

PRIVATE CLOUD Is STILL AN OPEN LOOP

eBook



A publication of

EVOLVEN

This eBook is presented to you by Evolven

Evolven. Configuration and Change Management for the Modern Data Center and Cloud.

Evolven's designed from the ground up to deal with the dynamics and complexity of the modern data center and cloud environments.

- ✓ **Dynamically crawls** and collects comprehensive environment configuration information
- ✓ Gives you a **consolidated view** of your **entire application stack**
- ✓ Captures and analyzes changes in configuration parameters down to the **most granular level**
- ✓ Applies **analytics** to deliver actionable information
- ✓ More: See our website to see Evolven in action

Free Trial



Contents

Demystifying Private Cloud	4
Private Cloud Promise: 100% Automation	5
Painful Real World Private Cloud Challenges.....	6
Rollout And Rollback – A Viable Strategy?	7
The First Private Cloud Decision You Have To Make.....	11
Change and Configuration: A Critical Part of the Picture	13
Closing the Private Cloud Management Loop.....	16



When cloud computing emerged as IT's main topic of discussion, the focus was on public clouds, mostly due to the high-profile nature of public-cloud vendors and accessibility. But now as we turn to implementation, many IT organizations have stepped away from public clouds because of concerns about security, storage of sensitive data outside the corporate firewall and basically the lack of control on the part of IT.

Tom Bittman, VP Distinguished Analyst at Gartner, says, **“2012 will be the year that private cloud moves from market hype to many pilot and mainstream deployments.”**

However, even with the knowledge of the power that private cloud holds, IT executives are now coming to see **the challenge of making the transition** to the new private cloud infrastructure. For most, provisioning private cloud is uncharted waters and a task that is one huge step into the unknown.

To bring more clarity to this effort, we illustrate and explore the key technical challenges that lay ahead for even the most well planned private cloud implementations.

Demystifying Private Cloud

Today industry experts and vendors apply different meanings for the term “private cloud”. Let's start with a definition for private cloud can serve as a basis for further discussion.

We will adopt the **NIST definition**, meaning that **private cloud is a deployment model where the cloud infrastructure is operated solely for an organization**. It may be managed **by the organization or a third party** and may exist **on-premise and off-premise**. The cloud itself is defined as a **model enabling on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction**.

Essential characteristics of the cloud include:

- On-demand self-service
- Broad network access
- Resource pooling
- Rapid elasticity
- Measured service



Private Cloud Promise: 100% Automation

The management for private cloud is expected:

- To operate a **shared pool of resources** to ensure their security, availability, performance and scalability
- To manage **provisioning and retiring** of computing resources
- To **define and deliver a catalog of services** to be activated and consumed by the users
- To **measure and report service utilization**
- To **manage service levels** provided to the organization

Today there is a variety of open source and commercial private cloud platforms aiming to address these requirements (Eucalyptus, Cloud.com of Citrix, VMware vCloud, RedHat RHEV etc.). In addition public cloud vendors deliver various forms of virtual private cloud off the customers' premises.

Traditional system management vendors and new entrants target cloud by augmenting cloud platforms with a suite of management tools for monitoring, provisioning, and reporting etc. Most of these tools are focused on **deployment automation** and **infrastructure monitoring**. There are certain attempts to introduce **server policy management** (e.g. Microsoft System Center or VMware Configuration Manager).

The result is that today's organizations can build a platform enabling a cloud management loop through almost any combination of free open source components, mid-range cloud management tools and commercial heavy weight system management solutions. However, this cloud management loop is **based on the assumption of complete, 100% reliable automation** as promoted by the vendors.



Painful Real World Private Cloud Challenges

Private cloud presents an opportunity to tremendously increase the agility of IT, allowing the IT organization to rapidly and efficiently respond to changing business requirements.

In the Forrester Report, *Market Overview: Private Cloud Solutions, Q2 2011*, James Staten and Lauren E Nelson explain the drive for private cloud saying, "Today's business executives are becoming more IT savvy, and most are demanding to have a "cloud strategy" for delivering more-efficient IT services. So it's no surprise that there's a rush by I&O professionals to get to "yes" on cloud computing, with a particularly strong desire to build a private cloud. Forrester surveyed enterprise hardware decision-makers in Q3 2010 and found that 6% stated that they have a private environment today, while another 25% stated that it was a high or critical priority for 2011." ([Market Overview: Private Cloud Solutions, Q2 2011](#))

However, at the same time, **private cloud introduces new management challenges to IT operations.**

