White Paper

Unlocking the Promise of Cloud Computing for the Enterprise

Achieving scalability, agility and reliability with cloud application platforms

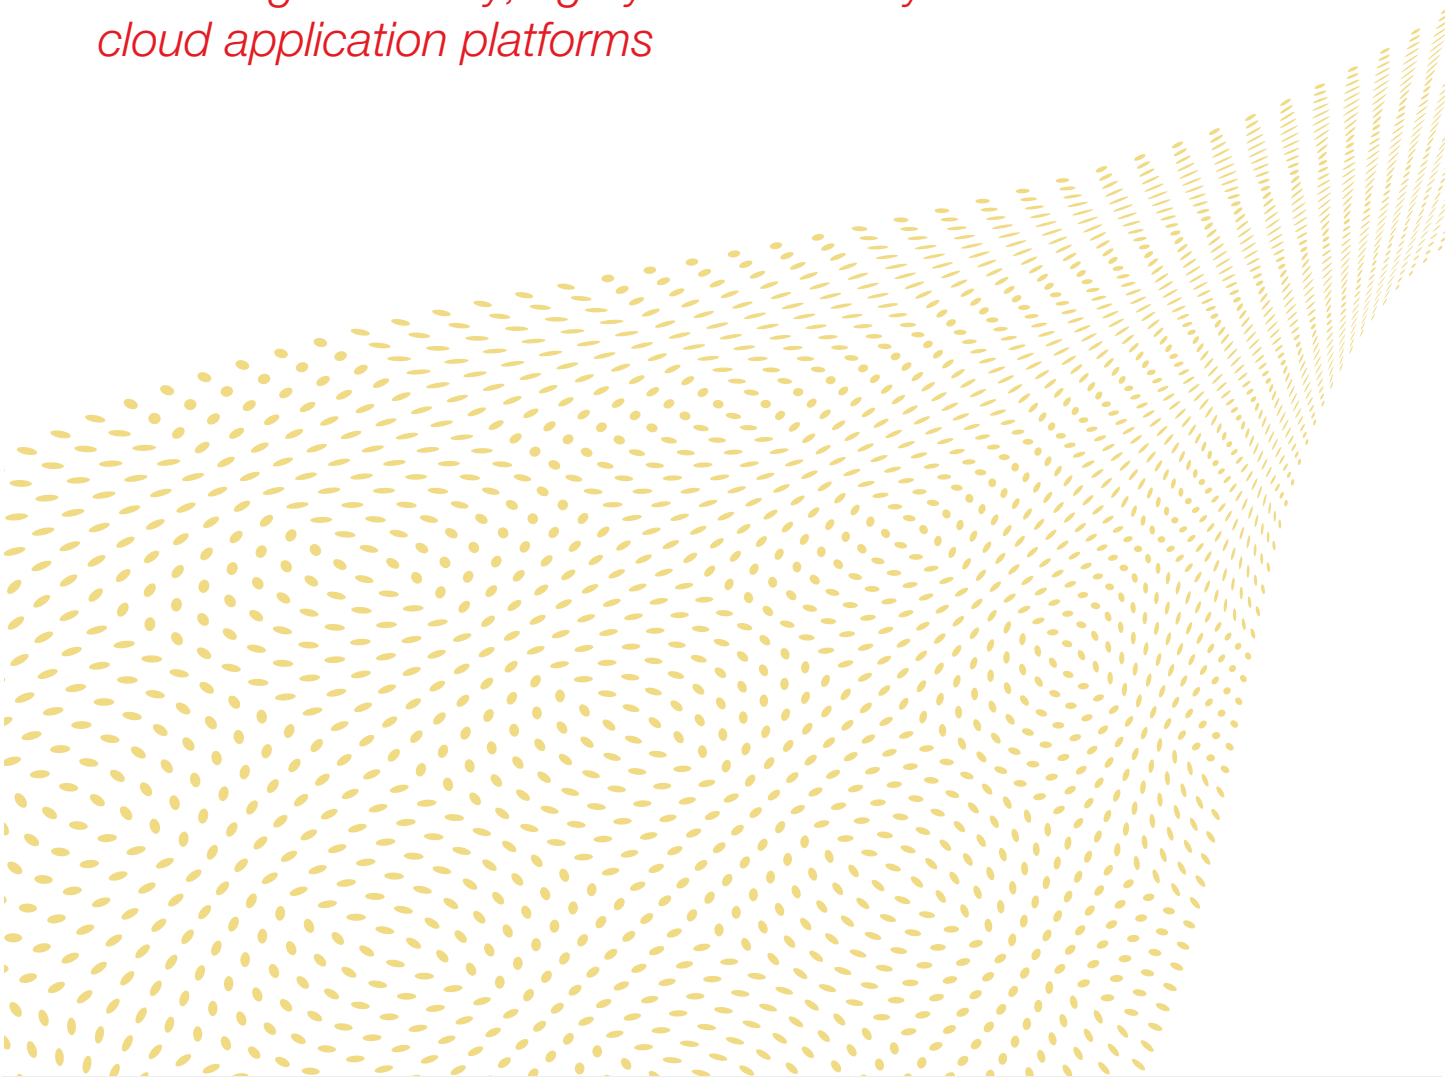

Contents

Executive Summary

The enterprise landscape is rapidly changing. Data is ubiquitous. Information is flowing into an organization's applications from more sources than ever before. Business expectations are also changing. Corporations today demand speed and flexibility from their applications. They want services that allow them to make better business decisions, create more satisfied customers, and react ever more quickly to evolving market conditions. Current economic circumstances and increased competition are also driving the demand for a more effective model to deliver applications and services.

This relentless push for a faster, better and more cost-effective technology delivery model has set the stage for new approaches to application development, deployment and management. Several technologies such as grid computing, virtualization, and service-oriented architecture (SOA) have offered partial solutions for enterprises that require applications with greater scalability, agility and easier management capabilities. However, these alone have not been enough.

Enter cloud computing, an innovative model for delivering IT infrastructure, applications and data that shifts the emphasis from static, stand-alone application silos to dynamic, shared environments, dynamically allocated among various tasks and accessed via a network.

Today, many forward-thinking enterprises are using cloud environments to take advantage of the increased scalability, agility, automation, and efficiency that this technology can deliver. Yet, because cloud computing has evolved so quickly, there are still many questions surrounding it. To help business and IT professionals truly understand the promise of cloud computing, this paper will examine its development and benefits from an enterprise perspective.

