



Enterprise Virtualization: Comparing Red Hat and Oracle Solutions

January 2013

**About
Sintre
Technologies**

Sintre Technologies is an authorized Red Hat product reseller that offers turnkey infrastructure products, expert Red Hat and Linux consulting, and custom training. Using creative product and service delivery strategies such as remote consulting and preconfigured services, Sintre is able to work within customer's physical and budgetary constraints. Red Hat products, custom Linux environments, Oracle VM deployments, and full environment designs and infrastructure deployments are all services in which Sintre specializes.

For additional information on Sintre services and consulting, please visit us at <http://www.sintre.com/>.

**About
the
Author**

As a Senior Linux Architect, Ryan Murray has helped many companies develop elegant, cost-effective Linux-based network environments. With production experience using Solaris Containers, multiple versions of VMware, oVirt, KVM over Red Hat Cluster Suite, Red Hat Enterprise Virtualization (RHEV) versions 2 through 3, and Oracle VM (OVM) versions 2 through 3, Ryan possesses a solid grasp of enterprise virtualization design and implementation.

Recently, the opportunity arose to simultaneously deploy RHEV 3.1 and OVM 3.1.1 for two different clients at each of their primary production data centers. That experience led him to create this whitepaper, in which he is more than excited to explore the technical difference between the two products.

Ryan holds the following certifications:

RHCA: Red Hat Certified Architect

RHCI: Red Hat Certified Instructor

RHCVA: Red Hat Certified Virtualization Administrator

RHCDS: Red Hat Certified Data Center Specialist

RHCE: Red Hat Certified Engineer

SCSA: Sun Certified System Administrator Solaris 9

SCNA: Sun Certified Network Administrator Solaris 9

SCSECA: Sun Certified Security Administrator Solaris 9

OPN (Oracle Partner Network) Certified Specialist

Oracle VM 3.0 for x86 Certified Implementation Appliance Specialist

Definitions and Abbreviations

Global Terms

| | |
|--------------------|---|
| VM: | A virtual machine that simulates the architecture and program execution of a physical computer. |
| Hypervisor: | The computer hardware that enables virtual machines as well as the abstract software or firmware layer of virtualization. Referenced in this paper generically as a system or process that enables virtualization. |
| Cluster: | A set of networked computers that work together such that they are viewed as a single system. |
| Socket: | Also known as a CPU slot, it is the mechanical component that provides connections between the microprocessor and the circuit board in a computer. Virtualization products are typically licensed by the number of CPU sockets available in the server. |

Red Hat Terms

| | |
|---------------------------|---|
| RHEL: | Red Hat Enterprise Linux. |
| RHN: | Red Hat Network, a hosted repository of current software packages and security updates. |
| RHEV: | Red Hat’s virtualization product, referenced here as the complete product inclusive of multiple components. |
| RHEV-H: | The physical servers or hypervisor responsible for hosting virtual machines in an RHEV environment. |
| RHEV-M: | The management server in an RHEV environment that hosts the Jboss application. It communicates with RHEV-H servers. |
| RHEV Admin Portal: | The web console interface provided by the RHEV-M server and used for administering the RHEV environment. |
| RHEV User Portal: | A web interface provided by the RHEV-M server. It is used for delegating subsets of administration tasks to specific users. |

Oracle Terms

| | |
|-----------------------|---|
| OEL: | Oracle Enterprise Linux, a ported version of RHEL. |
| OVM: | Oracle's virtualization product, referenced here as the complete product inclusive of multiple components. |
| OVM-S: | The physical servers responsible for hosting virtual machines in an OVM environment. |
| OVM-M: | The management server in an OVM environment. It hosts the Weblogic application and communicates with the OVM-S servers. |
| OVM-M console: | The web interface provided by the OVM-M server and used for administering the OVM environment. |

Table of Contents

| | |
|---|----|
| About Sintre Technologies | 2 |
| About the Author | 2 |
| Definitions and Abbreviations..... | 3 |
| Global Terms | 3 |
| Red Hat Terms | 3 |
| Oracle Terms | 3 |
| Executive Summary | 5 |
| Vendor Background | 5 |
| Building the Clusters | 6 |
| 1.1 Build Process Flow..... | 6 |
| 1.2 Installing the Management system..... | 7 |
| 1.3 Installing the Hypervisor System | 8 |
| 1.4 Configuring the Management Network..... | 8 |
| 1.5 Managing the Servers..... | 8 |
| Management Console Usability..... | 9 |
| 2.1 Administrative Interface..... | 9 |
| Core Technology | 10 |
| 3.1 Storage Support..... | 10 |
| 3.2 Hypervisor Implementation | 11 |
| 3.3 Cluster Fencing and Locking..... | 12 |
| 3.4 Role-Based Access Control..... | 12 |
| Additional Features | 13 |
| 4.1 Virtual Desktop Infrastructure | 13 |
| 4.2 Reporting | 13 |
| Product Support..... | 14 |
| 5.1 Customer Support..... | 14 |
| 5.2 Community Support..... | 14 |
| Product Migration | 15 |
| 6.1 Migrating VMs from Previous Versions..... | 16 |
| Product Pricing | 15 |
| 7.1 Pricing Data | 16 |
| 7.2 Sample Configuration Pricing | 16 |
| Conclusion | 16 |
| Architecture Diagrams | 17 |
| Disclaimer..... | 19 |

Executive Summary

Commonly considered one of the newer technology trends, virtualization was actually born back in the early sixties on mainframes by companies like General Electric, Bell Labs, and IBM. Support for the x86 platform occurred much later with the release of products like Virtual PC by Apple in 1997 and VMware Workstation in 1999. But it wasn't until 2001 that interest in the enterprise arena was sparked by the launch of the VMware ESX and GSX Server products.

