An Oracle White Paper
May 2010

Oracle Cloud Computing

Executive Overview

Cloud computing is a significant advancement in the delivery of information technology and services. By providing on demand access to a shared pool of computing resources in a self-service, dynamically scaled and metered manner, cloud computing offers compelling advantages in speed, agility and efficiency. Today, cloud computing is at an early stage in its lifecycle, but it is also the evolution and convergence of several trends that have been driving enterprise data centers and service providers over the last several years.

Cloud computing builds off a foundation of technologies such as grid computing, which includes clustering, server virtualization and dynamic provisioning, as well as SOA shared services and large-scale management automation. For the better part of a decade, Oracle has been the leader in these areas with thousands of customer successes and high level of investment. Today, Oracle offers the industry's most complete, open and integrated products and services to enable public, private and hybrid clouds.

Oracle aims to make cloud computing fully enterprise-grade and supports both public and private cloud computing to give customers choice. Oracle offers technology that enables organizations to build private clouds, leverage public clouds and provide cloud services to others. Oracle also offers a broad set of horizontal and industry applications that run in a shared services private cloud model as well as a public Software-as-a-Service (SaaS) cloud model.

This white paper provides an overview of Oracle's cloud computing strategy and how Oracle helps customers and partners plan their evolution and adoption of a cloud computing model.

Introduction to Cloud Computing

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This definition from the National Institute of Standards[1] has gained broad support from the industry.

The NIST definition of cloud computing describes five essential characteristics, three service models and four deployment models.

Five Essential Characteristics

- **On-demand self service** –Users are able to provision, monitor and manage computing resources as needed without the help of human administrators
- **Broad network access** – Computing services are delivered over standard networks and heterogeneous devices
- **Rapid elasticity** – IT resources are able to scale out and in quickly and on an as needed basis
- **Resource pooling** – IT resources are shared across multiple applications and tenants in a non-dedicated manner
- **Measured service** – IT resource utilization is tracked for each application and tenant, typically for public cloud billing or private cloud chargeback

Three Service Models

- **Software as a Service (SaaS)** – Applications delivered as a service to end-users typically through a Web browser. There are hundreds of SaaS service offerings available today, ranging from horizontal enterprise applications to specialized applications for specific industries, and

[1] NIST Definition of Cloud Computing v15.

also consumer applications such as Web-based email. *Oracle CRM On Demand* is an example of a SaaS offering that provides both multi-tenant as well as single-tenant options, depending on the customer's preference. Oracle also offers enterprise-grade enabling technology to Independent Software Vendors (ISVs) to build their own SaaS offerings. Oracle calls this enabling technology the *Oracle Platform for SaaS*. Hundreds of ISVs have built their SaaS offering on top of the Oracle Platform for SaaS.

- **Platform as a Service (PaaS)** – An application development and deployment platform delivered as a service to developers who use the platform to build, deploy and manage SaaS applications. The platform typically includes databases, middleware and development tools, all delivered as a service via the Internet. PaaS offerings are often specific to a programming language or APIs, such as Java or Python. A virtualized and clustered grid computing architecture is often the basis for PaaS offerings, because grid provides the necessary elastic scalability and resource pooling. Oracle offers a comprehensive PaaS product offering for public cloud service providers as well as enterprise customers to build their own public clouds. Oracle calls this the *Oracle PaaS Platform* (more on this later in this paper).
- **Infrastructure as a Service (IaaS)** – Compute servers, storage, and networking hardware delivered as a service. This infrastructure hardware is often virtualized, so virtualization, management and operating system software are also part of IaaS as well. An example of IaaS is Amazon's Elastic Compute Cloud (EC2) and Simple Storage Service (S3). Oracle does not offer IaaS cloud services, but Oracle provides hardware and software products to other IaaS providers to enable their public cloud services, and also offers the same technologies to enterprises for private use.

Four Deployment Models

- **Private Clouds** – For exclusive use by a single organization and typically controlled, managed and hosted in private data centers. The hosting and operation of private clouds may also be outsourced to a third party service provider, but a private cloud remains for the exclusive use of one organization.
- **Public Clouds** – For use by multiple organizations (tenants) on a shared basis and hosted and managed by a third party service provider.
- **Community Clouds** – For use by a group of related organizations who wish to make use of a common cloud computing environment. For example, a community might consist of the different branches of the military, all the universities in a given region, or all the suppliers to a large manufacturer.

- **Hybrid Clouds** – When a single organization adopts both private and public clouds for a single application in order to take advantage of the benefits of both. For example, in a "cloudbursting" scenario, an organization might run the steady-state workload of an application on a private cloud, but when a spike in workload occurs, such as at the end of the financial quarter or during the holiday season, they can burst out to use computing capacity from a public cloud, then return those resources to the public pool when they are no longer needed.

Comparing Public and Private Clouds

The two basic models of public and private clouds have a number of compelling business benefits, some of which are common to both public and private, while others are only for one or the other.

