Best Practices for Super-Charging your IT Strategy

Executive Summary

In order to deliver business value, IT organizations are increasing their adoption of new technology platforms, such as cloud computing and mobility while also implementing big data projects, rolling out new business applications and modernizing existing applications and infrastructure platforms. This white paper delves into best practices for super-charging your IT strategy to ensure the technical and business success of your IT initiatives.

Introduction

Driven by the need to provide measurable business value, the IT environment continues to evolve, usually building on what has gone before and sometimes rediscovering techniques that were once widely used. When the Internet first came to life as a research tool around 1970, the most common IT environment featured dumb terminals accessing mainframe computers over analog private lines. Driven by the expensive nature of mainframe computers, timesharing services, whereby companies rented time on a mainframe became quite popular, as was the use of operating systems that virtualized mainframe computers.

In the early to mid 1980s, the client/server model of computing became popular. In the initial deployments, the client devices were desktop PCs and the network that connected the clients and servers was built on a technology such as Frame Relay. Once the Internet started to be used for commercial purposes in the early to mid 1990s, client/server applications began to use the Internet for wide area connectivity. As the Internet became the dominant network service, it enabled the development of new forms of applications such as Web-based applications in which an application is accessed over the Internet and the user interface is a browser that resides on a PC. In 1995 the seeds of the e-commerce revolution were planted as in that year both Amazon and eBay were founded.

Starting in the mid 2000s, IT organizations began to virtualize their data center servers in a manner similar to how mainframe computers were once virtualized. Soon other forms of virtualization, such as desktop virtualization, became popular and users began to access applications using PCs, laptops and devices that resembled the dumb terminals of the mainframe era. Also during this time period, a new model of computing emerged, called cloud computing, some forms of which closely resembled the time sharing services of the mainframe era. Cloud offerings constituted large, highly virtualized data centers where the infrastructure is shared amongst hundreds or thousands of customers, known as tenants. End-users typically access these public cloud offerings over the Internet using a wide range of devices including a growing set of mobile devices; i.e., phones, tablets and phablets.

The vast majority of IT organizations are currently in the process of implementing a number of key initiatives that once again enable them to provide measurable business value. In addition to the increasing adoption of cloud computing and mobility, these initiatives

include big data projects as well as the deployment of new applications and the modernization of existing applications, infrastructure platforms, and private network equipment. Like the previously discussed initiatives, most IT organizations rely heavily on the Internet when delivering these new IT initiatives. Unfortunately, as discussed in this white paper, the use of the Internet creates significant performance, reliability, and security issues that IT organizations must overcome in order to ensure that their key IT initiatives yield the anticipated business value.

Key IT Initiatives

Below is a description of some of the key initiatives that IT organizations are in the process of implementing.

Mobility

The 2013 Application and Service Delivery Handbook[1] included the results of a survey in which the survey respondents were asked "What percentage of your organization's employees use a mobile device at some time during the day to access business related data and applications either from within a company facility or from an external site?" Their responses are shown in Table 1

	0%	1% to 9%	10% to 24%	25% to 49%	50% to 74%	75% to 99%	100%
Company Facility	6%	14%	26%	19%	22%	10%	4%
External Site	2%	23%	20%	20%	14%	15%	6%

Table 1: Amount of Mobile Access

One conclusion that can be drawn from the third row of Table 1 is that the vast majority (98%) of employees use a mobile device daily to access business related data and applications from an external site. This highlights the dual reality that today's workers are more likely than ever before to work outside of a corporate HQ office and that in order to be productive, today's workers must be able to access their business applications from wherever they are located.

