OSBR.ca

The Open Source Business Resource

APRIL 2010

APRIL 2010

PUBLISHER:
The Open Source Business Resource is a monthly publication of the Talent First Network. Archives are available at the website:
http://www.osbr.ca

EDITOR:
Dru Lavigne
dru@osbr.ca

ISSN:
1913-6102

ADVISORY BOARD:
Chris Aniszczyk
Tony Bailetti
Leslie Hawthorn
Chris Hobbs
Rikki Kite
Thomas Kunz
Michael Weiss

The editorial theme for this issue of the OSBR is Cloud Services. The articles in this issue help to clarify the meaning and value of the cloud, dispel myths around cloud services, and demonstrate the role open source can play as enterprises transition to the cloud computing model.

As always, we encourage readers to share articles of interest with their colleagues, and to provide their comments either online or directly to the authors.

The editorial theme for the upcoming May issue of the OSBR is Communications Enablement and the guest editor will be Chris McPhee. Submissions are due by April 20--contact the Editor if you are interested in a submission.

Dru Lavigne

Editor-in-Chief

Dru Lavigne is a technical writer and IT consultant who has been active with open source communities since the mid-1990s. She writes regularly for BSD Magazine and is the author of the books BSD Hacks, The Best of FreeBSD Basics, and the Definitive Guide to PC-BSD.

Cloud computing may be the biggest game changer within the enterprise since the adoption of the Internet in the 1990s and the personal computer in the 1980s. The economic climate, coupled with the rate at which technology is changing, is forcing companies to reduce information technology (IT) expenditures, increase productivity, and build more collaborative solutions as opposed to building everything internally. Cloud computing allows companies to outsource functions that are not core to their business or have become a commodity. Much of the technology driving cloud computing services is open source software. The LAMP stack (http://en.wikipedia.org/wiki/LAMP_(software_bundle)) has become widely adopted as the standard engine running much of the cloud services. With the exception of Microsoft's Azure cloud platform (http://microsoft.com/windowsazure), most cloud service providers have embraced open source software, allowing them to drive costs down while providing reliable services for their customers.

Pay-as-you-go is the new economic model for IT as we enter a new decade. Gone will be the days of making large purchases of commercial software with huge maintenance costs. In this new world, enterprises will consume only the services they need, only when they need them, and will pay for the resources required to scale up and down as necessary. This paradigm shift should force IT executives to focus more on building business functionality and less on managing and maintaining infrastructure and commodity applications. Open source software will play a huge role in making the shift to cloud computing economically feasible. At the same time, commercial software companies are racing to the cloud and are struggling to replace their expensive software licensing models with a pay-as-you-go model in order to make them an attractive alternative to open source software in the cloud.

It should be interesting to see how this plays out over the next few years. The articles that follow provide insights into what cloud computing is and how it is changing the IT landscape.

David Linthicum is the CTO of Bick Group, an author, and thought leader on cloud computing. Dave's article starts with a discussion on the causes of the shift to cloud computing. He then clarifies the meaning of the broad term of cloud computing by identifying its key characteristics, different deployment models, and various components. Dave finishes with an important discussion on the value and economics of the cloud.

Fred Waldner is an independent consultant and a Certified Professional Project Manager. His article discusses how the shift to cloud computing mirrors the same challenges that we witnessed with the shift to enterprise adoption of the Internet in the 1990s. He identifies some of the adoption obstacles that must be overcome.

Daniel Crenna is an award winning developer, author, and active entrepreneur. Daniel articulates the value of open source in hosted solutions as companies adopt cloud computing. He examines the reasons why companies use established open source software or give away their proprietary code as the seeds of new open source projects.

Ronald Schmelzer is a Managing Partner at ZapThink LLC. His article addresses the debate about private clouds. There are many companies that are afraid of deploying critical applications and/or data in a public cloud environment. Instead, many are contemplating building their own internal cloud which goes against much of the value proposition of cloud computing and keeps their dependency on expensive commercial software. Ron argues that the concept of private clouds is a marketing term used by vendors rebranding their products.

Tom Lounibos, CEO for SOASTA, discusses how the pay-as-you-go cloud computing model allows SOASTA to simulate performance tests with loads that were simply not feasible in the on-premise world. Low cost, on-demand cloud computing services allow companies to deploy more reliable and scalable solutions because they can now afford to test peak loads and fix bottlenecks before they happen in the production environment.

John Crupi is CTO and Chris Warner is Vice President of Marketing at JackBe Corporation. Previous articles discuss how the lower levels of the stack are becoming a commodity and that IT's focus should be on business functionality. John and Chris provide some examples of how organizations are leveraging mashups to aggregate data from various locations into composite applications.

Mike Kavis

Guest Editor

Mike Kavis is CTO of M-Dot Network and Vice President and Director of Social Technologies for the Center for the Advancement of the Enterprise Architecture Profession (CAEAP). Mike received his BS in Computer Science from Rochester Institute of Technology (RIT) and his MBA and MSIT from Colorado Tech.

"*The ability to apply cloud economics to core enterprise applications provides new ways for CFOs and CIOs to optimize and boost the cost efficiency of IT service delivery.*"

John C. Stame
http://jcstame.wordpress.com/2009/06/29/the-business-value-of-cloud-computing

The movement to cloud computing is the disruptive change that information technology (IT) departments will soon face as cloud computing begins to have an effect on the modern enterprise. IT managers must learn how to give as well as take information in this new, shareable environment, while still protecting their company's interests. Innovative companies will take advantage of these new resources, such as cloud computing, and reinvent themselves as unstoppable forces in their markets. Those who don't take advantage of this revolution will become quickly outdated, perhaps out of business.

This article will discuss what the cloud is, and how it can potentially have value within the enterprise. Rather than a hype-driven view of cloud computing, this article provides the guidance needed to begin defining the value of cloud computing.

Shift to Cloud Computing

The movement and direction are clear. Based upon current trends, IDC predicts a compound growth rate of 20% per annum for cloud computing, set against the overall software market which is only growing at around 6% per annum (http://www.idc.com/getdoc.jsp?containerId=prUS21480708). This leaves IDC in no doubt that there is a fundamental shift toward cloud computing as a delivery mechanism.

Moreover, a new Gartner report shows global revenue from cloud computing will top $150 billion by 2013 (http://www.itbusinessedge.com/cm/community/news/inf/blog/gartner-cloud-computing-poised-for-considerable-growth/?cs=31403). This includes the shift from on-premise to cloud-based providers, and thus the need for planning and architecture to make this happen.

What's causing this shift? There are five primary drivers:

1. Purchasers believe that the current cost of traditional enterprise software is disproportionate to the value that it creates. The traditional approach to paying software licenses is getting some push-back these days in light of open source, and emerging cloud computing providers such as Amazon.

2. In these budget-conscious times, there is intense pressure to reduce the cost of acquisition and maintenance of software solutions, where the on-going support and maintenance of solutions can often be four times the original capital cost.

3. Organizations are striving to reduce risk, and want a far more tangible relationship between software's benefit and its cost. This has also been apparent in the open source movement of the last ten years, and cloud computing is really an extension of that trend.

4. The drive for reduced risk demands a much greater predictability of the running costs of the organization's software solutions.

5. The value of solutions is no longer determined by the functionality available (in fact, most organizations only use a small subset of the functions available in their software products), but by the experiences of the users in the way that they use and interact with the solutions.

There is a huge resource being created on the Web by cloud computing providers. Not taking advantage of this resource will affect enterprises, much like those who ignored the rise of the Web in the early 1990s soon found themselves playing catch-up.

Defining Cloud Computing

While cloud computing is widely defined, it needs a standard definition that most can agree on. We will use the definition (http://csrc.nist.gov/groups/SNS/cloud-computing/index.html) provided by the National Institute of Standards and Technology, Information Technology Laboratory (NIST). Cloud computing is a pay-per-use model for enabling available, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is comprised of five key characteristics, three delivery models, and four deployment models.

The five key characteristics are:

1. On-demand self-service: a consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed, without requiring human interaction with each service provider.

2. Ubiquitous network access: capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms such as mobile phones or laptops.

3. Location independent resource pooling: the provider's computing resources are pooled to serve all consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. The customer generally has no control or knowledge over the exact location of the provided resources. Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

4. Rapid elasticity: capabilities can be rapidly and elastically provisioned to quickly scale up and rapidly released to quickly scale down. To the consumer, the capabilities available for rent often appear to be infinite and can be purchased in any quantity at any time.

5. Pay per use: capabilities are charged using a metered, fee-for-service, or advertising based billing model to promote optimization of resource use. Examples are measuring the storage, bandwidth, and computing resources consumed and charging for the number of active user accounts per month. Clouds within an organization accrue cost between business units and may or may not use actual currency.

Cloud software takes full advantage of the cloud paradigm by being service oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability.

In addition, NIST also defines three delivery models as infrastructure-as-a-service, software-as-a-service, and platform-as-a-service. We extend these a bit in this section to define delivery models in more detail.

All cloud computing approaches are not the same. There are several differing deployment models that all fall under the cloud computing umbrella. These include:

1. Private cloud: the cloud infrastructure is owned or leased by a single organization and is operated solely for that organization.

2. Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns.

3. Public cloud: the cloud infrastructure is owned by an organization selling cloud services to the general public or to a large industry group.

4. Hybrid cloud: the cloud infrastructure is a composition of two or more clouds that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability.

Each deployment model instance has one of two types: internal or external. Internal clouds reside within an organization's network security perimeter and external clouds reside outside that perimeter.

Many find that private clouds are a better solution for their situation because they leverage the benefits of cloud computing within their firewall, typically using open source solutions such as Xen (http://xen.org), Eucalyptus (http://eucalyptus.com), and Ubuntu Enterprise (http://ubuntu.com/cloud/private). Or, they may choose to leverage a mixture of public and private clouds, or a hybrid cloud. Finally, some may create semi-private or community clouds that are public clouds leveraged only by a closed group of companies or government agencies.

The Components of Cloud Computing

As cloud computing emerges, there is much discussion about how to define cloud computing as a computing model. Maturity models have been published and debated, and providers clearly have a model for their own products.

In an attempt to better define cloud computing, we created a stack that considers each component of cloud computing and how they interact. As seen in Figure 1, the components of cloud computing make up a wide range of services that can be leveraged over the Web through a subscription-type arrangement. Most services that can be leveraged from a data center can now be leveraged from the cloud.

While many debate the components, there are eleven major categories or patterns of cloud computing technology:

1. Storage as a Service (SaaS)

2. Database as a Service (DaaS)

3. Information as a Service

4. Process as a Service

5. Application as a Service (AaaS)

6. Platform as a Service (PaaS)

7. Integration as a Service

8. Security as a Service

9. Management/Governance as a Service (MaaS/GaaS)

10. Testing as a Service (TaaS)

11. Infrastructure as a Service (IaaS)

One should note that these categories are rapidly evolving. It is helpful to separate the various offerings considering that the term "cloud computing" is much too widely defined. Let's examine each component of cloud computing.

