

Realtime
publishers

Deploying and Managing Private Clouds
The Essentials Series

Steps to Migrating to a Private Cloud

sponsored by



Dan Sullivan

Introduction to Realtime Publishers

by **Don Jones, Series Editor**

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We’ve made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book’s production expenses for the benefit of our readers.

Although we’ve always offered our publications to you for free, don’t think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We’re an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I’m proud that we’ve produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you’ve received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you’re sure to find something that’s of interest to you—and it won’t cost you a thing. We hope you’ll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers..... i

Steps to Migrating to a Private Cloud 1

 Advantages of a Private Cloud 1

 Improved Hardware Optimization..... 2

 Reduced Support Costs with Self-Service Management 3

 Reduced Capital Expenditures 3

 Reduced Time to Deploy Applications and Services 3

When a Private Cloud Is Not the Right Option..... 3

Assessing the Current State of Readiness for a Private Cloud 4

Incrementally Moving to a Private Cloud..... 5

Summary 5

Copyright Statement

© 2011 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

Steps to Migrating to a Private Cloud

Advances in server hardware, network infrastructure, and virtualization allow for a range of IT architectures, including the use of computing clouds. Clouds are collections of computing, storage, and network services that can be rapidly allocated for particular jobs and then easily released to be employed for other tasks when the jobs are completed. The adoption and continued use of public cloud services demonstrate the effectiveness and efficiency of this computing model. Businesses and other organizations with underutilized server capacity may be well served by redeploying their hardware in a private cloud configuration.

This series provides an overview of private clouds, their advantages and disadvantages, and a suggested plan for migrating to a private cloud architecture. This article will begin the discussion with an assessment process and talk about:

- Advantages of a private cloud
- When a private cloud is not the right option
- Current state of readiness for a private cloud
- Incremental movement to a private cloud

An obvious question you must first answer in this process is, Why use a private cloud?

Advantages of a Private Cloud

The advantages of a private cloud stem from four common characteristics found in cloud environments:

- Improved hardware optimization
- Reduced support costs with self-service management
- Reduced capital expenditures
- Reduced time to deploy applications and services

Each of these characteristics is a product of the cloud architecture.

Improved Hardware Optimization

Traditional methods of deploying a single server to run a single application often result in underutilization of server capacity. Servers designed to maintain acceptable throughput during periods of peak demand waste computing capacity during non-peak periods. Running multiple virtual servers on a single hardware server can help to improve utilization, but because these configurations are relatively static, there may continue to be periods of low utilization.

Consider an example. A finance reporting server experiences peak demands in the middle of the night and early morning as daily reports are generated. A human resources application server experiences the greatest demand during normal business hours, especially the late morning and early afternoon. These two applications could run on a single server using virtualized hosts, improving the overall utilization rate. There are still, however, periods when the CPUs and other server resources are significantly underutilized.

As we see in Figure 1, even with virtualization, there are times when another virtual machine could make use of the physical server. A drawback of virtualization is that it requires significant effort to install, configure, and monitor each new virtual machine. Although we have available CPU capacity, the time and effort required to install another virtual machine outweigh the benefits of more efficient server utilization. Cloud computing software eliminates this type of administrative overhead and allows for rapid deployment of virtual machines.

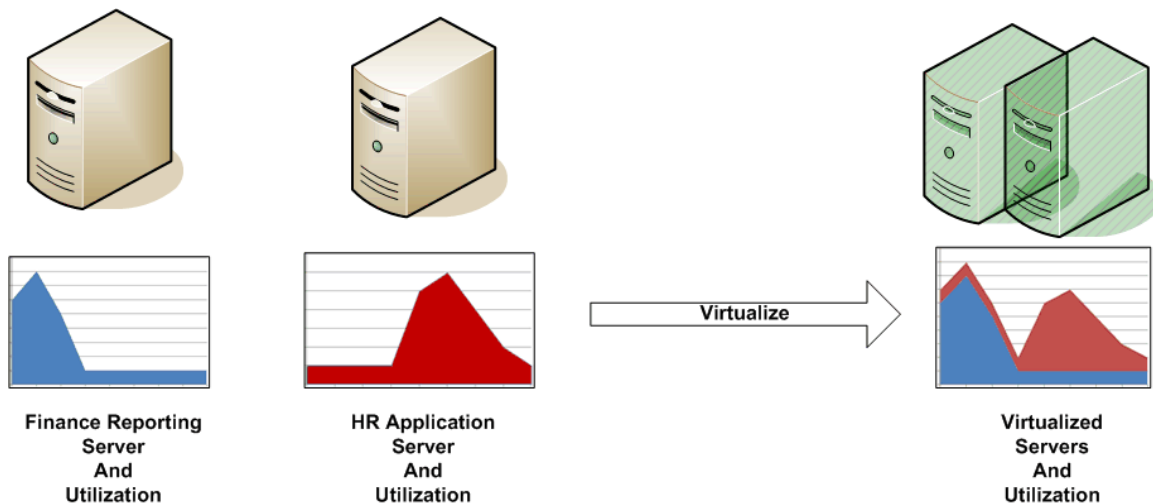


Figure 1: Virtualization can improve utilization but there may continue to be periods of underutilization.

Reduced Support Costs with Self-Service Management

With a private cloud, IT departments can establish catalogs of virtual machine images, sometimes referred to as a *service catalog*. End users can select an image from the service catalog and have it run on one or more virtual machines in the private cloud. A systems administrator is no longer needed to start a new virtual machine. End users decide when to start and stop their virtual machines. As a result, hardware can be more efficiently utilized while at the same time reducing support costs through self-service management.

