Design Considerations for Virtualizing Citrix Provisioning Services

Introduction

Today, IT architects strive to virtualize most server workloads in the datacenter. When designing a virtual desktop solution with Citrix XenDesktop or Citrix XenApp, there are conflicting schools of thought with regard to implementing Citrix Provisioning Services as virtual servers. As architects begin to standardize on a virtualized platform, the imminent question that always appears to arise is 'Should Citrix Provisioning Services be installed on a Physical or Virtual Server?' In the past, it was always considered a best practice to have Citrix Provisioning Services installed on a physical server. Now with advances in virtualization technology, there are several options available that make virtualizing Provisioning Services a reality in the enterprise.

Citrix Consulting Solutions has been involved in multiple scenarios where Provisioning Services was successfully virtualized within XenDesktop and XenApp environments on all three major hypervisors (Citrix XenServer, Microsoft Hyper-V, and VMware vSphere). The most notable reference of these successful implementations was accomplished in partnership with Cisco and is documented within the Cisco Validated Design documents. In each circumstance, the Provisioning Services virtual server was designed to ensure that it was adequately able to handle the assigned workload and ensure that the virtual Provisioning Services server was not the major restricting factor within the environment.[1]

Based on the experience of the Citrix Consulting Solutions team, this document provides detailed design considerations for virtualizing Provisioning Services such as:

- Ensure that the hypervisor host is able to distribute processing power across multiple CPUs.
- A 10Gbps network is the most conducive environment for virtualizing Provisioning Services and the respective Provisioning Services network traffic.
- If a 10Gbps network is not available, consider link aggregation at the hypervisor level to provide more available bandwidth for the virtual Provisioning Services machine.
- Consider utilizing SR-IOV or Pass-Through to minimize the virtualization overhead associated with network intensive virtual machines, such as Provisioning Services.
- Always configure Provisioning Services in a high-availability configuration with multiple virtual machines distributed across different hypervisor hosts.
- Virtualize Provisioning Services on an x64 version of Windows to take advantages of Windows System Cache.

In summary, Provisioning Services workloads can and have been successfully virtualized in both XenDesktop and XenApp deployments if the design considerations outlined in this document are thoroughly evaluated.

[1] The amount of physical memory available within the physical hypervisor host was the most restrictive factor in scaling the environment.

Considerations for Virtualizing Provisioning Services

The decision to implement Provisioning Services on a virtual machine is a decision that needs to be made based on several considerations with the balance of administration ease, cost savings, and end user performance all in mind. The following sections outline those key design considerations associated with deploying Provisioning Services on a virtual server:

Hypervisor

For wide spread adoption, Provisioning Services is hypervisor agnostic, able to be virtualized on XenServer, vSphere, and Hyper-V. Virtualizing the Provisioning Services server can provide administrative and cost benefits such as the ability to share resources with other virtual machines on the same infrastructure, create a simplified business continuity plan and quickly implement additional Provisioning Services servers to address demand. But along with the cost and administrative benefits associated with virtualizing Provisioning Services, IT administrators must consider the overhead that virtualization adds into the environment. The architecture of each of the major hypervisors is different, but the fact still remains that virtualization adds an additional layer of overheard that can affect a Provisioning Services virtual machine performance. Each hypervisor needs to handle the process of routing the network traffic requests between the virtual machine and the physical network adapters on the hypervisor server.

