



Why You Need a Cloud to Call your Own

Private cloud computing may not offer the same degree of cost savings as public cloud computing, but greater flexibility and time savings will outweigh less significant cost savings in the long term. A joint F5-IBM cloud computing solution based on proven technology and standards can help you get to the private cloud faster.

White Paper
by Lori MacVittie



Introduction

Cloud computing has been top of mind (and top of every technology trend list) for quite some time now. Early in the lifecycle of cloud computing its allure was "cheap" computing through shared resources. However, as organizations continued to raise concerns regarding security and control in public environments, the focus invariably turned to implementing cloud-computing architectures internally, as private cloud-computing environments.

While proponents of public-only cloud computing continue to portray private cloud computing as little more than virtualization, organizations moving forward with a private cloud-computing initiative will find there are tangible business benefits beyond merely "cheap" computing. These benefits not only improve the agility of IT but also better align IT with business concerns by making IT more efficient, responsive, and cost-effective. In fact, an April 2010 study cited by InformationWeek shows that organizations are seeking out cloud more for its flexibility than its financial incentives: "Forty-nine percent of companies are driven to the cloud for business agility reasons, while only 46 percent are motivated by cost efficiency," the report states. (IT Spending on Cloud Ratcheting Up, April 2010, InformationWeek) Ultimately, it is operational efficiency that enables the business flexibility and cost savings for which private cloud initiatives are begun and against which such implementations are measured for success or failure. Efficiency is at the heart of cloud computing, and it is the cornerstone upon which the entire model is based.

Adoption of private cloud services for IT workload processing or infrastructure is outpacing the use of public platform service providers.

"Privatizing the Cloud: Study on Enterprise Cloud Adoption", 2010 IOUG Survey Results

Achieving Operational Efficiency

Operational efficiency is not solely about resource utilization. There is certainly a financial incentive to squeeze out as much value as possible from hardware resources through consolidation and virtualization. However, the operational efficiencies gained from private cloud computing are much broader than simply increasing resource utilization across IT. Operational efficiency benefits of private cloud computing can be categorized as follows:

- **Financial**

Improved utilization and sharing of otherwise idle resources decreases the operational costs associated with the entire lifecycle of an application.



WHITE PAPER

Why You Need a Cloud to Call your Own

- **Time**

While often tightly coupled with financial considerations, the ability to rapidly provision resources can dramatically improve time-to-market, which, in turn, enables greater agility for both IT and business stakeholders in adapting to changing business conditions.

- **Process**

Many operational efficiencies are "built in" to the processes required by IT and business stakeholders to deliver applications. The automation and orchestration of such processes affords organizations opportunities to evaluate those processes and eliminate unnecessary, redundant, or otherwise disruptive steps. Automating repetitive tasks can mitigate the potential introduction of errors into those processes and reduce the time required to execute. And, in some cases, automation creates self-service opportunities that enhance the business stakeholder's desire for higher levels of control and visibility into IT processes.

Operational efficiency requires that organizations go beyond virtualization and begin to automate the data center, starting with processes that have a high impact on the business but are also highly automatable. These processes are ones that have well-defined steps that must be taken and which can be accomplished by leveraging application programming interfaces (APIs), remote scripting, and other automation-focused technologies across the entire infrastructure.

The implementation of automation and adoption of orchestration leads to greater operational efficiency—including the higher utilization rates that result in financial benefits.

Financial Benefits

Initially, "cheap" resources was the primary benefit associated with public cloud computing. This is not hype: Sharing resources and the infrastructure needed to leverage those resources does, in fact, reduce costs to a level with which enterprise IT organizations cannot hope to compete. This is simply a matter of mathematics and the scale on which public cloud-computing providers can share resources. There are, after all, only so many customers (departments, lines of business) within an organization that can "share" resources and, therefore, the costs of the entire infrastructure.

This does not, however, mean that there are no financial benefits to a private cloud-computing implementation. It is simply that the reduction in capital and operational expenses in terms of resources is not as great as it would be in a public cloud-computing environment. A lower price per resource can be achieved in the corporate data center, with private cloud computing, provided resources are shared across projects, departments, and lines of business.



