Deploying Database-as-a-Service with Pure Storage and VMware

vmware®   **PURE**STORAGE   **VLSS**

Table of Contents

Introduction

About This Test Plan

The purpose of this document is to describe the test plan for deploying a Database-as-a-Service (DBaaS) environment using VMware® vRealize™ Automation and a Pure Storage FlashArray. This reference architecture will provide a baseline to consolidate Oracle DB instances, reduce downtime, increase infrastructure resource utilization, and simplify management of the entire stack. The outcome of the testing will be documented in the reference architecture/deployment guide.

Business Problem

Deploying and managing an Oracle DB environment is not a trivial undertaking. Oracle infrastructure tends to have stringent performance, business continuity, and backup and recovery requirements. To support these requirements, IT departments often are tasked with maintaining sprawling infrastructure for production, disaster recovery, backup, quality assurance, test, training, development, and sandbox.

VMware and Pure Storage are actively working on developing solutions that will greatly simplify the deployment and management of a robust Oracle environment, while delivering higher levels of efficiency and flexibility for IT departments throughout the application lifecycle—implementation, migrations, consolidations, upgrades, and ongoing maintenance.

Enterprise customers have identified Database-as-a-Service as a key initiative that will enable higher levels of agility via rapid application development. The idea is to allow database consumers such as application developers, testers, and architects to provision databases easily using an on-demand, self-service platform. However, performance—and in particular storage performance—is the key issue that can hold back a DBaaS project. Traditional disk-based storage cannot meet the performance needs for DBaaS and is inherently too complicated to manage. This whitepaper will outline reference architectures for deploying high-performing, easy-to-manage DBaaS environments using VMware vRealize Automation on all flash storage.

Database-as-a-Service Definition

In the context of this reference architecture, DBaaS is defined as:

The ability to manage, monitor, and provision database resources on demand through self-service catalogs with minimal requirements for installation and configuration of hardware or software up front.

The idea is to allow database consumers such as application developers, testers, and architects and owners to provision databases easily and quickly using an on-demand, self-service platform.

Benefits of DBaaS

This section summarizes the benefits of deploying a DBaaS solution.

Reduce Application Time to Market
DBaaS gives you the ability to rapidly develop, test, and deploy new applications at a fraction of the time it used to take with legacy systems. Now, you can take new applications to market much faster and gain an edge over your competitors, and service your customers better.

Automate Database Lifecycle Management
Database sprawl is a major problem in many companies. IT departments have difficulty keeping track of all the DB instances that are deployed. This has serious implications from licensing, data governance, and security standpoints. With DBaaS you can ensure that you have complete visibility into your database environment and enforce governance through policy-based management. Developers and DBAs will no longer need to go to external public cloud services for their DB/application infrastructure needs.

Increase Operational Efficiency and Resource Utilization
DBaaS increases the efficiency of both DB/application and infrastructure teams. The storage administrator will no longer need to spend time and effort on provisioning storage resources for databases. DBAs will no longer have to wait for DB resources to be allocated and provisioned. This frees up time for both teams to focus on more strategic, higher-value activities.

Importance of Storage Infrastructure for DBaaS

Storage performance is often the limiting factor hindering DBaaS projects. Traditional disk-based storage cannot meet the performance needs for DBaaS and is inherently too complicated to manage. DBaaS requires

• High-performing storage infrastructure to ensure rapid provisioning of DBs

• Efficient data reduction to ensure storage capacity is not exhausted

• A high degree of resiliency to ensure minimal-to-no downtime

• Ease of use and management

Pure Storage's FlashArray meets all of these requirements and allows you to deploy a DBaaS environment in the most cost-effective manner and avoid expensive, complex systems like database appliances.

High-Level View of Environment

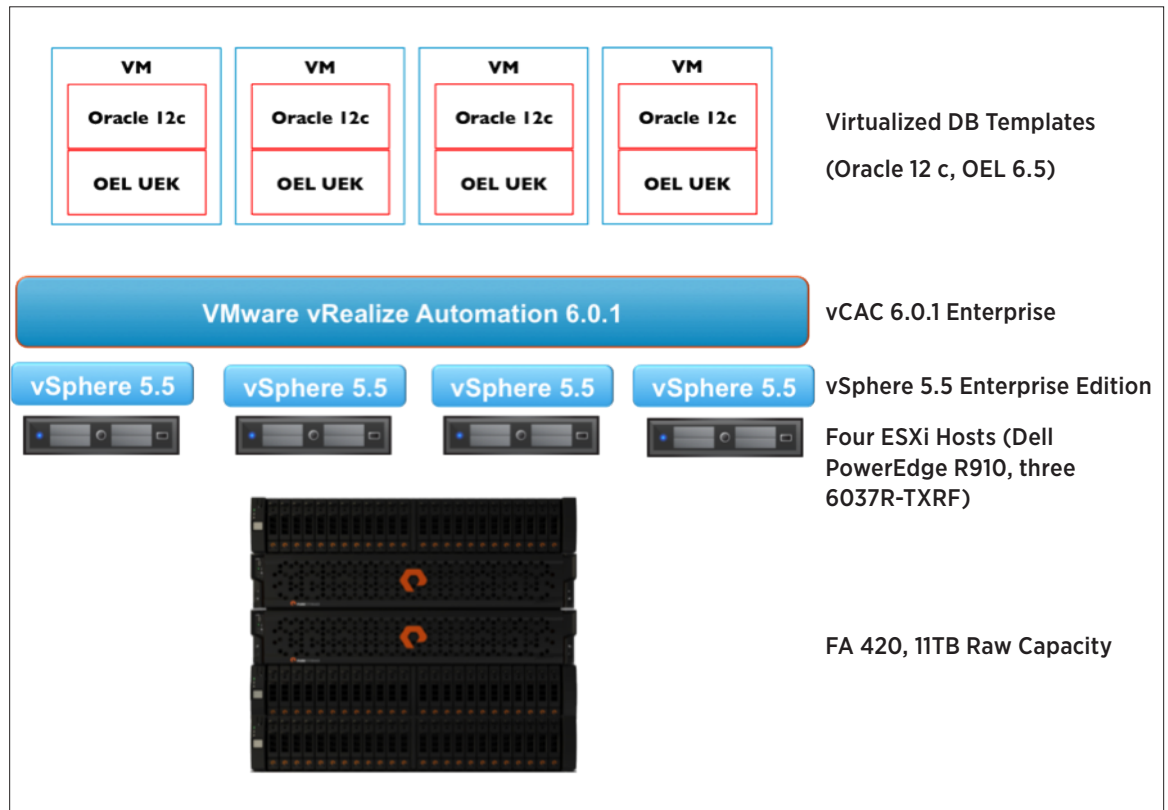Figure 1 provides a high-level overview of the solution architecture.

