



White Paper October 2012

By Jared Wray, Tier 3, Founder and CEO

Robert Miggins, PEER 1 Hosting, SVP of Business Development

www.tier3.com / www.peer1.com

Time to Move on From the Commodity Cloud

Identifying What Makes a Cloud Truly
Enterprise-Grade

Table of Contents

Executive Summary	1
Introduction: Doing More with Less	1
The Five Core Elements of the Enterprise-Grade Cloud	3
Element #1: A Robust Baseline of Infrastructure, Services, and Support	3
Element #2: Increased Automation	4
Element #3: Management	5
Element #4: Application Services	7
Element #5: Making it All Easier	8
A Timeline of Enterprise-Grade Cloud Attributes	9
Conclusion: Now Is the Time to Get on the Right Enterprise Cloud	10
About the Authors	11
Footnotes	12

Executive Summary

Today's enterprise IT leaders are looking to shift a large number of applications to the cloud to make their businesses more agile and efficient. In fact, a recent Forrester Research survey on IT budgets and priorities found that 35 percent of enterprises are focusing heavily on cloud computing, calling it a high or critical priority.¹

Despite these intentions to move applications to the cloud, CIOs will only do so with assurances and proof that their cloud service provider can deliver **truly** enterprise-grade capabilities in a manner that is not overly complex. Put another way, they expect similar levels of performance and manageability as they've come to expect from their own internal datacenter operations.

The next wave in cloud computing is happening now and is about more than just rapid deployment or spooling up more virtual machines. Enterprises have been asking for cloud services that deliver agile deployment as well as efficient, sustainable management of workloads across their internal and external computing environments. The next generation of cloud is about converged services, improved management, and higher quality infrastructure. This is where separation between cloud providers starts to emerge.

"Enterprise grade" has not been a term associated with multi-purpose, commodity public cloud, and for good reason. Most cloud service providers never architected their cloud to reach such exacting standards. Instead, they have focused on rapidly provisioning virtual machines as inexpensively as possible. As a result, enterprise IT leaders should tread carefully before committing their applications and workloads to a cloud service provider. The time and resources invested in migrating to an inadequate cloud - one that lacks enterprise-grade capabilities, services, and support - could have dire consequences for the business.

This paper describes the five core elements of an enterprise-grade cloud that IT leaders should look for **today** as they evaluate potential cloud service providers. It also explores future enterprise cloud capabilities that will evolve from these five core elements.

Introduction: Doing More With Less

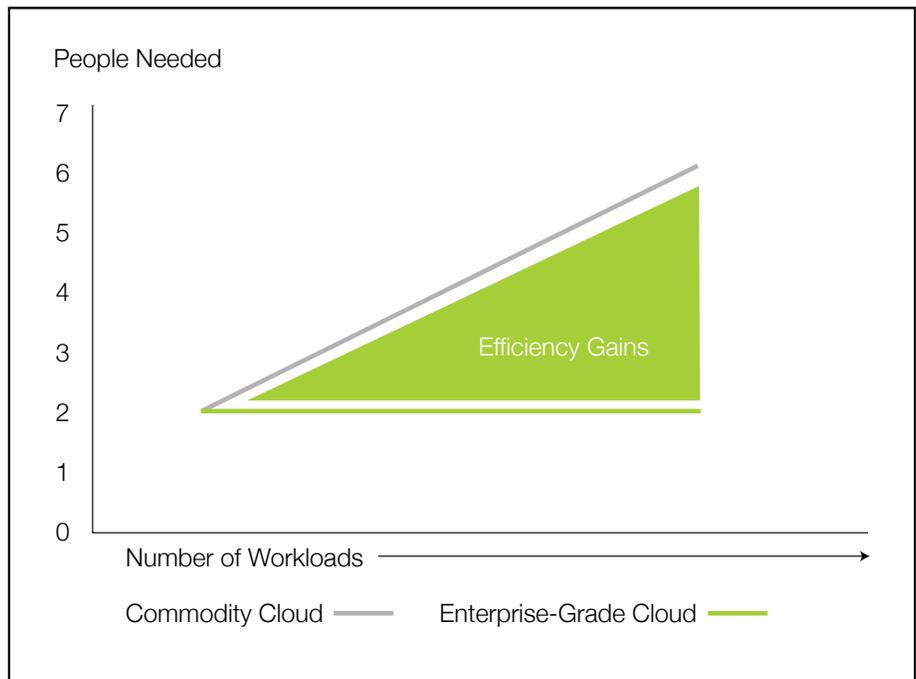
The first wave of cloud adoption touted the benefits of rapid deployment, and most clouds have delivered on that promise. But enterprise-grade cloud calls for something much more difficult to achieve - the combination of agility and management efficiency at scale. For enterprises to succeed, they need efficiency gains enabled by a cloud that can provide service automation and management control - even at scale.

