



WHITE PAPER

Preparing for Private Cloud and Hybrid IT with Red Hat Cloud Infrastructure

Sponsored by: Red Hat

Gary Chen
February 2015

IDC OPINION

Organizations are increasingly moving to private or hybrid cloud architectures. They are following the lead of cutting-edge Web companies that have built next-generation, scale-out infrastructures that are inherently different from traditional datacenters. It's important to note that change is happening not just in the infrastructure but also across the entire stack, from middleware to applications as well as the whole concept of how datacenters, IT, and people operate. A primary challenge in today's enterprise is how to adopt these new technologies and methodologies. Most organizations can't start completely anew because of the large investment required and, equally important, the continued reliance on existing infrastructure and traditional applications – many of which have existed for several technology generations. So while there will be pockets of completely new cloud computing platforms, there will also be an evolution of existing assets.

OpenStack, a rapidly evolving, next-generation open source cloud platform, is garnering enormous interest and momentum across the market, but it requires an entirely new approach, and applications need to be designed in a completely different manner to leverage its potential for massive scale. On the virtualization front, enterprises have grown their virtualization footprints over the past decade as vendors have made the technology secure and more than suitable for mission-critical workloads. Much of that footprint needs to be carried over into cloud, but using a more flexible approach, where new management technologies can be used in conjunction with virtual servers to provide end users the characteristics of cloud computing. These clouds can and will continue to run the traditional application portfolio. In addition, organizations are also beginning to use off-premise, public cloud resources. The public cloud represents yet another platform that needs to be governed and managed. It's not surprising then that enterprises will need a multi-faceted approach to move to the cloud era. More specifically, organizations will need to phase in existing infrastructure, build next-generation services, and integrate with public clouds all while reducing complexity.

Red Hat has developed a flexible, integrated solution that meets these requirements and helps organizations adopt a cloud infrastructure for both traditional and cloud-enabled workloads. Red Hat Cloud Infrastructure (RHCI) provides a datacenter virtualization platform for traditional applications and OpenStack for next-generation cloud applications, as well as a cloud management platform (CMP) to unify the management across both workload types, running on-premise or in public clouds. It is clear that enterprises will require a full solution to realize the many benefits of cloud while maintaining control and governance. RHCI is a prime example of a cloud solution that can help enterprises make the migration.

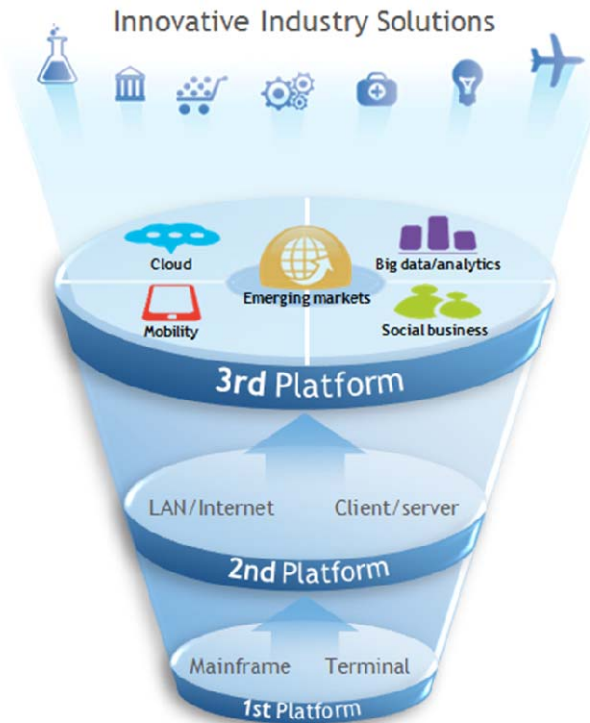
SITUATION OVERVIEW

The Rise of the 3rd Platform

The industry is transitioning to what IDC terms the 3rd Platform, being driven by new applications and service models such as cloud, mobile, social, and big data/analytics. This new era of computing is also on a different scale, hosting millions of applications and billions of users, compared with thousands of applications and millions of users for the 1st Platform, which was the mainframe era. The 2nd Platform is client/server infrastructure. The needs of the 3rd Platform are driving new requirements at every level of the stack, such as new scale-out cloud architectures for both applications and the underlying infrastructure (see Figure 1).

