V2V MIGRATION TO RED HAT ENTERPRISE VIRTUALIZATION ON THE DELL POWEREDGE R820

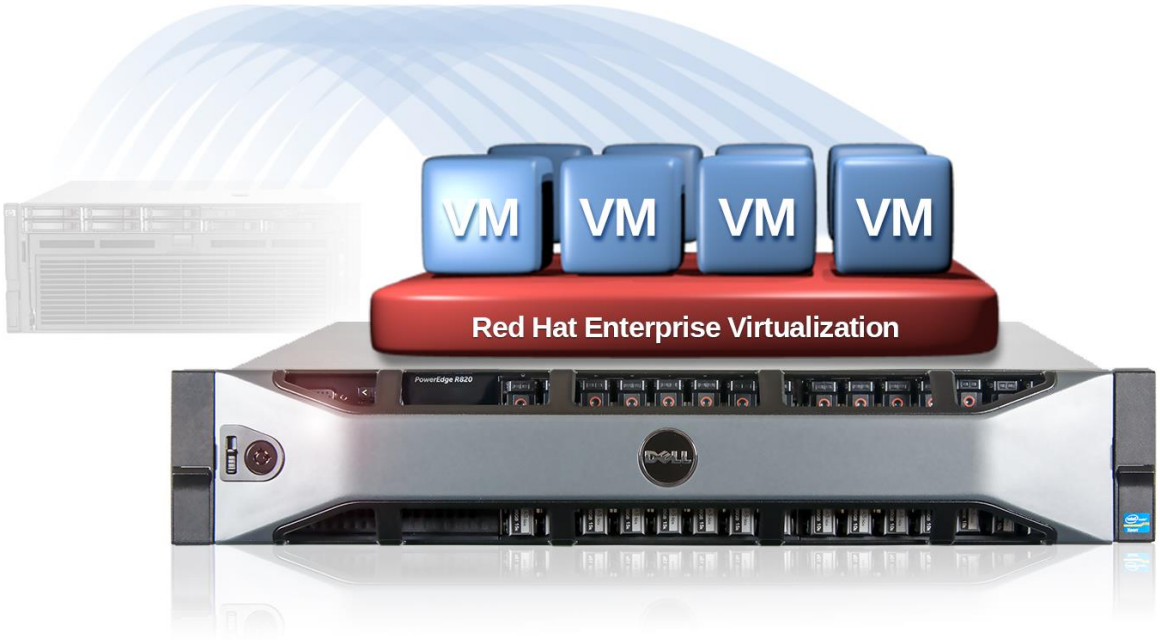

Principled Technologies®

redhat

DELL

A Principled Technologies migration guide commissioned by Red Hat, Inc.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Migration from one virtualized environment to another, called virtual-to-virtual migration (V2V), is a powerful method of reallocating or moving your virtualized workloads to meet business demands. Red Hat Enterprise Virtualization and the Red Hat virt-v2v migration tool make it easy for your existing virtual machines to migrate to the Red Hat Enterprise Virtualization platform. In our labs, we performed a migration of VMs on a legacy server to a Red Hat Enterprise Virtualization environment on a current-generation Dell™ PowerEdge™ R820 to show how straightforward the V2V migration process can be.

We also tested the database performance that the Dell PowerEdge R820 with Red Hat Enterprise Virtualization 3.1 delivered, and found that it delivered strong performance in a dense 2U form factor, which could allow you to reduce your rack and power footprint when you complete a V2V migration from similar performing 4U servers.

The Red Hat Enterprise Virtualization 3.1 Hypervisor is an open-source virtualization platform designed to support both Microsoft® Windows® and Red Hat Enterprise Linux virtual machines. The Red Hat Enterprise Virtualization subscription model may offer customers a lower TCO than other leading virtualization options on the market.

On the hardware side, the Dell PowerEdge R820 is a four-socket, 2U rack server, capable of supporting four processors with up to eight cores each and up to 1.5 TB of memory. Even with its slim 2U chassis form factor, the server can support 16 bays of 2.5-inch SAS, SATA, or SSD storage, with a maximum internal capacity of 16 TB.

In this document, we use this hardware and software combination to illustrate the simplicity of V2V migration from an environment with a different hypervisor to a Red Hat Enterprise Virtualization environment running on a dense 2U four-socket Dell PowerEdge R820.

Getting ready

Evaluating existing VMs for migration

The process of migrating your existing virtual machines from your existing hypervisor to Red Hat Enterprise Virtualization 3.1 is straightforward, but you must plan carefully prior to starting the migration. Considerations include, but are not limited to, the following:

- **Utilization levels and expected consolidation ratios.** Inventory your infrastructure, and gather details about the source virtualization

platform and hardware. Will your destination server have the resources to support the migrated VMs? Depending on the source servers, a platform such as the Dell PowerEdge R820 may be able to consolidate many source servers.

- **Supported source hypervisors.** Ensure that the virt-v2v migration tool supports the originating hypervisor.
- **Guest operating systems.** Ensure Red Hat Enterprise Virtualization 3.1 supports the guest operating systems targeted for migration. Currently, Red Hat Enterprise Virtualization 3.1 supports the following guests:
    - Red Hat Enterprise Linux 3 (32-bit and 64-bit)
    - Red Hat Enterprise Linux 4 (32-bit and 64-bit)
    - Red Hat Enterprise Linux 5 (32-bit and 64-bit)
    - Red Hat Enterprise Linux 6 (32-bit and 64-bit)
    - Microsoft® Windows XP® Service Pack 3 and newer (32-bit only)
    - Windows® 7 (32-bit and 64-bit)
    - Windows Server® 2003 Service Pack 2 and newer (32-bit and 64-bit)
    - Windows Server 2008 (32-bit and 64-bit)
    - Windows Server 2008 R2 (64-bit only)
- **VM documentation.** Document each target VM's specific networking details, application details, and other critical information.
- **Migration sizing.** Sizes of VM disks may vary, and the number of virtual disks may vary from VM to VM. Document sizing and ensure the network infrastructure is capable of handling the migration load during the allotted maintenance window.
- **Other infrastructure adjustments.** Evaluate other relevant items such as backup jobs, monitoring agents, and so on, to ensure that you follow critical business continuity policies.
- **Storage.** Note on what storage the target VMs currently reside. While the destination server hardware may provide a more powerful compute platform, the underlying storage subsystem plays a critical role in the performance stack. Ensure that there will be enough I/O on the destination storage connected to the Red Hat Enterprise Virtualization environment.
- **Networking.** Migrated VMs will obtain new virtual MAC addresses on the Red Hat Enterprise Virtualization platform, so ensure documentation exists for the physical and virtual switch environments, as well as guest operating system networking settings.

Prepared with this information, you can take steps to ensure a smooth V2V migration.

Preparing your environment

Red Hat Enterprise Virtualization, Red Hat Enterprise Virtualization Hypervisor, and Red Hat Enterprise Virtualization Manager

Red Hat Enterprise Virtualization is a virtualization environment based on Kernel-based Virtual Machine (KVM) technology, and designed by Red Hat to run enterprise-grade virtual workloads – whether the virtual machines are running Linux or Microsoft Windows. Red Hat Enterprise Virtualization features support for high-availability scenarios, live migration, live snapshots, live storage migration and management, extensive network management, and scheduling of system tasks. Red Hat Enterprise Virtualization is comprised of two different elements: Red Hat Enterprise Virtualization Manager and the Red Hat Enterprise Virtualization Hypervisor. Red Hat Enterprise Linux provides the foundation for both of these components.