Figure 1: Cloud Computing Components

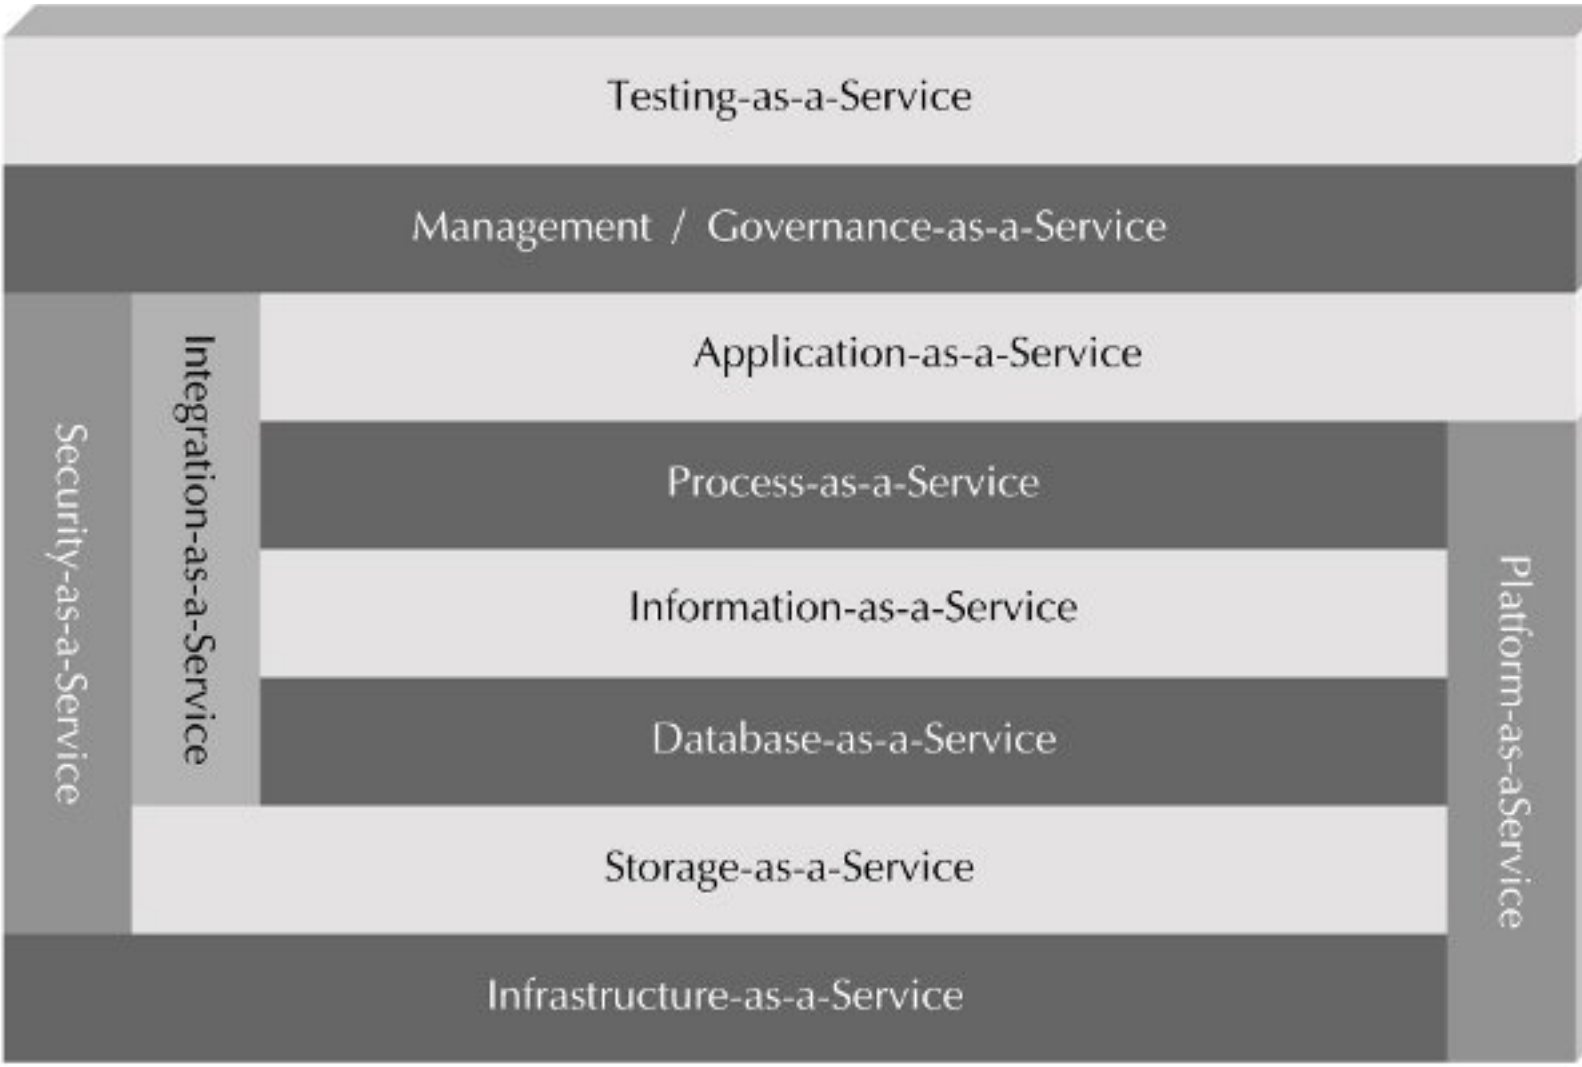

1. Storage as a Service: also known as disk space on-demand. Provides the ability to leverage storage that physically exists at a remote site, but is logically a local storage resource to any application that requires storage. This is the most primitive component of cloud computing and is leveraged by most of the other cloud computing components.

2. Database as a Service: provides the ability to leverage the services of a remotely hosted database, sharing it with other users, and having it logically function as if the database were local. Different models are offered by different providers, but the power is to leverage database technology that would typically cost thousands of dollars in hardware and software licenses.

3. Information as a Service: refers to the ability to consume any type of information, remotely hosted, through a well-defined interface such as an application programming interface (API). Examples include stock price information, address validation, and credit reporting.

4. Process as a Service: refers to a remote resource that is able to bind many resources together, such as services and data, either hosted within the same cloud computing resource or remote, to create business processes. A business process is a meta-application that spans systems, leveraging key services and information that are combined into a sequence to form a process. These processes are typically easier to change than applications,

and thus provide agility to those who leverage these process engines that are delivered on-demand.

5. Application as a Service: refers to any application that is delivered over the Web to an end user, typically leveraging the application through a browser. While many associate AaaS with enterprise applications such as Salesforce SFA (http://salesforce.com/crm/sales-force-automation), AaaS includes office automation applications such as Google Docs, Gmail, and Google Calendar.

6. Platform as a Service: refers to a complete platform, including application development, interface development, database development, storage, and testing, delivered through a remotely hosted platform to subscribers. Based upon the traditional timesharing model, modern PaaS providers provide the ability to create enterprise-class applications for use locally or on-demand for a small subscription price or for free.

7. Integration as a Service: is the ability to deliver a complete integration stack from the cloud, including interfacing with applications, semantic mediation, flow control, and integration design. Integration as a service includes most of the features and functions found within traditional enterprise application integration (EAI, http://en.wikipedia.org/wiki/Enterprise_application_integration) technology, but delivered as a service.

8. Security as a Service: is the ability to deliver core security services remotely over the Internet. While the typical security services provided are rudimentary, more sophisticated services are becoming available such as identity management.

9. Management/Governance as a Service: refers to any on-demand service that provides the ability to manage one or more cloud services. These are typically simple things such as topology, resource utilization, virtualization, and uptime management. Governance systems are becoming available as well, such as the ability to enforce defined policies on data and services.

10. Testing as a Service: is the ability to test local or cloud-delivered systems using testing software and services that are remotely hosted. While a cloud service requires testing unto itself, TaaS systems have the ability to test other cloud applications, Web sites, and internal enterprise systems, and do not require a hardware or software footprint within the enterprise.

11. Infrastructure as a Service: is actually Datacenter as a Service (DaaS), or the ability to remotely access computing resources. In essence, you lease a physical server that is yours to do with as you will, and, for all practical purposes, is your data center, or at least part of a data center. The difference with this approach versus more mainstream cloud computing is that instead of using an interface and a metered service, you receive access to the entire machine and the software on that machine. In short, it's less packaged. DaaS and other more primitive infrastructure-type services are often defined as IaaS.

The True Value of the Cloud

One would think that cloud computing would always be more cost effective than on-premise computing. However, the value that cloud computing brings to an enterprise is really dependent upon many variables and the dynamics of the business. Like anything requiring change, the business case must be worked first to see the ultimate worth of this approach.

The forces at work here include strategic and tactical issues that should be analyzed, including the ability to shift risk to cloud computing providers, the ability to drive down operating costs, the ability to fix inefficiencies within the existing architecture, and the value all of that brings. The problem is that most enterprises do not analyze these business opportunities properly, and thus make many mistakes when building the larger business case. Moreover, enterprise architects are notoriously bad at creating business cases.

It is one thing to say that cloud computing makes life better because everybody says it does, versus saying cloud computing will make this company more than $50 million dollars over the next 5 years. The ability to make money will always outsell something cool and popular. Keep that in mind as you create the business case and sell cloud computing to your executive team.

While most think cloud computing is really about the ability to save operational costs, this may or may not be the case and depends upon the enterprise or problem domain. Indeed, there are many dimensions to consider, including the value of:

• ongoing operational cost reduction

• preserving capital

• upsizing on-demand

• downsizing on-demand

• shifting the risk

• agility

• reuse

It is a good idea to figure out the actual cost reductions that cloud computing can bring to your enterprise IT. The trick here is not only to figure out how much money can be saved, but how much it will cost to save that money.

Let's say that you've decided that cloud computing is the way to go for a major enterprise application requiring some customization, and that the costs of an on-premise instance versus a cloud computing instance are as follows:

On-Premise:

Hardware $100,000

Software $100,000 for license, $20,000 per year software maintenance

Maintenance $200,000 per year for people to work the system

Data Center $50,000 per year

Over 5 years, the cash expenditure will look like this:

Year 1 $450,000

Year 2 $270,000

Year 3 $270,000

Year 4 $270,000

Year 5 $270,000

For a total cost of $1,530,000 over 5 years.

Now the cloud computing options, using the same 5 year cost horizon.

Cloud Computing:

Hardware $0

Software $10,000 per month subscription

Maintenance $0

Data Center $0

Over 5 years, the cash expenditure will be $600,000, or $120,000 times 5 years. Figure 2 shows the comparison of the two calculations. This example assumes a simple application-on-demand scenario. The cost will be more complex with a cloud computing solution that leverages a variety of cloud providers that need to interoperate and for a solution requiring application development.

Now What?

The cost of cloud computing is not free. Watch out for the manage-by-magazine effect, where cloud computing is so much a part of popular thinking around computing that you're not objective. It's okay to say "No" to cloud computing if there is no clear business benefit.

It will be difficult to create business cases as the world of cloud computing continues to emerge. We just don't know enough about cloud computing's long-term operational value, or many of the pitfalls, since most of the public cloud computing providers have not been around for that long. Consider this as a learning process, and make sure to examine the current state-of-the-art while leveraging the technology. The technology will constantly change while the architecture should be relatively stable.

We have a clear opportunity to change the way we do IT to have a more positive influence on the business when leveraging cloud computing. No longer will IT be a drain on company resources, and many enterprises will find that the new found efficiencies of IT have huge value on the bottom line. After all, IT is there to serve business, not the other way around.

Portions of this article are excerpted from the book "Cloud Computing and SOA Convergence in Your Enterprise…A Step-by-Step Guide".

David S. Linthicum is CTO of Bick Group, a company that focuses on making enterprises work well with the emerging cloud computing space. He's the cloud computing blogger for InfoWorld, Intelligent Enterprise, and Cloud Computing Journal, as well as a frequent speaker at technology events. In addition, he produces "The Cloud Computing Podcast" which can be found on iTunes. You can follow Dave on Twitter at twitter.com/davidlinthicum.

