WHITEPAPER

CLOUD INFRASTRUCTURE FOR THE REAL WORLD

Managing traditional and cloud workloads simultaneously

Gordon Haff

"To capture digital opportunities, CIOs need to deal with speed, innovation, and uncertainty. This requires operating two modes of enterprise IT: conventional and "nonlinear."

GARTNER, EXECUTIVE SUMMARY:
TAMING THE DIGITAL DRAGON:
THE 2014 CIO AGENDA

EXECUTIVE SUMMARY

Enterprise applications are changing. Traditionally, they could take years to write and they could live for decades. They were monoliths that had to be protected against failure at all costs. The decades of research and development that went into technologies such as failover clustering was ultimately in service of protecting the running state of this type of application; a failure of one of these applications could result in disaster for the business.

Newer, cloud-style workloads are different. They're disposable, stateless, and modular. They're often written quickly in discrete chunks that communicate over the network. They're frequently about collaboration and customer engagement. They have to adapt quickly to changes in the external environment.

IT organizations increasingly have to accommodate these two fundamentally different types of workloads, as well as applications that combine the two.

This accommodation involves bringing workloads to hybrid IT in an evolutionary way and managing them there. It requires optimizing for different types of workloads as they change over time, such as when traditional workloads are decommissioned in favor of new architectures. It means avoiding new points of proprietary lock-in and new silos, and breaking from proprietary software and hardware—at your own pace. This accommodation–a new IT model–requires a smooth transition.

This whitepaper discusses this new model in detail and how Red Hat® Cloud Infrastructure makes a smooth transition possible. It covers the way in which cloud Infrastructure-as-a-Service (IaaS), together with both the open hybrid management of that IaaS and the life cycle management of the applications running on the IaaS, can accommodate both workload types. The result is that you can transition to new application styles in the manner and at the pace of your choosing.

facebook.com/redhatinc
@redhatnews
linkedin.com/company/red-hat

redhat.com

INTRODUCTION

Early cloud computing talk about the inevitability of a public compute utility has mostly given way to more nuanced discussions. Cloud characteristics such as flexibility, rapid scalability, efficiency, faster application development, and agility are as desired as ever, and are seen as essential enablers for next-generation IT infrastructures. However, these infrastructures won't come in one shape and one size—and they won't pop into existence overnight.

There are a number of reasons for this variety. For example, it's well established that many organizations plan to continue running at least some of their applications and workloads[1] within facilities that they control, rather than in a public cloud (for real or perceived reasons of security, visibility, compliance, or control). It's also true that many existing IT investments can't and shouldn't transition to something new overnight; doing so would be neither practical nor cost-effective. As a result, clouds will often be hybrid blends of public and private resources and heterogeneous infrastructures.

Hybrid clouds—specifically, open hybrid cloud approaches that provide for application and data portability—are also the destination of choice because the workloads that organizations run are likewise hybrid. Hybrid workloads reflect the distinction between traditional enterprise workloads and new-style workloads. New application architectures and technologies are better suited for highly distributed, fluid IT infrastructures than was once the norm. But they also reflect the different roles that workloads play within an organization.

Author and consultant Geoffrey Moore (of Crossing the Chasm fame) draws a distinction between systems of record and systems of engagement.[2] The former are about data processing; it's no coincidence that term was once a synonym for what we today call IT. Systems of record are focused on secure and reliable operations. They're usually relatively static (with some exceptions, especially in financial services). They're the stuff that once ran on mainframes and big-iron UNIX servers; although many of these workloads have moved to x86, Linux®, and even virtualized environments.

By contrast, systems of engagement grew out of what Moore calls "the 21 century's build-out of PC apps and consumer games into the desktop PC Web of media and commerce into the 'appified' web of smart phones and tablets." They're correspondingly more focused on the user experience and more provisional, which is to say often short-lived and mutable.

At the beginning, these collaboration systems largely existed in parallel with and independently of IT; and IT folks sometimes disparaged them as shadow IT or rogue IT. But today, systems of engagement are increasingly important to organizations' missions, whether through engaging with consumers or leveraging cloud-scale services.

As a result, this formerly out-of-band activity is now often an integral part of IT, especially when it ties directly to back-end data and processes. Doing so touches on both technical matters and overall IT governance and compliance.

But accommodating both systems of record and systems of engagement within one infrastructure presents its own set of challenges. That's where open hybrid management of IaaS comes in.

