IDC
Analyze the Future

IDC TECHNOLOGY SPOTLIGHT

The Benefits and Significance of Private Platform as a Service

June 2013

Adapted from *The Pressures That Are Reshaping Platform as a Service* by Stephen D. Hendrick, IDC #240993

Sponsored by Red Hat

The dynamics of the platform-as-a-service (PaaS) market indicate that there is strong demand for both public and private PaaS. Demand for public PaaS is strong because of accessibility, affordability, and application utility needs. Demand for private PaaS is also strong because of security, sovereignty, and application utility requirements. Vendors that embrace the concept of public and private PaaS are also in favor of hybrid PaaS models where workloads can be directed to either public or private instances depending on how an enterprise sets application policy. Hybrid models provide the most flexibility where the private and public PaaS components are the same or have been specifically designed to work together. Vendors that support hybrid models will be the most successful because they provide an enterprise with the most flexibility. In addition, it is noted that open source platforms and open standards are very effective at eliminating enterprise concerns over lock-in. IDC has identified a number of benefits that accrue when using private PaaS, including service enablement, infrastructural independence, hybrid technology models, performance, security, and efficiency. This Technology Spotlight explores these trends and the role that Red Hat's OpenShift Enterprise plays in this strategic emerging market.

Introduction

Two profound evolutionary changes have occurred in application development and deployment (AD&D) over the past 25 years. The most significant change was the transition from developing every aspect of an application to the pervasive use of runtime components to deploy and manage applications. The core capabilities of deployment platforms include application deployment and management. Key issues that must be addressed include the ability to install or uninstall an application and then ensure that the application adheres to intended service levels while in production.

The benefits of deployment-centric development stem from the speed, reliability, and manageability that come from leveraging prebuilt runtime components. The downside from using prebuilt components is primarily vendor lock-in and the cost of acquiring, maintaining, and using these components. However, because enterprises are so focused today on process automation and the continuous delivery of application functionality, the advantages of platform technologies outweigh their concerns — especially when these concerns can be mitigated through the use of open source platforms.

Although the transition from development to deployment was the first and most significant change, the second important change is the transition from deployable software to software provided as a service. This is true even in the AD&D market, where even the application platform is being provided as a service. There is a transition from platforms that are managed by an enterprise to platforms that are highly configurable and largely self-managing and consumed as a service. PaaS products provide access, security, application hosting, session management, load balancing, scalability, and manageability. In some cases, PaaS products also provide support for application development as well as processes associated with life cycle, data management, and integration.

Platform-as-a-Service Market Trends

Vendor response to PaaS has been overwhelmingly positive. Nearly every leading AD&D software vendor has a public PaaS in the market, and those that don't will be making related announcements later this year. Private PaaS has taken a back seat to public PaaS because the design point for private PaaS is more demanding. This may sound counterintuitive, but delivering configuration-based deployment, scalability, multitenancy, and manageability into a private cloud environment with the same user experience as public PaaS is extremely challenging. In a public PaaS, users are shielded from how the vendor maintains the PaaS environment and its underlying infrastructural resources, which are abstracted away.

In a public cloud, as long as customers are being given a reliable service within stated SLAs, they are largely indifferent to what the service provider must do behind the scenes to deliver these services. However, in a private cloud, the presumption is that the PaaS services running on a virtualized infrastructure will be essentially self-managing, much like their experience in the public cloud. Most vendors simply have not had the time to retool enough to evolve their products to this state.

Many major cloud providers have no stated intent to provide their platform on a private cloud. This seems to fly in the face of IDC's 2012 survey data that suggests that much of the IT in large enterprises will not migrate applications to the public cloud either because of their legacy status or because their mission-critical characteristics mean that enterprises cannot afford to entrust any aspect of these capabilities to anyone else. Leading infrastructure-as-a-service (IaaS) providers do provide virtual private clouds and VPNs that attempt to bridge the gap by providing higher levels of hardware and network isolation. However, one major service provider correctly points out that that the scale at which it and other leading public platform vendors operate can be achieved only by highly optimizing the underlying infrastructure and providing security well beyond what is typically found in an enterprise datacenter. For this reason, some public clouds will eventually provide numerous benefits related to scalability, high availability, disaster recovery, and security that no private cloud will be able to match.

However, in the short run, concerns related to security, data sovereignty, governance, risk, and compliance are taking precedence over economies of size and scale, and interest in private PaaS is equal to and in some cases greater than interest in public PaaS. Security was the leading concern with public clouds, as reported in an IDC survey of 800 cloud-savvy IT professionals in 2012. It was also the third year in a row that security was the leading public cloud concern.