Benefits common to both public and private clouds include:

- **High efficiency** – Because both public and private clouds are based on a grid computing and virtualization, both offer high efficiency and high utilization due to the sharing of pooled resources, enabling better workload balancing across multiple applications.
- **High availability** – Another benefit of being based on grid computing is that applications can take advantage of a high availability architecture that minimizes or eliminates planned and unplanned downtime, improving user service levels and business continuity.
- **Elastic scalability** – Grid computing also provides public and private clouds with elastic scalability, the ability to add and remove computing capacity on demand. This is a significant advantage for applications with highly variable workload or unpredictable growth, or for temporary applications.
- **Fast deployment** – Because both public and private clouds can provide self-service access to a shared pool of computing resources, and because the software and hardware components are standard, re-usable and shared, application deployment is greatly accelerated.

Some benefits are unique to public cloud computing:

- **Low upfront costs** – Public clouds are faster and cheaper to get started, so they provide users with a low barrier to entry because there is no need to procure, install and configure hardware.

- **Economies of scale** – Large public clouds enjoy economies of scale in terms of equipment purchasing power and management efficiencies, and some may pass a portion of the savings onto customers.
- **Simpler to manage** – Public clouds do not require IT to manage and administer, update, patch, etc. Users rely on the public cloud service provider instead of the IT department.
- **Operating expense** – Public clouds are paid out of the operating expense budget, often times by the users' line of business, not the IT department. Capital expense is avoided, which can be an advantage in some organizations.

Other benefits are unique to private cloud computing:

- **Greater control of security, compliance and quality of service** – Private clouds enable IT to maintain control of security (data loss, privacy), compliance (data handling policies, data retention, audit, regulations governing data location), and quality of service (since private clouds can optimize networks in ways that public clouds do not allow).
- **Easier integration** – Applications running in private clouds are easier to integrate with other in-house applications, such as identity management systems.
- **Lower total costs** – Private clouds may be cheaper over the long term c to public clouds, since it is essentially owning versus renting. According to several analyses, the breakeven period is between two and three years.
- **Capital expense and operating expense** – Private clouds are funded by a combination of capital expense (with depreciation) and operating expense.

Cloud Benefits and Challenges

Recent surveys show that the top two benefits of cloud computing are speed and cost. Through self-service access to an available pool of computing resources, users can be up and running in minutes instead of weeks or months. Making adjustments to computing capacity is also fast, thanks to elastically scalable grid architecture. And because cloud computing is pay-per-use, operates at high scale and is highly automated, the cost and efficiency of cloud computing is very compelling as well.

The same surveys reveal that there are a number of issues and concerns that are holding some organizations back from rushing to the cloud. The top concern far and away is security. While one can debate the relative security of public clouds versus in-house data centers, the bottom line is that many organizations are not comfortable entrusting certain sensitive data to public clouds where they

do not have full visibility and full control. So some particularly sensitive applications will remain in-house while others may take advantage of public clouds. Another concern is quality of service, since clouds may not be able to fully guarantee service level agreement in terms of performance and availability. A third area of concern is fit, the ability to integrate with in-house systems and adapt SaaS applications to the organization's business processes.

To recap, cloud computing is characterized by real, new capabilities such as self-service, auto-scaling and chargeback, but is also based on many established technologies such as grid computing, virtualization, SOA shared services and large-scale, systems management automation. Cloud computing offers compelling benefits in terms of speed and cost, but also presents serious concerns around security, compliance, quality of service and fit. Organizations will likely adopt a mix of public and private clouds. Some applications will be appropriate for public clouds, while others will say in private clouds, and some will not use either.

With this view of cloud computing in mind, we now turn to the role that Oracle and its products play in the cloud era.

Oracle Cloud Computing Strategy

Oracle's overall corporate strategy is to provide the industry's most complete, open and integrated set of products from applications to disk. For cloud computing, Oracle's strategy is to:

- **Ensure that cloud computing is fully enterprise grade** – Oracle provides enterprise grade technology for high performance, reliability, scalability, availability, security and portability/interoperability (based on standards). Enterprises demand these characteristics before moving important workloads to a public or private cloud.
- **Support both public and private clouds to give customer choice** –Organizations are adopting different deployment models for cloud computing for different applications at different rates of speed, so Oracle supports customers no matter what type of cloud or non-cloud they choose.
- **Deliver most complete PaaS and IaaS product offerings** – Oracle provides the most complete portfolio of software and hardware products to enable organizations to build,

deploy and manage public and private PaaS and IaaS. A key element of Oracle's strategy is to offer the *Oracle PaaS Platform*, which is discussed in more detail later in this paper.

- **Develop and enable rich SaaS offering** – Oracle offers a very broad portfolio of horizontal and industry applications that are deployed in either a private shared services environment or in a public SaaS model.

While there is still significant debate on when and how the different dimensions of cloud computing will become viable for different industries and businesses, there is broad agreement that it will have an increasing impact on nearly every IT organization.

Cloud computing is driving a significant part of Oracle's product development plans – from enterprise applications to middleware, databases, servers and storage devices, as well as cloud management systems. Taken together, these developments are building off Oracle's grid computing architecture to create an out-of-the-box solution for cloud computing: *Oracle PaaS Platform*.

