

De-risking the Private Cloud Information Brief



Private Cloud Performance and Availability Management

What is unique about a private or internal cloud infrastructure that increases risk?

Compared with a 'traditional' data center, cloud environments frequently have more users, more applications, more data, and typically, a tighter set of operating standards to meet. But thanks to virtualized servers, storage and networking, they're usually running this increased level of activity on less hardware, which leads to lower overall cost. Unfortunately, one of the attributes of private clouds is the streamlined provisioning that can lead to unpredictable workloads, which negatively impacts performance.

The abstraction layers make it more difficult to address performance problems and troubleshoot root causes to find issues that can bring applications down. Additionally, the density increases the impact of a failure, as more users and applications are affected by these performance problems and downtime.

Typically, when an issue is reported, the user doesn't know where the problem is, it could be in the network, in the storage infrastructure or in the server hosting the virtual machine that's running their application. In the less virtualized physical environment, typically connected to FC SANs, the system administrator may know the cause of the problem or at least where to look. Thanks to the abstraction of the cloud environment, they are unsure where to begin.

What is the effect of this increased risk?

Like the life cycle of applications being deployed on virtual servers, the first applications being deployed in the private cloud tend to be tier 2 or less business-critical. Even when tier 1 applications are deployed, IT operations typically over-provision hardware and staff resources, in an effort to ensure adequate performance and availability.

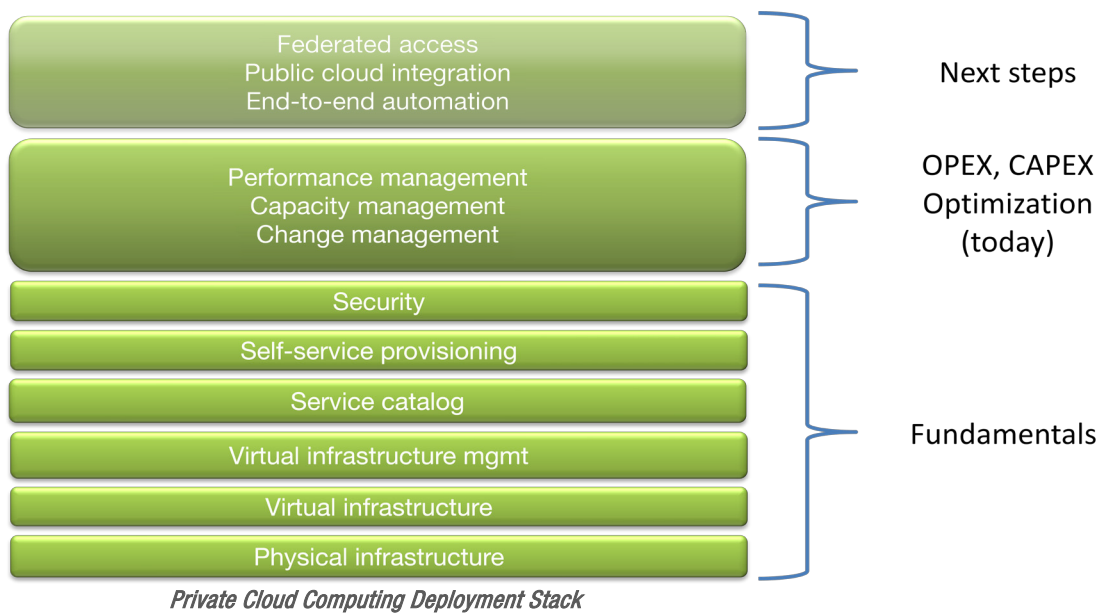
Summary:

- Compared with a 'traditional' data center, cloud environments frequently have more users, more applications, more data, and typically, a tighter set of operating standards to meet
- The most successful cloud implementations include a layer of management to handle change, capacity, and performance management
- Because of the dynamic workload nature of private clouds, performance management can have the greatest effect on maintaining service level agreements
- Most monitoring solutions don't provide a comprehensive, real-time, cross domain view of the environment, leading to ineffective troubleshooting
- Virtual Instruments provides the ideal performance and availability optimization solution for mission critical applications

What are key attributes of a successful private cloud deployment?