Beginning with the origins of cloud computing, this paper will help define exactly what cloud computing is and how the enterprise can benefit from it. In doing so, the paper outlines a number of "cloud characteristics" which together illustrate the true potential of cloud computing and provide a framework for assessing current and future cloud offerings. Finally, the paper draws a distinction between infrastructure-oriented clouds and platform-oriented clouds and explains how cloud platforms allow end-user applications to unlock the true promise of cloud computing.

Introduction

The term cloud computing has quickly grown from a little-known buzzword into one of the hottest topics in IT today. This surge in interest has led to a great deal of confusion about what cloud computing is and how to apply it to the enterprise. This is particularly true for large organizations, which have much to gain from cloud computing if done correctly.

Cloud computing describes a highly efficient way for an organization to deliver and consume IT infrastructure, applications and services. In its simplest form, cloud computing came from the computing model famously used by Web giants such as Amazon and Google to deliver their own services to customers. Yet, legitimate questions about this new technology still abound. For instance, how does cloud differ from virtualization or grid and utility computing? How does it differ from service oriented architectures (SOAs) or other Web platforms? And, what is meant regarding public or private clouds?

One thing becoming increasingly clear is that cloud computing offers compelling advantages for the enterprise. Many have turned to cloud computing for its ability to help bring new, more sophisticated applications to market more rapidly, and to scale those applications more easily—creating true competitive advantage in today's challenging economy. Yet cloud computing also offers good old-fashioned cost savings. For example, public clouds administered by a third party vendor can deliver large-scale enterprise projects at a fraction of traditional costs. Likewise, enterprises adopting private clouds have begun to realize reduced capital costs due to their ability to take better advantage of today's powerful commodity-grade hardware. At the same time, these organizations are lowering the total number of servers required due to increased efficiency. Enterprises also see lower recurring operational costs and administrative overhead due to the high levels of automation employed in cloud environments.