Figure 1: DBaaS Environment

Server and Storage Sizing

We have performed many tests in our lab and received commensurate feedback from customers about volume performance. The data demonstrates that whether it's a large LUN or a small LUN, the performance tends to be identical. All LUNs are created equal in all aspects, performance included. You can use a few large 64TB volumes (the new LUN maximum in vSphere 5.5) and datastores in the case of VMware or create a number of smaller thin-provisioned LUNs—there is no performance penalty. Indeed, using a thin-provisioned virtual disk on an already thin-provisioned volume is perfectly valid. For more information on this topic refer to How much storage should we provision for your VM? Never worry about it again.

VMware vSphere Template

The crux of this DBaaS is using a dedicated virtual machine template with a gold image of an Oracle 12c database already configured on it. We started with Oracle Enterprise Linux 6.5 as the OS image (`OEL65`). From there we followed the Oracle Databases on VMware Best Practices Guide to configure Oracle.

Note: This template is set up for a Test/Dev scenario, and did not follow all standard VMware storage recommendations. Because we were not greatly concerned with database performance, we created a simplified storage layout using VMFS instead of the recommended ASM on VMFS.

In the end, `OEL65` template had an Oracle 12.1.0.1 database running on Oracle Enterprise Linux 6.5 with a single listener configured on port 1521. The template uses 1 vCPU, 2048MB vRAM with two 25GB SCSI virtual disks. One disk was presented as /, which included the OS system files as well as the Oracle software install under `/u01` conforming to OFA standards. The second disk was mounted as `/u02` and was used to store 20GB of user data imported into the database.

From that first template, we copied a second template, `Oracle-1TB`, and then added a 1TB SCSI virtual disk. That disk was presented as `/u03` and then used to import 800GB of user data. The template was also modified to use 8 vCPU and 32GB of vRAM, of which 16GB was allocated to the Oracle DB.

VMware vRealize Automation (vRA)

A VMware vRealize Automation 6.0.1 environment was configured against the VMware vCenter™ environment using VMware installation documentation.

If you don't need to configure the entire environment and plan to use the existing tenants and business groups, follow these steps instead:

1. In vCenter, ensure your template is created. If not, convert your virtual machine into a template.

2. To log in to vRealize Automation, navigate to **Infrastructure > Compute Resources > Compute Resources**.

3. You'll see vRealize Automation listed under Compute Resources. Hover over vRealize Automation and a drop-down will display. Select **Data Collection**.

4. On the next screen, locate Inventory, and click **Request Now**. Wait for the inventory to complete updating.

5. Navigate to **Infrastructure > Blueprints > Blueprints**. The following window will appear.
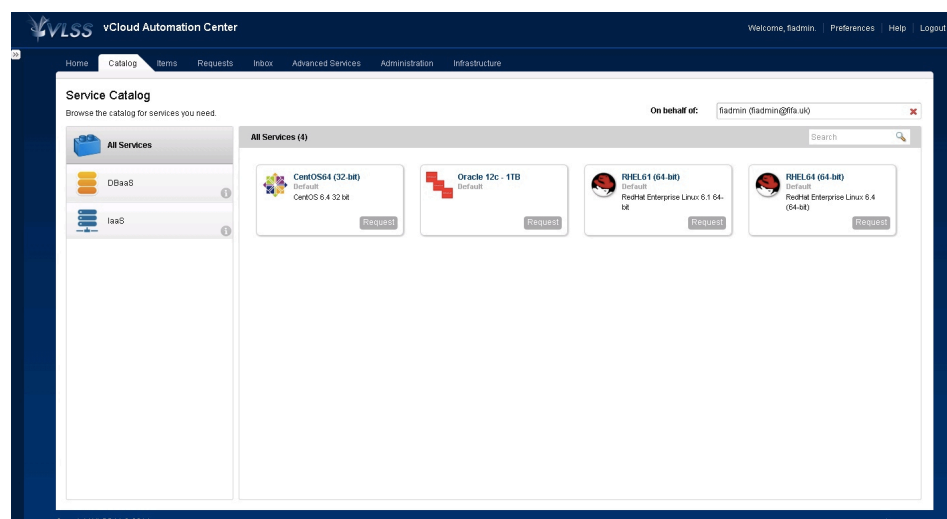

Figure 2: vCAC Service Catalog

6. To create a new blueprint, click **New Blueprint** and select **Virtual > vSphere (vCenter)**.

7. Enter the Name of the blueprint (for instance Oracle 12c – 50GB) and select a number of Archive (days).
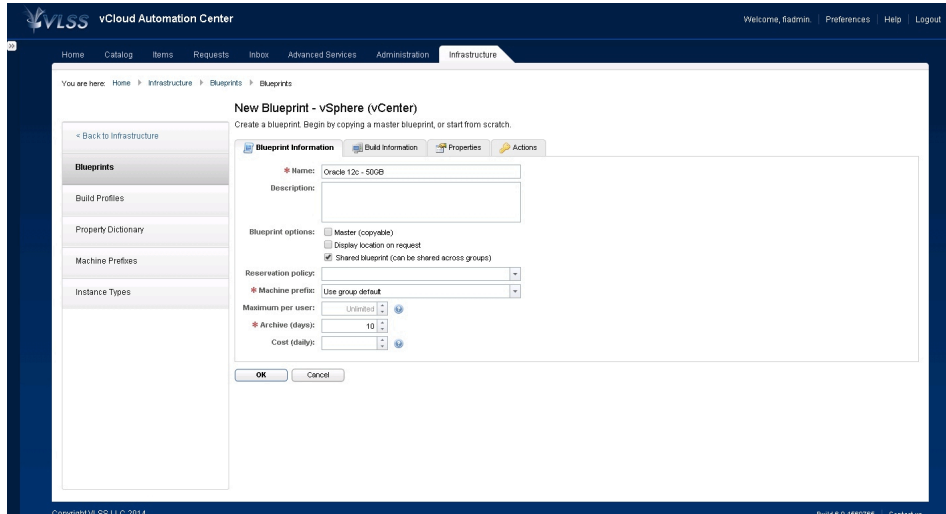