FIGURE 1

The 3rd Platform

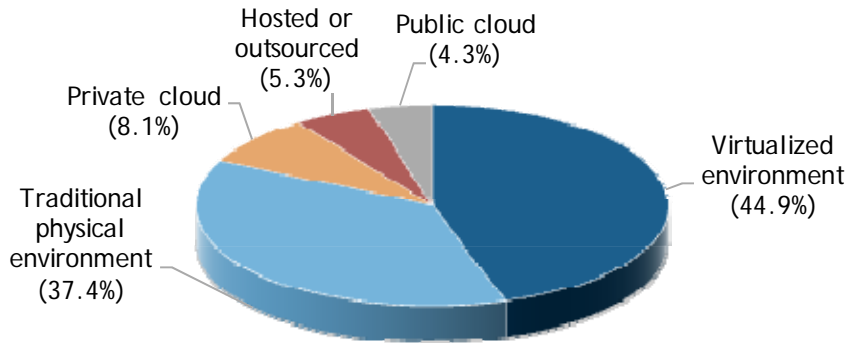


Source: IDC, 2015

While the 3rd Platform is growing, the long life cycles of enterprise applications have kept (and will continue to keep) 2nd Platform and even 1st Platform infrastructure around for a significant amount of time. Enterprises have made many significant investments in these applications, and it is not easy or cheap to migrate these applications to another platform. So over time, enterprises acquire many silos of technology across multiple computing eras and technologies. Figure 2 depicts the current distribution of workloads across these different infrastructures.

FIGURE 2

**A Diverse Range of Infrastructure for Workloads:
The Portion of Total Workloads Running on Traditional, Virtualized, Hosted,
Private, and Public Cloud Infrastructures**



