

DATABASES IN THE CLOUD

HOW VIRTUALIZATION AND AUTOMATION ENABLE TRUE DBAAS

Jeff Stonacek, House of Brick Technologies

Jim Hannan, House of Brick Technologies

Kris Cook, SITA



402.445.0764 | 877.780.7038 | www.houseofbrick.com
9300 Underwood Ave. Suite 300 | Omaha, NE 68114



Abstract

Everyone is talking about Cloud Computing, but there seems to be wildly differing perspectives on what it means, how it works, and the costs and benefits of such a model. When it comes to databases and other business critical applications in the cloud, the message is even more confusing. In this paper, we will share our lessons learned while helping some of the biggest companies in the world with their Cloud and Database as a Service (DBaaS) strategies. We will also share best practices and success stories around automating the provisioning and deployment of fully customized, complex application environments based on Oracle core database technology.

Authors:

- Jeff Stonacek, House of Brick Technologies - Principal Architect
- Jim Hannan, House of Brick Technologies - Principal Architect
- Kristopher Cook, SITA - Senior Manager, Database Services

Other Contributors:

- Nathan Biggs, House of Brick Technologies - CEO

Target Audience

- IT Architects
- Technical Administrators
- IT Management
- Business Unit Application Owners

Executive Summary

This presentation is not merely a description of tooling required to automate database deployments. DBaaS involves much more than that. DBaaS requires a culture change for most traditional IT shops. There are organizational and political concerns along with the licensing, security and technical automation components.

Organizational Concerns

- IT cannot operate in silos
- IT and business units must agree on a service catalog
- Need to embrace virtualization
- Business units have more infrastructure power



DbaaS Key Points

DBaaS has the following characteristics, public or private cloud:

- Consumer driven
- Utility pricing
- Elastic
- Self service
- Automation

Industry Tooling

Oracle Corporation has developed tools to enable DBaaS, as have other vendors. In this presentation we will compare and contrast Oracle's DBaaS offering with that of VMware. Each has its benefits and drawbacks.

Oracle

- Based on OEM 12c Cloud Control
- Oracle pushes their proprietary hardware platform, which is not virtualization friendly
- Tooling gap is on the IaaS side, difficult to provide IaaS with Exadata
- Oracle tooling revolves around Oracle products

VMware

Since the release of vSphere 4, VMware has been developing tooling to do both IaaS and DBaaS. The first product they introduced was VMware vFabric Data Director in August of 2011. This product was able to offer a service catalog of T-shirt sized Oracle Database appliances—small, medium and large. The word appliance was used due to the fact that the database virtual machine was difficult to modify and the only supported configuration was done through the web UI. With the release of vCloud Automation Center 2 years ago, VMware took the industry a giant step forward with IaaS. vCloud Automation Center has since been rebranded by VMware as vRealize Automation. It is now included in the vRealize Suite.

- Based on virtualized infrastructure
- Supports a hybrid mix of multiple cloud platforms (vCloud Air, Amazon Web Services, Open Stack, etc.)
- x86 hardware agnostic
- Relies heavily on vRealize Suite of products, as well as custom workflows for automation
- Tooling gap is after the deployment of the virtual machine or DBaaS



- VMware's DBaaS and IaaS vRealize strategies can be adopted for most applications (not just Oracle)

Background

Cloud computing is a hot topic in the IT industry. It has spawned a revolution in the fields of infrastructure and provisioning. Since House of Brick is heavily invested in providing services around virtualizing tier 1 workloads, it made sense to delve into the service model. In this paper we will discuss the automated service model for delivering IT components. In other words, providing Infrastructure as a Service (IaaS), Database as a Service (DBaaS) and Application as a Service (AaaS).

Definitions

Lets first provide some definitions for the common acronyms found today.

IaaS

Calling something Infrastructure as a Service (IaaS) can be a bit misleading. It is difficult to automate the provisioning of a core network switch, for example. When referring to IaaS, we are typically talking about the low level components necessary for running an application. Standing up a server is a prerequisite to running a database. Automating the provisioning of that server, including installing and configuring operating system (OS) components, constitutes IaaS.

Virtualization is **THE** key enabler of IaaS.

PaaS

Another acronym that can be a bit ambiguous is Platform as a Service (PaaS). This refers to the process of providing a computing platform on which to run an application. Things like Java middleware frameworks would be a good example of PaaS.

DBaaS

Database as a Service (DBaaS) seems like the easiest term to quantify, as anything required standing up a database. In the case of Oracle, that means satisfying prerequisites, adding storage, installing binaries and creating a database. DBaaS is really a component of PaaS.