Staff related costs often were 50% to 70% of the total cost over a period of three years. Cost of communications, power, cooling and facilities could add up to another 30% to 40% of the total. Hardware and software, when combined, usually represented somewhere between 20% and 25% of the costs.

—"Cloud Computing Could Cost More Than Using Your Own Systems," April 2010, ZDnet



WHITE PAPER

Why You Need a Cloud to Call your Own

Cost savings through reduced cost of resources is also realizable in environments that see low resource-utilization rates. Instead of acquiring new resources, reallocating resources that are otherwise sitting idle to other projects enables costs to be shared, and thus reduced, on a per-project basis.

Time Is Money

Financial benefits associated with more rapid provisioning and a reduced time-to-market are also feasible within a private cloud-computing implementation. By leveraging rapid provisioning and sharing of resources, IT organizations can respond more quickly to requests for resources—thus avoiding the lengthy acquisition and subsequent rollout processes required in traditional architectures. Because IT charges for most business projects based on time—which translates into costs—a reduction in time spent provisioning resources directly translates into financial savings for the project, as well as a shortened timeline to deploy.

For example, if provisioning a server takes 40 hours in a traditional environment but only 4 hours in a private cloud implementation, those "extra" 36 hours can be returned to the overall pool of hours available to business constituents' projects. In a traditional model, the only way to increase the pool of available hours is to hire more staff, which likely increases the cost of IT hours, or to somehow reduce the administrative overhead. Reducing administrative overhead can be helped by private cloud-computing models as it should, according to the experts, reduce the provisioning and maintenance costs associated with the business of IT. For example, "While a single system administrator can service approximately 140 servers in a traditional enterprise, in a cloud data center the same administrator can service thousands of servers."¹ Such a dramatic increase in the administrator-server ratio can result in a significant reduction in per-server administrative costs.

Process Automation

Leveraging a private cloud also leads to repeatable deployments through automation, which can further decrease time-to-market and increase project completion and success rates. Repeatable deployments can be leveraged as "templates" to further provide a firm foundation upon which application deployment becomes a "service" that is more easily used by operators and business constituents alike. Faster time-to-market at every level increases competitive advantages and makes the business more responsive to its customers—an often incalculable benefit.



WHITE PAPER

Why You Need a Cloud to Call your Own

Automating deployments leads to greater use of resources primarily because automation requires integration and interaction across the entire delivery infrastructure. This level of collaboration enables finer-grained control by orchestration systems, which results in better use of resources and often in a more timely fashion. For example, automating scalability on demand for web applications requires proactive monitoring of usage, user performance, and capacity. Orchestration systems incorporate this data into decision-making processes that ultimately lead to faster capacity provisioning, which helps eliminate the degradation or disruption of service often associated with a manual provisioning process. Higher levels of responsiveness help businesses better meet service-level agreements for key performance indicators (KPIs) such as more efficient call-center utilization or reduced customer abandonment rate.

A successful private-cloud implementation strategy will include an evaluation of provisioning, scalability, and even routine maintenance processes in order to streamline them and establish greater efficiency. This is especially relevant in well-established data centers in which processes have been built upon for years and may have become bloated with unnecessary steps or approval chains. In evaluating processes, it is also important to consider self-service, particularly if a process is tightly enmeshed with business needs—for example, resource provisioning and application access requests.

By streamlining and automating processes, organizations can gain more than just the cheaper resources and rapid provisioning associated with public cloud computing. Organizations that are serious about taking the opportunity to re-engineer processes through private cloud-computing initiatives can, in fact, achieve greater efficiencies by slimming down processes and creating a leaner, greener, and more responsive IT organization.

This is not, however, the case with public cloud computing, which requires that new processes—specified by the cloud-computing provider—must be adopted and integrated into existing IT processes. This conflation of externally driven processes with internally driven processes can confuse responsibility and accountability between IT and the provider, and it can leave the business stakeholder without the means to quickly resolve potential availability and performance issues.