One of the key challenges associated with supporting mobile workers is the fact that a large and growing number of companies have adopted BYOD policies. As a result, mobile workers are increasingly using devices such as a smartphone, a tablet or a phablet to access business critical applications and data. An unfortunate fact of life is that in most cases the adoption of BYOD has meant that IT organizations have lost control of the end user's device. In spite of losing that control, IT organizations are still expected to effectively

[1] http://www.webtorials.com/main/resource/papers/webtorials/2013-App-Serv-Handbook/2013_Application_and_Service_Delivery_Handbook-Complete.pdf

render information to those devices independent of complicating factors. These factors, each of which can have an affect on the end user experience, include the choice of operating system or browser; the screen size and resolution; or the type of network connection that a user has on their mobile device. An additional concern is that this new generation of end user devices is subject to the same malware and network intrusion attacks as are PCs, but they typically lack mature, robust products for malware protection (e.g. anti-virus software) and network intrusion protection; e.g., a personal firewall. Unfortunately, most IT organizations have responded to the growth in mobile usage by implementing somewhat limited initiatives such as an initiative that provides e-mail to mobile workers and that makes a limited set of existing applications and data accessible from mobile devices.

IT Best Practice: A best practice relative to supporting mobile employees is to develop a mobility strategy based on rethinking how employees will leverage a growing set of devices to share and consume information from anywhere, at any time. IT organizations need to implement solutions that address mobile device delivery holistically, and that can dynamically adapt to ensure end-users are able to have a positive user-experience accessing the business applications they rely on to do their jobs, regardless of the device or connection type they may be using.

Cloud Computing

There are three primary classes of cloud computing solutions: public, private and hybrid. Enterprises of all sizes are increasingly embracing these solutions because they provide business users with increased agility as well as self-service, on-demand access to a much richer range of services and applications at a lower cost and provide much quicker time-to-value than was traditionally possible.

The initial interest that enterprises had in cloud computing was focused largely on public cloud services, such as those that can be acquired from an Infrastructure-as-a-Service (IaaS) provider such as Rackspace or from a Software as a Service (SaaS) provider such as Salesforce.com. The second wave of interest that enterprises had in cloud computing was focused on IT organizations offering services and applications to their users based the IT organization adopting many of the technologies and techniques that were implemented by the providers of public cloud services; e.g., virtualization, automation, pay-as-you-go. This is typically referred to as private cloud computing. The phrase hybrid cloud computing refers to a solution which is a combination of a public cloud and private cloud solution. The hybrid cloud model offers organizations the ability to leverage the most appropriate computing environment for the given situation, and take advantage of the economies of scale offered by large public cloud providers as needed, while still retaining control over their applications that they run in a private on-premise cloud environment.

While there are many benefits associated with cloud computing, there are also some significant challenges, most of which are associated with public and hybrid cloud solutions. One of those challenges is that when an IT organization acquires an IaaS or an SaaS

solution, it no longer has control of the data center that houses the solution. When combined with BYOD, that means that the IT organization no longer has control of either end point, yet, as previously noted, the organization is still held responsible for the overall user experience. Because the IT organization no longer controls the data center that houses an application or service, the IT organization can no longer guarantee the reliability of the data center, the security of the data center, nor can it necessarily instrument the cloud provider's data center in such a way as to provide the same level of visibility that the organization is accustomed to having. In addition, the vast majority of public cloud services are accessed over the Internet. As is explained later in this white paper, the Internet has limitations that result in performance problems for any application or service that is accessed over the Internet.

IT Best Practice: Cloud Balancing is an example of an emerging best practice relative to the use of cloud computing. The phrase cloud balancing refers to the ability to intelligently load balance user traffic and failover across multiple data centers based on myriad criteria such as cost, geographic location of the user, and real-time performance characteristics of each data center. As shown in Figure 1, cloud balancing involves any number of public cloud and/or corporate data centers.

