White Paper

Akamai and Cloud Computing

A Perspective from the Edge of the Cloud

by Tom Leighton, co-Founder and Chief Scientist, Akamai Technologies

Table of Contents

Introduction

As one of the hottest concepts in IT today, "cloud computing" has been a subject of tremendous hype and much confusion. Even the term itself has been subject to multiple definitions and diverse interpretations. Only one thing seems clear: many different services, technologies, and strategies will be employed as businesses and IT vendors pursue cloud computing's potential.

And that potential is real. Cloud computing offers the promise of improved cost efficiencies, accelerated innovation, faster time-to-market, and the ability to scale applications on demand. For these reasons, SaaS applications and public cloud platforms have appealed to small and startup businesses in particular, as a way for them to gain easy, low-cost access to systems that would otherwise cost them millions to build. At the same time, cloud computing has also drawn the cautious but serious interest of larger enterprises.

This whitepaper provides a framework for evaluating the cloud computing marketplace by exploring its enabling technologies and current offerings, as well as the challenges it faces. We then consider Akamai's role as a provider of critical cloud optimization services that will help cloud computing fulfill its promise to deliver flexible, efficient, business-critical infrastructure for the enterprise.

Understanding the Cloud

Simply defined, cloud computing refers to computational resources ("computing") made accessible as scalable, on-demand services over a network (the "cloud"). And yet, cloud computing is far from simple. It embraces a confluence of concepts — virtualization, service-orientation, elasticity, multi-tenancy, and pay-as-you-go.

To facilitate our discussion, we will first lay out a cloud computing framework that is useful for categorizing the different offerings in the cloud computing space. Then, we will examine the characteristics of the various types of clouds being leveraged today.

2.1 The Cloud Computing Framework

Our cloud computing framework has five key components. The first, virtualization, can be thought of as an underpinning of cloud computing. By abstracting software from its underlying hardware, virtualization lays the foundation for enabling pooled, shareable, just-in-time infrastructure. On top of this technology base, we can then describe cloud computing's principal offerings (*Infrastructure-as-a-Service, Platform-as-a-Service, and Software-as-a-Service*) as a layered stack. Cloud optimization is the final, critical piece of the framework — encompassing the technologies and solutions that enable cloud computing to scale and deliver the levels of performance and reliability required for it to become part of a business's core infrastructure.

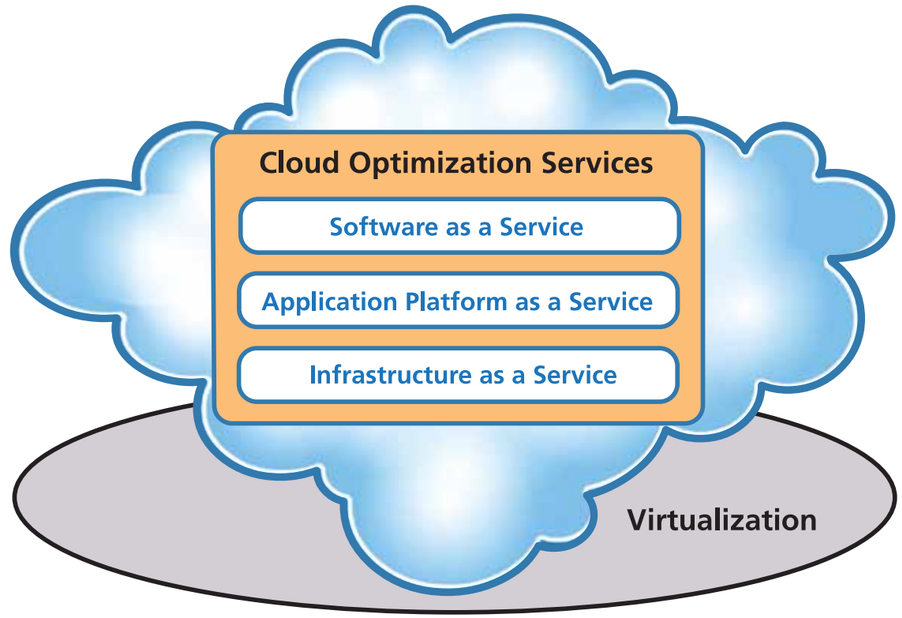

We will now look at each of these framework components in more detail.

