2009 EXECUTIVE REPORT

What Cloud Computing Means to You: Efficiency, Flexibility, Cost Savings

HOSTING.com
THE FIRST NAME IN HOSTING

ITBUSINESSEDGE

What Cloud Computing Means to You: Efficiency, Flexibility, Cost Savings

By Mike Stevens

Cloud computing is being heralded as the Next Big Thing in IT infrastructure, promising new levels of efficiency, flexibility and cost savings—particularly in the area of outsourced hosting, also known as infrastructure-as-a-service (IaaS). But because cloud computing is at the early stages of what Gartner calls the "hype cycle," there is widespread confusion about what the term actually means, not to mention questions about how this new technology can deliver practical business benefits.

The fact is, cloud computing has real value right now. This paper will explain how it works in the context of evolving technology trends and how, through emerging IT outsourcing models, cloud computing can function as an extension of a company's existing infrastructure, with enterprise-class features like redundancy, high availability and disaster recovery (fail-over) provided at a substantially lower cost than on-premises approaches.

One IT Web site recently collected 21 different definitions for cloud computing from various experts, but in the broadest sense the definition is simple. Cloud computing provides a remote service that users can access via the Internet. Under this broad definition, we are all familiar with cloud computing. Facebook, Flickr and the various Internet-based e-mail offerings such as Yahoo! Mail and Google's Gmail all store and process data remotely, and we take this for granted. Because of the close association of these familiar services with cloud computing, some analysts and trade-press editors are beginning to refer to all Internet-accessible services as cloud computing. This is both confusing and inaccurate. Such services more properly fall into the category of software-as-a-service (SaaS).

It is more accurate to look at cloud computing as a new approach to infrastructure—one that is a logical next step on a path to more efficient use of computing resources.

Clouds, Economics and Energy

The search for economic efficiency in IT infrastructure is a relatively new concern. Until recently, the only energy-related question that data-center managers worried about was whether enough of it would be available on a reliable basis. Three factors have changed this mindset:

- Rising energy costs
- A trend toward IT and Facilities Management sharing responsibility—and budget—for data-center energy usage
- Concern about the general ecological impact of IT, and its contribution to companies' carbon footprint in particular

Increasingly, data-center managers must meet not only SLAs but "performance per watt" goals as well. Two technologies have already emerged to meet this goal: virtualization and grid computing (also known as utility computing). Both embody principles that are central to cloud computing.

Virtualization: Key to Boosting Server Utilization

Server virtualization was one of the first and most important responses to demands for IT energy efficiency. It is specifically a response to the one-application-per-server mentality that has been prevalent in data centers for many years and that results in a situation where only 15 or 20 percent of their total computing capacity is in use at any given time. By enabling IT departments to run multiple applications on the same server, virtualization provides dramatic gains in server utilization. It is also an approach that non-technical senior managers can understand and may even demand (often without realizing that virtualization can be more difficult in practice than in theory).

Although virtualization clearly increases server resource utilization, it is not the ultimate answer to computing efficiency because it does not deal with the problem of usage spikes, such as the increased demand put on financial applications at the end of each financial quarter or the huge (but short-lived) spikes in e-commerce traffic following the airing of an infomercial. For example, if

application X exceeds the processing capability of the initially provisioned and available hardware during a peak transaction period, the application suffers a performance hit because virtualization cannot dynamically add additional hardware components and additional virtual machines on demand.

This is where technologies like VMware DRS (Distributed Resource Scheduler) come into play. VMware DRS can aggregate resources from multiple servers into pools of resources. By doing so, DRS can intelligently allocate available resources among virtual machines according to business needs. Highly demanding, short-lived applications will have access to all the necessary server resources when needed without isolating those resources specifically to those applications. When those highly demanding, short-lived applications are not demanding excessive resources, those resources can be dynamically reallocated to other virtual machines as needed.

Grid Computing vs. Usage Spikes

Another technology known as grid or utility computing was developed in the mid 1990s to address this issue. The central concept was to create a pool of geographically distributed computing resources that any application could draw on as necessary to execute a task—a pool large enough to easily absorb usage spikes by a single application. This pool or grid would provide computing power in the same way an interstate power grid provides electricity, with each "customer" (application) drawing as much power as necessary. Users would neither know nor care about the location of their data or the CPU that was processing that data.