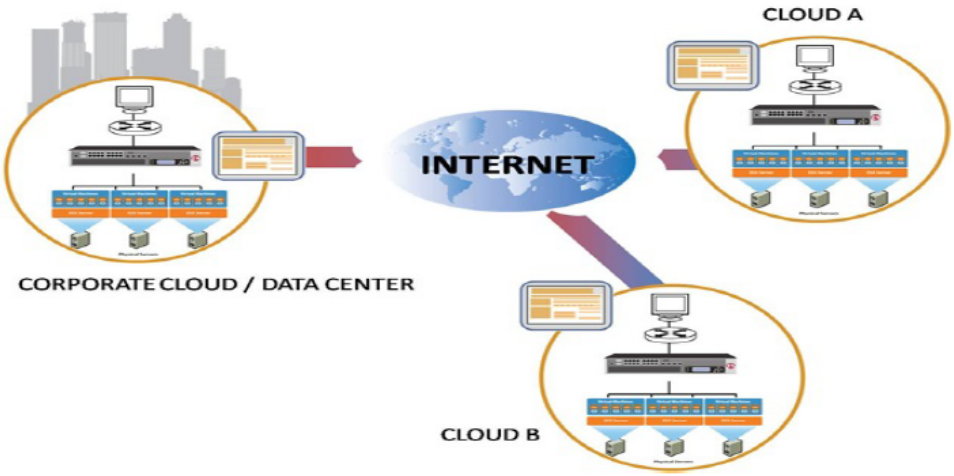

Figure 1: Cloud Balancing

One of the key advantages of cloud balancing is that if it is done correctly, cloud balancing maximizes the user experience by ensuring that users have 100% availability of the applications they are requesting, and that they have the best possible user-experience by ensuring they are accessing the most appropriate computing environment based on their unique device and network charactaristics. Another key advantage of cloud balancing is that it enables companies to build an infrastructure to support the average traffic load and to use public cloud services only to handle peak loads, seamlessly and automatically.

Big Data

Of the four initiatives discussed in this white paper, Big Data is the one that has emerged most recently. The phrase big data refers to working with a collection of data sets that is so large and complex that it is difficult, if not impossible to process them using traditional techniques. There are many drivers of the increasing size and complexity of data sets. One of these drivers is that unlike traditional companies, many companies that have come to life in the last five to ten years such as Google, Facebook and LinkedIn were built around the concept of analyzing data sets that were huge in size and growing. Another driver is that in addition to the structured data sets that have been the focus of business analytics for many years, companies are now analyzing new types of data including voice and video, as well as data gathered by a growing number of remote sensors.

There is, however, no hard and firm definition relative to how large and complex the data sets have to be in order to be considered to be big data. For some companies, facing hundreds of gigabytes of data for the first time may constitute big data, whereas for others, it may take tens or hundreds of terabytes of data to trigger a similar reaction. For some organizations, big data may be defined by the variety of data types that they are combining into a central analytics environment, or by the speed at which the organization aims to integrate and analyze their disparate data sets. Supporting massive scale, however, is just one of the challenges associated with big data. Another challenge is that increasingly big data is being applied in real time situations that that can only be supported by a system that exhibits end-to-end low, predictable delay, and that can address the inherent challenges associated with integrating datasets that are connected to the centralized analytics environment by a myriad of network connections.

IT Best Practice: To respond to these challenges, a best practice relative to the implementation of big data is for IT organizations to be able to accelerate the integration, and to be able to secure the delivery of myriad data sources into a centralized repository where sophisticated analytics are applied to the data. An example of a best practice in leveraging big data for business advantage is a company such as a hotel chain that runs a call center and uses big data analytics as a way to better predict customer dissatisfaction by analyzing the phone calls it receives. Another example of a best practice in leveraging big data for business advantage is a financial services firm that implements the ability to do monitoring of daily transactions, market feeds and customer service records to benefit in real time from fluctuations in the market and to develop new business models and services.

New Application Deployment and Application Modernization

In order to continue to provide business value, IT organizations are making investments that are large and growing both in the new applications that they implement themselves as well as in the applications that they acquire from a Software-as-a-Service (SaaS) provider. One measure of that investment is a recent analyst report[2] that indicated that the global SaaS market will grow from US $14.3 billion in 2012 to US $16.7 billion in 2013 and US $21.3 billion in 2015.

In addition to investing heavily in deploying new applications, IT organizations are also investing in modernizing their existing application postfolios. The goals of application modernization are to:

- Reduce cost by retiring legacy appliction infrastructure and network equipment;
- Create new business value from existing applications;
- Be able to deploy new applications and websites instantly to anyone, anywhere.