SaaS

Software as a Service (SaaS) applications are those that consumers access over the Internet, such as Salesforce.com. SaaS applications are sold “by the slice”, so that consumers pay for only those units that they use (such as number of users on Salesforce).

The Cloud Model

The service model has certain characteristics that make it work. However, many of them go against the traditional IT operating model, which must be addressed to make Cloud Computing possible.

Consumer Driven

Consumers drive the cloud model. Consumers are usually associated with the following groups:

- Development teams
- Application teams
- Operations
- Business unit owners

The idea behind Cloud computing is to satisfy the consumers without increasing infrastructure operations and related costs. The service provider, public or private, may own the infrastructure but the consumers own the application and business processes that consume the infrastructure.

Consumer driven implies self-service. In other words, the consumer makes decisions regarding which environments are needed, when they should be delivered, and the performance characteristics. Infrastructure teams must approve all requests, but the consumers run the show.

Workflow

The Cloud model is predicated upon an automated workflow engine. An automated workflow engine has the following characteristics:

- Communication mechanism (send a page, email, etc.)
- Approval mechanism
- Automated task initiation



IT organizations have a long history of utilizing workflow engines however; they were typically more informal in nature. For example, an application team would call the server team to request a new server for a software project. In the Cloud model, this communication stream must be automated or the system comes grinding to a halt. With infrastructure and operations teams potentially located around the globe, communications can come to a grinding halt if there isn't a framework in place to keep things moving along. Approvals automatically initiate the next level of communication, and so forth.

The critical piece to an automated workflow framework is task initiation. Once the DBA manager approves the request for a new database, the workflow engine should take on the task of creating the service. The only manual task necessary should be reading the email that states the environment is ready for use.

Service Catalog

A service catalog defines what offerings are available to the consumer. The hard part is defining the components (i.e. application server, database, etc.) and the configuration that will satisfy the request. The components of the service catalog must be very granular in nature. Once built though, the process of deploying becomes the job of the Cloud Management software.

While this is a technology solution, there is quite a bit of human intervention required to properly construct the service catalog.

Utility Pricing

The utility pricing model drives the Cloud. When requesting resources from a public cloud provider, the consumer pays by the unit. That unit may be CPU consumption, memory consumption, IOPs, etc.

In a private cloud, there needs to be a charge back model in place where groups in an organization have a budget for resources and are charged based on their consumption. If not, consumption runs rampant and licensing costs can skyrocket.

Elasticity

The Cloud Computing model relies on capacity on demand. Certain projects and activities require bursty capacity. A good example is patching. If an application requires a patch, it may be necessary to stand up a new environment (AaaS) and apply the patch for testing. Once the patch is promoted to production the environment can be destroyed.



Automation

Automation, again, is the key to making all of these components work. Automated workflow, automated provisioning, automated configuration, etc. Without automation there is no Cloud.

Conclusion

The Cloud Computing model is a way to increase efficiency with large, dynamic IT organizations. Without proper structure, the system will be unmanageable.

Technical Discussions and Examples

Organizational Impact

Many organizations that we talk to at House of Brick are excited about the vision and promise of Database as a Service (DBaaS). At a high level, this includes self-service automated provisioning of application-ready database servers that are fully configured and deployed on their infrastructure. What we have found, however, is that DBaaS involves much more than deploying new technology. There are key people and process concerns that, if not considered, introduce far more risk into your implementation plans.

Below are a few topics we routinely discuss with our customers. Their feedback on these topics let us work together to provide the consulting that they need in order to implement a rock-solid DBaaS strategy:

Is your organization DBaaS ready? (BTW...that is not a technology question!)

DBaaS is as much a people and process consideration (if not more so) than a technology implementation. Here are some intermediate questions to help you decide if your organization is ready.

- Do the silos in your IT organization work well together or is there infighting and territorialism?
- Do you have a strong, effective working relationship with the business units that you provide IT services for?
- Have you designed a service catalog for the business units that focuses on Service Level Agreements (SLA) rather than architecture specs?
- Have you committed to a 100% virtualization strategy?
- Are you ready to empower IT groups and business units to make self-service requests that are fulfilled automatically?



If you shoehorn DBaaS automation technologies into an organization that is not Cloud ready, then you will not achieve the vision, or experience the benefits of Cloud computing. Instead, it will continue to be business as usual, but with more headaches for people to work around.

Do you have the right people on staff for DBaaS?

