A Performance Comparison of AMD Opteron™ Processors with Microsoft® Hyper-V™ Server 2008 R2

A Proof of Concept

Solutions Innovation Center
Dell, Inc.

January 2010

Contents

List of Figures

List of Tables

Section 1

Executive Summary

Virtualization is becoming almost a necessity due in part to economic reasons but also in order to take full advantage of the sheer processing power offered by today's technology. This eventually leads to the question "how far will my servers scale in a virtual environment"? To demonstrate the relative scaling capabilities of Dell™ PowerEdge™ Servers with AMD Opteron™ processors running Microsoft® Hyper-V™ Server 2008 R2, Dell conducted a proof of concept to compare several server configurations.

The Open Source DVD Store benchmark was used throughout with Microsoft Windows Server 2003 and Microsoft SQL Server 2005, running in virtual machines on the Hyper-V Server 2008 R2 host system. Dell compared the performance difference of the Six-Core AMD Opteron models -- the 8435 and 2435 – with the 2384 Quad-Core, and 2222 Dual-Core in the Dell PowerEdge M905 and M605 blade servers.

The tests proved that Hyper-V Server 2008 R2 was able to continue to support more virtual machines as additional processor cores and memory were available to the server. Although the number of virtual machines each processor or system is able to support will vary with the actual workload, the relative performance comparisons presented in this paper can be used to help determine which system and processor configuration might best fit a particular need.

Section 2

Introduction

When sizing a new deployment or consolidation effort one of the questions that always comes up is how many virtual machines can be supported by a server? Unfortunately there is not always an easy answer to this question because the workload of each virtual machine (VM) may vary greatly depending on what type of application is running. However, conducting a proof of concept test using a repeatable process allowed Dell to create a relative comparison between several different processor configurations.

For this study, a standard VM was created as a template that could be clone multiple times. This ensured that each additional instantiation of VM would run very close to the same workload as the others. The DVD Store test suite was used to exercise the VMs so that they would consume CPU cycles and memory as well as cause disk I/O. To show the scalability potential of each configuration, additional copies of the VM were started until the host server under test ran out of physical memory.

All of the VMs were stored on an iSCSI SAN consisting of a pair of Dell EqualLogic™ PS6000S arrays with solid state disks. The SSDs provided a fast disk subsystem, capable of supporting the I/O needs of all of the VMs without becoming a bottleneck. One large volume was formatted with the NTFS file system and held all of the VM virtual disk (VHD) and configuration files.

Two different server platforms were used to accomplish the test – the Dell PowerEdge M905 and the Dell PowerEdge M605 blade servers. The M905 is a four socket system while the M605 is a two socket system. The M905 is also able to support larger amounts of RAM than the M605. The M905 was used with the AMD Opteron™ processor Model 8435, while the AMD Opteron processor Models 2435, 2384, and 2222 were tested in the M605.

VM Workload

For the VM template, Windows Server® 2003 SP2 and Microsoft SQL Server 2005 SP3 were installed on a virtual machine. After installation of the operating system and SQL Server, Windows Update was run to apply the latest patches. The base or template VM contained one virtual CPU, 2048MB RAM, one network connection and a 15GB virtual disk. The medium sized database installation script was then run, creating a database of about 1GB. The total disk space occupied by the template VM was around 10GB, which allowed for growth of the database and log files during test runs.

The DVD store client application, or load generator, was executed with the following options:

`--target=VMx`	*the unique name of each VM*
`--pct_new_customers=0`	*disables creation of new customers*
`--db_size_str=m`	*run queries for the "medium" size database*
`--n_threads=2`	*execute 2 threads*

`--run_time=15` *run for 15 minutes*

In addition to the run time of 15 minutes per test cycle, the DVD Store default also adds an extra minute of ramp up time at the beginning of the test.

It was determined that two threads per VM would allow the template VM to operate at around 80% average CPU utilization as verified inside the VM (using Task Manager). Eighty percent is a typical “maximum” load for a production physical server.