Oracle PaaS Platform

The Oracle PaaS Platform is a comprehensive portfolio of products to build an application platform delivered as a public or private cloud service. The Oracle PaaS Platform is based on Oracle grid technologies including Oracle Database with Real Application Clusters and Oracle application grid including WebLogic Server, Coherence in-memory data grid and the JRockit JVM. On top of this foundation of clustered middleware and database technologies, the Oracle PaaS Platform also includes components such as Oracle SOA Suite, Oracle BPM Suite, Oracle Identity Management and Oracle WebCenter.

The Oracle PaaS Platform is built on top of a robust Oracle IaaS offering consisting of Oracle Solaris, Oracle Enterprise Linux and Oracle VM for virtualization, Sun SPARC and x86 servers, and Sun storage. Both the Oracle PaaS and Oracle IaaS are managed by Oracle Enterprise Manager, which provides integrated systems management from applications to disk across the complete cloud deployment lifecycle.

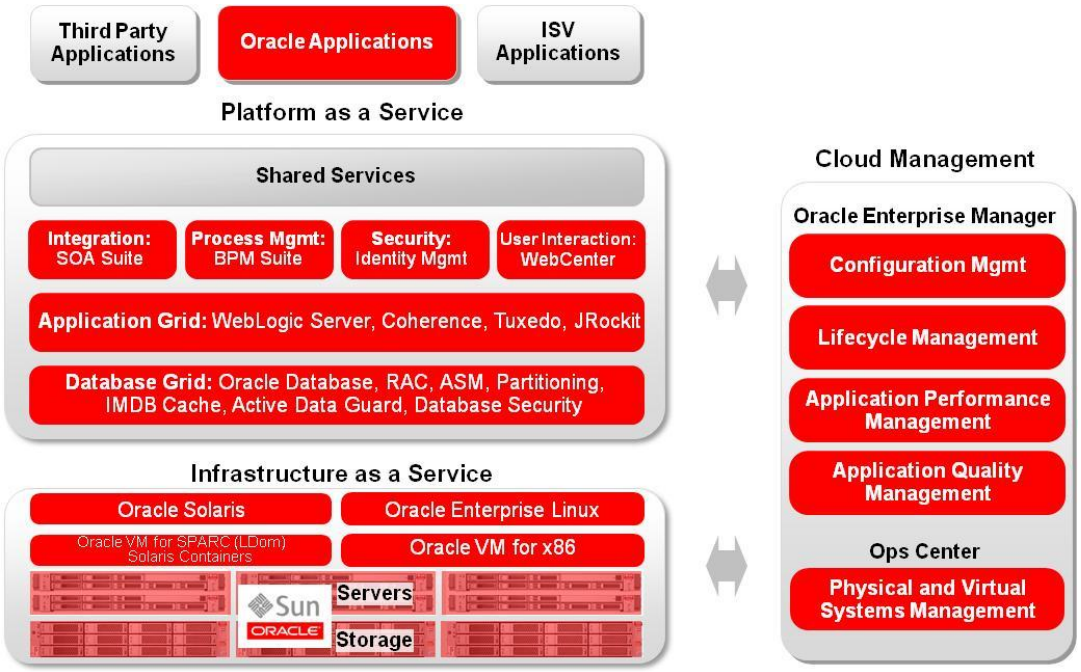

Figure 1. Oracle PaaS Platform

For many organizations deciding between building a PaaS or IaaS offering, the key issue is how much of a standardized, re-usable and shared platform do they want to provide their users. An IaaS provides the basic compute, storage and networking capacity, so it is the most flexible, but it requires users to provide the rest, including the application, middleware and database, resulting in greater development cost, time and heterogeneity. For many organizations, a private PaaS is a natural strategy that benefits users as well as the IT service provider. A PaaS gives users a standardized, re-usable and shared starting point for application development , providing faster and simpler development with sufficient flexibility. From IT's standpoint, a PaaS offering means greater manageability, security, consistency, efficiency and control.

Private PaaS Lifecycle

To illustrate how a private PaaS would work within an enterprise, the following is a brief description of the basic steps in the private PaaS lifecycle:

1. First, the IT department sets up the PaaS Platform, including a number of standard, shareable components (may be Web services, BPM processes or UI components) and a self-service application for users.
2. Second, a departmental developer goes through the self-service application to discover the available shared components, assembles the application from those components instead of

creating it from scratch, and deploy the application, all through self-service with no IT involvement.

3. Next, end-users begin to use the application.
4. To manage the application, the application owner can use the self-service application to monitor the application, adjust capacity if necessary, and track usage and how much they are being charged for use of the cloud resources.

Clearly the above lifecycle is an over-simplification of the real process, but is illustrative of the basic steps.

Enterprise Evolution to Cloud Computing

Most organizations will not jump immediately to cloud computing. Instead, most will evolve their current IT infrastructure to become more cloud-like characteristics over time. The available technology is rapidly evolving and advancing, but organizations will also need to change their policies and processes. In many cases, the technical building blocks for cloud computing are available in advance of enterprise readiness.