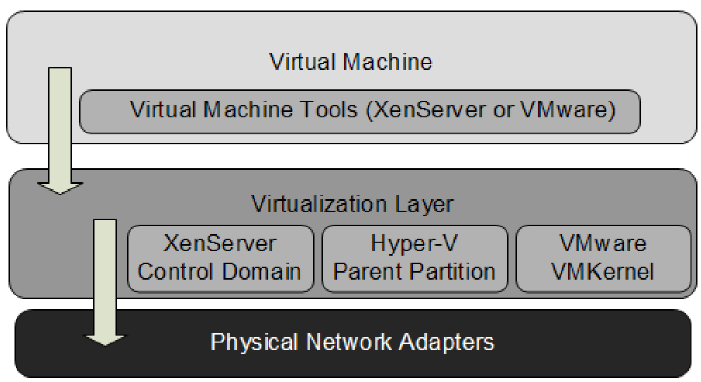

Figure 1: Hypervisor Architecture

A traditional hypervisor has a virtualization layer that acts as a broker between the virtual machines and physical hardware devices, such as the network adapter. If a virtual machine needs to send a packet across the network within the hypervisor architecture, the virtualization layer interprets the request from the driver on the virtual machine, processes the request against the driver for the physical resources and transfers the request or packet to the physical network adapter.   In instances where there is a high demand for network resources within the virtual machine, such as true with Provisioning Services, the hypervisor overhead may become more apparent as the amount of instances where the virtualization layer is required to broker requests between virtual to physical resources increases. The demand of processing these types of requests is typically reflected in more intense CPU cycles, therefore reflecting a spike in CPU utilization. The following sections will provide more design considerations to better address this concern.

Networking

As previously mentioned, Provisioning Services can tax networking resources during peak load such as simultaneous boot-up or logins; therefore, Provisioning Services will greatly benefit from increased bandwidth and direct access to the network card. For this reason, it is important to understand what existing options are available within a virtual environment to help reduce the hypervisor overhead and increase throughput.

The network demand for machines managed by Provisioning Services is greatest during the machine boot-up process. In order to stream the boot-up of a Windows 7 image, Provisioning Services needs to send approximately 210 MB across the network. Individually this data volume is not significant. However, in the event that a large number of desktops are simultaneously booted (e.g. a restart after a datacenter maintenance window) the network can become temporarily saturated. For example, multiply 210 MB for a single Windows 7 boot by the number of desktops that could be simultaneously starting in a XenDesktop environment to begin to understand the potential network demand. If the Provisioning Services write cache is configured for server-side versus the more common client-side or 'cache on device's HD' option, then the network traffic is further increased. If the network in the datacenter does not support 10Gbps Ethernet, architects may choose to aggregate multiple 1Gbps network adapters to provide more through-put. Link aggregation must be configured at the hypervisor level to be leveraged by a virtual Provisioning Services machine. Within each hypervisor, the following methods are supported for NIC bonding:

- Citrix XenServer: XenServer provides Source Load Balancing (SLB) in an active/active mode, but only supports load-balancing of virtual machine traffic across the physical NICs. (http://support.citrix.com/article/CTX124421). XenServer will load balance multiple VMs across multiple NICs, but a single VMs traffic is not split across two physical NICs (i.e. A single VM cannot get the combined data throughput of both NICs).

- Microsoft Hyper-V: Microsoft does not officially support link aggregation or NIC teaming within Hyper-V since it's a third party technology, but Microsoft does recommend working directly with that hardware vendor to determine interoperability with Hyper-V. (http://support.microsoft.com/kb/968703)

- VMware ESX: VMware ESX 4.x supports 802.3ad and LACP in static configurations. LACP has both a static and dynamic mode, but only the static mode is supported in ESX. (http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1001938)

Alternatively, a virtual Provisioning Services machine running within a blade environment can be configured so that it only provisions virtual machines within its own individual blade enclosure.

In most cases, the blade chassis will have a dedicated 10Gbps switch shared among the blades servers within the chassis. If a single Provisioning Services virtual machine is configured within the blade chassis, that machine can utilize the backplane in the chassis without ever sending traffic across the datacenter network.

When utilizing a virtualization solution, network traffic passes through the hypervisor layer so design decisions such as number of simultaneous boot-ups and write cache configurations all need to be evaluated as part of the network demand on Provisioning Services. If the choice is made to virtualize Provisioning Services on XenServer, one of the design considerations that may alleviate the networking limitations is SR-IOV.

