WHITEPAPER

SEVEN CONSIDERATIONS FOR RUNNING YOUR WORKLOADS IN A PUBLIC CLOUD

Gordon Haff

EXECUTIVE SUMMARY

Loading an application onto a public cloud is easy. A credit card number, a few clicks, and you're good to go. That on-demand, self-service access is a big upside of using a public cloud.

But that incredibly short and fast on-ramp can be a downside as well.

The reality is that, for most organizations today, running a small, simple application on one public cloud may be a first step but it isn't the ultimate—or even a mid-term—goal. That goal is more nuanced, more complex. That goal is to reliably run large, scalable applications made up of many loosely-coupled components. That goal is to maintain flexibility around where applications run—whether that means switching providers or running them on a mixture of in-house, dedicated external, or multi-tenant cloud resources. That goal is to make use of public clouds as part of a broader hybrid IT environment rather than a new silo of vendor lock-in.

So how do you move a workload (or write a new application) for a public cloud and meet that goal?

Many different factors come into play. However, based on a wide range of conversations with customers, partners, analysts, and others, we've developed a list of seven key considerations that are especially important and applicable to a broad range of situations:

1. Developing an appropriate application architecture
2. Maintaining portability of applications and workloads
3. Ensuring data is portable
4. Understanding legal and regulatory compliance
5. Enabling hybrid cloud management, policy, and governance
6. Isolating workloads as needed in a public cloud
7. Providing access to standard languages, frameworks, and tools

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

redhat.com

INTRODUCTION

"The cloud is now a viable option for a broad range of enterprise workloads and is actually preferred for a growing class of new customer-facing web and mobile applications."

FORRESTER RESEARCH, INC.
"CLOUD MANAGEMENT
IN A HYBRID WORLD"
JULY 2013

Public clouds promise nearly unlimited capacity, attractive economics stemming from large-scale infrastructures and associated operational best practices, and a pay-as-you-go pricing model that can make up-front server purchases a thing of the past.

In practice, the decision whether to run a particular workload on a private or public cloud depends on a wide range of factors–from control to cost to compliance to available capital. The reality is that most organizations, especially large ones that adopt public clouds, will do so as part of a hybrid IT environment in which resources and services are procured from a variety of sources—both internal and external.

There are many aspects to designing and managing such an environment. It requires understanding where the responsibility lies for the different layers of IT infrastructure. Regulatory requirements and other legal or risk management considerations will often constrain where particular workloads run—as may well the technical architecture of the workloads themselves. Running workloads across a hybrid environment, or even (as is often the case) keeping options open for where they may run in the future, introduces additional considerations.

Learn more about The Cloud Security Alliance's 98 "control areas" at **www.isaca.org/chapters2/kampala/newsandannouncements/Documents/IT%20contro%20objectives%20for%20Cloud%20computing.pdf**

Given that cloud computing is essentially next-generation IT, a complete examination of all the factors that can affect building and operating clouds would require an equally complete exploration of IT governance in the broadest sense. The Cloud Security Alliance's Cloud Control Matrix alone examines 98 "control areas" in 11 categories, each mapped to its area of relevance and relevant regulations and certifications. And that's just the factors most directly related to the security concepts and principles that apply to assessing the overall security risk of a cloud provider (whether an internal or an external one).

For the purposes of this whitepaper, we have distilled cloud adoption considerations into a short list. These considerations come up repeatedly in our conversations with customers, partners, and providers. We selected considerations that are particularly relevant to choosing a public cloud, especially in the case where it will be used in concert with a broader hybrid IT infrastructure. By all means, also consider your own detailed IT governance plans and best practices—as well as the external risk assessment and security controls documents of your choice. But we think these are a good place to get started if you're taking a look at public clouds.

KEY CONSIDERATIONS

"The time is ripe for I&O pros to rise to the cloud challenge and chart a path to hybrid cloud success."

FORRESTER RESEARCH, INC.
"CLOUD MANAGEMENT
IN A HYBRID WORLD"
JULY 2013

DEVELOPING AN APPROPRIATE DISTRIBUTED APPLICATION ARCHITECTURE

Clouds, even public clouds, don't require an overnight wholesale replacement of an organization's IT infrastructure and application portfolio. Indeed, one of the fundamental tenets underpinning the idea of hybrid clouds is to make it possible for organizations to adopt cloud at their own pace while continuing to make use of their existing IT investments.

