ORACLE®
OpenStack

Guide to Database Cloud using Oracle OpenStack

Automated Deployment of Oracle Database 12c and Oracle RAC in
Oracle OpenStack Release 4

ORACLE®

Table of Contents

Introduction

This guide answers important questions regarding automated deployment of Oracle Database and Oracle Real Application Clusters (RAC) that often arise when planning, designing, deploying and maintaining a private database cloud in OpenStack. It will help you understand the advantages of this service-oriented approach and address specific issues associated with this cost-effective IT strategy.

Questions

Q1 – What is automated deployment of Oracle Database in OpenStack?

Automated deployment of Oracle Database is the ability for users to request database environments by choosing from a predefined service catalog using a self-service provisioning framework. The key benefits are agility and faster deployment of database clouds for faster time to market of your applications. As databases are provisioned and de-provisioned, the associated computing, networking and storage resources in OpenStack are consumed and then released. Database resources can be consumed for the duration of a project, and then be automatically de-provisioned and returned to the resource pool. Computing costs can be tracked and charged back to the consumer.

Q2 – What drives organizations to automated deployment of Oracle Database and Oracle RAC in OpenStack? What are some popular use cases?

Organizations are drawn to implementing database clouds in OpenStack because they can quickly deliver database functionality to many users from the same OpenStack hardware and software infrastructure, while automating manual provisioning processes. Provisioning new databases in traditional environments can take days or weeks. Automated deployment of Oracle Database can compress that time down to a couple of hours, which can enable greater agility for the business, with less risk and lower costs, as shown in here:

AGILITY ↑	COST ↓	RISK ↓
SELF-SERVICE	SHARED RESOURCES	HIGHER AVAILABILITY
SERVICE CATALOG	METERING	TIGHTER SECURITY
ELASTICITY	AUTOMATION	GREATER CONTROL

Database clouds are quickly gaining traction with Dev/Test organizations since they permit developers to set up and tear down databases on an as-needed basis. They can provision new databases on demand without IT assistance. If they are using a storage environment that has snapshot storage, such as Oracle ZFS Storage Appliance, users can clone an existing pluggable database quickly.

Q3 – How do traditional database provisioning processes differ from automated deployment of Oracle Database and Oracle RAC in OpenStack?

In a traditional database environment, different administrators may approach common configuration and provisioning tasks in different ways, so there may be little consistency or standardization in IT processes.

Whereas, Oracle OpenStack uses the Murano application catalog service to invoke standard Oracle Database templates. Creating database configurations based on standard templates, that have been extensively tested and hardened, reduces the risk introduced by manual configurations. IT defines different "tiers" or classes of database service based on the standard templates. An online catalog defines varying sizes and service levels for different applications. Authorized users can select the databases they need from this catalog, consume those services for the duration of a project, and then automatically de-provision the databases and return them to the resource pool. Once the database cloud environment is established, users can perform simple provisioning activities without engaging a DBA to manually configure, provision and tear down databases in response to each individual request.

Q4 – How does automated deployment of Oracle Database save time or simplify activities for DBAs?

When a user, department, or line of business needs a database in a traditional database environment, they must make a formal request to the IT department, wait for that request to be approved, and then wait for IT to install the necessary software and provision the database resources. This cycle involves management approvals and a series of discrete tasks for DBAs and system administrators.

Automated deployment of Oracle Database and Oracle RAC enables a self-service model that relieves IT professionals and DBAs from having to manually configure, provision and tear down databases in response to each individual request.

Automated deployment of Oracle Database in OpenStack represents a paradigm shift in IT, allowing DBAs and administrators to move away from routine tasks of manual provisioning to more valuable work.

Q5 – How do you sell the concept of automated deployment of Oracle Database in OpenStack to the IT department and management?

Currently, many OpenStack deployments are for Dev/Ops, IT services, Test, QA and web services. Enterprises are looking to support their corporate business applications with OpenStack, to increase their ROI. Oracle Database is used by the majority of enterprise business applications. Automated deployment of Oracle Database simplifies deployment of these enterprise applications, provides operational efficiencies and helps accelerate business innovation.

IT is continually asked to do more with less. IT leaders are looking for ways to increase agility while reducing cost and risk. Charging back for database services can transform the IT environment from a cost center to a revenue generator. Not only does the database cloud enable more agile database services, it also helps the business to allocate costs more accurately. User organizations pay to the degree that they consume each service, incenting

those organizations to be more efficient. They can commission databases for particular purposes, then decommission them and return them back to a resource pool that all users share.