It now appears that the grid vision of computing power as a utility was somewhat ahead of its time. A number of significant issues must be resolved before today's mission-critical applications can run on a grid, primarily related to the difficulties of simultaneously running one instance of an application on multiple physical servers. At present, the cost of resolving these issues and getting a grid to work is so high that it destroys the business case for grids. But by propagating the idea that users and the computing resources they use don't need to be in the same physical location, grid computing paved the way for a practical, technologically viable approach to computing based on shared, distributed resources that may be remote: cloud computing.

New Attitudes Toward Servers

In simple terms, cloud computing exploits two new attitudes toward servers:

- Servers can run more than one application (virtualization).
- Servers need not be located on the premises (distributed computing).

These new attitudes, combined with technical advances in CPU capacity and greater bandwidth availability, have led to this new mode of delivering computing resources. Cloud computing utilizes virtualization technology to maximize efficiency, but it avoids the complexity of grid computing issues by not attempting to parse out single workloads over multiple physical servers. In other words, cloud computing retains the one-application-per-machine approach, but it does so using virtual machines.

One of the best-suited business models for exploiting this new technology is infrastructure outsourcing, which is the second theme of this paper.

The Evolution of Infrastructure Outsourcing

The business case for infrastructure outsourcing has two components, and cloud computing strengthens both of them. The first is cost. Acquiring the requisite hardware to meet the requirements of a growing company is a capital expense that many companies would like to avoid. Furthermore, the acquisition-through-purchase model is

Figure A

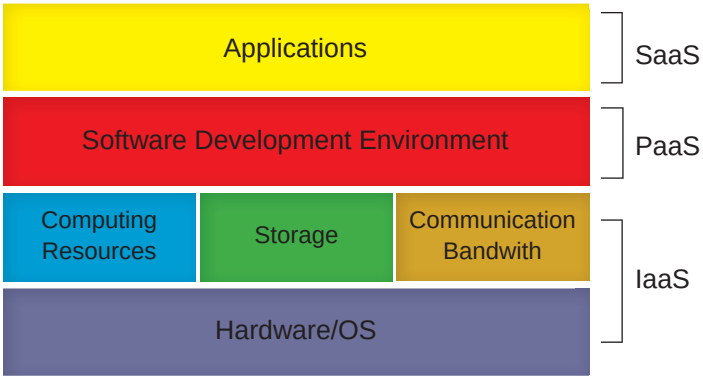

Currently, managed hosting, or infrastructure-as-a-service (IaaS) is one of the most important and fastest growing uses of cloud computing. Other uses include platform-as-a-service and software-as-a-service.

inherently wasteful. Servers and memory systems are almost always purchased to meet projected capacity needs, not current needs. As a result, a substantial percentage of acquired resources are always wasted until the company grows into them, which may take as long as a year.

In contrast, an outsourced model such as managed hosting offers a pay-as-you-grow model that enables companies to utilize only the capacity they need. In addition, the actual cost of providing these resources is lower. Hosting companies reduce their own costs through a combination of superior technology and economies of scale, and they can pass on a portion of their savings to customers.

Strategic Focus, Extreme Flexibility

The second argument for infrastructure outsourcing is that it enables IT departments to focus on strategic initiatives that add value and create competitive advantage, rather than getting mired in the everyday details of managing servers (plus the support technology and technicians required to keep those servers up and running).

Infrastructure outsourcing was able to deliver these benefits even in its early days, when servers were run on the old one-application-per-server model. As virtualization technology became available, efficiency increased dramatically. Cloud computing takes efficiency (and its cost benefits) to a whole new level and adds a benefit that virtualization alone cannot provide: extreme flexibility. A hosting company using cloud technology can literally double a customer's capacity in half an hour.

Security: Comparable to Physical Environments

When IT managers begin to consider infrastructure outsourcing, one of their major concerns is always security, particularly when a virtualized environment is involved. But as cloud computing has matured, the security issues have been resolved. For networks, unified threat management (UTM) technology combined with virtual domains (VDOMs) can provide the same isolation and hardened firewall/IPS security policies as traditional physical environments. Data security can be ensured when each customer has its own isolated virtual SAN disk, with no sharing between customers.