Figure 2: Comparing the Cost of On-Premise with Cloud Computing

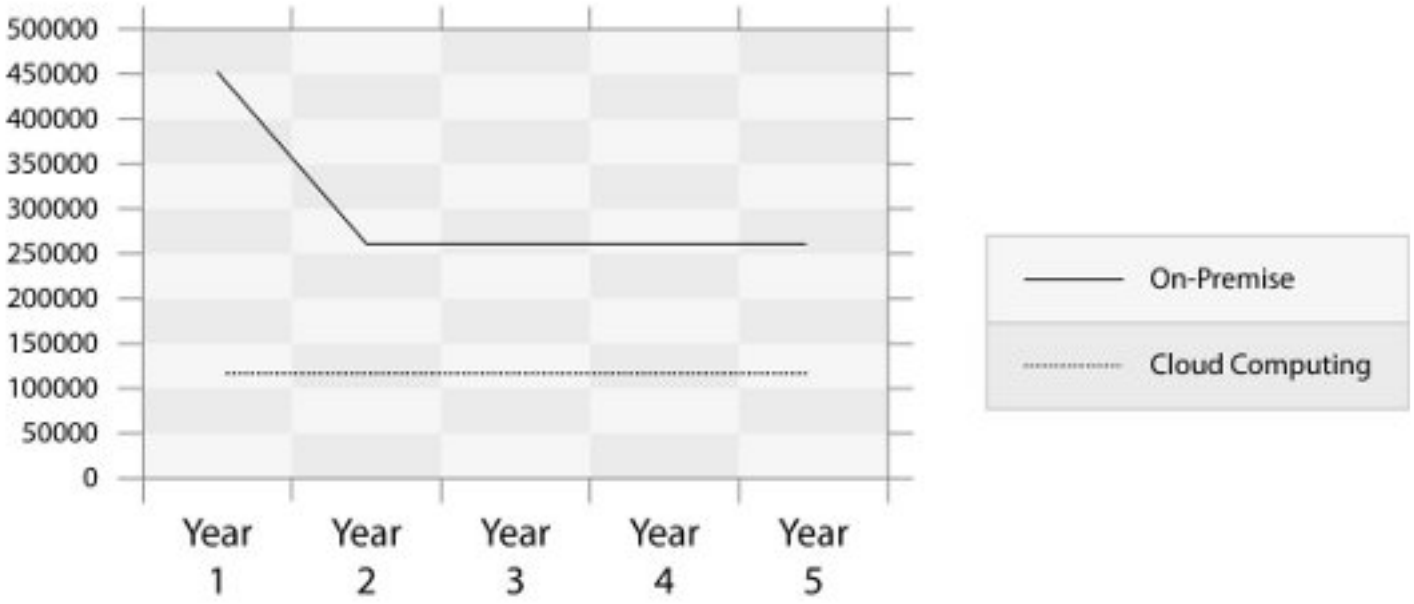

HOW WILL CLOUD COMPUTING AFFECT ORGANIZATIONS?

"As we now navigate one of the most of severe recessions of the last century, it is worth asking what comes next, post recession. During this trying time, a revolution is underway in computing. It is a revolution that promises to shake the foundations of how technology is delivered to all organizations. This revolution is cloud computing and it promises to reshape the structure of the knowledge economy."

http://tinyurl.com/yal7bzn

We believe that, in the not too distant future, most people will be making use of cloud computing. How they use the cloud will vary, but everyone will be using the cloud. It is expected that cloud computing will follow a similar path as the Internet into the corporate environment:

- there will be concern about security associated with data
- there will be concern about being dependent on services that are controlled by third parties
- many organizations will be challenged by existing corporate processes
- some organizations will be leaders in the charge while other will lag

For entrepreneurial businesses, this means that cloud computing can represent business opportunities to: i) reduce costs by leveraging cloud environments to provide infrastructure and services; ii) bring an expanded portfolio of products to a global marketplace through reduced costs to access large complex infrastructure; and iii) start small and expand services quickly and economically.

This article presents extracts from the literature on cloud computing and provides a brief outline of a common cloud computing taxonomy.

After reading this article you should have a better understanding of the basics behind cloud computing, the reasons for the large potential savings, the key challenges organizations face when implementing cloud computing, as well as the paradigm shift that cloud computing represents to organizations.

Cloud Computing Basics

A recent report (http://eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf) by the University of California at Berkeley stated, “Cloud Computing has been talked about, blogged about, written about and been featured in the title of workshops, conferences, and even magazines. Nevertheless, confusion remains about exactly what it is and when it’s useful, causing Oracle’s CEO to vent his frustration: 'The interesting thing about Cloud Computing is that we’ve redefined Cloud Computing to include everything that we already do. I don’t understand what we would do differently in the light of Cloud Computing other than change the wording of some of our ads.' These remarks are echoed more mildly by Hewlett-Packard’s Vice President of European Software Sales: A lot of people are jumping on the [cloud] bandwagon, but I have not heard two people say the same thing about it.”

While there is talk that cloud computing is a rebranding of what already exists, we also see many comments about cloud computing representing a major paradigm shift. We see articles talking about vertical, horizontal public, private, and hybrid clouds, with each representing an approach that matches cloud computing concepts to specific market requirements.

HOW WILL CLOUD COMPUTING AFFECT ORGANIZATIONS?

What appears to be emerging is a taxonomy of cloud computing that includes: Infrastructure as a Service (IaaS); Platform as a Service (PaaS); and Software as a Service (SaaS).

There is a consistent theme regarding cloud computing data centres: they are large, efficient, and highly automated. The University of California at Berkeley report outlines that commercial cloud computing centers will be capital intensive, costly, and will require high business volumes to operate. This allows these organizations to gain greater economies of scale, while also growing their business revenue. The report goes on to indicate that cloud computing providers with extremely large footprints will have the economy of scale that will allow them to purchase hardware, network bandwidth and power for 1/5 to 1/7 the price offered to medium-sized data centers. These organizations can leverage existing investments made to host internal data center requirements and expand them for revenue growth by being a leader in this emerging solution.

In addition to purchasing power, cloud computing providers will support the pooling of resources, alleviating the need for innovators to find their own resources. Instead of each organization needing resources to support its organizational needs, there will be a pool of resources supporting the cloud environment that will provide services to a global community. The implications will be significant changes to resource needs by organizations that move to cloud computing. There have been numerous testimonials of organizations scaling back their information technology (IT) organizations as they take advantage of SaaS offerings.

Cloud Computing Taxonomy

While there are many proposed taxonomies for cloud computing, the most prevalent taxonomy and the one we will use is the three tier taxonomy:

- Infrastructure as a Service (Iaas)
- Software as a Service (SaaS)
- Platform as a Service (PaaS)

Wikipedia defines IaaS as "the delivery of computer infrastructure (typically a platform virtualization environment as a service). The service is typically billed on a utility computing basis and amount of resources consumed (and therefore the cost) will typically reflect the level of activity. It is an evolution of web hosting and virtual private server offerings." For many organizations, this simply represents a low cost server rental arrangement.

Wikipedia defines SaaS as a "model of software deployment whereby a provider licenses an application to customers for use as a service on demand. SaaS software vendors may host the application on their own web servers or download the application to the consumer device, disabling it after use or after the on-demand contract expires." A large volume of cloud computing offerings being marketed on the Internet are directed towards SaaS. The marketing material highlights the significant opportunity to use SaaS to reduce costs. However, many organizations will be challenged to realize the cost benefit when compared against their internal cost metrics because future release upgrades are not built into the cost of operating the system.

Wikipedia defines PaaS as "the delivery of a computing platform and solution stack as a service. It often goes further with the provision of a software development platform, that is designed for cloud computing at the top of the cloud stack." PaaS really becomes critical when you start to think about the complexities associated with licensing.

HOW WILL CLOUD COMPUTING AFFECT ORGANIZATIONS?

Cloud Computing – Paradigm Shift

Fifteen years ago when the Internet was gaining popularity, many organizations had rules governing Internet use because of the concern around security. This ranged from only allowing Internet access from computers that were not connected to the company network to having limited access by only a few critical individuals. Similarly, security is the top concern raised regarding cloud computing.

The Big Shift: Measuring the forces of deep change, 2009 SHIFT Index (http://johnseelybrown.com/shiftindex.pdf) notes that: “Until now, companies were designed to get more efficient by growing ever larger, and that is how they created considerable economic value. The rapidly changing digital infrastructure has altered the equation, however: as stability gives way to change and uncertainty, institutions must increase not just efficiency but also the rate at which they learn and innovate, which in turn will boost their rate of performance improvement." Scalable efficiency, in other words, must be replaced by scalable learning. Innovation will be a priority for companies either to grow their business or cut costs.

Buyya et al (http://gridbus.org/papers/hpcc2008_keynote_cloudcomputing.pdf) outline that software organizations will be facing numerous challenges in creating software that is capable of supporting large volumes of individuals versus being designed for individual computers. Organizations will face challenges shifting from their existing conservative IT strategy governed by internal processes and controls to an external services model. Conservative regulatory guidelines may cause barriers moving to a public shared model.

In addition to the business and organizational challenges, there will be a number of IT challenges.

While many of the organizational challenges may be IT centric, we are separating the organizational challenges from the IT design challenges. Hagel III et al (http://hbr.org/2009/07/the-big-shift/ar/1) indicate that companies must design and track metrics showing their participation in knowledge flows to reinforce the criticality of this operational activity and the need to collaborate to improve performance and accelerate the adoption of new products and services. The Berkeley report provides a top 10 list of obstacles:

1. service availability

2. data lock-in

3. data confidentiality and auditability

4. data transfer bottlenecks

5. performance unpredictability

6. scalable storage

7. bugs in large-scale distributed systems

8. scaling quickly

9. reputation sharing

10. software licensing

Concerning the issue of cloud computing security, most concerns involve the associated risk of moving private data to a public cloud. Organizations that are leading the move to cloud computing are strategically selecting applications and functionality where the perceived security concerns pose little risk. Organizations need to select cloud service providers not only based on what they provide today, but also where the organizations will be in the future.

A History of Cloud Computing (http://tinyurl.com/c3pjh2) discusses different aspects of security in cloud com-

puting: “once the security issues are resolved, cloud computing services can enable an enterprise to expand its infrastructure, add capacity on demand, or outsource the whole infrastructure, resulting in greater flexibility, a wider choice of computing resources and significant cost savings. Considerations such as security, data privacy, network performance and economics are likely to lead to a mix of cloud computing centers both within the company firewall and outside of it."

The Berkeley report talked about the ability to scale up or down the hardware systems in cloud computing. This ability is dependent on having sufficient hardware available when required, but it is also dependent on having automated provisioning tools, tools with automated workflows to detect and make the allocation adjustments automatically, and having service level agreements with customers that allow proper prioritization of services. Cloud computing service providers need to use these tools and processes as well as make sure that they are robust with proper procedures. It also states that developers should start designing new systems to leverage cloud computing features, thus paving the way for more change. The general consensus is that cloud computing will represent a large amount of change, even if only implementing at the IaaS layer. Going to full SaaS, which appears to be the hot topic presently, represents a greater amount of change, largely through the introduction of new products and services.

The Berkeley report outlines the significant cost advantage potential of cloud computing which, combined with elimination of capital requirements and pay-as-you-go concepts, creates an environment conducive to innovative startups and thus the potential for disruptive innovation. It emphasizes the elasticity feature of cloud computing which enables organizations to manage the risks of over- and under-provisioning.

Summary and Conclusions

Cloud computing has recently been gaining in popularity and acceptance. There are many different perspectives on what cloud computing is. Regardless of whether it is rebranding or new services, the cost benefits of cloud computing services appears to be gaining acceptance in today’s economy where organizations are focused on addressing financial challenges. IaaS, PaaS, SaaS represent a cloud computing taxonomy that can be used by organizations looking to move to on-demand hardware and software services.

Cloud computing is much more than just making use of on-demand servers if organizations want to capture the full benefits. It means leveraging economies of scale through automation. It means moving to global applications, not just enterprise applications. It means a major shift towards knowledge sharing versus in-house expertise. For entrepreneurial organizations, cloud computing represents a major shift which can either represent an opportunity or a risk.