Test Methodology

Each test cycle was run for 15 minutes. Several performance statistics were monitored and averages taken during the middle of the test run, giving the test plenty of time to stabilize so that fairly consistent results were achieved from run to run. When adding more VMs, at least 5 minutes was allowed for the VMs to start up and stabilize before running the next iteration of tests. After testing the maximum number of VMs for a particular hardware configuration, the cleanup script that is provided in the DVD Store package, was run against each VM. This script will essentially revert the data to its default state. Finally, a dbcc shrinkfile script was also run on every VM to shrink the logs. This allowed the VM recover the disk space where the logs had grown during the test.

After completion of a test cycle, the volume containing the VMs was detached and then reattached to the next test subject. This ensured the storage sub-system was consistent throughout the test variations.

When monitoring the system utilization with Hyper-V Server 2008 R2, it’s no longer simply a case of opening Task Manager and checking the performance tab. This is because the Hypervisor creates “partitions” and the desktop actually runs in a child partition just like every other VM that is launched. So to get accurate results, Performance Monitor must be used to monitor some new counters that are only available with Hyper-V Server 2008 R2. Specifically, the following counters can be used:

```
HyperV Hypervisor Virtual Processor\%Guest Run Time\_Total

HyperV Hypervisor Logical Processor\%Total Run Time\_Total

Processor\%Processor Time\_Total
```

The `%Guest Run Time` (%GRT) counter shows the percentage of CPU utilization inside the VM. This is roughly equivalent to logging on to each individual VM, opening Task manager and checking the performance tab except that the `_Total` counter shows the average of all running VMs.

The `%Total Run Time` (%TRT) shows the percentage of CPU utilization that each partition is using – or in other words, how much of the host CPU each VM is using. By monitoring the `_Total` counter, it can be determined how much of the host systems CPU is being consumed.

Section 3

DVD Store Database Test Suite

The DVD Store 2 database test suite (DS2) is an open-source database application that simulates an online e-commerce store – typical of one that might sell movies on DVD media. DVD Store has a backend database component, a web application tier, and a load generator program.

DVD Store was designed to utilize many advanced database features (transactions, stored procedures, triggers, referential integrity, etc.) while keeping the database easy to install and understand. In the case of this proof of concept, the DS2 workload is being used as a stress tool to generate CPU cycles, disk I/O, and consume memory in the virtual machines.

For these tests, only the backend database component was used and not the web tier. The "medium" data set which creates a database of approximately 1 GigaByte (GB), was installed on the virtual machines. The command-line interface was used to execute the client workload against each VM workload. All of the client load generators were executed from an M905 server running Windows Server® 2008 R2.

DVD store is downloadable (for free under the GNU open-source license) from:

http://linux.dell.com/dvdstore

Section 4

Microsoft Hyper-V Server 2008 R2

Microsoft Hyper-V Server 2008 R2 builds on the architecture and functionality of Windows 2008 Server Hyper-V, adding new features and flexibility. Hyper-V Server 2008 R2 enables creation of a virtualized server computing environment using Hypervisor technology included with Windows Server 2008 R2. A virtualized computing environment may be used to improve the efficiency of computing resources by utilizing more of the available hardware resources. Each virtual machine is a virtualized computer system that operates in an isolated execution environment. This allows Hyper-V to run multiple operating systems simultaneously on one physical computer.

Hyper-V provides the software infrastructure and management tools to create and manage a virtualized server environment allowing businesses to:

- Reduce the costs of operating and maintaining multiple physical servers by increasing the efficiency and hardware utilization of physical servers.
- Improve availability by configuring VMs on redundant failover host servers.
- Reduce the amount of time it takes to install, configure, and update systems by using standardized virtual hardware and drivers which may lead to improvements in development and deployment efficiency.

Furthermore, reducing the number of physical servers may also lead to savings in physical space (floor space, rack space, etc.), and power and cooling.