n = 405 worldwide respondents

Source: IDC's *Virtualization and the Cloud MCS*, 2013

Virtualization Trends

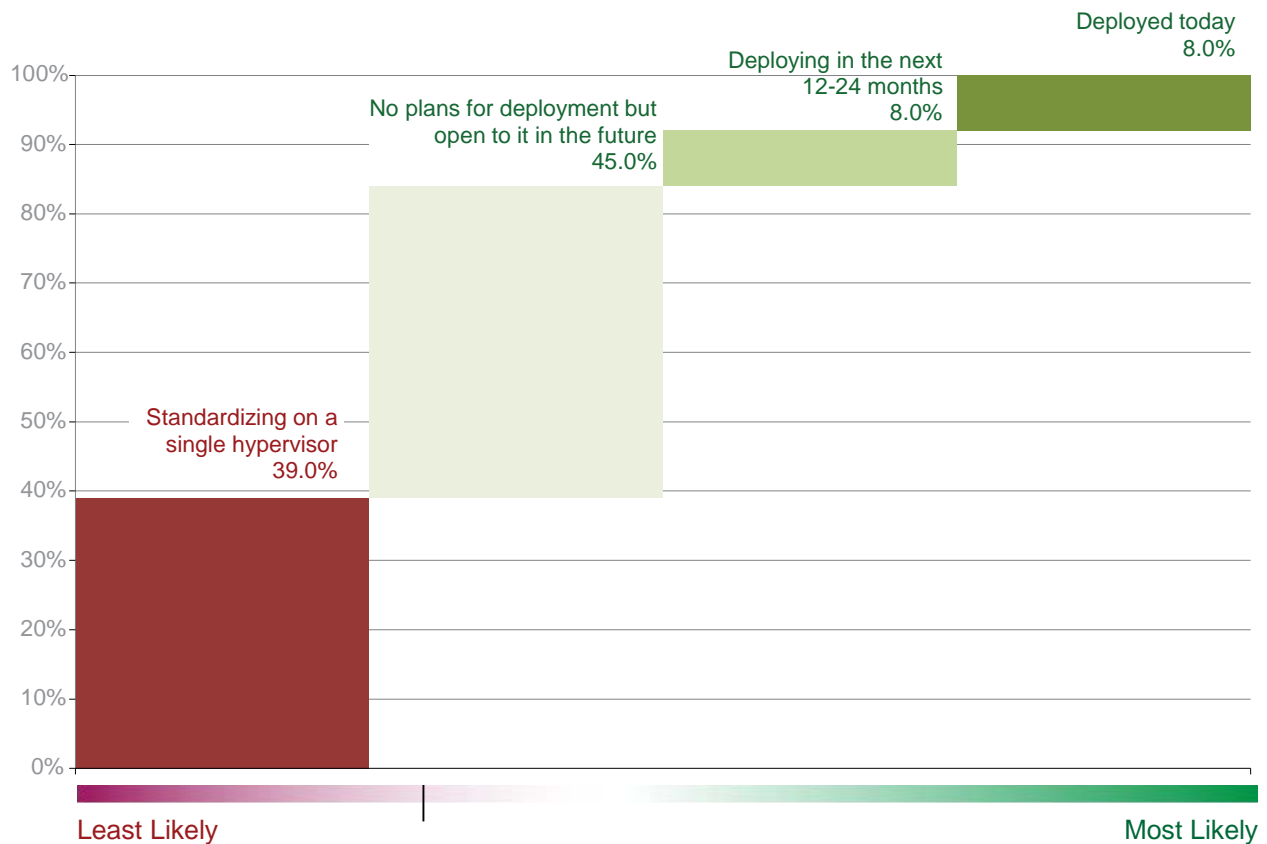
Virtualization is an important technology that started at the tail end of the 2nd Platform and remains a key foundation as the industry transitions to the 3rd Platform. Much of the 2nd Platform infrastructure has already been virtualized, a trend that began in the server market a little more than a decade ago. However, there remains a significant installed base of physical servers as well. While hypervisors are increasingly attached to private cloud systems, traditional server virtualization is still growing healthily as customers continue to deploy these systems to support 2nd Platform applications.

One of the newer developments in the virtualization market is the growing interest in multi-hypervisor deployments. Alternative hypervisors have improved significantly over the past several years, reaching enterprise-grade readiness and thus increasingly attracting the attention of enterprise buyers. It's still a relatively small market, with 16% of customers having deployed or planning to deploy multiple hypervisors. However, the opportunity is in the 45% of customers that don't have definite plans to deploy today but are open to deploying in the future. The hypervisors are ready today, but the key to multi-hypervisor success lies in the management. Adding a second hypervisor increases the complexity of the environment and also requires new skills. This has to be balanced against the benefits of licensing costs, reducing lock-in and commitments to a single vendor. There are pros and cons to both sides of a single hypervisor or multi-hypervisor, which will vary based on individual requirements and environments. The market is currently split on the issue, but clearly, there are potential customers out there that consider a multi-hypervisor approach the right solution. Any shift in the market to new hypervisors will happen with new workload deployments because migrating existing workloads from hypervisor to hypervisor can be difficult and expensive, making it cost prohibitive for workloads already successfully virtualized. Thus the ramp of alternative hypervisors, as depicted in Figure 3, will be slower and steady over time as new workloads are deployed.

Private cloud may be a tipping point for multiple hypervisor adoption. Often, these clouds are built with a completely different architecture to host a different set of applications than traditional virtualization, so backward compatibility with existing virtualized applications is not an issue. For many of these clouds, alternative hypervisors such as KVM are the norm. In addition, many of these new cloud platforms manage hypervisors in a different way than traditional virtualization management frameworks, allowing better abstraction and management of multiple hypervisor platforms, which can address the complexity problem of a diverse environment.