100% Automation Is Not Achievable

While the amount of manual changes drastically decreases, 100% automation is difficult to achieve. Take the scenario for when a particular system experiences a critical performance issue. While investigating the environment experiencing the issue, you detect that certain infrastructure parameters need to be adjusted. As the result of the investigation, most probably these parameters will **be changed manually** on the servers experiencing the issues. Then the deployment automation platform will be updated. **In a very dynamic and complex environment, it becomes very difficult to track such manual changes which are frequently made at a very granular level for individual parameters.**

Moving Targets in the Dark

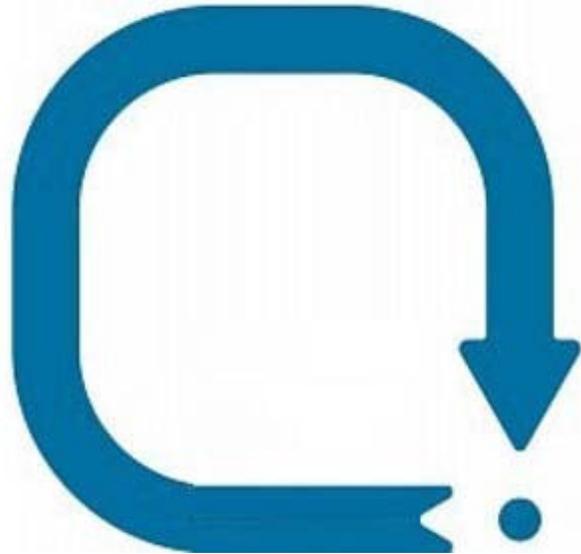
Components facing rapid changes. Following administrative tasks, new user or modified user requirements, issue resolution etc., **components of the deployment automation platform (scripts, templates, images etc.) continuously change.** At the same time, automated deployment encapsulates the spun off environments, offering the operations staff **limited visibility into the actual environments configuration.** The result is that it becomes very **difficult to correlate** actual changes in the environments whose automation assets were modified. Such disconnect causes environmental drift, configuration inconsistencies, longer issue resolution time and other risks and inefficiencies.



Rollout And Rollback – A Viable Strategy?

Rollout and rollback – is this a viable approach in large environments?

The idea that at any time virtual server images can be reinstalled, rolled back or re-imaged makes cloud computing an exciting option and suggests that it can dramatically reduce the overhead on the IT department. With the ability to roll back, you can recover rapidly from the deployment of an invalid, incorrect, or corrupt change, that compromises environment performance or availability, and puts the organization at risk. Yet, while automation of deployment makes cloud



infrastructure flexible and efficient, managing configuration is still a major challenge. The reason is that **in large environments, we've found that the management of configuration based on the rollout and rollback approach of the entire server images is less effective and problematic.**

Maintaining Environment Consistency

With the increased pace of changes streaming into the business system environment, environment management based solely on virtual image usage presents risk simply too high to tolerate. Due to the dynamic nature the Private Cloud, the many interdependencies and limited visibility into the actual environment configuration, **rollback of a system based on virtual images reset requires that the entire setup must be updated, and then synchronized, while making the upgrade and roll back of the entire business service is a challenge.** We observed numerous organizations leveraging private cloud for rapidly scaling up and down particular business systems while running these systems continuously.

Greater Risk from Retirement of Cloud Servers for Upgrade

Let's say that a certain business system is performing at the expected service levels, but you need to deploy a change to the software infrastructure. **Unnecessary operational risks** are introduced, when existing working virtual servers are retired, and you roll out new, replacement servers based on stored images that include the changes. We are seeing leading companies operating in the cloud **changing their practices to upgrade existing servers without retiring them.**

This landscape is going to become more intense as Thomas Bittman, VP and distinguished analyst at Gartner Research, anticipates, saying " We'll see about a 10X increase in private cloud deployments in 2012. Enterprises will find where private cloud makes sense, and where it's completely over-hyped.



We'll see successes – and there will also be a number of failures (we've seen some already)." ([Top Five Private Cloud Computing Trends 2012](#))

Servers Not Retired

Most often is the case that servers are not retired unless there is a need to scale the environment down, raising stability problems. This means that servers are provisioned using some base virtual image and then enhanced and upgraded using automated deployment on top of the base images. Then, exactly like in a traditional static data center, it is critical to maintain an ongoing assurance of environment consistency, as well as control over environmental changes. It is required to validate results of automated deployment (see [Automated Application Deployment Is Not Enough! 3 Reasons Why You Absolutely Still Need To Validate Your Releases](#)) and ensure that manual changes that can happen are detected and rolled into the deployment automation platform.