Where Did Cloud Computing Come From?

The concept of cloud computing can be traced back to ideas such as utility computing, which envisioned a future where packaged information technology services would be metered and delivered to customers much like electricity, gas and water. Over the past decade, the rise of pervasive network computing has provided plenty of working examples that approach the utility computing vision. With the commoditization and standardization of hardware and operating systems, the growing deployment of virtualization, the rise of service oriented architectures (SOAs), and the Web's astounding growth, lessons learned about providing common enterprise services delivered through hardware, software and people have become more easily repeatable.

At the same time, perpetual licensing models are beginning to fall by the wayside as many services such as network security management, data center hosting or even departmental billing are now delivered on a pay-per-use basis. Complex capabilities of all types are now available on demand to enterprises and individuals alike. Newly formed business models, such as software-as-a-service (SaaS), have allowed forward-looking software firms to build their companies on delivering these shared services.

Today, examples abound of companies offering utility services falling under the "cloud" banner. Amazon's Elastic Compute Cloud (EC2) offers computational services that enable people to use CPU cycles without buying more computers. Storage services such as those provided by Amazon's Simple Storage Service (S3) or companies like Nirvanix allow organizations to store data and documents without adding a single on-site server. SaaS companies like Salesforce.com deliver CRM services, so clients can manage their customer information without installing specialized software.

All of these offerings have been described as cloud computing. True, they all seek to deliver massively scalable and elastic IT services for a decreasing cost. Yet they differ from each other in many significant ways. So, the question remains: what is cloud computing?

What is Cloud Computing, Anyway?

Cloud computing is a computing paradigm based on the use of shared IT resources, dynamically allocated to specific users and tasks and accessed as a service via an application program interface (API). A cloud environment enables users to run applications by deploying them to the cloud, which acts like a virtual datacenter. The physical cloud resources may reside in a number of locations, the details of which are not typically known to the service's users. Cloud resources are offered as a service on an as-needed basis, and delivered by IP-based connectivity. The cloud itself typically consists of large numbers of commodity-grade servers, harnessed to deliver highly scalable and reliable on-demand services.

It's important to understand that cloud computing is an evolving paradigm ultimately defined by developers and early adopters. While it's important to have a good working definition of cloud computing, a more complete understanding of cloud computing requires an examination of its functional attributes. It is these characteristics that define the true value of cloud computing for the enterprise.

The Promise of Cloud Computing

Cloud computing offers a new level of capability, flexibility and manageability for enterprise IT infrastructure and the applications that run on it. The enhanced capabilities promised by the cloud can be understood as a set of ideal characteristics that cloud users want to ultimately achieve for their applications. These characteristics include:

- **Elastic Scalability.** Cloud environments allow businesses to serve larger audiences; solve bigger, more challenging problems; access incremental compute resources on-demand; and reduce the risk of new projects by starting small and growing as the need develops.
- **Agility.** As a shared resource, the cloud improves efficiency by allowing users to have multiple applications share the cloud environment and by providing flexible, automated management to distribute the computing resources among the cloud's users and their applications.

Fig 1. Cloud computing takes the best of a number of related technologies to provide a new style of computing

	Key Feature	What's Missing?
Traditional Grid Computing	Job scheduling across many machines for computational applications	Difficult to administer; Lacks agility, robustness and broad applicability of cloud computing
Virtualization	Virtual machies decouple operating systems from hardware	Virtual machines fail to fundamentally solve scale and reliability
Hosting	Computational infrastructure available for rent	Lacks on-demand services
Software-as-a-Service (SaaS)	Application availability through the cloud	SaaS is an application on the cloud, not a cloud by itself
Utility Computing	Packaged computing, application, and storage sold as a service	Describes a business model, not a technology approach or architecture

