Oracle Private Cloud Appliance Backup Guide

ORACLE PRIVATE CLOUD APPLIANCE

ORACLE WHITE PAPER | MARCH 2017

ORACLE

Introduction

The Oracle Private Cloud Appliance (PCA, formerly called Virtual Compute Appliance) is an integrated hardware and software system designed to reduce infrastructure complexity and deployment time for virtualized workloads. It is a complete platform for a wide range of application types and workloads, with built-in compute, storage and networking. PCA automatically discovers components and configures them to work with one another, reducing design and administrative effort, eliminating potential errors, and speeding time to application deployment. It enables virtualization life cycle management using Oracle VM.

This document reviews Private Cloud Appliance architecture, describes automated internal system backups of software components, and describes how to backup data, including PCA system data, Oracle VM repositories and database, and virtual machine contents.

Scope

This document covers backing up data on the Private Cloud Appliance (PCA). The Private Cloud Appliance software stack consists of the Private Cloud Appliance Controller, Oracle VM Manager and their repository and data objects. PCA system data is located on the built-in (internal) ZFS storage appliance. Virtual machine data may also reside on the internal ZFS and on optional and recommended external storage. This document is intended for use with Private Cloud Appliance 2.0 and later.

Because the Private Cloud Appliance is an engineered system and is managed as a unit, infrastructure data restore (as distinguished from user data in VMs) is restricted to specific use-cases. In general it is expected that Oracle Support would be contacted to assist in restoring system infrastructure data. Virtual machine data, in particular contents of Oracle VM repositories, is in most regards backed up the same way on PCA as on non-engineered systems running Oracle VM, with PCA-specific considerations described below.

Private Cloud Appliance Availability Concepts

The Private Cloud Appliance (PCA) is designed with redundant management, compute, network and storage components to ensure that failure of any single component does not affect overall system availability. PCA takes automated internal backups of key infrastructure data to capture configuration and system state. These measures protect against individual hardware component failures and accidental deletion or data corruption. This document will discuss the different types of data and how they are backed up.

Internal backups do not protect against catastrophic failures such as lost power or system damage, so this document also discusses how to backup data to storage external to the PCA system. Taking regular backups is part of standard operating procedures for all production systems.

The Private Cloud Appliance contains a ZFS storage appliance residing on the PCA internal networks. This ZFS appliance serves as the PCA “system disk” and also as the default location for Oracle VM data. Data must be copied from the ZFS appliance using hosts that are on both datacenter and internal PCA networks. A recommended method is to replicate ZFS shares to an external InfiniBand-connected ZFS storage appliance using the ZFS appliance share replication feature. Alternatively, data can be copied to external storage using the PCA management nodes or an appliance VM, or (starting with PCA 2.2) a compute node with a custom host network.

Oracle recommends that critical systems be protected using disaster recovery (DR) practices and technologies. An example of this is a PCA system located at a different site, with content kept up to date with changes made on the primary system. That would provide standby capability in the event of failure of the primary system. Disaster Recovery for Oracle VM 3, including PCA, is described in the MOS note “Oracle VM 3: Getting Started with Disaster Recovery (Doc ID 1959182.1)”.

Architecture and Components

This section summarizes the architecture of the Private Cloud Appliance, with specific attention to the data on each component. For further information, please consult the *Oracle Private Cloud Appliance Release 2.2 Administrator's Guide*. Each Private Cloud Appliance consists of:

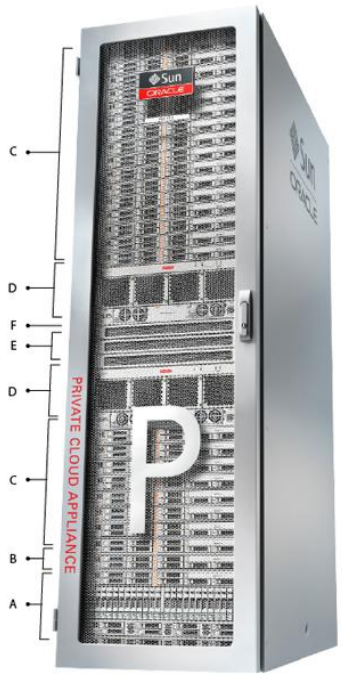

Figure 1. Components of Private Cloud Appliance Rack

TABLE 1. PRIVATE CLOUD APPLIANCE COMPONENTS LEGEND

ITEM	QUANTITY	DESCRIPTION
A	1	Either Oracle ZFS Storage Appliance ZS3-ES or Sun ZFS Storage Appliance 7320
B	2	Either Oracle Server X5-2, Sun Server X4-2, or Sun Server X3-2, used as management nodes
C	2 -25	Either Oracle Server X6-2, X5-2, Sun Server X4-2, or Sun Server X3-2 used as virtualization compute nodes
D	2	Oracle Fabric Interconnect F1-15 Director
E	2	NM2-36P Oracle Datacenter InfiniBand Expansion Switch

F	2	Oracle Switch ES1-24

Management Nodes

The heart of each Private Cloud Appliance is a pair of dedicated management nodes, arranged in an active/standby cluster for high availability. Both servers can run the same services and have equal access to the system configuration, but one operates as the master while the other is in standby mode. The standby node automatically takes over if a failure occurs. The master node runs the full set of services, in particular Oracle VM Manager, the PCA Controller which handles provisioning, the Dashboard, failover, and other system services. These services are briefly unavailable during a failover. This does not affect Oracle VM Server or virtual machine operation on the compute nodes, which continue without interruption. The standby node runs a subset of services until it is promoted to the master role, at which time the previous master node assumes the standby role..

