

BEYOND VIRTUALIZATION TO CLOUD

Cut through the cloud clutter. Refocus your initial virtualization efforts.

Design cloud services, not systems. Changing the form of what IT offers.

Optimize and automate IT in the cloud. Cloud is built, run, and governed differently than the static virtualized datacenter.

Accelerate business results with your cloud. Users turn to IT as the preferred service provider.

EXECUTIVE SUMMARY

Cloud isn't virtualization. It's more dynamic. It's hybrid. But, most fundamentally, it's a different mindset. Cloud isn't about servers, even virtualized ones. It's about services. Building a cloud involves designing a catalog of standardized services—think of them as application or development environments—and offering them to consumers, such as developers, through a low-touch, self-service interface. Access to these services is controlled by policy, as is the runtime management (such as patching) of these environments after they are deployed.

Although hybrid clouds can be adopted in an evolutionary way, this change in mindset means that cloud deployments, like other strategic IT projects, benefit from a degree of process—even if it's lightweight process. This whitepaper outlines one such approach based on research from the IT Process Institute. It consists of four steps:

- 1. Cut through the cloud clutter.** Refocus your initial virtualization efforts on skills and competencies that support hybrid cloud deployments.
- 2. Design cloud services, not systems.** The key to cloud success and to minimizing shadow IT is not just speeding up delivery of servers, network, storage, and other computing resources, but also changing the form of what IT offers.
- 3. Optimize and automate IT in the cloud.** While a cloud typically utilizes virtualized resources, it is built, run, and governed differently than the static virtualized datacenter.
- 4. Accelerate business results with your cloud.** Broad adoption occurs when users have enough confidence and trust in the cloud solution that they turn to IT as the preferred service provider.



facebook.com/redhatinc
[@redhatnews](https://twitter.com/redhatnews)
linkedin.com/company/red-hat

INTRODUCTION

Some new technologies make their way into organizations at the periphery. Perhaps they perform some specific task that is outside of the day-to-day concerns of IT management. Perhaps they're an ad hoc tool of some sort—useful but not part of any formal workflows or procedures. Perhaps they increase efficiency in a way that can be adopted incrementally one server or one group of applications at a time.

In the early days, virtualization mostly fell into this latter category. During the early 2000s, many companies anxiously sought ways to avoid purchasing servers and other IT gear. Server virtualization fit the bill perfectly. As virtualization has become more widespread, IT shops have started approaching it more strategically. But it started out as a tactical, cost-cutting move.


By contrast, for reasons that this whitepaper will address, cloud computing is fundamentally more strategic. And, therefore, usually best approached systematically and with a degree of rigor and process. That's not to say heavyweight process. However, by bridging IT silos, automating actions, and providing self service to users, a cloud delivers a powerful tool to make your IT infrastructure more flexible and responsive to the business. Wielding that tool most effectively takes some upfront planning.

This paper focuses on a framework developed by the IT Process Institute. It's fairly lightweight, is based on discussions with organizations that have begun implementing clouds, and dovetails nicely with the experiences of Red Hat's services organization.

It consists of four steps:

1. Cut through the cloud clutter.
2. Design cloud services, not systems.
3. Optimize and automate IT in the cloud.
4. Accelerate business results with your cloud.

The methodology presented in this whitepaper summarizes material from a series of whitepapers that Kurt Milne of the IT Process Institute prepared for Red Hat. It is based on an approach published in *Visible Ops Private Cloud: From Virtualization to Private Cloud in 4 Practical Steps* by Kurt Milne, along with Andi Mann and Jeanne Moran.



1. The focus of this whitepaper is on “private clouds” being built within an organization’s own datacenter or otherwise on dedicated servers that they control. However, everything discussed here also applies to hybrid clouds, both in the sense of clouds that span both public and on-premise resources and in the sense of clouds spanning heterogeneous technology platforms within an organization.

STEP ONE: CUT THROUGH THE CLUTTER

The goal in this phase is to refocus your initial virtualization efforts on skills and competencies that support hybrid cloud deployments.¹ The initial discovery pilot phase will enable the identification of challenges, requirements, and key metrics that will prepare you for the larger cloud implementation. Your mantra for these activities is, “Get ready for dynamic workloads.” You should set end goals for virtualization and private cloud deployment. You should start laying the groundwork for building shared resource pools and for managing mobile and transitory workloads.

Getting ready to build a cloud consists of five steps. The first three primarily involve planning.