FIGURE 3

Plans to Support Multiple Hypervisors



n = 405 worldwide respondents
 Source: IDC's *Virtualization and the Cloud MCS*, 2013

The Growth of Private Cloud

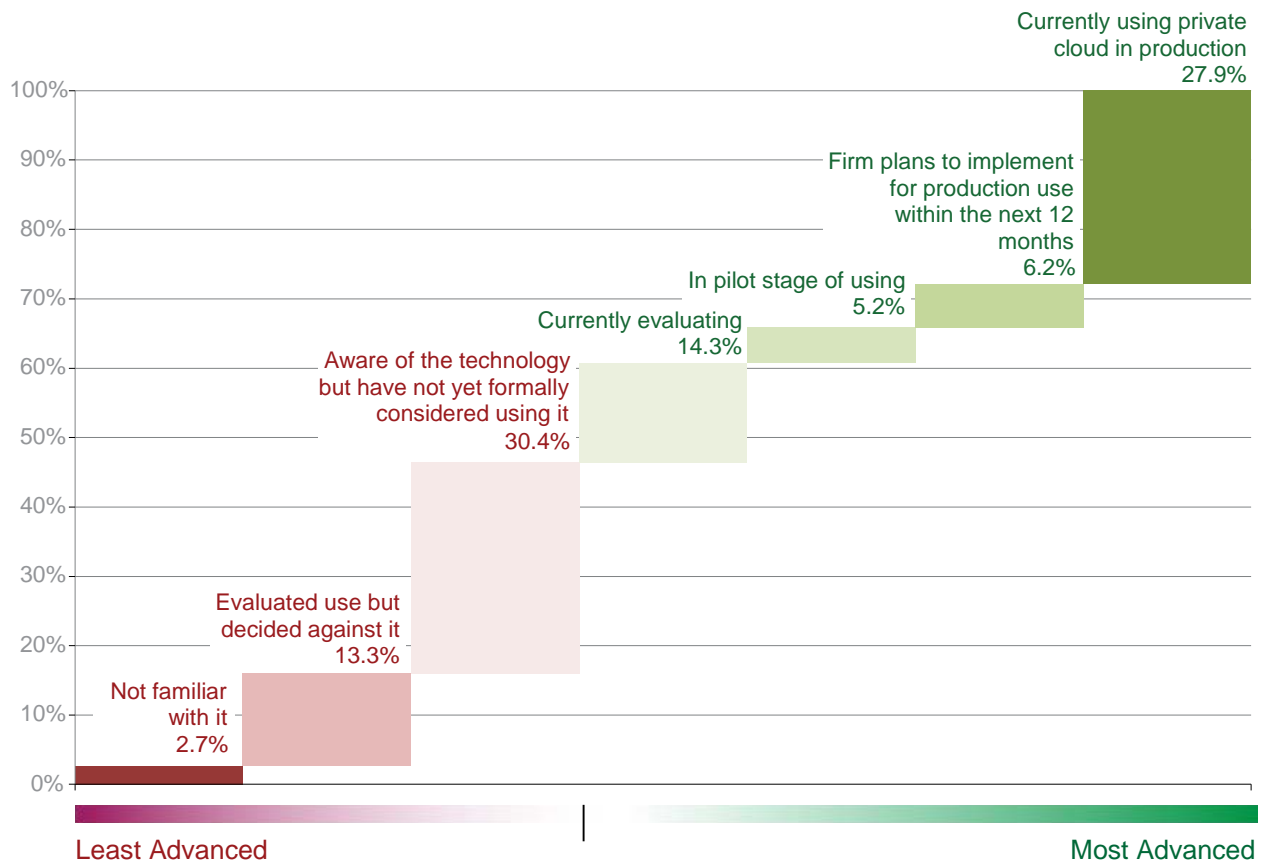
The cloud service model and technologies were developed and matured by public cloud providers. As grassroots usage among enterprise users grew for public cloud, users also began to demand similar services from internal IT, leading many corporations to develop initiatives around adopting a service provider model backed by cloud technology. Private clouds run on-premise and are built and operated by IT.

While private cloud accounts for only 8% of enterprise computing today, 28% of enterprises are already using it for production, with another 11.4% in the pilot or planning stages, as depicted in Figure 4. This means that private cloud has moderate enterprise penetration today, but with only small-scale deployments at each site. IDC research shows that early private cloud adopters generally limit private cloud deployment to relatively few targeted use cases/applications and only select groups of users.