- **Reliability and Fault-Tolerance.** Cloud environments take advantage of the built-in redundancy of the large numbers of servers that make them up, enabling high levels of availability and reliability for applications—providing a better customer experience and eliminating disruption to business due to infrastructure failure.
- **Shared Multi-tenancy.** By enabling IT resources to be consolidated, multiple cloud users and/or applications share a common infrastructure, allowing for improved efficiency and costs to be more effectively managed without sacrificing the security of each user's data.
- **Service-oriented.** The cloud is a natural home for service-oriented applications, which need a way to easily scale as services get incorporated into other applications. This allows new applications to be brought to market more quickly, based on capabilities that have already been developed as services.
- **Utility-based.** Users only pay for the cloud resources they use, either by subscription or transaction-based models.
- **SLA-driven.** Clouds are managed dynamically based on service-level agreements that define policies like delivery parameters, costs, and other factors, so enterprises can rely upon a service/application once it has been published.
- **APIs.** Because clouds virtualize resources as a service they must have an application programming interface (API). The API extends the cloud and its applications to meet unforeseen business needs.

Not every cloud environment is able to achieve all of these characteristics simultaneously. In fact, many of these characteristics cannot be achieved through cloud infrastructure alone but rather must be addressed at the application level, either by its developers or by a cloud application platform upon which it runs.

Related Technologies

Because cloud computing is at the convergence of many independent computing trends, early attempts to define cloud computing have tended to do so relative to existing technologies, such as grid computing, virtualization, hosting, SaaS and utility computing. To better understand the cloud and the characteristics that cloud users want to achieve for their applications, it is helpful to examine these technologies. While these technologies have contributed to the development of the cloud, they also fall short of the ultimate promise of cloud computing.

- **Grid Computing.** Grid computing has become synonymous with scalability, and most grid environments are shared resources. However, the complexity of traditional grid computing tools severely limits their agility. And while they are typically shared, they are usually not multi-tenant in that multiple applications are not able to simultaneously run on the grid. Traditional grids are also lacking with regard to their ability to provide reliability and meet SLAs.
- **Virtualization.** By decoupling virtual machines from physical machines, server virtualization can dramatically increase the flexibility and agility of an IT organization. Yet virtualization alone fails to solve the fundamental scalability and reliability problems for applications. In addition, virtualization typically lacks the automation required for the self-service, self-healing functionality of a cloud.
- **Hosting.** Hosting services provide outsourced space on servers in a datacenter and, in fact, many hosting companies have entered the cloud computing marketplace. However, these companies typically differentiate their cloud offerings from traditional hosting in that their hosting solutions are delivered in an on-demand, self-service, and API accessible manner.
- **SaaS.** Although software-as-a-service offerings are often hosted in cloud environments, SaaS itself primarily refers to the business model of delivering applications as-a-service.
- **Utility computing.** In the utility computing model, computing infrastructure and applications are packaged and sold as a service, like electricity. In many respects, cloud computing is closely tied to the utility computing business model, with cloud computing being a broader concept that relates to the underlying architecture in which the services are designed and delivered.

While all of these technologies have helped to enable cloud computing, attempts to view cloud based solely on these individual solutions misses many of its main points. Cloud computing is in fact the evolution and convergence of these seeming independent computing trends, resulting in a new style of computing in which the whole is greater than the sum of its parts.