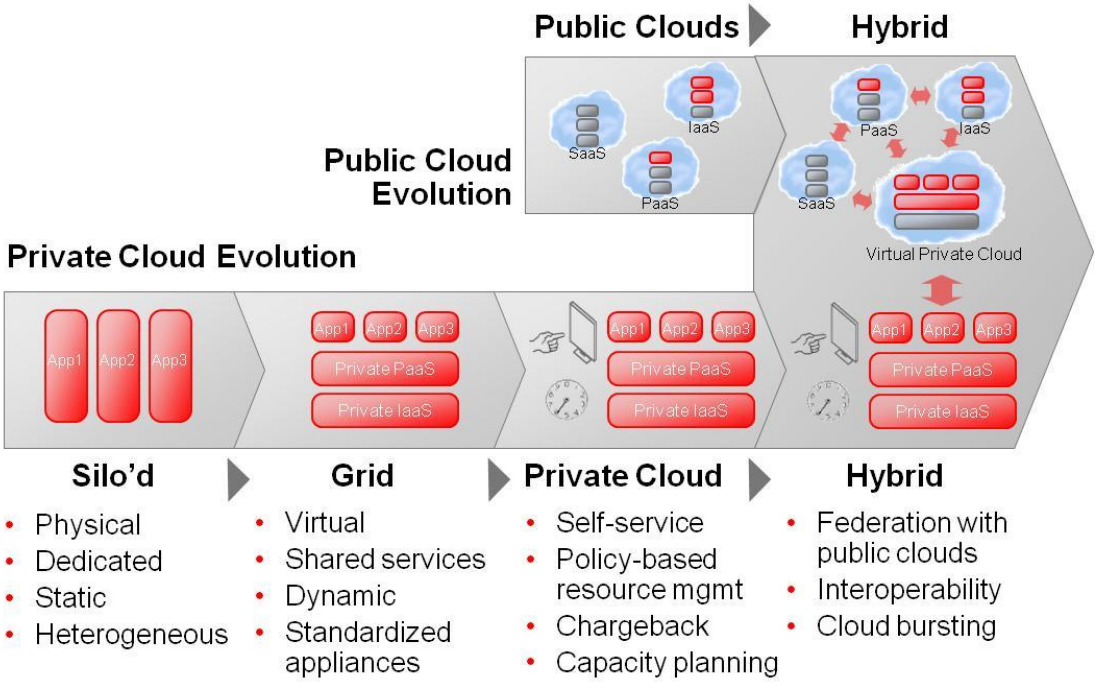

Figure 2. Enterprise Evolution to Cloud Computing

The first step in the evolution that many enterprises are already taking is to move from a siloed environment to a grid or virtualized environment –moving from a dedicated, rigid, physical structure for each application to a virtual environment with shared services, dynamic provisioning and

standardized configurations or appliances. This trend is very strong right now. Many enterprises are leveraging grid and virtualization technologies to consolidate and reduce costs. Oracle has a very strong and complete offering for grid, with products in the database and middleware layers, such as Real Application Clusters, TimesTen, WebLogic and Coherence, plus Oracle VM for server virtualization and Enterprise Manager for managing the entire stack.

From here, enterprises can evolve to cloud by adding self-service and pay-per-use to the envoniment. A user goes to the employee portal, signs in, makes a request for a virtual machine with a certain amount of CPU, memory and disk, picks a VM image for database or middleware, then clicks "submit." If that employee's role and entitlements allow her to have that amount of IT resource, then it is automatically provisioned without an IT administrator being involved. If not, the request may be routed to a manager and/or IT for approval. In just a few minutes, they are up and running with an instance of the PaaS platform. After the application is deployed, the system has policy-based resource management to automatically make capacity adjustments, and the employee's business unit gets an internal charge every month based on how much IT resources they consumed. To make all that happen, the enterprise must have policies and processes defined, and the technology must be able to support it.

Meanwhile, public clouds are also evolving. There are already many different public cloud offerings at all the layers of the cloud: SaaS, PaaS and IaaS. Many began as ISVs or hosting service providers. These offerings are typically highly specialized and isolated.

Ultimately, the evolution will move to a hybrid cloud model where a single application can span both private and public clouds and is managed in a federated manner through a single "pane of glass." For this to happen, there need to be standards for interoperability and portability, and there needs to be technology to support such interoperability. The notion of "cloud bursting" is very compelling to CIOs. This is when an application's steady-state workload runs on the in-house private cloud, on hardware and data centers owned by the enterprise. But when there's a peak in the workload, it can dynamically burst out to a public cloud and take advantage of that capacity. When the peak is over, it can return that capacity back to the pool and shed that cost.

Organizations will evolve through these basic stages at different rates of speed, and organization will occupy several stages all at once. In other words, organizations will leave some stable legacy applications in a siloed mode. They will consolidate other applications to a virtualized, grid environment. And for some types of applications, they will move to a full self-service private cloud and ultimately to a hybrid cloud model. All of this will take time, as we are still in the early days of cloud computing.