Red Hat Enterprise Virtualization Manager is the management platform for the virtualization environment that runs directly on Red Hat Enterprise Linux and provides a Web-driven console for central management of the virtualized infrastructure. Red Hat Enterprise Virtualization Manager allows administrative control of the complete Red Hat Enterprise Virtualization feature set, including high availability (HA), live migration, live snapshotting, live storage migration, storage management, network management, and the system scheduler.

Red Hat Enterprise Virtualization Hypervisor is the hypervisor foundation of the Red Hat Enterprise Virtualization Manager solution. Red Hat Enterprise Virtualization Hypervisor, based on KVM, leverages hardware-based CPU virtualization extensions to improve performance.

Installing Red Hat Enterprise Virtualization Hypervisor and configuring the hypervisor

Red Hat Enterprise Virtualization Hypervisor 3.1 is straightforward to install. Installation of the hypervisor itself is as simple as booting to the install disc, choosing the Install Hypervisor option and target disk, setting a password, and rebooting on completion. For complete details on Red Hat Enterprise Virtualization Hypervisor installation, see Chapter III.8 (Installing Red Hat Enterprise Virtualization Hosts) in the Red Hat Enterprise Virtualization Installation Guide.[1] For the purposes of this document, we assume that Red Hat Enterprise Virtualization Manager installation and

[1] https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Virtualization/3.1/html-single/Installation_Guide/index.html

configuration are complete. For complete details on Red Hat Enterprise Virtualization Manager installation, see Appendix B in this document, or Chapter II (Installing Red Hat Enterprise Virtualization Manager) in the Red Hat Enterprise Virtualization Installation Guide.[2]

To configure the Red Hat Enterprise Virtualization Hypervisor server after installation, log into the Red Hat Enterprise Virtualization Manager Web management console, and add the new host. Enter the configuration information for the host, such as the hostname, IP address, and root password for the hypervisor machine, and click Approve. For complete details, see Appendix B. Next, configure your networking and set up your master storage domain, which can be Internet Small Computer System Interface (iSCSI), Fibre Channel, Network File System (NFS), or local storage. See the Red Hat Enterprise Virtualization Administration Guide for complete details on networking and storage domain configuration.[3]

Virt-v2v

The **virt-v2v** utility performs the V2V migration. Virt-v2v is capable of converting KVM, Xen, and VMware® vSphere®-based VMs running Red Hat Enterprise Linux 4, 5, and 6; Windows XP, Vista®, and Windows Server 2003 and 2008 to a Red Hat Enterprise Linux Open Virtualization Format (OVF) format. The converted VMs are set to use para-virtualized drivers when available.

Virt-v2v runs on any Red Hat Enterprise Linux 6 server and installs in three steps:

1. Subscribe to the appropriate channel on the Red Hat Network (RHN) – either the Red Hat Enterprise Linux Server (v.6 for 64-bit x86_64) or Red Hat Enterprise Linux Workstation (v.6 for x86_64) channel.
2. Install appropriate Windows-related packages, if you plan to convert Windows guests. The packages libguestfs-winsupport and virtio-win are critical and required for Windows guests' conversions. The libguestfs-winsupport for Red Hat Enterprise Linux Server 6 is located in the Red Hat Enterprise Linux Server V2V Tools for Windows (v. 6) channel. The virtio-win package is located in the Red Hat Enterprise Linux Server Supplementary (v. 6) channel. After connecting to the appropriate channel, run the following command to install the packages:
```
#yum install libguestfs-winsupport virtio-win
```
3. Install virt-v2v:

[2] Ibid.
[3] https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Virtualization/3.1/html-single/Administration_Guide/index.html

```
#yum install virt-v2v
```

Performing the virtual-to-virtual migration

Here, we detail how to perform a V2V migration from a legacy server running a different hypervisor to a new Dell PowerEdge R820 running Red Hat Enterprise Virtualization Hypervisor 3.1 managed by a server running Red Hat Enterprise Linux 6.3 and Red Hat Enterprise Virtualization Manager 3.1. For complete details, see Appendix C.

Pre-migration

Because a managed virtualization environment frequently requires guest agents, custom drivers, and specialized management tools to maximize performance, be sure to do the following:

- Remove any legacy drivers or tools required for your old hypervisor.
- Remove any management agents that are specific to your source environment.
- Create an NFS export domain, which will be the intermediary storage location before importing it into Red Hat Enterprise Virtualization Manager.
- Optionally, configure advanced network mappings and create a conversion profile in virt-v2v.conf. See Red Hat documentation for more details.[4]

When converting a Linux VM, there are additional steps that you need to take. The conversion process may install a new kernel and matching drivers on the VM. If the guest previously registered to Red Hat Network, the package management is auto-configured.

For a Windows migration, ensure the libguestfs-winsupport and virtio-win packages are installed on the host and running the virt-v2v conversion. In addition, for best integration with Red Hat Enterprise Virtualization, Red Hat recommends using the guest tools, so you should upload the installation ISO to Red Hat Enterprise Virtualization Manager. When the conversion process is complete, Red Hat Enterprise Virtualization Manager will use the guest tools ISO to install the Windows drivers.

Executing virt-v2v

Once you have done the prep work, executing virt-v2v is a simple command-line process. The example below assumes a v2v conversion from VMware ESXi to Red Hat

[4] https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Virtualization/3.1/html-single/V2V_Guide/index.html#sect-Preparing_to_Convert_a_Virtual_Machine

Enterprise Virtualization. Xen-based virtual machine conversions are performed in a similar manner. See Red Hat documentation for details.[5]

1. Edit the `/root/.netrc` file and add the IP address of your source host, where "hostname" is the hostname of the machine, and 192.168.1.100 is the IP address of the machine:
   ```
   hostname 192.168.1.100 login root password Password1
   ```

2. Run virt-v2v. In our environment, we used the following command line:
   ```
   virt-v2v -ic esx://192.168.1.100/?no_verify=1 -o \
   rhev -os RHEV-M:/vol/rhev_export --network rhevm VM01
   ```

This initiates the conversion process, or the migration of the existing VM from a different hypervisor to Red Hat Enterprise Virtualization. For further information on this process, see Appendix C or the Red Hat Enterprise Linux 6 V2V guide at https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Virtualization/3.1/html-single/V2V_Guide/index.html#sect-Converting_a_Virtual_Machine.

During the migration

For Linux guests, virt-v2v creates a new libvirt domain for the guest using the same name as the original. Virt-v2v cannot reconfigure a guest's network configuration at this time, so ensure that the configuration is valid on the first boot of the VM after migration. Virt-v2v will automatically complete the remaining driver changes for the storage and network components on the Linux VM.

When a Windows guest imports into Red Hat Enterprise Virtualization Manager, the driver install process completes upon system restart.

Due to the virt-v2v's almost entirely automated process, there is little for the administrator to do once the migration process is underway. Under the hood, the Linux and Windows guest operating systems reconfigure themselves to use virtio para-virtualized drivers. Figure 1 depicts the V2V migration process.

