



NetApp®
Go further, faster



This article was originally published in the NetApp Tech OnTap Newsletter. To receive the newsletter monthly and enjoy other great benefits, sign up today at www.netapp.com/us/communities/tech-ontap.

Lessons from the Trenches:

Deploying Hyper-V R2 and NetApp

By Patrick Cimprich, Avanade

PATRICK CIMPRICH
VICE PRESIDENT AND
CHIEF ARCHITECT
GLOBAL TECHNOLOGIES
AND SOLUTIONS
AVANADE

Patrick Cimprich has 19 years of IT experience in the design of large-scale application and storage infrastructures. At Avanade, Patrick is the global lead for data center solutions and also looks after Avanade Dynamic Computing Services, a thriving global development, testing, and hosting platform often used as the main proving ground for recommended customer and partner configurations.

Before going into virtualization, let me provide a corporate background on Avanade. Avanade offers business technology services that connect insight, innovation, and expertise in Microsoft® technologies to help customers realize results. The company applies Microsoft expertise from its global network of consultants. This means that we often put new technology through its paces earlier than many other folks.

That's especially true for one of my teams, Avanade's Dynamic Computing Services (DCS). DCS is a global platform for development, stress testing, and proof of concept to support both internal efforts and customer engagements. Our work often leads to Avanade's recommended customer and partner solutions and configurations. I want to offer some guidance based on our recent experience using Microsoft virtualization technology (Microsoft Windows Server® 2008 R2 with Hyper-V™ and Microsoft System Center suite) and NetApp® shared storage. This article covers:

- Information about our current virtual server environment
- Our experiences moving to Hyper-V R2
- Building for performance

- Planning for live migration
- Storage tips for cluster shared volumes (CSVs)
- Server and storage management and virtual machine (VM) self-service portals

CURRENT VIRTUAL SERVER ENVIRONMENT

When I say that we run a “dynamic” computing services environment, here's an example of what I mean. In the time a new provisioning request or change might flow through Avanade's internal production environment, our group might do 10 or 20. We may have 30 to 40 different projects running in parallel. Our customers have just about every enterprise application in operation, so our VM environment also supports nearly every Microsoft and non-Microsoft application available. Our current Microsoft virtualization environment includes:

Virtual Machines

- 350 virtual machines:
 - 225 VMs running on Windows Server 2008 Hyper-V R2
 - 125 VMs running on Microsoft Virtual Server 2005 R2
- Plans to add another 200 to 300 VMs under Hyper-V R2 in the next two months