Figure 3: Create New Blueprint from vSphere / vCenter Template

8. Click **Build Information**, change the Action field to **Clone**, and click the ellipsis button to select the machine to clone from.
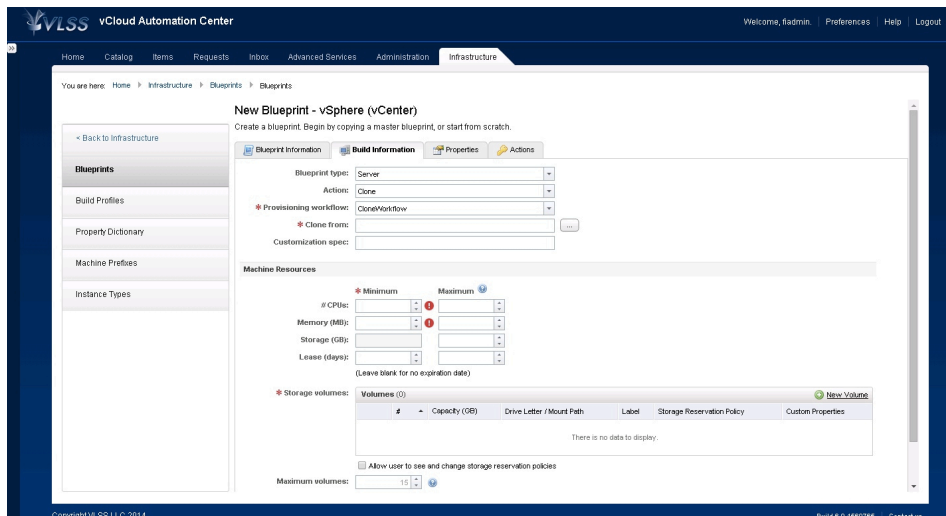

Figure 4: New Blueprint — Set Provisioning Workflow

9. Choose the virtual machine template that you want to clone and click **OK**.
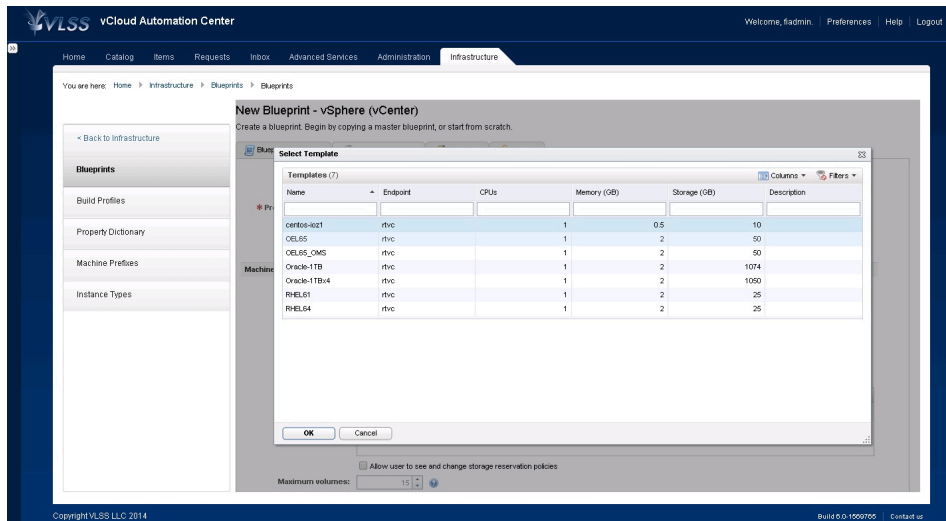

Figure 5: New Blueprint — Select vSphere / vCenter Template

10. The rest of the fields should populate automatically. Click **OK** at the bottom of the screen.

11. You'll now see the `Oracle 12c – 50GB` new blueprint. Hover over it, and select **Publish** from the drop-down menu. Click **OK**.

12. Click **Administration** in the main menu.

13. Click **Catalog Management > Catalog Items**.

14. Click the drop-down next to `Oracle 12c – 50GB`, and click **Configure**.

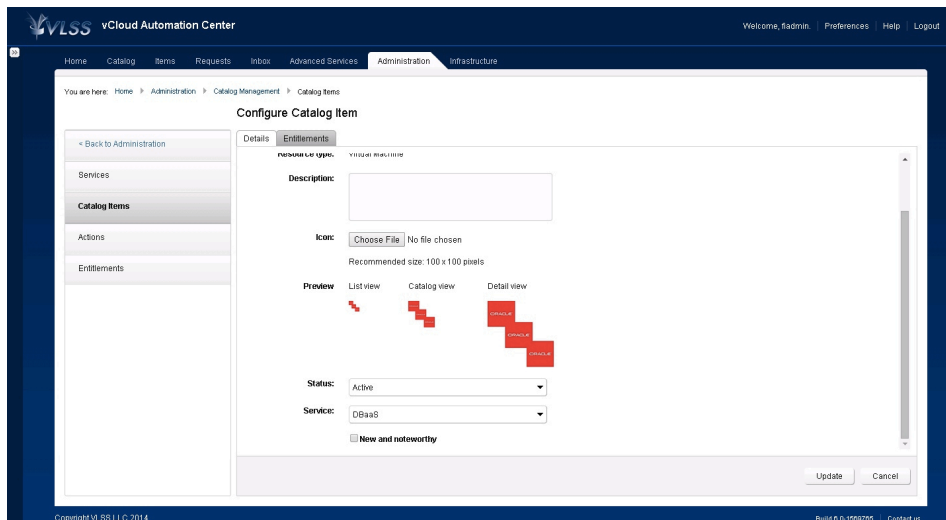15. In the Service field, select **DBaaS** and click **Update**.