One way that companies modernize their applications is to re-write them in order to move away from a legacy programming language or database. However, similar to the situation with big data, there is not a hard and firm definition of what constitutes a legacy programming language or database. For example, within some companies the phrase legacy programming language refers to a language such as COBOL whereas in other companies, it refers to languages such as Java.

Another way that IT organizations modernize their applications is by implementing Web-based user interfaces and utilizing Web-specific protocols such as HTTP. While HTTP is not a chatty protocol[3], it is used to download web pages. It is common for a Web page to have fifty or more objects. Hence, although HTTP is not chatty, downloading a web page may require hundreds of round trips. In addition, it is also common for a Web page to rely heavily on dynamic content which must be accessed uniquely by each end-user from the origin datacenter. Having end users access dynamic content also results in requiring multiple round trips in order for the dynamic content to be transferred to the end-users. Due to the large, variable delay associated with the Internet, requiring multiple round trips to download a Web page results in significant negative impact to the overall end-user experience and to user adoption of applications.

A third way that IT organizations modernize their applications is to leverage public cloud services. In one scenario, the IT organization moves the application to an IaaS provider after potentially making some modifications to the code. In another scenario, the IT organization decides that the existing application is at end of life, retires the application and acquires equivalent or enhanced functionality from a SaaS provider.

[2] http://saasmarkets.com/report-global-saas-market-growing-driven-by-crm-mobile-deployments/
[3] A chatty protocol requires tens or hundreds of round trips to complete a single transaction.

In most cases, the opportunity to reduce IT cost and complexity by modernizing an organization's application portfolio is compelling. However, these modernization efforts typically drive more of the organization's application traffic off of a legacy private network infrastructure and onto the public Internet. This creates significant challenges for many organizations, because as described in a subsequent sub-section of this white paper the Internet is not a business-ready platform.

IT Best Practice: A best practice relative to the modernization of applications is to do a thorough risk/reward analysis prior to starting the modernization project. Part of the analysis is to quantify the specific benefits that the project will produce. The other key part of the analysis is to identify the challenges that the project will likely create. An example of that is that if the project will result in the use of additional services from a public cloud provider, the analysis and strategy development should identify the associated performance, reliability, and security challenges, and how the organization plans to proactively address these challenges. The analysis should also identify ways to mitigate the risk associated with the increased use of the Internet as a core component of the organizations overall IT delivery platform.

Increasing Reliance on the Internet

The IT initiatives discussed in this white paper have accelerated the previously mentioned shift in the types of network services that IT organizations utilize. Ten to fifteen years ago, the dominant network service was a private WAN based on a technology such as Frame Relay, ATM or MPLS. Today, the Internet is the dominant networking service used for the delivery of IT services. The continually increasing use of the Internet was documented in an analyst report[4] that indicated that while 50% of IT organizations anticipate increasing their use of MPLS based services, 84% of IT organizations anticipate increasing their use of the Internet.

Limitations of the Internet

The Internet provides a number of significant benefits including nearly universal availability and relatively low cost. Unfortunately, as explained below, the Internet also has a number of significant limitations. Those limitations are:

Performance

One of the major risks associated with leveraging the Internet for any of the IT initiatives discussed in this paper is the fact that the Internet is inherently slow and innefficient, thereby resulting in a negative impact to the performance of business applications being delivered over the Internet. The impact of poor application performance was quantified in a recent analyst report entitled The 2013 Application and Services Delivery Handbook[5].

[4] http://www.webtorials.com/content/2011/11/2011-cloud-networking-report.html
[5] http://www.webtorials.com/content/2013/06/2013-application-service-delivery-handbook.html

That report contained the results of a survey in which the survey respondents were given a set of outcomes that could result from poor application performance. They were asked to indicate the type of impact that typically occurs if one or more of their company's business critical applications are performing badly, and they were allowed to indicate multiple impacts. The impacts that were mentioned most often are shown in Table 2.