SR-IOV

SR-IOV (Single Root I/O Virtualization) is a hardware function of the network card that enables an adapter to offer virtualization-safe hardware virtual functions (VF). These VFs can be directly assigned to a virtual machine utilizing Intel's Virtualization Technology for Direct I/O (VT-d). This technology is crucial because it allows a single network adapter to communicate directly with multiple virtual machines, eliminating the need for multiple physical adapters. As previously mentioned, there is a CPU overhead associated with routing network traffic through the hypervisor layer. With SR-IOV there is no hypervisor penalty for I/O since the virtual machine interacts directly with the physical device itself, therefore removing the traditional hypervisor overhead associated with brokering virtual to physical network adapter requests. Although SR-IOV is extremely beneficial to a virtualized Provisioning Services deployment, it is a new technology that is only available in XenServer with Intel VT-D processors and a limited number of network adapters. At the time of this document, VMware vSphere and Microsoft Hyper-V did not currently support SR-IOV nor was the roadmap for support announced. While adoption for SR-IOV is growing and could greatly improve a virtualized Provisioning Services workload, driver support for the VF adapters is still limited.

The following graphs show how the CPU utilization of the hypervisor can be conserved with SR-IOV within XenServer:

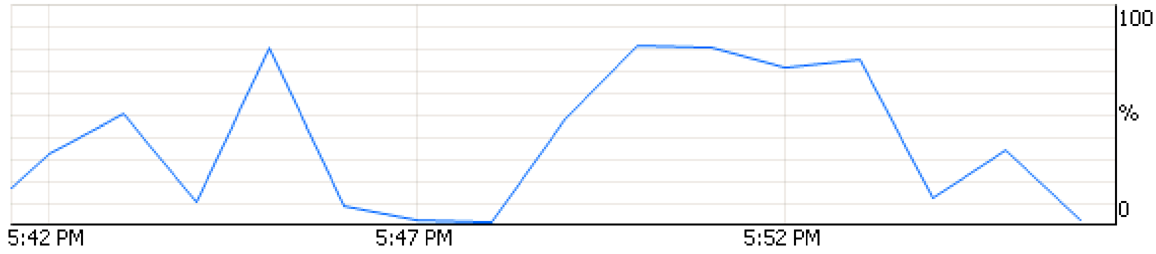

Figure 2: CPU Utilization Streaming 300 Desktop VMs with SR-IOV OFF

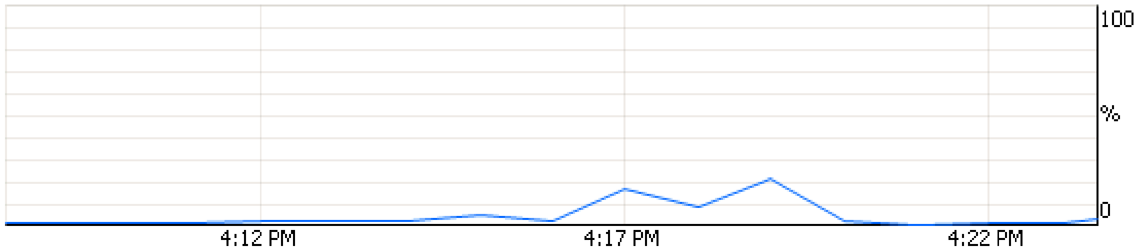

Figure 3: CPU Utilization Streaming 300 Desktop VMs with SR-IOV On

It should be mentioned that although SR-IOV is very beneficial in conserving hypervisor resources, there are some other design considerations that must be taken into account. The SR-IOV network adapters' virtual function (VF) is directly associated to an individual virtual machine. Within XenServer, VFs are configured independently within the XenServer CLI while the configurations for the virtual interface functions and bridges in dom0 (XenServer control domain) are contained in the virtual machine template. While this may be problematic for moving virtual machines between hosts, Provisioning Services includes its own High Availability mechanism that can be utilized to overcome this issue.