That said, "cloud-native" applications follow different design patterns than those associated with traditional enterprise apps. Individual cloud application instances are stateless. If an instance fails, you just start more instances. Or, if demand goes down, you can shut off instances. No single instance is anything special. 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 into a single monolithic application.

Applications written in this new style, often in concert with new approaches for storing and distributing data, are a better match for the scaling and availability mechanisms used in clouds generally and public clouds in particular. For example, failover clustering and other traditional approaches to server availability are simply not available in public clouds; rather, availability is ensured at the service level in software.

Likewise, with a limited ability to scale up individual instances in a public cloud, applications must scale-out to add capacity.[1] Doing so is especially important when one considers that being able to respond elastically–even to large and rapid shifts in demand–is one of the key benefits that public clouds can provide.

If you're going to run a production application on a public cloud, it's also important to understand the most appropriate ways to protect that application from failures. This might include taking advantage of cloud provider features such as availability zones that can insulate an application from certain types of failures in a cloud provider's infrastructure. It also includes developing an understanding of how making certain choices (e.g., using block versus object storage) can potentially create (or protect against) particular failure modes.

Many types of applications can be readily transplanted to a public cloud. But taking full advantage of a public cloud's ability to rapidly scale computing capacity up and down in response to changes in demand and to rapidly provision new services as needed requires applications that have been optimized to take advantage of such features.

MAINTAINING PORTABILITY OF APPLICATIONS AND WORKLOADS

Portable applications can run in multiple environments without changes. This is the true test of portability. With true portability, developers can write once and deploy anywhere, thereby keeping their options open to preserve flexibility. Furthermore, application portability ensures that IT organizations can deploy workloads on new and existing infrastructures or on the public cloud of their choice based on business and IT governance requirements.

A number of different aspects play into workload portability. One of the most important is having the ability to deploy certified operating systems and middleware across a variety of providers both internal and external. These common runtimes effectively act as a container for applications running in different environments and help insulate the applications from platform specifics, which they might otherwise have to explicitly deal with. (This sort of abstraction has historically been one of the most important functions of an operating system.)

Of course, ensuring portability of applications running in a given public cloud is not solely a technology issue. Maintaining portability also requires that applications not be written in a way that locks them into a single platform—for example, by using platform-specific features or interfaces that are specific to a particular public cloud provider. While it sometimes will make sense to take advantage of technology unique to a single public cloud provider, one should do so with a full understanding of the tradeoffs involved in doing so.

*1 See the Red Hat whitepaper **Cloud Infrastructure for the Real World** for a more detailed discussion of the differences between traditional enterprise workloads and cloud-style workloads.*

ENSURING PORTABILITY OF DATA

Discussions of portability tend to focus on applications and the compute services that they run on. However, it's not very helpful to move applications if you can't also move the data those applications are operating on. There are several aspects to consider when moving applications and their data to a public cloud.

One is the "laws of physics" or, less flippantly, latency and bandwidth considerations associated with applications operating on that data or moving it around. For performance reasons, it's often desirable to keep applications and their associated data nearby each other. As a result, if you want to work with a large set of data that you're generating locally, it may make sense to do the analysis or processing of that data locally as well. The converse is also truel. If your data is in a public cloud, perhaps you should run your Hadoop cluster or other computing tools there as well.

There's also an economic angle associated with large volumes of data. While pay-by-use pricing is often one of the features that attract users to public clouds, it's not always going to be the cheapest option, especially when large amounts of data are going to stored for an extended time. These are some of the conceptual aspects of data portability.

From a technology perspective, a unified data and scale-out storage software platform can accommodate files and objects deployed across physical, virtual, public and hybrid cloud resources. Having a common distributed storage approach across your entire infrastructure helps to break down pools of capacity and simplifies moving data from one platform to another if you want to do so for either business or technical reasons.

UNDERSTANDING OF LEGAL AND REGULATORY COMPLIANCE

LEGAL AND REGULATORY QUESTIONS MIGHT INCLUDE:

- Does the provider comply with PCI DSS or HIPAA standards?
- What notification processes does the provider have in the event of a data breach or a subpoena of data?
- Are there laws restricting certain types of data to a particular country or other locality?
- What types of audit information does the provider make available to customers?

Data also looms large in any discussion of legal, regulatory, or broader governance concerns involving public clouds. It's not so much a matter of classical security concerns like the "Is the cloud safe?" question posed by a headline writer looking for clicks. Rather, it's about considering a wide range of privacy, regulatory compliance, legal, availability, durability, and confidentiality concerns—many of which are primarily about data rather than other aspects of compute infrastructure and operations.