This cloud maturation will help enterprise IT better align with the needs of the business. As the cloud infrastructure handles more common management tasks, strategic IT resources can focus on their core competencies and better respond to business demands.

Improved efficiency and control in the cloud will help enterprises overcome the currently high operational costs associated with running their own datacenters, which are exacerbated by a lack of parity between their internal and external cloud environments. Mass-market commodity clouds simply cannot perform such critical tasks as auto-provisioning, orchestration, and management. Without them, CIOs have been hesitant to shift the majority of workloads to the cloud. Thus, businesses seeking to reap more cloud benefits across the organisation are now imposing new requirements on cloud vendors.

Consider the issue of cloud efficiency in terms of IT's highest cost – headcount. Unlike commodity cloud services, enterprise-ready clouds do not need to add higher levels of human support as the breadth and depth of workloads in the cloud continue to grow. Figure 1 below demonstrates the resource advantage enterprises stand to gain by using enterprise-grade cloud over commodity cloud.

Figure 1. Resource growth without workload management growth



IT has the opportunity to deliver higher value to the business – but only if cloud service providers are prepared. Enterprises that want to take advantage of the next generation of cloud need to partner with a cloud service provider that understands, and has the capabilities to deliver, a more sophisticated set of services for enterprise use. When reviewing a cloud service provider's capabilities, IT leaders should constantly ask themselves two core questions:

- 1. Does the service provider offer a cloud with simplified management and higher levels of automation to free up my personnel from mundane infrastructure tasks?** (i.e. capabilities for simultaneously securing and monitoring all machines at once, building and managing complex large-scale environments, performing multi-machine management and provisioning across large-scale deployments)

2. Can my application developers and IT administrators consume cloud services in new ways to enhance agility? (e.g., leveraging converged services, such as Infrastructure-as-a-Service, Platform-as-a-Service, Database-as-a-Service, etc.)

A “no” answer to either question could mean getting locked into a cloud that does not meet long-term needs. Yet IT leaders may find it difficult to obtain definitive answers to these questions. As such, CIOs and other IT-focused leaders need to recognize five elements of enterprise-grade cloud that do not exist in commodity cloud.

The Five Core Elements of the Enterprise-Grade Cloud

Most enterprises will leverage a mix of internal and external cloud in their environment for years to come. This results from their distinct needs to have complete control over certain mission-critical applications while seeking agility, efficiency, and scale for all their applications. To realize efficiencies, enterprises need their external cloud provider to make easier the interplay between the cloud and internal datacenters – all while continuing to deliver uptime and cost savings. Five core elements of enterprise-grade cloud make this possible.

Element #1: A Robust Baseline of Infrastructure, Services, and Support

Following the adage that you must walk before you can run, a cloud service provider that aspires to deliver enterprise-grade services must first have established global infrastructure, services, and support that correspond to the exacting demands of enterprise IT.

Infrastructure

Before delving into a cloud provider’s stack, IT leaders need to understand the provider’s global infrastructure capabilities. The commodity cloud operating model of pushing traffic over uncontrolled public networks and running workloads in shared server farms will not meet the service level agreements (SLAs) that enterprises require. IT leaders must ensure they partner with a hosting provider that offers robust infrastructure and quality support. In particular, the hosting provider should stand firmly behind SLAs for network, compute, and storage uptime, with SLAs actually calling for payment to the customer in case of failure.