Figure 6: Enable New Blueprint in Service Catalog

16. You can now navigate to **Catalog > DBaaS** to see the new blueprint ready for use.
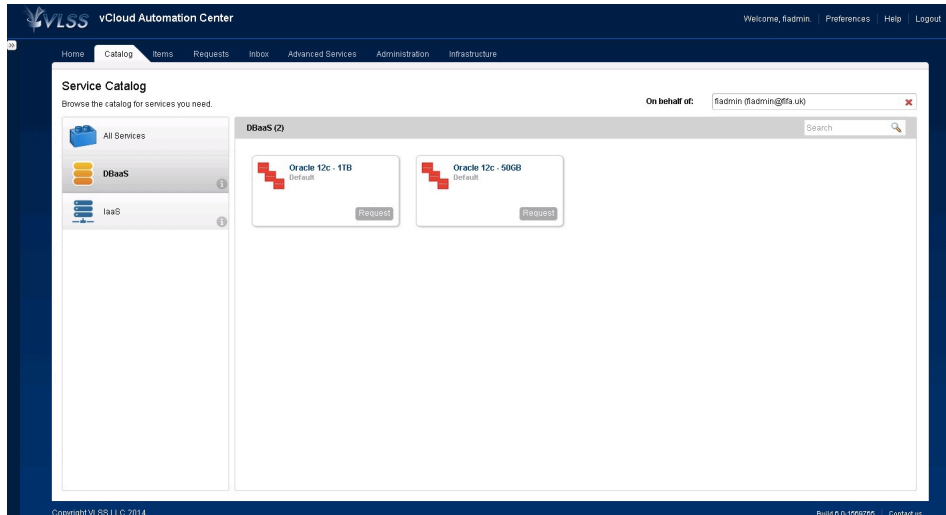

Figure 7: vCAC DBaaS Service Catalog

Database-as-a-Service (DBaaS)

At this point, we have a working environment for DBaaS. A user with the privileges to access either the `OEL65` or `Oracle-1TB` service in the vRealize Automation catalog can now provision as many databases as allowed.
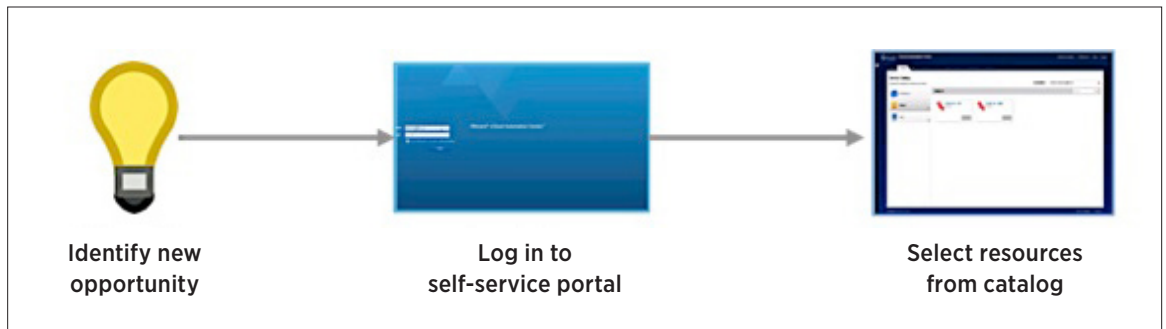

Figure 8: DBaaS Configuration Workflow

Test Scenarios

This section outlines test scenarios and results.

Test Time to Mass Provision Oracle 12c Databases

• 50 OEL65 Databases (1 vCPU, 2GB vRAM and 50GB storage)
   - **23 minutes** through vRealize Automation
• 1 Oracle-1TB Database (8 vCPU, 32GB vRAM, 1TB storage)
   - **Under 6 minutes** to complete through vRealize Automation

Notable Findings

• Tests found no impact on storage performance during provisioning.
• 7:1 data reduction on for Oracle 12c database workloads
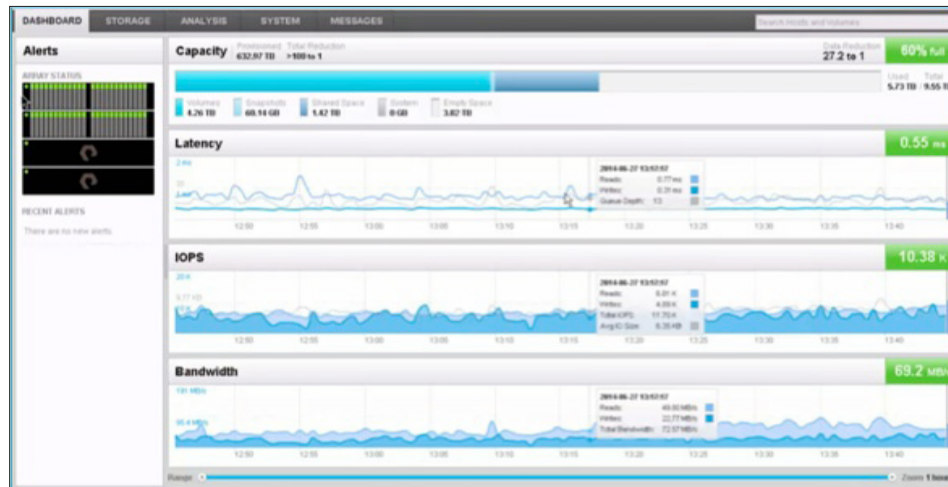• 27:1 overall data reduction on array also running numerous other workloads

Figure 9: Capacity Test Results

• Pure Storage uses metadata copies to satisfy VMware `XCOPY` requests, so there is very little resource requirement and no physical data movement being done during the provisioning.
• The FlashArray, as previously noted, does not actually copy the data described in a `XCOPY` transaction—it just moves or copies metadata pointers. Therefore, for the most part, the bottleneck of any given virtual machine operation that leverages `XCOPY` is not the act of moving the data (since no data is moved), but how quickly an ESXi host can send `XCOPY` SCSI commands to the array. Hence, copy duration depends on the number of commands sent (dictated by both the size of the virtual machine and the maximum transfer size) and correct multipathing configuration. If more data can be described in a given `XCOPY` command, fewer commands overall must be sent, and it therefore takes less time for the total operation to complete. For this reason, Pure Storage recommends setting the transfer size to the maximum value of 16MB. The following commands provide for retrieval of the current value, and for setting a new one.

```
esxcfg-advcfg -g /DataMover/MaxHWTransferSize
esxcfg-advcfg -s 16384 /DataMover/MaxHWTransferSize
```

Day 2 Performance

• Pushed 925K TPM against provisioned DB OEL65

• Pushed 150K IOPS using HammerDB against 5 OEL65 databases

Day 2 Operations

In order to make the scenario more realistic, we decided to script in some enhancements to the DBaaS that can be leveraged using vCenter Orchestrator. Specifically, we wanted to rename the cloned database from VLSSPOC to align with the hostname. Additionally, we wanted to install an Oracle 12c agent and register the agent, host, listener, and database with the Oracle 12c EM OMS.