Management nodes boot off local disks, arranged in RAID pairs for media resiliency, and access all system-wide data stored on the built-in Oracle ZFS Storage Appliance. The master role is determined via OCFS2 Distributed Lock Management on an iSCSI LUN which both management nodes share on the Oracle ZFS Storage Appliance. The management node that acquires the lock assumes the master role.

Management nodes access the Private Cloud Appliance internal networks using pre-defined network addresses, and access customer datacenter networks via addresses defined by the system administrator at install time. Each management node has an IP address on the customer datacenter network, and a virtual IP (VIP) address is assigned to the node currently owning the master role. This provides external connectivity between management nodes and the customer network for browser access to the Private Cloud Appliance Dashboard and Oracle VM Manager, and for data transfer. The process for assigning these addresses is described in the Installation Guide.

Compute Nodes

Oracle Server X3-2, X4-2, X5-2 or X6-2 compute nodes in the Private Cloud Appliance constitute the virtualization platform. They run Oracle VM Server and provide processing power and memory capacity for virtual machines under Oracle VM Manager's control.

Compute nodes are automatically configured into the Oracle VM environment by a provisioning process orchestrated by the management nodes. Private Cloud Appliance software installs Oracle VM Server software on each compute node, defines their network configurations, and places all compute nodes into an Oracle VM server pool. Starting with PCA 2.2, administrators can define “tenant groups”, which isolate compute and storage resources in separate server pools which can be assigned to different customers.

Each compute node has local storage to host the Oracle VM Server boot environment. Local disk capacity can optionally be used as Oracle VM repositories for virtual machines running exclusively on that compute node, and can be used for applications that do not require being able to run on another compute node. This is not a typical use case, and virtual disks typically reside in shared repositories as described in the section on storage, below.

Network Infrastructure

The Private Cloud Appliance relies on a combination of Ethernet connectivity and an InfiniBand network fabric using “wire once” Software Defined Networking (SDN) using Oracle Virtual Networking (OVN). The appliance contains redundant network hardware components, pre-cabled at the factory, to help ensure continuity of service in case a failure should occur.

The PCA uses private, “internal” networks which are not exposed to the customer’s datacenter network. This provides isolation, security, and the ability to use pre-defined IP address ranges for each networked component without duplicating datacenter network addresses. The internal Ethernet network relies on two interconnected Oracle ES1-24 switches, to which all other rack components are connected. This network serves as the appliance management network, in which every component has a predefined IP address in the 192.168.4.0/24 range. All Oracle Server management and compute nodes have a second IP address in this range (the OS IP address with 100 added to the last octet), used for Oracle Integrated Lights Out Manager (ILOM) connectivity.

One of the Oracle ES1-24 switches has an Ethernet cable attached to port 24 or port 19, which the administrator uses during initial setup to connect a workstation with fixed IP address 192.168.4.254. From this workstation, the administrator opens a browser connection to the web server on the master management node at `https://192.168.4.3:7002/dashboard`, in order to monitor the initialization process and perform the initial configuration steps when the appliance is powered on for the first time. This is an example of an external bastion host with access to the PCA internal network.

The Private Cloud Appliance rack also contains two Sun Datacenter InfiniBand NM2-36P Expansion Switches. Each management node, compute node and storage head connects to both switches for redundancy. Both InfiniBand switches, in turn, have redundant cable connections to both Oracle Fabric Interconnect F1-15 directors in the rack. These components form a physical InfiniBand backplane with a 40Gbit (Quad Data Rate) bandwidth.

Private Cloud Appliance network connectivity is managed through two Oracle Fabric Interconnect F1-15 directors. Data is transferred across the physical InfiniBand fabric, but connectivity is implemented in the form of Software Defined Networks (SDN). SDNs dynamically connect virtual machines and bare metal servers to networks, storage and other virtual machines, maintaining the traffic separation traditionally provided by hard-wired connections and exceeding their performance.

Compute nodes are connected to the internal networks and to the customer datacenter networks. Oracle VM Server on each compute node communicates over Private Cloud Appliance internal networks for management, storage, heartbeat and live migration. By default, compute nodes do not have IP addresses on the customer datacenter network, which increases their isolation and reduces attack surface. Starting with PCA 2.2, custom networks can be created to give compute nodes IP addresses on the customer network.