The Rollback Challenge

Frequent changes, fixes, and improvements, this is the nature of systems that are implemented in cloud platforms. That is until a problem is found. So what do you do?

You can fix the problem manually, which requires an investigation to figure out the cause, and that also means putting the system on hold. This takes time, and during that time the system is dormant, harming the business.

The alternative is to roll back to the state prior to the faulty change. Previously the idea of redeploying all of your servers just for a one-line configuration change was unthinkable. The cloud has changed that, with its fast and easy provisioning feature.

However, problems are often only discovered after a good amount of time after releasing to production. The more time that passes, the harder it is to roll back. Users get accustomed to new features, so you can't just take them off. Business and customer data is accumulated in updated schemas making it difficult to roll the database back to support application rollback. Or sometimes the deployed change was supposed to address a critical issue. Rolling back just brings the issue back.

Manual Rollback

The other option is to try to rollback manually, specifically selecting what areas to rollback. Not so fast! There are numerous gotchas to this scenario. Since instances can differ in many ways – say, contain different application data – you would need to configure those differences manually. Clearly, it's impractical to approach Cloud instances manually, the way that servers used to be managed on-premise. The complexity of the cloud rules out this option, making this too labor and time intensive.



SaaS-like Approach to Changes

In SaaS, when a new version is deployed, typically only a percentage of the users get new version. As the stability of the new version is verified, then gradually the base of users is expanded. So can changes to private cloud infrastructure be deployed similarly?

This option is not ideal either, as it adds more complexity and room for error to cloud management. This would mean that you have to manage at least two versions of the application and underlying infrastructure in production and design mechanisms, merging data coming from two versions.

Different Rollout Processes and Their Problems

Changes and updates to Private cloud can take place through several different deployment models.

Image Deployment

Golden images has been an attractive approach, with minimal effort and time that is required to start a copy of an existing image. Downside: The disadvantage of this method is that this is just combining machines 'AS IS'. End users have limited options; the image catalog likely will only contain the commonly used images but not less common combinations of components that might be required for specific user purposes.

Template-based Deployment

By following the template, a fully configured cloud instance can be deployed to any of a number of cloud environments. RightScale and Kaavo, for example, provide template-based cloud deployment. Downside: This means building and configuring templates that may miss ad hoc changes that took place.

Script based Deployment

This means using a script to build the image, that ensures the rollout has the latest version of necessary components. Downside: An operator needs to configure these scripts to ensure that the rollout will be error-free, just like in software development. So just as bugs can creep in during coding, you can have a rollout script that contains bugs, even after you used it to deploy.

Combination

There is no single way to assemble and deploy systems efficiently in a private cloud environment. Any of these approaches can be mixed and combined to roll out changes in the cloud. And all of the problems can be present in the combination approach as well.

Staying On top of the Configuration

Configuration management tools have been used by groups running big infrastructures with lots and lots of systems to manage. Yet the dynamism of the cloud brings more problems, even if they are only using a couple of server instances to run their systems. This means configuration management and change management tools need to be able to dynamically stay on top of the different states of the private cloud based servers, to know what changed and what is the impact.



Most of the existing cloud vendors are **missing this critical configuration management element, and are focusing almost solely on deployment automation**. James Staten, Vice President and Principal Analyst Serving Infrastructure & Operations Professionals at Forrester explains, "IT pros have most of the basic ingredients to cook up their own cloud-like infrastructure — but there's no recipe, and many ingredients just don't combine well. Complicating the story are the traditional infrastructure silos around servers, networks, and storage that must work together in a new, truly integrated way. Vendors like Cisco, Dell, EMC, HP, and IBM know you need packaged solutions that just work, but until recently they left too much of the burden on their customers." ([Are Converged Infrastructures Good For IT?](#))

Solution for Managing change in the Private Cloud

From this, we can conclude that **management of the actual environment configuration and the changes that happen in the environment is critical for successfully operating in private clouds**. To fully realize private cloud in the enterprise, organizations need a solution that can identify changes in near real-time, at a comprehensive and detailed level, to facilitate these deployment approaches . One thing is very clear: If your IT organization is not willing to make this investment for whatever part of its data center is transitioned to a private cloud, then it will not have a cloud that exhibits agile provisioning, elasticity and lower costs per application.



The First Private Cloud Decision You Have To Make

Today, most organizations are in the process of considering, testing or piloting private cloud for limited applications. But only very few organizations run fully operational private cloud platforms. Eric Damage, [programme manager for EMEA software and service group at IDC](#), expects that "the decision in the next year or two will only be about the private cloud. The bigger the company, the more they will consider the private cloud. The enterprise cloud is locked down and totally managed. It is the closest replication of virtualisation."