[5] https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html-single/V2V_Guide/index.html#sect-V2V_Guide-Converting_a_Virtual_Machine-Converting_Virtual_Machines

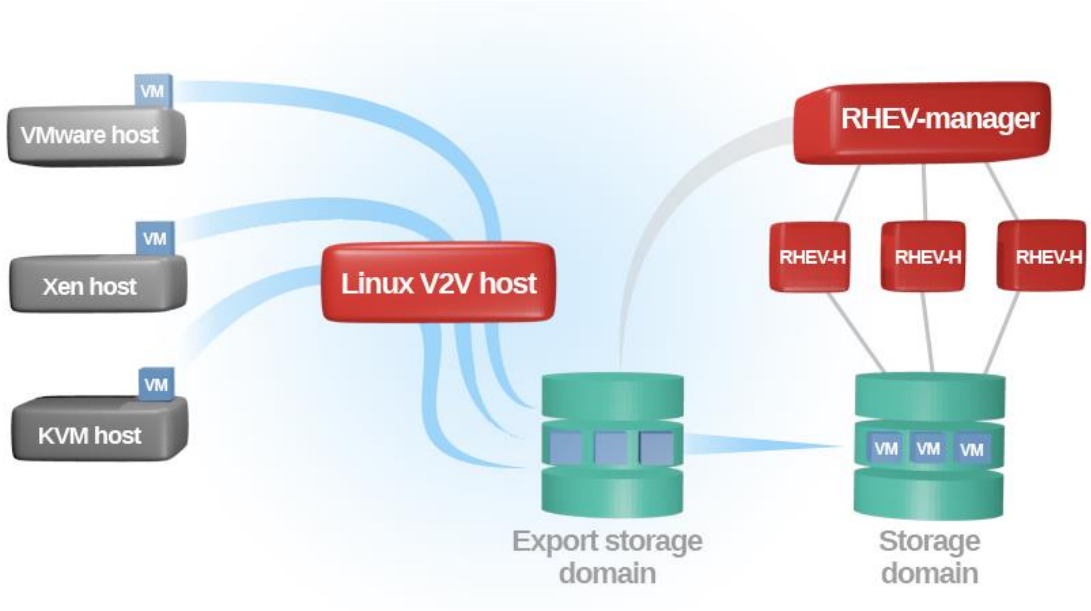

Figure 1: The flow of virtual machines through V2V process.

Importing the VM from the Export storage domain

Once the virt-v2v process completes, import the VM from the Export storage domain by completing the following steps:

1. Log into the Red Hat Enterprise Virtualization Manager Web console and navigate to the Storage tab. Each newly migrated VM will now be located under the VM Import tab in the Export domain used in the virt-v2v command.
2. To import the VMs into your Red Hat Enterprise Virtualization environment, select the VM you wish to import, and click Import.

Figure 2 shows the layout of the VM Import tab in Red Hat Enterprise Virtualization Manager.

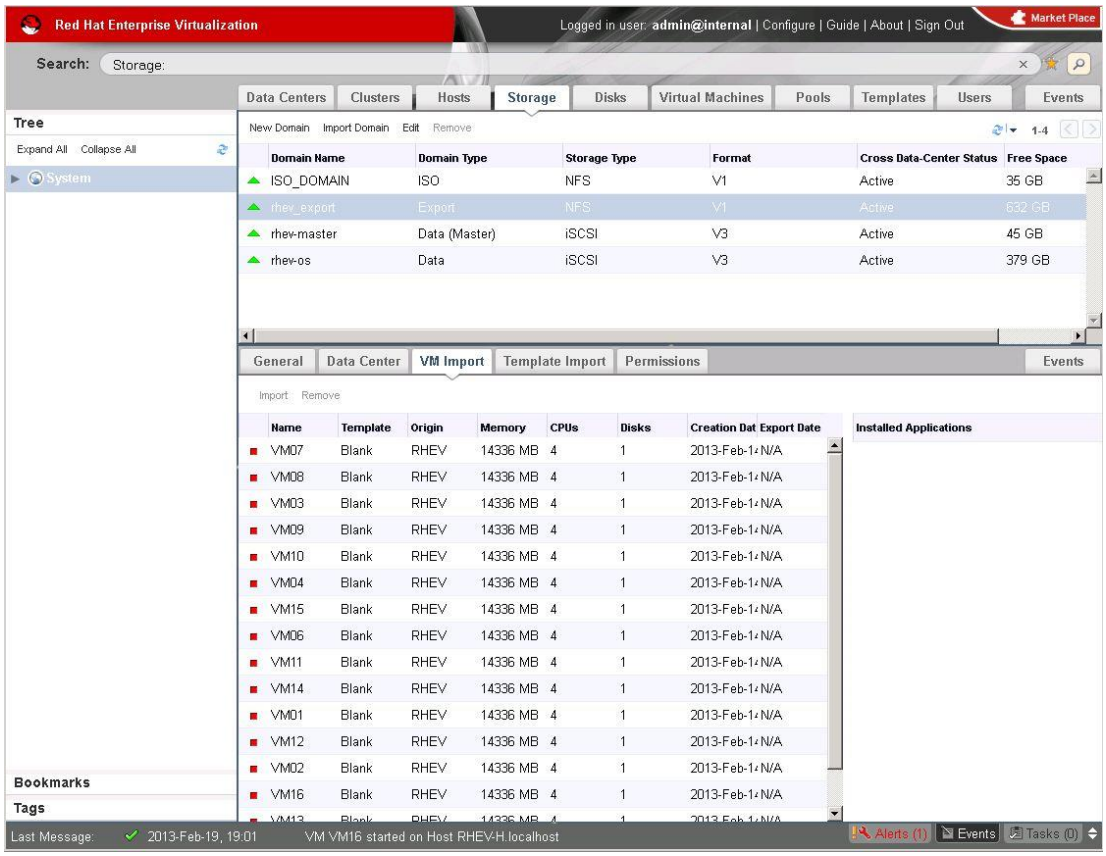

Figure 2: Viewing the VM Import tab under the Export storage domain in Red Hat Enterprise Virtualization Manager.

In the import wizard screen, you can view general information about a VM you select, including the installed OS, attached virtual disks, and virtual network configuration. Before initiating the import, select the Default Storage Domain, select the cluster, and select the storage domain for any attached virtual disks. By default, the disks will be placed in the VM's default storage domain, but in Red Hat Enterprise Virtualization 3.1 you can now attach multiple disks from separate storage domains to a single VM.

After you select the appropriate settings, click OK to import the VM. The VM will appear in the desired storage domain and cluster with a status of Image Locked, and a status of Down once the import process has completed. Multiple imports can be queued at once, limiting the required wait time for the administrator to import many VMs. Figure 3 shows the Import Virtual Machine(s) screen.