Traditionally, IT staff could pick a silo in which to work exclusively. The DBA would request services from the network, storage, and infrastructure teams, but would never cross into those domains.

Cloud computing and DBaaS are bringing a convergence in many areas. We are using converged infrastructure for the host, and we are using converged tooling for management. Next you need to consider converged IT staff.

In order for Cloud Computing and DBaaS to be successful, you need at least one, and ideally a whole group of professionals, who have converged expertise across all disciplines (DBA, Storage, Sys Admin, Architecture, Network, Security, Virtualization, etc.). This convergence of expertise will allow them to quickly diagnose and solve problems without finger pointing from other groups. It will also help in architecting solutions that are appropriate for the business because they consider all aspects of IT.

You can do this in your own datacenter without paying a public cloud provider

Database as a Service does not require swiping your credit card at Amazon, or any number of other public Cloud providers. You can start in your own datacenter. The first step is to make sure that you have a virtualization strategy that includes your business critical systems. Automation and self-service can then be done incrementally as the organization proves ready.

Despite what many software vendors are saying, DBaaS tools are still immature

Geoffrey Moore identified a chasm in the technology adoption lifecycle between the early adopters and the early majority customers. We look at a lot of tools and a lot of customers, and in my opinion, nobody has yet crossed the DBaaS chasm. If you want to know why we think that has not happened yet, re-read the topics in this white paper.

Oracle's Cloud management suite of tools provides the necessary framework for deploying databases. As is Oracle's custom, many individual tools are bolted together to form a suite of tools (See Oracle Streams as an example). To create a database there are a few different options. EM 12c can employ DBCA and DBCA templates to create fresh databases, or RMAN can be used to clone databases.

Other vendors struggle when it comes to the database portion of Oracle service catalog items. VMware corporation has some of the best (if not the best) tooling available for Cloud computing.



However, when it comes to DBaaS the tooling falls short of fully automating the provisioning and configuration process. It is then up to the infrastructure team to fill in the gaps with scripts and logic.

Have you figured out how your licenses will be impacted by your DBaaS implementation decisions?

Your Oracle, Microsoft, or other software licenses probably cost you more every year (by several times) than all of your hardware put together. If you're not careful, a poorly architected Cloud strategy will absolutely blow your software licensing costs sky-high. We have good news though – that does not have to happen. Spend some time understanding your actual agreements to see where you have to pay licensing fees. Then carefully architect your DBaaS plans to minimize those license costs. House of Brick can help you with consulting in this area, or you can read the excellent blog posts on <http://www.houseofbrick.com> that deal with licensing issues (especially those concerning Oracle).

DBaaS is an exciting new idea that is getting a lot of buzz. As with all new technology trends, there are some innovators and early adopters that are paying the price of figuring out the hard stuff so that the mainstream customers can have it easier later on. I hope that in reading this white paper, however, you understand that DBaaS, and Cloud Computing in general, involves far more than just technology. In order for your DBaaS strategy to be successful, you must also consider the people and process components with equal if not more weight.

Oracle Enterprise Manager 12c - Cloud Services

Oracle's DBaaS offering revolves around its Enterprise Manager suite of tools. Oracle Enterprise Manager (OEM) has been around for a long time and Oracle has continued to improve and expand the functionality of the tooling.

OEM Packs

OEM is licensed by packs and each different pack does different things. For DBaaS, the Oracle Cloud Management Pack is required. The Cloud Management Pack requires licensing of the Database Lifecycle Management as well.

The Enterprise Manager packs are licensed by the underlying hardware of any targets that are managed by EM. The calculation for number of licenses needed, in a CPU based licensing model, is number of processor cores of the target servers.

As of the time of this white paper, the costs for the DBaaS packs for EM are as follows:

- Lifecycle Management Pack - \$12,000 per processor
- Cloud Management Pack for Oracle Database - \$5,000 per processor



As you can see, this is not a cheap solution. Be sure to do your homework before embarking on the Oracle Cloud model.

OEM Cloud Services Hierarchy

Now that the bean counting is out of the way, let's discuss the nuts and bolts of how Oracle implements DBaaS.

Zones

Anyone familiar with Solaris will recognize the term “zone”. Oracle has reused this terminology with respect to DBaaS. From an EM 12c perspective, a zone is a way to segregate infrastructure resources for later consumption. This is conceptually similar to a Solaris Zone. EM refers to these as Platform as a Service (PaaS) infrastructure zones. A zone can be a physical host, an Exadata compute node or an OVM virtual machine. Think of a zone as compute resources.