Fig 2. The cloud computing ecosystem

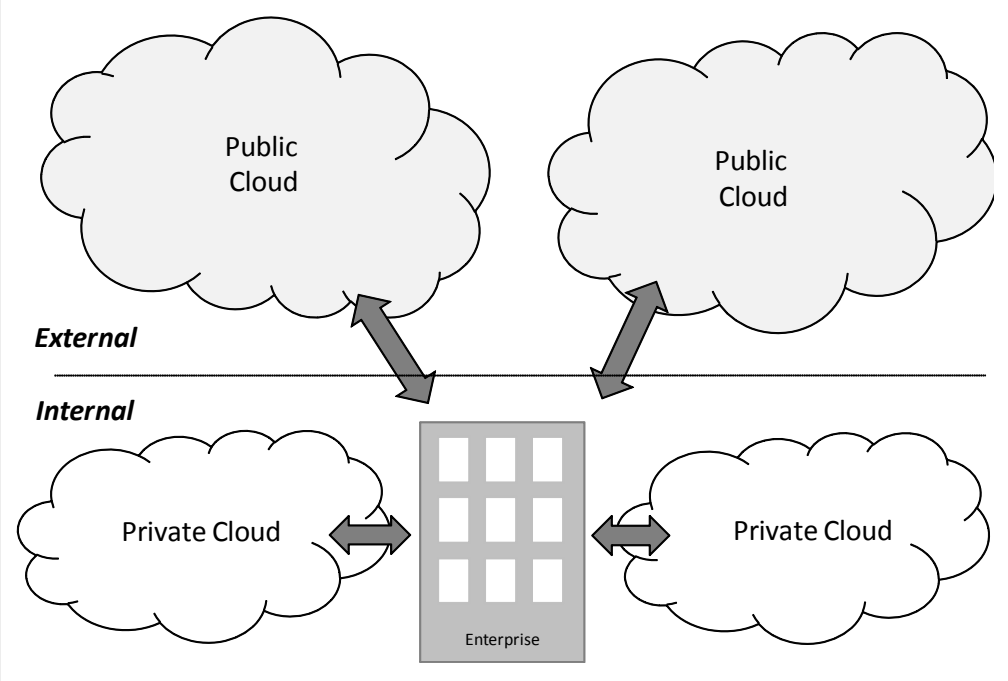

Public vs. Private Clouds

One commonly disputed point about cloud computing is the notion of public vs. private clouds. Some cloud computing practitioners, particularly those affiliated with cloud service providers, assert that in order for an IT environment to be a true cloud, it must be delivered via a third party and accessed via the Internet. Still others believe, as we do, that the true measure of a cloud is the extent to which it achieves characteristics such as those outlined above.

As enterprise leaders look to cloud computing as a way to accelerate growth and gain a competitive advantage they must choose whether they'd like their cloud administered by a third party offering 'public' cloud services, or to deploy their own 'private' cloud that resides within their organization's firewall. Here's a brief explanation of the advantages and disadvantages of each:

Public Cloud

Public clouds appeared first on the horizon and have helped define cloud computing best practices. As its name implies, a public cloud computing environment is open for use by the general public, which could include individuals, corporations or other types of organizations. Amazon EC2 is an example of a public cloud because it is available for use by the general public. However, due to the shared nature of cloud environments, security is a serious issue with public clouds. Concerns about compliance, governance and reliability also plague public cloud service providers. As an example, Amazon's S3 service was recently down for eight hours due to authentication issues.

At the same time, public clouds are successfully employed by many enterprises today. For example, public clouds are widely used to affordably accelerate enterprise application development and testing. Likewise, they are often used for large-scale projects, such as the New York Times archive project which used 100 Amazon EC2 instances and 5.5TB of S3 storage to generate PDFs of 11 million articles for the paper's archives, at a small fraction of traditional costs.

Private Cloud

In contrast, a private cloud computing-like environment exists within the boundaries of an organization, typically for its exclusive usage. To make a non-technical analogy, private cloud computing is like owning your own car, instead of using public transportation. Public transportation, while extremely efficient and cost effective, has not completely replaced personal automobile ownership for many obvious reasons.

Often, the creation of private clouds is driven by C-level executives' concerns around public clouds, such as security and compliance. Interest in private clouds is also driven by IT departments who have long sought utility-like IT environments, where computing resources and applications can be provisioned with greater efficiency. To this end, today's enterprise IT organizations are increasingly looking to private clouds as a way to tie together investments in system management software, grids and virtualization, and move them closer to achieving the cloud characteristics they seek. However, a move towards private cloud computing has one major drawback: IT departments have to buy, build and manage the cloud themselves.

It's clear that both public and private environments have their advantages and disadvantages. But both types of cloud computing offer the major benefits of reduced costs, simplified maintenance and management, along with greatly extended scalability.

Infrastructure vs. Platform Services

Beyond the issue of who owns and manages the cloud, two main variations of clouds have emerged, each distinguished by the level of abstraction offered by the cloud environment and the extent to which it is primarily concerned with providing raw infrastructure versus higher-level application services.

Infrastructure-as-a-Service (IaaS)

Infrastructure-oriented approaches to cloud, including Infrastructure-as-a-Service offerings, seek simply to provide access to virtualized computing resources in an on-demand manner. Typical of this approach is Amazon's EC2, through which any user can request Linux virtual machine instances that are created on the fly and billed based on actual usage. The user of cloud infrastructure always knows how many virtual machines they have and what their individual IP addresses are and, in the case of Amazon's EC2, know the "sizes" for each instance. Regarding EC2, clients don't know where the machines are geographically or what kind of hardware is being used. This is what makes the service cloud-like.