Oracle PaaS Platform Components

The following describes the various components of the Oracle PaaS Platform from the bottoms-up. Refer to Figure 1.

Storage

Sun's Open Storage products combine open-source software with industry-standard hardware to deliver an open, scalable storage platform that can be optimized to provide public or private cloud storage services at scale. Indeed, Sun's ground-breaking Sun Fire X4500 Server helped the industry see the benefits of combining server and storage technology in the same system. Sun delivers virtual networking for large-scale computing through InfiniBand to massive-scale compute grids with the Sun Datacenter Switch 3456, scaling up to 13,834 nodes.

Providing a glimpse of what synergies are possible when software and hardware are designed together, *Oracle Exadata Storage Servers* also feature software intelligence, but with a particular affinity with the Oracle Database. Oracle Exadata Storage Servers combine Oracle's smart storage software and Oracle's industry-standard Sun hardware to deliver the industry's highest performance database storage.

Oracle Exadata Storage Servers use a massively parallel architecture to dramatically increase data bandwidth between the database server and storage. In addition, smart storage software offloads data-intensive query processing from Oracle Database 11g servers and does the query processing closer to the data. The result is faster parallel data processing and less data movement through higher bandwidth connections. This massively parallel architecture also offers linear scalability and mission-critical reliability.

Servers and Operating Systems

The modern enterprise data center is a highly optimized, complicated engine that drives and defines how a business functions. The most popular vision of cloud computing has everything running on farms of blade servers. While this vision is becoming more and more viable, there remains broad demand for severs of varying sizes and configurations that are optimal for certain particular application requirements.

Oracle offers Sun server product lines (Netra, Blade Servers, SPARC Enterprise, X64, et al.) and industry standard operating system support (Solaris, Linux, Windows), which will continue to provide a wide range of physical infrastructure optimized for the highly virtualized and distributed nature of a cloud implementations.

For operating systems that power the cloud, there is a fundamental requirement of enterprise class security, performance and reliability. Oracle's has a long track record of investment and leadership in the development and support of Linux with *Oracle Enterprise Linux* and *Unbreakable Linux Support*. Oracle has also supported the Linux community by providing code for a clustered file system, data integrity, optimized IO and virtualization that has been accepted into the main line Linux kernel. Oracle remains focused on continuing to enhance and extend the enterprise-class capabilities of Linux.

The added dimension of offering the best Unix technology with *Oracle Solaris* provides numerous key benefits for those wanting to build cloud infrastructure for IaaS or PaaS. The many Solaris innovations in networking, security, clustering and failover, file system, virtualization and self-diagnostics directly contribute to a high performance cloud computing implementation.

Virtualization

The combination of Sun and Oracle's virtualization products offers an unparalleled portfolio that enables users to virtualize from the desktop to the data center.

Server Virtualization

Virtualization, by isolating the software from the hardware, permits rapid software deployment with minimal to no required physical hardware provisioning, thus dramatically reducing the time necessary to applications get up and running. The key enabling characteristics of virtualization are low system overhead for optimum performance and integrated management capabilities in order to deploy cloud application components quickly and flawlessly.

Oracle VM offers support for both x86 and SPARC architectures, making it possible to deploy in a heterogeneous environment. Customers can leverage Oracle VM to consolidate servers, rapidly deploy software, recover quickly from system failure, and match resource capacity to workloads.

Oracle now offers Sun's virtualization products under the Oracle VM brand.

Oracle VM Server for SPARC supersedes Solaris Logical Domains (LDoms) and continues to provide highly efficient hypervisor for Sun SPARC *Chip Multithreading* (CMT) servers and enables multiple, independent Solaris instances on a single server. *Oracle VM Server for SPARC* is still included as part of Oracle Solaris. Oracle also supports Solaris as a guest operating system on top of *Oracle VM Server for x86*.

Oracle will also continue investment in Solaris Containers technology which is available for all Solaris across SPARC and x86. The virtualization capabilities of Solaris Containers enable customers to safely consolidate multiple Solaris applications onto one system and increase utilization rates with the advanced functionality of Solaris 10. This enables customers to host up to thousands of applications on a single system, with a single OS instance. This hardware independent virtualization can dynamically adjust to business goals and uses less than 1% system overhead. Solaris Containers can start and restart in seconds and are easy to create, replicate, rename and clone which can greatly simplify and accelerate cloud administration.

System administrators can use the cloning capability to quickly provision a new Container based on changes in workload demand or type. They can also move Containers from system to system or disk to disk within the same system as capacity or configuration needs change using the attach/detach features.

VM Templates and Assemblies

Rapid application deployment in cloud environments requires highly automated, mature application packaging and provisioning capabilities. There is a need to quickly manipulate applications, to deploy, grow, shrink, move, and clone them.

Oracle VM Templates are virtual machine images containing pre-installed and pre-configured enterprise software that can be used to develop, package and distribute applications for faster deployment. Oracle VM Templates can speed and simplify application deployments and help reduce the risk of errors in production, development, or test environments. Each VM Template is essentially a software appliance because just like hardware appliances, they are pre-built and very easy to deploy.