---

1 *The term "workload" is generally used in this paper to connote the set of operating system instances and the other software running within them that provide a specific service to users or other workloads. Examples of a workload are a datastore or a social game.*

2 *http://www.linkedin.com/today/post/article/20121112224653-110300724-the-future-of-enterprise-it-systems-of-record-meet-systems-of-engagement*

HOW WORKLOADS DIFFER

The prior discussion considered two different classes of enterprise workloads from the perspective of the business. Now, let's shift our view to consider them in the context of associated workload characteristics and corresponding infrastructures.[3]

TRADITIONAL ENTERPRISES

By way of analogy, a traditional enterprise application is like a pet.[4] It's a unique individual that can't just be replaced. If the pet clumsily careens into a tree and injures itself, you take it to the vet. You can also think of these applications as craftwork, with each application environment lovingly assembled and carefully tuned.

From a technical perspective, traditional enterprise applications are defined by their persistence and statefulness. Lose an application instance and an arbitrary copy can't just take over; each instance holds unique information about users and transactions and often is tied into other centralized resources such as an SQL database. These fundamental characteristics have driven a huge amount of computer systems R&D, spanning decades. Failover clustering, fault tolerance, hot-plug, and redundant hardware components are all the product of needing to protect persistent and stateful applications.

In more recent times, these characteristics have significantly influenced some of the directions that enterprise virtualization has taken. Live migration allows a running virtual machine (VM) to transfer from one physical server to another for hardware maintenance and other purposes. Storage snapshots create a restore point of a VM's operating system, applications, and data.

All this, in turn, leads to generally more monolithic applications. If you're going to have to spend the money and effort to protect individual application instances, it often makes sense to pack more services and capabilities into an individual instance. This was often fine for relatively static systems of record, and may even be necessary for certain types of applications that are inherently single-threaded or have low latency or high bandwidth requirements. But it works against flexible and scalable IT infrastructures.

To summarize, in traditional enterprise applications:

- VMs are stateful. The identity of the individual VM matters and applications generally map to a specific VM.
- VMs are often large, with multiple CPUs and large amounts of memory and storage.
- VM life cycles can be measured in months or even years.
- Application availability is tied to the heartbeat of a specific VM. Therefore, technologies like clustering and live migration are important to maintaining uptime.
- Applications often scale up, requiring adding resources (such as compute, memory, and storage) to a running VM.

3 It's worth noting that many newly written applications apply modern design patterns and technologies to implement reliability, availability, and other aspects of systems of record on cloud-style IaaS. However, most systems of record today have traditional enterprise workload characteristics, while most systems of engagement are more cloud-like.

4 The original notion of pet and farm animal workload types is commonly attributed to Bill Baker.

CLOUD-STYLE IAAS

By contrast, an application designed for the cloud[5] is like a widget coming off an assembly line. If it breaks, you get another one—and you won't be able to tell the difference. Therefore, replacement is often a more sensible economic choice than repair.

Individual cloud application instances are stateless. If a VM instance fails—or even an entire physical server's worth of VMs fails—you just start more VMs. Or, if demand goes down, you can shut off VMs. No single instance is anything special.

This isn't to say there's no need for high-availability architectures or that those architectures will necessarily be simple or straightforward. In some respects, providing reliability across a distributed architecture is more challenging than in the traditional case where availability often came more or less out-of-the-box from the operating system or system vendor.

Lack of state and persistence have important design implications. With no need to protect single instances, it's often better to encapsulate fine-grained services and have them communicate through lightweight protocols rather than to combine everything.[6] The result is a collection of services that 1) can be scaled up and down as needed and 2) can often be reused as part of an organization's complete application portfolio. Furthermore, while this approach certainly doesn't eliminate testing, it does tend to get away from the traditional "don't touch the production server" mentality.

To summarize, in cloud-style workloads:

- VMs are stateless and, while groups of VMs have different characteristics, they don't have individually unique identities.
- VMs are small, with modest CPU and memory, and storage often resides elsewhere.
- VM life cycle is measured in hours to months.
- The application itself provides availability mechanisms that don't depend on the health of an individual VM, which can be destroyed and recreated in the event of a problem.
- Applications scale horizontally by starting new VMs to add capacity.

MIXED WORKLOADS

Applications can also consist of a mixture of workload types. In fact, mixed workloads are a natural consequence of systems of engagement consisting of reusable discrete services interacting with and consuming data from back-end systems of record, which mostly still consist of traditional enterprise workloads. Mixed workloads only highlight the need to manage across different workload types and corresponding infrastructures, rather than creating new isolated cloud or virtualization silos.