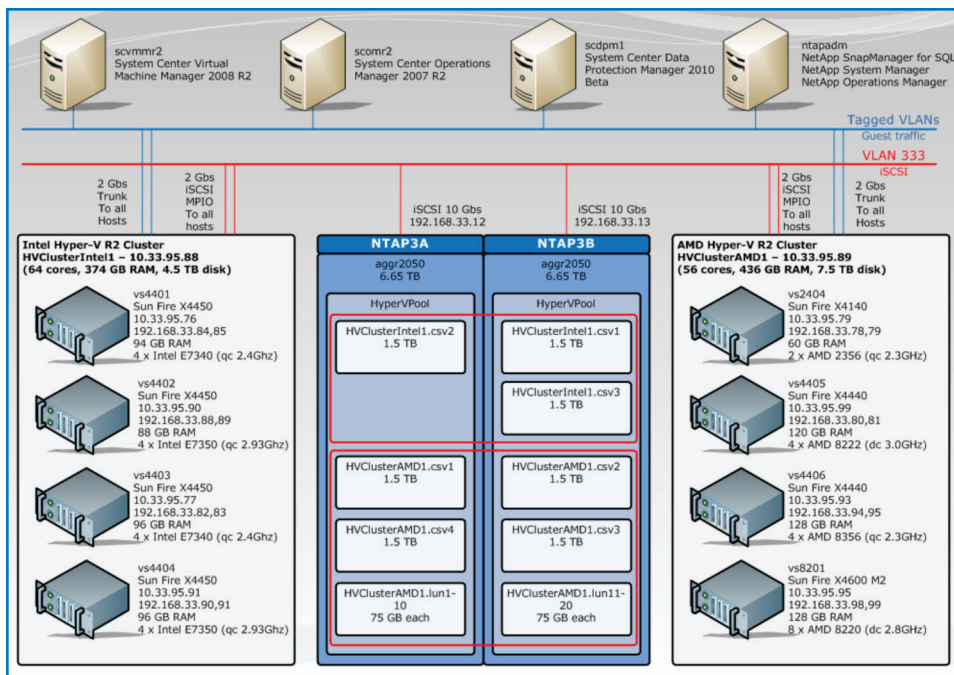


Figure 1) Avanade DCS Hyper-V R2 environment.

Source: Avanade

Servers

- 8 Sun Fire quad-core servers, each averaging 100GB RAM

Storage

- NetApp FAS3170C

We manage the environment with Microsoft System Center Virtual Machine Manager 2008 R2 and System Center Operations Manager 2007 R2. Figure 1 shows a conceptual view of this architecture, along with the layout of the underlying NetApp aggregates (shown in light blue at center), volumes (gray) and LUNs (white). We'll discuss more about our storage configuration later in this article.

OUR EXPERIENCE MOVING TO HYPER-V R2

We'd been looking forward to the Hyper-V R2 high-availability features like live migration and cluster shared volumes. Those were two critical features I wanted to see. I had already decided that when R2 became available we would go there, and go there fast. Within 24 hours of getting our hands on the R2 bits, we had our first cluster up with virtual machines running on it. Two weeks after, we'd moved our entire Hyper-V R1 system (and several hundred VMs) to R2.

The R2 migration experience was fantastic. Our daily use of CSVs and live migration

has already made quite a difference in our operations. Some of the top benefits we've seen since the R2 move include:

- **Easier, reduced maintenance.** Live migration lets us make all systems in the environment highly available. Now, instead of waiting to do hardware maintenance at night or on the weekend, live migration lets us do it midday without affecting production.
- **Simplified storage management.** Use of cluster shared volumes on the NetApp system has eliminated a lot of the activities we had with R1 (like assignment of GUIDs or the need to keep a VM on its own dedicated LUN). Instead, we can put our virtual servers in a big NetApp storage pool and let things sort themselves out more on their own.
- **Better service levels.** With live migration, we can treat all Hyper-V VMs as equal citizens. It's very rare now that we need to ask our customers for downtime. Live migration lets us just move things to another node, allowing customer systems to keep functioning.
- **Better performance.** Using CSVs has allowed us to harness the power of NetApp and its spindles for optimal performance. Other Hyper-V R2 features have also helped boost performance. Our R2 environment now runs between

20% and 30% faster than our R1 environment did (more about performance in the next section).

We've been doing large-scale performance testing of Microsoft applications on Hyper-V since it was in beta. There are obvious performance differences between applications, and careful storage design is still essential. But so far we've seen excellent performance. For example, we've had success running Microsoft SQL Server® in excess of 5,000 transactions per second on Hyper-V with NetApp. (For high transaction loads, dedicating an iSCSI connection to the virtual machine hosting SQL Server is one way to make sure that it gets the I/O it needs.)

BUILDING FOR PERFORMANCE

Several factors contributed to the 20% to 30% performance boost we've seen in our Hyper-V R2 environment. There were performance improvements in Hyper-V itself. Plus, as part of our recent move to new corporate headquarters, we were able to upgrade a number of components, including:

- **Upgraded network gear.** We put in place 10Gbs Ethernet connections for our NetApp FAS3170 as well as new, nonblocking connections to all servers. This new network provides