The next level of this type of application packaging is the concept of *VM Assemblies*. While software appliances are useful, enterprise applications are not always self-contained, single-VM entities but are sometimes complex, multi-tier applications spanning multiple VMs. There might be multiple VMs in the Web Tier, other VMs in the middle tier, and other VMs in the database tier. There needs to be a way for these multi-VM applications to be packaged for easy deployment.

Oracle Virtual Assembly Builder is a tool that takes such a multi-tier, distributed application and packages it up into an assembly that can be reused in a way similar to the way appliances are used. The assembly, like an appliance virtual image, is essentially a file that contains the images of the constituent appliances as well as metadata about appliance configuration, connections, and start up sequence. This technology will be a critical element for creating a library of applications and shared services in either public or private cloud environments.

Database and Storage Grid

On top of Oracle's IaaS are the higher level software components that are needed to build modern applications in cloud environments. A highly optimized PaaS layer requires high performance database and middleware services delivered as quickly and seamlessly as physical and virtual cloud infrastructure.

Oracle Database has offered grid computing capabilities since the release of Oracle Database 10g in 2003. Since then, Oracle has made continued to enhance the grid capabilities of the database in the areas of clustering with *Oracle Real Application Clusters (RAC)*, storage virtualization and manageability with *Automatic Storage Management (ASM)* and database performance with *In-Memory Database Cache*. When lighter weight database services are needed, *Oracle Berkeley DB* and *MySQL* are also possible options that are actively developed and supported by Oracle.

Now fully ten years after the introduction of Oracle RAC, none of the other database vendors offer a database product with equivalent capabilities such as predictable performance, scalability, availability in a database cluster. RAC is unique in database technologies in its support of grid architecture and mechanisms, particularly with its ability to adjust dynamically add and remove nodes. Not only does this support the elastic capacity required by private cloud, it substantially enhances ongoing management tasks such as upgrades, patches, and migration without compromising business continuity.

Oracle RAC is supported on all Oracle operating systems and virtualization platforms and features online addition and removal of cluster nodes without needing to halt the database or the application. The ability to run database (and as we will see below, middleware) clusters on either physical or virtual cluster nodes gives the ultimate in choice of deployment options in a cloud environment.

Another new RAC innovation that is relevant to cloud implementations is *Oracle RAC One Node*. Oracle RAC One Node is a one-node version of Oracle RAC which enables customers to standardize on a single deployment model for all their database needs. It enables the consolidation of multiple databases into a single cluster with minimal overhead while providing the high availability benefits of failover protection, rolling upgrades, et al.

Automatic Storage Management creates a grid of storage for a database that can also be adjusted dynamically and automatically. It has become the most common method of deploying storage under Oracle RAC and provides significant manageability and productivity gains for DBAs. This simplified administration is also ideal for those wishing to offer database services in a cloud.

Oracle In-Memory Database Cache enables certain tables, rows and columns from Oracle Database to be cached in the memory of the middle tier servers, delivering very low latency and high throughput. Data remains synchronized with Oracle Database and is accessed through a standard SQL interface. Oracle In-Memory Database Cache also supports clustering for elastic scalability and high availability.

Application Grid

Of the key characteristics of cloud computing identified earlier, elastic capacity is often a high priority for many applications. Having an infrastructure automatically adapt to the evolving needs of various applications means the IT shop can do more with less: it can provide better performance against application SLAs across the enterprise without having each application individually provisioned for its own worst-case workload.

Similar to the grid architecture in Oracle Database and storage, Oracle Fusion Middleware also supports robust grid functionality in the middle tier with a group of products called Oracle *application grid*. Like database grid, application grid enables a pool of middle tier servers to operate like a single large server. Capacity can be added and removed by simply adding nodes on an as needed basis, and the whole cluster is highly available.

The key technologies that make up Oracle's application grid are *Oracle WebLogic Server* as the flagship application server; *Oracle Coherence* providing in-memory data grid services, *JRockit JVM* providing a lightweight, lightning fast Java runtime environments; and transaction monitoring and management with *Oracle Tuxedo*.

Oracle WebLogic Server is the core Java EE application server technology within Oracle Fusion Middleware and the overall market leader among application servers. Its clustering capabilities, with support for automated load balancing and failover as well as dynamic addition and removal of nodes, serves as the fundamental mechanism for adjusting capacity. WebLogic Server clustering supports both console-based and script-based automation of capacity adjustment, providing flexibility for a wide range of needs while maximizing ease of use. WebLogic Server clustering management easily plugs

into the Oracle Enterprise Manager framework, enabling unified management of the entire private cloud infrastructure from a single console.