Resource Pools

A pool is a specific resource that is assigned to a zone. A pool resource is anything that allows for the running of a database – for example, the appropriate version and patch level of Oracle binaries.

A pool can be categorized in the following ways:

- **Pools for databases** – resources needed to create a new database
- **Pools for a container database** – used to plug in a tenant pluggable database (PDB)
- **Pools for a schema** – old school database consolidation

Service Catalog Template

A service catalog template is a list of options available to the end consumer. In this case, a service catalog could be:

- DBCA Template for the building of a new database
- RMAN backup of an existing database to be used for a clone
- Schema level export to be used to clone a schema
- Linked Clone operation to make a Change on Write clone of a database



These service catalog items are pre-defined and stored in OEM. Consumers are granted access controls to zones and then are able to provision service catalog items.

Infrastructure as a Service

The main gap with regards to Cloud Computing in the Oracle model is a lack of automation around Infrastructure as a Service (IaaS). Oracle is pretty good at automating the provisioning of database resources once the infrastructure is up and running. However, the process for standing up server resources is manual and typically requires substantial lead-time. It is hard to imagine having spare Exadata capacity just lying around waiting for future consumption.

As stated earlier, virtualization is THE key enabler of IaaS. The Oracle gap with regards to IaaS revolves around Oracle's virtualization stack. The following points highlight why Oracle's Cloud model struggles to deliver IaaS:

- VMware is best of breed with regards to virtualization
- OVM has historically had issues with performance and stability
- Exadata, Oracle's best of breed hardware platform is a big, mainframe like computer that does not embrace virtualization
- Tailored primarily to Oracle workloads, what do I do with everything else?

By contrast, fully virtualized platforms like VMware have IaaS built in from the ground up. VMware vSphere has full-featured template functionality built in, as well as automated provisioning suites like vRealize, to manage the Cloud Computing environment.

VMware vRealize – Cloud Services

The core of VMware's IaaS and DBaaS is part of the vRealize suite of products. vRealize is sold in Advance and Enterprise editions. In this white paper we will focus on the Enterprise edition. vRealize Enterprise Suite includes the following products:

- VMware vRealize Automation
- VMware vRealize Operations
- VMware vRealize Log Insight
- VMware vRealize Orchestration

vRealize Suite has two licensing options by per processor or OSI (OS Instance). The OSI model as suggested by VMware is friendlier to public cloud solutions.

vRealize Operations

Formerly named vCenter Operations Management Suite, the rebranded product has been



placed under the vRealize umbrella and renamed to vRealize Operations. vRealize Operations is the centralized one-pane of glass interface for your complete enterprise including virtual and physical hardware. At the heart of tool it allows for management hosts, virtual machines and monitoring capabilities. VMware vRealize Suite can also manage other hypervisor platforms such as Microsoft Hyper-V and Redhat KVM. It works with external cloud service providers such as VMware vCloud Air and Amazon Web Services or OpenStack-based private and public clouds. vRealize Operations 6.0 can manage up to 3 vCenter Servers, 360 hosts and 18,000 virtual machines. Additionally, you can purchase third party management packs, some notable ones are listed below:

- Brocade
- NetApp
- EMC
- Blue Medora for Oracle OEM
- IBM Tivoli

House of Brick predicted a tool like this would become available to the industry at some point. It was clear while working with many of our customers that the hybrid environment is the industry norm. Most IT shops have a combination of physical and virtual hardware. Additionally, it is common for IT organizations to run two different hyper-visors, most commonly Hyper-V and vSphere. vRealize Operations helps to bridge the management of those hybrid environments.

vRealize Log Insight

vRealize Log Insight is the log management tool for vRealize Suite. In the theme of centralizing the hybrid environment for both physical and virtual, this tool allows for log management and analysis in one tool. Out of the box the monitoring is very vSphere centric, but additional third party Content Packs can be added to the tool. Those tools can be found here: <https://solutionexchange.vmware.com/store/loginsight>. Log Insight offers an intuitive interface for creating custom reports and allows for drill downs on specific objects like a virtual machine.

vRealize Automation

Formerly named vCloud Automation Center, vRealize Automation is the core component for doing DBaaS and automated deployments within vRealize. The solution is based on building a service catalog. The service catalog contains prebuilt virtual machine templates for agile deployment or DBaaS. Fundamentally this approach has some nice advantages, it is quick and with good patching processes and tools you can get the virtual machine updated with little effort. The gap however exists within the Oracle stack, especially with Oracle software versions and database configuration. Most Oracle DBAs understand how this can be unique from one build to the next. Because of this, most IT organizations that use DBaaS choose to



run scripts after the deployment of the virtual machine to setup the database. The biggest challenge is for software vendors like Oracle and VMware, is that the point of entry for customer customization varies across the industry. For example, a software company hosting a banking application for small banks to offer online banking would be very standardized on builds and configurations. In this case the templates would be near of 100% complete and ready. An organization supporting lots of internal applications would have many versions and Oracle parameters that differ across builds. Taking a deeper look at the challenges templates present, we discovered the following issues:

- Templates are out of date within days of creation
- A large catalog must be maintained for every Oracle binary version and unique database setting or parameter
- Deployment of Oracle OEM agents needs to be done after the host name is defined

In House of Bricks experience, scripting agent installs and databases installation fills the gap with the VMware tooling. Pure Storage has a great example of this in their white paper [“Deploying Database-as-a-Service with Pure Storage and VMware”](#)¹.

In this white paper they describe a technique very common in the industry, scripting the OEM agent after the new Oracle virtual machine has been deployed from the vRealize Server Catalog. The script runs through a sub-component of vRealize Automation called vRealize Orchestrator. Orchestrator is a product that has been around for more than 5 years. It’s a workflow engine that can automate certain tasks. We touched on OEM packs a few times in this section. For clarity, most customers under a VMware solution license only Oracle Diagnostic Pack and Tuning Pack and not the previously discussed Lifecycle Manager and Cloud Management Pack discussed in the Oracle Enterprise Manager 12c - Cloud Services section of this paper.

vRealize Orchestrator as VMware states “simplifies the automation of complex IT tasks...2”, which we find to be a very good summary of the product. Orchestrator is able to automate tasks at the infrastructure level, virtual machine and OS layer.

Case Study

About SITA

Serving the air transport industry since 1949, SITA provides services to over 450 members and 2,800 customers in more than 200 countries and territories. Airlines, airports, governments, air cargo, aerospace and international organizations all use SITA's IT and communications solutions. Almost every airline and airport in the world does business with SITA.

They provide the broadest portfolio for the air transport industry, including:



- Managed global communications, infrastructure and outsourcing services
- Services for airline commercial management, passenger operations, flight operations, aircraft operations, air-to-ground communications, airport management and operations, baggage operations, transportation security and border management, cargo operations and more

SITA has 3 datacenters in Atlanta and 2 more in London that all host databases. They were in the middle of a multi-year data consolidation effort, but the time that it was taking to build physical servers to migrate on to was just not efficient for their purposes.

Because of these struggles, in 2012 SITA started a journey to redefine the core components of their IT infrastructure. The SITA IT leadership team decided to virtualize 100% of their workloads, both existing and new, using VMware vSphere. A critical step SITA took towards successfully virtualizing workloads was to define standards for their technology services. The standards produced started with hardware and continued all the way to defining the application layer (Oracle, SQL Server, WebLogic). In reality, this was not only a virtualization journey for SITA, but also a journey to redefine IaaS and DBaaS. The process of defining the standards, starting with building Reference Architecture (RA) for the infrastructure components, became the blueprint for deploying and operationalizing their new environment. The Reference Architecture plan covered the following areas:

- Storage
- Network
- Blade Chassis
- VMware vSphere
- Red Hat Linux
- Appliance RAs: Oracle Database/SQL Server/WebLogic

A Reference Architecture review board was put together that included IT managers and SITA architects who reviewed each of the RA documents. This review process, would in some cases, require many revisions before the documents could be officially published. While not a rigidly defined timeframe, approximately once a year the RA docs are revised and reviewed using the same process. Depending upon technology changes within each of the stacks, the revisions may take more time, as in the case of the integration of Oracle 12c. After an RA document has been completed, a standards document is then created.

The early adoption of Reference Architecture proved to be a significant one for SITA and provided them with a foundation on which to successfully build Automation and DBaaS. With clear standards in place, it was predictable how a database creation script would run. The DBaaS is fully automated today through scripts and a tool called Process Automation Manager from CA. SITA went live on the first Production Oracle RAC virtualize build in January 2014.



References

¹ White Paper: [Deploying Database-as-a-Service with Pure Storage and VMware](#)

² [Product Information: vRealize Orchestrator](#)

Additional Resources

DBaaS Blogs:

[DBaaS Overview](#)

[5 Things a Business Person Should Think About When Considering DBaaS](#)

Oracle Licensing Blogs:

[Soft Partitioning on VMware is Your Contractual Right](#)

[Managing Oracle Licensing in a Shared Storage Environment](#)