The cloud needs to provide consistent network performance, enterprise-grade hardware, and a globally redundant infrastructure. This will enable developers to architect applications for optimal conditions — a highly available environment with deeply integrated server management and sophisticated deployment options.

Services and Support

Unlike many commodity cloud customers, enterprises need high-touch support. As such, they need to look beyond clouds that provide 100 percent of customer management, services, and support via an online panel, with minimal or no expert personnel available for live support.

An enterprise-grade cloud must be backed by frontline experts who are experienced enough to resolve performance and availability matters as quickly as possible. Poor workload performance or outages can cost an enterprise millions of dollars and damage a company’s reputation. In fact, a recent study by Aberdeen places the

average cost of enterprise application downtime at over \$160,000 per hour.² A cloud provider's datacenter support personnel must demonstrate knowledge of core infrastructure support (i.e., backup, disaster recovery, network, security, etc.) as well as application/workload performance optimization to avoid and resolve potential issues.

Enterprise-grade service and support is not just something a cloud provider can decide to begin doing tomorrow. IT leaders should investigate the depth of experience a cloud provider has, including proof points of having supported large-scale, global deployments for an extended period of time. Enterprises do not wish to become the test case of a service provider's early forays into enterprise cloud.

Element #2: Increased Automation

An important driver of cloud adoption continues to be automation of common tasks. Automation for deployment, monitoring, and software installation is critical to maintaining the efficiency gains referenced earlier (see Figure 1). Such automation of the cloud consists of two main functions — orchestrating deployment and streamlining the human activity required for everyday administration.

Orchestrating Deployment

With a growing number of applications and workloads to deploy, enterprises will benefit greatly from repeatable, consistent deployments that occur virtually and automatically. This will save time and resources while also enabling IT to meet expectations across the business when responding to requests for new services.

Orchestration depends upon automated deployment of best-practice-optimized templates for complex environments and applications. A template would include such items as server provisioning, configurations of operating systems and software, and application installation — all wrapped into design “blueprints” for repeatability and consistency. Leveraging the cloud provider's orchestration blueprint, IT can give individual businesses a one-click process for establishing their own application-specific environments, such as a Microsoft® SharePoint® environment.

Advanced orchestration capabilities enable an IT team to extend their expertise throughout the business for agility and scale. Repeatable orchestration between IaaS and Web applications allows a variety of enterprise users to easily deploy and retire workloads without involving IT, thereby reducing operational support costs. No longer will IT need an assembly line of people to perform such tasks as adding machines and packaging configurations. Rather, they can expose each step in the provisioning cycle as an individual component that is preconfigured for fast implementation. Once deployed, the entire ecosystem should have automated monitoring and remediation capabilities, such as healing and auto-scaling when a server reaches pre-set capacity thresholds.

Whereas a next-generation enterprise cloud service integrates all these capabilities, commodity clouds use bolt-on services in an attempt to deliver efficiency. Many of their tools are black-boxed, leaving IT people struggling to deploy new environments while application developers cannot optimize code to tools they cannot see. An enterprise-grade cloud provides better, transparent tools for addressing implementation and performance.

Consider an IT team that is performing basic patch management across a virtual machine group of 300 servers. Doing so involves simple, mundane tasks. In this case, a total of nine such tasks are required, as follows:

1. Notify the operational team about servers going out of rotation
2. Log the event to an internal tracking system
3. Pause monitoring for all servers
4. Pull load-balanced servers out of rotation and drain sessions
5. Apply the patch
6. Turn on monitoring to verify everything works as intended
7. Place load-balanced servers back into rotation
8. Log that everything is not yet completed with maintenance
9. Notify the operations team of servers “back into rotation” status

An automated system would transform this nine-step process so that the IT person must only perform three steps: initiate the job (system then handles steps 1-4); apply the patch; return system to normal mode (system then handles steps 6-9).

By reducing the number of steps from 9 to 3, the enterprise saves approximately 850 manual labor hours per year for a 300-server deployment.

Streamlining and Simplifying Common Tasks

An enterprise-grade cloud can automate common tasks at scale, including: auto-scaling, patch management, and running scripts across servers. Once again, the advantage an enterprise-grade cloud holds over commodity cloud is deeper integration of these utilities within the entire technology stack.