Private cloud adoption today is also primarily based on existing enterprise virtualization because that is the installed base of infrastructure that has been built out over the past decade. Organizations developed CMP-like capabilities in conjunction with these virtualized servers to provide cloud features such as metering, chargeback, and self-service. A much smaller, but fast-growing, type of cloud uses a next-generation architecture similar to the hyperscale public cloud providers. It generally does not use traditional enterprise hypervisors or run legacy enterprise applications. Instead, new cloud-native apps need to be developed to take advantage of the new style of infrastructure. OpenStack is a prime example of this newer-generation cloud architecture and has generated a lot of buzz in the industry.

FIGURE 4

Private Cloud Adoption



n = 405 worldwide respondents
 Source: IDC's *Virtualization and the Cloud MCS*, 2013

OpenStack Builds Strong Momentum

OpenStack is an open source project with tremendous momentum and participation across the industry. OpenStack is primarily an infrastructure-as-a-service platform, though recently the project has expanded into some platform-as-a-service areas. The goal is to create programmable infrastructure that can be accessed through a set of open APIs or GUI management tools that access those APIs. The project spans the core infrastructure disciplines including compute, storage, and networking. The code is modular, and the project is divided into subprojects that cover each module. OpenStack is also extremely pluggable, in many ways acting as a framework that interfaces with components through drivers. For example, OpenStack supports every major hypervisor in the market through drivers. But it can also support other compute destinations such as OS containers such as Docker and even bare metal. Likewise, the storage and networking functions are supported on a variety of storage devices and networking systems through a pluggable driver interface. The open, modular, and pluggable nature of OpenStack is what has drawn a strong community and the interest of customers that are encouraged by the efforts to foster more standardization and interoperability in a market of fragmented platforms.

OpenStack Adoption

OpenStack has been initially adopted by service providers, Web companies, and academia/research. They are typically advanced organizations for which technology is the core business, and they feel comfortable using code that is in the early stages of maturity and often customize and support the software themselves. This is very typical of how open source software starts in the market. These adopters had to work through the early issues, but they also demonstrated that OpenStack could be successfully used to build clouds at scale. Today, OpenStack is looking to the commercial enterprise as the next frontier and to a much broader mainstream audience. To expand into this commercial enterprise market, the project and commercial vendors are working on several areas key to this success. They are building capabilities and interface characteristics that make OpenStack easier to deploy and manage by enterprise IT organizations. These include:

- Focusing not just on software features but on QA to develop better software development and testing processes
- Developing interoperability guidelines, certifications, and tests
- Reducing the number of disruptive changes and providing a better procedure for rolling out the essential new releases
- Smoothing enterprise deployment issues such as installation and non-disruptive, rolling upgrades
- Growing the knowledge and skills base by developing better documentation and training resources
- Figuring out how to align OpenStack's fast six-month project release cycles with longer commercial product release cycles
- Providing management functionality such as capacity planning, control policies, chargeback, and rich self-service functionality

Vendors also have a large opportunity to build on top of OpenStack because OpenStack provides only the base infrastructure for clouds. Several key management ingredients are still needed to build a full cloud service that is not part of the core OpenStack project. Vendors are developing and integrating a variety of cloud management platform functions with OpenStack to create a full enterprise cloud solution. OpenStack does not include platform-as-a-service functionality either, another opportunity for vendors to further extend the stack.

OpenStack is a cloud system software, but it still requires an operating system to run. The vast majority of OpenStack deployments use Linux and leverage many of Linux's features, such as KVM, cGroups for service-level management, container functionality such as kernel namespaces, and core networking functions for software-defined networking. The tight integration between Linux and OpenStack has naturally led to the commercial Linux vendors becoming key OpenStack contributors and distributors because of their experience in developing, testing, and supporting Linux and their knowledge of the open source development model.