Fred Waldner is a graduate from University of Toronto's Engineering Program with a Masters in Business Administration from York Univeristy's Schulich Business School. He is currently studying Technology Innovation Management from Carleton University's Department of Systems and Computer Engineering. He provides independent Consulting and Project Management services and is a Certified Professional Project Manager from the Project Management Institute. His previous experience includes managing projects at Xerox, Metropolitan Life and IBM, including the migration of data centres, providing traditional hosting environments to On-Demand hosting centres, and leveraging virtualized environments.

"The whole point of open source-- the reason any open source project exists-- is to save us time. To keep us from rewriting the same software over and over. Puppies are cute and fuzzy and sweet, but they're also giant timesinks. To imply that an open source project is as labor intensive as a puppy is reinforcing the very worst stereotypes of open source: software that's only free if your time is worthless."

Jeff Atwood
http://tinyurl.com/gadf7

The use of open source in hosted solutions is undoubtedly widespread. However, it is seldom considered important in its own right, nor do the majority of hosted solutions providers contribute to or create open source as natural artifacts of doing good business. In this exploration of the nature of hosted solutions and their developers, it is suggested that consuming open source, as well as creating and disseminating it to collaborators and competitors alike, is essential to success.

In an open source ecosystem, hosted solutions can compete on differentiation rather than lose time and money to concerns that do not add value. Hosted solution providers can reduce the cost of their solution, the time it takes to deliver new ones, and improve their quality without additional resources.

Architecture of Hosted Solutions

A hosted solution, like most software, is designed for an end user that, for the most part, is unconcerned about the internal workings of the solution. Once the end user selects a solution, the characteristics of security, reliability, and redundancy are taken for granted. The value add of hosted solutions, due to their nature as web-based applications and the imposed limits of ubiquitous web browsers, is not usually derived from intellectual property, but rather from the solution's ability to interpret, predict, or act on the user's intent. Therefore, hosted solutions are usually "execution plays" rather than proprietary engines. This is not true in all cases, but even when a hosted solution delivers substantial proprietary value, the apparatus that it uses to communicate those results to users is still limited to "forms over data". All hosted solutions must employ some means of capturing a customer's money for services, and usually use subscription-based pricing to correspond with the nature of a hosted solution's pay-per-use business model and software-as-a-service delivery model.

Superficially, one potential beneficial side effect of hosted solutions is enabled by the definition of what constitutes a source or binary distribution. Since a web application's binaries are not distributed, but rather executed, with the resulting output transmitted to the end user, these web solutions do not have to disclose through attribution the open source components used. However, this does not extend to resources such as images and scripting, which are distributed as source in their entirety for interpretation by the web browser. Whether this effect is the result of a temporary lapse in the state of the art, or a widely accepted misinterpretation of open source licensing law, the fact remains that you can build web applications entirely out of open source components without acknowledging that fact. From a consumption standpoint, this helps hosted solution providers shop for open source components freely and evaluate them fairly against commercial alternatives. Investigating open source is part of a due diligence process for the cost-conscious developer.

Consider the architecture of a typical hosted solution: an often relational persistence engine, layers of infrastructure code responsible for assigning dependencies and provisioning data services, customer membership services to handle activations, subscription levels, cancellations, and marketing efforts, and billing architecture to collect for services rendered. While this is hardly the definition of a hosted service, it does constitute a large portion of the development effort. Every hosted solution must have these key components, and none of them require a particular finesse: customers are subscribed or they aren't, and credit card transactions are valid or void. The only remaining consideration, after infrastructure, is a service's core business and differentiators, which should account for the majority of a developer's effort. Luckily, open source provides the means to realize this balance.

Open Source Advantages

We believe that companies should use established open source, or give away their own proprietary code as the seeds of new open source projects, to build the vast majority of their hosted solutions. Through promotion and financial support of open source communities that mandate the componentization of web architecture, hosted solutions can leverage the principle of reusability to greatly reduce time to market, effectively beginning the development cycle at the service delivery stage rather than the infrastructure stage. The company can then expend its resources to solve the specific challenge their hosted solution provides to their customers. This means that a service's key differentiators combined with the execution of service delivery become the only true measures of success when infrastructure concerns are equalized by unilateral investments in open source architecture.

Because the general cost and complexity of deploying web applications has decreased dramatically, many developers of hosted solutions are smaller teams with tighter deadlines. These teams have a limited capacity for "quality bandwidth", or the sensitivity to detect and act on infrastructure issues to improve code reliability. While we know based on instinct that a customer will not respond well to an ill-conceived idea well executed, nor a brilliant idea that's impossible to take for granted, we don't realize that there is nothing to be gained competitively by advancing an internal state of the art at the cost of time, money, and quality.

It seems counter to our competitive nature that we would benefit directly from elevating our competitors to a certain level of core competency, or through implication reveal our infrastructure vulnerabilities. To quote a friend, most of a hosted solution's genetics are a "solved problem". Submitting to our tendency to believe that in-house development cannot be duplicated by others or that our conscientious investment in private application infrastructure is a differentiator in itself is the kiss of death for a hosted solution.

It benefits the entire ecosystem when everyone is using the same open source components. First, the need for technical expression that was previously invested in a low quality bandwidth strategy (in-house development) can be directed at evolving the open source offerings that are now actively developed by everyone. Because ecosystem members are motivated financially, they are incented to innovate. This process could not occur in isolation, and its continued health ensures that hosted solution development always begins near the end of the development cycle, right where the differences made matter most to overall success and profit.

The Advantages of Contributing Back

In software development there is a long-established tradition of identifying and eliminating dependencies as a marker of quality. Dependencies reflect the hierarchy of components and the code that calls into each to perform tasks. The greater the number of components, the greater overhead needed to understand, maintain, and extend that hierarchy.

Leveraging open source components acts out a law of decelerating returns, where a hosted solution could topple from sheer overlap and inefficiency. If fewer dependencies means more efficient systems, then this would add weight to the decision to procure all software development internally. Furthermore, software development cultures thrive on quality and performance, and the state of the art within an open source community may vary greatly in quality, elegance, or completeness compared to the capability of the team that would implement it.

The open source equivalent of in-house code could be inferior engineering, depending on the maturity of the project and the activity of its primary team. However, there is an opportunity for the hosted solution provider to produce rather than consume open source, and take the lead in producing the same software they would have in-house, but for a larger audience. While it is possible that developing public code rather than private code can result in longer lead times, this only occurs when the provider misinterprets the purpose of open source as competitive in its own right, rather than fulfilling the need of producing repeatable solutions with lower cost of ownership and a higher quality bandwidth.

The provider ends up focusing on competitive factors in the open source community itself, such as mindshare or the developer marketplace, instead of the market for their hosted solution. This may be an effective bootstrapping play. If the open source creator can sell solutions based on their open source technology, it could fund the development of the consuming application, though this would largely be a departure from the original intent.

Competitive factors exist, but should be left to the greater community. It is important to focus on the community surrounding your stewardship, and not interpret the open source movement as a zero-sum game that requires a winner and loser. For example, many open source projects suffer from a lack of documentation. If an open source project gains traction without adequate documentation, then the cost of improving the documentation is justified by the quality improvement the community provides, but is not accounted for prior to community growth or success.

This has a two-fold positive effect. First, it increases the quality bandwidth of an equivalent solution if the project is successful in attracting supporters. Second, it forces competition to either accept your prevalent solution, spend equal or greater resources improving it for their own needs, or recreate similar infrastructure for their own projects without reaping the reward of an open source community. Even in an extreme case where a potential competitor decides to launch their own equivalent open source resource rather than collaborate with your community, the introduction of choice into an area previously underserved or dominated by commercial players only helps your efforts. It helps in two ways: by supplying alternative views on hosted solution concerns that can be replicated safely under open source licensing, and by bolstering community through network effects.

By no means should a development team aspire to open source its "secret sauce", but many hosted solutions are developed by bootstrapping web development firms whose primary means of growth, prior to launching a hosted solution startup, has been through contracting for clients. These are companies that are accustomed to thinking that their core business is software development rather than the simplified accounting, online image editing, or other service showcased in their hosted solutions. They correctly prescribe a mandate not to share their unique value proposition and techniques, but apply that thinking to rudimentary functions of their hosting infrastructure, the same area where they are normally poised to deliver value as consultants.

Conclusion

Consider that Ottawa's Shopify (http://shopify.com), a success story in a flooded hosted eCommerce market, is the curator of no less than eleven open source projects. Each is easily integrated into competing products and represents countless hours of development, integration with dozens of third parties, and thorough testing. While these projects do generate some amount of promotional value to the company, Shopify is not in the software development business, they are in the hosted eCommerce storefront business. On the surface, there is very little for them to gain from developer goodwill. However, it is clear that they steer the development choices of competitors, inspire innovation around their mission-critical systems, and increase their quality bandwidth from public exposure to critical payment processing code. This is the case with ActiveMerchant (http://activemerchant.org), their open source project responsible for handling credit card transactions from dozens of merchant services.

Shopify is the market leader because they differentiate and execute their unique selling proposition, and they leverage open source code they could have easily kept behind closed doors, to help ensure they stay competitive.

While there are many ancillary benefits to promoting open source, such as developer goodwill leading to personnel attraction and retention, the effects of open source are far more integral. By refining the focus of a hosted solution provider to its true core business and repurposing all non-essential infrastructure development through open source consumption and stewardship, companies that create hosted solutions stay competitive by maximizing the potential of the code they write, the code they don't write, and the communities they build.

Daniel Crenna is an award winning developer, author, and active entrepreneur from Ottawa, Canada. He is the lead developer of several open source projects including TweetSharp (http://www.tweetsharp.com), the most popular Twitter development library for .NET developers, and DotNetMerchant (http://dotnetmerchant.org), a hosted eCommerce platform built on Microsoft's ASP.NET technology. He is currently developing Lunarbits, a hosted eCommerce solution for small to medium publishers selling digital goods. He recently became a graduate of the Lead to Win (http://leadtowin.ca) program.

"If you see a bandwagon, it's too late."
James Goldsmith

Product marketers love a bandwagon, and no bandwagons have been more appealing in the past ten years of information technology (IT) as the Service Oriented Architecture (SOA) and cloud computing ones. Much of the challenge of marketing products is getting the attention of the target customer in order to create an opportunity to pitch products or services to them. Of course, if it doesn't work with one bandwagon, as the old adage goes, try, try again. This is why we often see the same products marketed with different labels and categories applied to them. Product vendors will insist that they have developed some new add-on or tweaked a user interface to include the new concept, but at the very core, the products remain fundamentally unchanged.

It is particularly frustrating when product marketing gets in the way of implementing what otherwise would be a valuable concept. Competing vendor, consultant, and individual implementer messages on the meaning of a specific term interferes with realizing real value. This is especially the case with the emerging concept of private clouds. While the term could potentially have real meaning and lasting value, the product and consulting marketers have turned any potential meaning into mush that hides that value.

Cloud computing is primarily loosely-coupled, location-independent virtualized services run on abstracted infrastructure with the primary intent of reducing IT expenditures, increasing flexibility, or improving overall system robustness. Given that this is the general cloud concept, is there any value in a new concept called private clouds? How does the addition of the word private add value to the service-oriented cloud computing that has been discussed for a handful of years? Is it a valuable term, or mere marketing spin? This article first examines the range of definitions being applied to the private cloud concept, then offers a summary on the value provided by private clouds.