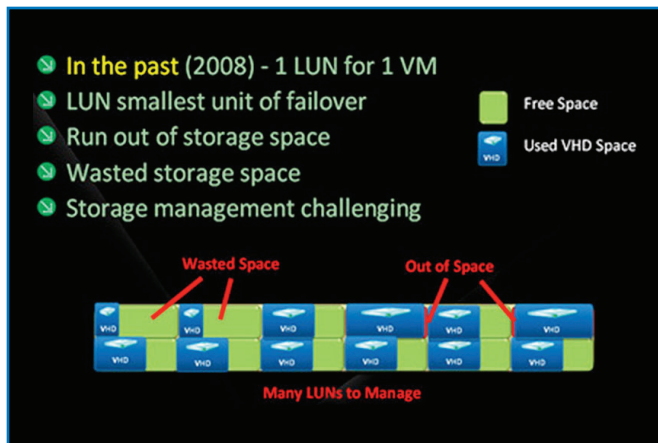


Figure 2) LUN to VM mapping prior to R2.
Source: Copyright Microsoft, TechEd 2009



Figure 3) LUN to VM mapping after R2.
Source: Copyright Microsoft, TechEd 2009

ample performance for the I/O-intensive demands of iSCSI storage traffic as well as live migrations.

- **Server consolidation.** We went from 15 to 20 smaller servers supporting Hyper-V R1 to 8 larger, beefier servers for R2.
- **Storage upgrade.** We deployed a NetApp FAS3170 system with a NetApp Performance Acceleration Module (PAM). This allowed us to speed I/O throughput and significantly reduce disk latency. We can now apply more disk spindles to I/O and also use PAM caching to speed random I/O requests.

PLANNING FOR LIVE MIGRATION

Some planning is definitely required to support live migration. This is part of a bigger discussion surrounding the design of your network and how many physical network adapters your servers need in order to properly support different types of traffic. Microsoft (www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=3fac6d40-d6b5-4658-bc54-62b925ed7eea) and NetApp (www.netapp.com/us/communities/tech-ontap/tot-hyperv-best-practices-0911.html) both offer excellent guidance on this topic specific to Hyper-V. Here are a couple of tips:

- **Design principle #1.** Treat your iSCSI storage network with the same care as you'd treat a Fibre Channel environment. That means taking the time to plan dedicated networks and network adapters like VLANs and subnets to handle storage-related traffic. If you do that, things are likely to work out great. If you don't, you'll definitely be in for some pain.
- **Design principle #2.** This principle fits closely with the first one. The rule of thumb is to make sure you aren't starving one network link in favor of another. If you're using iSCSI, this means segregating live migration traffic from other storage and end-user VM traffic. You do this by assigning live migration traffic its own dedicated network adapter. This also means that you need to monitor your network infrastructure to validate that links aren't overloaded. (See the sidebar at right for more about the number of dedicated network adapters recommended for use with Hyper-V.)

Live migration will absolutely dominate the network. When we live migrate virtual machines from one node to another, we see our network links go to near 100% utilization during the migration. Saturating trunked

gigabit NICs (or even single gigabit NICs) is generally difficult to do, but because so much of the live migration activity involves copying the contents of RAM from one system to another, extremely high network utilization is the norm. That kind of load puts a new level of demand on your servers and network infrastructure. It also exposes any weaknesses fairly rapidly.

We chose to segregate our iSCSI traffic onto two separate, dedicated links (see Figure 1 on the previous page). Now, during live migrations we're assured that storage communications are not affected.

HOW MANY ADAPTERS DO YOU NEED FOR HYPER-V?

NetApp reference architect Chaffie McKenna recently summarized the recommended number of physical network adapters needed for different Hyper-V R2 configurations. She also gives other R2 configuration guidance (www.netapp.com/us/communities/tech-ontap/tot-hyperv-best-practices-0911.html).