Q6 – What are some of the primary planning considerations for implementing Oracle Database in OpenStack?

There are multiple phases to adoption – from discovery and evaluation, through running pilots, test and development, and initial implementation to broad adoption. When planning and designing automated deployment of Oracle Database in OpenStack, focus on:

1. Thorough planning to help ensure that you can meet user needs. Ask yourself: what do users need for their applications? What are their workloads? Not all database workloads are created equal, of course. Some are more mission critical than others. Some are transient workloads like development and test for example, as well as backup and recovery. What are the performance requirements of the applications? Development teams, application teams, business unit owners and IT units should agree on the DB service catalog for the targeted workloads. Future technology requirements should also be considered. Once you have database cloud capabilities and people can request databases, the demand may be higher than you initially expect.

2. Selecting the right technology. OpenStack has its own release schedule and individual services within these releases can also undergo several updates. Oracle Database 12c and Oracle RAC each have their own release schedules as well. The architecture and technology used for automated deployment of Oracle Database needs to be able to decouple the release schedules of OpenStack and release schedules of Oracle Database so the user can take full advantage of the new capabilities that each release offers.

One existing OpenStack project provides database cloud capability. However, it is tightly integrated with OpenStack services such as Nova, Cinder and Swift and therefore any updates to these OpenStack services with a new OpenStack release impacts the database services. This can severely limit its usefulness in the enterprise applications that need to take advantage of the newer features and patches in Oracle Database and Oracle RAC.

Oracle OpenStack uses OpenStack Murano service for automated deployment of Oracle Database and Oracle RAC. Murano provides an application catalog service that enables end users to easily deploy applications. It provides an abstraction layer, as shown in the figure below, on top of the OpenStack Orchestration Service (Heat), decoupling Oracle Database and OpenStack release schedules.
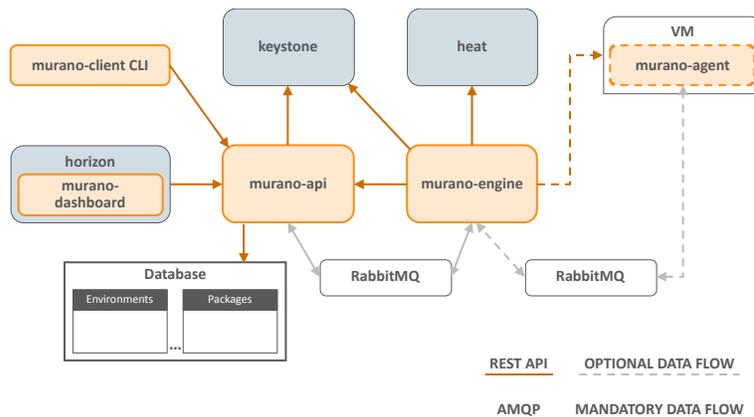
Q7 – How can Enterprise Architecture (EA) help with automated database deployments?

For successful adoption and sustainable execution, it is important to understand how database clouds fit into the overall environment. You can achieve the greatest value when the database cloud strategy is considered along with your organization's overall enterprise architecture and IT strategy. EA helps structure database cloud projects through a series of overlapping steps, or iterations. A systematic process determines how these environments are created, with attention to the business, technical, and operational implications.

For example, if the HR department wants to implement a new payroll service and share it with other departments, the EA team begins by creating use-cases that reflect the needs of each department. They consider the profiles of the various departments and develop a service catalog that lists the functions needed. Then, each department can provision as little or as much of each service as is needed and also specify the database size, degree of availability, disaster recovery, and uptime required.