Guest virtual machines access the customer datacenter network using a separate Oracle VM network named "`vm_public_vlan`", which can be used for both VLAN and non-VLAN environments. An additional virtual machine network named "`vm_private`" is internal to the Private Cloud Appliance and used for private, high-performance, low-latency network traffic between virtual machines. Both networks can be optionally used with VLANs for network separation. A network named "`mgmt_public_eth`" is provided to access the Oracle VM management console, but can optionally be used for virtual machine traffic too. These networks are defined in Oracle VM Manager as only having the "Virtual Machine" function (also called a "channel") - indicating they are used for guest VM TCP/IP traffic and not cluster management, storage, or live migration. This ensures that guest VMs do not see infrastructure network traffic. VLANs can be defined to isolate VMs from one another.

These essential SDNs are configured during the Private Cloud Appliance initialization process:

- The storage network is a bonded (based on redundant links) IP over InfiniBand (IPoIB) connection between the management nodes and the ZFS storage appliance, using the `192.168.40.0/24` subnet.
- The Oracle VM management network is a Private Virtual Interconnect (PVI) that connects management and compute nodes in the `192.168.140.0/24` subnet. It is used for all network traffic by Oracle VM Manager on the active management node and Oracle VM Server on the compute nodes.

» The virtual machine networks are the `vm_public_vlan` and `vm_private` networks described above. They are bonded Ethernet connections on each compute node. Oracle VM uses the networks for virtual machine connectivity, internal to the appliance and on the data center public network. Only guest virtual machines have IP addresses assigned on these networks, ensuring isolation from the Private Cloud Appliance system components. The subnets are configurable, and `vm_public_vlan` connects the public interfaces of the compute nodes to the data center network. A default VLAN segment of 1 is used for non-VLAN tagged networks. Additional VLAN networks can be defined for multi-tenancy and network isolation.

Starting with PCA release 2.2, administrators can create custom networks, described in the section “Network Customization” in the PCA Administrators Guide. Custom networks may be internal to the PCA or external. A “host network” is an external network that provides network connectivity to the compute nodes, so compute nodes are plumbed with datacenter network IP addresses. This can be used for connectivity to external storage or backup servers.

Default networks are represented in Oracle VM Manager as shown in Figure 2:

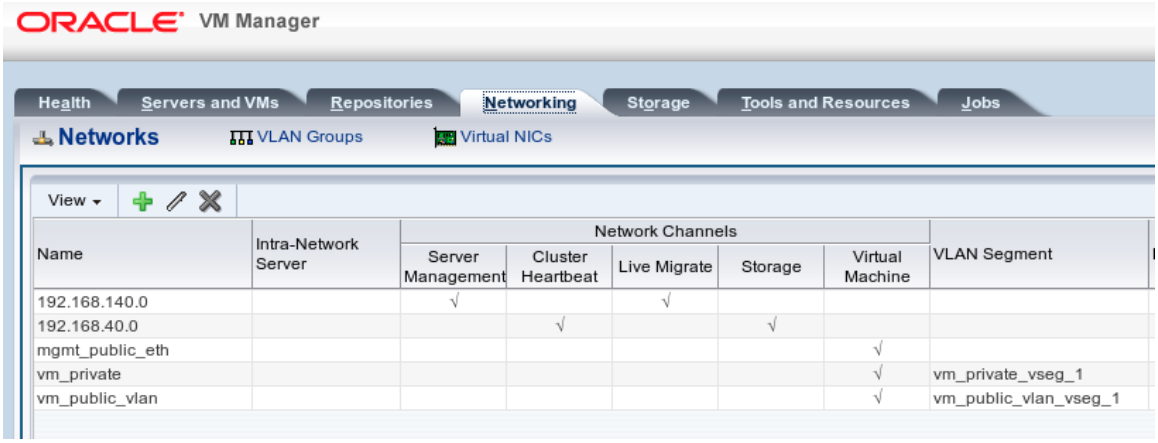

Figure 2. Oracle VM Manager Networking view

The Oracle Fabric Interconnect F1-15 directors also manage the physical public network connectivity of the Private Cloud Appliance. Two 10GbE ports on each F1-15 **must** be connected to next-level data center switches as described in the Private Cloud Appliance Installation Guide. During the installation process, the administrator assigns three reserved IP addresses from the data center (public) network range to the management node cluster of the Private Cloud Appliance: one for each management node, and a Virtual IP owned by whichever management node is currently the master. The Virtual IP is used to connect to the master management node's web server, which hosts both the Private Cloud Appliance Dashboard and the Oracle VM Manager web interface.

Configuration data for each of these essential network infrastructure components is backed up by the PCA's automatic internal backup process as described in the section "Private Cloud Appliance Internal Backup".

Storage

The PCA comes with built-in storage, using the Oracle ZFS Storage Appliance (originally a ZFS 7320, now a ZS3 model). This “internal ZFS” contains PCA system data, and should be considered a 'system disk' for the entire PCA. It also contains a default Oracle VM repository which can store VM disk images and templates. Backing up this repository is similar to backing up any Oracle VM repository. The internal ZFS is on PCA private networks and not exposed to the customer’s Ethernet network. Steps to back it up to external storage are described below.