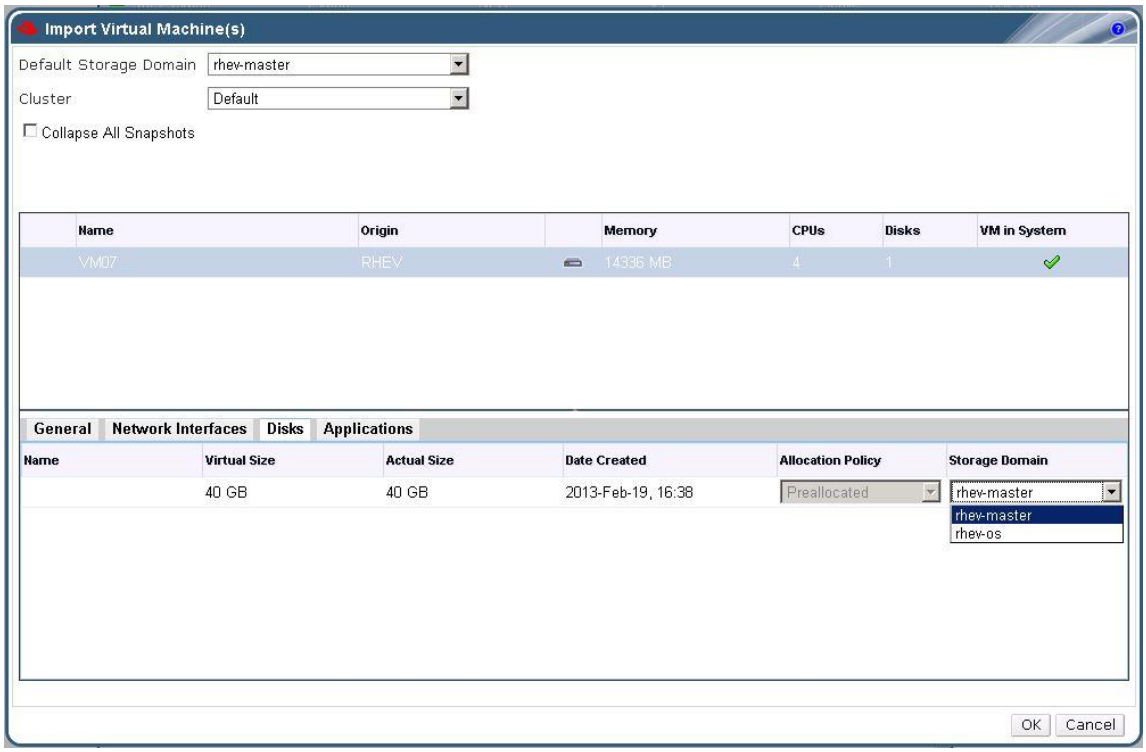

Figure 3: Importing the virtual machine and setting storage options.

Post-migration considerations

After the import steps are complete, verify that IP addresses, host names, clusters, and services are all configured correctly and communicating as expected. The only tasks that remain are to confirm that your management infrastructure is intact and ready to resume daily business operations. Tasks may include the following:

- Replacing drivers, guest tools, and management agents as needed on your VMs if necessary.
- Checking backup jobs, and verify that they are running properly in the new environment.

ABOUT OUR ENVIRONMENT

About the Dell PowerEdge R820 rack server

The Dell PowerEdge R820 (see Figure 4) is a quad-processor-capable 2U rack server, powered by the Intel Xeon processor E5-4600 product family. The R820 incorporates scalable network and I/O options designed to allow you to customize the server to meet the specific needs of your business.

Supporting up to four processors with up to eight cores each, the PowerEdge R820 can support up to 32 physical cores and 64 logical processors using Hyper Threading. In addition, the server supports up to 1.5 TB of memory, through 48 DIMM slots. With up to 16 2.5-inch hot-plug drive bays supporting SAS, SATA, or SSD technologies, the PowerEdge R820 has a maximum internal storage capacity of 16TB.

The I/O customization capabilities are flexible as well, supporting a variety of PCIe cards, PERC RAID controllers, and network controllers. The Dell iDRAC7 management add-in, combined with the Lifecycle Controller, allows for enhanced remote management of the server.

For configuration information for the PowerEdge R820 we used in our migration and test, see Appendix A.

Figure 4: The Dell PowerEdge R820 server.

About Red Hat Enterprise Virtualization for Servers

Red Hat Enterprise Virtualization for Servers is a virtualization environment based on KVM, designed to run enterprise-grade workloads – whether the virtual machines are running Linux or Microsoft Windows. Red Hat Enterprise Virtualization features support for high-availability scenarios, live migration, live snapshots, live storage migration and management, network management, and scheduling of system tasks.

About Virt-v2v

The **virt-v2v** utility performs the V2V migration. Virt-v2v is capable of converting KVM, Xen, and VMware vSphere-based VMs running Red Hat Enterprise Linux 4, 5, and 6; Windows XP, Vista, and Windows Server 2003 and 2008 to a Red Hat Enterprise Virtualization OVF format.

GREAT PERFORMANCE IN JUST 2U OF SPACE

When you are consolidating servers, perhaps legacy 4U servers to denser solutions such as the PowerEdge R820, you want to ensure that your new server/hypervisor solution can deliver the performance you require while taking advantage of new power-efficient technologies and space-saving design. A more compact, power-efficient solution such as the 2U Dell PowerEdge R820 solution could reduce the operational expenses you will incur over time.

To show that the Dell PowerEdge R820 solution could deliver strong performance using just 2U of space, we tested the Dell PowerEdge R820 with Red Hat Enterprise Virtualization running 16 VMs, with each VM containing a single 10GB PostgreSQL™ (version 9.2) database. Each VM ran Red Hat Enterprise Linux® 6.3 as the guest OS. We created a real-world e-commerce workload using the DVD Store 2.1 benchmark, which calculates results in orders per minute (OPM) that the solutions can process (see Appendix G for details). We used two Dell EqualLogic™ storage arrays for testing: a Dell EqualLogic PS6110X tray and a Dell EqualLogic PS6110XS tray.

As Figure 5 shows, the 2U Dell PowerEdge R820 with Red Hat Enterprise Virtualization 3.1 achieved a total of 512,522 OPM. If a comparable 4U solution delivers similar OPM, then the denser Dell PowerEdge R820 solution could double your performance per rack.

	Dell PowerEdge R820 solution
VM 1	30,592
VM 2	32,108
VM 3	32,696
VM 4	32,342
VM 5	30,618
VM 6	32,608
VM 7	33,159
VM 8	32,614
VM 9	31,257
VM 10	32,132
VM 11	32,308
VM 12	32,222
VM 13	31,303
VM 14	31,801
VM 15	32,173
VM 16	32,589
Total	**512,522**

Figure 5: OPM the Dell PowerEdge R820 solution processed per VM.

As the cost of powering and cooling servers rises, reducing ongoing operating expenses by investing in more power-efficient servers makes good long-term business sense. We measured the power consumption of the Dell PowerEdge R820 solution during the test, and found that it used an average of just 461W throughout our test duration due to power-efficient design. Figure 6 shows the average CPU utilization of the solution throughout our test and the power in watts it consumed. We measured CPU utilization by calculating the inverse of the idle percentage counter.

	Dell PowerEdge R820 solution
CPU utilization	79%
Average power draw in watts during the test window	461

Figure 6: CPU utilization and power consumption throughout our tests.

For complete details about our performance testing, see Appendices D through G.

SUMMING IT ALL UP

Migrating from your existing hypervisor and legacy hardware to Red Hat Enterprise Virtualization on a Dell PowerEdge R820 may be a great solution to maximize performance and lower costs, but many organizations consider such a migration too challenging. By choosing a hypervisor with an easy-to-use automated migration process, you can eliminate many of the concerns about a conducting a V2V migration. As we have shown through completing our own sample migration, performing a V2V migration from another hypervisor (including Xen, KVM, or vSphere) to a Red Hat Enterprise Virtualization environment on a new Dell PowerEdge R820 is a straightforward process that could possibly provide significant benefits for your server infrastructure.

APPENDIX A – SERVER CONFIGURATION INFORMATION

Figure 7 provides detailed configuration information for the test server.