WebLogic Server is complemented and enhanced by Oracle Coherence for an in-memory data grid cache for Java, .NET and C++ objects. Coherence significantly increases performance, reliability, and scalability of data-intensive transaction processing or analytical applications by caching data—potentially very large data amounts of data—in memory, accessing that data at memory speeds rather than disk access speeds. This provides the obvious benefit of performance improvement due to the memory access speed, but has additional performance benefits based on the ability to parallelize computation across the data grid as well reliability benefits deriving from the way objects are replicated in the grid, eliminating single points of failure. Coherence automatically "repartitions" or redistributes data objects optimally across the data grid as nodes are added to or removed from the grid, again supporting the elastic capacity requirement of application grid and cloud. Management of this mechanism also plugs cleanly into Oracle Enterprise Manager.

WebLogic Server and Coherence are packaged with additional application grid components such as the JRockit Java Virtual Machine (JVM) in an offering called WebLogic Suite. With the automatable, dynamic clustering mechanisms of WebLogic Server and Coherence managed by Enterprise Manager, consolidation of applications onto a centralized, shared application grid is one of the key steps on the path to private cloud computing.

SOA and Business Process Management (BPM)

With a basic mechanism for elastic capacity in place such as application grid, the next consideration in setting up a private cloud that is not just sharing infrastructure but also sharing services. Service-oriented architecture (SOA), with its approach of modularizing applications into reusable components accessible through standardized interfaces using XML, SOAP, and the various WS-* specifications, is the obvious starting point.

Oracle SOA Suite provides a comprehensive yet easy-to-use basis for creating the reusable components at the heart of your PaaS private cloud. Rich drag-and-drop SOA component features in JDeveloper and the SCA designer enable rapid creation of components and subsequent composition of those components into applications. Oracle Service Bus provides a simple way to make components available to department application creators using your PaaS cloud. End-to-end instance tracking and Oracle Business Activity Monitoring provide a range of metrics visualizations supporting both the central IT function charged with keeping the PaaS up and running and the departmental application owners concerned with business-level performance indicators.

In addition to SOA components, many enterprises will want to include business process components managed within a unified BPM framework as part of their PaaS. Oracle BPEL Process Manager provides the federation capability to create BPEL process components out of new as well as legacy assets while also supporting the flexibility to enable multiple departments to incorporate PaaS-based BPEL components into their respective workflows.

User Interaction

Like SOA and BPM components, standard and shared user interface (UI) components are great candidates for inclusion in an enterprise's PaaS. A centrally-managed library of UI components can give department application owners a great head start in composing their solutions and also gives the central IT function a desirable level of control over consistency across the enterprise's UIs. At the same time, a robust, standards-based UI framework can give the departments the flexibility they need to accommodate their specific functionality, customization, and personalization needs for applications and portal solutions.

UI technologies play an additional role in a PaaS environment as the basis of the self-service interface for the cloud. In many cases this will be a fairly extensive portal that must work closely with an identity management system to authenticate users, filter their access based on roles, and present the platform's shared components for application development and composition.

Oracle WebCenter Suite provides a number of portal and user interaction capabilities that are ideal for creating reusable UI components as part of a PaaS. Themes and skins provide powerful facilities for tailoring the look and feel of applications in a tiered way—for an entire Web interface or for portions of a Web interface associated with a department. This enables consistency in look and feel while consolidating deployment. The Advanced Personalization Framework provides the ability to further tailor the usage of the UI and the information delivered to the UI based on users' activities. Powerful mashup integration capabilities enable business users to further personalize the information they want to see while maintaining enterprise information security. Common enterprise metadata services provide a revolutionary way to store and manage all look-and-feel changes, personalizations, and mashups via uniform metadata that enables in-place customizations at runtime and insulates the UI from changes to the base application.

Identity Management

A high-priority concern for many enterprises in creating a private cloud is identity and access management. Particularly for organizations in domains with a high level of regulation and/or sensitive customer data, cloud's self service can be a significant challenge. Balancing rich mechanisms for identity and access management with convenience features such as single sign-on is a must for cloud environments.

Implementing PaaS with a high degree of self service in a security-critical environment requires an approach where security pervades the entire architecture rather than being bolted on as an afterthought. An important strength of Oracle Fusion Middleware is that, in addition to each of the products having best-of-breed security in their respective categories, their security mechanisms are well integrated, enabling ease of deployment, ease of change, and high reliability.

Oracle Identity and Access Management Suite provides an ideal facility for managing access and security in a PaaS environment. Within the suite, *Oracle Access Manager* supports corporate directories and single sign-on. *Oracle Entitlements Server* provides centralized access control policies for a highly decentralized PaaS environment.

Oracle Identity Manager is a best-in-class user provisioning and administration solution that automates the process of adding, updating, and deleting user accounts from applications and directories. It improves regulatory compliance by providing granular reports that attest to who has access to what. Oracle Identity Federation provides a self-contained and flexible multi-protocol federation server that can be rapidly deployed with your existing identity and access management systems. With its support for leading standards-based protocols, it ensures the interoperability to securely share identities across vendors, customers, and business partners without the increased costs of managing, maintaining, and administering additional identities and credentials.

Cloud Management

Cloud computing succeeds or fails based on the quality of its systems management. Enterprise Private PaaS requires highly sophisticated automation in order to manage the vast amount of computing power, huge data sets and highly virtualized IT services. Oracle Enterprise Manager provides full lifecycle management for cloud computing and manages the complete Oracle stack from applications to disk.