Q8 – How do you create service catalogs and establish tiers of service?
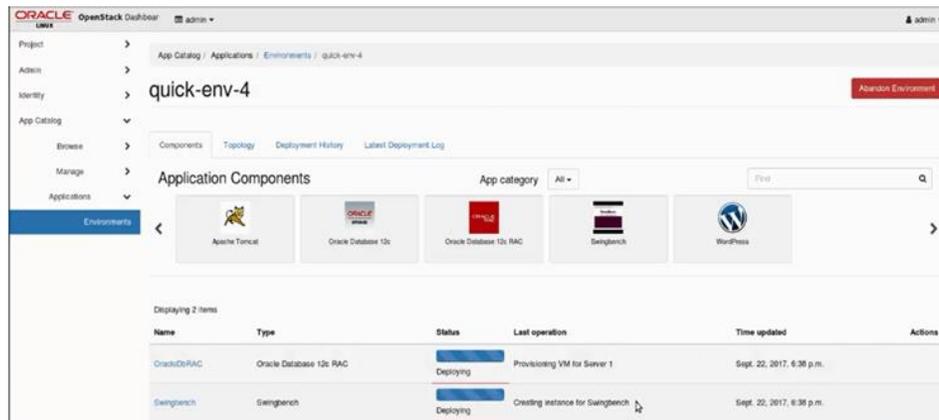
The Murano application catalog in Oracle OpenStack lets you specify how to set up instances of databases and applications and then configure the environment according to the business need. For example, a bronze tier might include basic database services using pluggable databases within Oracle Database 12c. A silver tier might also include Oracle Real Application Clusters for instantaneous failover to another server in the case of an outage, and a platinum tier could use Oracle Maximum Availability Architecture to enable automatic failover among clusters in two data centers. Other parameters that may come into play when defining service tiers include geography, disaster recovery levels, and security. Most organizations create simple service catalogs with a few basic tiers, as shown in the figure below:

Sample Business Catalog

		Bronze	Silver	Gold	Platinum
Local Outage	RTO	Best effort	24 hours	8 hours	1 hour
	RPO	0	0	0	0
DR	RTO	Best effort	2 weeks	1 week	3 days
	RPO	80 hours	40 hours	24 hours	1 hour
Reserved capacity for DR		N/A	N/A	50%	100%
Cost Model		Service Tier x Footprint = monthly charge			

DB Size	CPU	RAM GB	Storage GB	Adders
S	2	12	200	Extra RAM, Storage
M	4	24	400	
L	8	36	800	

Q9 – Can Oracle Data Guard and Oracle GoldenGate be used after launching Oracle Database 12c and Oracle RAC?

Today, the above services cannot be automatically launched from Oracle OpenStack Release 4 (R4). However, you can manually associate them with a newly launched Oracle Database or Oracle RAC instance from R4. With Oracle Linux you can install and launch Oracle Data Guard or Oracle GoldenGate from scratch and integrate them with the Oracle Database instance that was provisioned via Murano.

Q10 – How do you track costs and charge them back to consumers of database services?

Customers can choose from several third-party billing and chargeback software solutions offered for OpenStack. Oracle OpenStack is certified as "OpenStack Powered" and is compatible with a large number of applications in the OpenStack ecosystem that support OpenStack standard APIs.

Q11 - What is the supported configuration for automated deployment of Oracle Database and Oracle RAC?

The following configuration is supported for automated deployment of Oracle Database and Oracle RAC in Oracle OpenStack Release 4:

» Single instance Oracle Database 12c or Oracle RAC
» Hypervisor: Oracle VM Server via OpenStack Nova libvirt-xen
» Network: Several options via OpenStack Neutron ML2 plugin
» Storage: Oracle ZFS Storage Appliance via OpenStack Cinder

Oracle OpenStack also includes third-party plugs-ins (networking, storage) that are made available by the third-parties in the upstream OpenStack release. Contact Oracle when a specific partner plug-in is required in your configuration.

Q12 - How do you set up automated deployment of Oracle Database 12c and Oracle RAC in OpenStack?

Oracle OpenStack Release 4 is the only OpenStack distribution available today with support for the automated deployment of Oracle Database 12c and Oracle RAC. Refer to the documentation listed below.

Release notes, Installation and Deployment Guide, Configuration Guide and Application Deployment Guide

The Application Deployment Guide for Release 4 provides step-by-step instructions for setting up automated deployment of Oracle Database 12c and Oracle RAC. Areas covered include:

» Chapter 1: Creating and loading Murano-Enabled Images for Oracle Linux
» Chapter 2: Setting up an Environment for Oracle Applications
» Chapter 3: Preparing images, creating and deploying Oracle Database 12c Application
» Chapter 4: Creating and deploying the Oracle Real Application Clusters 12c Application

Conclusion

In this guide we have answered common questions related to planning, designing, and deploying a private database cloud in the OpenStack environment. Many customers that adopt database cloud report lower overall costs, greater autonomy for business users, and reduced risk within IT planning scenarios. Other important benefits include greater predictability and agility, which enables the business to innovate and change. Finally, IT professionals and application developers can experience new levels of creativity and innovation.

ORACLE®

CONNECT WITH US

blogs.oracle.com/openstack

facebook.com/OracleLinux

twitter.com/OracleLinux

oracle.com/openstack

Integrated Cloud Applications & Platform Services

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