Virtualization

Gartner Group describes virtualization as the "highest impact trend changing infrastructure" — a fitting description for the technology that gave birth to the current cloud computing frenzy.

By abstracting server software from its underlying hardware, server virtualization improves the efficiency and availability of resources and applications running on that server. Consider that, according to Gartner, roughly 80% to 90% of enterprise computing capacity is unused at any given time. Virtualization enables these once-idle CPU cycles to be utilized.

Taking this concept to the cloud means extending it — beyond making better use of a single physical machine or machine cluster — to the aggregation of computing resources across multiple data centers, multiple applications and multiple tenants, allowing each to scale up or down on demand. For example, it allows application vendors like Salesforce.com to efficiently manage multiple customers from a centralized data center. Likewise, it enables cloud computing providers to build large, automated server and storage services available as on-demand services over the Internet.[1]

Companies offering important cloud virtualization technologies include:

- **VMWare –** VMWare is the market leader in enterprise virtualization software. VMware's ESX server runs directly on server hardware and provides a completely virtualized set of hardware (a VM or virtual machine) to the guest operating system. VMWare also provide a range of other products and tools to improve the reliability and manageability of virtualized server infrastructure.
- **Xen –** The Xen hypervisor has fueled market adoption with an open source model (www.xen.org). It has been embedded into virtual machine server offerings by firms like Sun (xVM server) and Oracle (VM server), as well as being used by Amazon in their EC2 cloud computing service offering. Xen was initially commercialized by XenSource, which was purchased by Citrix in 2007 and is now available as Citrix XenServer Enterprise Edition.
- **Microsoft –** The newest entrant into the hypervisor market is Microsoft's HyperV. Originally known as Windows Server Virtualization, HyperV was released in 2008 and primarily supports Windows operating systems as guests, although it does support SUSE Linux Enterprise Server as well. At this time, Microsoft HyperV is less mature than VMWare's offerings. It currently does not support "live migration" of a VM from one hardware server to another (as VMWare can), but this functionality will be available in the second half of 2009.

Infrastructure-as-a-Service

A key component of cloud computing's core offerings, *Infrastructure-as-a-Service (IaaS)* makes basic computational resources, like storage, disk space, and servers, available on demand as services. Rather than using physical machines, customers use IaaS to get access to virtual servers on which they deploy their own software, generally from the operating system on up. Multi-tenancy, pay-as-you-go and extensible virtual machine instances are key attributes.

Currently, the compelling values IaaS brings to businesses and ISVs are cost savings and risk reduction. IaaS eliminates the substantial capital expenditures required when deploying infrastructure and large-scale applications in-house. Cloud providers generally offer a pay-as-you-go business model that allows companies to scale up and down in response to their real-time business needs, rather than having to pay up front for infrastructure that may or may not get used, or having to overprovision infrastructure in order to address occasional peaks in demand. To date, IaaS has seen heaviest adoption among small to mid-sized ISVs and businesses who don't have the resources or economies of scale to build out large IT infrastructures.

Examples of cloud IaaS vendors include:

- **Amazon –** Amazon has a number of IaaS offerings. Running on Xen's virtualization technology, Amazon's Elastic Compute Cloud (EC2) allows customers to "rent" virtual servers by the hour for rates as low as $0.10 per hour. Amazon Simple Storage Service (S3) is a virtual storage service that allows customer to "rent" storage space of as little as $0.15 per GB per month (traffic into and out of S3 is charged separately). Additionally, Amazon's CloudFront offers basic file-caching content delivery services for as little as $0.17 per GB in the US. All of these services are priced very aggressively, however Amazon's cloud offerings have had several outages and are not considered enterprise class at this time.
- **GoGrid –** GoGrid is a cloud startup that delivers virtual servers (Windows and Linux) and storage on demand. It is similar to Amazon in concept, but charges based on server RAM hours (the number of server hours times the amount of RAM reserved on each server). GoGrid competes with Amazon by touting a superior Web management portal, better support, and lower prices on volume deployments.
- **Joyent –** Another IaaS startup, Joyent provides virtualized servers running the OpenSolaris operating system for as little as $45 per month for 256 MB of RAM. Claiming to be the world's largest OpenSolaris installation, Joyent differentiates itself by supporting open standards and touting superior scalability, as well as by providing professional services to help customers performance tune their Joyent infrastructure. Joyent also provides storage on demand using the Sunfire X4500 storage system, with 282 TB of storage currently under management. Joyent counts companies like ABC/Disney, LinkedIn, and Major League Baseball among its clients, and it supports over a quarter of the daily active application usage on Facebook.

- **Akamai –** Brought to market nearly a decade ago, Akamai's NetStorage is an on-demand, multi-petabyte storage offering, purpose-built to provide dynamic scaling and provisioning for Web application needs. NetStorage, as well as Akamai's CDN services, are IaaS offerings that work as part of a comprehensive suite of solutions for improving the performance and reliability of site and application delivery.
- **Content Delivery Networks (CDNs) –** The many CDN vendors in the marketplace, led by Akamai, can also be categorized as IaaS providers, as they reduce the need to scale out Web servers.

Platform-as-a-Service

Another emerging component of the Cloud Computing Framework is *Platform-as-a-Service (PaaS)* offerings, which provide environments for easily developing and deploying scalable Web applications — without needing to invest in or manage any underlying infrastructure. By providing higher-level services than IaaS — such as an application framework and development tools, PaaS generally provides the quickest way to build and deploy applications, with the trade off being less flexibility and potentially greater vendor lock-in than with IaaS.

The PaaS landscape is broader and includes vendors such as:

- **Force.com of Salesforce.com –** Built on the same proprietary virtualization and multi-tenant management technology as Salesforce.com's flagship CRM application, Force.com is a development and deployment platform that is tightly-coupled with the Salesforce.com application stack, requiring customers to use the same underlying data structures and application server. Force.com applications are built using Apex (a proprietary Java-like programming language) and Visualforce (an XML-like syntax for building user interfaces in HTML, AJAX or Flex). Force.com also provides AppExchange, a directory of applications which users can purchase and add to their Salesforce environment. As of September 2008, there were over 800 applications available from over 450 ISVs.
- **Google –** Targeted at content-centric (as opposed to transactional) applications, Google's App Engine aims to enable rapid development of Web applications that will run on Google's infrastructure. At this time, Google App Engine supports both the Python and Java languages for application development. The Google App Engine environment is intended for Web application development and is generally not suitable for enterprise applications.
- **Microsoft –** Microsoft's Azure Services Platform is a cloud application platform that allows applications to be hosted and run at Microsoft datacenters. Windows Azure is the cloud operating system that serves as the runtime for hosted applications and provides a set of services that allows development, management and monitoring. Microsoft Visual Studio will be used to develop applications that will run on Azure. Microsoft Azure is currently in Community Technology Preview, which is set to end in Q2 2009.
- **Akamai –** Akamai's EdgeComputing service distributes J2EE applications to edge servers that create virtual application instances on demand. Akamai's unique edge distribution offers extreme scalability and low application latency due to proximity to the end user. Like NetStorage, EdgeComputing is part of a suite of solutions that work together to comprehensively support a site or application.
- **Oracle –** Oracle's SaaS Platform delivers an Oracle-centric suite of database, middleware, management and VM technologies on which ISVs can build SaaS applications. Oracle provides technology training and support as well as joint marketing opportunities to promote SaaS with Oracle.
- **RightScale and Elastra –** These two startup companies can be thought of as providing an application platform environment for Amazon's cloud. They provide pre-configured application servers and database servers, as well as application development and management/monitoring environments for Amazon's EC2 infrastructure.