It is hard to imagine that public PaaS vendors will easily address the security issue. To keep prices down, PaaS service providers need to maintain high application densities on servers and at the same time distribute application instances over multiple servers to ensure reliability. The result is that, in a public PaaS, an enterprise's application and data are likely to reside on heavily shared resources, which at worst could compromise security and at best exposes an enterprise to noisy neighbors.

The solution to this problem lies in isolation. If you can ensure isolation between virtual machines (VMs) and across regions within a VM, then you effectively solve the security problem without impacting application densities. However, this problem is not easy to solve and tends to require a combination of platform and operating system modifications.

The dynamics of the PaaS market indicate that there is strong demand for both public and private PaaS. Demand for public PaaS is strong because of accessibility, affordability, and application utility needs. Demand for private PaaS is also strong because of security, sovereignty, and application utility requirements.

Vendors that embrace the concept of public and private PaaS are also in favor of hybrid PaaS models where workloads can be directed to either public or private instances depending on how an enterprise sets application policy. Hybrid models provide the most flexibility where the private and public PaaS components are the same or have been specifically designed to work together. In fact, we believe that vendors that support hybrid models will be the most successful because they provide enterprises with the most flexibility.

Another decision that enterprises will need to make regarding PaaS is whether to choose an open source or closed source platform. The decision regarding open source versus closed source is relevant only to the private PaaS market. Enterprises are licensed to use a public PaaS, but they don't own the PaaS. The degree of private PaaS "ownership" can vary widely depending upon how a vendor chooses to license it. However, open source private PaaS licensing tends to follow the traditional open source model that emphasizes just maintenance and support.

Open source platforms and open standards are very effective at eliminating enterprise concerns over lock-in. This is one of the two most appealing characteristics of open source products. While open source products have the advantage of soliciting community involvement, vendors that pursue open source products are generally seeking to deliver less differentiated products that exhibit cost leadership. This clearly contrasts with feature-laden highly priced products that claim to be well differentiated. This philosophy is responsible for the other appealing characteristic of open source platforms: cost leadership.

The Benefits of Private Platform as a Service

IDC has identified a number of benefits that accrue when enterprises use private PaaS:

- **Service enablement.** Private PaaS consists of platform technologies delivered as a service. Service enablement of these platform technologies means that many capabilities have high levels of abstraction and automation. For example, private PaaS products should excel at deployment, security, availability, multitenancy, dynamic scalability, integration, and management. These features should be easily configurable and abstract away infrastructural decisions that traditionally reside further down the stack. The benefit of service enablement is that complex application hosting activities, including deployment, availability, and scalability, can be configured and implemented in minutes instead of hours, days, or weeks (depending upon infrastructure availability). This accelerates time to market and enables enterprises to focus more on process development and deployment instead of how to address hosting needs.

- **Infrastructural independence.** Private PaaS is designed to run on a collection of servers. These servers are typically virtualized, but this does not always have to be the case. Therefore, the infrastructural bar is set fairly low, so the servers can reside anywhere within a VPN. This means that some of the servers could be in the enterprise's datacenter and some could reside in a third-party virtual private cloud environment. The advantage is that enterprises have the flexibility to make infrastructural capex or opex decisions without having to worry about the impact on systems running on the private PaaS.

- **Hybrid technology models.** Where a vendor delivers both private PaaS and public PaaS, there is an opportunity for providing a compelling hybrid PaaS. This is because most vendors strive to have the usage metaphor for the public and private versions of their PaaS nearly identical, which allows workloads to span both environments (contingent on policy). This facilitates bursting and can protect against noisy neighbors.

- **Performance.** The hallmark of PaaS is automated provisioning and dynamic scalability, both of which apply to the more advanced public and private PaaS offerings. Automated provisioning shrinks the time needed to implement an application instance from days or weeks to seconds or minutes. Dynamic scalability monitors workload against capacity and provisions or deprovisions instances to ensure that application SLAs are always met. Also included under performance are availability and the ability of more advanced public and private PaaS offerings to distribute workload across available instances in ways that optimize performance and enable true 24 x 7 x 365 availability.
- **Security.** Security encompasses identity and access as well as isolation. Private PaaS identity and access management (IAM) services provide identity, access, and single sign-on (SSO) services that are typically role based and designed to integrate easily with LDAP directory services (or an equivalent). However, the more advanced private PaaS products provide much higher levels of control over how servers are virtualized, segmented, and isolated. The challenge is that security requires isolation and isolation works against the desire to increase application density. Achieving high application densities without compromising isolation is best accomplished through the careful coordination of controls within a security-enhanced operating system.
- **Efficiency.** Private PaaS delivers efficiency to IT activities in a variety of ways. Advanced private PaaS offerings provide development environments, frameworks, or life-cycle integrations that provide efficient ways to build applications. Private PaaS automated provisioning, configuration-based availability, and dynamic scalability together eliminate most of the heavy lifting required when deploying an application and managing it after it is in production. The concatenation of these development and deployment services enables continuous delivery of application functionality and allows the enterprise to focus on a more streamlined approach to application development and deployment. As a result, enterprises can reduce time to market, cycle applications faster, and deliver applications that have much higher levels of quality, reliability, and availability.