In the decade since that launch, a number of additional vendors have entered the enterprise virtualization market with a diverse set of competitive offerings. Two of the leading Linux technologies, available from Red Hat and Oracle, deliver a mostly robust set of features at compelling price points. Following a recent and concurrent deployment of those solutions at two independent Sintre client sites, we created this technical comparison of each vendor's implementation in the hopes of providing valuable insights to those companies contemplating a new or enhanced virtualization deployment strategy.

The topics covered here should not be considered an extensive feature comparison. Instead, this paper is more a highlight of the key features and a discussion of some of the issues we discovered while building multiple production environments. If a section doesn't go into detail for one of the solutions, this only means that the required action was performed without incident or with no caveats worthy of mention.

Vendor Background

Red Hat

Red Hat has long been recognized as an industry leader and a champion of the open source community. Founded in 1993 with corporate headquarters in Raleigh, North Carolina, the company entered the world of virtualization in 2008 when it purchased Qumranet.

Qumranet was the creator of desktops hosted on Kernel-based Virtual Machines (KVM) on servers linked with their SPICE protocol (Simple Protocol for Independent Computing Environments). Released as open source in December 2009, SPICE has evolved to deliver a high-quality user experience that provides video streaming and 3D graphics while maintaining low CPU consumption.

Red Hat Enterprise Virtualization (RHEV) 3.0 was released in January 2012, and the latest 3.1 version was released in December 2012.

Oracle

Established in 1977, Oracle's main office is located in Redwood Shores, California. Best known for its enterprise database products, Oracle went on a \$20 billion buying spree of companies in 2006. It now owns Sun Solaris, as well as multiple virtualization solutions (including Solaris Zones, LDoms, and xVM) and the SPARC server hardware platform.

Oracle VM (OVM) is Oracle's implementation of the Xen hypervisor that also leverages intellectual property that Oracle acquired from Sun Microsystems and the 2009 purchase of Virtual Iron. Those technologies were integrated with the Oracle VM 3.1 release in May 2012. Oracle continues to converge on Enterprise Manager 12c to manage its virtualization and infrastructure portfolio.

Building the Clusters

1.1 Build Process Flow

A high level overview of the steps required to deploy a virtual environment. See a general diagram of the architecture components for each solution on page 17.

- RHEV**
1. Build the RHEV-M node by installing RHEL 6.2. Then install the RHEV-M software and configure with the *rhev-setup* command.
 2. Build at least one RHEV-H node using the provided ISO image. Once a rhevm network is configured, a request to add the system to the RHEV cluster is made through the console screen.*
 3. Through the Admin Portal, the pending request can now be accepted.
 4. Any logical networks and vlans can be created.
 5. Once the networks are attached to the Data Center, they can be presented to the RHEV-H node, including any bonding configurations. The host is now ready to activate.
 6. Add a primary DATA storage domain, and attach it to the Data Center.
 7. Add an ISO storage domain, and attach it to the Data Center.
 8. If needed, upload ISO images from terminal using *rhev-iso-uploader --iso-domain=ISO_DOMAIN upload /path/to/iso*.
 9. A VM can now be created and deployed.

* The request step from item 2 and the accept step from item 3 can be consolidated by adding the RHEV-H directly from the console.

- OVM**
1. Build the OVM-M node by installing OEL 5.5. Then install OVM-M from the ISO image provided.
 2. Build at least one OVM-S node using the provided ISO image. During the install, the ovmm network is configured.
 3. From the OVM-M console, a discovery request is used to add the new OVM-S system.
 4. Edit the OVM-S node with network bonding and configuration information. The network names and functions can be modified after the fact on the Network tab.
 5. Discover SAN or NAS storage devices for the central storage repository that will host the VM image files.
 6. Discover SAN or NAS storage device for the cluster configs and voting files.
 7. Create a server pool to which the OVM-S nodes will be assigned. This requires a virtual IP address and storage for the clustering voting disks.
 8. A new repository can now be created by selecting an already discovered SAN or NAS storage container.
 9. If needed, upload ISO images either through HTTP/FTP/NFS shares, or transfer the iso file using SSH onto a OVM-S node under */OVS/Repositories/<UUID>/ISOs/*. After the transfer, the storage repository must be refreshed to display the new ISO image.
 10. A VM can now be created and deployed.

1.2 Installing the Management system

RHEV The required software is provided through a RHN software channel allowing the install to take place with the *yum* command. RHEV uses a PostgreSQL databases and it supports this configuration in a production environment, although an external database can be used if desired.

OVM The required software is provided as an ISO download from Oracle's site, and the installer script must be used. An embedded Oracle database is not supported, and it's listed only as a "demo" option in the install selection. If an external database is used, be aware that the installer is not RAC (Real Application Clusters) aware. This is surprising, since RAC is arguably Oracle's best product.

To install using a RAC cluster, a direct connector SID must be used—not the RAC SID. After the install, stop the ovmm service and edit the file `/u01/app/oracle/ovm-manager-3/machine1/base_adf_domain/config/jdbc/OVMDS-6373-jdbc.xml` as noted below. Change any references of the direct database SID to the RAC-aware SID. Also, make sure to use the exact formatting for the jdbc thin client, or the database connection will not be made.