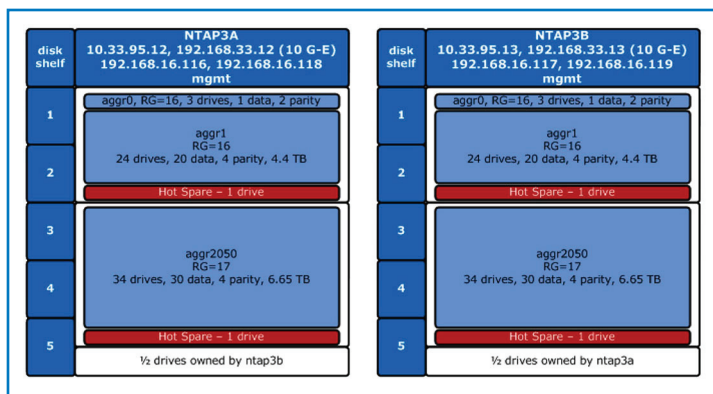


Figure 4) Layout of NetApp aggregates.
Source: Avande

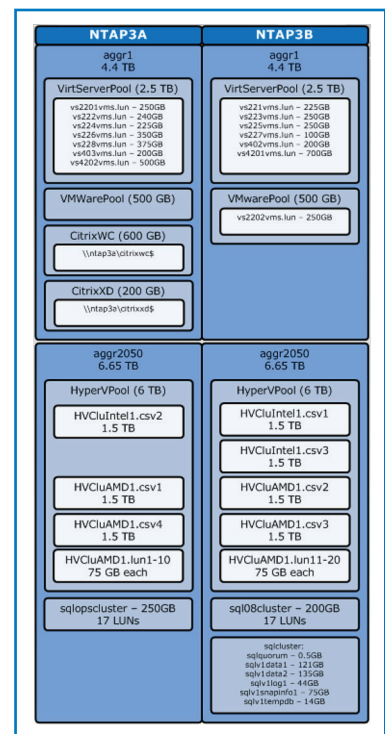


Figure 5) Layout of NetApp volumes and LUNs on each aggregate.
Source: Avande

LEARN MORE FROM AVANADE

Patrick and his colleague Andy Schneider have deep experience working with NetApp storage, Microsoft, and the world of server virtualization. Read about their experiences in these resources:

- Avande's 16-Node Virtual SQL Server Cluster (www.microsoft.com/caseStudies/Case_Study_Detail.aspx?casestudyid=4000006429)
- Avande Delivers Rapid, Turnkey VM Provisioning (www.netapp.com/us/library/customer-stories/avanade-virtualization.html)
- The Firm's Early Success with NetApp and Microsoft Hyper-V (www.netapp.com/us/library/customer-stories/avanade-story.html)
- Avande's Original "Lessons Learned" with NetApp and Microsoft Hyper-V (<http://www.netapp.com/us/communities/tech-ontap/avanade-hyperv.html>)

STORAGE TIPS FOR CLUSTER SHARED VOLUMES

Cluster shared volumes are another new feature in R2 that simplifies storage configurations and allows better support of Windows Server 2008 R2's Failover Clustering with Hyper-V. Before, in Hyper-V R1, there was a design limit that required the creation of dedicated LUNs to support highly available VMs. This meant you either had to sacrifice storage complexity for availability or vice versa. With R2, CSVs allow you to house many VMs on the same LUN while also making them highly available.

Microsoft senior program manager Steve Ekren did a great job explaining this difference in his TechEd 2009 session (<http://download.microsoft.com/download/8/C/2/8C21BAFE-3432-48D1-962A-F7A9DD54A2AC/Innovating%20High%20Availability%20with%20Cluster%20Shared%20Volumes.pptx>). An excerpt from his discussion appears in Figures 2 and 3 on the previous page.

Cluster shared volumes, however, offer both good news and bad news for storage system designers.

Here's the bad news: CSVs require a storage architecture that can be shared among all

cluster nodes. Collapsing multiple VMs multiplies the I/O load placed on your storage volume. This means you need to make sure that your physical storage system is designed to accommodate the extra load. Otherwise, bad things can happen, including reduced performance, application timeouts, and possibly even system crashes or data corruption.

Now here's the good news: Using the right storage system with careful design can overcome these issues. I've seen a big difference in this area with NetApp. We've had great experience with NetApp's integration with Microsoft Windows Server 2008 R2 with Hyper-V. Combining NetApp storage with clustered Hyper-V hosts has been very easy to do, much easier than with a lot of other tools.

Supporting multiserver access with NetApp is easy because of the tight integration NetApp has with Microsoft. Things like multiprotocol support and direct storage integration through SnapDrive® make Windows® cluster builds a breeze.

