like a small to medium-sized business, SMB feature sets are too limited for ROBO, which may require features such as automated deployment and compliance tools. Likewise, packages for enterprise datacenters are too full featured for ROBO, and you would pay for features that are not needed. Ideally, you want a virtualization product that is somewhere between a full-strength datacenter product and an SMB product.

- **Licensing model.** The main consideration with licensing is the lower VM density of ROBO sites. It's common to see infrastructure software in enterprise datacenters being licensed per physical CPU. The assumption with virtualization is that you will get more out of that license by stacking as many VMs on that server as you can, which is a big driver for efficient consolidation. With ROBOs, the VM density is much lower and there just aren't enough workloads to increase utilization, so you really need to look at the overall cost per VM with whatever licensing model you're considering rather than the cost per server.

- **License distribution.** Some ROBO or SMB virtualization licenses have limitations on physical servers or CPUs per site. However, customers may have ROBO sites that vary greatly in size, for example a flagship store location versus a small rural store. Be careful to check that use right limitations won't be a problem, and look for flexibility to accommodate various future needs as well as economies and simplification of licenses that can be applied aggregately across all ROBO sites rather than on a site-by-site basis, especially if you're dealing with a large number of sites.

Q. Have customers explored a zero footprint configuration for ROBOs, utilizing public and/or private clouds? Is this a viable option?

A. A zero footprint scenario is intriguing for many customers as it would centralize infrastructure and eliminate the distributed management problem. While IDC certainly sees some workloads moving to private or public clouds, there are many challenges and many reasons why some workloads need stay onsite at the ROBO.

The primary challenges to using remote public or private cloud services for ROBOs are the speed, latency, and reliability of the WAN connection. Many ROBO sites, like retail store locations or manufacturing plants, are in areas with average to poor Internet service, with few options for improved service. Beyond performance, the availability and reliability of the connection would become critical, serving as a single point of failure. Many ROBOs today don't have redundant WAN connections, and those that do are generally only partially redundant. While the "last mile" network connection to the ROBO site itself may have multiple redundant connections, locality and provider monopolies often funnel these connections to the same network back end, leading to a single network dependency. This is the primary reason that workloads are run locally at ROBOs today. Many of these workloads, like a manufacturing control system, really need to run close to the physical location. Running a control system remotely would be quite risky as any WAN disruption could halt factory operations, not to mention the questions about performance and security about going over the Internet. Another similar example is a retail point-of-sale system in a store, where running locally is a much better model, allowing sites to function independently.

In some scenarios, like a corporate branch office with information workers, you may have workloads that could more effectively be centralized, particularly when you consider the mobility trends already transforming user behavior. Information workers tend to have viable options to reestablish connectivity should their primary network access point become unavailable for any reason.

And certainly there has been workload creep in many ROBOs, with nonessential workloads being stood up there that could be centralized or run in the cloud. The most likely scenario that IDC sees is a hybrid cloud deployment, where each workload will be examined for its appropriateness for cloud given its specific requirements and the available ROBO infrastructure and WAN. IDC believes a small number of existing ROBO workloads could be farmed out to a public or private cloud, but the majority will continue to run locally for the previously discussed reasons of availability, security, and performance.

Cloud adoption is undeniably an important trend that can improve services, save customers money, and simplify IT, and customers today need to begin to examine their options. Many customers have already started this journey, and while it isn't always easy to make blanket recommendations, there are some types of ROBO locations and workloads that could make sense for cloud. However, ROBO infrastructure is slow to change, so ROBO computing is going to be impacted by cloud much less and much later than datacenters. Because of the nature of ROBO computing, much of it is still going to run on-premises for the foreseeable future. Virtualization of ROBO infrastructure is going to offer much more immediate benefits with the current infrastructure limitations, and that's where you're going to see most of the ROBO transformation focused today.

ABOUT THIS ANALYST

Gary Chen is research manager for IDC's Cloud and Virtualization System Software. In this role, Mr. Chen focuses on server virtualization software and its transformational effect on the larger datacenter ecosystem of systems management, networking, storage, and security. His research also focuses on cloud system software, an extension of virtualization that provides the fundamental infrastructure for clouds.

ABOUT THIS PUBLICATION

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

COPYRIGHT AND RESTRICTIONS

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the Custom Solutions information line at 508-988-7610 or gms@idc.com. Translation and/or localization of this document require an additional license from IDC.

For more information on IDC, visit www.idc.com. For more information on IDC Custom Solutions, visit http://www.idc.com/prodserv/custom_solutions/index.jsp.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

©2014 IDC