PCA supports compute node connection to external storage for Oracle VM repositories and LUNs presented to virtual machines. Storage arrays like the Oracle FS1 Flash Storage System can be connected via Fibre Channel, and a ZFS storage appliance can be connected via InfiniBand. If host networks are configured on PCA 2.2 and later, then Ethernet-based NFS and iSCSI storage on the customer network can also be used. Oracle strongly recommends using external storage for the scale, performance, and management and backup capabilities provided by enterprise-grade storage arrays like Oracle ZFS. The PCA's internal ZFS is suitable for moderate performance and capacity requirements. Oracle staff can help size storage to meet application requirements.

Guest virtual machines generally use virtual disks from a repository. VMs can also use iSCSI or Fibre Channel LUNs presented directly to the VM for optimized performance. These are described as "physical disks" in Oracle VM documentation, and appear to the VM as local disks. Virtual machines can additionally connect to a datacenter's networked storage based on NFS, CIFS or iSCSI protocols. This is done under virtual machine operating system control just as with physical server environments, and is transparent to PCA or Oracle VM operation..

PCA Internal ZFS Details

Compute and management nodes boot off their own local disks and access system-wide data stored on the ZFS appliance. The ZFS appliance provides storage space for the Private Cloud Appliance software, including the Oracle VM Manager, and OS boot images for network-installing compute nodes. It also contains backup copies of the configuration data for every system component. Compute and management nodes access the ZFS appliance via the IP over InfiniBand (IPoIB) network described previously. This network is private to the appliance, and provides high performance, isolation, and security.

The Oracle ZFS Storage Appliance contains two clustered storage heads for redundancy, so data remains available even if a head fails. Data is arranged in a redundant RAIDZ2 configuration that is optimized for hardware fault tolerance, detects and corrects data errors, and can tolerate media failures without data loss. Total usable space is approximately 11TB. Cache and log SSD devices are used to improve performance. This provides a balance between performance, available space, and redundancy for data protection. Approximately 6TB is available for a Oracle VM storage repository containing virtual disk image files, ISO images and templates with moderate performance requirements. External storage should be used for higher capacity and performance requirements.

The storage pool contains projects named OVCA (from the original product name) and OVM, with predefined iSCSI LUNs and NFS file systems. The OVCA project contains all LUNs and file systems used by the PCA software. The most significant of these are:

- LUNs:
  - `Locks` - used exclusively for cluster locking of the two management nodes. The two nodes contend for the lock to determine which assumes the master role.
- File systems:
  - `MGMT_ROOT` - used for storage of all files specific to the Private Cloud Appliance, and mounted via NFS on each management node at `/nfs/shared_storage`. Its contents and backup are described below.

The OVM project contains all LUNs and file systems used by Oracle VM. The iSCSI LUNs are automatically deployed for use by Oracle VM, and the file systems are made available for customers preferring NFS:

- LUNs:
  - `iscsi_repository1` – the default Oracle VM storage repository
  - `iscsi_serverpool1` – the server pool file system for the Oracle VM clustered server pool.
- File systems:

- `nfs_repository1` - available to be used as an additional Oracle VM storage repository in case NFS is preferred over iSCSI.
- `nfs_serverpool1` - available for a server pool file system for the Oracle VM clustered server pool. In practice unused since there is one pool and its pool file system is already allocated on an iSCSI LUN.

As mentioned above, Oracle recommends using external storage to increase performance, capacity and management. The storage can be a ZFS storage appliance connected by InfiniBand, or other storage connected by Fibre Channel via a Fibre Channel switch, or NFS and iSCSI over a host network. The Private Cloud Appliance documentation describes how to attach and configure external storage. Virtual Machines can themselves use the 10GbE Ethernet datacenter network to connect to existing networked storage using NFS, CIFS or iSCSI just as if they were running on physical servers, without requiring administrative support from PCA or Oracle VM. Backup of external storage is identical to backup on non-PCA Oracle VM deployments.

Data Backup Categories

Data components on the Private Cloud Appliance have different backup and restore requirements, depending on their purpose, where they reside, and how they are created. For the highest level of data integrity, Oracle recommends using application aware backup technologies consistent with Oracle Maximum Availability Architecture (MAA), just as on physical servers. This leverages well-established backup and recovery practices, enables transactional awareness, provides data integrity superior to "crash consistency" and is more efficient than relying exclusively on virtual disk image backups.

- Private Cloud Appliance system data containing system state and configuration for PCA itself. This data is backed up internally from PCA and Oracle VM components (including Oracle VM Manager MySQL backup and ZFS Storage Appliance configuration) to a directory on the internal ZFS appliance. This protects from accidental removal or corruption of the backed up components, but not loss of the ZFSSA. A few simple commands can copy data to external storage (illustrated below). It is expected that Oracle Support would assist in any restore operation to ensure the Private Cloud Appliance is in a consistent state.
- Compute nodes, running Oracle VM Server, have local per-node storage which is not accessible by other servers and is not backed up. Oracle VM Server contents are recreated when a node is re-provisioned so backing them up is not required. If local disk repositories are used, their contents should be copied using documented Oracle VM repository backup procedures
- Virtual machine data resides in repositories that contain virtual machine metadata, virtual disk image files, and pre-built virtual machines in the form of assemblies and templates. It may also include LUNs presented directly to virtual machines (referred to as "physical disks" in Oracle VM documentation). This data can be backed up and restored by the customer using documented Oracle VM recovery practices. Guest virtual machines are on the datacenter network, and their file contents can be backed up by scripts or backup product clients exactly as if they were physical machines using existing backup methods from Oracle and third party vendors. This can be used in addition to backing up entire virtual disk images for MAA-consistent data integrity.
- The Oracle ZFS Storage Appliance is the storage location for the above data categories, except for the optional external storage. Backing up ZFS storage appliance contents is the primary mechanism for backing up PCA.