Software-as-a-Service

The best enterprise-ready example of Cloud Computing is found in *Software-as-a-Service (SaaS)*, where complete end-user applications are deployed and managed by a third party and delivered over the Web. SaaS continues the cloud theme of low-cost off-premise systems, on-demand availability, and pay-for-use models, while further eliminating development costs and lag time. This enables organizations to bring services to market quickly and frees them from dependence on internal IT cycles. The speed and ease with which SaaS applications are purchased and consumed has made this component of the cloud computing framework the most widely-adopted today.

Important cloud SaaS vendors and services include:

- **Salesforce.com –** Salesforce.com popularized the Software-as-a-Service model for delivering business applications in the cloud with their flagship CRM offering. They followed with the Force.com PaaS to further strengthen their ecosystem.
- **NetSuite –** A strong Salesforce.com competitor, NetSuite offers Accounting, ERP, CRM, and e-Commerce software.
- **Demandware –** An emerging SaaS-based e-Commerce engine, Demandware uses a metadata-driven architecture to allow customers to run on a shared platform but still customize the look and feel of their shopping experience.
- **Workday –** Workday is a Human Resources SaaS provider that is seeing traction in larger enterprises with its Human Capital Management, Financial Management, Payroll, Procurement, Resource Management and Business Intelligence application services.

- **SAP Business ByDesign –** Targeted at small and medium business, Business ByDesign offers HR, Finance and other ERP applications hosted by SAP and available in an on-demand model.
- **Hotmail, Yahoo! Mail, GoogleMail, Cisco WebEx, Adobe Web Connect –** These early Web-based communications applications have evolved to dominate the largest segment of the SaaS market according to Gartner.[2]

Cloud Optimization Services

Cloud optimization services provide scale and reliability for the previously-described components of cloud computing. They enable cloud offerings to operate across an unpredictable and unreliable Internet while delivering the robust levels of service required by enterprises.

The value of cloud optimization services can be understood as a direct function of application adoption, speed, uptime and security. For example, a site leveraging IaaS components that fail to scale for a flash crowd will lose customers and revenue. Likewise, a SaaS application that is slow or unresponsive will suffer from poor adoption. Thus cloud optimization is critical in order for cloud computing services meet enterprise computing requirements.

Akamai is uniquely positioned as a provider of multiple cloud optimization services:

- **Acceleration –** Akamai's Web Application Accelerator and Dynamic Site Accelerator solutions (for HTTP-based applications) and IP Application Accelerator (for all other IP-based applications) leverage unique technologies to deliver global acceleration and scalability for SaaS and other Cloud-based applications.
- **Security –** Akamai's intelligent edge network, along with its Web Application Firewall, SiteFailover and Global Traffic Manager offerings, provide real time, in-the-cloud security and business continuity services for SaaS applications and cloud-based application components.
- **Distributed computing –** Akamai's cloud-based NetStorage and EdgeComputing services offer distributed data storage and application execution, respectively. Both work as part of Akamai's comprehensive suite of optimized site and applications delivery solutions.

2.2 Types of Clouds

Cloud type is another important distinction to consider when evaluating the cloud computing landscape. Cloud type refers to the type of networks over which the cloud services are accessed: the public Internet or private internal networks.

Public Clouds

Most of the early spend and traction for cloud computing have been focused on public cloud services — those that are accessed over the public Internet. Public cloud offerings embody the flexible, pay-as-you go benefits that have driven the cloud computing hype. But by their very nature, public cloud services suffer from their reliance on the Internet, and they often have serious performance, security and reliability challenges. For this reason, cloud optimization services, like those provided by Akamai's portfolio of security, acceleration, and business continuity solutions, are critical to the success of public cloud offerings. We will talk more about this in section 3.1 below.