Red Hat's Approach to Private Platform as a Service

Red Hat is best known as an open source software vendor that supports the Red Hat Enterprise Linux (RHEL) operating system. However, a variety of acquisitions and significant R&D spending have enabled Red Hat to be recognized as one of the leaders in enterprise application platforms (with the JBoss Middleware family) and cloud computing technologies (OpenStack and CloudForms) including PaaS (OpenShift).

Red Hat launched OpenShift Online, its public PaaS, in 1Q 2011. This was a controlled release with generous terms to encourage developer use and feedback. In 1Q 2012, Red Hat "open sourced" the technology behind OpenShift in the form of OpenShift Origin under the Apache 2 open source license. The commercial release of OpenShift Online occurred in 2Q 2013. The experience that Red Hat gained from OpenShift Online was instrumental in the development of OpenShift Enterprise, and now Red Hat has delivered on a strong hybrid PaaS strategy with both public and private PaaS offerings.

Overview

OpenShift Enterprise is Red Hat's private PaaS that was launched in 4Q 2012. It is a multilanguage, auto-scaling, self-service, and elastic cloud application platform that is built on RHEL, which is essentially the only infrastructural requirement. A key subsystem of RHEL that is leveraged in the multitenant architecture of OpenShift is the Security-Enhanced Linux (SELinux) subsystem developed by Red Hat and the National Security Agency (NSA). Red Hat describes SELinux as follows:

> SELinux is an implementation of a *mandatory access control* mechanism in the Linux kernel, checking for allowed operations after standard *discretionary access controls* are checked. It was created by the National Security Agency and can enforce rules on files and processes in a Linux system, and on their actions, based on defined policies.

Discretionary access control (DAC) is inadequate if security is paramount because DAC access decisions are based only on user identity and ownership, ignoring other security-relevant information such as the role of the user, the function and trustworthiness of the program, and the sensitivity and integrity of the data.

In SELinux (and RHEL), all processes and files are labeled with a type. The type defines a domain for processes and files. Each domain and its processes run independently, and policy rules define how processes in the domain interact with each other and with files. Access is allowed to a domain only if there is a specific rule that allows it. This goes beyond the DAC that is part of traditional Unix and enables the fine-grained access control that allows OpenShift Enterprise to ensure process and file isolation.

OpenShift Enterprise runs on one or more instances of RHEL, each of which is termed by Red Hat as a "node." Each node can be segmented into secure containers, each of which runs an instance of a user application. These secure containers are termed by Red Hat as "gears." Consequently, the node and gear constructs provide OpenShift Enterprise with the ability to segment a RHEL instance into as many isolated containers as are practical given the available memory and overhead associated with managing the operation of these containers.

OpenShift Enterprise coordinates the operation of these gears through the use of a broker. Each node has an OpenShift Enterprise agent running on it that communicates with the broker. The function of the broker is to make decisions about how to best provision and scale end-user applications based on existing demand and capacity.

Today, decisions regarding spinning up additional nodes to add capacity to the PaaS are vested with datacenter operations. Part of this rationale is based on the desire of enterprises to retain policy control over how to size gears and how many gears to permit on a node. This means that while the act of deploying another instance is at the discretion of datacenter operations, once the instance is deployed, OpenShift Enterprise (with the support of JBoss Enterprise Application Platform [EAP]) will optimize how application capacity is leveraged to support application demand.

Red Hat also realized in the development of OpenShift that many languages, platforms, and databases drive enterprise applications. Therefore, the company needed another abstraction layer that would enable support for the heterogeneous mix of assets that reside in the IT domain.

Red Hat uses the term "cartridge" to define how OpenShift Enterprise installs languages and middleware products into a gear. A cartridge maps the API of a specific language or middleware component into the corresponding OpenShift Enterprise API class for a language, platform, or database. As of 2Q 2013, Red Hat provides the following cartridges for OpenShift Enterprise:

- Languages: Java, PHP, Python, Ruby, and Perl
- Databases: PostgreSQL, MySQL
- Platforms: JBoss EAP 6 and JBoss Enterprise Web Server (EWS)
- Other: Custom

The custom cartridge provides a do-it-yourself cartridge for developers to add their own language, database, or middleware component.