Internal Backup

The configuration data of all components within Private Cloud Appliance is automatically backed up and stored on the Oracle ZFS Storage Appliance as a set of archives. Backups are named with a timestamp based on when the backup is run to make it easy to identify backup dates. A `crontab` entry on the active management node starts a backup job every day at 9am and 9pm.The primary purpose of this backup data is to be restored in case of emergency, in coordination with Oracle Support.

Backups are stored on the MGMT_ROOT filesystem on the Oracle ZFS Storage Appliance and are accessible on each management node at `/nfs/shared_storage/backups`. The backup process creates a compressed tar file named with the timestamp for the current backup job. For example, the backup taken at 9am on September 8, 2015, is named "`2015_09_08-09.00.01.tar.bz2`". The compressed tar file contains several subdirectories:

- `nm2`: contains the Sun Datacenter InfiniBand NM2-36P Expansion Switch configuration data
- `opus`: contains the Oracle ES1-24 switch configuration data. Contents are saved as XML files named `OpusSwitchConfig192.168.4.200.bak` and `OpusSwitchConfig192.168.4.201.bak`
- `ovca`: contains the configuration information relevant to the deployment of the management nodes such as the password wallet, the network configuration of the management nodes, configuration databases for the Private Cloud Appliance services, and DHCP configuration.
- `ovmm`: contains the most recent backup of the Oracle VM Manager database, the source data files for the current database, and the UUID information for the Oracle VM Manager installation. The backup process for the Oracle VM Manager database is handled automatically from within Oracle VM Manager and is described in detail in the *Oracle VM Installation and Upgrade Guide* in the section *Oracle VM Manager MySQL Backup and Restore*
- `xsigo`: contains the configuration data for the Oracle Fabric Interconnect F1-15 units. The data is in XML format and is equivalent to logging into the directors and issuing commands "`system export testexport.xml`" followed by "`file copy testexport.xml scp://user@bastion-host/`" The active PCA management node also makes equivalent backups of the Oracle Fabric Manager (OFM), formerly called the Xsigo Management System (XMS), to `/opt/xsigo/xms/xms-backups/` and `/opt/xsigo/xms/director-backups`.
- `zfssa`: contains all of the configuration information for the Oracle ZFS Storage Appliance ZS3 as described in section *Configuration Backup* in the *Sun ZFS Storage Appliance Customer Service Manual*. Configuration data can be imported into the ZFS Storage Appliance and then restored to active state.

The backup process collects data for each component in the appliance and stores it in a way that makes it possible to restore that component to operation in the case of failure. Note that there is no automated process to remove old backup tar files from `/nfs/shared_storage/backups`, so this directory will consume increased disk space over time, but still small relative to the ZFS appliance capacity. Customers should remove old files consistent with their retention policies.

Restore should be undertaken in conjunction with Oracle Support to ensure proper procedural steps are used. The exception to this rule is the `ovmm` subdirectory, which can be used to restore an Oracle VM Manager database using documented Oracle VM methods.

Exporting Internal Backup Data to External Storage

The previous section describes data stored on the PCA's ZFS Storage Appliance. This should be copied to external storage on a regular basis to handle a catastrophic loss that would make the storage appliance unavailable.

Backup data tar files can easily be copied from the management node with `scp` or `rsync` commands. For example, commands like the two below can be issued on an Oracle Linux or Oracle Solaris host on the customer network (assuming the public hostname for the management node is as shown below, and adjusting destination names):

```
# scp root@mn1.sample.com:/nfs/shared_storage/backups/2015_09_08*bz2 .
```

```
# rsync -azP root@mn1.sample.com:/nfs/shared_storage/backups/ /home/pcabackups
```

This first command copies the internal backup created on September 8, 2015 (assuming the datacenter network hostname of the management node); the second command synchronizes the entire backup directory. Backups can also be initiated from the node by issuing a command like:

```
[root@ovcamn05r1 ~]# scp /nfs/shared_storage/backups/* backupuser@backuphost:
```

Alternatively, the management node could mount an NFS server on the datacenter network and simply copy the directory tree under `/nfs/shared_storage/backups/` onto it.