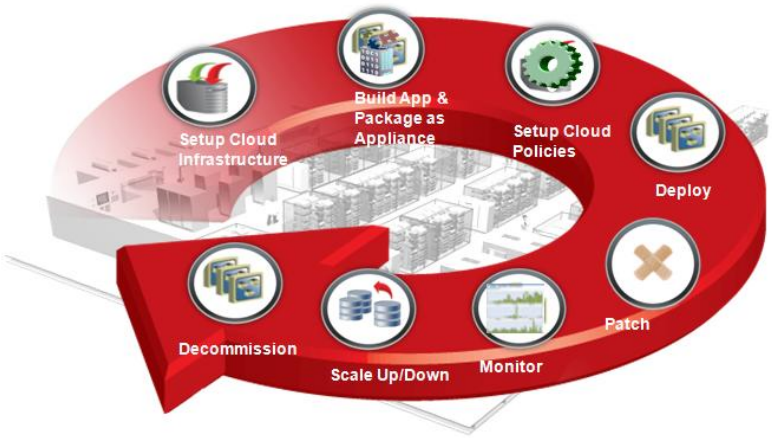

Figure 3. Cloud Management Lifecycle

Setup Cloud: First of all, when deploying a virtual environment, Enterprise Manager enables easy provisioning Oracle VM server software on bare metal machines. It simplifies the creation of virtual server pools as well as the configuration of these pools to be deployed in a highly available manner. Once this setup has been done, administrators can use Enterprise Manager to simplify the creation and registration of virtual machines so that they can be centrally monitored and managed. A rich set of options is available for the creation and provisioning of guest virtual machines – including Oracle VM templates, ISO images and PXE bootable guest virtual machines that boot over a network.

For applications deployed in physical environments, Enterprise Manger automates the bare metal deployment of operating system and additional software in an efficient manner. Administrators can provision the complete stack on bare metal server to simplify and accelerate the deployment cycle.

Build and Package as Application: Second, Oracle makes it easier to create application appliances or multi-tier assemblies that can be deployed into a cloud environment. The Oracle VM Template Builder and Oracle Virtual Assembly Builder tools offer easy ways to create, package and distribute the complete application environment which can easily be imported and deployed into a cloud environment. Deployment time for applications can be reduced significantly. Once imported into a cloud environment, these applications can be used as gold images to clone them to create multiple instances of the application running in the cloud.

Setup Cloud Policies: Cloud administrators can set up policies using Enterprise Manager that dictate how the VM and its applications are placed and started within the server pool in the cloud. Different policies can dictate whether this placement on a server happens automatically based on algorithms that determine which server has the best resources available or based on user-defined policies, or even based on a preferred-server policy that specifies a named subset of servers that must be used to host a particular application and its virtual machine instance.

Deploy: Oracle Enterprise Manager offers a very rich set of capabilities for deploying not just the cloud infrastructure but also applications that run in the cloud. Oracle Enterprise Manager can be used for automating the deployment of a wide variety of software packages including Oracle Database, Oracle Fusion Middleware, and both packaged and custom applications. It supports a variety of mechanisms for deploying including both cloning, fresh installs, and template based creations. A customizable framework for deployment procedures is available to adapt the out-of-box mechanisms to your private cloud or develop completely new procedures useful for any enterprise.

Patch: Enterprise Manager provides comprehensive patching capabilities that include understanding the customer configuration, advising customers on what patches are available and should be applied, planning patch deployment and testing in an highly automated and scalable manner, and verifying that patches were deployed in compliance with IT procedures and processes. Patching automation is available for the underlying VM and operating systems as well as the database, middleware and applications that run in the cloud. Patching of HA environments can be done in a rolling manner as well to ensure high quality of service.

Monitor: Once the applications are deployed to the cloud, it is critical to monitor its functions on an ongoing basis. Enterprise Manager offers rich features to monitor not only the cloud infrastructure – the hardware, the VM, and operating system but also the applications deployed in the cloud, including the end user experience and service levels.

Scale Up/Down: During the operation of these applications deployed to the cloud, unpredictable demand patterns may require automatic scaling up or scaling down of resources for a given application. Enterprise Managers simplifies the way users can automatically add or remove capacity to support these workloads. Some of the mechanisms include – cloning and live migration of virtual machines, as well as adding/removing nodes to clusters of RAC, In-Memory Database Cache, WebLogic and Coherence.

Decommission: Finally, once an application is not being used any longer or has been scaled back, Enterprise Manager can be used to uninstall the application so that the resource can be used for other purposes.

Summary

Oracle aims to make cloud computing fully enterprise-grade and supports both public and private cloud computing to give customers choice. Oracle provides the industry's most complete, open and integrated portfolio of products to build, deploy and manage public and private clouds.

Disclaimer

The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle’s products remains at the sole discretion of Oracle.

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

Oracle Cloud Computing
May 2010
Authors: George Demarest, Rex Wang

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com