After install and using the direct database SID:

```
<url>jdbc:oracle:thin:@HOSTNAME:PORT:LOCALSID</url>
```

```
and <value>LOCALSID</value>
```

After editing file for RAC aware SID:

```
<url>jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST(ADDRESS=(PROT  
OCOL=TCP)(HOST=DR-SCAN-HOSTNAME)(PORT=PORT)))  
(CONNECT_DATA=(SERVICE_NAME=SID)))</url>
```

```
and <value>SID</value>
```

Install Design Notes

Red Hat's approach follows a software appliance model with the assumption that internal components, such as the database, don't need manual administration. This removes the potential need to use another department's computing or personnel resources for limited administrative tasks. Basic stop or start actions and viewing log files of the small PostgreSQL instance are well within most Linux administrators' skill set. In fact, simply running *rhevm-log-collector* will capture the database logs. If any additional database administration tasks need to be performed, Red Hat provides the necessary commands.

While Oracle's approach to separate job responsibilities by offloading OVM to an external Oracle database mirrors best practices, it creates a possible limitation for many clients. Customers that don't have an approved version of Oracle (10.2.0.4 or greater, 11.1.0.7 or greater, and 11.2.0.1 or greater), or don't want application data intermixed with production or client databases, might be out of options. Installing the embedded database with OVM will result in an unsupported environment.

1.3 Installing the Hypervisor System

RHEV If you install over an older major version of RHEV-H, the installer assumes an upgrade is desired, and the install fails. To proceed, you must first delete or uninstall the old version of RHEV-H on the system before re-starting the install process. When asked, Red Hat states this limitation is the expected behavior. However, this is not the default behavior of the Linux Anaconda installer. If you install RHEV 3.1. over an older version, such as RHEV 2.2, this will add a couple of steps to the install process.

OVM The Oracle OVM installer asks if an install or upgrade is wanted. This is the expected install default behavior on Linux systems, and was verified with OVM by installing a fresh version of OVM-S 3.1.1 over an older OVM-S 2.1 without incident.

1.4 Configuring the Management Network

RHEV During the server installation, network configuration can be skipped unless you choose to configure it. Networking is not required during the install, because RHEV-H supports a Server to Manager initiation process. The request is initiated from RHEV-H using the text-based GUI screen in the console. This screen also has a tab for configuring the network, so the settings are easily modified later. The join request will show up in the RHEV-M console, though it requires an administrator to approve the request.

Red Hat also supports a Manager to Server initiation process. However, if this method is chosen, the network will not be immediately available after completing the RHEV-H install, unless the management network is configured during the install process.

OVM During the installation, network configuration is required. Networking is required because of the manner in which the initial join request is created. The OVM-M node is responsible for reaching out and discovering the OVM-S node. If network settings were not required during the install, an administrator would have to manually configure the network just as they would for any Linux system. However, this option is not advised by Oracle as inconsistent changes could cause functionality issues later.

1.5 Managing the Servers

RHEV After the install, logging in as the admin user presents a text-based GUI screen. All the settings that are normally necessary to fully configure and manage a RHEV-H node are accessible from this point. Although it is not a command-line interface as many Linux engineers might expect, it is very fast and easy to setup networking (just the rhevm network, as the rest can be done from the RHEV-M console), NTP, SSH access, logging, storage, RHN, and RHEV-M connection details. Full command-line access is available, but it is not recommended.

OVM After the install, a text status screen is displayed. Switching to another virtual console allows direct shell access. Users are faced with a warning stating that manual changes could cause the system to stop functioning. Normal Linux administrator procedures are followed to edit any needed changes. Caution should be used as any manual changes or deviations from OVM configuration standards could cause the system to behave abnormally.

Install Summary Comparison

Both RHEV-H and OVM-S are simple to install and can be configured to allow the administrator to close the server console window as soon as the “Reboot Now” option is selected from the install. Theoretically, the server console is never needed again, and all configurations are managed from the central administrative console or portal.

Management Console Usability

2.1 Administrative Interface

RHEV The visually appealing interface and intuitive process flow of RHEV’s management console are key strengths of the solution. The color-coded status of cluster resources makes identifying on and offline resources immediately obvious. Multiple VMs can be selected at the same time, and stop or start actions are available.

The search bar on the top of the screen enables quick access to specific resources and VMs. RHEV also allows VMs to be assigned multiple tags. Tags are created to better organize virtual servers and enable faster searching. For instance, using the search bar query “Vms:tag=QA” would list only systems associated with the QA tag.

Any search query can be saved as a Favorite and accessed directly from the Bookmarks side menu. Once VMs are tagged and bookmarks are created, an administrator can rapidly find any system in the console. This not only helps speed administration tasks, it also helps ensure critical actions are not executed on the wrong systems. For example, if every DEV system needs to be shut down, the DEV bookmark can be selected and actions taken on just this subset. Without grouping, an administrator scrolling through multiple VMs could accidentally select the wrong system.

OVM The management console in OVM 3 is a dramatic improvement over OVM 2. Every action is available now as a button, or by right-clicking the resource. One of the remaining negatives is the fact that status is still indicated by only a text string. In an environment with more than twenty VMs, attempting to scan the list for the current state of your resources is somewhat visually tiring.

Oracle also does not support grouping or tagging of virtual systems. In large enterprises, this makes the management of virtual servers more time consuming and could lead to accidental actions taking place on the wrong system. Only one VM may be selected at a time. Starting or stopping multiple systems takes three mouse clicks per VM—an untenable interaction if a hundred or more VMs are deployed.