Examples of vendors pursuing public cloud strategies include:

- **Business applications providers –** Salesforce.com, NetSuite, and WebEx
- **E-Commerce application or component providers –** Demandware, BazaarVoice, Recommind and Adobe (Scene 7)
- **Consumer-focused application providers –** Yahoo! Mail, Google Apps, and hundreds of similar offerings
- **IaaS/PaaS offerings –** Akamai NetStorage and EdgeComputing, Google App Engine, Amazon EC2, and EMC's Mozy

Private Clouds

The concept of *private clouds* has emerged more recently as a way for enterprises to achieve some of the efficiencies of cloud computing with an infrastructure internal to their organization, thus increasing perceived security and control. By implementing virtualization and other cloud computing technologies on single-tenant hardware infrastructure behind their firewall, enterprise IT teams can enable pooling and sharing of compute resources across different applications, departments, or business units.

Private clouds require up-front development costs, on-going maintenance, and significant internal expertise, and therefore provide a much different benefit profile compared to public clouds. Private clouds are most attractive to enterprises that are large enough to achieve economies of scale in-house and where the ability to maintain internal control over data, applications, and infrastructure is paramount.

Vendors pursuing or supporting private cloud strategies include:

- **VMWare –** VMWare is focused on leveraging its stronghold in enterprise virtualization to heavily promote the adoption of private clouds. This is a logical strategy for the company, as most public cloud SaaS, PaaS and IaaS providers tend rely on either proprietary virtualization software or the open source Xen hypervisor, as opposed to paying for commercial VMWare software.
- **IBM –** While IBM does provide pay-as-you-go access to a suite of services like DB2, WebSphere Portal, and Lotus Web Content Management on Amazon EC2, its main cloud computing focus is on the private cloud. Private clouds provide a ready market for IBM's services business as well as a more natural fit with IBM's hardware, middleware, and infrastructure offerings. IBM's initial foray into cloud computing includes building out a practice for developer cloud platforms as well as thin-client desktop virtualization capabilities for enterprises
- **HP –** HP is supporting private cloud strategies through thought leadership in Dynamic Cloud Services, which are personalized based on location, preference and the communities a user participates in. HP believes that such services will depend on 4 components: more intelligence in the client, more intelligence in the network, next generation data centers, and software to automate and manage the cloud. HP is actively developing in all four areas. HP recently announced HP Cloud Assure, a new Software-as-a-Service (SaaS) offering designed to help businesses safely and effectively adopt cloud-based services.
- **Cisco –** Cisco has recently positioned itself to become the vendor of choice for building the next generation of enterprise data centers through its Data Center 3.0 strategy. In March 2009, Cisco unveiled their Unified Computing System, a new data center architecture that integrates virtualization technologies with Cisco's carrier-class networking infrastructure to create an efficient, low-cost, cloud infrastructure that meets the security and service level requirements of the enterprise.

Hybrid Clouds

Hybrid clouds combine cloud computing and optimization services across private, public and even non-cloud environments. Companies like Akamai, HP, Cisco, Microsoft, EMC and IBM support hybrid clouds as the most logical path to enterprise adoption of cloud computing. Hybrid clouds allow enterprises to approach cloud computing in a modular way, fitting the most applicable features and benefits of cloud and non-cloud environments to each specific business use case. For example, a company may leverage public cloud offerings for certain application components to achieve cost-effective scalability, while running other, highly-sensitive components of the same application within an on-premise, non-cloud environment.

2.3 Where is the Cloud?

As the confusion around cloud computing begins to settle, and the market begins to understand the benefits and tradeoffs of different types of clouds and cloud offerings, we should expect to see a shift in focus from the hype around cloud computing's potential benefits to the reality of implementing and using cloud-based solutions. As this reality sets in, questions that have been conveniently abstracted away (e.g., Where exactly is the cloud where these services running?) become critically relevant.

Despite the broad variety in cloud computing offerings, their underlying deployment infrastructures can be categorized into two basic architectures — namely, centralized versus highly-distributed. These two network architectures existed long before the cloud computing phenomenon; they are in fact the same architectures that underlie all Web-based infrastructures.

So while cloud computing may revolutionize the way infrastructure and applications are consumed, its underlying deployment infrastructure is nothing new.