Element #3: Management

Management is a critical feature of next-generation clouds for the enterprise when one considers how rapidly an enterprise cloud can grow through automated deployment of a large number of workloads to the cloud. Two aspects of cloud management set enterprise-grade clouds apart from commodity clouds: Infrastructure-based management and organization-based management:

Infrastructure-based Management

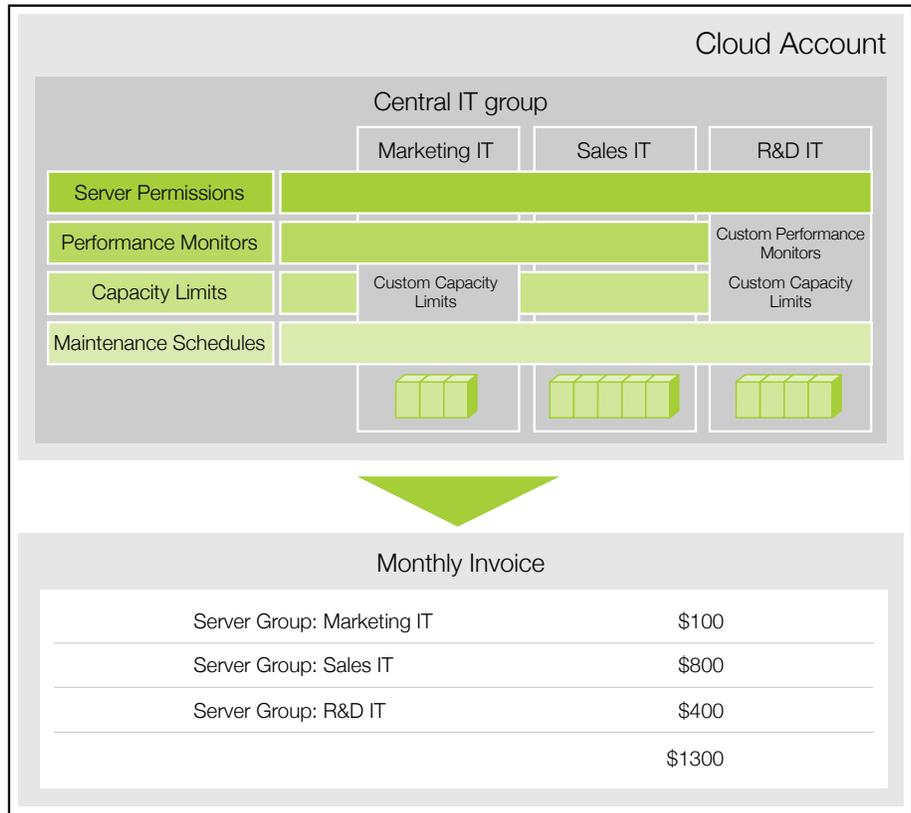
Virtual Machine sprawl, or VM sprawl, has plagued many enterprise virtual environments, and VM sprawl will occur in enterprise clouds as well unless there is robust cloud infrastructure management. To keep an environment in check, an enterprise cloud should provide grouping-type features so that IT can easily invoke policies that help manage an environment of multiple servers and virtual instances. Through one unified administrative interface, organizations can expand resource usage and/or place limits on it per user-defined thresholds. This ensures optimal performance of newly added workloads, while simultaneously preventing unchecked resource usage and overly high billing are added, the cloud should automatically manage provisioning, scaling with patch management, security, and monitoring — all together and at the same time.

Organization-based Management

Typically, each business unit within an enterprise wants its own cloud for control of its own specific workloads. Otherwise, business unit leaders cannot accurately track billing and resource utilization to make people accountable and accurately forecast future needs and budgets. They must have the ability to set policies (e.g. security and compliance, auto-scaling, resource limiting, etc.) and apply them across the business unit's cloud. This puts IT in the position of delivering IT-as-a-Service to each individual business unit, providing management of specific accounts and sub-organizations. This represents a vast departure from the way commodity cloud service providers view their customers and deliver cloud services. Commodity clouds do not concern themselves with internal IT-as-a-Service capabilities within each business unit of an enterprise. Instead, they look at the enterprise in aggregate, never differentiating services between business units.