Full Device Pass-Through

As mentioned earlier, SR-IOV provides a virtual device that is able to be directly assigned and passed-through to a VM. Full Device Pass-Through works in the exact same manner as SR-IOV, allowing a physical device to be directly assigned to a VM. Unlike SR-IOV though, Full Device Pass-Through only allows a physical network adapter to communicate directly with a virtual machine in a one-to-one ratio (SR-IOV supports a many virtual machines-to-one physical network card ratio). This form of pass-through capability dedicates a single virtual machine to a specific card or port on the network adapter and eliminates the virtualization layer involvement in the I/O for that specific virtual machine. While Full Device Pass-Through does provides similar functionality as SR-IOV, bypassing the virtualization layer, there are a limited number of slots that a physical host can support therefore restricting the number of virtual machines that can utilize Full Device Pass-Through in a single host. With virtual machine density per server steadily increasing, the use of a dedicated network adapter per virtual machine is limited and still dependent on hardware specific drivers within the virtual machine. In addition, once a virtual machine is directly connected to a physical device, the virtual machine will lose the ability to dynamically move between hosts in a high availability configuration. To overcome this limitation, an administrator must have multiple virtual Provisioning Servers for High Availability across multiple hosts. In the event that a Provisioning Services virtual machine with Full Device Pass-Through enabled has a failure, then the secondary Provisioning Services virtual machine will manage the requests until the failure can be resolved. This functionality is available within VMware ESX 4 through the technology known as VMDirectPath I/O supported with Intel's VT-D or AMD's IOMMU technology and needs to be configured on the ESX host.

Virtual Machine Queue

Virtual Machine Queue (VMQ) is a feature available within Hyper-V. Unlike Pass-Through there is no direct relationship between the physical device and the virtual machine, but there is a performance enhancement to gain with the process of routing packets between the virtual machine and the physical network adapter. With a traditional hypervisor, the virtualization layer is responsible for sorting the incoming packets based on their virtual machine destination. This routing process introduces additional CPU cycles therefore resulting in a virtualization overhead. Virtual Machine Queue establishes a dedicate queue on the network adapter and the sorting process happens on the network adapter reducing the overhead on the hypervisor. VMQ must be enabled on the network adapter and configured within Hyper-V.

vDisk

When virtualizing Provisioning Services, the overall design of the environment should attempt to minimize network traffic when possible. One other item that needs to be carefully considered is the Provisioning Services vDisk configurations. The first decision is the physical location of the vDisk which will typically be on the local disk of the hypervisor server or a shared storage device. By having the vDisk on local storage, latency and network traffic are reduced within the environment and the virtualization overhead of processing the network traffic is also decreased.

Alternatively, the vDisk could also be placed on shared storage to offer increased accessibility by multiple servers. The best practices design recommendation is to ensure that the Provisioning Services environment is leveraging the Windows System Cache. With the ability to use the System Cache and a shared vDisk, the shared vDisk is only read from physical disk for the first request. Once a second request is made for the same vDisk, Provisioning Services will retrieve the data from the System Cache memory as long as there is enough memory available within the virtual machine (more information on memory configuration is available in this Citrix Consulting article CTX125126). Therefore, the design recommendation is that when virtualizing Provisioning Services, it is important to ensure that an x64 version of Windows is used for the virtual machine hosting Provisioning Services to leverage the Windows System Cache. If shared storage is preferred within the virtual machine environment, the actual storage location of the vDisk can be on either a block level storage device or a CIFS share (more information on vDisk storage locations is available in this blog).

Administrators must weigh the options carefully to determine the solution that works best in their environment, but either way the disk I/O from the vDisk requests on local or shared storage can affect other virtual machines if their data is located on the same disk. Therefore, it is considered a best practice to have the vDisk on a disk that is dedicated to the Provisioning Services virtual machine.