Concept #1: Company-Owned and Operated, Location-independent, Virtualized (Homogeneous) Service Infrastructure

Some are of the opinion that a private cloud consists of infrastructure owned by a company to deploy services in a virtualized, location-independent manner. What differentiates private clouds from simply implementing clustered applications or servers is that the cloud is not built for a specific service or application. Rather, it is an abstracted, virtualized environment that allows for deployment of a wide range of disparate services. In practical terms, companies will most likely not implement this vision of private clouds using a diversity of heterogeneous infrastructure. Indeed, it is in their best interests to control the costs and complexity of support, training, and administration by implementing their private clouds using a single vendor stack. So, this vision of private clouds is often a single-vendor (homogeneous) cluster of virtualized infrastructure that enables location-independent service consumption. Implementing any sort of homogeneous stack reduces the need for loosely-coupled services, and thus weakens the service-oriented cloud computing value proposition as a whole for that company. For most companies, what they used to call grids are now being relabeled as private clouds.

Concept #2: Virtualization Plus Dynamic Provisioning (Elasticity)

Others assert that the above definition doesn't go far enough. In order for the company-owned and implemented infrastructure to be considered a private cloud, it must include the concept of

elasticity. Elasticity means that the hardware and software resources must be provisioned in a dynamic manner, scaling up and down to meet changes in demand, thus enabling a more responsive and cost-sensitive approach to IT provisioning. This idea of private clouds sounds a lot like the Utility Computing (http://utilitycomputing.com) concept sold as part of IBM's decade-old vision of On Demand Computing. From this perspective, a private cloud is company-owned On Demand Utility Computing implemented with services instead of tightly coupled applications.

Concept #3: Governed, Virtualized, Location-Independent Services

There is a notable contingent that argues that the private cloud is a response to some of the security and governance issues raised by the (public) cloud. Specifically, a private cloud provides more control over what and how. In this regard, a private cloud is a governed cloud that enables virtualized, governed, location-independent services. There has been much consternation over the fact that the most popular public clouds share infrastructure between customers and require that data and communications cross the company firewall. Many IT administrators and managers want the technological benefits of cloud computing without the governance risk of having it reside in someone else's infrastructure. Basically, they want the virtualization, loose coupling, and location-independent benefits of cloud computing without the economic benefits of leveraging someone else's costs and investments. In essence, they would rather own a version of the Amazon EC2 (http://aws.amazon.com/ec2) cloud than use it, solely for reasons of governance.

Many are concerned about the governance and security draw-backs of cloud computing. Rather than simply dismissing the economic benefits of public clouds, why not approach private clouds as a veneer placed on top of public clouds? Couldn't companies impose their governance and security requirements on third-party infrastructure, using company-owned governance tools and approaches to manage remote services? Couldn't organizations demand that public clouds provide greater governance and security control? Does the addition of the term private provide the same sort of value as it does in the context of the Virtual Private Network (VPN)? We didn't throw out the Internet because it was insecure and replace it with individual private Internets. So, why should we replace public cloud computing with private clouds?

Concept #4: Internal Business Model for Pay on Demand Consumption of Location-Independent, Virtualized Resources

Some individuals and companies insist that the primary value of any cloud, whether implemented privately or acquired from a public vendor, is the business model of pay-as-you-go service consumption. From this perspective, a private cloud is an internal business model that enables organizations to consume and procure internal, virtualized, loosely-coupled services using a pay-on-demand model similar to a chargeback mechanism. Rather than an IT organization paying for and supporting the costs of the business users in an aggregate fashion, they can provide those resources using the same business models employed by Amazon, Google, Salesforce.com and others in their public clouds. In order to realize this vision of private clouds, companies need a means to enable transactional service purchases, auditing of service usage, and organizational methods for enabling such inter-departmental charges. At the most fundamental level, this vision of the private cloud treats IT as a business and a service provider to

the rest of the organization.

Concept #5: Marketing Hype, Pure and Simple

TechTarget offers the most cynical view of the private cloud. In their words (http://searchcloudcomputing.techtarget.com/sDefinition/0,,sid201_gci1333074,00.html), a private cloud is a "marketing term for a proprietary computing architecture that provides hosted services to a limited number of people behind a firewall. Marketing media that uses the words "private cloud" is designed to appeal to an organization that needs or wants more control over their data than they can get by using a third-party hosted service." Basically, they opine that the term has marketing value only. Where does this place IT practitioners? Reading between the lines, they encourage us to ignore the usage of the term.

What Does it Mean, If Anything?

The fact that there's no single perspective on private cloud might indicate that none of the definitions really warrant separating the private cloud concept from that of cloud computing as a whole. One reasonable perspective is that the definitions discussed above are simply differing infrastructural and organizational approaches to implementing service-oriented cloud computing. Those approaches do not warrant a whole new term and certainly not millions more in infrastructure expenditure. Trying to create a new concept of private clouds from any of a number of perspectives--architectural, infrastructural, organizational, governance, and business model--seems to introduce more confusion than clarification. After all, shouldn't all clouds, private or not, provide the benefits described above? Doesn't the concept of a private, company-owned cloud in some ways weaken the cloud value proposition?

Who really benefits from this private cloud discussion: IT practitioners or vendors with products to sell?

ZapThink (http://zapthink.com) has had many positive things to say about cloud computing, and we do believe that as a business model, technological approach, and service-oriented domain it will have significant impact on the way companies procure, develop, deploy, and scale their applications. We're starting to see hundreds of companies that develop whole products and services without procuring a penny of internal IT hardware or software resources. This is the bonanza that is cloud computing.

The point of any new term should be to clarify and differentiate. If the term does neither, then it is part of the problem, not the solution. When vendors start promoting their products with new terms, ask yourself: does this change what you are doing now, or is this the beating of the bandwagon's marketing drum? The goal is not to buy more stuff; the goal is to provide the business increasing value from existing IT investments. This is the purpose and goal of enterprise architecture and the reason why IT exists in the first place.

Ronald Schmelzer is a Managing Partner at ZapThink LLC. ZapThink's Licensed ZapThink Architect (LZA) SOA Boot Camps provide four days of intense, hands-on architect-level SOA training and certification. Advanced SOA architects might want to enroll in ZapThink's SOA Governance and Security training and certification courses. Or, are you just looking to network with your peers, interact with experts and pundits, and schmooze on SOA after hours? Join us at an upcoming ZapForum event. Find out more and register for these events at http://zapthink.com/eventreg.html.

"Anything you build on a large scale or with intense passion invites chaos."
Francis Ford Coppola

Mr. Coppola was, of course, talking about the making of Apocalypse Now. But the general case is also true, and perhaps particularly relevant when applied to information technology (IT). From organization to architecture, managing the mayhem is increasingly difficult as the scope grows. It takes a very good architect to harness the chaos of the Internet and the vast array of available options, and design systems that support thousands or even millions of users.

This article discusses the opportunities and challenges when leveraging the cloud to performance test large-scale websites and applications. Getting performance right, particularly at web-scale, requires a level of passion that results in both a view of the big picture and an attention to detail. We'll describe how to use the scale of the cloud to gain confidence when deploying sites servicing potentially massive amounts of web traffic. We'll start by describing what we mean by the cloud in order to set context. We'll then describe the components that comprise SOASTA's CloudTest (http://soasta.com/product), focus on how we deliver the offering to the market, and examine the key tenets of the cloud testing methodology based on SOASTA's experiences.

Scalability Issues

Scale introduces the potential for chaos. The number of things, from hardware to software to network, that can impact performance grows exponentially as an application scales and the inter-dependencies between the various components grow increasingly complex.

Testing at web-scale requires:

- an infrastructure to support the traffic
- a means to manage the deployment and execution of the test
- analytics that can massage the massive amounts of data and deliver actionable information
- Performance Engineers with the experience to navigate through complex environments
- a methodology to do it all as efficiently and effectively as possible

This article describes how SOASTA addresses these issues and the considerations for anyone trying to test at web-scale.

At SOASTA, we help ensure that websites and applications are highly reliable and scalable. This has become an increasingly dicey proposition. External events such as the Super Bowl, Cyber-Monday, Mother's Day, a significant drop in the interest rate, or even news about a celebrity can suddenly drive unforeseen traffic to a website. The impact of social media can exacerbate the situation, turning the success of driving traffic to a website into an embarrassment if a site is slow, or worse, if it crashes. The proliferation of mobile devices and the effort by businesses of all types to expand their reach has driven a level of traffic to websites and applications that were previously only experienced by a few, high profile sites.

SOASTA helps companies respond to these unprecedented spikes in traffic by delivering performance intelligence: capturing data about all the elements that impact a website's responsiveness and reliability and immediately turning it into actionable information. SOASTA delivers this service through a combination of a test application, services, and a Performance Test methodology that recognizes

the unique challenges of testing live websites at peak traffic loads.

The Cloud

Most IT professionals tend to describe the cloud in terms of services, typically categorized as:

• Software-as-a-Service (SaaS)

• Platform-as-a-Service

• Infrastructure-as-a-Service

Software-as-a-Service is characterized by providing application capabilities, such as backup, email, customer relationship management (CRM), billing, or testing. Platform services, such as Force.com, Google Apps (http://google.com/apps), Microsoft Azure (http://microsoft.com/windowsazure) and Engine Yard (http://www.engineyard.com), are designed to make it easier to build and deploy applications in a specific multi-tenant environment without having to worry about the underlying computing stack. Infrastructure services allow companies to take advantage of the compute and storage resources delivered by vendors, such as Amazon, Rackspace (http://www.rackspace.com) and GoGrid (http://www.gogrid.com). The differentiator from traditional hosting or managed service providers is that with cloud computing you have greater flexibility provided by virtualization and elastic application programming interfaces (APIs).

Software and Infrastructure services are the farthest along in terms of adoption. While SaaS has often been viewed as an alternative to open source, offerings from companies like Salesforce.com have evolved to open source solutions such as SugarCRM (http://sugarcrm.com/crm) and Zimbra (http://www.zimbra.com). Similarly, the early infrastructure providers developed proprietary APIs and used commercial tools for virtualization to build out their services. Today, products such as libCloud (http://incubator.apache.org/libcloud) and Eucalyptus (http://www.eucalyptus.com) are emerging open source alternatives for providing elasticity. Open source virtualization solutions are also available.

Clouds are often categorized as private, public, or hybrid. Internal clouds are often about optimizing a private infrastructure and are usually referred to as private clouds. Public clouds provide access to universally accessible resources. Today, public cloud vendors are offering hybrid alternatives that leverage their compute and storage resources, yet require the proper authorization for access. Extending an internal or managed infrastructure by renting virtualized resources on-demand has become an increasingly viable option. As a result, the line between public and private clouds is starting to blur.

With a private cloud, instead of renting the infrastructure provided by others, companies use commercial or open source products to build their own cloud. Of course, more control means more responsibility. These companies need to address challenges such as repurposing servers, choosing a virtualization platform, image management and provisioning, and capturing data for budgeting and charge back.

The Killer App for the Cloud

It has become clear that the infrastructure services provided by the cloud are being driven by applications that are particularly well suited to taking advantage of virtualization and elasticity. Elasticity is most commonly associated with changes in supply and demand based on price. The definition for the cloud is not much different.

An elastic API refers to an infrastructure vendor's ability to quickly respond to demand by allowing customers to quickly spin up servers, and just as quickly take them down. For applications such as performance testing this is incredibly important.

The best-known examples are companies that leverage the cloud to accommodate dramatic swings in traffic or short-term application requirements. For instance, pharmaceutical and financial services companies have heavy compute requirements that may last for hours or weeks at a time. These apps are a great match for the ability to deploy hundreds or thousands of compute cores, on-demand. As a specific example, last year Intuit.com deployed servers in the cloud for their Taxcaster application since it only needed to support peak traffic for the few weeks out of the year surrounding tax deadlines.

Testing and development has become a killer app for the cloud. Development and test beds are often deployed for short periods of time, have increased use based on business cycles, need to scale on-demand to respond to specific requirements such as duplicating production infrastructure or generating load, have limited privacy and security concerns, and are simple to deploy.