Set cloud goals based on business objectives. Building a cloud designed specifically for your enterprise has to start with a business discussion. If your infrastructure group is starting a cloud project without developers on the team, stop them. Round out the team with developers, users—and, more importantly, externally facing product, marketing, and sales managers. Engage all stakeholders in a discussion about how cloud can accelerate business processes or transform business offerings. Establish clear objectives and success criteria in business terms.

Adopt a portfolio view of your infrastructure. As you move ahead with a cloud strategy, you’ll most likely be managing a mix of physical, virtual, and cloud resources. As a result, you will allocate a portion of the datacenter as a pool of shared, virtualized, and scalable resources. Many IT executives plan to put 30-50% or more of workloads in their private cloud environments. However, private cloud resources will be managed in an environment with physical servers and mainframes, as well as static virtualized resources. To put it in real-estate terms, building the cloud-centric datacenter of the future will be a remodel, not a tear-down. Doing so requires understanding the key attributes of current workloads, scoping the mix of heterogeneity of current environments, and examining how requirements change as you progress from development through test/QA to production.

Target workloads for the cloud environment. Assess your current workloads to identify those that are a good fit for a hybrid cloud. This snapshot will be used to set long-term targets for the percentage of overall workloads targeted for the cloud. In the short term, it will also be used to identify workloads for initial cloud deployment.

Then move beyond planning. Get hands-on with two critical activities:

Evaluate cloud computing models. Evaluate different models in the context of your objectives. Be sure to consider agility, service quality, cost, and security and compliance. Consider hybrid computing models that utilize internal and external cloud resources. Note that private cloud resources may include resource pools hosted by an external service provider (but under your control). A hybrid model may include features that allow movement of workloads from private cloud to external public cloud service providers.

Deploy a proof of concept based on a standard architecture. Deploy vendor solutions in-house and determine how higher levels of automation and standardization integrate with your existing infrastructure, processes, and skills. The overall goal of a proof of concept is to demonstrate success with a working reference implementation based on business requirements. To get there, you must test the assumptions you made during your evaluation.

STEP TWO: DESIGN SERVICES, NOT SYSTEMS

Hybrid clouds offer users fast access to computing resources similar to those offered by public cloud providers. (Indeed, an overarching requirement for any on-premise cloud has to be, to users, it's as easy and flexible as using a public cloud would be.) However, deploying raw compute resources, either on internal private or external public resource pools, is the lowest common denominator in cloud. The key to cloud success and to minimizing shadow IT is not just speeding up delivery of servers, network, storage, and other computing resources, but also changing the form of what IT offers.

Users are thrilled to get self-service access to cloud services within 15 minutes. But success for hybrid cloud initiatives requires joining self-service cloud access with the traditional enterprise IT need for governance, security, compliance as well as world-class service delivery and business continuity. A thoughtful service-design approach that shifts focus from resources to delivery and consumption of IT as a service can help meet both user and IT requirements.

A service-design approach includes understanding business objectives, detailing specific user needs, defining services that meet those needs, and defining the functional and technical specifications needed to deliver those services. It also includes creating an IT “factory” to build and deploy workloads in simple or complex cloud environments both at internal and external resource locations. These processes require clearly defined policies that specify what, how, where, and when workloads are deployed, whether to deploy in public or private clouds, static virtual environments, or even physical dedicated servers.

Key activities related to building the right thing include:

DESIGNING BUSINESS-OPTIMIZED SERVICES

It is not sufficient to replace physical servers with virtual servers in a cloud environment.

Hybrid clouds should include fast access to applications and workloads that are fully configured and functional when deployed. The service definition and process to deploy the service can then be improved and evolved based on customer feedback, lessons learned, changing requirements, and maturity of technology.

Basic service design frameworks include:

- Defining core services such as a mobile development or a ready-to-go drug research environment.
- Specifying supporting services such as backup, high availability, security, and network configurations.
- Providing service-level options that address such factors as performance, resource allocation (CPU, memory, I/O, network, and storage), business impact, business continuity and disaster recovery.

SPECIFY AND CERTIFY TEMPLATES

Once you have defined services for specific users, you need a way to break down services into components that can be assembled to enable those services in a deterministic and predictable way.

A common virtualization approach has been to deploy monolithic images that include components from the operating system through to applications. With this approach, images quickly get out of date. If you change one component, you must recreate all images containing that component. Additionally, you may need different images for different phases of the application lifecycle. For

example, when you launch an application server for testing software, you want the resources and configuration to match the production environment. However, you don't want the server configured to send email alerts to the operations team. You may also need variants of images for provisioning in different private or public cloud environments.