Documents such as the Cloud Controls Matrix from the Cloud Security Alliance are useful sources for understanding the types of controls that may be relevant to a cloud provider and their customers and how they relate to specific standards and regulations.[2] A variety of cloud-specific standards and certifications aim to better clarify best practices and even requirements for particular uses. For example, FEDRamp provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud services and products across the entire U.S. federal government.

For reasons of greater visibility and perceived control (if nothing else), many organizations still feel more comfortable hosting their more-sensitive data internally. We're not going to argue for or against that position here but, rather, note that data breaches happen everywhere. Wherever and however data is stored, it's increasingly important to understand and mitigate the risks associated with storing it and using it.

2 https://cloudsecurityalliance.org/research/ccm

ENABLING HYBRID CLOUD MANAGEMENT, POLICY, AND GOVERNANCE

In a 2012 Forrester report, James Staten and John R. Rymer noted that "Management in the future of cloud shifts to providing, measuring, and using portfolios of services."[3] This observation is consistent with trends that we see around the adoption of public cloud services of many types. Public clouds may have initially entered organizations as ad hoc services paid for using departmental credit cards. However, in this future hybrid IT world, they're often becoming more strategic elements of an overall portfolio.

Realistically, the management of these portfolios won't always be pulled into a single master console. Infrastructure-, Platform-, and Software-as-a-Service abstractions (IaaS, PaaS, and SaaS respectively) often benefit from different management approaches. Furthermore, federating the management of services will often be a better fit organizationally than trying to centralize everything.

That said, it's certainly desirable to bring the management of hybrid infrastructure—whether the heterogeneous virtualization platforms and physical servers of a datacenter or the IaaS sourced from a public cloud provider—under consistent provisioning, monitoring, and metering controls. Furthermore, as a broader principle, consistent policies and risk management principles should be applied to sourcing IT services in general, in whatever form they take and wherever they come from.

ISOLATION OF WORKLOADS IN A PUBLIC CLOUD

It's worth briefly taking note of a subtlety that we've largely ignored to this point. We, like most discussions on the topic, have painted private and public clouds as if they were a binary choice—dedicated systems in a datacenter you own and control or a multi-tenant public cloud in which your applications could be sharing physical servers, disk drives, and networks with an arbitrary number of others.

That, of course, is a gross oversimplification. Many organizations today run owned or leased servers in shared datacenter facilities and understand that complete control over a private cloud is something of an idealization in many cases.

From the perspective of public clouds, however, one particular approach is of interest in this context. A virtual private cloud (VPC) lets you provision a logically isolated section of a public cloud in which you have complete control over your virtual networking environment, such as subnets and configuration of routing and network gateways. This allows you, for example, to create a public-facing subnet for your web servers that has access to the Internet, while placing backend systems such as databases or application servers in a private-facing subnet with no public Internet access. Furthermore, you can then create a VPN connection between your corporate datacenter and your VPC—making it effectively an extension of your in-house datacenter.

Although not all providers offer the option, a VPC is a powerful tool for providing consistent and integrated management across both a private and a public cloud—as well as providing an additional level of isolation within the resources of a public cloud provider. In addition to VPCs, many cloud providers also offer dedicated servers, allowing their customers to further isolate their workloads, either for performance or security reasons. This approach may be more costly than using multi-tenant servers, but in some cases it is an absolute customer requirement.

3 Forrester Research, Inc. Cloud Keys An Era Of New IT Responsiveness And Efficiency, November 2012.

ACCESS TO STANDARD LANGUAGES, FRAMEWORKS, AND TOOLS

Cloud application development is multi-lingual and needs to be able to integrate and extend existing enterprise applications and processes. Red Hat has conducted surveys at several events that consistently reveal that most respondents intend to develop software for cloud environments using similar languages as they use today.[4] PaaS platforms, for example, that limit developers to a specific language on a specific hosting platform are frequently criticized by developers because they constrain this choice. It's telling that a number of language- and framework-specific PaaS solutions have shifted toward a more polyglot (multiple languages/frameworks) approach.

It's also important to deliver a consistent development and deployment platform across hybrid cloud environments. Portable programming lets developers create software with the tools of their choice without being tied to a specific vendor's toolset or platform. There's no need for developers to rewrite applications or learn new skills when moving to a new environment. This also gives developers access to the broadest range of innovation and the ability to choose lightweight frameworks, tools, and middleware for fast and easy application creation, while also retaining access to everything needed for enterprise-class development.