System	Dell PowerEdge R820 solution
Power supplies (per chassis)	
Total number	2
Vendor and model number	Dell E1100E-S0
Wattage of each (W)	1,100
Cooling fans (per chassis)	
Total number	6
Vendor and model number	San Ace 60 9GA0612P1J611
Dimensions (h x w) of each	2.5" x 2.5"
Volts	12
Amps	1.50
General	
Number of processor packages	4
Number of cores per processor	8
Number of hardware threads per core	2
System power management policy	Performance per Watt (DAPC)
CPU (current-generation solutions)	
Vendor	Intel
Name	Xeon
Model number	E5-4650
Stepping	C2
Socket type	FCLGA2011
Core frequency (GHz)	2.70
Bus frequency (GT/s)	8.0
L1 cache	8 x 32 KB (instruction), 8 x 32 KB (data)
L2 cache	8 x 256 KB
L3 cache (MB)	20 MB
Platform	
Vendor and model number	Dell PowerEdge R820
Motherboard model number	Dell UL94V-0
BIOS name and version	Dell 1.3.7
BIOS settings	Defaults
Memory module(s)	
Total RAM in system (GB)	256
Vendor and model number	Samsung® M393B2G70BH0-YH9
Type	PC3-10600
Speed (MHz)	1,333
Speed running in the system (MHz)	1,333
Timing/Latency (tCL-tRCD-tRP-tRASmin)	9-9-9-24
Size (GB)	16

System	Dell PowerEdge R820 solution
Number of RAM module(s)	16
Chip organization	Double-sided
Rank	Dual
Operating system	
Name	Red Hat Enterprise Virtualization Hypervisor 6.3
Build number	6.3 - 20121212.0.el6_3
File system	ext4
Kernel	2.6.32 - 279.19.1.el6.x86_64
Language	English
Graphics	
Vendor and model number	Matrox® G200eR
Graphics memory (MB)	16
RAID controller	
Vendor and model number	Dell PERC H710P Adapter
Firmware version	21.1.0-0007
Cache size	1 GB
Hard drive	
Vendor and model number	Seagate® ST9146852SS
Number of disks in system	4
Size (GB)	146
Buffer size (MB)	16
RPM	15,000
Type	SAS 6 Gb/s
Ethernet adapter 1	
Vendor and model number	Intel Gigabit 4P I350-t
Type	Integrated
Ethernet adapter 2	
Vendor and model number	Intel Ethernet Converged Network Adapter X520-SR2
Type	Discrete
USB ports	
Number	4 external, 1 internal
Type	2.0

Figure 7: Configuration information for the test server.

APPENDIX B – MIGRATION TESTING – INSTALLING THE UNDERLYING INFRASTRUCTURE

Installing Red Hat Enterprise Virtualization Hypervisor 3.1

We performed the following steps to install Red Hat Enterprise Virtualization Hypervisor on our migration target system. For further information, refer to the Red Hat Enterprise Virtualization installation documentation.

1. Boot to the Red Hat Enterprise Virtualization Hypervisor stand-alone installation disk.
2. Allow the installation to automatically boot to begin the installation process.
3. Once the installer loads, select Install Hypervisor.
4. Select the proper disk to install Red Hat Enterprise Virtualization on, and click Continue.
5. Select the disk again, and click Continue.
6. Enter and confirm a new password, and click Install.
7. Once the installation is completed, click Reboot.

Installing Red Hat Enterprise Virtualization Manager

We performed the following steps to install Red Hat Enterprise Virtualization Manager, starting from a basic Red Hat Enterprise Linux 6.3 installation. The following steps require an active Internet connection and a subscription to the Red Hat Network service.

1. Log into the Red Hat Enterprise Virtualization Manager as root.
2. Sync to local time source.
3. Set up /etc/hosts.
4. Type `rhn_register` to begin the registration on the RHN and accept defaults, making sure to select your appropriate region.
5. To add channels for Red Hat Enterprise Virtualization, enter the following commands:
```
rhn-channel --add --channel=rhel-x86_64-server-6-rhevm-3.1
rhn-channel --add --channel=jbappplatform-5-x86_64-server-6-rpm
rhn-channel --add --channel=rhel-x86_64-server-supplementary-6
```
6. To turn off Red Hat Enterprise Linux firewalls, enter the following commands:
```
service iptables save
service iptables stop
chkconfig iptables off
service ip6tables save
service ip6tables stop
chkconfig ip6tables off
```
7. To update Red Hat Enterprise Linux, type `yum upgrade`
8. To start the Red Hat Enterprise Virtualization Manager install process, type `rhevm-setup`.
9. Press Enter to accept the default HTTP port.
10. Press Enter to accept the default HTTPS port.
11. Enter a password for the admin@internal username (we used `Password1`).
12. Enter a password for the internal database (we used `Password1`).
13. Choose your default storage type (we chose iSCSI).

14. Allow the installer to configure an NFS share on the local server for use as an ISO Domain.
15. Choose a mount path for the ISO Domain share.
16. Choose a display name for the ISO Domain.
17. Allow the installer to configure the firewalls automatically.

Connecting to Red Hat Enterprise Virtualization Manager

1. To connect to Red Hat Enterprise Virtualization Manager, type the fully qualified domain name of the Red Hat Enterprise Virtualization Manager server in a Web browser.
2. At the login screen, type `admin` as the username and enter the password you created earlier in the password field.
3. Click the Hosts tab. Select New to add a new host.
4. Type the name, IP address, and root password into the appropriate fields, then click OK to add the host. If there is no configuration for Power Management, click OK to continue. This may take a few minutes to complete.
5. Wait for the status to show Pending Approval and click Approve.
6. Once the Edit and Approve Host screens appear, click OK. If there is no configuration for Power Management, click OK to continue.
7. Configure networking and storage domains according to Red Hat Enterprise Virtualization documentation.

APPENDIX C – MIGRATION TESTING – PERFORMING THE MIGRATION

Configuring V2V and performing the migration

1. Connect to the Red Hat Enterprise Virtualization Manager using ssh.
2. Run the following command to install virt-v2v:
   ```
   #yum -y install virt-v2v
   ```
3. If migrating from VMware vSphere, edit the `/root/.netrc` file and add the information below. If targeting Xen or KVM VMs, syntax may vary. Replace hostname with your host name, IP address with your ESX host IP address, and Password1 with your password.
   ```
   hostname 192.168.1.100 login root password Password1
   ```
4. Run the following command to copy the VM from your legacy host to your Red Hat Enterprise Virtualization Manager Export Storage Domain:
   ```
   #virt-v2v -ic esx://192.168.1.100/?no_verify=1 -o rhev -os \
    RHEV-M:/vol/rhev_export --network rhevm VM01
   ```
5. Import the VM inside Red Hat Enterprise Virtualization Manager.

APPENDIX D – PERFORMANCE TEST TOPOLOGY

We used 15 virtual client machines for our test, with each client targeting one VM on the server under test. Each client machine was a VM running Windows Server 2008 R2 SP1 Enterprise with the following:

- The latest ds2pgsqlserverdriver.exe, Mono.Security.dll, Npgsql.dll files, and the necessary run scripts.
- .NET framework 4, required for the DS2 PostgreSQL driver and Npqsql .Net Data Provider.

We ran one physical workstation-class client machine, bringing the total to 16 clients, to ensure there were no performance differences between running virtual and physical clients.

For the client virtual machines, we cabled all four of the client host's 1Gb network interface cards (NICs) to a Dell PowerConnect™ 5448 switch. We then created four virtual machine networks and dedicated each physical NIC to its own virtual machine network. We distributed the virtual NICs round-robin amongst the VMs.