KVM and OpenStack

As discussed previously, OpenStack supports all the major hypervisors, but KVM currently dominates as the hypervisor of choice for OpenStack deployments. OpenStack Foundation user data suggests that KVM accounts for 70%+ of OpenStack compute. The reasons for KVM's tight association with OpenStack are as follows:

- KVM is part of the Linux kernel and thus the default hypervisor for Linux.
- OpenStack requires a base operating system for the control plane servers, and Linux is the obvious choice because of its open nature.
- A fully open software stack, from hypervisor to operating system to cloud system software, is desirable to allow open software development without licensing limitations as well as to troubleshoot problems.

For these reasons, KVM became the default virtualization environment for developers and early adopters, and subsequently, the KVM driver became the most robust, and many reference architectures and documents use KVM as the default configuration. While success for OpenStack will also mean success for KVM, it may also open up opportunities for KVM outside of OpenStack as users are introduced to and become comfortable with the technology and begin to apply it to traditional server virtualization environments.

Pets Versus Cattle

If OpenStack can support all the major hypervisors and even bare metal, why couldn't OpenStack become the framework to control the entire datacenter? Significant architectural differences between OpenStack and traditional infrastructure make that difficult. A commonly used analogy is the pets versus cattle comparison.

Traditional workloads in this analogy are "pets." Pets are unique and valuable, and owners spend a lot of money to keep them happy and healthy. Traditional applications are all unique in their infrastructure requirements, with each application dictating a specific set of needs. These applications also expect that the infrastructure will be highly available, and IT administrators spend a lot of time and money to make sure they never go down. Pet workloads are typically stateful applications and scale up. In short, the level of service a traditional application achieves depends on the infrastructure on which it runs.

Cloud-native applications were pioneered in the public cloud and made different assumptions about the infrastructure based on the cloud services that were available. This created a new approach to applications. The application could no longer dictate unique infrastructure requirements as the cloud presented a certain set of infrastructure services that were universal for all applications. Many of these clouds also had a different availability model, with high availability focused at the regional level and only moderate availability at the individual server level. The idea was to shift some of the burden of availability to the application, enabling the creation of highly available applications on moderately available infrastructure for reasons of cost. Thus cloud-native applications are distributed, built to scale out, and often stateless. "Cattle" is the term often used to describe these applications because they have little individual personality, and if one application has a problem, it is often just ignored or killed and restarted, with the rest of the "herd" masking the failure.

OpenStack was built originally to be optimized for cattle workloads. Over time, OpenStack has added some "pet" features and likely will add more as the enterprise market for OpenStack grows. However, most traditional workloads are already virtualized, and there is already an optimized environment built to run those. OpenStack will manage some traditional workloads, but it is difficult to reconcile two very different architectures, and perhaps, there is no real reason to try and do so. Likewise, traditional virtualized server environments are adding cloud features, such as layers of automation and self-service portals as well as some scale-out features. However, traditional virtualized server environments will probably not be able to fully meet next-generation cloud application requirements simply by extending the traditional architecture with these added cloud features. While there will be some overlap between these two environments, the differences are too great to be unified, and the reality is that the industry and customers will likely be managing both environments for the foreseeable future.

In addition, the continued use of bare metal physical environments and off-premise public clouds leads to a very hybrid model of IT that enterprises will be managing for quite some time. The reality in enterprise is that applications tend to stay where they were first deployed. Migration is difficult and expensive, and unless customers are being forced off the platform, it's not cost effective to migrate for migration's sake. Organizations have made significant investments in 2nd Platform applications, and completely rebuilding these applications to run on the 3rd Platform is a monumental task that is nearly impossible in many cases and will take a generation to complete at best. While clouds – both private and public – have tremendous benefits and growing footprints, the traditional workloads and their architectures are not going away soon, and this will create a challenge for IT to manage an increasingly diverse mix of environments.

RED HAT CLOUD INFRASTRUCTURE

Red Hat Cloud Infrastructure is a robust cloud solution that integrates several Red Hat products and can be used to deploy and manage both traditional environments and new cloud environments, giving organizations the ability to reduce complexity while meeting the needs of cloud applications. It consists of several components, all of which are detailed in this section.