Table 2: Impact of Poor Application Performance	
Impact	Percentage
The Company Loses Revenue	62.0%
IT Teams Are Pulled Together	59.8%
Company Loses Customers	45.1%
CIO Gets Pressure from his/her Boss	45.1%
Harder for IT to get Funding	44.6%
CIO Gets Other Pressure	42.9%

One observation that can be drawn from Table 2 is that if a business critical application is performing poorly, it has a very significant business impact and it also has a very significant impact on the IT organization.

Reliability

The Internet is not a single network, but rather thousands of networks interconnected to appear as a single network. The individual networks that compose the Internet exchange information between each other that describes what IP address ranges they contain (a.k.a., routes). Within a single network - called a routing domain - a specialized networking protocol is used to communicate IP address ranges to all the routers within the individual network. Routing protocols within a network can detect a network link or an equipment failure and update the routing table on all routers within a few seconds when properly designed. For the exchange of information between networks - called inter-domain routing - a special routing protocol called the Border Gateway Protocol (BGP) is used. The size and complexity of the Internet as well as the inherent characteristics of BGP mean that a failure to a network link or piece of network equipment and the resulting routing path changes may take several minutes before all of the routing tables are updated. To put this into perspective, traditional voice circuits take milliseconds to reroute voice calls when a network link fails.

The impact of a network link failure and the time it takes for the Internet to update its routing tables and find an alternative path varies according to type of application involved. For a simple Web application, a brief outage may go unnoticed if users are not loading Web pages or performing some similar activity during the outage. For real-time applications like Voice-over-IP (VoIP) or IP Video, an outage of even a few seconds will cause interrupted calls and video sessions. In addition, there are two primary types of communication over the Internet: TCP and UDP. With TCP-based communication, lost packets are retransmitted until the connection times out. The good news is that if TCP is

used, packets don't tend to get lost. The bad news is that the associated packet retransmission negatively impacts application performance. With UDP-based communication, there isn't a built-in mechanism to retransmit lost data and UDP-based applications tend to fail rather than recover from brief outages.

Network Latency

As described by Moore's Law of Internet Latency[6], Internet latency is typically greater than the latency in a private WAN. That law references the business model used by the Internet and it states, "As long as Internet users do not pay for the absolute (integrated over time) amount of data bandwidth which they consume (bytes per month), Internet service quality (latency) will continue to be variable and often poor." Another cause of the relatively large amounts of delay that is associated with the Internet is that BGP does not consider delay or packet loss when making routing decisions.

Packet Loss

As noted, one of the reasons why there is more packet loss in the Internet than there is in a private WAN is due to the way that BGP makes routing decisions. Another part of the reason why there is relatively high packet loss in the Internet is because the only vendors that receive revenue from the Internet are the providers of the first and the last mile. One of the affects of that business model is that there tends to be availability and performance bottlenecks at the peering points which occur in the "middle mile".

Inherent Characteristics of TCP

Within TCP one or more missing packets are re-transmitted based on TCP's retransmission timeout parameter. If this parameter isn't set correctly it can cause either unnecessary delay or additional congestion. Another TCP parameter that impacts performance is the TCP slow start algorithm. This algorithm calls for the initial data transfer between two communicating devices to be severely constrained and to increase very slowly. The slow start algorithm is also applied in those situations in which a packet is dropped.

Security

In spite of the fact that the Internet is now the dominant network service, there are a number of security vulnerabilities associated with the use of the Internet. The issues associated with the lack of Internet security were mentioned in a recent report entitled IBM X-Force Mid-Year Trend and Risk Report [7], which summarized the work the IBM X-Force team did in the first half of 2013 to discover, analyze, monitor and record security threats and vulnerabilities. According to that report, "The majority of [security] vulnerabilities that the X-Force team documents are those in Web application programs." The report also stated that "While nearly 93 percent of the Web contains ordinary content, every 20th Website shows pornography" and added that "Twenty-three percent of all malicious links hosted on the Internet are located on pornographic sites." The report also pointed out that in early 2013 Internet Explorer was the target of multiple sophisticated attacks and it

[6] http://www.tinyvital.com/Misc/Latency.htm
[7] http://www-03.ibm.com/security/xforce/

commented on the growth in security incidents associated with mobility and with the increased use of social media and concluded that "…for a variety of reasons, companies are struggling with a commitment to apply basic security fundamentals."