Fig 3. Cloud infrastructure vs. cloud platform

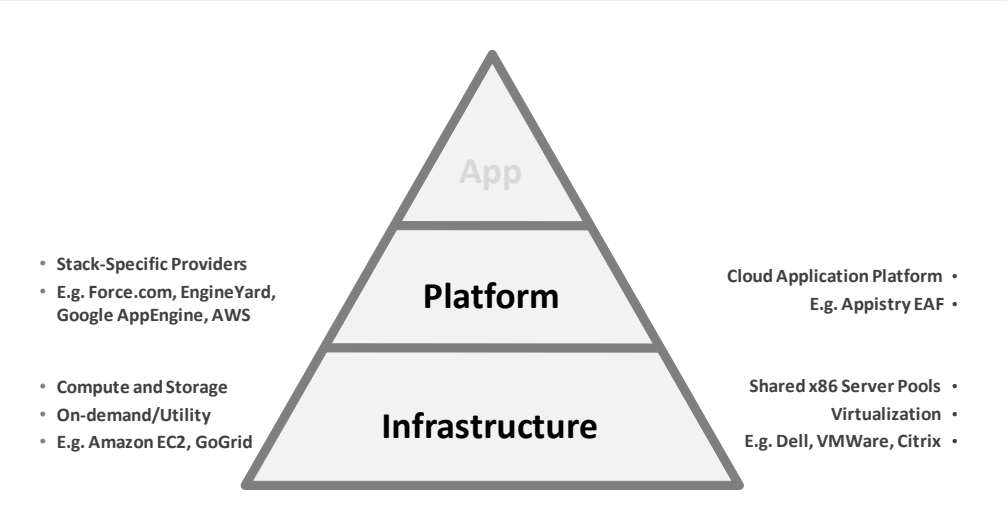

Platform-as-a-Service (PaaS)

Platform-oriented approaches to cloud are distinguished by the higher level of abstraction they provide as well as the supporting services they make available to the applications that run on them. Salesforce.com's Force.com or Google's AppEngine both typify the PaaS approach and the distinction between infrastructure-and platform-oriented cloud types. The AppEngine user is solely concerned about the application they are creating to run on the platform. To deliver an application they simply package it and deploy it to AppEngine. The deployment happens in a single step and the end-user doesn't know whether the application is being run on one virtual machine or 10 at any given point in time. In addition, the application can take advantage of special services provided by the AppEngine platform, such as authentication or data access.

The distinction between platform and infrastructure is a key consideration for organizations exploring cloud computing. Platform- and infrastructure-oriented approaches to cloud computing measure up against the cloud characteristics listed above in very different ways.

For organizations seeking the simplest cloud computing experience and the fastest time to market for applications, a cloud-based platform offers a distinct advantage. In addition to hosted PaaS offerings, stand-alone cloud application platforms are available which run atop infrastructure-oriented cloud services (IaaS) as well as private cloud environments. These platforms allow enterprises to reap the benefits of a platform-based approach with several advantages over hosted PaaS offerings, particularly with regard to increased portability and standards support, reduced lock-in, increased visibility and broader support for enterprise technologies such as Java and .NET.

Delivering the Promise of Cloud Computing for the Enterprise

Much of the excitement around cloud computing stems from the strong tangible and intangible benefits it offers to those companies that adopt it. Enterprises large and small, in a variety of industries, are implementing cloud computing platforms for their ability to be a significant source of both competitive advantage and cost savings.

Competitive Advantage

Cloud computing allows organizations to deliver highly-scalable and reliable applications more quickly and cost effectively. This yields tremendous advantage for enterprises that depend on or benefit from either large-scale or bursty applications, whose load tends to arrive in short spikes. With cloud-based applications in place, these enterprises are freed to deliver new products to their customers more rapidly. Because they are able to 'think big,' unencumbered by the constraints of scale, they are also able to make better business decisions by applying sophisticated analyses to the wealth of data available within them.