Red Hat Enterprise Virtualization (RHEV) is an enterprise virtualization solution that consists of the KVM hypervisor and a virtualization management console. RHEV is a full-featured virtualization solution that has the typical features that enterprises expect, such as:

- Scalability to 160 CPUs and 2TB RAM for both host and guest
- Agility features such as live migration and storage live migration
- High availability for guest virtual machines (VMs)

This environment is optimized for traditional enterprise or "pet"-type workloads. Red Hat is a leading contributor to KVM, now part of the Linux kernel, and created the oVirt open source project, the first open source enterprise virtualization management solution upon which RHEV is based.

For next-generation cloud applications (aka "cattle"), Red Hat Cloud Infrastructure also includes Red Hat Enterprise Linux OpenStack Platform, a commercial, enterprise-focused distribution of OpenStack. It is a hardened, tested distribution that is built with the same principles that made Red Hat a success in commercial open source and Linux:

- Red Hat Enterprise Linux, a mature and hardened operating system, underpins the OpenStack control infrastructure, and KVM is used for the compute nodes.
- Red Hat provides an ecosystem of hardware, software, and services partners to help customers build and deploy robust OpenStack-based clouds.
- Red Hat backs the Red Hat Cloud Infrastructure products with enterprise-level support and training/certification materials.
- Red Hat's support life cycle for OpenStack is three years for a given version, which is currently the longest available for a commercial OpenStack product.

Red Hat is a leader in the OpenStack community and has consistently been a top code contributor to the project since it joined. Red Hat is committed to the open cloud and is applying the same tried-and-true community building and open source commercialization processes that it has developed over the many years of building Linux into a commercial enterprise solution.

With RHEV and OpenStack providing the core infrastructure for traditional and next-generation applications, the rest of the RHCI suite is about managing the hybrid environment.

Red Hat Satellite is a systems management solution used to manage individual elements in both environments. It provides essential features such as provisioning, configuration management, software management, and subscription management to both the host machines and the guest VMs.

CloudForms is a cloud management platform that can span diverse environments. It can manage:

- Traditional server virtualization environments, both Red Hat and non-Red Hat
 - CloudForms enables advanced management and cloud features for these environments, adding features such as orchestration, capacity planning, monitoring reporting, chargeback, policy compliance, and self-service portals.
- Cloud environments such as OpenStack private clouds and public clouds such as Amazon
 - CloudForms provides higher-level cloud management across both environments for consistent functionality across the hybrid cloud, regardless of the underlying cloud platform.

In addition, Red Hat Cloud Infrastructure contains Red Hat Enterprise Linux, which forms the basis of Red Hat Enterprise Linux OpenStack Platform and also Red Hat Enterprise Virtualization. In addition, organizations can opt to purchase Red Hat Cloud Infrastructure with Red Hat Enterprise Linux, which provides unlimited guests for virtual machines for those tenants that wish to run their applications in the organization's cloud.

As customers evolve their traditional server virtualization while building private clouds and integrating public cloud resources, they are struggling to juggle the various technologies, all of which continue to provide value and support key applications. Red Hat Cloud Infrastructure is a hybrid IT solution that has its feet in both the traditional virtualization world and the next-generation cloud world and ties both worlds together with cloud management tools to help manage these diverse environments. Red Hat Cloud Infrastructure works alongside existing infrastructure investments and integrates with numerous virtualization, management, and public cloud technologies. For customers taking the next step into cloud, it offers a solution for today's needs as well as a path for tomorrow, helping customers on their cloud journey.

CHALLENGES/OPPORTUNITIES

Challenges

- There is a large installed base of non-Red Hat server virtualization and a very competitive hypervisor market in the enterprise. Vendors that own these hypervisors will seek to own the management stack on top of it as well as the cloud stack. KVM came later to the market, and while it has a good attach to OpenStack, it doesn't have a large traditional virtualization installed base to build off of.
- The cloud platform market is extremely competitive and fragmented. Other major enterprise system software vendors are also building private cloud solutions, seeking to capitalize on their installed base of system software. The top major hyperscale cloud providers all use proprietary platforms, with OpenStack being used for mid- and smaller-scale clouds.
- The OpenStack community has grown large and fast, and the challenge continues to be not to collapse under its own weight. It can be tough to achieve consensus and vision and act swiftly upon decisions with such a large and diverse community whose members often gravitate toward their own selfish, commercial interests rather than the good of the community. There is also a danger of fragmentation with commercial versions of OpenStack. Various implementations may not be fully compatible with other OpenStack clouds, or the OpenStack software can be used to make a closed solution. It is still very possible for vendors to build a very closed solution using open source.