Reduced Capital Expenditures

More efficient use of server hardware can quickly lead to reduced capital expenditures. In the past, a new business service may have required additional hardware to support that service. Using a private cloud, an IT department can pool its server resources and allocate virtual servers as needed. New business services can tap into the same pool of servers without having to purchase hardware based on peak demand expectations. Cloud administrators monitor usage and plan for expansion based on the aggregate needs of all users—not on the peak demand periods of each individual application.

Reduced Time to Deploy Applications and Services

Eliminating or reducing the need to tightly couple hardware to applications can also reduce the time required to deploy new applications and services. Application owners and system architects can design and deploy new services to the cloud without having to wait to procure, install, and configure new hardware.

The advantages of cloud computing are compelling reasons to shift to this model, but is it always the right choice?

When a Private Cloud Is Not the Right Option

Ideal solutions in some circumstances are not ideal in others. Cloud computing is no exception. In general, cloud computing can be an efficient and cost-effective way of deploying IT resources when the following conditions are met:

- The cloud provider has standardized hardware
- The cloud provider has standardized operating system (OS) and application stacks
- Sufficient networking infrastructure is in place to support cloud computing
- Management tools, such as server and network monitoring, are in place
- The distribution of jobs lends themselves to distributed computing platforms

A private cloud may not be the best option if the hardware you plan to deploy is not standardized. A common hardware foundation will make deployment and management much easier. If different platforms require different drivers and separate versions of machine images, the cloud will require additional management and undermine some of the cost advantages of a cloud.

If jobs are not easily divided among virtual machines, a private cloud may not be appropriate. For example, a business intelligence application that generates a large number of reports can be split over several virtual machines with each instance generating reports for different departments. In contrast, an optimization problem that requires a large amount of shared memory on a single machine or a graphics rendering application that requires specialized graphics hardware are best run on dedicated hardware.

Another requirement that is easily overlooked is having proper management tools in place. Clouds require monitoring to ensure services are running as expected in the short term and to provide data for longer-term capacity planning. If appropriate monitoring tools are not in place already, they should be included in the plans for deploying a private cloud. This is just one of the requirements that you should assess before migrating to a private cloud.

Assessing the Current State of Readiness for a Private Cloud

Planning is the process of identifying how to get from where you are to where you want to be. The first step in the case of private cloud deployments is understanding where you are in terms of existing IT infrastructure, including an assessment of:

- Existing infrastructure
- Resource utilization levels
- Policies and procedures governing IT operations
- Reporting and cost allocation procedures

Each of these represents important elements for successfully delivering private cloud services.

Existing infrastructure includes servers, storage systems, and networking infrastructure. Ideally, private clouds make use of similar servers. They have the same CPU cores, amounts of memory, types of power supplies, and network interface components. As you inventory your hardware, determine how many of your existing servers can be used in a cloud. Servers that are at or near the end of their useful life should not be included in a cloud. Older hardware may still be useful for less demanding applications, such as print servers.

Similarly, plan to deploy storage arrays based on similar technologies. The more we use standardized hardware, the more we can reduce the management overhead of the cloud.

Assess network capacity. Are there segments of the network with high latency? Is existing bandwidth sufficient for current needs? It is important to understand if any parts of the network infrastructure will need to be upgraded prior to moving to a private cloud. It is possible that a shift to a cloud model will not significantly alter the demand for network services, but monitoring is essential to understand the requirements for your particular environment. For example, it may be difficult to estimate the growth in demand for network services when deploying a private cloud. With reduced costs to deploy new services and greater self-management, business units may begin to experiment with new applications and increase the number of jobs they run.

Formal policies and procedures should be in place prior to migrating to a private cloud. Many aspects of these policies will be embodied in systems that implement the cloud. User authentication policies, for example, must be in place in order to control access to cloud resources. Similarly, billing and cost recovery policies must be defined so that users can make choices about the types and amounts of cloud services to use. Procedures should be in place to automate as much as possible the accounting and billing aspects of delivering cloud services.

Incrementally Moving to a Private Cloud

Deploying a private cloud does not have to occur in a short period of time; it may be advantageous to take an incremental approach to implementing a cloud. This can be done following a basic three-part process:

- Identify hardware and applications that can be moved to a private cloud. Again, not all jobs are suitable for a cloud environment, but many business applications are.
- Determine the utilization levels needed to support an initial set of cloud applications. Use monitoring logs on existing servers to understand peak demand levels, the duration of peak demand periods, as well as average CPU, memory, storage, and network utilization over an extended period of time.
- Establish a monitoring program to capture data about how the cloud infrastructure is used.

Once a private cloud is deployed, you will still likely have non-cloud resources operating as part of the broader IT infrastructure. It is important to continue to monitor these resources as well. As business users migrate their applications to the cloud, additional resources may be needed in the private cloud. Some of this need may be met by servers that operate outside the cloud. A comprehensive monitoring program will help to optimize the allocation of resources between cloud and non-cloud deployments.

Summary

Private cloud computing environments offer a number of advantages for optimizing the use of servers, storage, and network services. Although cloud architectures are not appropriate for all requirements, many business applications are well served by cloud computing. Organizations planning a move to a private cloud should conduct a readiness assessment prior to deploying a private cloud. As part of the implementation, ensure policies and procedures are in place to effectively manage the cloud and have monitoring tools in place to collect the data that will be needed to ensure both the short-term availability and long-term efficiency of the private cloud.