Core Technology

3.1 Storage Support

RHEV Both file-based images on NFS and block level storage on iSCSI/FCP are supported in RHEV. If block level storage is selected, RHEV-H nodes use LVM (Logical Volume Manager) to allocate smaller physical devices to present directly to a VM. An SPM (Storage Pool Manager) is used in RHEV. Since the Linux kernel LVM is not clustered itself, only one RHEV-H node can alter LVM allocations, such as when a VM is provisioned or a snapshot is created. An internal RHEV process, the VDSM (Virtual Desktop and Server Management) daemon, replicates these LVM layout updates on each system in the cluster.

The SPM responsibility can migrate to other RHEV-H nodes if there are errors, or if the server is put into maintenance mode. Internal locking processes are used to ensure only one system can write metadata as the SPM.

When logical hard drives are created with a new VM, two formats are available: RAW and QCOW2. RAW, as the name implies, is a direct mapping to the underlying storage. Using QCOW2 as the storage format allows snapshots to be created. This is accomplished by having an additional layer between the logical drive and the raw storage. Although this layer does have some processing overhead, it allows a new QCOW2 image to be created in which mappings are directed to the original raw storage. This appears to duplicate the original data since the mappings are transparent to the VM. Any files that are changed get written to the new image while the original data is left untouched. New with RHEV 3.1 is the ability to perform a live snapshot without first shutting down the VM.

When attaching an iSCSI based storage domain, multiple initiator IPs can be used. The RHEV-H nodes will use multipath to failover if one of the links to storage is down. Bonding can be used instead of multipathing, but this is not ideal as there will be delays in the failover, and it could result in corrupt or slow data.

A data center may have only one storage type, meaning you can not mix iSCSI and NFS storage domains. This will be a limitation until RHEV 4.0 is released. ISO images are handled by the ISO storage domain. This domain is available for the entire data center, and it can service VMs from different storage domains. RHEV handles CD mounts from the ISO domain very well. New VMs can use the *run once* option to select the desired CD or PXE boot image. CDs can also be easily changed by selecting the *change cd* option.

OVM Oracle utilizes file-based images regardless if NFS or iSCSI/FCP is selected as the cluster storage type. iSCSI/FCP is used to create an OCFS2 volume to which every cluster node has access. A file image is then created on the shared volume and presented to the VM. Snapshots are available as a feature of OCFS2. In theory, they work similarly to QCOW2 in that when a snapshot is created, a new file or image is created that links to the original data. New data is then written to the new file or image.

Since OCFS2 is a file system, not block storage, all snapshots are created at a file level. Even though OCFS2 supports file level snapshots, OVM does not support snapshots of VMs. It is interesting that Oracle left this feature out, considering

that deploying a new VM from a template does take advantage of OCFS2 snapshots. Currently, the only way to backup a VM is to shut it down, and then create a template from the VM. If the backup is needed, a new VM must be created from this template.

Only one initiator IP can be used with iSCSI domains through the management console. Bonding must be used on the OVM-S nodes to provide redundancy for network issues, or multipathing can be configured on each OVM-S server manually. Although OVM's auto-discover of multiple paths does not work with iSCSI, it works well when using direct attach storage, such as FCP.

A data center may have any storage type available, and as long as the OVM-S nodes has access to both storage repositories, image files can migrate between them.

OVM doesn't have a global ISO domain concept. Instead, each storage repository has its own path to store ISOs. This creates a little extra work and a small amount of duplicated data as administrators have to upload the same ISO to every repository they plan to install from. When attaching a CD to a VM, the VM must be edited and added under the Disks tab. When the CD is no longer needed, it must be detached in the same way. This creates an awkward process flow, especially when building new VMs from an ISO or PXE boot.

3.2 Hypervisor Implementation

RHEV Guests are virtualized using KVM and qemu with the virtualization extension features found on most modern processor chips. This level of virtualization is considered full virtualization in that the virtual system is not fully aware that its hardware requests are passed through a kernel module emulating hardware and not directly to the hardware. Although KVM does not support full paravirtualization of hardware system calls, Red Hat provides optional device drivers to improve performance, specifically virtio for network interfaces and hard drives. Using these drivers would then technically classify the virtual guest virtualization type as hardware virtualization with paravirtualized drivers.

OVM Guests are virtualized using XEN with three main configuration options: HVM, HVM with PV drivers, and PVM. The first mirrors KVM in the fact that hardware calls are handled by the virtualization extension features on most modern processor chips. Hardware virtualization with paravirtualized drivers, also called PVHVM or HVM with PV drivers, is essentially the same as KVM with Red Hat's virtio drivers installed. The virtual system is able to utilize the virtualization extensions on the processor and some paravirtualized hardware calls.

The last method is paravirtualization, or PVM. This type requires the guest system to be fully aware that it is not running on direct hardware. For this option to work, the guest operating system must be built with a custom kernel, as hardware system calls are modified to pass directly to the hypervisor. PVM can provide slight performance increases over HVM as system calls go directly to the hypervisor, and not to an abstract hardware layer that requires further processing. Because PVM guest systems must be modified, this method will not work with closed, proprietary OS kernels such as Windows.

3.3 Cluster Fencing and Locking

Fencing and locking helps ensure that a failed cluster node does not access shared resources which can lead to data corruption.