One of the first decisions that organizations must make about private cloud is: what is the scope of private cloud services to be offered?

The desired scope of service will affect the complexity of implementation, robustness of the service and ease of service management. As with many other things, it's a tradeoff between complexity of the service and ease of its operation.

Thomas J. Bittman, VP Distinguished Analyst, Gartner, explained "The business case for private cloud really requires the business to be involved. If the business, for some reason, sees no value in speed and agility, private cloud is likely a wasted investment. I've seen examples of private clouds deployed by IT without business involvement, and then – surprise, surprise – no one used it. Cloud Fail." ([Private Cloud and Hot Tubs](#))

From organizations we observed, there are several examples of services they are running in private clouds.

On-Demand Provisioning Of a "Clean" Virtual Machine

There is a number of standard operating system images that a user can choose to spin off. Typically IT narrows down user requirements to no more than just a handful of images. The machines are used by internal teams for development, testing and production of non-critical systems. The users are actually responsible for the setup and support of software infrastructure and applications running on top of the operating system.

Provisioning of a Software Infrastructure On-Demand

An example of such service could be spin off of a database server, an application server or some kind of software component and server combination. Organizations find that a number of required



configurations are much greater than in the operating system case. This happens due to the specific requirements imposed by existing business applications on the underlying infrastructure.

Transition of a Business System to the Elastic Infrastructure

This setup is practically transparent to the business users that just enjoy better performance and stability. However implementation of a business system in private cloud has direct impact on the associated operational costs.

Software-as-a-Service (SaaS)

A user is able to get application access on-demand either through a dedicated or shared application instance. We saw several examples of such services in organizations – SharePoint, BI systems, IT management tools etc. Architecturally such services are very sophisticated as automatic scalability management is required to provision on-demand and further operate the applications.

Decide What's Best for You

Typically we see companies starting at the level of "clean" virtual machines and evolving from there over time to SaaS. Clearly SaaS provides the ultimate value to business by delivering rapidly and cost efficiently robust and scalable business systems. Some systems will be developed from scratch providing native support for the cloud.

"There are a lot of misperceptions about what it takes to get your private cloud investments right and drive adoption by your developers...your cloud should be smaller than you think, priced cheaper than the ROI math would justify and actively marketed internally," says James Staten, VP Forrester ([Getting private cloud right takes unconventional thinking](#)).

In summary, it's critical for organizations to look beyond and not just get caught up in the hype of private cloud, but make informed decisions so that private cloud solutions that are deployed will be suitable for your business's situations.



Change and Configuration: A Critical Part of the Picture

Private cloud infrastructure enables IT to serve the business by allowing new services to be quickly provisioned and by adding elasticity to better handle service availability and demand spikes. Yet realizing this capability isn't so easy.

In transitioning to private cloud services, change and configuration management, one of the key areas of IT management, is strongly affected. **There are many configuration management challenges introduced in transitioning to a private cloud-computing infrastructure exacerbated by the state of abstraction and pace of dynamic change in private cloud environments.**



Migration Management

When transitioning a working application from a physical data center to private cloud, it is essential to preserve key configuration information. You need a detailed inventory of configuration information of your applications and a clear understanding of all the components, and their relationships. You need to be able to understand and identify what configuration parameters will need to be updated when moving to the private cloud, and how the configuration impacts performance, availability and security of the application.

Today the added challenge is that this information gathering process is largely manual, making it harder to gather detailed knowledge of the application, its configuration and dependencies.

When this configuration needs to be changed, you need to know what the changes are and why they are made. There needs to be a way to detect existing configuration to plan migration to private cloud, and then compare configurations across environments while carrying out the migration. This requires gathering the essential and critical, configuration information, which is a huge undertaking, going through literally thousands of configuration parameters per technology – some critical and some not.

Image Library Management

One of the keys to making a large scale private cloud practical is the ability to rapidly provision and deploy new virtual machines. This entails creating sets, or libraries, of images on which additional deployments will be based. Virtual images libraries are the foundation for running many private clouds. Yet, effectively and efficiently managing this collection of images is not a small undertaking. A centralized repository of images allows you effective control over image versions, user access, changes



and much more. This library provides easy integration to the component(s) that drive provisioning into your private cloud.

Transitioning to image based provisioning however shifts the focus of configuration management to the image library. Virtual servers must be correlated with the images used to spin off the servers. To maintain control over the environment, you must analyze the consistency between configurations of the same purpose servers.