OpenShift Enterprise primarily focuses on providing support for deployment and management of production applications. OpenShift Enterprise approaches application life-cycle management in a traditional open source way. Developers build application code, quite possibly using a Git Repo (or equivalent) to coordinate team development. Maven can be used for builds and Jenkins for continuous integration. The ready availability of these tools makes them a classic choice for use with Red Hat products and services. Red Hat supplies enough cartridges to ensure that all of the tiers of a Web application (database, application server, and Web server) can be deployed on OpenShift Enterprise.

Benefits

Red Hat is well known for RHEL and its open source business model for RHEL support. The company is now finding that its reputation is effective leverage to open the enterprise middleware door. Many enterprises have adopted Red Hat's JBoss EAP because it is a full-function modular application server that competes effectively because of its open source roots and lower price points. This halo effect will also serve Red Hat well as it enters the PaaS market.

While many Red Hat customers are quick to mention price as a key factor in their decision to adopt Red Hat middleware products and services, others, such as large enterprises, cannot afford to rely on products that are not industrial strength. This enables Red Hat to compete effectively on features while maintaining a cost leadership position.

Red Hat is the first of the major application server vendors to provide both a public PaaS and a private PaaS. Because of the strong enterprise interest in hybrid cloud models, Red Hat will have first-mover advantage in the high-growth PaaS market.

Although most enterprises are not keen on modifying open source platform code, because of the upgrade and maintainability risks that are created, there is no doubt that Red Hat's open source policy provides these enterprises with a useful insurance policy that simply is not available from closed source vendors.

Challenges

The public and private PaaS markets reflect high levels of immaturity because of the retooling complexity that vendors face in providing features as a service along with new features such as multitenancy, automated provisioning, dynamic scalability, and security. Consequently, no vendor has yet provided a complete public and/or private PaaS offering.

Red Hat's acquisition of ManageIQ is a case in point. Red Hat already has a modular application server (JBoss EAP 6) that can be spun up in a few seconds. Therefore, why not offer the ability to have policy-driven application scalability? Naturally this makes perfect sense. However, because of the recent acquisition of ManageIQ (4Q 2012) and the joint objective of being able to leverage resources that reside in other IaaS environments, the engineering work to accomplish this is considerable and therefore takes time. Since Red Hat is an open source vendor, its business model and economic scale are very different from those of other IT systems suppliers. Consequently, R&D resources are always tight, which impacts time to delivery.

Another area where Red Hat needs to improve OpenShift Enterprise is application management. The current approach is to provide an API that will send KPI data to an enterprise's own monitoring and management tools. While this approach may be palatable to large enterprises in the short run, OpenShift Enterprise really needs its own management console that comes equipped to provide policy-driven monitoring and management of OpenShift Enterprise resources and applications.
An OpenShift Enterprise management console is a prerequisite to attracting small and medium-sized enterprises, and this point is not lost on Red Hat. The company already has JBoss Operations Network (JON), which is an effective management console, and is currently working to integrate it with OpenShift Enterprise.

Conclusion

Red Hat surprised the vendor community by delivering a public PaaS and a private PaaS in less than two years. It is also the only leading middleware vendor that currently provides both. The importance of delivering a hybrid PaaS alternative cannot be underestimated.

Most of the vendors with material market share in public PaaS have no intention of providing private PaaS because it's not part of their DNA. For three years in a row, enterprises have held the line on wanting private PaaS because of concerns about security. Some of these security concerns are simply due to poor messaging on the part of vendors, and some are legitimate.

Multitenancy creates a challenging environment that is difficult to secure. Even if organizations effectively address this multitenancy concern, they still may have noisy neighbors that compromise their resource needs. These issues, along with governance, risk, compliance, and sovereignty, are driving enterprise interest in private PaaS, and the situation is not likely to change anytime soon. IT is also under increasing pressure to do more with less. The pragmatic way to address this issue is to focus on products that exhibit cost leadership and enable developers to work more efficiently.

Red Hat's private PaaS does an effective job of mitigating security concerns while allowing enterprises to maintain high application densities and utilization rates and improve developer productivity. These characteristics enable OpenShift Enterprise to provide an optimal combination of security and efficiency at exactly the right time.

The appeal of PaaS is that middleware capabilities focused on application deployment and management are provided as a service, which means highly configurable instead of code centric. The architecture of OpenShift Enterprise, with its nodes, gears, and cartridges, is an effective foundation for delivering abstracted deployment services that have the right blend of deference to large enterprise needs while providing an efficient environment for deployment and ongoing operations.

OpenShift Enterprise provides a crucial capability demanded by enterprises today and has a clear road map that ensures that high levels of automation and management will find their way into the offering by the end of 2013. To the extent that Red Hat can address the challenges described in this paper, IDC believes that OpenShift Enterprise is well positioned for success.

ABOUT THIS PUBLICATION

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

For more information on IDC, visit www.idc.com. For more information on IDC GMS, visit www.idc.com/gms.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com