To do this, we created another vCenter template, `OEL65_OMS`, containing a shell script `oms.sh`. It requires two parameters: the name of the OMS host and the name of the provisioned virtual machine (for example, `/bin/sh /oms.sh oms.fifa.uk FIFA0389`).

Once the script was developed, we used an existing vCenter Orchestrator workflow called `Run Script in Guest_v2` to execute the script on provisioning. Creating that workflow is beyond the scope of this paper. However, once the workflow is created, it is easy to have a provision request from vRealize Automation call a vCenter Orchestrator workflow:

1. Create a new blueprint as described above, publish it, and associate it with the DBaaS service.

2. Under **Infrastructure > Blueprints**, select the new blueprint **Oracle 12c – 50GB + OEM** and click **Edit**.

3. Switch to the Properties tab and add properties to match the vCenter Orchestrator workflow. In this case, there are three set properties based on the template:

```
my.command = /bin/sh
my.script = /oms.sh
my.custom.field = oms.fifa.uk
```
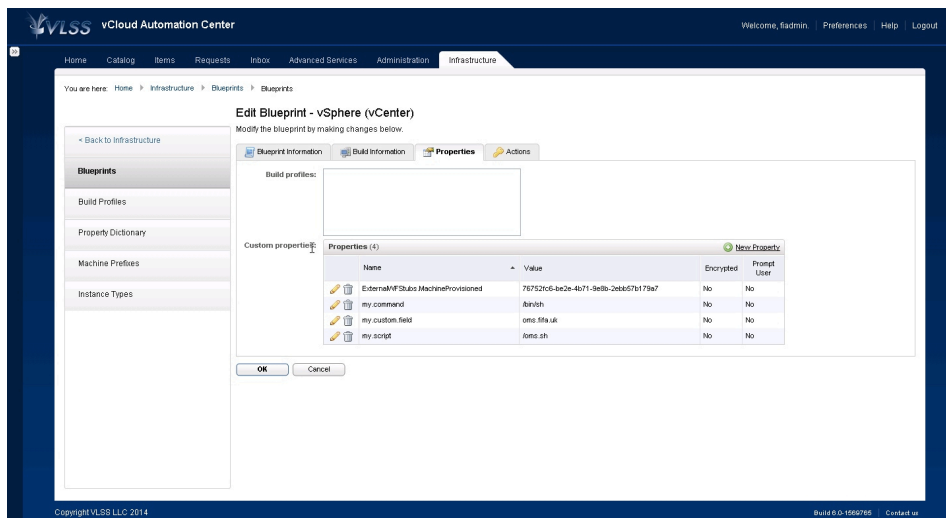

Figure 10: Add Custom Properties to vCAC Blueprint

4. Click **OK** to save.

5. Open the vCenter Orchestrator client, and run the workflow **Library > vCloud Automation Center >
Extensibility > Assign** to assign a state change workflow to a blueprint and its virtual machines.
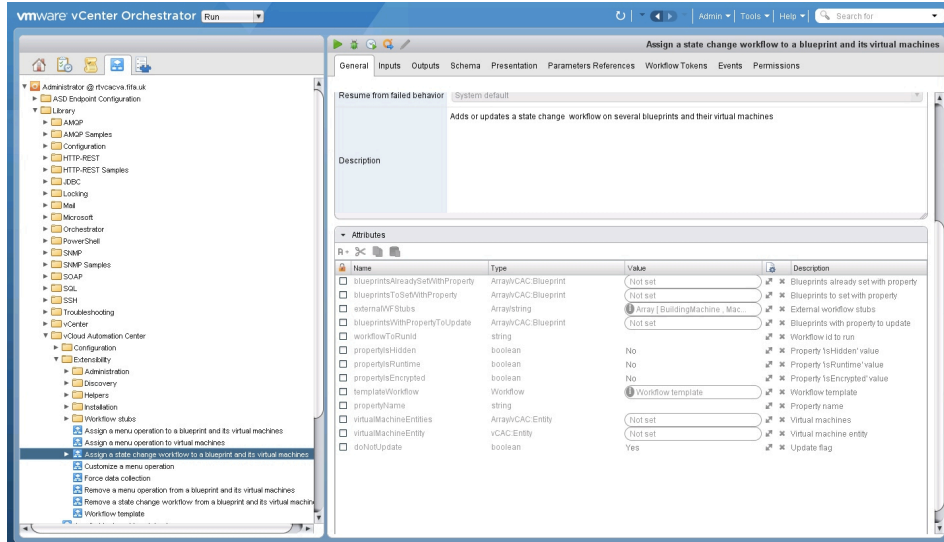