We attached one NIC in the Dell PowerEdge R820 to the PowerConnect 5448 switch. We initiated and controlled all testing sessions from a separate controller machine. Storage traffic used a separate switch, which we describe in Appendix E.

APPENDIX E – PERFORMANCE TESTING – STORAGE CONFIGURATION

Setting up the storage

Dell EqualLogic PS6110X and PS6110XS storage array configuration overview

Our complete storage infrastructure consisted of four internal drives in the system under test, a Dell PERC H710P storage adapter, and one Intel Ethernet Converged Network Adapter X520-SR2 for iSCSI network traffic to the external storage.

We configured the four internal drives in one RAID 1, dedicated to the hypervisor and backup storage. For external storage, we used one Dell EqualLogic PS6110X tray and one Dell EqualLogic PS6110XS tray interconnected through a Dell Force10™ S4810P switch. For redundancy, we cabled both controllers on each tray to the Force10 switch. We cabled the server to the Force10 switch via both ports on the Intel Ethernet Converged Network Adapter X520-SR2.

Each storage array contained 24 disks. The Dell EqualLogic PS6110XS contained seven 400GB SAS SSDs and 17 600GB 10K SAS HDDs in a RAID 6 (accelerated). The Dell EqualLogic PS6110X contained 24 900GB 10K SAS HDDs in a RAID 10. We assigned each tray to a separate storage pool, using the Dell EqualLogic PS6110XS SSD/HDD hybrid tray for our virtual disks containing database data and the Dell EqualLogic PS6110X tray for our virtual disks containing operating system and PostgreSQL log data. We used virtual disks for the OS and created a 1TB volume on the PS6110X for a Red Hat Enterprise Virtualization iSCSI storage domain. For our data and log drives, we created individual volumes for each VM and directly attached them using iSCSI from within the guest operating system. For the data volumes, we created 16 25GB volumes and limited access to each VM's iSCSI initiator IP addresses. Similarly, for the log volumes, we created 16 30GB volumes and limited access to each VM's iSCSI initiator IP addresses.

Dell Force10 S4810P switch configuration

Before beginning, make sure both switches are running the latest version of firmware available. For our testing, we used version 8.3.10.3. With a serial cable and terminal utility, connect to the switch and complete the steps below:

Configuring the out-of-band management port

1. In the switch's terminal, type `enable` and press Enter.
2. Type `config` and press Enter.
3. Type `interface ManagementEthernet 0/0` and press Enter.
4. Type `no shut` and press Enter.
5. Type `ip address 192.168.1.1 255.255.255.0` and press Enter.
6. Type `no shut` and press Enter.
7. Type `exit` and press Enter.

Configuring login credentials

1. Type `username admin privilege 15 password 0` *`password`* and press Enter, where *`password`* is the password you wish to set.
2. Type `enable password level 15 0` *`password`* and press Enter.

Configuring the remaining ports

1. Type `interface range tengigabitethernet 0/0 – 43` and press Enter.
2. Type `portmode hybrid` and press Enter.
3. Type `switchport` and press Enter.
4. Type `flowcontrol rx on tx on` and press Enter.
5. Type `no shut` and press Enter.
6. Type `exit` and press Enter.
7. Type `no dcb enable` and press Enter.
8. To save these settings, type `copy run start` and press Enter.
9. To apply these settings, type `reload` and press Enter. When asked if you want to proceed with the reload, type `yes` and press Enter. After the switches reboot, the settings will be in effect.

APPENDIX F – PERFORMANCE TESTING – SERVER AND VM CONFIGURATION

We performed the following steps to configure our management and host servers. We configured the host server with default BIOS settings and ensured that we used the most balanced system-power profile settings available, Performance per Watt (DAPC) .

Installing Red Hat Enterprise Virtualization Hypervisor on the Dell PowerEdge R820

1. Boot to the Red Hat Enterprise Virtualization Hypervisor stand-alone installation disk.
2. Allow the installation to automatically boot, and begin the installation process.
3. Once the installer loads, select Install Hypervisor.
4. Select the proper disk for Red Hat Enterprise Virtualization installation, and click Continue.
5. Select the disk again, and click Continue.
6. Enter and confirm a new password, and click Install.
7. Once the installation is completed, click Reboot.

Configuring Red Hat Enterprise Virtualization Hypervisor on the Dell PowerEdge R820

1. Run these commands to enable SR-IOV support:
```
mount -o remount,rw /dev/.initramfs/live
vi /boot/grub/grub.conf
```
2. Add the following to the default kernel boot line:
```
intel_iommu=on
```
3. Download the latest SR-IOV and NUMA hooks to /tmp (www.ovirt.org/VDSM-Hooks_Catalogue).
4. Navigate to the root folder and extract the RPM files:
```
cd /
rpm2cpio /tmp/vdsm-hook-sriov-4.10.0-0.442.git6822c4b.el6.noarch.rpm | cpio -idmv
rpm2cpio /tmp/vdsm-hook-numa-4.10.0-0.442.git6822c4b.el6.noarch.rpm | cpio –idmv
```
5. Make the files persistent using the persist command:
```
persist /etc/sudoers.d/50_vdsm_hook_sriov
persist /usr/libexec/vdsm/hooks/after_vm_destroy/50_sriov
persist /usr/libexec/vdsm/hooks/before_vm_migrate_source/50_sriov
persist /usr/libexec/vdsm/hooks/before_vm_start/50_sriov
persist /usr/libexec/vdsm/hooks/before_vm_start/50_numa
```
6. Modify /etc/rc.local to include the lines below.
```
#configure regular->huge page conversion rate
echo 100 >
/sys/kernel/mm/redhat_transparent_hugepage/khugepaged/scan_sleep_millisecs

#configure SR-IOV
#remove driver
modprobe -r igb
modprobe -r ixgbe
#readd driver with new configuration
modprobe igb max_vfs=4
modprobe ixgbe max_vfs=16
```

```
sleep 5
#configure VM mac addresses for auto configuration
#1Gb NICs
ip link set em1 vf 0 mac 00:50:56:b4:2b:eb
ip link set em1 vf 1 mac 00:50:56:b4:0a:28
ip link set em1 vf 2 mac 00:50:56:b4:02:0a
ip link set em1 vf 3 mac 00:50:56:b4:21:36
ip link set em2 vf 0 mac 00:50:56:b4:42:bf
ip link set em2 vf 1 mac 00:50:56:b4:61:a1
ip link set em2 vf 2 mac 00:50:56:b4:5c:a0
ip link set em2 vf 3 mac 00:50:56:b4:15:21
ip link set em3 vf 0 mac 00:50:56:b4:40:01
ip link set em3 vf 1 mac 00:50:56:b4:38:c7
ip link set em3 vf 2 mac 00:50:56:b4:0a:ee
ip link set em3 vf 3 mac 00:50:56:b4:02:81
ip link set em4 vf 0 mac 00:50:56:b4:54:14
ip link set em4 vf 1 mac 00:50:56:b4:31:3a
ip link set em4 vf 2 mac 00:50:56:b4:73:3d
ip link set em4 vf 3 mac 00:50:56:b4:75:33
#10Gb NICs
ip link set p2p1 vf 0 mac 00:50:56:b4:0b:54
ip link set p2p2 vf 0 mac 00:50:56:b4:00:5e
ip link set p2p1 vf 1 mac 00:50:56:b4:33:b0
ip link set p2p2 vf 1 mac 00:50:56:b4:2c:ca
ip link set p2p1 vf 2 mac 00:50:56:b4:57:82
ip link set p2p2 vf 2 mac 00:50:56:b4:7e:2c
ip link set p2p1 vf 3 mac 00:50:56:b4:54:32
ip link set p2p2 vf 3 mac 00:50:56:b4:58:54
ip link set p2p1 vf 4 mac 00:50:56:b4:43:4c
ip link set p2p2 vf 4 mac 00:50:56:b4:24:d4
ip link set p2p1 vf 5 mac 00:50:56:b4:51:c7
ip link set p2p2 vf 5 mac 00:50:56:b4:7c:e3
ip link set p2p1 vf 6 mac 00:50:56:b4:75:39
ip link set p2p2 vf 6 mac 00:50:56:b4:59:80
ip link set p2p1 vf 7 mac 00:50:56:b4:1f:0a
ip link set p2p2 vf 7 mac 00:50:56:b4:3c:cb
ip link set p2p1 vf 8 mac 00:50:56:b4:2f:70
ip link set p2p2 vf 8 mac 00:50:56:b4:22:bc
ip link set p2p1 vf 9 mac 00:50:56:b4:0d:de
ip link set p2p2 vf 9 mac 00:50:56:b4:0f:a6
```

```
ip link set p2p1 vf 10 mac 00:50:56:b4:7a:6c
ip link set p2p2 vf 10 mac 00:50:56:b4:50:b8
ip link set p2p1 vf 11 mac 00:50:56:b4:66:98
ip link set p2p2 vf 11 mac 00:50:56:b4:58:d1
ip link set p2p1 vf 12 mac 00:50:56:b4:76:57
ip link set p2p2 vf 12 mac 00:50:56:b4:0d:d4
ip link set p2p1 vf 13 mac 00:50:56:b4:45:d7
ip link set p2p2 vf 13 mac 00:50:56:b4:04:d7
ip link set p2p1 vf 14 mac 00:50:56:b4:70:19
ip link set p2p2 vf 14 mac 00:50:56:b4:11:38
ip link set p2p1 vf 15 mac 00:50:56:b4:54:2f
ip link set p2p2 vf 15 mac 00:50:56:b4:30:eb
```