Section 5

Solution Architecture

The architecture of the systems that were tested is shown in Figure 1. It included the following elements:

- 2 x Dell PowerEdge M905 – DVD store Load Generator and Hyper-V Host
- 3 x Dell PowerEdge M605 – Hyper-V Host
- 2 x Dell/EqualLogic PS6000s – Data store
- 1 x Force10 Networks S50N – iSCSI SAN Network switch

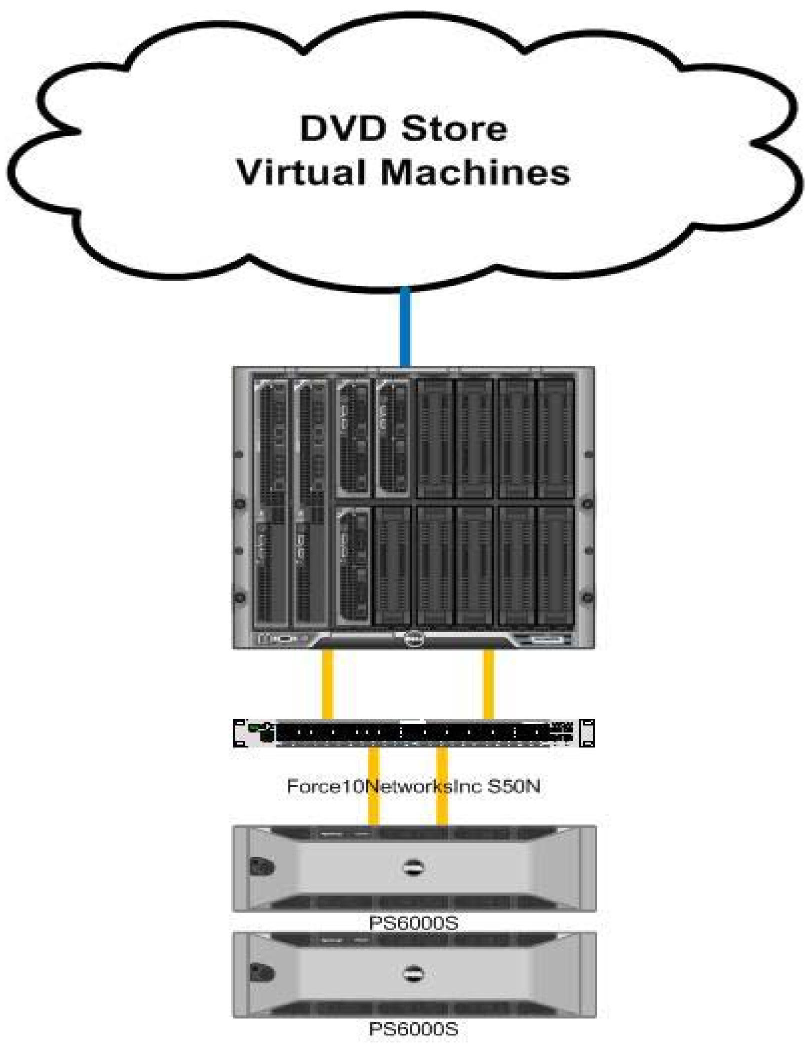

Figure 1: Hyper-V POC Solution Hardware

On all systems, the internal disks booting the Windows operating system consisted of two 73GB 15K rpm SAS drives attached to the onboard SAS 6i/r and configured as a RAID 1 volume. Each server ran Windows Server 2008 x64 R2. On all of the Hyper-V hosts, the Hyper-V role was installed and configured. The Dell/EqualLogic Host Integration (HIT) Kit 3.3.2 was installed so that the iSCSI volume containing the virtual machines could be attached and mounted.