Figure 2 below depicts the management flexibility of enterprise-grade cloud to vary services by business unit or department. In this case, the enterprise has chosen to receive services specific to major business functions (e.g., Marketing, Sales, and R&D). The marketing organization operates fine with generic settings for server permissions, performance monitors, capacity limits, and maintenance schedules. But as Figure 2 shows, the cloud provider enables the R&D organization to save costs by customising capacity limits and performance.

Figure 2. Enterprise-grade cloud enables internal IT-as-a-Service



The organization-based management approach is particularly important for complex businesses that need to:

- **Break down sub-accounts with features for billing and reporting purposes.** Just as mature virtualized environments in the datacenter can perform department-level chargeback, so too must the enterprise cloud.
- **Manage access to cloud-based workloads from a parent account, mainly for security purposes.** Enterprises often collaborate with third-party vendors or bring in contract personnel to utilize workloads in the cloud. In these circumstances, the cloud must provide for security isolation and account isolation through roles and permissions that govern individual access.
- **Improve planning.** Business units and departments must be able to review data regarding historical usage patterns and their associated charges. This will help them when budgeting time and resources for future workloads in the cloud.

Referring back to Figure 2, all of these IT-as-a-Service capabilities should be integrated into the cloud stack of a next-generation cloud service provider. If they are not, the cloud will not scale to meet the growing management needs of the enterprise's cloud-based initiatives – leaving IT to manage multiple vendors with very different service levels and no guarantees on security.

Element #4: Application Services

IT leaders must seek a high level of integration between a cloud provider's IaaS, PaaS, and its other platform services (I.e. Database-as-a-Service, Messaging-as-a-Service, Caching, etc.). Tight coupling of these services do not exist in commodity cloud. In enterprise-grade cloud, by contrast, such coupling enables customers to manage utility-based services as they assemble new applications and as their existing workloads consume the services.

Combined IaaS and PaaS in the cloud means that enterprise-grade clouds can manage applications on the platform side far more easily. For example, IT can decide to use managed services instead of taking on such tasks as installing and maintaining their own database, running SQL Server, or backing up SQL Server. The cloud platform should be able to spin up a Database-as-a-Service and automatically handle all associated infrastructure and operational tasks, with no requirement for customer intervention. The capabilities within such a Database-as-a-Service offering include the following:

- Conducting basic database tuning
- Scaling resources for compute, network, and storage
- Executing backups
- Providing global disaster recovery
- Ensuring high availability

Most organizations do not have enough database administrators (DBAs) to manage all the mundane tasks that will be required by new cloud-based workloads, making Database-as-a-Service quite useful. By offloading these tasks to the cloud, IT leaders can enable DBAs to focus on more higher-value activities, such as performing advanced database tuning and developing new schema architectures.

The types of services found within enterprise-grade cloud will continue to flourish over time. Those few cloud providers that offer enterprise-grade cloud today have already set their sights on delivering multiple services: Messaging-as-a-Service, Queuing-as-a-Service, and Identity-as-a-Service.