RHEV Multiple power management vendors are supported for fencing such as ilo, drac5, and APC devices. Because VMs are presented as LVM volumes, and are not image-based files on an underlying file system, failure to properly fence a failed node could result in some limited amount of data corruption. Specifically, only the VMs that the node was running before the failure would be impacted. Even though the node has access to all the LVM volumes, it should only attempt writes to the volumes that correlate to the VMs it had been running prior to the failure event.

If the failed node was acting as the SPM (Storage Pool Manager), data corruption is averted by having the VDSM process register a lock with SANLock. When the old SPM node fails, the lock is removed and another node in the cluster will contend for SPM status. Once successful, the new node will have its lock set and attempts from the old SPM to write metadata to that new node are prevented.

OVM The primary fencing mechanism OVM employs is the OPROCD and Hangcheck Timer processes provided by the Oracle Clusterware Suite. Both processes monitor system availability through a series of checks. If the system is deemed unstable because of a timeout or check failure, OPROCD will halt the system and initiate a restart.

With all images on OVM residing on a file system, OCFS2, there are extra checks and risk involved. OCFS2 utilizes voting disks to help maintain cluster quorum by adding heuristics to test for availability.

When a server pool is created in OVM, a storage volume is needed of more than 10GB to act as the cluster voting disk. Each node in the cluster writes a heartbeat file to the voting disk on a regular schedule. A kernel thread, o2hb, will set a timer to fence the system. If the heartbeat file is updated within the time allotted, the fence action will be canceled and a new timer will be set.

If a node were to reboot from a fence, another cluster node will lock and replay the failed nodes file system journal. If this process were not to complete correctly, the failed node could attempt to write to its journal concurrently with another node. This would result in data corruption, and the entire file system would need an fsck run to remove the inconsistent inodes. OVM also supports hardware fencing with IPMI which helps protect resources if Oracle Cluster Suite were to fail.

3.4 Role-Based Access Control

RHEV The command *rhevmanage-domains* can connect the RHEV-M service to either IPA or AD. This doesn't open access to every directory user, the RHEV administrator still has control over which users can log in. A selection of predefined roles are available as well as the functionality to create custom roles. The user based control is even extended to resources such as VMs, templates, server pools. Production systems, for example, could be hidden from developers and only visible to desired users.

OVM Although the OVM service runs on Weblogic, which is fully capable of connecting to user directory services, the integration to OVM itself has not been completed. This is a major concern for enterprise customers and limits the usability of OVM to a single administrator or department. Multi-tenant access or the future implementation of a Virtual Desktop Infrastructure are not possible without these basic controls. When questioned, Oracle support provided no insight on the development of these features.

Additional Features

4.1 Virtual Desktop Infrastructure

RHEV Red Hat provides VDI-ready features using the SPICE protocol for usb pass-through and GPU offloading. VM pools can also be created allowing for fast provisioning of new desktop systems. This is accomplished by creating a new VM and customizing it as desired, then creating a template from this VM.

A VM pool can then be created using this template. When a user logs in to the User Portal, they can start a new VM from this pool, and the system becomes available to the end user. Once the user is done with the VM, any changes made will not be saved and the VM will revert back to its original state. To allow for user data to be saved across reboots, the original template should be configured to use a centrally located home directory or mount point. Two great use cases for this feature are outlined below.

1) Using either thin clients for desktops, or a PXE booted minimal OS, employees will only be presented approved desktop images. Any of their data can be stored through autofs or mapped network drives. Even when the client moves to a different thin client system, they can boot into the same exact image and access the same centrally stored user data. For an organization, this can lower desktop licensing, strictly enforce OS standards and prevent unapproved alterations.

2) A team of developers, each with their own desktop, might need access to multiple versions of operating systems to test code or applications. Creating a VM pool for each OS type allows developers to start up the VMs they need, when they need it and shut it down when they don't. This on demand, self service interface helps save resources, money, and time spent during the development and testing release cycles. All development and testing tools can be preinstalled on the VMs as well as a central storage mount for home directories.

OVM OVM does not provide a VDI solution, nor is there any information for such a feature in the future.

4.2 Reporting

RHEV By installing the package *rhev-m-dwh*, a customized JasperServer instance is created containing multiple preconfigured reports. The reporting services track historic data for RHEV resources. These reports are also integrated with the RHEV-M console under the drop down menu "Show Report". This included functionality for reporting allows administrators to easily track historic

performance of the infrastructure and of their application servers. Detailed reports are very helpful for resource and capacity planning.

- OVM** OVM does not provide any reporting or historic data, nor is there any information for such a feature in the future.

Product Support

5.1 Customer Support

- RHEV** Since the RHEV 2 initial release, Red Hat went on a hiring frenzy to bring talented developers and staff on board to support their new product offering. This effort is noticeable when calling or emailing support. Even when faced with unsupported limitations, such as the Firefox admin portal tech preview on RHEV 3.0, Red Hat has been willing and eager to assist with workarounds. It is also quite possible to contact Red Hat support and get a technical engineer on the phone within 10-15 minutes.

In Sintre's experience, the engineers seem happy to be at work, and that attitude goes a long way when faced with complex technical challenges. If a Red Hat product does not have the desired functionality, the support engineers will happily convert the support ticket into a feature request. We have seen multiple enhancements requested on RHEV 2.2 added to later versions of the product.

- OVM** Even with a premier support contract, the level of Oracle support is lacking. When calling Oracle directly, calls are routed to a non-technical representative that either opens a new case or sends an existing case owner a message requesting an update. Speaking directly to an engineer knowledgeable with regards to the OVM product does not seem to be possible, at least in our experience. Despite listing a phone number as the primary contact mechanism, all communication is done through the support ticket and email.