Details of the server hardware used are shown in Table 1:

Qty	PowerEdge™ Server	AMD Opteron™ Processor	CPU Speed (GHz)	# Cores	RAM (GB)	Usage
1	M905	8435	2.6	24	96	DVD store load generators
1	M905	8435	2.6	24	96	Hyper-V Host
1	M605	2435	2.6	12	32	Hyper-V Host
1	M605	2384	2.7	8	32	Hyper-V Host
1	M605	2222	3.0	4	16	Hyper-V Host

Table 1: Hyper-V POC Server Details

To store the virtual machine data that was used during the tests, two Dell EqualLogic PS6000S iSCSI arrays were configured in a single pool, and as one large volume. Details are shown in Table 2:

Array	Number of Disks	Disk Type	RAID	Firmware Version	Total Storage Available
PS6000S	16	50GB SSD	10	4.21	309GB
PS6000S	16	50GB SSD	10	4.21	309GB
Total					**618GB**

Table 2: Hyper-V POC Storage Details

The Dell/EqualLogic iSCSI arrays were connected to a Force10 Networks S50N. The PowerEdge blade servers also connected to this switch through a pass-through module on Fabric B of the blade chassis. Each server had two iSCSI initiators. The iSCSI SAN Switch details are shown in Table 3:

Switch	# ports	Firmware Version
Force10 Networks S50N	48	7.8.1.3

Table 3: Hyper-V POC iSCSI SAN Switch Details

Section 6

Test Results and Analysis

M905 (96GB) – AMD Opteron™ processor Model 8435 (24 cores)

With 96GB RAM a maximum of 45 virtual machines were able to start up, at which point the system consumed a total of 94.3GB of RAM and Hyper-V would not allow any additional VMs to start.

# VM	%Total Run Time	%Guest Run Time	Avg Disk IOPs	Total Memory Used	VM Manager % Avg CPU used per VM
1	7	68	297	5.8	2
10	36.9	65.6	2148	24	3
20	67.6	62.5	3950	44	3
30	81.6	49.8	5184	64.5	2
40	85.3	39	6439	84.7	2
45	89.3	35.7	6878	94.3	2

Table 4 – M905 (24 cores) Results

After 30 VMs, the `%Total Run Time` (%TRT) starts to flatten out slightly which indicates that the host CPU is starting to get saturated. At this point the `%Guest Run Time` (%GRT) also starts to decrease, indicating that each virtual CPU is getting less CPU cycles from physical host. Compared to the results from the M605 with 12 cores, the peak capacity is double, if not slightly more (possibly due to the fact that the M905 has more than 2X the RAM over the M605).