Figure 11: Adding vCO Workflow to vCAC Blueprint

6. Change the workflow stub to **MachineProvisioned** and select the correct vRealize Automation
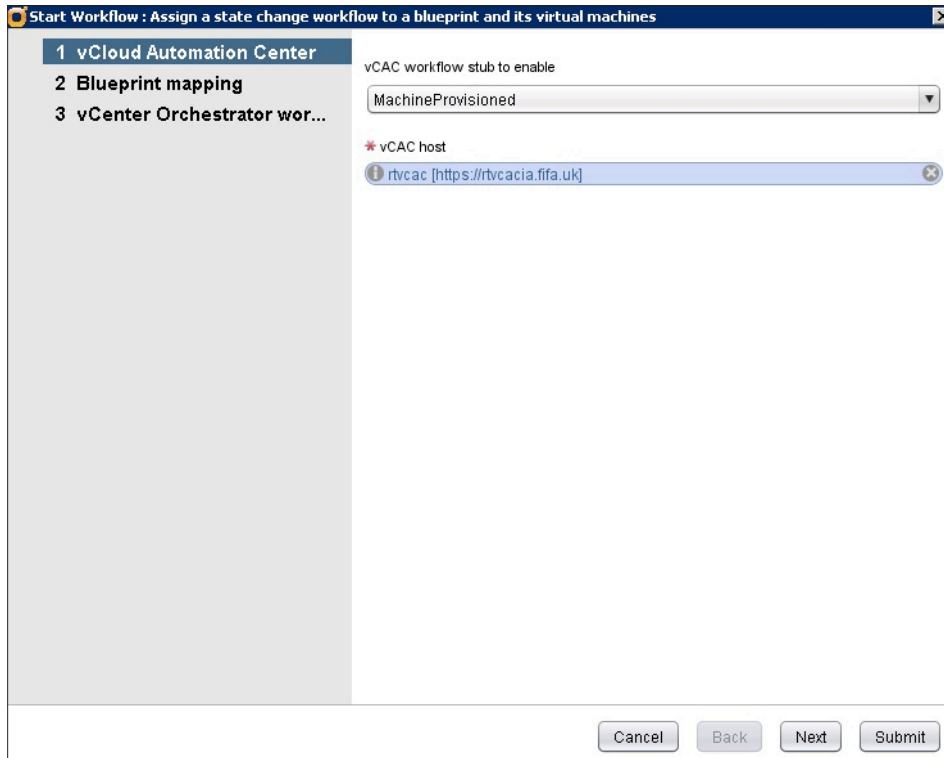host before clicking **Next**.

Figure 12: Add vCO Workflow to Machine Provisioning Event

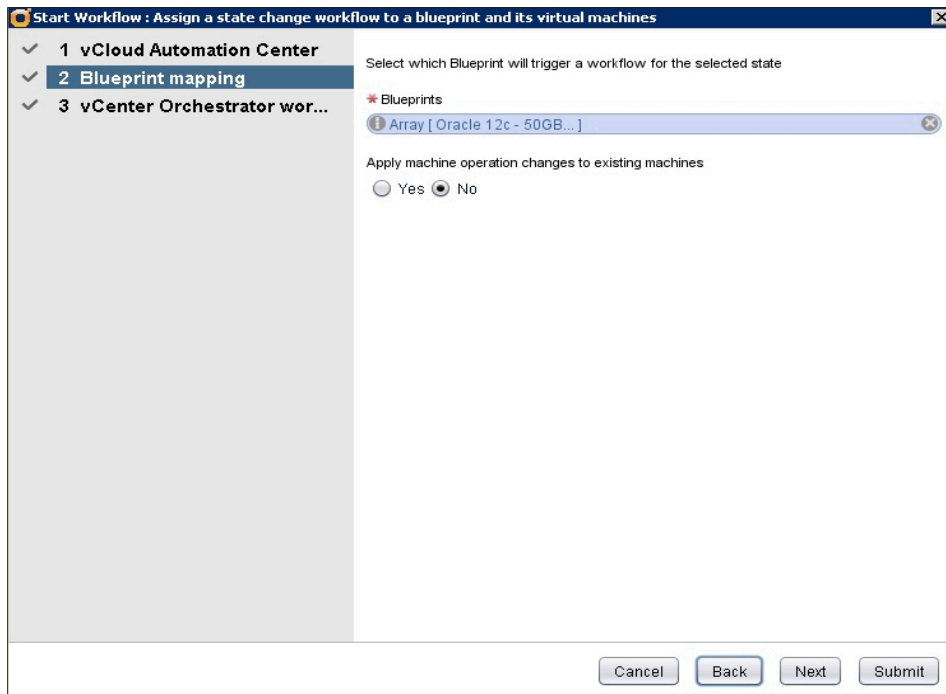7. Select the correct blueprint from the clickable lists and click **Next**.

Figure 13: Select vCAC Blueprint to Add vCO Workflow

8. Select the correct workflow from the clickable lists, and click **Submit**.
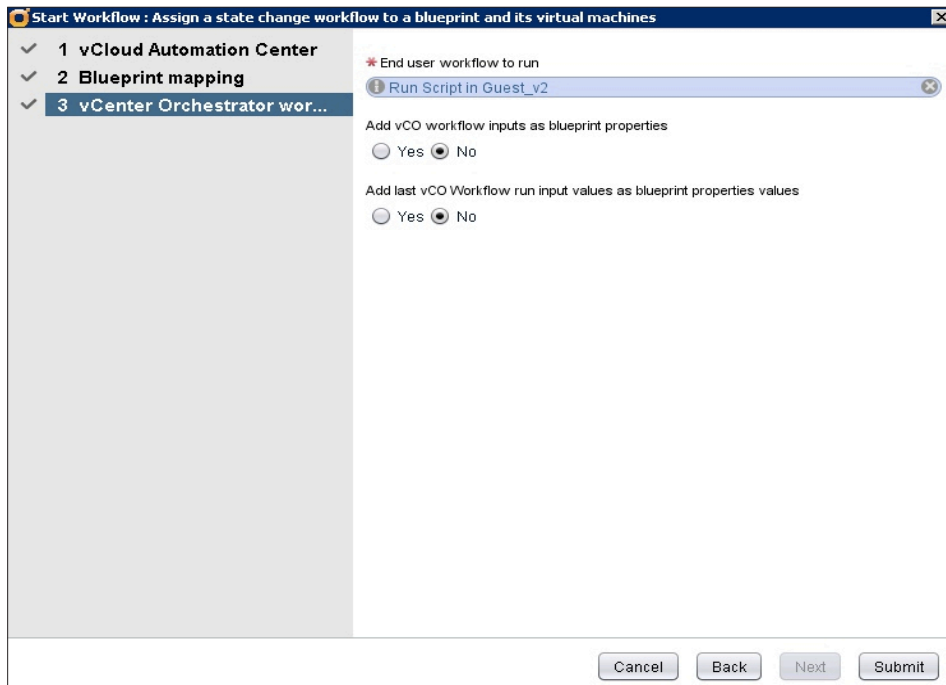

Figure 14: Select vCO Workflow

9. After the job has successfully completed, check the Blueprint properties in vRealize Automation. There should be a new property called `ExternalWFStubs.MachineProvisioned`.