Testing Metrics

Citrix Consulting recommends that even with all the best practices for virtualizing Provisioning Services deployed, the environment should still be tested to ensure that end user performance is optimized. When testing the environment key statistics must be gathered from both the hypervisor and inside the Windows host of the Provisioning Services machine to evaluate performance. From the hypervisor, the following generic components should be monitored:

- CPU Usage and/or Utilization
- Memory Available, Used, & Ballooning (if applicable)
- Network Send, Receive, & Retransmits
- Disk Reads, Writes, & Queue Length

Within the Provisioning Services machine, the following performance monitoring metrics should be monitored (Ensure that the performance monitoring process is not affecting the virtual machine performance):

- CPU - % Processor Time, % Usage, % Interrupt Time
- Page File - % Usage
- Memory - Available & Committed
- Physical Disk – Read, Writes & Queue Length
- Network – Send & Receive
- TCP - Retransmits

Conclusion

In conclusion, Citrix has long recommended that Provisioning Services be installed on a physical server as opposed to a virtual machine. The physical server installation does provide a much simpler solution and mitigates risks associated with critical virtual machines all battling for the underlying hypervisor physical resources, but the decision to virtualize Provisioning Services on a large scale has proven successful in several Citrix Consulting deployments. To achieve optimal results, the best practices discussed in this paper should be implemented to ensure a virtualized Provisioning Services is deployed in a way that benefits both administrators and end users. The design implemented must converge with the ecosystem of the overall virtual machine environment to ensure acceptable performance on both the hypervisor and virtual machine layer. Ultimately, the decision to virtualize Provisioning Services should be made based on internal testing and an in-depth evaluation of the environment's configurations and workload.

References:

- Using SR-IOV with Intel® Ethernet Server Adapters: http://www.intel.com/support/network/adapter/pro100/sb/CS-031492.htm
- How to Create a Separate Provisioning Services Network http://support.citrix.com/article/CTX120955
- Provisioning Services Memory and Storage Considerations http://support.citrix.com/servlet/KbServlet/download/24250-102-647628/Advanced%20Memory%20and%20Storage%20Considerations%20for%20Provisioning%20Services%20v3_0208.pdf
- Provisioning Services and CIFS Stores - Tuning For Performance http://community.citrix.com/display/ocb/2010/11/05/Provisioning+Services+and+CIFS+Stores+-+Tuning+For+Performance
- VMware Configuration Examples and Troubleshooting for VMDirectPath http://www.vmware.com/pdf/vsp_4_vmdirectpath_host.pdf
- Best Practices for Configuring Provisioning Server on a Network http://support.citrix.com/article/CTX117374
- Microsoft Using Virtual Machine Queue http://technet.microsoft.com/en-us/library/gg162704(WS.10).aspx
- Blog on Virtualizing Provisioning Services http://community.citrix.com/pages/viewpage.action?pageId=162955756

Revision History

Revision	Change Description	Updated By	Date
0.1	Document Created	Consulting Solutions	January 25, 2011
0.2	Implement Feedback	Consulting Solutions	January 27, 2011
0.3	Additional Feedback	Consulting Solutions	January 28, 2011
1.0	Final Draft	Consulting Solutions	February 2, 2011

About Citrix

Citrix Systems, Inc. (NASDAQ:CTXS) is the leading provider of virtualization, networking and software as a service technologies for more than 230,000 organizations worldwide. Its Citrix Delivery Center, Citrix Cloud Center (C3) and Citrix Online Services product families radically simplify computing for millions of users, delivering applications as an on-demand service to any user, in any location on any device. Citrix customers include the world's largest Internet companies, 99 percent of Fortune Global 500 enterprises, and hundreds of thousands of small businesses and prosumers worldwide. Citrix partners with over 10,000 companies worldwide in more than 100 countries. Founded in 1989, annual revenue in 2008 was $1.6 billion.