Installing Red Hat Enterprise Linux and Red Hat Enterprise Virtualization Manager on the management server

1. Mount the Red Hat Enterprise Linux 6.3 DVD.
2. Boot the machine to the Red Hat Enterprise Linux 6.3 installation DVD.
3. Press Enter to install using graphical mode.
4. At the Media test screen, select Skip and click Next.
5. At the Red Hat Enterprise Linux 6.3 title screen, click Next.
6. At the Choose a Language screen, select English and click Next.
7. At the Keyboard Type screen, select U.S. English and click Next.
8. At the Installation Devices screen, select Basic Storage Devices and click Next.
9. Enter the hostname and click Next.
10. At the Time zone selection screen, select the appropriate time zone and click Next.
11. Enter the root password in the Root Password and Confirm fields and click Next.
12. At the type of installation screen, select Use All Space and click Next.
13. Click Next. If a warning appears, click Write changes to disk.
14. At the Boot loader screen, click Next.
15. At the Default installation screen, select Basic Server and click Next. Installation will begin automatically.
16. At the Congratulations screen, click Reboot.
17. Log into the server as root.
18. Disable SELinux by modifying the `/etc/selinux/config` file and changing the SELINUX line to `SELINUX=disabled`.
19. Synchronize to local time source.
20. Set up /etc/hosts
21. Type `rhn_register` to begin the registration on the RHN and accept defaults, making sure to select your appropriate region.
22. To add channels for Red Hat Enterprise Virtualization, type:
```
rhn-channel --add --channel=rhel-x86_64-server-6-rhevm-3.1
rhn-channel --add --channel=jbappplatform-5-x86_64-server-6-rpm
rhn-channel --add --channel=rhel-x86_64-server-supplementary-6
```
23. To turn off Red Hat Enterprise Linux firewalls, type:
```
service iptables save
service iptables stop
chkconfig iptables off
```

```
service ip6tables save
service ip6tables stop
chkconfig ip6tables off
```

24. To update Red Hat Enterprise Linux, type `yum upgrade`
25. Type `rhevm-setup` to start the Red Hat Enterprise Virtualization Manager installation process.
26. Press Enter to accept the default HTTP port.
27. Press Enter to accept the default HTTPS port.
28. Enter a password for the admin@internal username (we used Password1).
29. Enter a password for the internal database (we used Password1).
30. Choose your default storage type (we chose iSCSI).
31. Allow the installer to configure an NFS share on the local server to use as an ISO Domain.
32. Choose a mount path for the ISO Domain share.
33. Choose a display name for the ISO Domain.
34. Allow the installer to configure the firewalls automatically.
35. Run the following command to enable SR-IOV on Red Hat Enterprise Virtualization Manager:
    ```
    rhevm-config -s UserDefinedVMProperties='sriov=.*;numa=.*' --cver=3.1
    ```
36. Type `service ovirt-engine restart` to restart Red Hat Enterprise Virtualization Manager.

Configuring iSCSI on the Red Hat Enterprise Linux VM

These steps can be performed prior to or after the V2V migration.

1. Log onto the VM as root.
2. Install the Dell EqualLogic Linux HIT kit. See support.equallogic.com for details.
3. Use the provided Dell EqualLogic commands to discover the iSCSI portal and login to the target:
   ```
   # rswcli --add-group groupname
   # ehcmcli login --target <data targetname>
   # ehcmcli login --target <log targetname>
   ```
4. Format the disks:
   ```
   # mkfs.ext4 -L data /dev/eql/data
   # mkfs.ext4 -L log /dev/eql/log
   ```
5. Edit the /etc/fstab file to include
   ```
   /dev/eql/data /var/lib/pgsql/9.2/data ext4 nobarrier,noatime,_netdev 0 0
   /dev/eql/log /var/lib/pgsql/9.2/data/pg_xlog ext4 nobarrier,noatime,_netdev 0 0
   ```
6. Mount these file systems, and reset the tuned profile:
   ```
   # mount -a
   # tuned-adm profile enterprise-storage
   ```

Installing PostgreSQL on the Red Hat Enterprise Linux VM

We used the PostgreSQL database server, version 9.2, as the database software on the VMs. These steps can be performed prior to or after the V2V migration.

1. Download the rpm for PostgreSQL 9.2 from http://yum.postgresql.org/repopackages.php and upload the rpm file to the VM.
2. Log onto the VM as root.
3. Install the PostgreSQL database server and client:
   ```
   # yum install postgresql92-server postgresql92-contrib
   ```
4. Initialize the PostgreSQL database cluster for the first time:
   ```
   service postgresql-9.2 initdb
   ```

5. Tune the database engine. Modify the following lines of the `/var/lib/pgsql/9.2/data/postgresql.conf` file:
```
checkpoint_segments = 1024
checkpoint_timeout = 1h
checkpoint_completion_target = 0.9
shared_buffers = 4GB
effective_cache_size = 10GB
```
6. Add one line to the PostgreSQL configuration file `pg_hba.conf` to permit SQL queries from the client network:
```
host all all <test bed IP subnet>.0/24trust
```
7. Modify the listen_addresses line in the PostgreSQL configuration file `postgresql.conf`:
```
listen_addresses = '*'
```
8. Run `service postgresql-9.2 restart` to restart PostgreSQL with the new settings.
9. Create the DS2 user for the OS on the Red Hat system:
```
# useradd ds2
```
10. Set user password to ds2:
```
# passwd ds2
```
11. Download the latest DVD Store distribution for PostgreSQL from http://linux.dell.com/dvdstore/ and upload the file to the VM.
12. Unpack the DS2 distribution into `/ds2`.
13. Change directory to `/ds2/pgsqlds2`.
14. Follow the DVD Store instructions for generating the text file data.
15. Log on as postgres user:
```
# su postgres
```
16. Run the shell script `pgsqlds2_create_all.sh`

Converting and importing the gold VM to Red Hat Enterprise Virtualization environment

Our source machine was running VMware vSphere – For Xen and KVM source hosts see Red Hat documentation for syntax.