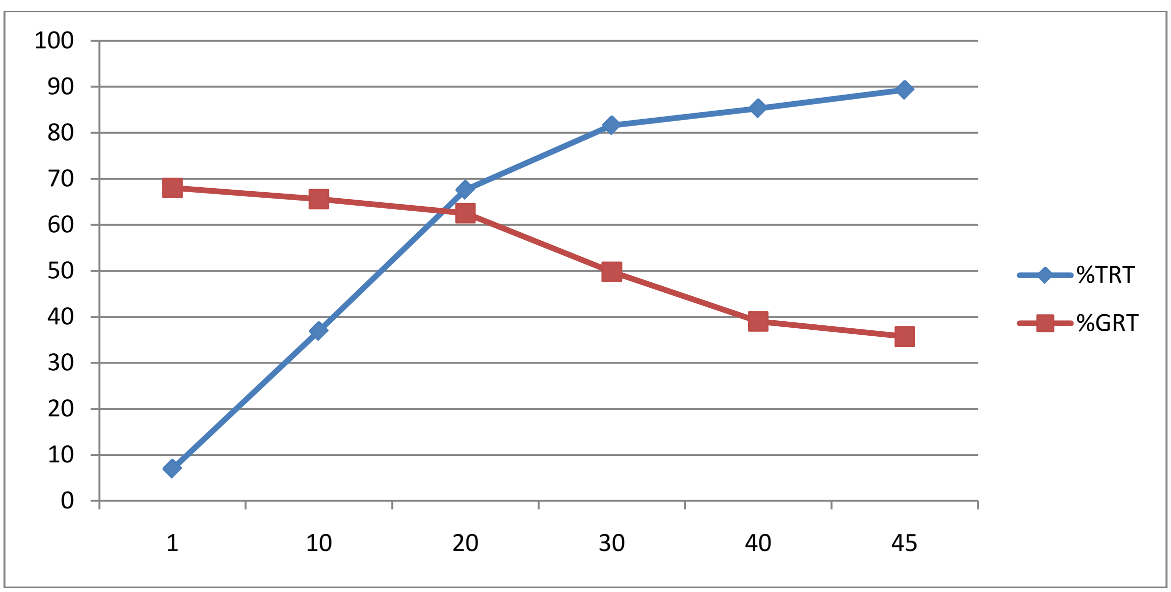

Figure 2 – M905 (24 cores) CPU Performance

Examination of the average disk I/O performance shown in Figure 3 reveals that the disk subsystem continues to scale nearly linearly up to 40 VMs. This confirms that the volume that the VMs were located on was not a bottleneck in the tests. Because each VM had 2GB of RAM and the DVD Store database was approximately 1GB, which allowed the VM to cache much of the database into memory. As a result, the I/O actually reaching the disk was predominantly writes (roughly 4X the amount of read I/O hitting the disk subsystem).

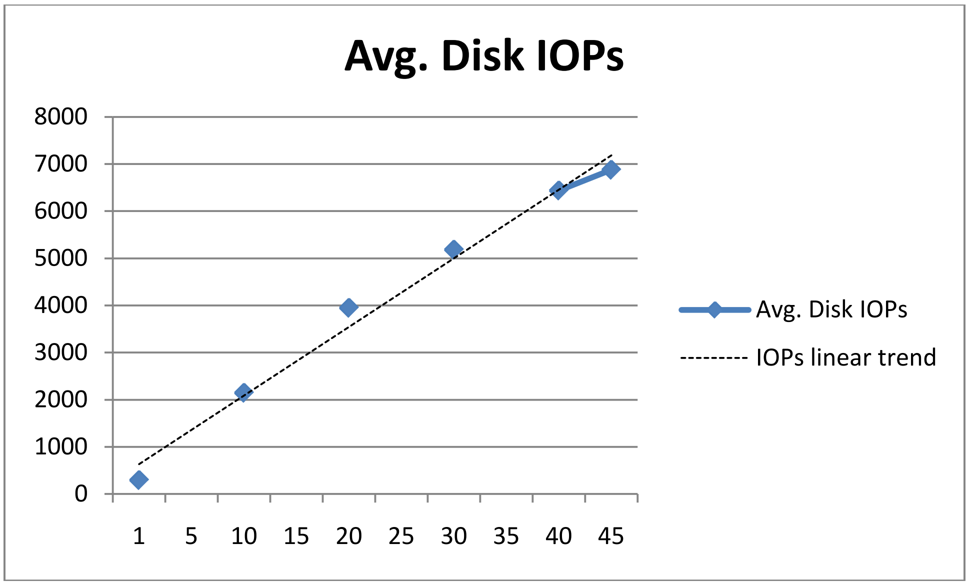

Figure 3 – M905 (24 cores) Disk Performance

M605 (32GB) – AMD Opteron™ processor Model (12 cores)

With 32GB RAM in the M605, 14 virtual machines were able to start, at which point the system consumed a total of 30.1GB of RAM and Hyper-V would not allow any additional VMs to start.

# VM	%Total Run Time	%Guest Run Time	Avg Disk IOPs	Total Memory Used	VM Manager % Avg CPU used per VM
1	8.1	73.3	252	3.9	6
5	36	66.5	1260	11.9	6
10	69.4	66.5	2115	22.2	5
14	81.9	57.7	2500	30.1	4

Table 5 – M605 (12 cores) Results

At 14 VMs, the %TRT is at 81.9. Again, this is an indicator that the host CPU is starting to get saturated. Figure 4 also shows that the %GRT flattens out with 5-10 VMs, but then starts to drop again by 14 VMs, indicating that each virtual CPU is getting less CPU cycles from physical host. If the results are compared to the M905 with 24 cores, it can be determined that our peak capacity is slightly less than half of the point where the M905 hits its peak. If the maximum of 64GB RAM had been installed in the M650, it would have been able to start a few more VMs, however with the %TRT already above 80 it is typically above the average load that would be recommend load for this configuration.