Centralized Datacenters — New opportunity, Old approach

As with traditionally-architected Web sites, SaaS, PaaS and IaaS providers typically host their applications and services in a single location or a small number of datacenters.

For example, Amazon hosts EC2 in just three US datacenters and a single European datacenter. Similarly, Salesforce.com runs its main CRM application, as well as all the applications running on its Force.com platform, in just four US datacenters, plus a single point-of-presence in Singapore to serve international customers. Private clouds, whether collocated or on-premise, often run from a single location.

This approach is adequate when application users are very close to the application host location. For example, a single location can be adequate for a private cloud serving an on-premise group of employees.

However, for applications with distributed users, highly-variable demand, or shared service platforms the centralized datacenter approach is insufficient. Content and services delivered over the Internet rely on the proper functioning of and interplay between more than 13,000 competing networks. This means network outages, peering point congestion, routing inefficiencies, and other Internet middle-mile bottlenecks frequently cause application performance and reliability to fall short of expectations.

Highly Distributed Networks — Getting close to end users

By locating the cloud computing infrastructure in a highly-distributed manner, it is possible to overcome the challenges posed by the Internet's middle mile. Akamai is unique in taking this approach. While a number of other large cloud providers — including content delivery networks — do run multi-datacenter operations, these are fundamentally different from the highly-distributed infrastructure approach used by Akamai. With a centralized or multi-datacenter infrastructure, data and applications must still travel over the Internet's problematic middle mile to reach end users. A highly-distributed architecture — where servers are located at the edge of the Internet, close to end users (e.g., directly within the end user's ISP and in the end user's city) — is the only way to avoid middle mile bottlenecks and deliver LAN-like responsiveness for applications running over the global Internet.[3]

The Akamai Perspective on Cloud Computing

Akamai has built the world's largest distributed cloud optimization network, comprised of more than 42,000 servers in 1,500 locations, across nearly 1,000 networks worldwide. The Akamai network is uniquely capable of transforming the Internet into a reliable, high-performance platform, just as it has done for online businesses and Web sites over the last ten years. As computing moves into the public cloud, and as private clouds scale to provide global access, Akamai's network at the edge of cloud will become the unique enabler that sustains the cloud computing movement and drives its success. By developing new technologies to address inefficiencies in application, transport, and routing layer protocols as well as new products to address security and continuity of services, Akamai has established a network with the commercial-grade performance needed for enterprises to recognize the promise of cloud computing.

3.1 Accelerating Cloud Computing Applications

The fundamental enterprise requirements of cloud computing are performance and scalability. Akamai's services address the challenges of delivering interactive and dynamic applications and content at LAN-level performance through the cloud. Although caching services can help with origin offload and performance for static sites, it takes substantially more capabilities than a CDN to reach the enterprise requirements for cloud computing performance.

Akamai's Web Application Accelerator, Dynamic Site Accelerator, and IP Application Accelerator leverage Akamai's highly-distributed network and unique communications, routing, and application-layer optimizations to deliver global reach, responsiveness, and scalability for SaaS, PaaS, and IaaS applications. By routing around Internet trouble spots and delivering content and applications from servers located close to end users, Akamai's solutions dramatically improve the user experience and accelerate application adoption.

In Figure 2 below, for example, we see an illustration of the type of performance gain achieved when using Akamai's service for an application running on Amazon's EC2.

Global SaaS Application running on Amazon EC2 Infrastructure accelerated by Akamai's Web Application Accelerator