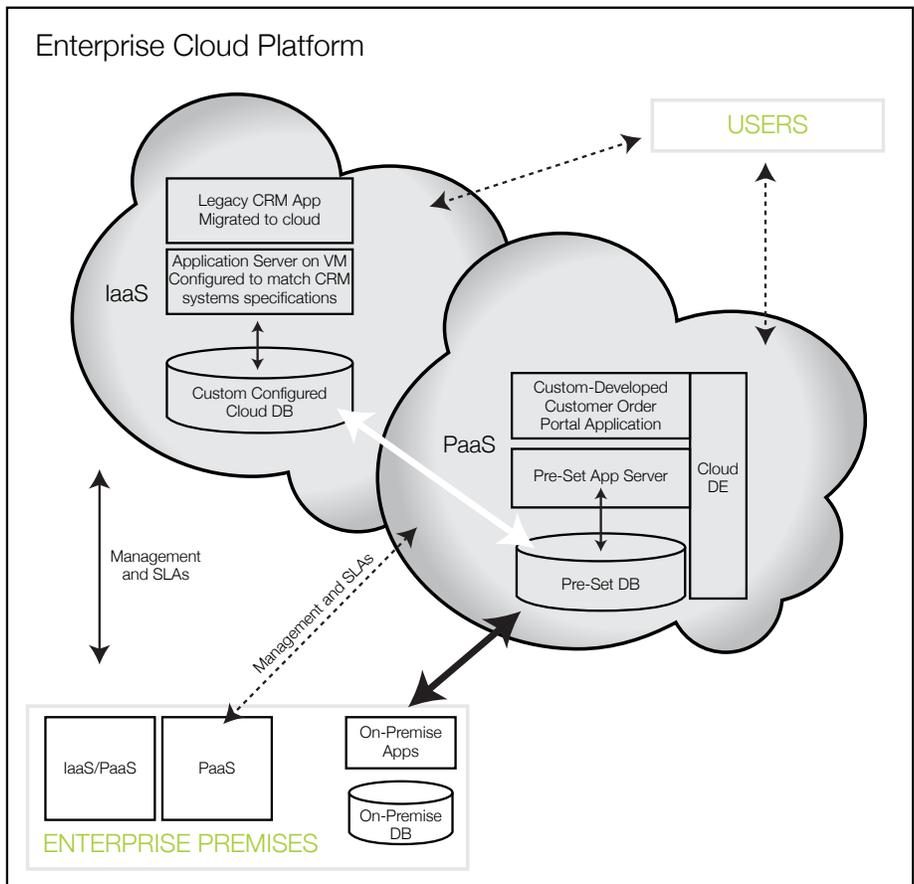
In the future, monitoring and automation will lead to yet another shift in the enterprise cloud. Today, most clouds provide threshold-driven alerts that notify users when troublesome conditions exist, such as a hard drive nearing its storage limit or workloads approaching levels that adversely impact server performance. Eventually, enterprise-grade clouds will include “self-awareness” features, whereby users can specify custom rules to dictate how individual cloud workloads should behave under a variety of conditions. This will bring about further automation and efficiency gains.

Element #5: Making it All Easier

“Easy” in the enterprise cloud can mean many things. Considering the first four elements already discussed, easiness in the enterprise-grade cloud comes about by delivering all elements with as little effort as possible. This requires a cloud stack that integrates everything the enterprise needs to bring together internal and external resources at scale: global infrastructure, automation, management, and services.

“Integrated” does not mean that different elements and tools are loosely integrated to deliver the whole cloud. Such a cloud would never scale with ease. A truly enterprise-grade cloud, as depicted in Figure 3 below, has woven IaaS, PaaS, automation, management, orchestration, and services into a single unified platform that was intended specifically for the enterprise.

Figure 3. Making internal/external cloud easier



Only with such a unified platform can enterprise IT leaders ensure that they have a highly consumable cloud that will not break down as workloads and complexity grow. The enterprise-grade cloud must make everything easier. Otherwise, enterprises will not gain the agility and efficiency that they seek from the cloud.

A Timeline of Enterprise-Grade Cloud Attributes

The push for enterprise-grade cloud is already serving to tier the cloud service provider market. The fact is that only a small number of cloud service providers exist with the infrastructure and unified platform capable of delivering on the enterprise cloud promise. The table below depicts the timing for key attributes to be found in enterprise-grade clouds. Some exist today, and others are coming within the next 1-3 years.