Another method is to add the Virtual Machine role to the management network defined in Oracle VM Manager, and create an appliance virtual machine with a virtual network interface on that network. This would let the appliance VM access the internal backup on the ZFS Storage Appliance and the Oracle VM repositories, while also having network access to the public datacenter network. This is described in MOS note "How to Create Service Virtual Machines on the Private Cloud Appliance by using Internal Networks (Doc ID 2017593.1)". The virtual machine would use methods similar to those described above, and would have the advantage of letting the customer choose their own management software, network MTU, and retaining customization over a PCA upgrade. Customers should evaluate the security and isolation implications of using a specially configured VM that has access to internal PCA data, and the possibility of accidentally providing virtual NICs on the management network to other VMs, versus using the management node for backup functions. Such a VM can enhance security compared to using the management node by being protected with firewalls and using the installation's authentication standards.

Oracle VM Repositories and Backup Data

The preceding section discussed infrastructure configuration data, including the Oracle VM Manager's MySQL backup. It is also important to backup Oracle VM Manager repositories, which contain virtual disks, templates, assemblies, and virtual machine descriptor files. Successful recovery of any Oracle VM environment requires backup and recovery of repositories. Most practices for backing up Oracle VM inside the Private Cloud Appliance are the same as in the general Oracle VM environment. For a full discussion of Oracle VM backup and recovery, please read the *Oracle VM User's Guide* and *the Oracle VM 3: Backup and Recovery Best Practices Guide*.

Repository Locations and Types

The Private Cloud Appliance provisions an Oracle VM repository per PCA rack, named `RackN-repository` where *N* is the rack number. The repository is available to all compute nodes on the rack and resides on either iSCSI (the default) or NFS. This is the default internal ZFS repository for virtual disks, assemblies, ISO images, templates and virtual machine descriptions for VMs. Additional repositories would be created if tenant groups are deployed.

An external ZFS appliance can be connected to the InfiniBand storage network and shared by all PCA hosts. This is the recommended deployment architecture for scale, performance, and ease of management. The external ZFS appliance would be connected to the PCA InfiniBand storage network for virtual machine operation, and to the customer datacenter Ethernet network for backup, replication to other sites, for data access by non-PCA clients, and for management access to the ZFS appliance browser interface. Contents on the internal ZFS appliance can be replicated to the external one, and vice versa. The PCA can be connected to other external storage via Fibre Channel or Ethernet. This also enhances scale and performance, and permits using the customer's existing storage infrastructure and standard backup methodology. **Backing up repository data on external storage is identical to backing up Oracle VM repositories in non-PCA environments**.

The PCA system also creates local disk repositories on the internal disks of each compute node when the node is provisioned. This extends the disk capacity of the system by making use of the approximately 800GB to 1TB of disk space remaining on each compute node's disks after Oracle VM Server is installed. The actual amount depends on whether the compute node is an X3-2 or the later X4-2, X5-2, or X6-2. As more compute nodes are installed, more disk capacity becomes available. Each local repository is named `ovcacnXXrY-localfsrepo`, where `ovcacnXXrN` is the name of the compute node (*XX* is the rack unit position, and *N* is the rack number).

Local repositories can only be used by virtual machines on the one compute node they are attached to, and become unavailable if the compute node fails. This potential single point of failure should be considered when deciding where to store virtual machine data. Virtual machine contents that can be recreated can be usefully kept there, improving the overall storage capacity of the PCA environment.

Virtual Machine and Repository Data Backup

Data can be protected at the file, virtual disk and virtual machine, and repository levels. Oracle recommends that application data on virtual disks be backed up on a file basis by the virtual machines when file level granularity and transactional semantics are needed. This can be done using application software, command line utilities or networked backup products just as on a bare metal physical system. This provides transaction and data integrity and permits application control. For example, a virtual machine running Oracle Database can use RMAN to copy its contents to an external backup destination. Virtual machines can also use third party backup vendor products to do file-level incremental backups to backup servers.

Application-level backup within a virtual machine is ideal for transaction-level control, but is not a substitute for backing up a VM's complete virtual disk images. VM level "online backups" can be made by cloning a VM using the Oracle VM Manager's user interface or CLI. Oracle VM supports "thin clones" of running VMs whose disks are on OCFS2 format iSCSI or fibre-channel LUNs repositories, which are used with PCA. Clones can be used as crash-consistent VM snapshots. Data integrity is ensured by taking VM clones when the VMs are not running. This does not provide protection against loss of a PCA system or a repository, but can be used in case of data corruption or to back out an OS or application change in the VM.

Repositories themselves should be backed up to protect against storage loss or corruption. They can be backed up using ZFS share replication from one ZFS appliance to another, or by using block level replication features available with various enterprise-grade storage arrays. When any external storage array is used for Oracle VM contents, data can be backed up using standard Oracle VM methods.

Internal ZFS appliance repository data can be backed up using procedures in the Oracle VM User's Guide section "Enabling Storage Repository Back Ups". This creates an NFS share that exports a repository's contents. The mount is exported to a specified host from a compute node in the pool. A local repository is exported from the compute node owning it, and the shared repository can be exported by any compute node in the rack. The NFS share can be mounted on a bastion host, and virtual machine configured to the internal network as described previously at MOS note 2017593.1, or the management node, which has network access to the internet network available to the compute node, and to the customer's datacenter network. This can be done by performing the following steps:

First, navigate in the Oracle VM Manager User Interface to the Servers and VMs tab. Under `Server Pools`, expand `Rack1-ServerPool` to show the server names, and highlight one of the servers to select it. Then select `Repository Exports` from the `Perspective` view in the management pane as shown in Figure 3.