The message here is that your chosen public cloud providers should provide the tools you want to use rather than constrain your choices to what they happen to make available.

HOW RED HAT CAN HELP

Launched in 2009, the Red Hat Certified Cloud Provider Program assembles the solutions cloud providers need to plan, build, manage, and offer hosted cloud solutions and Red Hat technologies to customers. A Red Hat Certified Cloud Provider offers a trusted destination for Red Hat customers, independent software vendors (ISVs), and partners to benefit from Red Hat offerings in public clouds under innovative consumption and service models delivered by the cloud providers.

The Red Hat Certified Cloud Provider designation is awarded to Red Hat partners following rigorous validation by Red Hat. Each provider meets testing and certification requirements to demonstrate that they can deliver a safe, scalable, supported and consistent environment for enterprise cloud deployments. The program provides customers, ISVs, and partners with the confidence that Red Hat product experts have validated the solution so that implementations begin with a solid foundation.

4 While there's much ongoing innovation in middleware, tools, and languages, there's also considerable inertia in practices and established skills which any organization has to take into account when planning development methodologies.

In addition, Red Hat Cloud Access enables qualified enterprise customers to migrate their current subscriptions for use on Red Hat-certified clouds. This gives customers the ability to make use of Red Hat support, relationships, and technology on certified clouds, while maintaining a consistent level of service and support across all certified deployment infrastructures with consistent and predictable pricing.The program works in concert with a broad range of Red Hat open hybrid cloud products, including:

Red Hat Enterprise Linux remains the operating system of choice for enterprises seeking a portable, open, scalable, customizable system that is ideal for building applications in the cloud. It is the cornerstone of our vision for the open hybrid cloud–a vision that started long before the concept of the cloud became fashionable and is fueled by an operating system that was born in the cloud and serves the companies moving to cloud. Some of the largest clouds in the world are built on Red Hat Enterprise Linux and other Red Hat technologies.[5]

Red Hat CloudForms supports web-based access to your specified service catalogs with role-delegated, automated provisioning, quota enforcement, and chargeback across a hybrid cloud environment. With CloudForms, resources are automatically and optimally used via policy-based workload and resource orchestration, ensuring service availability and performance. You can simulate allocation of resources for what-if planning and continuous insights into granular workload and consumption levels to allow chargeback, showback, and pro-active planning, and policy creation. CloudForms allows you to consolidate management of both on premise and AWS virtual machines by provisioning Amazon Machine Instances (AMIs) in a policy controlled manner, through enterprise defined self-service portals and service catalogs.

OpenShift Online by Red Hat is Red Hat's public cloud application development and hosting platform that automates the provisioning, management and scaling of applications so that you can focus on writing the code for your business or next big idea. A choice of programming languages, including Java, Ruby, PHP, Node.js, Python, and Perl, and a complete set of developer tools are available within OpenShift Online to increase developer productivity and accelerate application delivery. And integrated development tools and intuitive interface enable you to get started quickly. No new programming models, no app changes, and no cloud lock-in.

Red Hat JBoss Middleware provides a portfolio of tools to help developers and companies build applications in public and private PaaS environments. With the advent of xPaaS, users can also move beyond the limits of simple application development to the next generation of PaaS technologies and capabilities. xPaaS is a rich set of application development and integration capabilities that will enable users to build and deploy complex enterprise-scale applications.

Red Hat Storage Server brings enterprise-class features to big-data environments, such as geo-replication, high availability, POSIX compliance, disaster recovery, and management—without compromising API compatibility and data locality. Customers now have a unified data and scale out storage software platform to accommodate files and objects deployed across physical, virtual, public, and hybrid cloud resources.

---

*5 See the whitepaper **Red Hat Enterprise Linux — The original cloud operating system** whitepaper for more information on Linux adoption. trends for cloud deployments.*

CONCLUSION

Contrary to the scare headlines of the week, public clouds aren't inherently unsafe or risky. Rather, like any other IT asset or service, they should be evaluated, adopted, and operated in accordance with the same sort of due diligence and attention to benefits and costs (including legal or regulatory exposure) that should apply when sourcing any type of service.

This means that choosing a public cloud provider on which to run your enterprise workloads is no longer an isolated tactical decision that can be readily undertaken without consideration for the bigger picture. Rather, public clouds increasingly should be viewed through the lens of the portfolio of services IT brings together and manages in order to meet business needs while managing risk to appropriate levels.

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

redhat.com
#11602217_v1_1113

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