10. Provision a new DBaaS from the self-service portal using that blueprint. Within a few minutes, the newly created database should appear in Oracle 12c EM.
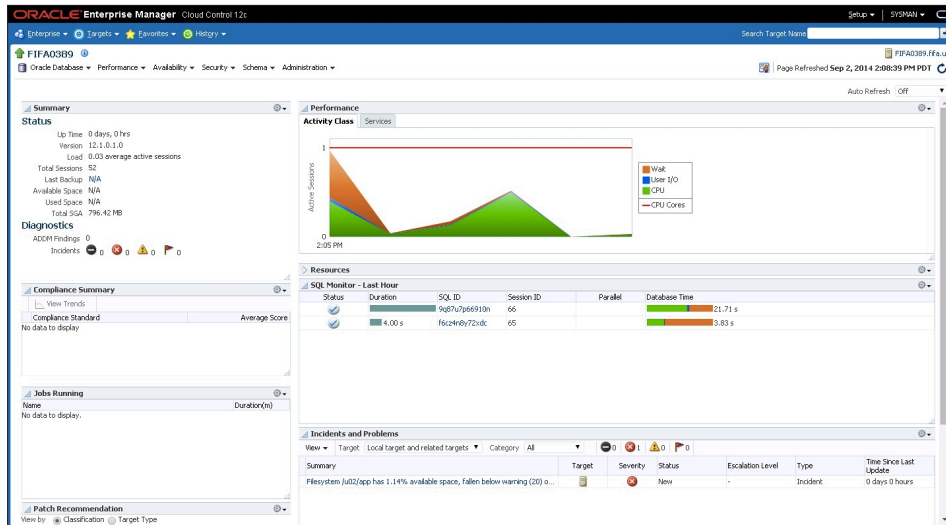

Figure 15: Oracle EM Summary for Newly Provisioned DBaaS Database

About the Authors

The following authors contributed to this paper:

• Amjad Afanah, Senior Product Manager, VMware

• Rajeshkumar Easwaramoorthy, Senior Engineer, VMware

• Dean Bolton, Chief Architect, VLSS

• Avinash Nayak, Product Marketing, Pure Storage

Appendix

Hardware and Software Details

This section provides specific details of the test deployment hardware and software.

Pure Storage FlashArray FA-420
The deployment included Pure Storage FlashArray FA-420, configured as detailed in Table 1.

COMPONENT	DESCRIPTION
Controllers	Two active/active controllers providing highly redundant SAS connectivity (24Gb) to two shelves, interconnected for HA via two redundant InfiniBand connections (56Gb).
Shelves	Two flash memory shelves with 22 SSD drives, 22 x 256GB for a total raw capacity of 11TB (10.3TiB).
External Connectivity	Four 8Gb Fibre Channel ports per controller, total of eight ports for two controllers.
Management Ports	Two redundant 1Gb Ethernet management ports per controller. Three management IP addresses required to configure the array, one for each controller management port and a third one for virtual port IP address for seamless management access.
Power	Dual power supply rated at 450W per controller and 200W per storage shelf or approximately 9Amps of power.
Space	Entire FA-420 system hosted on an eight rack unit (8 RU) space (2 RU for each controller and 2 RU for each shelf).

Table 1: Pure Storage FlashArray FA-420

Server Configuration
The Dell PowerEdge R910 used in this deployment was configured as detailed in Table 2.

COMPONENT	DESCRIPTION
Processor	24 cores Intel® Xeon® E5-2697 v2 @ 2.70GHz
Memory	384 GB

Table 2: Server Configuration

Software Configuration

Table 3 details the software configuration for the test deployment.

COMPONENT	DESCRIPTION
VMware vSphere	vSphere 5.5 Enterprise Plus
VMware vRealize Automation	vRealize Automation 6.0.1 Enterprise Edition
Oracle Database	Oracle 12c Enterprise Edition (plus Oracle Enterprise Manager 12c)
Operating System	Oracle Enterprise Linux 6.5

Table 3: Software Configuration

Best Practice Tuning for Pure Storage

Pure Storage FlashArray is certified on vSphere 5.x platform (see High-Level View of Environment) and supports the following VAAI primitives: `XCOPY`, `WRITE SAME`, Atomic Test & Set and `UNMAP`.

Before beginning your project, make sure you adhere to Pure Storage best practices. Refer to the Oracle Databases on VMware Best Practices Guide to get the best out of your infrastructure.

Boot from SAN Recommendations

Pure Storage FlashArray storage provisioning works at a host-cluster-group level or at host level. LUNs provisioned to hosts at host level (LUNs numbering from 1 through 10) are dubbed private volumes that can be used for boot from SAN LUNs. The host cluster group has a direct co-relation with ESXi clusters and hosts. The best practice is to zone the ESXi hosts and Pure FlashArray and create a host cluster corresponding to ESXi clustered hosts as this ensures common shared storage across the cluster.

VMware SATP and PSP

Pure FlashArray is an ALUA-compliant array with all paths optimized for doing active IOs (Active/Optimized). When a Pure FlashArrray is plugged into the ESX farm, the array gets claimed as ALUA array with the `VMW_SATP_ALUA` as the Storage Array type. However, the default path selection policy is `Fixed` path. This needs to be changed to `Round Robin` in the storage manage path option in vCenter. This can also be accomplished using the following command line on a per device basis:

```
esxcli storage nmp device set -d naa.6006016055711d00cff95e65664ee011
--psp=VMW_PSP_RR
```

Instead of doing this every time, it is convenient to set the default PSP for VMW_SATP_ALUA as `VMW_PSP_RR` (Round Robin) by running this command on the ESX host or through a vMA (VMware Management Assistant):

```
esxcli storage nmp satp rule add -s "VMW_SATP_ALUA" -V "PURE" -M "FlashArray" -P
"VMW_PSP_RR" -O "iops=1"
```

This command will create any new PURE LUNs with Round Robin policy. Please note this command assumes that there are no other ALUA arrays that are connected to the host as this will affect all the array connectivity.