A better approach is to provision cloud services from a set of templates that are generic resource and configuration definitions. Templates are assembled to deploy services based on rules for each environment and application lifecycle phase.

CLARIFY DEPLOYMENT AND BUILD POLICIES

Part of the problem with rogue purchase behavior is that IT loses control of the technology and information for which it is ultimately responsible. For example, with private cloud, a user may request four web servers to put sensitive data in an unsecured network zone. Policies should prevent that from happening. To control self-service cloud deployments, you need two types of policies to guide cloud provisioning: deployment policies and build policies.

It may help to think about policies in terms of building a decision tree. Input to the codified policy workflow should include predefined information that is collected during the service request. Part of the service request process should, therefore, include asking questions that collect the information needed to satisfy the policy.

AUTOMATE REPEATABLE BUILD AND DEPLOY

A trustworthy on-ramp to cloud for simple or complex workloads ensures that services are deployed the right way every time. Cloud provisioning requires an IT factory model in which machines build machines based on a specific bill of materials. The process should be automated and highly standardized. Variations should be the exceptions.

What does a cloud IT factory model look like? It combines templates and deployment policies in a way that deploy workloads exactly the same way every time. Workload deployments must be described in a structured format that may include the bootable operating system, any software components, configuration provided or required, and specific targeting information to instantiate the workload.

CREATE A SELF-SERVICE ORDER MECHANISM

Once you have implemented an automated, repeatable way to deploy cloud services, you can then add self-service access for users. Self service gives users fast access to technology while maintaining IT control. A one-touch order allows users to select a bundle or build a service request from a list of service offerings, supporting services, and service levels.

The primary focus of the self-service mechanism is to offer business-optimized services that are designed for the user. Offer service packages using terms that make sense to requestors. For example, a developer LAMP stack, marketing website, or collaboration toolset package should include everything needed for the most common use cases.

STEP 3: OPTIMIZE AND AUTOMATE IT IN THE CLOUD

What about after the workload is deployed? Who maintains and updates the cloud? How does IT ensure ongoing security and compliance?

While a cloud typically utilizes virtualized resources, it is built, run, and governed differently than the static virtualized datacenter. As a result, IT must address unique run-time challenges such as shared resources, massive scalability, standardized systems management, and hybrid and heterogeneous solutions.

Understanding and addressing these differences is critical for cloud success. So is taking the right steps to optimize run-time activities. These steps are:

EXPANDING AUTOMATION

Automated provisioning provides on-demand access to services in service catalog. But automating build and deploy is only part of managing a cloud environment. Ad hoc and manual management of labor-intensive, error-prone cloud management activities doesn't address cloud scalability requirements, nor does it allow you to optimize service delivery and resource utilization. To address those goals, you'll have to ratchet up the level of automation.

Automation spans many areas of IT operations. It includes workload moves, resource scaling, backup and disaster recovery, application lifecycle management, and retiring resources that are no longer needed. The bottom line is this: cloud run-time management requires automation. And automating time-consuming, error-prone maintenance activities is essential for efficient and reliable cloud service delivery.

MANAGING HYBRID HETEROGENEOUS CLOUD ENVIRONMENTS

Adding cloud software to a pool of computing resources enables self-service deployment of services and allows IT to respond to changing usage levels. But achieving very high degrees of scalability in a dedicated, on-premise cloud environment may result in underutilized resources that sit idle during normal usage levels.

To optimize utilization and achieve extreme scalability, consider a hybrid cloud strategy in which workloads are deployed across both internal resource pools and resources managed by third-party IaaS cloud providers. A hybrid approach can offer more options for scalability while maximizing the utilization of internal computing resources.

UPDATING SERVICE MANAGEMENT PROCESSES AND DOCUMENTATION

Because clouds are built, run, and governed differently than static virtual environments, the processes you use to manage the run-time environment must be updated for cloud. Provisioning a normal workload into shared resource pools changes capacity planning, which, in the past, was typically tied to a static project funding and planning cycle. Automation of resource changes and workload moves poses tracking, monitoring, and support issues not found in static environments. Giving users self-service access to production resources violates traditional controls that require advisory board review of every production change.

ENABLING CONTINUOUS COMPLIANCE

In cloud environments, you should strive for a state of operations in which machines build and maintain machines, and continually sense and respond to unauthorized changes. A "fire and forget" approach for deploying cloud services can jeopardize cloud goals.