At the same time, the dramatic reduction in the complexity of the IT environment, the reduced provisioning times, and the ability to shift the allocation of cloud resources in response to changing market conditions creates a more flexible IT organization and a more agile enterprise. And because the enterprise is empowered to align IT infrastructure investments with project benefits, overall risk is reduced and the enterprise is able to invest in more strategic initiatives.

Cost Savings

Cloud computing offers public cloud users the ability to forgo capital investment for many IT projects, instead acquiring compute resources on-demand from a 3rd party cloud provider. Even those companies building their own clouds stand to benefit though, as a robust cloud computing environment allows enterprises to significantly drive reduced capital expenditures by enabling mission-critical applications to be deployed on inexpensive, commodity-grade servers. They further benefit by increasing the efficiency and utilization of the company's computing resources. Cloud also helps the enterprise reduce operational costs by increasing the level of automation in the data center and reducing the cost to deploy applications.

Cloud computing offers enterprises a new level of scalability, agility and cost-effectiveness for their applications. But without a cloud platform, an organization's applications will not reap the full benefit of cloud computing. Ultimately, the greatest advantages of cloud computing can only be realized if an organization's applications are able to take advantage of the cloud's inherent capabilities. For example, just because an organization is running its applications on Amazon's EC2 does not mean those applications can scale like amazon.com. In order to take advantage of the full scalability, agility and reliability of the cloud, applications must be built on a cloud-enabled platform.

Appistry: Enterprise Cloud Pioneer

Appistry is the leading provider of cloud application platform software today. The company's products unlock the power of next-generation cloud computing for organizations of all sizes and across a wide variety of industries, offering a more agile and scalable IT environment. Appistry's cloud application platform addresses the complex challenges of building, deploying and managing a wide variety of highly-scalable applications and services for both public and private cloud environments.

Cloud computing offers enterprises a new level of scalability, agility and cost-effectiveness for their applications. But without a cloud platform, an organization's applications will not reap the full benefit of cloud computing.

Appistry's products are designed specifically for the cloud, delivering transparent scalability, application-level fault tolerance, and automated management for new and existing applications. With a broad customer list that includes FedEx, GeoEye, Lockheed Martin, Northrop Grumman and others, Appistry's platform supports mission-critical applications for some of the world's leading organizations.

Appistry's flagship product, Appistry CloudIQ Platform, provides a run-time cloud application platform that complements an organization's existing technology choices and allows businesses to rapidly create highly scalable, service-oriented applications with minimal investment and risk. In addition to enabling linear application scalability in a cloud environment, Appistry CloudIQ Platform also simplifies development, reduces project risk, enhances reliability and ensures applications survive infrastructure failure.

Appistry's cloud computing application platform is at the convergence of grid computing, virtualization and SOA. Together, these technologies enable a new world of "real-time" IT applications and infrastructure, which offers great promise to the enterprise, coupling the scalability benefits of grid computing, the architectural benefits of SOA, and the operational benefits of virtualization.

Real-World Deployments of Appistry's CloudIQ Platform

Appistry helps customers in a variety of industries successfully take advantage of public and private cloud environments. Here are a few real-world examples of companies that have utilized Appistry CloudIQ Platform to successfully deliver cloud-enabled applications:

GeoEye Accelerates Development and Delivery of New Capabilities

Today's global intelligence market relies on super-accurate, reliable technological tools. GeoEye, the leading provider of map-accurate commercial satellite imagery to the U.S. military and intelligence community, chose Appistry CloudIQ Platform to provide a cloud application platform to help it accelerate innovation and deliver new capabilities. By simplifying the development, deployment and management of GeoEye's mission-critical software applications, Appistry CloudIQ Platform enabled the company to reduce capital and operational costs, while enhancing its ability to meet the evolving needs of its customers.

Running on 50+ Linux-based computers, Appistry CloudIQ Platform provided a self-managing and self-healing cloud environment that easily scales up or down to meet GeoEye's business requirements. For GeoEye, Appistry's cloud application platform

eliminated scalability limitations, enabling incremental application growth and eased the burden of relentless data inflow. It also reduced hardware and software acquisition and recurring costs nearly 80%, while simplifying application management through decreased operational complexity.