A Joint F5-IBM Solution

One of the fears often associated with choosing a private cloud deployment over a public cloud environment is initial costs. Organizations are rightfully concerned about the new solutions and infrastructure required to support a private cloud initiative, as well as the potential need for new or retrained employees who have the skills necessary to successfully implement even the most basic of cloud-computing frameworks.

F5 Networks and IBM have joined forces to provide a solution that mitigates these risks by leveraging well-understood technologies to create the closed-loop, dynamic-provisioning environment that is the foundation of cloud-computing environments. By taking advantage of proven and often existing enterprise-class components, a joint F5-IBM solution addresses concerns regarding skill sets and the oft-forgotten need to integrate with existing management frameworks and technologies. It is this necessary relationship between management of shared resources and network-based application delivery functions that is critical to properly scaling a large private-cloud implementation.

F5 and IBM provide a standards-based integration between IBM Service Management offerings and F5 BIG-IP Application Delivery Controllers. This integration leverages universally compatible protocols such as Simple Network Management Protocol (SNMP) and Domain Name System (DNS) to achieve a dynamic, closed-loop provisioning environment for applications that are deployed in a private cloud-computing architecture, which is, itself, highly scalable, such that individual components can be deployed as "cloud" components themselves.

Thresholds for provisioning based on variables such as CPU utilization or application performance can be set based on organizational policies or preferences. The F5-IBM solution is the integration and the framework for acting upon those variables to enable a dynamic provisioning environment. The integration removes the need for organizations to implement such a framework, and it provides an optimized process through which applications can be efficiently scaled out using the automated scaling techniques associated with cloud computing.



In a traditional data center, a network administrator maps the addition of a new server to the network, assigning it switch and router resources; then a security and compliance administrator checks the configurations and installs any additional protections needed for the new server. With an internal cloud, those three tasks can be collapsed into one—the creation of a VM [virtual machine] that's met with the approval of all three. IT departments need to put work into the process of constructing VMs so that can be accomplished in an automated fashion without disrupting IT operations or creating security risks or data privacy breaches.

—Source: "Why 'Private Cloud' Computing Is Real – And Worth Considering," April 2010, InformationWeek



WHITE PAPER

Why You Need a Cloud to Call your Own

Implementing a private cloud-computing environment based on IBM Service Management offerings and F5 products such as F5 BIG-IP Local Traffic Manager (LTM) Application Delivery Controller, gives customers the ability to lay the foundation for a highly scalable and resilient set of cloud offerings upon which more complex services can be delivered. Simultaneously, this F5-IBM-based environment helps eliminate the business impact of needing new skills to design, deploy, and manage a private cloud-computing environment. Additionally, an integrated F5-IBM solution leverages what are likely existing components— hardware and software—to reduce the cost impact of the potential infrastructure acquisition required to embark on a private cloud-computing initiative.

Conclusion

The benefits ascribed to cloud computing can be realized in both public and private cloud-computing implementations. While in some respects there are greater benefits in public cloud-computing environments—higher cost savings resulting from better economies of scale, for example—there are also benefits such as process optimization and elimination of redundancies that can only be fully realized with a private cloud-computing environment.

A private cloud-computing implementation gives organizations the opportunity to realize benefits without compromising on the architectural control required to assure integrity of data, systems, and processes. A private cloud-computing implementation based on a joint F5-IBM solution preserves that architectural control while eliminating the need to architect. And, the joint solution enables organizations to implement a cloud-computing foundation that is sustainable, scalable, and built on existing, proven technology.

A joint F5-IBM solution can jump-start a private cloud-computing strategy by providing an integrated, flexible foundation upon which the organization can work toward the goal of providing IT as a service.

¹ The Economics of the Cloud, Microsoft, November 2010.

|| <http://www.microsoft.com/presspass/presskits/cloud/docs/The-Economics-of-the-Cloud.pdf>

WHITE PAPER

Why You Need a Cloud to Call your Own



F5 Networks, Inc.
401 Elliott Avenue West, Seattle, WA 98119
888-882-4447 www.f5.com

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

Europe/Middle-East/Africa
emeainfo@f5.com

Japan
f5j-info@f5.com

©2015 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5. CS01-00039 0113