Compliance is achieved through preventive, detective, and corrective controls that make it hard to do the wrong thing, that immediately detect when right thing hasn't been done, and then alert staff and restore conditions to the desired state. To accomplish this, you need a way to check compliance not based solely on audit output or artifacts of build routines, but also on audit blueprint and automation rules that produce the artifacts. It may sound more complicated, but it's actually good news. Instead of checking every server to verify that it is at patch level, you audit to see that every server matches your blueprint and verify that the blueprint is at the desired patch level.

STEP 4: ACCELERATE BUSINESS RESULTS

A hybrid cloud can remove much of the typical IT friction associated with growth and innovation efforts. But cloud offers more than speed. Cloud can improve utilization of computing assets. And, cloud can increase workflow efficiency for a wide range of IT operational processes.

But the cloud's "better, faster, cheaper" value proposition has a critical dependency: broad adoption within the organization. Building something better is a wise use of IT resources only if users adopt what you build. Otherwise, you may be leaving money on the table and undercutting the value IT can offer the business.

Broad adoption occurs when users have enough confidence and trust in the cloud solution that they turn to IT as the preferred service provider. Key activities related to maximizing utilization include optimizing the economics, reshaping user behavior towards IT as a service, streamlining processes to increase collaboration, and shifting towards service-oriented accounting. Some of these steps will happen sooner in some organizations than others. And some will be more important in some organizations than others. But these sorts of steps lead to better aligned IT activities and business outcomes.

OPTIMIZE PRIVATE AND HYBRID CLOUD ECONOMICS

The simplest way to optimize cloud economics is to minimize the cost of building and maintaining the solution, and then move as many workloads as feasible into this new, optimal environment. Admittedly, not all workloads belong in the cloud. But by using the same tools to manage physical, virtual, and cloud environments, you can derive additional value from your investment. The point here is that the economics improve with higher utilization of your cloud environment, your processes, and your tools. This leveraging of heterogeneous and hybrid environments is an important part of how open hybrid clouds are different from a more silo-ed approach.

RESHAPE USER BEHAVIOR TO CONSUME IT AS A SERVICE

Cloud creates a unique opportunity to remake how IT value is delivered as a service. Getting users to consume IT as a service drives full adoption so you realize maximum business value out of cloud.

Cloud does more than speed up the provisioning process. It can fundamentally change how resources are consumed. With on-demand access to predefined services, users can get what they need when they need it. With policy- and automation-based provisioning, the service works the same way every time. Speed and consistency create opportunities for users to order and consume in a fundamentally different way.

OPTIMIZE PROCESSES TO STREAMLINE CROSS-FUNCTIONAL COLLABORATION

A critical point in the application lifecycle is the handoff from those who write code to those who support datacenter operations. How that handoff is handled can inhibit or accelerate business results. Inherent in this handoff is a conflict between the developers who create a product (code) and the system administrators who deliver services. The conflict represents a gap between groups that deliver value in different ways. For one group, value is measured by speed and agility. For the other, it is measured by efficiency and stability. One group is responsible for meeting functional requirements while the other is responsible largely for meeting nonfunctional requirements.

In siloed organizations that do not address this conflict, the gap can slow down application and patch release cycles, reduce quality measured by both code and service levels, and cause excessive overhead work for both development and operations personnel.

ADOPT SERVICE-ORIENTED COST ACCOUNTING

Traditionally, IT has been funded through a combination of business projects and annual budget allocation. Large upfront costs are typically tied to project funding. Ongoing management costs are typically tacked on as overhead. Alternatively, IT might allocate ongoing fixed costs across business units based on revenue or headcount.

In contrast, hybrid clouds create the opportunity to shift to service-oriented costing. With on-demand, self-service access you eliminate procurement headaches and IT operations bottlenecks. Instead, you provide resource pools that scale as needed and, from the user perspective, don't require capacity planning. The result: IT can allocate costs based on services delivered. This presents an opportunity to tie costing to services and give service consumers and service funders visibility into allocation and usage—even if the organization doesn't opt for full chargeback costing.

CONCLUSION

While delivering the right cloud services in a reliable and repeatable manner is critical for cloud success, it isn't enough. It takes a smart mix of process improvement and visibility to create truly compelling hybrid cloud solutions. IT organizations that achieve the right mix are able to gain the confidence of business users and, consequently, affect behavior change that unleashes the value of cloud computing.

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.



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

redhat.com
#10781557_v2_0413

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
latammktg@redhat.com