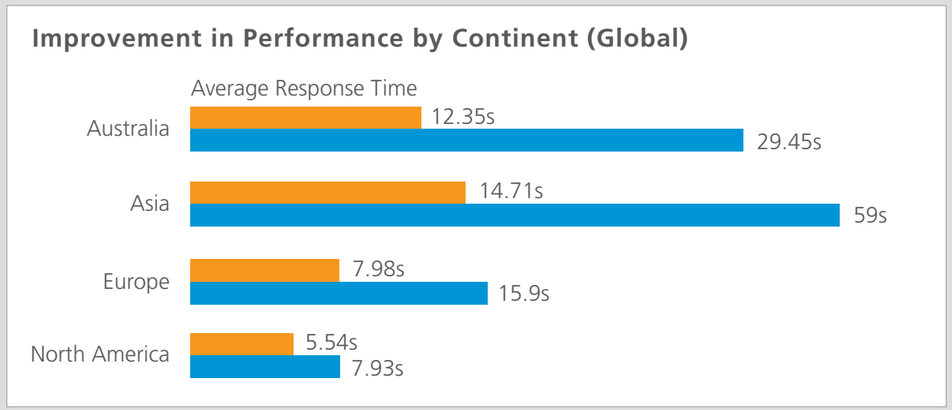

Application response times were at least twice as fast in Asia, Europe, and Australia with Akamai, compared to the origin third-party Cloud-based infrastructure

Customer Case Study: SaaS Acceleration

Bullhorn, a SaaS provider of front-office staffing and recruiting software, serves more than 1,000 firms, including some of the world's largest staffing firms.

Akamai's Web Application Accelerator accelerates the Bullhorn staffing application over the public cloud to provide LAN-like responsiveness for end users across the globe. Bullhorn credits Akamai with contributing to the astounding 1000% growth in quarter-over-quarter business it has seen from its international customer base.

3.2 Distributing Application Components to the Edge

The greatest possible application performance and scalability are achieved when the application itself can be distributed to the edge of the cloud, close to the end users, something Akamai has been doing since EdgeComputing and NetStorage were first introduced nearly a decade ago.

Akamai EdgeComputing enables companies to deploy and execute J2EE applications or application components on Akamai's edge servers, while supporting multiple forms of optimized back-end communications. Application instances are automatically created in different cities and regions based on real-time demand — something that cloud services such as Amazon EC2 and Google App Engine cannot provide. This allows EdgeComputing customers to enjoy truly maintenance-free scalability and unparalleled end user performance.

By deploying content-centric application components–such as site search, surveys and contests, or page assembly — at the edge, while running transaction-oriented application components at the origin (or public or private cloud), the application experience is optimized while running over a hybrid cloud.

Customer Case Study: Distributed Computing

With Akamai EdgeComputing, Sony Ericsson avoided the costly build out of several regional data centers, while successfully supporting its worldwide mobile phone customer base. Sony Ericsson deployed some application components — including a phone configurator, shopping cart, and dealer locator application — to EdgeComputing, while other application components ran in a centralized datacenter.

This hybrid cloud strategy reduced dealer locator application response time by over a factor of four and increased online application availability from 92% to 100%. Sony Ericsson offloaded nearly 100% of Java application server processing to the Akamai network, reducing its cost per user session by 33%.

3.3 Securing Cloud Applications and Platforms

Because of their reliance on Web infrastructure, SaaS and other applications running on public cloud platforms are as vulnerable to Internet threats and service attacks as traditional Web sites and applications. The Akamai network acts as a "secure perimeter" that eliminates public entry points to cloud infrastructures, helping to keep malicious DDoS attacks, Internet worms, hacker threats, and attacks on application vulnerabilities outside the origin data center. Akamai also applies technologies such as DNS security, IP layer protection and access control, HTTP origin cloaking, and application request checking as well as a suite of Web application firewall filters to protect cloud services right from the edge of the cloud.

Customer Case Study: Cloud Security

Recently, a major US government Web site came under a sustained 12-hour Distributed Denial of Service (DDoS) attack. The attack originated from computers in 7 different Chinese provinces using 11 different IP addresses. Akamai's network successfully defended against the attack, resulting in no detrimental impact to end-user performance or functionality.

For this site, a typical 12-hour period generates approximately 250,000 origin request;, however, during this attack there were approximately 1.8 million requests. This more sophisticated DDoS attack targeted a dynamic portion of the customer site by making various search request calls, initiating connections, and making full GET requests before issuing aborts. Due to the design of the Akamai Edge network and the customer's use of the Dynamic Site Accelerator's (DSA) "SiteShield" module, the vast majority of requests were absorbed by Akamai Edge servers in Asia, and never reached the origin servers in the United States.