Highest Resolution

The dynamic nature of the private cloud can significantly increase the rate of change and complexity. This makes understanding and tracking actual versus desired configuration states very challenging. Virtualization in the private cloud introduces abstractions to the data center encapsulating virtual machines in a kind of black box. In addition the automated deployment layer hides changes to the actual environment configuration.

Limited visibility to the underlying configuration makes it difficult to identify the root cause of faults, leading to slower application deployment times, impaired service levels, and more risk from human error.

Understanding the change requires visibility into the configuration at the highest resolution level of individual configuration parameters: registry keys, data schema fields etc. This makes it quite challenging to correlate a change in a deployment script to the actual impact this change has in the deployed environment. Furthermore manual changes might still happen in the private cloud environment in emergency cases. Typically these will be very granular, carried out to resolve a specific issue. Configuration tracking needs to take a deep dive, at the highest resolution level, into the application configuration to catch subtle manual changes.

Validation of Automated Activities

It seems obvious that any software change should be tested before being rolled out into production. Automation of cloud management tasks such as provisioning and deployment is no different. You should validate the results of such automated activities. Organizations need to leverage automated technology to help ensure that all private cloud deployments are properly vetted so that failures in the course of provisioning do not produce unsynchronized environments.

One of the most penetrating of the automated management actions is a change in the existing environment. Change validation integrated with automated activities should become a regular part of change management process for private cloud for effective change management, configuration management, and release management processes, allowing organizations to better ensure that environments run smoothly, without disruption or outages, even when updates are introduced into production. Much time can be saved in the problem analysis phase by having a definitive list of what has or has not changed in the configuration of a service and infrastructure.



Single Point Of View

In many cases the environment running in private cloud is essentially managed by two parties:

- The team operating the private cloud
- The users

The extent of users' involvement depends on the type of service offered. In any case, when several parties are involved it is essential to have a single point of view IT teams need to centrally manage application, server, and on the end-to-end environment configuration and changes from a single console, with a summary dashboard that displays the current state.

A typical situation where these two parties need to coordinate is when users and the cloud operations team argue over responsibility for an incident. The ability to know what changes which party did can help to avoid finger pointing and significantly reduce investigation time.

Private Cloud Is a Highly Dynamic Environment

The high pace of change is due to the built-in scalability management and user triggered on-demand services. You cannot rely on manual processes to register and track changes since most of the activities are run automatically.

To get the environment under control, the configuration should be automatically monitored. Proactive monitoring is the only way to ensure that your team knows what's going on before the phone rings. Administrators need to be in a position to monitor the data center infrastructure and drill down to manage any risks or faults in the environment as they occur. At the same time, it is critical to provide visibility across the entire stack of private cloud layers at the highest level of resolution.

While the amount of resulting information can be overwhelming, change management teams need to rely on analytics. This analysis should be capable of interpreting the results of automatically collected granular configuration information, and produce actionable findings for driving operational decisions.

An example of such analysis could be when a certain base image is updated. The base image is used to spin off a set of new servers, while, still up and running, the servers are produced from the older version of the image. Analysis will produce a specific list of critical configuration differences introduced by the image drift. Such information will help to make a decision if the old servers need to be retired or upgraded. This will reduce operational risk for the entire set of servers.



Closing the Private Cloud Management Loop

The requirements for rapid elasticity and on-demand self-service are addressed (to various degrees of success) by existing private cloud platforms, and their implementations, . However they frequently ignore the challenges private cloud introduces for change and configuration management. As a result they do not provide new capabilities required to support environment agility without losing control. This is the gap that leaves the cloud management loop wide open.

To close this loop, it is necessary to automate change and configuration monitoring, implement analytics on top of the collected information and integrate the new process into the cloud management, starting from the private cloud setup phase, through migration of existing systems and to ongoing operations.

We intentionally sidestepped the whole subject of organizational changes that IT needs to adopt to run private cloud efficiently. Transformation of an organization is a complex, and time consuming effort. This subject deserves a separate white paper. At the same time we found that the new approach to change and configuration management both compensates for deficiencies of existing organizational resources in the context of private cloud and helps to drive the change exposing data essential for transformation.



About Evolven

Evolven redefines configuration management and change management with its groundbreaking Change & Configuration Monitoring solution. Evolven's software-as-a-service solution enables companies to dramatically increase the stability of their IT environments, reduce the risk of production outages, lower operating costs, and cut environment incident investigation time and effort.

Evolven's Change and Configuration Monitoring eliminates production outages. Really.

Start a free trial

www.evolven.com/sign-up.html