Visibility

In part because the Internet is a network of networks, it is difficult, if not impossible, for an IT organization to get detailed visibility into the end-to-end performance and user-experience of the applications they are delivering over the Internet. Some of the initiatives discussed in this white paper further complicate that situation. For example, most IT organizations won't be able to load management agents onto the mobile devices that are used by the company's employees and in virtually all instances, IT organizations won't be able to host their management software at a SaaS provider's site. The result is a total, or near total loss of visibility into the performance of any business critical application that is accessed over the Internet. This makes proactive management impossible and greatly increases the amount of time it takes to perform reactive network management.

Scalability

There are two key components to the issue of scalability and the Internet. One component is that the Internet doesn't provide any functionality to enable scalable compute solutions; e.g., the Internet doesn't provide intelligent load balancing that enables cloud balancing solutions. The second component is that while it is possible to overcome many of the performance, security and management challenges associated with the Internet by implementing premise based solutions, that approach doesn't scale due to the complexity and overhead associated with that approach. The deficiencies of that approach are highlighted in the previoulsy mentioned quote from IBM: "…for a variety of reasons, companies are struggling with a commitment to apply basic security fundamentals." The struggles that companies are currently having would be magnified by an approach that required multiple, premise-based solutions be implemented and managed.

The Akamai Solution

Akamai offers a service called Terra Alta, which allows organizations to overcome the limitations of the Internet previously discussed and to leverage the Internet as a standard platform for delivering their business applications from any data center to any user, on any device, anywhere in the world. At its core, Terra Alta is a solution made up of several advanced technologies that have been developed specifically to meet the demands of businesses that are running any or all aspects of their business over the Internet. The Terra Alta service runs on the Akamai Intelligent Platform (Figure 2), which is a global platform deployed in over 80 countries, spanning the most important networks within the Internet, and is a single network hop away from over 90% of Internet users. The Intelligent Platform is controlled by Akamai's software and it is constantly monitoring Internet conditions in real-time. Also, the globally distributed nature of the Terra Alta service provides a natural defense mechanism against web-based security attacks.

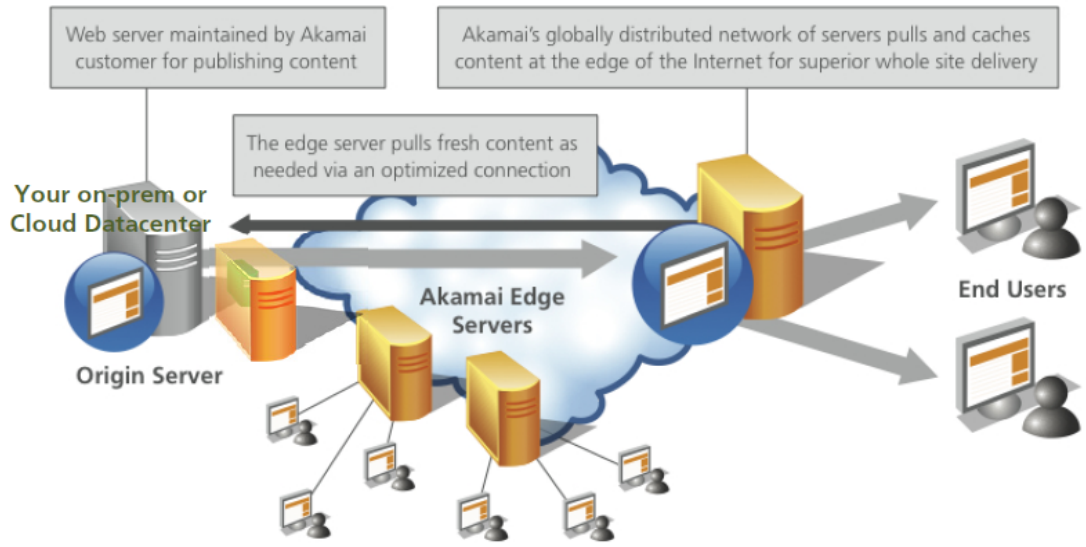

Figure 2: The Akamai Intelligent Platform

Below is a discussion of some of Terra Alta's functionality that allows organizations to overcome the previously discussed limitations of the Internet.