Leveraging the cloud for testing yields a number of benefits:

- it's now possible to test at both typical and peak traffic levels, from hundreds of users to millions
- generating geographically dispersed load provides the most accurate representation of real-world traffic
- the lower cost enabled by renting hardware, and using an on-demand service, allows testers to respond to accelerated development cycle times by making agile performance testing a realistic alternative
- measuring both internal and external tests, and using both the lab and production environments, provides the most efficient and effective results
- testing live web-based applications in production and from outside the firewall is the only way to gain complete confidence in the application

This last point usually causes folks to sit up and take notice. There's risk involved in testing a live, production site. Historically, companies would test a fraction of expected load in the lab and then extrapolate what those results would mean when actual load was hitting the site. So where's the greater risk: testing a live site or not knowing if your site will scale to meet demand?

According to JP Garbani, Vice President and Principal Analyst at Forrester Research, 74% of application performance problems are still reported by end-users to a service/support desk rather than found by infrastructure management. That may have been somewhat acceptable when most applications were internal facing and had hundreds of users. But web applications are exposed to a much larger potential audience, often made up of customers. And the failure or poor performance of an application has a direct impact on the perception of the brand and, potentially, revenue.

The SOASTA Performance Test Methodology

To address this gap, SOASTA leveraged its experience deploying nearly 300,000

on-demand, load testing servers to develop a methodology that leverages existing best practices to extend traditional approaches and address the new opportunities and challenges presented by cloud testing. The following sections provide a high level view of this methodology.

Testing in the Performance Lab

Cloud testing does not obviate the need or eliminate the benefits of testing in a lab environment as well as the production environment, and it's important to have continuity between the two. Ongoing performance testing in a lab allows application engineering teams to assess performance over time, and helps catch any show-stopping performance bugs before they reach production. In addition, the lab provides a place to performance test code and configuration changes for performance regression, before releasing changes to production and outside of the normal build cycle. This could include things like a quick bug fix in a page, or a seemingly minor configuration change that could have a performance impact and should be tested before it is deployed. Often, these kinds of changes are deployed with little to no testing and come back later to cause performance issues.

Testing in Production

Testing in production is the best way to get a true picture of capacity and performance in the real world. Testing in production is the only way to ensure that online applications will perform as expected. There are many things that SOASTA's production testing approach typically catches that cannot be found with traditional test methods. These include:

- batch jobs that are not present in the lab (log rotations, backups, etc.) or the impact of other online systems affecting performance
- load balancer performance issues, such as mis-configured algorithm settings
- network configuration problems such as 100MB settings instead of 1GB on switches and routing problems
- bandwidth constraints
- latency between systems inside and outside of application bubbles

SOASTA's production testing methodology helps identify the invisible walls that show up in architectures after they move out of the lab. Traditionally, testers have been limited to making extrapolations over time about whether small tests on a few servers in a lab can support exponentially higher amounts of load in production. Without proper testing, these types of assumptions always result in hitting unexpected barriers after multiple years of consistent traffic growth. We have seen that successful companies are using production testing to learn things about the performance of their sites that they could have never learned in a lab.

Strategy and Planning

This approach to performance engineering calls for an umbrella strategy with associated individual test plans. Test plans roll up into an overall strategy that ensures confidence in the ability of key revenue generating applications to perform as expected. The result is an ongoing performance engineering strategy throughout an application's evolution. It includes a number of test plans centered on individual objectives, such as holiday readiness, a major architectural change, or the release of a major version of code.

Having a well-defined strategy, with explicit test plans, provides business and engineering leaders with a high degree of confidence in operational readiness.

Using this approach gives greater insight into an application's performance and readiness.

Using an iterative process within test plans to achieve defined goals allows for a stream of continuous improvement in the applications being tested. A cycle that starts with the test definition and ends with obtaining actionable intelligence results in a continuous cycle of improvement.

The process of creating a test plan starts with the define phase. During this phase, the flows to be tested throughout the site are defined, metrics to be monitored are established, and success criteria for the tests are agreed upon.

In the design phase, the user scenarios are written and test parameters are set up. Things such as the mix of users executing different parts of the application, the virtual user targets, and the ramp-time are modeled.

The test phase is where the execution of tests takes place, and where data is collected for assessment.

Finally, the assess phase, parts of which may occur during the test execution, is when the data collected throughout test execution is used to provide actionable intelligence.

Types of Tests

The following are common test types included in a plan, which, when taken together, make for a well-rounded view of application performance and reliability. The most successful online application companies are executing on well-defined performance and readiness plans that include a mix of these tests.

Baseline: the most common type of performance test. Its purpose is to achieve a certain level of peak load on a predefined ramp-up and sustain it while meeting a set of success criteria such as acceptable response times with no errors.

Spike: simulates steeper ramps of load, and is critical to ensuring that an application can withstand unplanned surges in traffic, such as users flooding into a site after a commercial or email campaign. A spike test might ramp to the baseline peak load in half of the time, or a spike may be initiated in the middle of steady state of load.

Endurance: help ensure that there are no memory leaks or stability problems over time. These types of tests typically ramp up to baseline load levels, and then run for anywhere from 2 to 72 hours to assess stability over time.

Failure: ramps up to peak load while the team simulates the failure of critical components such as the web, application, and database tiers. A typical failure scenario would be to ramp up to a certain load level, and while at steady state the team would pull a network cable out of a database server to simulate one node failing over to the other. This would ensure that failover took place, and would measure the customer experience during the event.

Stress: finds the breaking point for each individual tier of the application or for isolated pieces of functionality. A stress test may focus on hitting only the home page until the breaking point is observed, or it may focus on having concurrent users logging in as often as possible to discover the tipping point of the login code.

Diagnostic: designed to troubleshoot a specific issue or code change. These tests typically use a specially designed scenario outside of the normal library of test scripts to hit an area of the application under load and to reproduce an issue or verify issue resolution.

SOASTA CloudTest

SOASTA CloudTest is deployed as an on-demand service, leveraging the cloud to generate the load. It is comprised of the methodology described above, the services provided by our experienced load testers, and the Global Cloud Test Platform that provides a cross-cloud infrastructure for generating load.

Within the application, open source libraries are a fundamental part of the offering, used throughout the product for providing various functions. SOASTA provides the software as part of the service. As seen in Figure 1, CloudTest is deployed using a distributed architecture in the cloud, complemented by an appliance for testing behind the firewall.

While customers can use SOASTA's application for test creation and execution, the Global Cloud Test Platform is built to support additional tools, including Apache JMeter (http://jakarta.apache.org/jmeter), the most popular open source load-testing tool.

Figure 1: SOASTA Architecture

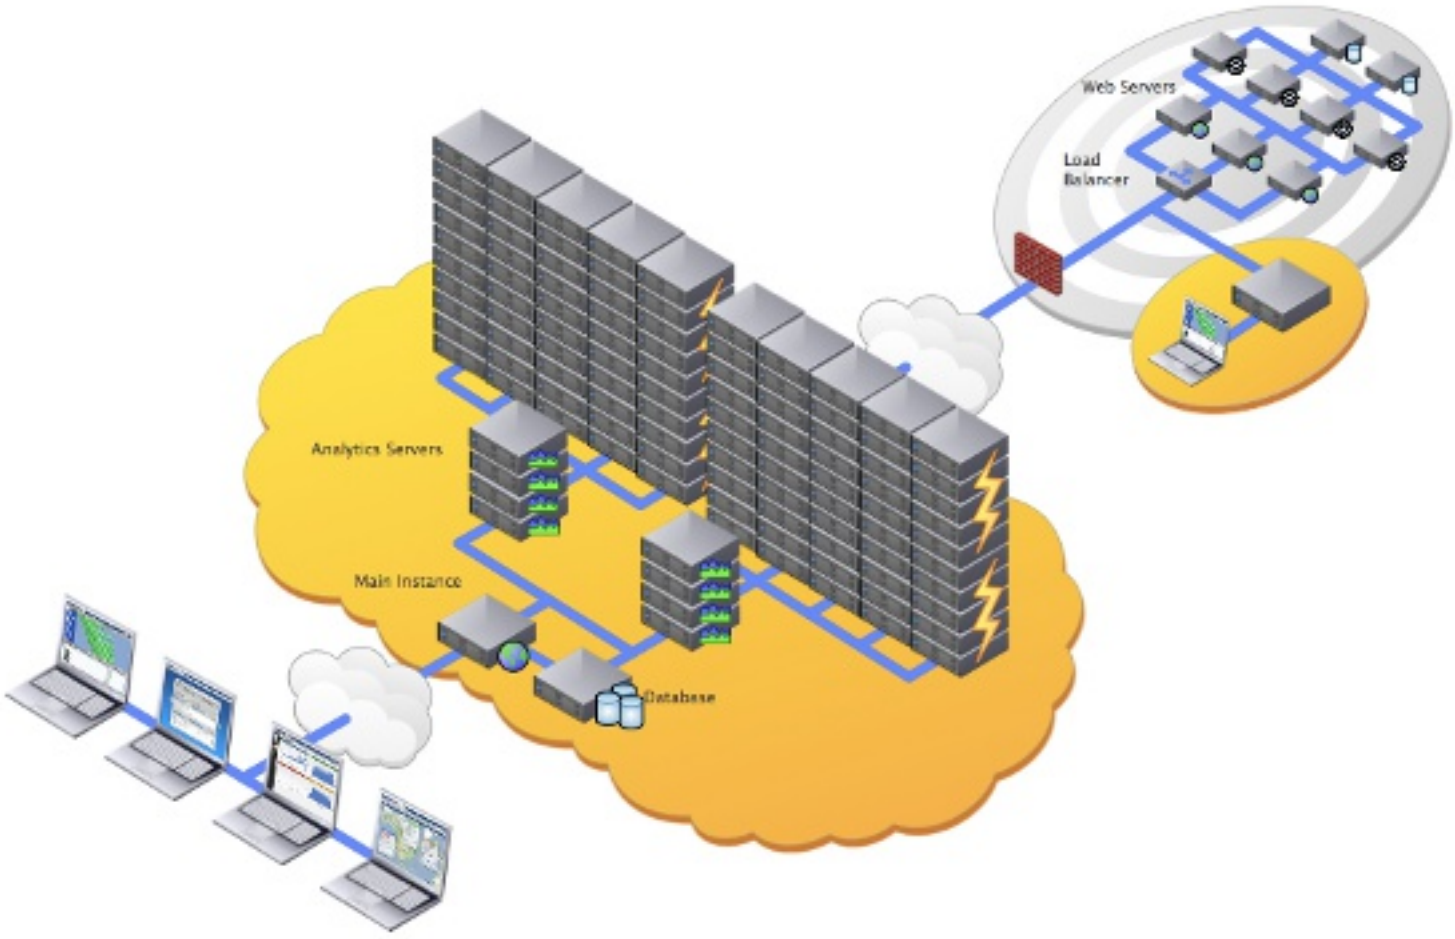

The SOASTA platform reduces the complexity and time of deploying JMeter scripts to the cloud, making it dramatically easier for the JMeter community to create, deploy, execute and analyze web-scale load and performance tests. JMeter scripts run without modification. Once the test is built, SOASTA takes care of managing and provisioning servers and executing the test.

The key capabilities we've built into this approach come as a result of our experience deploying to the cloud. The first deployment environment was Amazon EC2 (http://aws.amazon.com/ec2). Because the requirements for load and performance testing fit almost all of the characteristics described above, Amazon's implementation of a cloud infrastructure was a perfect match. EC2 was the first to provide a platform that dramatically changed the cost equation for computing resources and delivered an elastic API for speed of deployment.

As the application depends on the swift provisioning and releasing of servers, SOASTA had to quickly identify bad instances and bring up replacement instances. The provisioning technology in SOASTA's implementation is one of the key features of the platform. As new APIs, including open source alternatives such as libCloud become available, SOASTA will use them to expand the reach of the Global Test Cloud.