5 To be clear, it's not suggested that only applications written from scratch that use the most leading-edge distributed application technologies can use an IaaS. But "cloudier" apps can better use the capabilities an IaaS affords them and they are less likely to miss the features related to persistence that often are not in an IaaS.

6 Lightweight, encapsulated, single-function services are often referred to as "microservices." Over time, these microservices will often run in containers within an operating system instance.

THE PATH TO THE CLOUD

These different types of workloads and everything in between—even within a single hybrid application—are the reality of enterprise IT. That IT is also hybrid in other respects, whether in the sense of heterogeneous, on-premise infrastructure or because it includes both dedicated capacity and multi-tenant public clouds. Furthermore, it is not a static environment.

One of the hallmarks of the shift to cloud computing is that it makes computing resources available more dynamically in response to changing business needs.

It should also be clear that hybrid IT is not a forklift upgrade that abandons the existing IT investments, processes, and applications that are deeply wired into how the business operates. Such an upgrade would be costly and, even more importantly, risky. The vast majority of companies evolve their IT; they don't blow it up. Even the handful of large firms that are the most visible public cloud consumers got to that point over time.

The type of resources an application consumes and how it consumes them is fundamentally an IT-policy decision based on a variety of criteria. Total cost, how capacity is charged for, predictability, reliability, performance, auditability, security processes, and existing investments will all affect where IT decides to run various workloads over time.

However, such infrastructure decisions should generally be abstracted from users. Put another way, the user shouldn't know or care where an application is running. This is an important point because private clouds that don't provide the ease-of-use and flexibility of public clouds are an invitation for shadow IT. Just calling on-premise IT a "cloud" won't make it one, as far as users are concerned, if it doesn't solve their problems the same way public clouds do.

This new model for IT involves bringing workloads to hybrid IT in an evolutionary way and managing them there. It requires optimizing for different types of workloads as they change over time, such as when traditional workloads are decommissioned in favor of new architectures. It means avoiding new points of proprietary lock-in and new silos and breaking from proprietary software and hardware—at your own pace. The new model requires a smooth transition.

The answer is an open hybrid cloud approach that uses existing investments in virtualization and other software and hardware, and provides you with a single view into the management of this hybrid infrastructure.

HOW RED HAT IS MAKING HYBRID A REALITY

Red Hat Cloud Infrastructure lets organizations build and manage an IaaS cloud based on enterprise datacenter virtualization, while simultaneously providing an on-ramp to highly scalable, public-cloud-like infrastructure based on Red Hat OpenStack® technology. Furthermore, Red Hat Cloud Infrastructure complements existing virtualization and other IT investments, as well as public clouds, and manages the entire infrastructure through a single, centralized administration console using Red Hat CloudForms and the software running on that infrastructure with Red Hat Satellite.

Thus, Red Hat Cloud Infrastructure provides infrastructure suitable for the gamut of workloads that exist in organizations today. And it provides comprehensive, unified management for those workloads across an open, hybrid cloud, doing so without the cost or lock-in of proprietary solutions.

The solution is built on the enterprise-hardened and trusted Red Hat Enterprise Linux platform and uses Red Hat's history of delivering enterprise-grade solutions based on open standards and projects. For example, the Red Hat subscription model allows organizations to take advantage of the community innovation that is driving the fast-changing OpenStack project, while providing a supported life cycle for production deployments.

Red Hat Cloud Infrastructure includes:

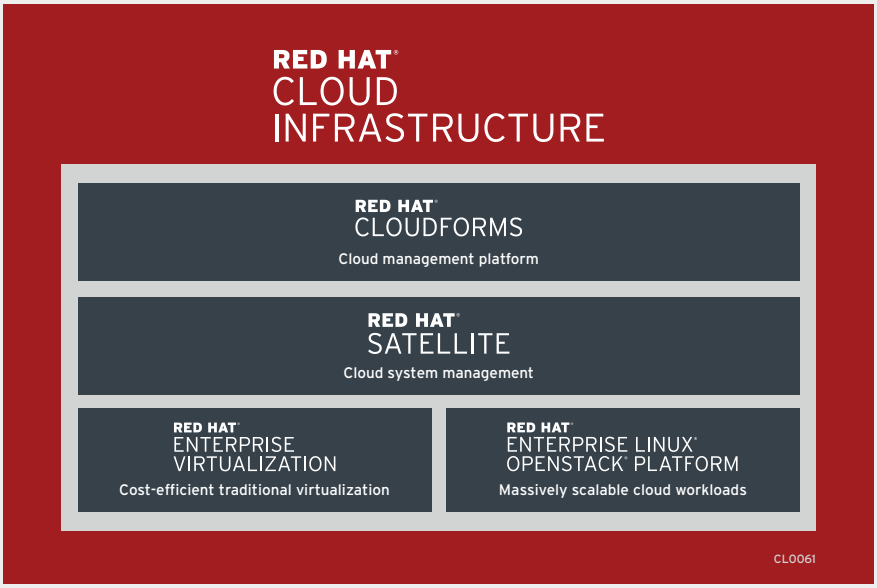

Red Hat Enterprise Virtualization, a complete datacenter virtualization solution for Linux and Windows workloads that lets you build an agile, secure virtualization foundation with the features needed for more traditional enterprise application workloads, such as live migration of running applications and storage snapshots.

Red Hat Enterprise Linux OpenStack Platform, a massively scalable IaaS offering that delivers an open, flexible, and enterprise-ready private cloud foundation built on Red Hat OpenStack technology and the Red Hat Enterprise Linux platform. As a top contributor to the community OpenStack project, Red Hat uses its long-standing open source community leadership to create trusted, enterprise-supported private cloud infrastructure. This infrastructure is ideally suited for organizations as they start to evolve their workload portfolios to architectures that are scale-out, latency-tolerant, and designed for availability at the application level.

Red Hat CloudForms, open hybrid cloud management that provides visibility and control over existing heterogeneous virtual infrastructures and lets you deploy, monitor, and manage cloud services. It enables integrated life cycle management of server and storage provisioning, workload optimization, policy-based compliance and control, and analytics across heterogeneous virtualized datacenters and clouds. It thereby provides a single point for viewing and maintaining policy-based control of all your workloads across your entire infrastructure—and provides you with the means to smoothly transition to the new IT.

Red Hat Satellite complements CloudForms by providing application life cycle and advanced system management for Red Hat Enterprise Linux applications running on both cloud and traditional infrastructures, including physical servers. Red Hat Satellite provides simple tools to provision new systems, manage their updates and configuration changes, and monitor system performance. These capabilities allow for scaling the management of cloud environments, giving system administrators more power and freedom to manage their ever-growing infrastructures. Furthermore, these capabilities enable the automation of standard operating environments on which cloud flexibility and dynamism depends.

CONCLUSION

The cloud is a new pattern for computing. It's an abstraction atop infrastructure that was once tightly bound to an application, often for the lifetime of that infrastructure. Even virtualization didn't change that model much; it just swapped out physical servers for virtual ones. Beneficial? Absolutely, and often an economic game changer. But it didn't really change the application in any fundamental way.

That changes with the cloud. Cloud computing touches on many different aspects of development and operations. But one of the most fundamental areas it touches is the application itself. The self-service, automation, and embedded policy associated with open hybrid clouds are big wins. But, to fully take advantage of the architecture, application architectures often need to change to take the maximum advantage of the cloud's dynamism. Applications need to lose state. They need to become more horizontally scalable. They need to decompose into services. They need to provide reliability and availability mechanisms in new ways.

But applications don't change overnight in large organizations. Indeed, other than data, applications may be the longest-lived component of an IT architecture. Certainly, new applications are getting written at what seems like a record pace as enterprises respond to burgeoning demands brought on (and business opportunities created) by data analytics, mobility, and social. But existing applications continue on as well, often serving as the back-end for the new stuff.

As a result, most organizations should evolve to the cloud in a way that takes advantage of their existing infrastructures and applications. That's the approach Red Hat makes possible with Red Hat Cloud Infrastructure.

Red Hat Cloud Infrastructure provides you with infrastructure options suitable for both your traditional and cloud-style applications, as well as open hybrid cloud management spanning these options and more. It lets you move to new application architectures on your terms, at your pace, with your IT policies, and under your control. It's the bridge between the infrastructure and applications that your business runs on now and the infrastructure and applications that will provide new revenue and other opportunities today and in the future.

ABOUT RED HAT

Red Hat is the world's leading provider of open source solutions, using a community-powered approach to provide reliable and high-performing cloud, virtualization, storage, Linux, and middleware technologies. Red Hat also offers award-winning support, training, and consulting services. Red Hat is an S&P company with more than 70 offices spanning the globe, empowering its customers' businesses.

NORTH AMERICA
1 888 REDHAT1

EUROPE, MIDDLE EAST AND AFRICA
00800 7334 2835
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com

facebook.com/redhatinc
@redhatnews
linkedin.com/company/red-hat

redhat.com
#11015587 / INC0190867_1014
DOC139116