Conclusion

The success of companies offering infrastructure-as-a-service (IaaS) based on cloud computing is a strong indication that cloud computing will become increasingly more important over time. Meanwhile, as the debate over the exact definition of cloud computing continues in academic circles and technical chat rooms, the reality of cloud computing is giving companies cost efficiencies and flexibility that have never before been possible.

Candystand.com: a Case Study

The experience of an online gaming company is a perfect example of how a managed hosting company can use cloud computing to make it easier for its clients to better serve the needs of their customers.

Candystand.com is the Web's premier destination for free-to-play, high-quality online games suitable for the entire family, serving approximately 5 million unique visitors per month. Recently, when the company's ownership structure changed, it needed to migrate to a new managed hosting provider to replace the current provider, which would no longer be available.

As a company that depends on online advertising as its source of revenue, literally every minute of down time affects Candystand.com's bottom line. In addition, Candystand.com enjoys substantial loyalty and user confidence among its players, and any hint of unreliability could seriously affect that valuable business asset. Because of these two factors, high availability was the company's most important criterion in choosing a new hosting company. In addition, Candystand.com has very aggressive growth targets, which made scalability a second important criterion. A third criterion, of course, was cost.

Cloud technology provided by Hosting.com was able to meet all these criteria, and its ability to do so can be specifically attributed to its use of cloud computing. The virtualized environment facilitates the load balancing that is absolutely necessary for a site like Candystand.com to ensure high availability, and no technology provides faster scalability than cloud computing. Perhaps most important, Hosting.com was able to meet these criteria at a cost that was between 30 and 40 percent lower than competitors using conventional "pre-cloud" technology.

This IT Business Edge Executive Report Sponsored by

Hosting.com provides colocation, cloud computing, dedicated and managed hosting, and business continuance services to a global customer base demanding a high level of security, reliability and responsiveness. Hosting.com monitors, manages, and enhances the Web-based platforms for companies whose Web presence is crucial and high availability mandatory.

Hosting.com introduced its Cloud Computing Hosting Solution, CloudNine, in the fall of 2008. The dynamic CloudNine infrastructure recognizes a business's evolving needs and reacts instantly to provide the most effective and secure solutions.

CloudNine is an innovative solution providing flexibility and scalability in a controlled and secure environment. Key components include: root access to private and secure virtual servers, a self-serve Web-based Customer Portal, hosting in Hosting.com's fully redundant and geographically-dispersed data centers, and guaranteed 100% uptime and 24x7 support.

Hosting.com's recent Cloud Computing Trends Report and interactions with consumers show that businesses are likely to deploy cloud computing environments due to the reduced costs, scalability, improved service-level guarantees, flexibility, and overall performance associated with cloud-based solutions. Due to these advantages, a number of companies have moved their entire infrastructure to a cloud platform, while others are testing the platform with a portion of their hosted infrastructure.

CloudNine's unique pricing model and standards-based infrastructure provide companies with the ability to do both. Specifically, CloudNine provides value to businesses that:

- Are utilizing the Web to grow to offset the recession
- Have fluctuating needs based on business, technical or seasonal requirements
- Require on-demand growth without the worry of managing hardware
- Seek budget-friendly disaster recovery

IT Business Edge

IT Business Edge delivers the information, analysis and context business technology decision makers need to maximize returns on IT investments and align IT initiatives with business objectives. As a Technology Intelligence Agent, IT Business Edge provides content different from that of a traditional IT publisher, news service or analyst firm. Our editors monitor all these sources — plus many others — for critical IT information that they translate into actionable advice for high-level IT and business managers. Subscribers access our practical content and useful decision-making tools through a rich Web site, targeted e-mail newsletters and varied RSS feeds. All these outlets feature our business-focused blogs, exclusive interviews with field experts and industry insiders, plus our database of more than 20,000 abstracts summarizing content from 3,000-plus sources. Visit our Web site at: http://www.itbusinessedge.com.

About the Author

Mike Stevens began his career as a technical writer in semiconductor manufacturing and then switched to marketing. At his own Silicon Valley–based agency, he worked with an impressive list of clients, including HP, EMC, Fujitsu and Microsoft. His primary focus for the last seven years has been enterprise software.