1. Connect to the Red Hat Enterprise Virtualization Manager using ssh.
2. Running the following command to install virt-v2v:
```
yum -y install virt-v2v
```
3. Edit the /root/.netrc file and add the following (replacing IP address with your source host IP address):
```
machine 192.168.1.100 login root password Password1
```
4. Run the following command to copy the VM from your source host to your Red Hat Enterprise Virtualization Manager export storage domain:
```
virt-v2v -ic esx://<vsphere-server-IP>/?no_verify=1 -o rhev -os RHEV-M:/vol/rhev_export --network rhevm VM01
```
5. Log into the Red Hat Enterprise Virtualization Manager Web console, and navigate to the Storage tab.
6. The newly migrated VM will now be located under the VM Import tab in the Export domain used in the virt-v2v command.
7. To import the VM into your Red Hat Enterprise Virtualization environment, select the VM you wish to import and click Import.
8. Select the desired default storage domain, cluster, and disk location, and click OK to begin the import.

Cloning the VM on Red Hat Enterprise Virtualization

1. Once the VM has been imported, right-click the VM, and select Make Template.

2. Enter a name, description, select the Default cluster, and select the proper storage domain.
3. Uncheck Make Private, and click OK.
4. Select VMs, and click New Server.
5. Select the new template for the VM to be based on.
6. Type a name for the new VM.
7. Select the proper storage domain.
8. Click OK.
9. Continue deploying each VM until you have deployed all 16.
10. Assign the proper IP addresses and hostnames to each VM.
11. Remount the proper iSCSI LUNs in each VM.

Configuring the Red Hat Enterprise Virtualization environment and VM settings

We performed the following steps on the Red Hat Enterprise Virtualization environment.

1. In Red Hat Enterprise Virtualization Manager, right-click the first VM and click Edit.
2. Select Host.
3. Assign the VM to specifically use the only Red Hat Enterprise Virtualization Hypervisor host, and check the box Run VM on the selected host (no migration allowed).
4. For CPU Pinning topology, use the information in Figure 8 to pin each VM four vCPUs to the appropriate physical resources. Repeat for each VM.

VM1	0#0_1#32_2#4_3#36	VM9	0#2_1#34_2#6_3#38
VM2	0#8_1#40_2#12_3#44	VM10	0#10_1#42_2#14_3#46
VM3	0#16_1#48_2#20_3#52	VM11	0#18_1#50_2#22_3#54
VM4	0#24_1#56_2#28_3#60	VM12	0#26_1#58_2#30_3#62
VM5	0#1_1#33_2#5_3#37	VM13	0#3_1#35_2#7_3#39
VM6	0#9_1#41_2#13_3#45	VM14	0#11_1#43_2#15_3#47
VM7	0#17_1#49_2#21_3#53	VM15	0#19_1#51_2#23_3#55
VM8	0#25_1#57_2#29_3#61	VM16	0#27_1#59_2#31_3#63

Figure 8: CPU pinning topology information.

5. Click Custom Properties.
6. Click the + sign, and use the drop-down menu to select numa.
7. To assign the VM to the one NUMA node with assigned CPUs enter “strict:0” in the field to the right.
   VMs 1 through 4 --- > strict:0
   VMs 5 through 8 --- > strict:1
   VMs 9 through 12 --- > strict:2
   VMs 13 through 16 --- > strict:3
8. Click the + sign and use the drop-down menu to select SR-IOV.
9. For SR-IOV NIC configuration, use the information in Figure 9 to pin each NIC in each VM to the appropriate physical resources.

VM1	em1_0,p2p1_0,p2p2_0	VM9	em3_0,p2p1_8,p2p2_8
VM2	em1_1,p2p1_1,p2p2_1	VM10	em3_1,p2p1_9,p2p2_9
VM3	em1_2,p2p1_2,p2p2_2	VM11	em3_2,p2p1_10,p2p2_10
VM4	em1_3,p2p1_3,p2p2_3	VM12	em3_3,p2p1_11,p2p2_11
VM5	em2_0,p2p1_4,p2p2_4	VM13	em4_0,p2p1_12,p2p2_12
VM6	em2_1,p2p1_5,p2p2_5	VM14	em4_1,p2p1_13,p2p2_13
VM7	em2_2,p2p1_6,p2p2_6	VM15	em4_2,p2p1_14,p2p2_14
VM8	em2_3,p2p1_7,p2p2_7	VM16	em4_3,p2p1_15,p2p2_15

Figure 9: NIC pinning topology information.

APPENDIX G – PERFORMANCE TESTING – RUNNING THE TEST

About our DVD Store 2.1 workload

To create our real-world e-commerce workload, we used the DVD Store Version 2.1 benchmarking tool. DS2 models an online DVD store, where customers log in, search for movies, and make purchases. DS2 reports these actions in orders per minute that the system could handle, to show what kind of performance you could expect for your customers. The DS2 workload also performs other actions, such as adding new customers, to exercise the wide range of database functions you would need to run your e-commerce environment.

We ran DS2 with 16 threads, and with a 0.01-second think time. For the specific flags we used, see below. For more details about the DS2 tool, see http://www.delltechcenter.com/page/DVD+Store.

We created a series of batch files and shell scripts to automate the complete test cycle. DVD Store outputs an orders-per-minute metric, which is a running average calculated through the test. In this report, we report the last OPM reported by each client/target pair.

Each complete test cycle consisted of the general steps listed below. We ran three test cycles and chose the median outcome.

1. Clean up prior outputs from the host system and all client driver systems.
2. Drop all databases from all target VMs.
3. Restore all databases on all target VMs.
4. Shut down all VMs.
5. Reboot the host system and all client systems.
6. Wait for a ping response from the server under test (the hypervisor system), all client systems, and all VMs.
7. Let the test server idle for 10 minutes.
8. Begin the test counters and power meters.
9. Start the DVD Store driver on all respective clients.

We used the following DVD Store parameters for testing the virtual machines in this study:

```
ds2pgsqldriver.exe --target=<target_IP> --ramp_rate=10 --run_time=10 --
n_threads=16 --db_size=10GB --think_time=0.01 --detailed_view=Y -warmup_time=20
```

Measuring power

To record power consumption during the test, we used an Extech Instruments (www.extech.com) 380803 Power Analyzer/Datalogger. We connected the power cords from the servers under test to the Power Analyzer's output-load power outlet. We then plugged the power cord from the Power Analyzer's input voltage connection into a power outlet.

We determined wattage measurements with the Power Analyzer's Data Acquisition Software. On a separate machine connected to the Power Analyzer via an RS-232 cable, we installed the software. We captured power consumption at one-second intervals. Once rebooted, each server idled for 10 minutes prior to recording power measurements. For average power usage, we measured use during the 10-minute test run_time window.

ABOUT PRINCIPLED TECHNOLOGIES

Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.