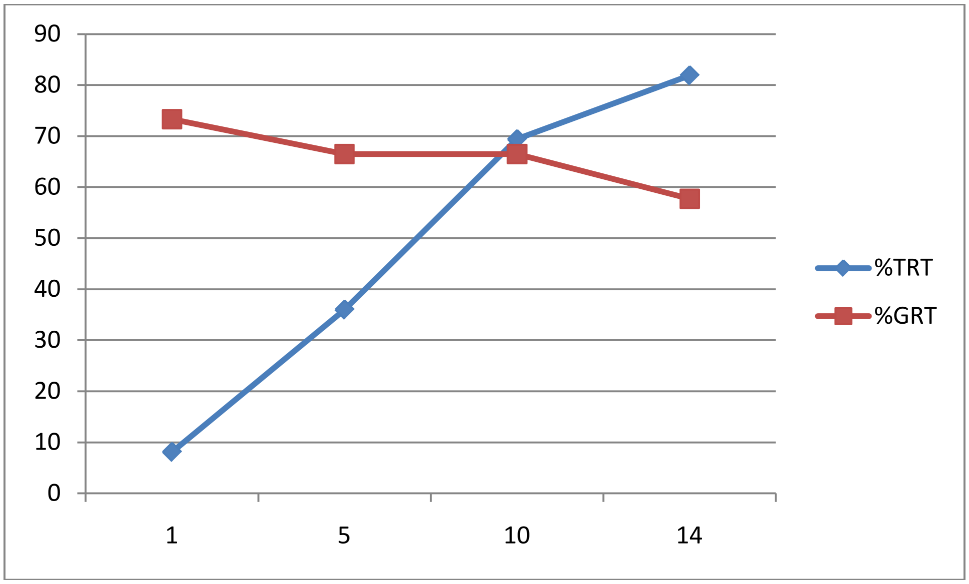

Figure 4 – M605 (12 cores) CPU Performance

M605 (32GB) – AMD Opteron™ processor Model 2384 (8 cores)

With again 32GB RAM in the M605, 14 of the virtual machines were started, at which point the system consumed a total of 30.1GB of RAM and Hyper-V would not allow initialization of any additional VMs.

# VM	%Total Run Time	%Guest Run Time	Avg Disk IOPs	Total Memory Used	VM Manager % Avg CPU used per VM
1	13.5	72.8	284	4	10
5	53	66.5	1146	12	9
10	83.3	54.5	1967	22.1	7
14	85.8	40.1	3174	30.1	6

Table 6 – M605 (8 cores) Results

By 10 VMs, Table 6 shows that the %TRT already exceeds 80%. Again, this is an indicator that the maximum CPU usage at the physical host is near saturation. It also shows that the %GRT is already dropping steadily by the time 5 VMs are running. A comparison of the results of the M605 with 12 cores (Figure 4) shows that %TRT and %GRT intersect at around 10 VMs, while the M605 with 8 cores shown in Figure 5 has its %TRT and %GRT intersecting at around 6 VMs. By comparison, this shows us that the M605 scales fairly linearly when additional cores are available.