Clearent Improves Access to Insight for Merchants and VARs

In a highly competitive industry—the credit card processing business—the ability to bring products to market more quickly while maintaining high-quality service is a major competitive advantage. Conventional approaches to system development and deployment in the payment processing industry depend largely on legacy mainframe technology or, at best, fault-tolerant servers. However, this approach is anything but agile, and creates a steep perceptual barrier for new entrants to the market. Not only are up-front costs high, but unwieldy management regimes and steep growth costs come along for the ride as well.

Clearent chose to build its core applications on the Appistry cloud application platform and deploy them to a private cloud hosted in a secure data center, thus laying a strong technology foundation for the company's future growth. For Clearent, the advantages are clear—the cloud-enabled platform provided by Appistry CloudIQ Platform simplifies product development and allows the company to deliver high levels of scalability and reliability.

Communications Giant Reduces Costs by 94%

Global telecommunication is another highly competitive industry. One of the top U.S. providers of personal communication services differentiates itself from the competition by offering a wide range of products and solutions, including wireless, long-distance voice and data transport, global Internet protocol (IP), as well as local and multi-product bundles.

In order to introduce applications that support subscriber growth and deliver new services, the company also wanted to reduce operational complexity, ensure application reliability, and limit costs associated with building, deploying and maintaining network applications. For this purpose, the company utilized Appistry CloudIQ Platform to develop a private cloud for a highly transactional operational support system (OSS) application. Based on internal company estimates, the cloud-based solution reduced hardware and software acquisition costs by over estimated 94% and three-year recurring costs by in excess of 40%. It also insulates the application layer from physical infrastructure frailties, ensuring uptime. In response to growth in demand, the cloud platform effortlessly assimilates additional capacity, easing planning while it simplifies application management by decreasing operational complexity.

These real-world examples demonstrate that Appistry CloudIQ Platform can deliver highly scalable, highly agile large-scale applications with fewer layers of infrastructure than traditional deployment approaches. This has the potential to transform the way that organizations build, deploy and maintain software applications.

The Future of Cloud Computing for the Enterprise

Today's enterprises demand more effective approaches to deliver applications and services. While technologies such as grid computing, virtualization, and service-oriented architecture (SOA) have offered partial solutions for enterprises that require applications with greater scalability, agility and easier management capabilities, these approaches have not been enough.

Cloud computing, on the other hand, is at the convergence of these technologies, bringing a new level of capability, flexibility and manageability for enterprise IT infrastructure and the applications that run on it. However, not all cloud environments are created equal. For an enterprise's applications to truly achieve the full promise of the cloud, they require more than a cloud infrastructure alone. They require a cloud application platform.

Stand-alone cloud application platforms like Appistry CloudIQ Platform run atop infrastructure-oriented cloud services (IaaS) as well as private cloud environments. Appistry CloudIQ Platform also provides enterprises with several advantages over hosted PaaS offerings, including increased portability and standards support, reduced lock-in, increased visibility and broader support for enterprise technologies such as Java and .NET.

Going forward, there is little doubt that various forms of cloud computing will play an increasingly important role for organizations seeking a competitive advantage. Enterprises that employ cloud application platforms—in either public or private cloud environments—will benefit from the increased scalability and agility of their cloud-based applications. Cloud platforms will also help them to significantly reduce time-to-market, realize substantial cost-savings and react more quickly to changing market conditions. Because of these clear benefits to IT departments and to an organization's business goals, cloud computing is here to stay. And as it evolves from a buzzword at conferences to the key topic at corporate board meetings, cloud computing will continue to grow and deliver dramatic results for enterprises.

About Appistry

Appistry simplifies cloud computing for the enterprise, opening the door to a more agile and scalable IT environment. Appistry's application platform addresses the complex challenges of building, deploying and managing a wide variety of applications and services for both public and private clouds. Appistry's products are designed specifically for cloud environments, delivering transparent scalability, application-level fault tolerance, and automated management to new and existing applications. Appistry customers include FedEx, GeoEye, Lockheed Martin and Northrop Grumman.

For more information about Appistry, please visit **www.appistry.com**.

Appistry 10845 Olive Blvd., Suite 260, St.Louis, MO 63141 **main.** 314.336.5080 **fax.** 314.336.5086