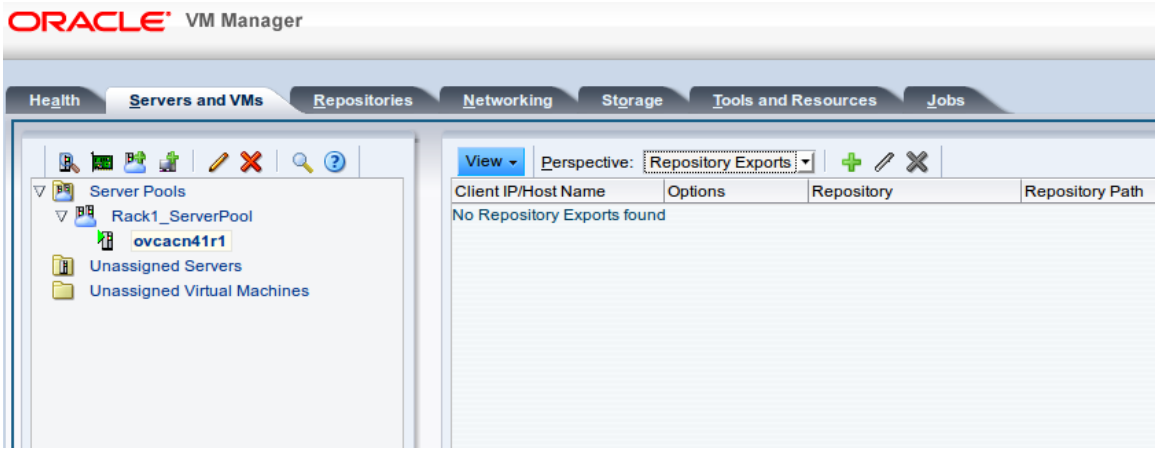

Figure 3. Oracle VM Repository Exports view

Next, click on the  icon to create a repository export, which opens the following dialog. Select the repository to be exported, and the IP address or hostname of the host that will mount the NFS share. In this example, we use a management node, since it has network access to the compute nodes and to the datacenter network; it could also be a service virtual machine. Specify "`ro,no_root_squash`" to ensure a read-only mount that makes all of the repository files visible, as shown in Figure 4.

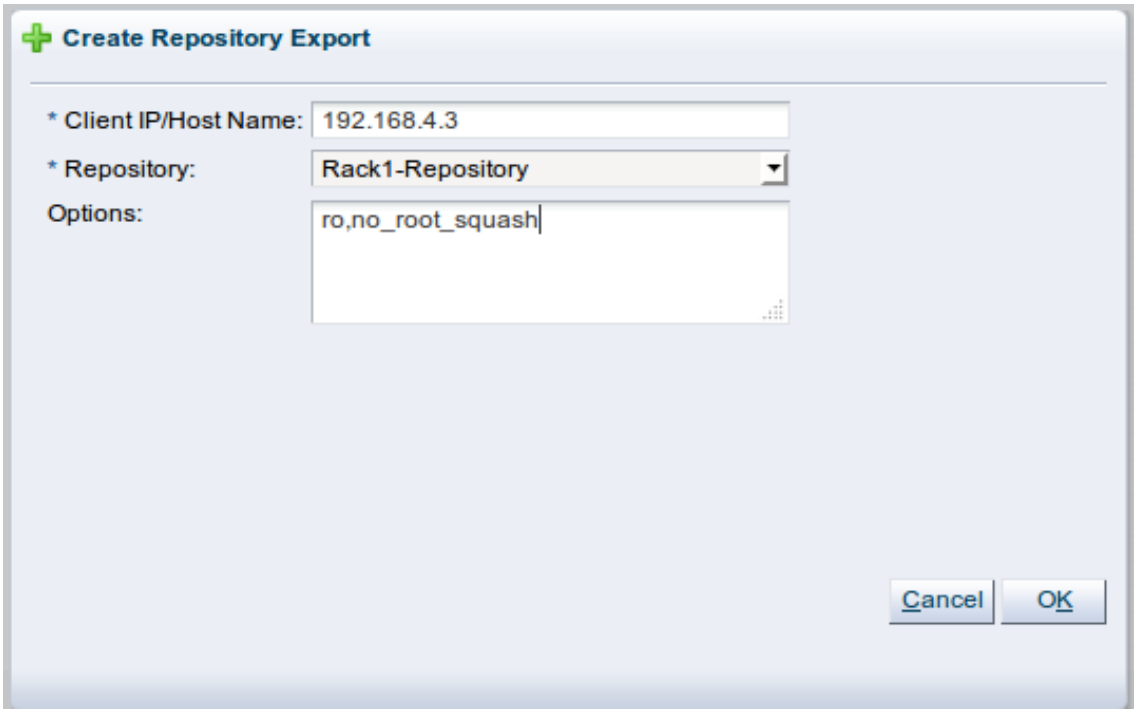

Figure 4. Create Repository Export

After filling in the dialog boxes, click OK. The user interface will show the name of the exported path, as shown in Figure 5. This path can be cut-and-paste into a command line to issue an NFS mount.