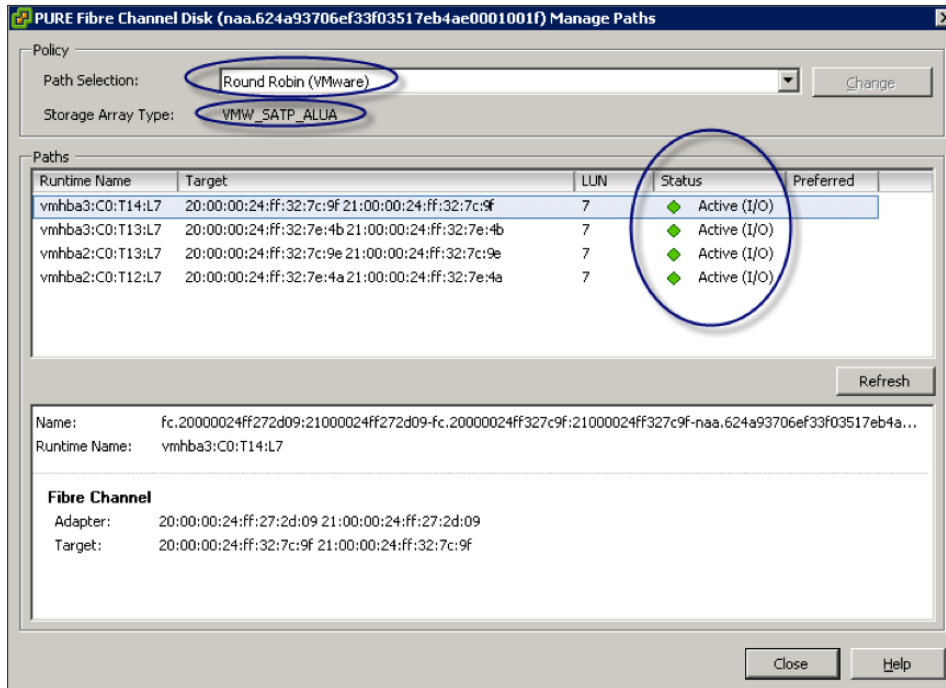Figure 16 below shows the correct settings of the PURE LUN in vCenter.

Figure 16: Correct Settings for Pure Storage FlashArray LUN

ESXi Tuning

We suggest changing the Disk parameter `SchedNumReqOutstanding` (DSRO) to a larger value than the default value of 32 and disk `schedQantum` to a larger value as shown.

After changing the DSRO, make sure you tune the underlying FC HBA driver queue depth with the same value (256 in our case).
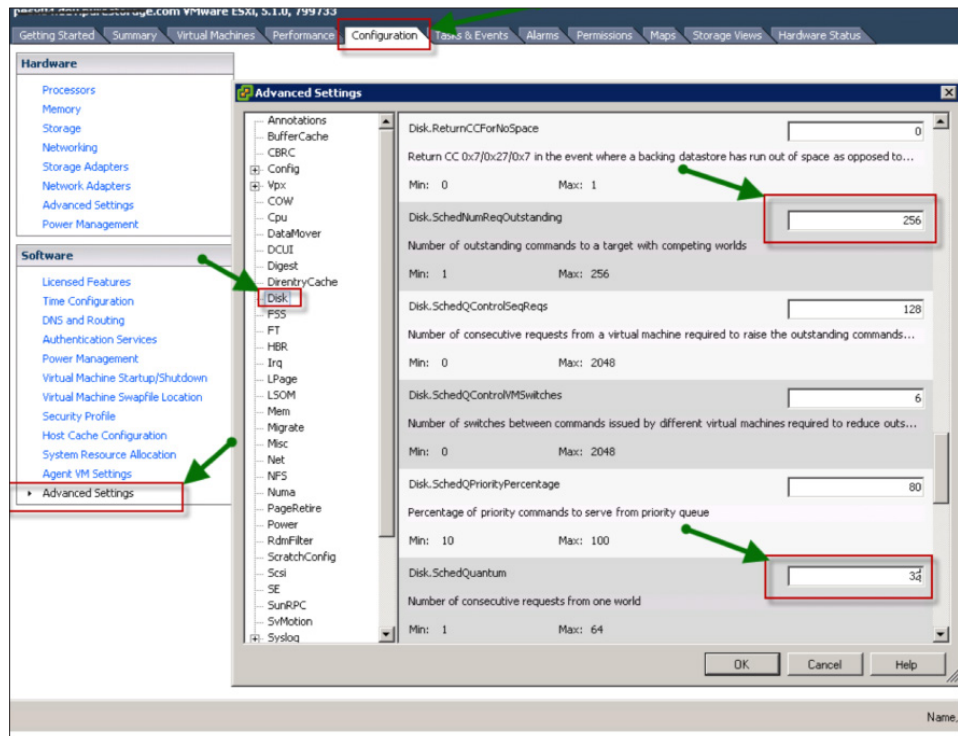
Figure 17: ESXi Tuning

Guest-Level Tuning

There are no special configurations required for guests to use Pure Storage FlashArray in customer environments. All the best practices that customers have currently deployed are equally applicable. Best practice are summarized below:

1. Configure the guest OS in accordance with the installation guidelines.

2. Provision vCPU and memory as per the application requirements.

3. Install and configure VMware Tools. This is a crucial step, as clock sync, VMware paravirtual driver and graphic support are part of the tools.

4. Configure a paravirtual virtual SCSI adapter for the guest. In our testing, we found paravirtual driver outperforms all other adapter types. If paravirtual driver is not on the supported guest OS, LSI Logic SAS is a good alternative.

Some guest-specific performance tuning best practices are listed below:

Linux
• Use two or more Fibre Channel ports per initiator.

• Set the block device scheduler `for [deadline] or [noop].`

```
for disk in `lsscsi | grep PURE | awk '{ print $6 }'`
do
  echo noop > /sys/block/${disk##/dev/}/queue/scheduler
done
```

- To avoid interrupt contention on a particular CPU, use a Linux kernel that supports setting `rq_affinity` so that it will schedule I/O on the core that originated the process.

```
for disk in `lsscsi | grep PURE | awk '{ print $6 }'`
do
   echo 2 > /sys/block/${disk##/dev/}/queue/rq_affinity
done
```

- Stop the system from using the block devices as a source of entropy, reducing any impact that has on system resources.

```
for disk in `lsscsi | grep PURE | awk '{ print $6 }'`
do
   echo 0 > /sys/block/${disk##/dev/}/queue/add_random
done
```

- Check your HBA queue depth.

```
cat /sys/module/qla2xxx/parameters/ql2xmaxqdepth
echo 64 > /sys/module/qla2xxx/parameters/ql2xmaxqdepth
```

- With a thread-heavy workload, a larger file descriptor limit may be needed. (This is functional and not performance related.)

```
ulimit -n 819200
```