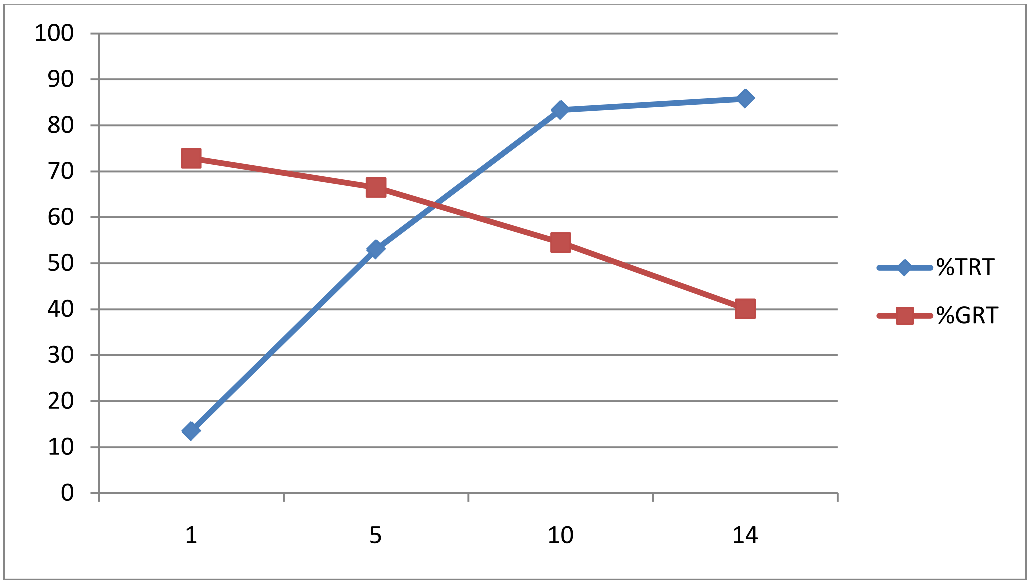

Figure 5 – M605 (8 cores) CPU Performance

M605 (16GB) – AMD Opteron™ processor Model 2222 (4 cores)

With 16GB RAM in the M605, only a maximum of 7 virtual machines could be started, at which point the system consumed a total of 15.1GB of RAM and Hyper-V could not initialize any additional VMs.

# VM	%Total Run Time	%Guest Run Time	Avg Disk IOPs	Total Memory Used	VM Manager % Avg CPU used per VM
1	26.3	64.1	250	3.2	19
3	64.4	61.1	897	7.2	19
5	83.3	46.9	1168	11.3	14
7	87.8	35.1	1486	15.1	11

Table 7 – M605 (4 cores) Results

After 5 VMs, the %TRT is well above 80% indicating the host system is nearing the maximum CPU usage at the physical host. Table 7 also shows the %GRT is dropping steadily after 3 VMs. Looking at the results of the M605 with 8 cores in Figure 5, shows that %TRT and %GRT intersect at around 10 VMs while the M605 with 4 cores shown in Figure 6 has its %TRT and %GRT intersecting right around 3 VMs. Again, this shows that the M605 scales fairly linearly when additional cores and memory are added.