In the diagram below, we outline the fundamental attributes of a private cloud in the ‘fundamentals’ tier. Nearly all deployments have some level of server virtualization, though technically it’s not a requirement. And most do some level of self-service provisioning, though the actual physical steps may include a manual process involving traditional provisioning steps. The most successful cloud implementations add a layer of management to handle change, capacity, and performance management, the ‘OPEX CAPEX’ tier. With those layers optimized, companies can mitigate the tendency to over-provision, and see real improvements in OPEX and CAPEX. Because of the dynamic workload nature of private clouds, performance management can have the greatest effect on maintaining service level SLAs. Capacity and change management are important, but less so because the rate of change of the infrastructure itself is far less than the dynamic nature of users’ interaction with applications.

Going forward, the ‘next steps’ are to more fully automate the management of the infrastructure, to seamlessly integrate with public clouds, and to provide federated access to users. But successful private cloud infrastructures are being built today with those steps still out in the future.



What are some of the typical solutions to this challenge?

There are tools available from most of the device manufacturers, like SRM software from the array vendor, fabric monitoring programs for the switches and HBAs, and virtual server monitors from literally dozens of vendors. But these ‘element’ tools require that you know the general location of the problem, which is often not the case. These tools also don’t provide a comprehensive, real-time, cross domain view of the environment, and troubleshooting with them can be ineffective, like hearing only half of a phone conversation. The result can be extended periods of trial and error, hunting for the root cause, or creating routines to track specific links and waiting for the problem to reoccur. In a business-critical cloud environment, where more data is at risk and more users can be affected, this may not be good enough.

What are the key attributes of performance management for private clouds?

- Ideally, the solution should be agentless and easy to deploy, completely non-intrusive to the existing infrastructure components, and it should present a simple GUI for quick analysis and use by general administrators.
- If the solution is vendor independent, it’s much easier to support a heterogeneous cloud, you avoid vendor lock-in, and you reduce the number of tools you need to deploy and support.

- It should support a cross-domain scope, meaning that it should monitor as much of the data path as possible - from the server to the switches to the storage arrays.
- A true real-time solution with millisecond granularity will offer a troubleshooting approach that ensures even hard to find, intermittent problems are quickly discovered and remediated.
- Proactive alerts allow for quick problem identification, and predictive metrics help avoid problems from affecting applications.
- The best solutions scale to the world's largest enterprises to tens of petabytes, tens of thousands of servers and thousands of switch ports, ensuring that no matter how large the application gets, you won't outgrow your performance optimization platform.
- Lastly and perhaps most importantly, you must monitor and analyze performance from the point of view of the application. It does no good to know that you're pushing thousands of IOPS through a pipe if you can't determine how that flow is affecting application response time latency.


Conclusion

Private cloud computing, while offering advantages in reduced IT OPEX and CAPEX and in flexibility and speed of application deployments, brings additional risks to performance and availability.

Virtual Instruments can provide the ideal performance and availability optimization solution for your virtualized mission-critical applications, whether they run on purely physical, or on the most advanced private cloud infrastructures. We call the solution 'Virtual Infrastructure Optimization'. Virtual Infrastructure Optimization holistically assesses the entire private cloud infrastructure, and provides the administrator with the real-time and historical data necessary to make intelligent decisions about I/O capacity, utilization, and performance for every layer of the infrastructure - network, server, storage, and applications.

Other solutions in this area peer into single dimensions of the infrastructure, and while these solutions have a place in routine planning, they do not address the critical dynamic cloud computing infrastructure management needs: real-time *data about what is happening in the moment*. Virtual Instruments' VirtualWisdom provides the real-time granular monitoring and analysis necessary to mitigate the risk of deploying tier 1 applications on private clouds.

For more information on how VirtualWisdom supports private cloud computing, please go to: <http://www.virtualinstruments.com/cloud>

	Corporate Headquarters 25 Metro Drive Suite 400 San Jose, CA 95110 Phone: 408-579-4000 Fax: 408-579-4001	Sales sales@virtualinstruments.com Phone: 408-579-4081	Support support@virtualinstruments.com
--	---	---	--

©2011 Virtual Instruments. All rights reserved. Features and specifications are subject to change without notice. VirtualWisdom, Virtual Instruments, SANInsight are trademarks or registered trademarks in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or servicemarks of, and are used to identify, products or services of their respective owners. 09/11