The other key capability is a real-time analytic engine built exclusively for testing web and mobile applications, enabling quality assurance and development teams to test and monitor their websites under both typical and extreme traffic conditions. Given the massive amounts of data generated in web-scale tests, including the resource being monitored as the test is executed, a cloud-based, highly scalable engine is required to provide actionable information in real-time.

Summary Thoughts

The cloud has approached that point in the hype cycle where its value is being questioned by various pundits because the benefits don't necessarily conform to their specific requirements. The reality is that many companies have found tremendous value. When combined with web-based technology, experienced people and a new methodology, it is clear that performance testing from the cloud can help tame the chaos associated with large scale.

Tom Lounibos, CEO for SOASTA, brings more than 30 years of experience building early stage software companies, leading two companies to successful IPO's. Tom is a regular speaker at both Cloud and testing events, and has become a leading advocate in using the Cloud to empower individuals and accelerate changes in how applications are built, tested and deployed. Most recently, Tom served as President and CEO of Kenamea. Prior to Kenamea, he was CEO of Dorado Corporation, a financial services software provider. Previous to Dorado, he was EVP of Sagent Technology through its 1999 IPO, entrepreneur-in-residence at Crosspoint Venture Partners, and held executive positions at Digitalk Corporation, Knowledgeware (KWI) and Encore Financial Services. Tom also serves on several boards in the Silicon Valley.

Recommended Resources

10,000 Hours in the Cloud
http://opensource.sys-con.com/node/1150203

A Tale of Clouds
http://tinyurl.com/ylr268q

Cloud Computing: New Wine or Just a New Bottle?
http://tinyurl.com/ykkvyrn

"Lots of users are looking at lots of data and doing eyeball correlation. Ideally, we'd like to automate that."

Dave Mihelcic, CTO, Defense Intelligence Systems Agency

After Google Maps Mashups launched the mashup craze in 2005, this consumer-based idea quickly evolved, matured and migrated into the enterprise. Structured tools and platforms emerged that codified more formal, repeatable approaches to mashups, including products like IBM Mashup Center (http://www-01.ibm.com/software/info/mashup-center), JackBe Presto (http://jackbe.com/products), and WSO2 Mashup Server (http://wso2.com/products/mashup-server). Finally, the emergence of cloud-based hosting platforms and cloud-based data services have given many enterprises access to a richer set of options for mashup creation and deployment.

Where are mashups in the enterprise today? The goal of this article is to document and share mashup use-cases. It is also an effort to expand upon the previous OSBR articles Open APIs, Mashups and User Innovation (http://osbr.ca/ojs/index.php/osbr/article/view/653/615) and Mapping Mashup Ecosystems (http://osbr.ca/ojs/index.php/osbr/article/view/860/830). It provides practical mashup implementation examples and patterns collected through empirical and anecdotal research conducted with commercial and government mashup practioners.

In the last three years we have performed many formal interviews and surveys with commercial and government mashup practioners in all 50 United States and over 50 countries around the world. The examples in this article are based on the practical implementation examples and patterns collected through this empirical research as well as anecdotal information from our own implementation experience of 'inside the firewall' and cloud-based mashups. Much of the research cited is based upon data collected from the 3,890 members of JackBe's Mashup Developer Community (http://jackbe.com/enterprise-mashup) and that community's Presto Cloud mashup platform running on Amazon EC2 (http://cloud.jackbe.com).

From our research we've learned many things about the ways in which mashups are (or are expected to be) applied. One macro-trend appears to unite all other enterprise mashup trends and patterns: enterprise mashups are most applicable in data-intensive areas of an organization where i) information needs are dynamic; and ii) data must reach the user. Given the explosive rate of data growth in enterprise data in recent years (http://economist.com/surveys/displaystory.cfm?story_id=15557443), enterprise mashups can have meaning to just about every organization.

The cloud plays a critical role in the relevance and value of enterprise mashups. The value of a mashup increases exponentially as more data sources are incorporated. It is easy to assume that the onus is upon data sources inside the enterprise firewall and in the past this was generally true. The recent rapid growth of enterprise-quality cloud-based data sources makes it much easier to establish genuinely meaningful mashups from third-party data services.

Organizational and Industry Trends/Uses

From an organizational and industry perspective, mashup use-cases and usage patterns are decidedly difficult to generalize. Anthony Bradley, Group Vice President at Gartner Research, said it this way: "...you can't build a general business case for mashups...although you can build a

specific mashup-centric business cases." From our own implementation experiences and research, we believe this generally true due to the breadth of applications and industries. However a few patterns have emerged with modest frequency.

Within the realm of government, mashups are gaining a credible reputation within military agencies as a situational awareness solution (http://www.eweek.com/c/a/Web-Services-Web-20-and-SOA/Mashups-Give-Defense-Department-Strategic-Edge-391656/1/). These mashup-driven dashboards can help the military planner become better "aware of what is happening around [them] to understand how information, events, and [their] actions will impact [their] goals and objectives, both now and in the near future" (http://en.wikipedia.org/wiki/Situation_awareness).

Mashups expedite what is often a labour-intensive process. As one commander put it, they reduce the "gather time" and thus "increase my window for decision and response". Many Situational Awareness Dashboard implementations rely upon cloud-based open sources as an important element in the dataset. An example Situational Awareness Dashboard, based upon an effort being completed for the Defense Intelligence Systems Agency (http://www.eweek.com/c/a/Web-Services-Web-20-and-SOA/Mashups-Give-Defense-Department-Strategic-Edge-391656), is shown in Figure 1.

Mashups are also being employed by some federal, state and local agencies to support transparency requirements resulting from the Open Government Initiative (http://whitehouse.gov/Open), as well as reporting requirements resulting from the Emergency Economic

Figure 1: Mashup-Driven Situational Awareness Dashboard

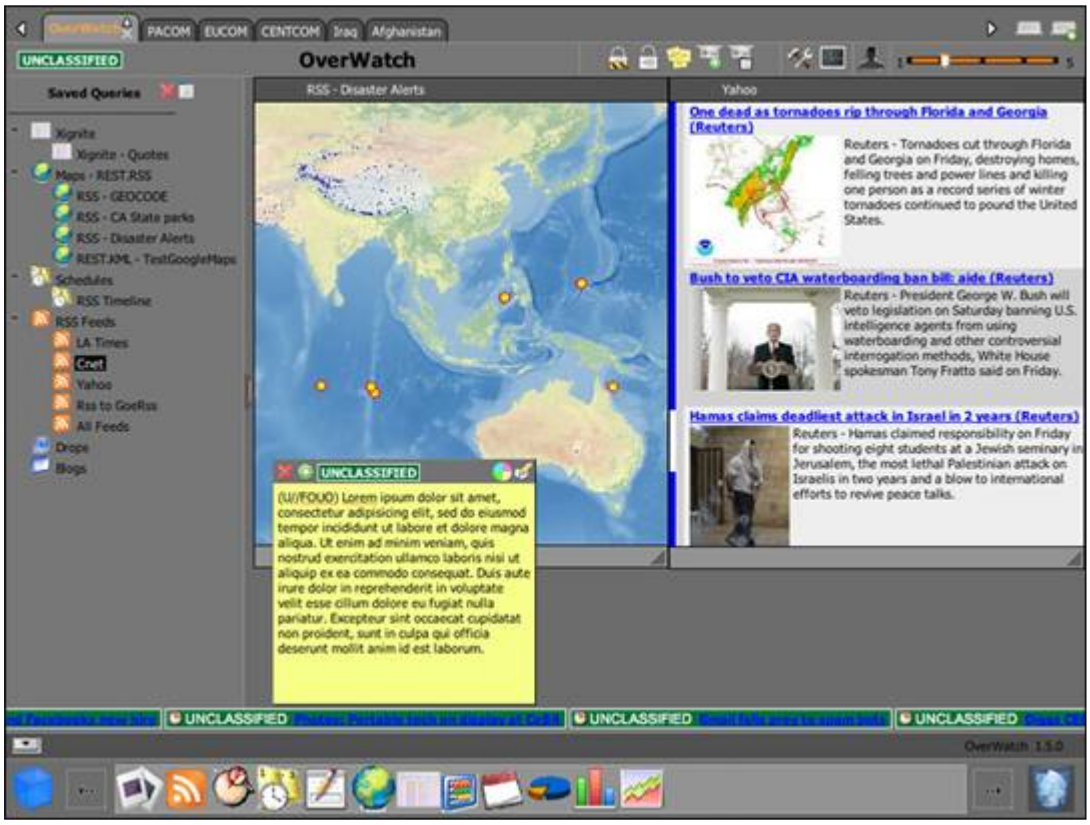

Stabilization Act of 2008 (http://en.wikipedia.org/wiki/Emergency_Economic_Stabilization_Act_of_2008) and the American Recovery and Reinvestment Act (ARRA) of 2009. In this case the mashups produce machine-readable feeds for consumption by other cloud-based systems/applications, as well as employee- and citizen-facing Web-based dashboards. The mashups assemble a wide array of information from internal agency systems that manage financials, grants, grantees, recipients, awards, and projects, as well as cloud-based data from other government entities with public application programming interfaces (APIs).

With the commercial sector, mashup-driven operations dashboards appear to be particularly popular within companies that deal in the manufacture or maintenance of physical goods or assets. Other trackable trends in the commercial space include mashups for online product marketing, project management, network operations, network cyber-defense (http://youtube.com/watch?v=2Q_9WtEOY1c), disaster planning and response (http://youtube.com/watch?v=6xB0psBjpjI), and content assembly/syndication. Many of these use cases rely upon cloud-based data services as an important input to the mashup; many of these mashups also subsequently publish standardized data services for consumption by other cloud-based systems/applications.

These major mashup application trends can be summarized as:

Operations dashboard: similar to situation awareness dashboards, mashup-driven operations dashboards provide management and key decision-makers with a broad view of actionable information sooner than traditional data warehouse and business intelligence tools. For example, one bank created an operational dashboard for senior executives with real-time operational visibility in the bank's key performance indicators, including financial, human resources, and customer metrics.

Online product marketing: the mashup suggests which products to promote in online sales campaigns by correlating up-to-the-minute news stories and online search trends with the company's product catalog.

Project management: consolidates project expenses, budgets, timelines and other relevant data from databases, spreadsheets, and project files for analysis and reconciliation.

Network operations: dynamically connect and unify data from network devices, application logs, topology systems, support knowledge bases, and trouble ticket systems, giving network managers better information that can reduce issue resolution times.

Network cyber-defense: a mashup dashboard that reconciles digital network attack signatures with helpdesk trouble tickets and allows analysts to identify and respond to threats more effectively.

Disaster planning and response: an assembly of open/public news and information sources related to natural disaster can provide a unique picture for responders and aid workers to plan, coordinate and integrate their efforts, ultimately reducing the impact of the disaster. A sample Disaster-Planning Mashup, created from the CDC's 2009 state-by-state H1N1 Flu Statistics (http://cdc.gov/flu/weekly) in tandem with staffing levels for a theoretical nation-wide company, is shown in Figure 2.