Cloud Attribute	Today	1 year	1-3 years
Robust Cloud Infrastructure	<ul style="list-style-type: none"> Architected for risk mitigation/ fault tolerance Isolated networks SSAE16 Disaster recovery (DR) built in 	<ul style="list-style-type: none"> Audit-based environment DDoS mitigation High- performance offerings 	<ul style="list-style-type: none"> DR integrated testing and validation Complete stack isolation Performance-based QoS Auto failover between environments
Automation	<ul style="list-style-type: none"> Auto-scaling Automatic infrastructure monitoring Orchestrating deployment 	<ul style="list-style-type: none"> Application awareness Data-driven scaling Application marketplace Auto-patching & updating 	<ul style="list-style-type: none"> Autonomics (self-healing) Data-driven scaling Automatic notification
Management	<ul style="list-style-type: none"> Infrastructure-based Organization-based Server grouping and bulk update capabilities 	<ul style="list-style-type: none"> Identity federation Role-based permissions 	<ul style="list-style-type: none"> Workflow management
Services	<ul style="list-style-type: none"> IaaS PaaS 	<ul style="list-style-type: none"> IaaS & PaaS in same network DBaaS (SQL & NOSQL) Hadoop as a Service Messaging as a service Application marketplace 	<ul style="list-style-type: none"> Extract, transform, and load (ETL)
Simplicity	<ul style="list-style-type: none"> Pay by hour for resources consumed Choice in sizes API-driven Public or virtual private 	<ul style="list-style-type: none"> Federation Interoperable Public, virtual private, private or hybrid 	<ul style="list-style-type: none"> Supports open standards

IT leaders need to discuss these attributes with potential service providers today, because commodity clouds and enterprise clouds are becoming easier to spot. This will help ensure that the cloud providers they engage are planning to enable the enterprise for years to come.

Conclusion: Now Is the Time to Get on the Right Enterprise Cloud

Enterprises are making plans to shift major workloads to the cloud over the next 12-18 months. As they do, IT leaders will need to demand of cloud providers the same type of agility and sustainable efficiency from the cloud that they have realized from internal virtualization initiatives.

By asking the right questions of cloud providers regarding the architecture of their platforms and the baseline infrastructures on which they operate, organizations can quickly narrow down their enterprise-grade cloud provider choices.

The ideal cloud provider will demonstrate a history of global infrastructure with enterprise-class service and support as well as a unified cloud stack that simplifies and streamlines automation and management of the cloud environment, even as it grows in size and complexity. That defines enterprise cloud done more easily.

About the Authors:



Jared Wray
Founder and Chief Executive Officer
Tier 3

Jared Wray is founder and chief executive officer of Tier 3, the enterprise cloud platform company. An early cloud pioneer, Wray founded Tier 3 in 2006 to address the need for enterprise on-demand services. Today he is the visionary chief architect of Tier 3's enterprise-class cloud which Wray designed for IT and developer operational agility via advanced management, innovative orchestration and fabric services all paired with a secure, highly available, highly automated platform. As the originator of the Iron Foundry open source PaaS project, Wray is a thought leader in the evolution of cloud services for the entire enterprise application portfolio. A serial entrepreneur, Wray previously founded Dual, an interactive development firm with clients such as Microsoft and Nintendo.



Robert Miggins
SVP of Business Development
PEER 1 Hosting

Robert Miggins is the senior vice president of business development for PEER 1 Hosting. He has worked for more than 10 years in IT infrastructure, with all aspects of the industry including sales, marketing, product development, and operations. Robert has been with PEER 1 for more than four years and is responsible for developing the company's product roadmap, strategic partnerships, and customer loyalty initiatives. He also heads up PEER 1 Hosting's programs targeting vertical markets, such as SaaS and hosting resellers. Prior to holding his current position, Robert was the vice president of ServerBeach, where he led sales, marketing, tech support and operations for the dedicated hosting company. Prior to ServerBeach, Robert worked for four years at Rackspace, holding roles such as vice president of sales, vice president of product management and general manager of Rackspace UK. Robert received a Bachelor of Arts in Economics from Washington and Lee University and an MBA from the University of Texas.

Footnotes

¹ Staten, James, "Are you opportunistic or strategic with your cloud investments?" Forrester Research (Forrester Blogs), July 18, 2012. Accessed September 2012 from the following online source: http://blogs.forrester.com/james_staten/12-07-18-are_you_opportunistic_or_strategic_with_your_cloud_investments

² Csaplár, Dick, "The Cost of Downtime is Rising," Research with Results (a blog from the Aberdeen Group), February 28, 2012. Accessed October 2012 from the following online source: <http://blogs.aberdeen.com/it-infrastructure/the-cost-of-downtime-is-rising/>