For example, when a Sintre engineer opened a case asking if OVM could support Active Directory or LDAP integration or if virtual machines could somehow be organized in groups through the console, the only response (a week later) was a statement that OVM does not have those features. No insight was provided as to when the product might have these basic enterprise features, nor was there a suggestion that our requests would be added to a bug or feature enhancement list.

Obviously, Oracle does have backline engineers that know the product. Upload a full OVM vmp log report, and a technically sound response will be delivered within the week. Keep in mind, though, that this experience was seen with a \$100K OVM support contract with an overall \$1M Oracle support contract. The quality of support provided to an organization with only \$10K in support contracts seems likely to engender even greater frustration.

5.2 Community Support

- RHEV** When searching for public support information on issues, the results were limited. There seems to be little information about RHEV on the Internet. A YouTube search resulted in mostly sales pitch slides or high level overview demos. It's disappointing to see so little community interaction and general product enthusiasm.

RHEV itself can only be downloaded with a paid subscription or a formal request for a demo. The free version of RHEV is available through ovirt.org, which mirrors the same type of relationship shared by Fedora and Red Hat OS distributions. oVirt is upstream to RHEV, contains separate support forums, and is on a different release version. While it's great that a completely free version is available, the separation between the free and the enterprise version prevents a unified user base. Searching forums for oVirt tips doesn't help an RHEV 3.1 user much, as that product code is based on an older release.

OVM There were many more helpful results found online for OVM. Oracle Customer Support may be lacking, but with the help of other engineers on the public support forums, issues can be quickly resolved. Searching on YouTube will result in pages of technical tutorials and user uploaded support content. Whether the user generated problems and resolutions are from a large production environment or a small home test lab, the code and product cycle are the same and the information is directly related.

Product Migration

6.1 Migrating VMs from Previous Versions

RHEV To migrate VMs from RHEV 2.2 to RHEV 3.1, a storage medium, called an export domain, must be visible to both clusters. First, the export domain is attached to RHEV 2.2. Then the VMs must be shut down and exported to the storage domain. Afterwards, the export domain must be deactivated, then detached from the cluster. The storage domain can then be attached to RHEV 3.1 and the VMs imported.

OVM Oracle's import process is more straightforward. It has the flexibility to handle migrations between remote data centers that do not share common storage domains. From OVM 2.2, a VM can be exported as a template which can then be moved to the OVM-M 3.1.1 node through SCP, FTP, NFS, or physical transfer.

OVM also supports HTTP import methods for templates and ISOs, and the OVM-M system can be configured to host an http repository for import purposes. From there, the console is used to import the VM by providing the full http path to the file. Once the template import is complete, a new VM can be created from that template.

Product Pricing

7.1 Pricing Data

RHEV RHEV standard: \$499/socket
RHEV premium: \$749/socket

OVM OVM premier limited: \$599/up to 2 sockets
OVM premier: \$1199/unlimited sockets

7.2 Sample Configuration Pricing

Example 1:

A 3-node cluster, 2-socket system:

RHEV standard: \$2,994

RHEV premium: \$4,494

OVM premier limited: \$1,797

OVM premier: \$3,597

Example 2:

3 node-cluster, 4-socket systems:

RHEV standard: \$5,388

RHEV premium: \$8,988

OVM premier limited: N/A

OVM premier: \$3,597

Conclusion

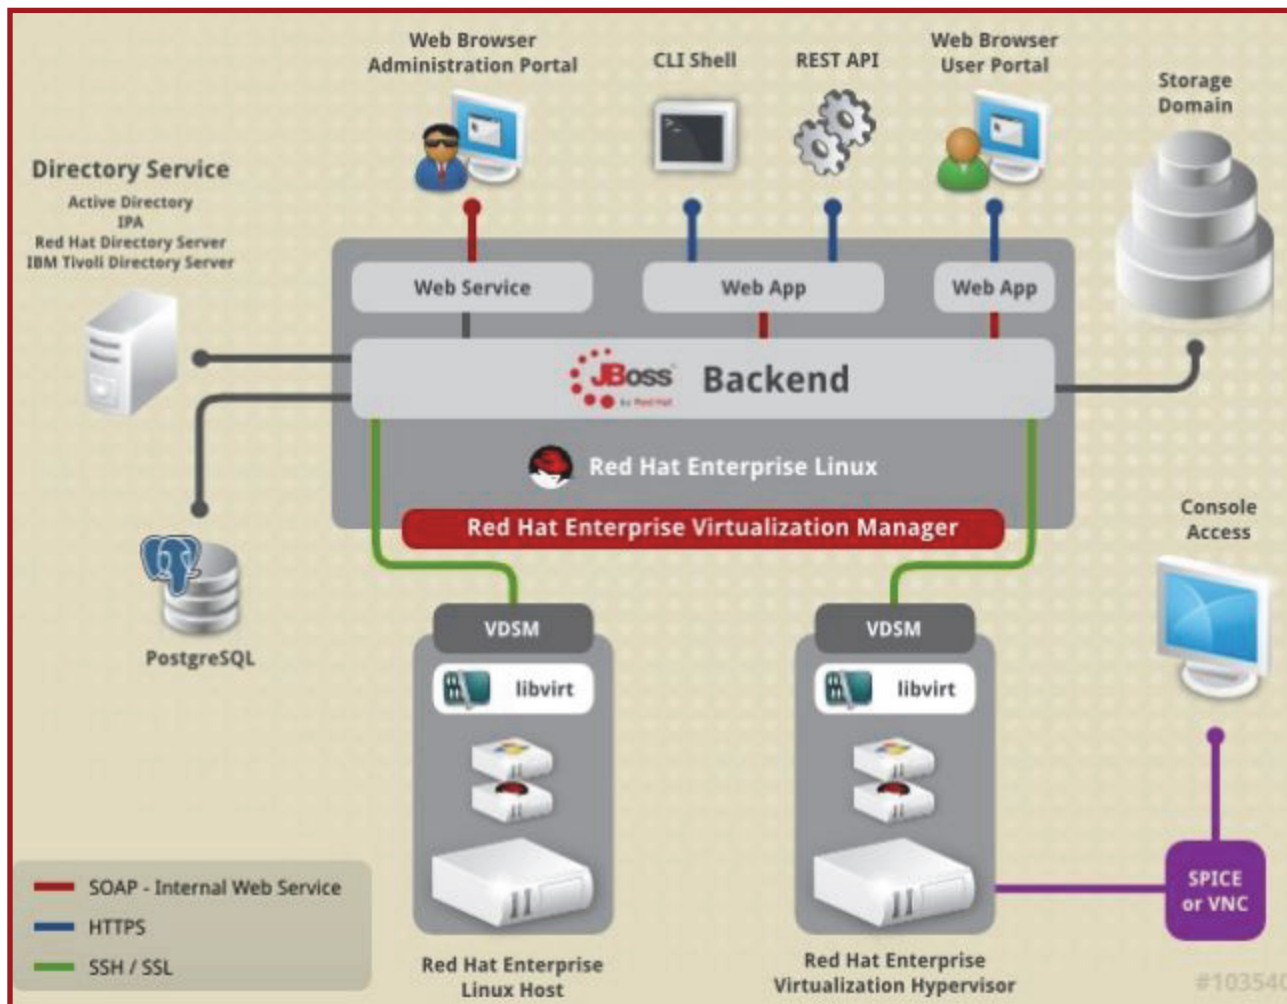
According to estimates by Gartner, nearly half of the x86 architecture workloads worldwide are now virtualized. While the so-called Big 3 virtualization vendors (Citrix, Microsoft, and VMware) dominated the enterprise market in the early days, advances made by the key Linux solution vendors are gathering momentum and market share.

Red Hat and Oracle each delivered mature third generation products in 2012 that bear careful consideration for those companies looking to initiate or extend their virtualization strategy this year. While both are cost-effective implementations, Oracle VM may be better suited to those companies that have already invested heavily in Oracle technologies and for whom price is an overwhelming criteria. For everyone else, the better choice is the Red Hat Enterprise Virtualization solution due to its superior usability, robust feature set, and responsive customer support.