3.4 Insuring Site and Application Availability

Akamai's services address enterprise business continuity and reliability requirements in several ways. Designed from the ground up to recover on its own from all types of failures — whether at the machine level or Internet-wide, the Akamai network offers a zero-downtime infrastructure.

For the enterprise running multiple data centers, multiple cloud instances, or using multiple cloud providers, Akamai's Global Traffic Management (GTM) is a cloud-based, highly-scalable, on-demand service that allows an enterprise to balance traffic between those entities based on a variety of business policy and Internet performance factors. For example, those policies can include automatic failover, weighted load balancing, or IP-based routing.

Akamai offers a number of other business continuity options in case of origin server failure. Akamai's SiteFailover includes the ability to monitor an origin server and, based upon setup, reroute origin requests to Akamai services, or an alternate origin data center. In addition, Akamai has a cloud-based, globally-distributed, on-demand service available to improve the performance, availability, and resiliency of an enterprise's mission-critical DNS infrastructure.

These services, combined with a 100% uptime SLA, enable enterprises to leverage cloud computing while maintaining the rock-solid availability their businesses demand.

Conclusion

As one of today's hottest IT topics, cloud computing is covered daily across the press, from academic journals to technology blogs and even in the travel section of the *New York Times*.[4] Most of the hype has focused on offerings in the public cloud, where, currently, centralized architectures are common. The drawbacks of this type of architecture have already begun to surface, as many of the major cloud vendors have suffered outages and downtime over the last year. As cloud computing moves out of the hype and experimentation mode into more mainstream adoption, businesses running applications on cloud platforms will look to Akamai's cloud optimization services to make the cloud responsive, scalable and secure.

Cloud Computing won't have a single vendor or a single cloud answer; its incarnations will be as varied as the applications and services it supports. As *The Economist* recently stated, "The computing sky will probably always be cloudy, meaning that there will be many private and public clouds, and they will come in all shapes and sizes. And most of them will be interconnected."[5] But regardless of the path the cloud computing evolution takes, Akamai's cloud optimization services will play a critical role in driving its growth, with innovative solutions that enable success for both cloud computing providers and the enterprises that use them.

[1] Gartner Group, Virtualization Changes Virtually Everything, March 2008.

[2] Gartner Group, Market Trends: Software as a Service Worldwide, 2007-2012, Sept. 2008.

[3] For a more in-depth examination of Internet bottlenecks and the benefits of highly-distributed networks, please see this author's article, "Improving Performance on the Internet," published in the February 2008 issue of Communications of the ACM

[4] Cohen, Billie, "In the Cottage, Yet Industrious," New York Times, April 16, 2009.

[5] "Gathering Clouds," The Economist, March 19, 2009.

The Akamai Difference

Akamai® provides market-leading managed services for powering rich media, dynamic transactions, and enterprise applications online. Having pioneered the content delivery market one decade ago, Akamai's services have been adopted by the world's most recognized brands across diverse industries. The alternative to centralized Web infrastructure, Akamai's global network of tens of thousands of distributed servers provides the scale, reliability, insight and performance for businesses to succeed online. Akamai has transformed the Internet into a more viable place to inform, entertain, interact, and collaborate. To experience The Akamai Difference, visit www.akamai.com.

Akamai Technologies, Inc.

U.S. Headquarters
8 Cambridge Center
Cambridge, MA 02142
Tel 617.444.3000
Fax 617.444.3001
U.S. toll-free 877.4AKAMAI
(877.425.2624)

International Offices
Unterfoehring, Germany
Paris, France
Milan, Italy
London, England
Madrid, Spain
Stockholm, Sweden
Bangalore, India
Sydney, Australia
Beijing, China
Tokyo, Japan
Seoul, Korea
Singapore

www.akamai.com