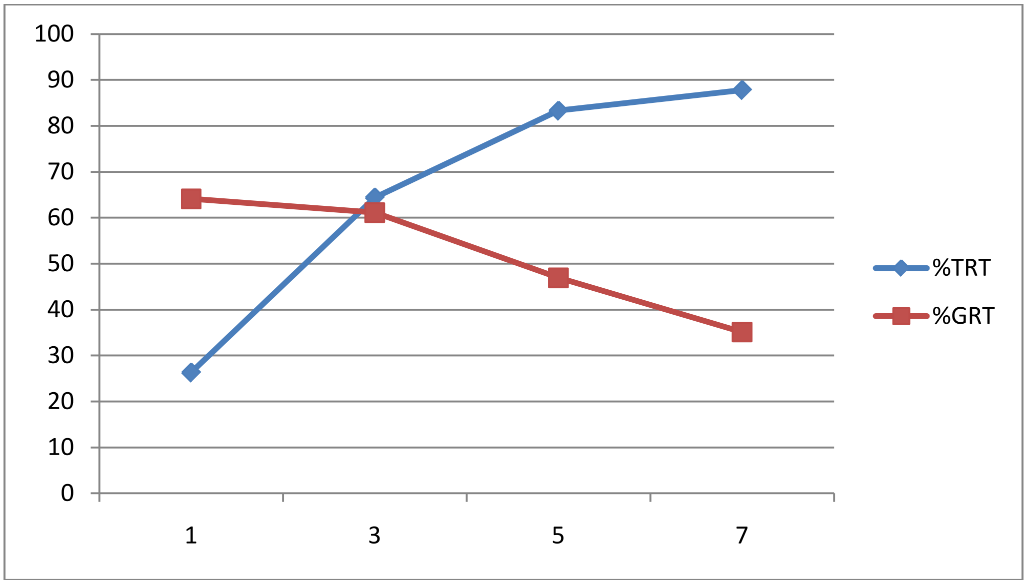

Figure 6 – M605 (4 cores) CPU Performance

Combined Results

Finally, the combined graph in Figure 7 shows all four test configurations side-by-side. The x-axis (horizontally) represents the number of virtual machines running on the Hyper-V host, while the y-axis (vertically) represents the CPU performance collected from the %TRT counter.

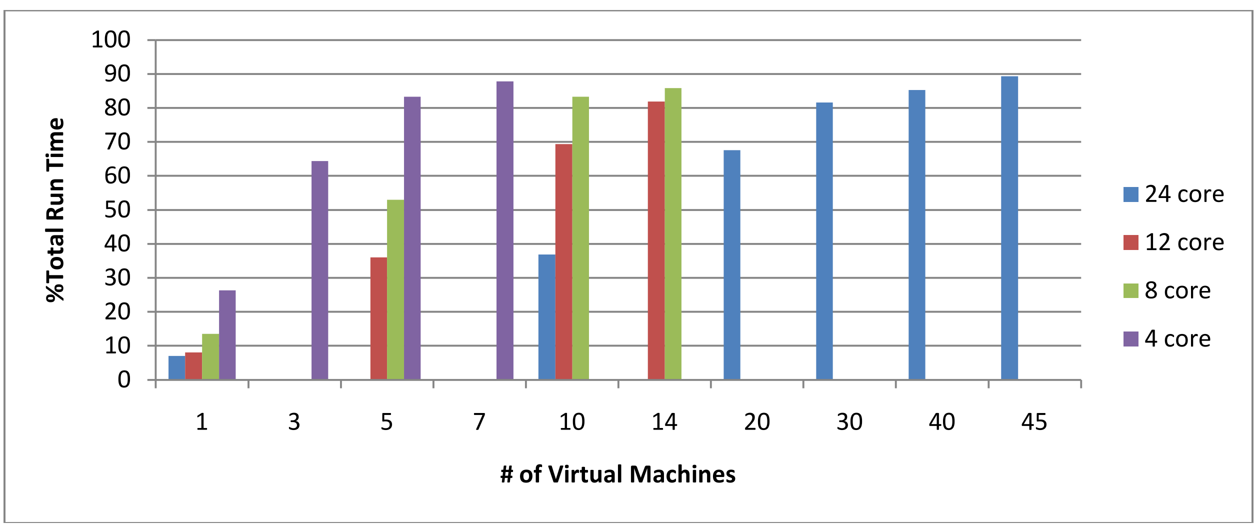

Figure 7 – Combined CPU Performance

Section 7

Conclusions

The tests performed in the proof of concept showed that as more processor cores and memory were available to the physical server, Hyper-V Server 2008 R2 could continue to scale and support additional VMs. Although the number of virtual machines each processor or system is able to support will vary with the actual workload, the relative performance comparisons presented in this paper can be used to assist in sizing of new deployments, processor and memory upgrades, or consolidation efforts.

For example, a prediction could be made with some confidence that if an existing Dell PowerEdge M605 with 4 cores was running 5 virtual machines at 80% average CPU utilization, a processor or system upgrade to one running 8 or 12 cores would most likely be able to support twice as many VMs.   Or for further consolidation, upgrading to the Dell PowerEdge M905 with 24 cores should be able to support 20 or more similarly tasked virtual machines assuming that the system had adequate memory and disk I/O capacity to support the additional VMs.