For us, looking at how the load would be balanced across NetApp aggregates and spindles was very important. Figures 4 and 5 show details of our configuration.

Here are a few more tips for storage design with CSV:

- **Make the most of NetApp technology.** We've leveraged NetApp technology (especially NetApp aggregates) to its fullest extent. Doing so has allowed us to virtualize our storage along with our servers. NetApp does a really good job of making this type of storage virtualization technology easy to use.

- **Break down the walls.** For many years, the storage industry has debated the value of segregating different workloads onto different disk spindles. When you have NetApp aggregates to work with, many of those principles don't necessarily apply. We were able to design a few large LUNs (or CSVs) on a couple of large NetApp aggregates. I encourage people to consider combining storage together to harness the power of aggregates.

In our case, one of the biggest aggregates contains about 35 physical disk spindles. This allows all those disk drives to work together in concert to provide aggregate performance that's shared across the CSVs. NetApp makes it easy to combine the power of those individual disk drives to support greater averaged I/O loads. NetApp also makes it very easy to modify the configuration later to support additional load if necessary.

- **Combine CSVs with NetApp deduplication.** Since rearchitecting our NetApp aggregates and volumes for cluster shared volumes, we've noticed

a big improvement in capacity savings by using NetApp deduplication. We are seeing an average deduplication ratio of 50% across our entire virtualization storage infrastructure. Given the size of our VM environment before deduplication (around 7 to 8TB), that's been a significant savings of 3 to 4TB, which supports our plans for adding many more VMs. Our use of the NetApp system with CSVs provided the right combination to really let deduplication shine.

CONCLUSION

Our continued experience with both Hyper-V and NetApp remains very positive. As we move forward, we will be exploring other integrated NetApp solutions for Microsoft virtualization. This includes the prospect of augmenting our data protection with NetApp SnapManager® for Hyper-V. We'll be looking at use of NetApp ApplianceWatch™ PRO 2.0 as well, for more centralized monitoring of NetApp storage via Microsoft System Center Operations Manager.

We have also been working with NetApp to develop an internal self-service portal that allows someone to request a virtual machine for a short period, such as three weeks. This portal leverages NetApp FlexClone® technology to rapidly provision the virtual environment. We'll be implementing that through a combination of Microsoft Virtual Machine Manager and Microsoft PowerShell™ scripting.

COMPLETE HYPER-V BEST PRACTICES

If you're deploying Hyper-V with NetApp storage, access to the latest best practices information is indispensable.

Download the latest detailed guides:

- Best Practices: SnapManager for Hyper-V (www.netapp.com/us/library/technical-reports/tr-3805.html)
- NetApp Storage Best Practices for Microsoft Virtualization (www.netapp.com/us/library/technical-reports/tr-3702.html)
- NetApp Implementation Guide for Microsoft Virtualization (www.netapp.com/us/library/technical-reports/tr-3733.html)

ADDITIONAL RESOURCES

Sometimes there's no better way to learn about NetApp efficiencies with server virtualization than to see them in action for yourself. These YouTube videos from NetAppTV show NetApp storage in action with Microsoft Server 2008 R2 with Hyper-V:

- Honey, I Shrank the LUNs (www.netapp.com/us/company/events/webcasts/hyperv-netapp-demo.html)
- NetApp SnapManager for Hyper-V (www.netapp.com/us/company/events/webcasts/netapp-smhv-demo.html)

Tech OnTap

SUBSCRIBE NOW

Tech OnTap provides monthly IT insights plus exclusive access to real-world best practices, tips and tools, behind the scenes engineering interviews, demos, peer reviews, and much, much more.

Visit www.netapp.com/us/communities/tech-ontap to subscribe today.



www.netapp.com

NetApp creates innovative storage and data management solutions that accelerate business breakthroughs and deliver outstanding cost efficiency. Discover our passion for helping companies around the world go further, faster at www.netapp.com.

© 2010 NetApp. All rights reserved. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, ApplianceWatch, FlexClone, SnapDrive, and SnapManager are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Hyper-V, Microsoft, PowerShell, SQL Server, and Windows are trademarks or registered trademarks of Microsoft Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.