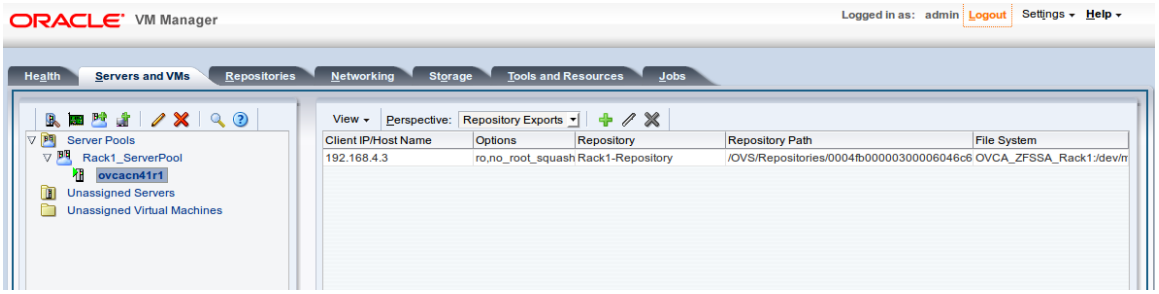

Figure 5. Show Repository Export

Once mounted, the exported file system can be browsed and copied to an external location. Log into a management node, mount the share, and copy its contents to a network destination external to the PCA environment:

```
# mount 192.168.4.5:/OVS/Repositories/0004fb00000300006046c6c5fbcb91e9/ /mnt
# ls -l /mnt
drwx------ 3 root root 3896 Oct  6 04:27 Assemblies
drwx------ 2 root root 3896 Oct  5 04:28 ISOs
drwxr-xr-x 2 root root 3896 Oct  5 04:28 lost+found
drwx------ 3 root root 3896 Oct  6 05:53 Templates
drwx------ 2 root root 3896 Oct  8 02:37 VirtualDisks
drwx------ 5 root root 3896 Oct  8 02:37 VirtualMachines
# scp -rp /mnt/* backupuser@backuphost:
```

This provides a complete backup of the Oracle VM repositories. This procedure uses one of the management nodes as a bastion host to access the repository and copy it to external storage, but the same method could be used with a service virtual machine, as discussed previously. This process is used with standard Oracle VM deployments, and can be managed using documented Oracle VM procedures. PCA systems using software level 2.2.1 or later can create a host network providing network connectivity to the compute nodes, and then directly copy repository contents to external storage without needing an NFS export or bastion host.

Storage Array Snapshot, cloning, and replication

The preceding methods are host-based and “Oracle VM centric”, using Oracle VM configuration administrative interfaces to expose important data and create backups. A very effective approach is to use the capabilities available in many storage arrays to provide data protection. This is largely out of the scope of this white paper, as details depend on the different storage array features. This applies to Oracle VM in general rather than being specific to the Private Cloud Appliance, and there are many attractive options.

For example, LUNs or NFS shares can be cloned to provide copies to use for fallback, or to be replicated between different storage devices. Data can be replicated between the internal ZFS appliance and an external ZFS, and then backed up to a remote ZFS appliance or to tape. Similar replication can be done with Fibre Channel devices. This could be used for local backup or for disaster recovery architectures. Combined with application and VM level backup under Oracle VM, this can be a robust approach for data protection.

Conclusion and further reading

This whitepaper reviews the architecture of the Private Cloud Appliance, with a special focus on data elements in it that need to be backed up to protect them against outage. PCA creates internal backups of key data components to protect them against accidental erasure or corruption. Additional steps can be used to copy important data to external storage, providing insulation against catastrophic failures.

The first source for further reading is the PCA product page http://www.oracle.com/technetwork/server-storage/private-cloud-appliance/overview/index.html This has links to the official PCA documentation and PCA whitepapers, including the guide to expanding PCA with ZFS storage at http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/expand-opca-using-ozfssa-2613723.pdf

The reader is encouraged to review the Oracle VM documentation and whitepapers as most of their content applies to the Private Cloud Appliance. The page http://www.oracle.com/technetwork/server-storage/vm/overview/index-160875.html links to these papers, such as http://www.oracle.com/technetwork/server-storage/vm/ovm3-disaster-recovery-1872591.pdf The reader is also encouraged to review MOS note “Oracle VM 3: Getting Started with Disaster Recovery (Doc ID 1959182.1)”..

ORACLE

CONNECT WITH US

blogs.oracle.com/oracle

facebook.com/oracle

twitter.com/oracle

oracle.com

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

Hardware and Software, Engineered to Work Together

Oracle Private Cloud Appliance Backup Guide
March 2017

Author: Jeff Savit
Contributing Authors: Premal Savla, Greg King, Satinder Nijjar]

Oracle is committed to developing practices and products that help protect the environment