SureRoute

SureRoute chooses the end-to-end path that has the least delay and the least packet loss. This functionality can improve the performance and effective availability of the Internet itself by ensuring that viable routes are found to circumvent outages, peering innefficiencies, or congestion issues.

Cloud Balancing

Customer-specific policy rules are analyzed against a real-time evaluation of Internet traffic conditions to direct a user's request to the most appropriate data center which can reduce cost and/or improve performance. Cloud Balancing also allows end-users to maintain session persistence with a data center they have an active transaction established with, and failover between multiple data centers or virtual hosts in real-time if there is an application downtime event.

Cloud Monitor

This provides real-time access to granular levels of information about all application activity, including metrics on complete request/response cycles as well as origin response times. This information can be integrated with existing reporting and analytics tools.

Akamai Instant

This designates the most likely next pages to be visited by users and starts the process of gathering content, making Web service calls, or doing database lookups. As a result, appropriate content is pre-fetched close to the users prior to the users requesting it, which reduces delay when the requests are made.

Enhanced Akamai Protocol
This protocol overcomes the negative performance impact of the previously mentioned TCP characteristics; i.e., the TCP retransmission timeout parameter and the TCP slow start algorithm.

Object Pre-fetching
This functionality minimizes the amount of time it takes for a browser/client to load and render an application, including the embedded objects.

Application Layer Optimization
Performance can be improved by utilizing application layer optimization techniques such as compression and caching at the edge Akamai server with the cache performing intelligent pre-fetching from the origin.

Content Offload
Static content can be offloaded out of the datacenter to caches in Akamai servers and through persistent, replicated in-cloud storage facilities. This reduces the amount of time it takes to access this content.

Front End Optimization (FEO)
FEO makes a web page download and render faster, by recognizing the device type that the end user is using to access the application and then making real-time optimizations to the delivery of that application for that specific device.

Real User Monitoring (RUM)
Collects and reports on application response time.

Akamai Best Practices

As noted, the vast majority of IT organizations are rolling out initiatives that are fundamentally changing how IT services are provided and leveraging the Internet in new and exciting ways. While this is happening, businesses are continuing to evolve. In many cases that evolution includes expanding geographically and when faced with the need to support geographic expansion, many IT organizations believe that before the business can expand, they need to set up new data centers around the globe. This approach adds significant cost and can dramatically elongate the amount of time it takes before the company is operating in the new markets.

Based in part on Akamai's ability to optimize the performance of Internet traffic and in part on Akamai's ability to effectively route traffic from one geographic region to another, or from one site to another, the use of Akamai's services enables the company to be operating in new markets very quickly. Customers can also utilize Akamai as an extension of their own IT architecture by leveraging Akamai's global platform's scale for infrastructure offload.

The preceding discussion about using Akamai services to reduce the time and the cost of moving into new markets demonstrates that when IT organizations are doing something new, such as implementing one of the initiatives discussed in this white paper, or supporting emerging business models, they should analyze a variety of service delivery models. While engaging Akamai at any stage in the lifecycle of an initiative adds value, a best practice is to engage Akamai early in the lifecycle, preferably in the planning stage, to leverage Akamai experts in helping to design and architect a strategy that positions your organization to leverage the Internet with confidence and ensure the success of your IT initiatives. It is also important to actively engage with Akamai throughout the lifecycle. For example, many IT organizations follow a "set it and forget it" approach to using a service from a third party. Given the dynamic nature of business and of applications, IT organizations will maximize the benefit of using Akamai services if they regularly use the reporting functionality provided by Akamai to make changes to the service.