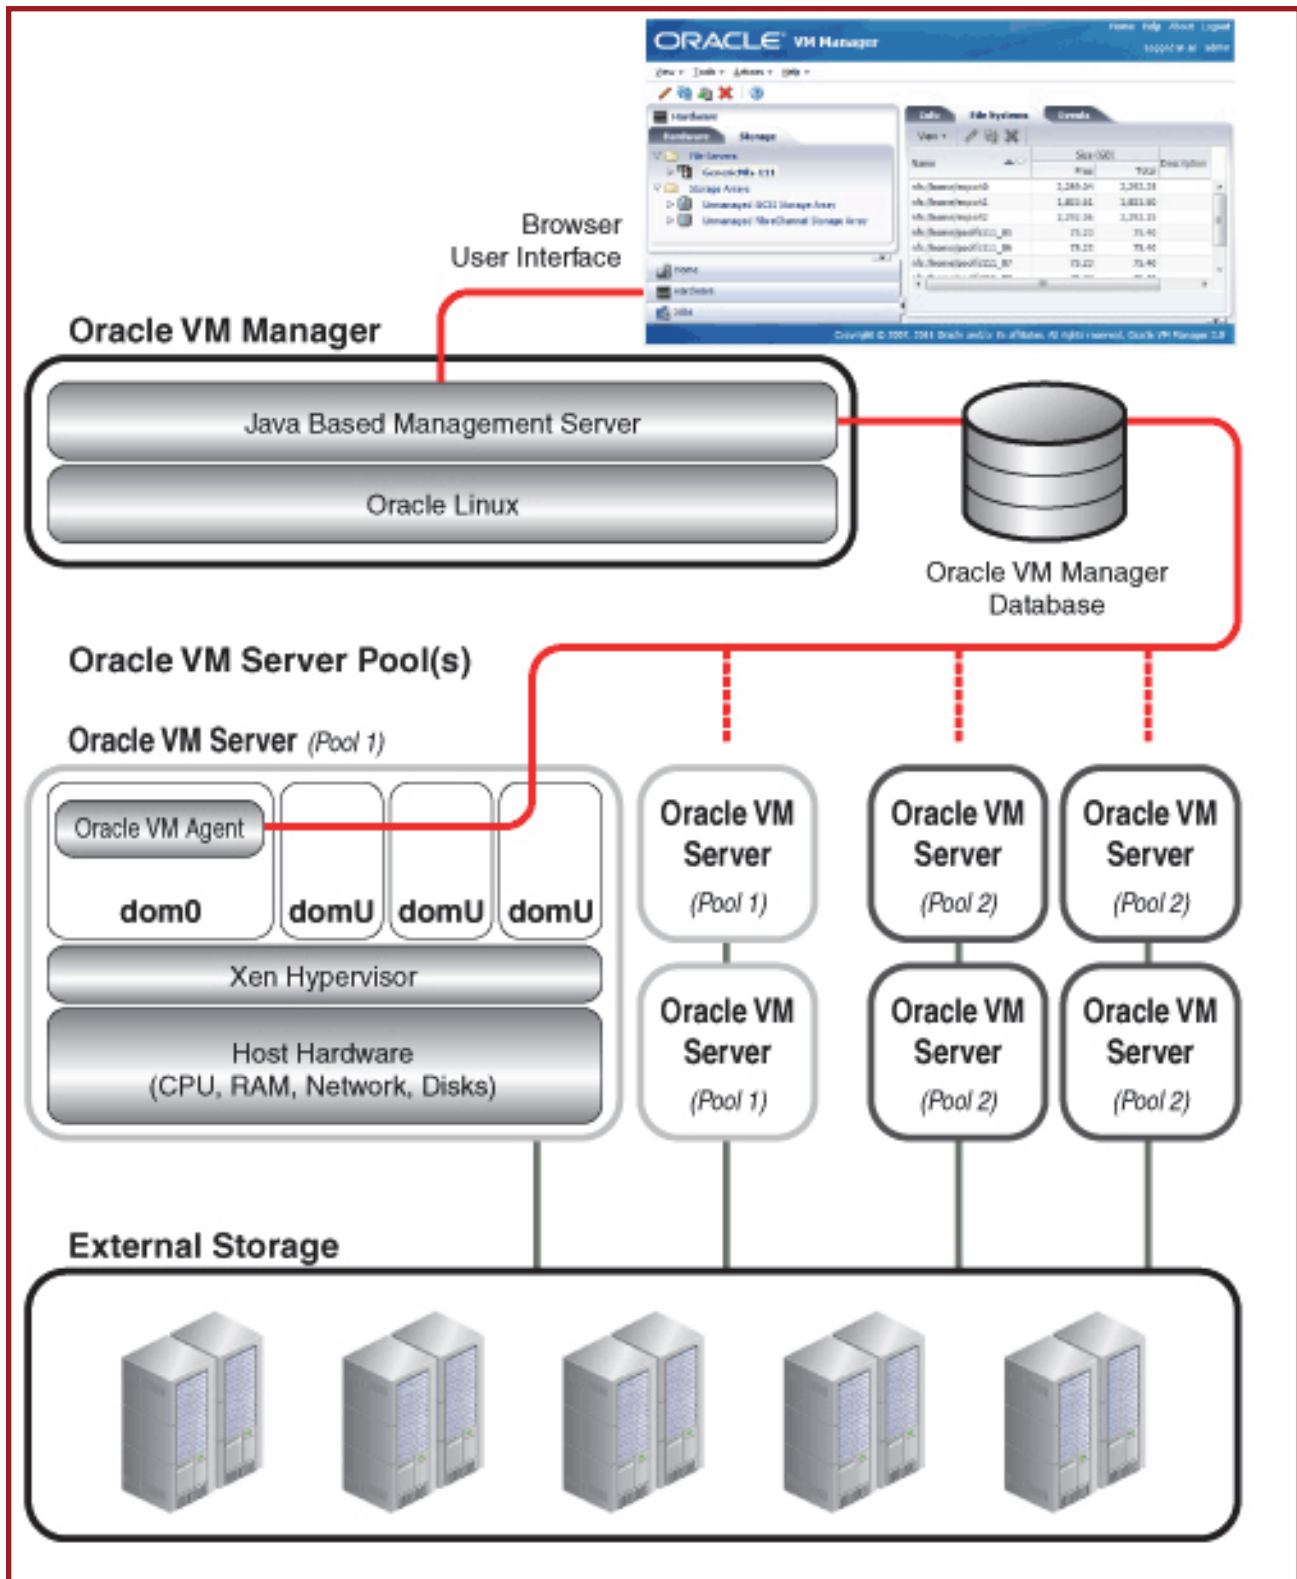
Red Hat RHEV and Oracle VM statements and technical features were based on live enterprise production experience and publicly available documentation and images at https://access.redhat.com/knowledge/docs/Red_Hat_Enterprise_Virtualization/ and http://docs.oracle.com/cd/E27300_01/index.html respectively.

Architecture Diagrams

Red Hat



Oracle



Disclaimer

This document contains information which is the proprietary property of Sintre and its affiliates. This document is received in confidence and its contents cannot be copied without prior written consent.

Nothing in this document constitutes a guaranty, warranty, or license, express or implied. Sintre disclaims all liability for all such guaranties, warranties, and licenses, including but not limited to: Fitness for a particular purpose; merchantability; non infringement of intellectual property or other rights of any third party or of Sintre; indemnity; and all others. Sintre products are not intended for use in medical, lifesaving, or life-sustaining applications. The reader is advised that third parties can have intellectual property rights that can be relevant to this document and the technologies discussed herein, and is advised to seek the advice of competent legal counsel, without obligation of Sintre.

Sintre retains the right to make changes to this document or to related product specifications and descriptions at any time, without notice. Sintre makes no warranty for the use of this document, and it assumes no responsibility for any errors that can appear in the document nor does it make a commitment to update the information contained herein.

Copyright © 2012 Sintre. All rights reserved.

*Other brands and names are the property of their respective owners.