- OpenStack still has essential issues to resolve to become a commercial product. The process of installing, maintaining, and upgrading an OpenStack cloud can be intimidating and overwhelming. To achieve commercial success, OpenStack must become easier to use and more accessible to the masses. This will require both technical work and fundamental project changes such as release cycle management and interoperability testing. OpenStack-to-OpenStack compatibility is still a work in progress, and this is a key value proposition of OpenStack that must be delivered.

Opportunities

- With all the fragmentation in the cloud platform market, OpenStack has arisen as the hope for a unifying standard. With the broad participation by the industry, OpenStack technology will hopefully be brought to market in various forms and be supported by the IT industry. The open nature of the APIs and the development and healthy nurturing of a large, growing community will be key to OpenStack's continued success.
- As the industry moves to a cloud service model, enterprise IT must adapt to this new world or face irrelevance. This has led to the adoption of the private cloud. While the public cloud has its uses, some enterprise workloads just can't be moved off-premise. For the foreseeable future, hybrid cloud will remain the dominant model, and enterprises must build cloud infrastructure and services to support its users and develop modern applications with a hybrid world in mind.
- While OpenStack-to-OpenStack interoperability, portability, and compatibility are still works in progress, they hold the key to broad market adoption. Yes, OpenStack is a great platform with many features, but much of the value to customers is the ability to choose from a wide variety of OpenStack private cloud implementations, ecosystem solutions and add-ons, and OpenStack public clouds and have everything work together seamlessly. OpenStack, by its open and collaborative nature, is a cloud initiative in the industry with the reach and power to accomplish that vision, and enterprises are hoping that they can bank on that promise for an open cloud.

Red Hat is a leading Linux vendor and also has significant experience with the open source model. OpenStack needs a stable operating system underneath, and many OpenStack functions rely on Linux, thus requiring engineering work on both sides to ensure good integration. Red Hat's position in Linux lends significant engineering and business model experience, which is a competitive advantage. In addition, Red Hat's work in open source virtualization (RHEV) and open source cloud management (CloudForms) is also key, as these additional elements are needed to build a full cloud solution as OpenStack provides only the core infrastructure.

CONCLUSION

Cloud is a simple concept, but in reality, it surfaces in a variety of ways. Enterprises will need to tackle a matrix of:

- Cloud at multiple layers (e.g., infrastructure layer, platform layer, application layer) and management of all of these layers together
- On-premise private cloud assets as well as off-premise public cloud assets
- Multiple cloud architectures that could be completely next generation or an evolution of traditional infrastructure such as virtualized servers

While new technologies such as OpenStack and Docker are hot key areas going forward, not every workload can or will be migrated to OpenStack, and Docker won't replace every hypervisor that is out there today. Enterprises are accumulations of many generations and platforms of technology, and next-generation cloud will be just one more. The existing server virtualization environment will still be around for a very long time, and it will gain cloud features too, albeit in a different way than OpenStack. The juggling act that is enterprise IT will become only more difficult!

To succeed, enterprises will need to approach their cloud strategy from multiple angles and will require a diverse set of technologies. Red Hat is a vendor that is developing solutions to address these needs as a leading provider of open solutions. IDC has seen that openness become a larger factor for customers in a cloud world. Red Hat's legacy in enterprise Linux and KVM virtualization provides the foundation for other emerging areas such as OpenStack and CMPs. By tying it all together as an integrated solution and with Red Hat's experience in supporting open source for the enterprise, Red Hat has a tremendous opportunity to provide cloud solutions that will be key for enterprises to make the transition into the cloud era.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2015 IDC. Reproduction without written permission is completely forbidden.