As was previously discussed, companies are beginning to use several data centers, including public cloud providers such as Amazon, to deliver their overall application portfolio. A related best practice is to use Akamai to provide cloud balancing. The idea being that IT organizations are using multiple cloud providers, or multiple Amazon availability zones, for many reasons. Using Akamai, service requests are sent to a cloud provider based on myriad criteria including geographic proximity or performance information. Another best practice is to leverage Akamai's ability to support dynamically migrating sessions between clouds without dropping any sessions.

Whether they are operating in a traditional IT environment, or in an environment with mobile workers accessing public cloud services and applications, IT organizations are still held responsible for meeting a combination of implicit and explicit service level agreements (SLAs). One best practice relative to ensuring that IT organizations can meet their SLAs is to use Akamai's Terra Alta platform which offers a guaranteed SLA. Another best practice is to use the previously described functionality provided by Terra Alta in order to overcome the performance impacting limitations of the Internet and to accelerate content delivery.

In addition to using the previously described functionality provided by Terra Alta there are a number of best practices related to specific Akamai provided functionality or techniques that can enable IT organizations to better support initiatives such as the ones described in this white paper. Those best practices include:

Differentiate Dynamic vs. Personal Content

The application server generates dynamic content in order to respond to an end user request. In some cases, the response is valid for a period of time and should be cached and reused. In contrast, personalized content can only be served to satisfy a single request.

Consolidate Personalization into Groups
If database access is required for personalization, then a best practice is to aggregate the personalization into one area of the page if possible and treat the rest of the page as being cacheable.

Minimize Round Trips
One way to minimize round trips is to cache non-personalized content wherever possible. Another way is to realize that some pages can be personalized on the client side with java script based on simple cookie values.

Timout Idle Connections
Set timeouts so that idle connections do not consume resources.

Use Net Storage for Large files
Akamai's Net Storage service offers highly scalable, highly available, geographally distributed, mirrored storage regions that are optimized for performance.

Conclusion

The last several decades has seen a steady, predicatable evolution in terms of how IT services are provided. This evolution started in the mainframe era of the early 1970s and progressed through client server computing to Web based applications to the initial implementations of mobility, server virtualization and cloud computing. The key learnings from this long evolutionary period are:

1. IT organizations need to continually find new ways to add value;
2. Change is a constant and needs to be planned for;
3. The Internet is the dominant WAN service;
4. The Internet has limitations that negatively impact availability, performance, management and security.

The vast majority of IT organizations are currently in the process of providing business value by extending existing initiatives such as mobility and cloud computing. In addition, many IT organizations are driving other initiatives such as big data, the deployment of new applications and the modernization of existing applications. All of these initiatives increase the amount of traffic that transits the Internet.

However, as mentioned, one of the lessons that was learned over the last several decades is that the Internet has limitations that negatively impact availability, performance, management and security. The Akamai Terra Alta service, which is built on the Akamai Intelligent Platform, overcomes the limitations of the Internet and enables organizations to leverage the Internet as a standard platform for delivering all their business applications. The Terra Alta service achieves this goal by providing a range of functionality including the ability to choose the end-to-end path through the Internet that has the lowest delay, by

pre-fetching content before users request it and by offloading static content out of the data center to caches in Akamai servers and through persistent, replicated in-cloud storage facilities.

Some of the best practices associated with using Terra Alta apply to virtually all IT initiatives. For example, in order to both reduce the time to value and to eliminate unnecessary capital costs, a best practice is to engage Akamai early in the lifecycle of a project, preferably in the planning stage. However, given the dynamic nature of business and of applications, actively engaging with Akamai throughout the lifecycle of the project is also a best practice. Another best practice that applies to a wide range of IT initiatives is to use Akamai's SLA test objects to test performance and hence position IT organizations to ensure that they can meet their SLAs.

Some of the best practices associated with using Terra Alta relate to specific initiatives such as cloud computing. A best practice in the use of cloud computing is for companies to use multiple data centers, including public cloud providers such as Amazon, to deliver their overall application portfolio. A related best practice is to use Terra Alta to provide cloud balancing. Relative to cloud balancing, Terra Alta provides value by sending service requests to a cloud provider based on myriad criteria including geographic proximity or performance information.

More information can be found at: www.akamai.com