Figure 2: Disaster-Planning Mashup

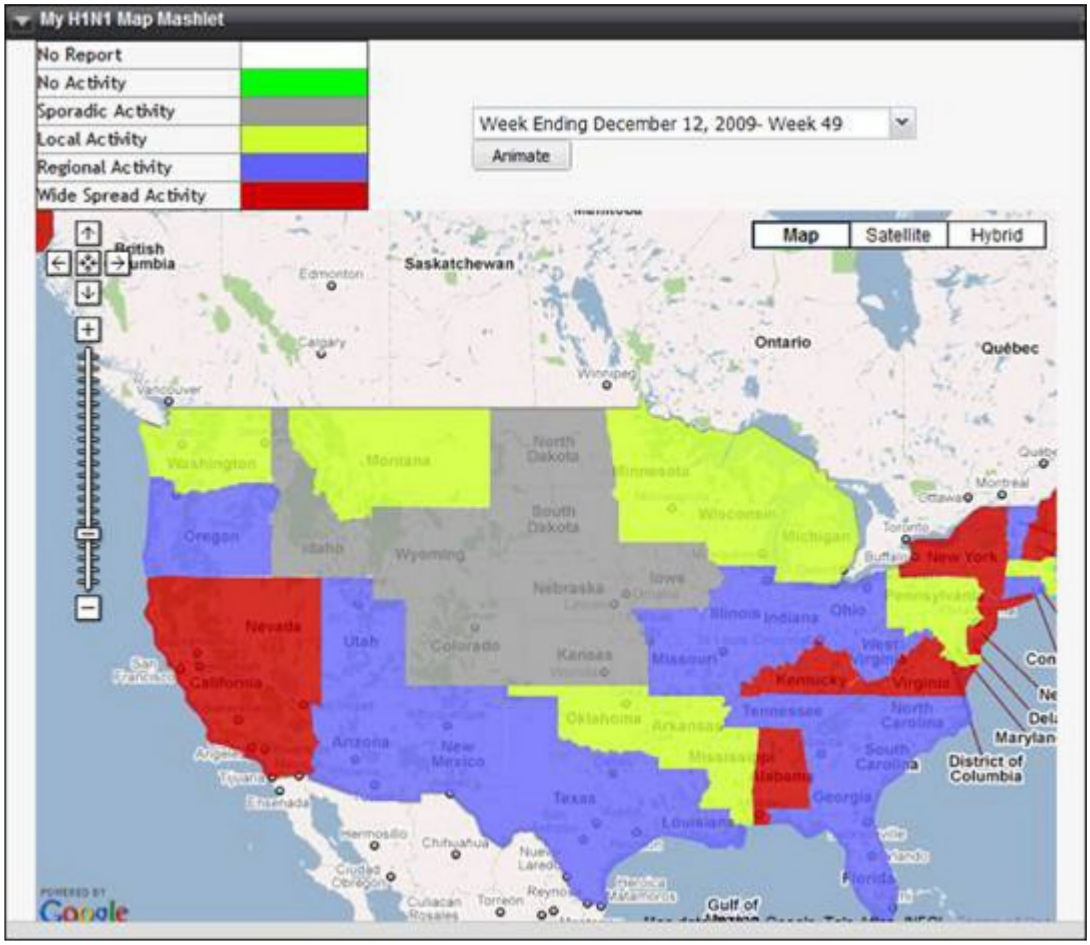

Content assembly/syndication: mashup-driven widgets of semi-structured content (pieces of text/documents with descriptive metadata, like pages/chapters of books) matched to user preferences that can be shared as well as embedded into web applications and used in mobile devices. This content mobility means greater reuse and a wider consuming audience.

Since mashups are being used in just about every major industry and across a wide variety of functions, the most compelling story is told by the broad range of applications.

In addition to the major trends above, other notables examples include:

Investment decision-support: a customizable watchboard for investment managers that consolidates investment pricing data, analysis, news, and other details necessary to make investment decisions more effectively. A sample Investment Decision-Support Mashup, created as a working prototype for the trading division of the largest diversified financial institution in Mexico (http://www.jackbe.com/case_studies/accival_study.php), is shown in Figure 3.

IT portfolio management: a mashup-driven portfolio management dashboard of key performance indicators and commentary/analysis made up of structured data from project plans/timelines, trouble ticket systems, and spreadsheets along with unstructured content from IT sources like blogs and wikis.

Figure 3: Investment Decision-Support Mashup

Data aggregation/redistribution: similar to content assembly/syndication, mashups can consolidate and republish data originating from multiple internal and external sources for use and/or license by third-parties, for use internally in new information products and externally by business partners for supply chain integration.

Credit management: similar to investment decision-support, mashups can provide consolidated information for analysts in order to provide better insight into loan recommendations and collections processing.

Metadata management: data quality can be improved by consolidating metadata from disparate databases and metadata repositories, improving data reuse and reducing software development/maintenance costs.

Single view of the customer: a single consolidated view of customer information can aid in customer service and product up-selling/cross-selling, ultimately increasing customer satisfaction, retention and revenue.

Technology Trends/Uses

Looking at mashups from a technology perspective, the research shows that mashups have some affinity with existing software infrastructure. For example, when potential mashers were asked "Which of the following is the MOST relevant to your interest in mashups?", the majority of respondents connected their enterprise mashup effort with Web user interface (UI) or Data Integration solutions. Figure 4 is based on 1073 responses.

Figure 4: Responders Interest in Mashups

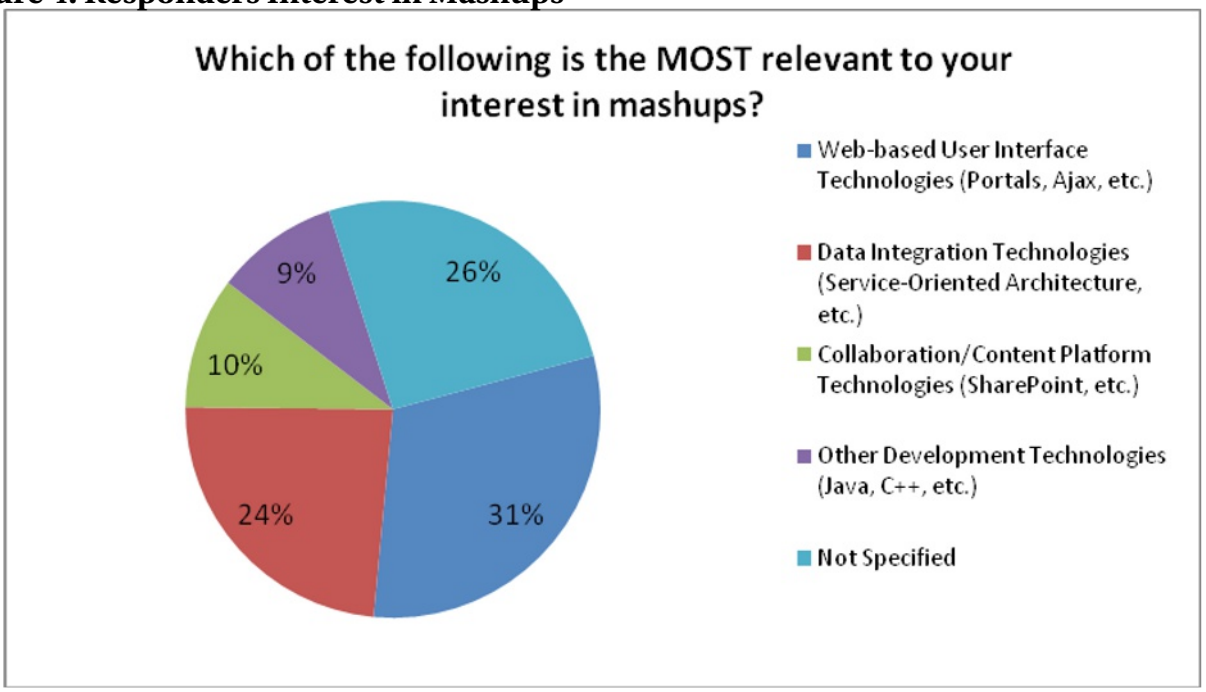

In all of these cases the software/application may be cloud-hosted or derive some or all of its data from cloud-based services.

When considering data-focused technologies such as service oriented architecture (SOA), mashup creators placed the most importance on the presentation component of the mashup that puts the data from those systems in a user-facing format. When considering user-centric technologies that already have user-facing elements such as portals, mashup creators placed more emphasis on the data-assembly/manipulation aspects of mashups.

The survey responses provided a number of enlightening insights on these synergies:

1. Web-based User Interface Technologies

- use mashups to expose data from enterprise resource planning (ERP) systems, combine the data into useful information, and publish this information on portals.
- ...we are using mashup technology to support interaction between data layers and various UIs.
- ...provide some simplified tools for the public to use to explore our data warehouse portal data...

2. Data Integration Technologies

- ...our SOA business services are finished...we'll be developing the presentation tier in enterprise mashups.

3. Collaboration Platform Technologies

- ...streamline and mash [Microsoft] SharePoint site with external data to make it a more information packed and dynamic environment...
- ...creating a SharePoint dashboard to represent data from disparate datasources including home-grown applications (Oracle-based), Siebel CRM, Primavera, Sabre, as well as feeds from news and stock sites.

4. Other Development Technologies

- ...mashup data provided to several application toolsets, providing coherence between them.

- ...an integrated reporting functionality from databases storing related and/or redundant data.

Conclusions

When asked "Why do you use mashups, what drives your interest?", one survey participant had this visionary response: "...to create and optimize Information Systems and Interactive Business Processes to support the growth of companies and organizations that need to collect, organize and share value added information with their clients, partners and suppliers using data from different sources."

The applications of enterprise mashups far exceed the humble consumer-driven beginnings of just a few years ago. Enterprise mashup applications are as varied as the users that use them, the industries they use them in, and problems they are applied to. Every data-dependent organization should consider this emerging technology when considering potential approaches to their dynamic information needs.

As enterprise-quality cloud-based data sources continue to grow and enterprise-grade cloud-hosting options mature, mashup creators will find it increasingly easier to establish genuinely meaningful mashups from external sources and on external platforms.

John Crupi is CTO at JackBe Corporation (http://www.jackbe.com). John has 20 years experience in OO and enterprise distributed computing. John spent eight years with Sun Microsystems, serving as a Distinguished Engineer and CTO for Sun's Enterprise Web Services Practice. He is co-author of the highly popular Core J2EE Patterns book (http://www.corej2eepatterns.com), has written many articles for various magazines and is a well-known speaker around the globe. He is a frequent blogger and was recently featured on Fox Business Network (http://jackbe.com/crupi_fox.php). John was also named as a member of the Software Development Magazine's Dream Team (http://ddj.com/architect/184415303) and a Washingtonian Magazine Tech Titan (http://ctovision.com/2009/05/washington-tech-titan-reception). John can be found on Twitter (http://twitter.com/johncrupi).

Chris Warner is Vice President of Marketing at JackBe and he co-manages JackBe's 3,800-member Mashup Developer Community (http://jackbe.com/dev). Chris has more than 18 years of experience in all types of enterprise software environments, private and public, big and small. He has published articles and blogs in many notable journals including SAP Netweaver Magazine, Ajax World Magazine, The SOA Magazine, and SOA World Magazine. Chris can be found on Twitter (http://twitter.com/chriswarner2000), LinkedIn (http://www.linkedin.com/in/chriswarner), Facebook (http://facebook.com/people/Chris_Warner/538484257), YouTube (http://youtube.com/user/jackbemashups), and a few trout streams in Montana.

March 31

Congratulations to OSBR Editor

Ottawa, ON

We congratulate our OSBR Editor, **Dru Lavigne**, for the launch of her latest book, The Definitive Guide to PC-BSD. The book was published by Apress on March 22, 2010 and is the first book on this open source, Unix-like, desktop friendly operating system. It contains everything you need to move from PC-BSD beginner to power user.

http://apress.com/book/view/1430226412

March 31

Congratulations to BBB

Ottawa, ON

We congratulate the contributors to Big Blue Button, particularly Fred Dixon and Richard Alam of Blindside Networks. Big Blue Button is an open source software that provides the web conference services required to deliver excellent educational experiences to learners in remote locations.

http://www.bigbluebutton.org

A Guide to Distributed Digital Preservation

Copyright: K. Skinner, M. Schultz

From the Abstract:

Authored by members of the MetaArchive Cooperative, A Guide to Distributed Digital Preservation is the first of a series of volumes describing successful collaborative strategies and articulating specific new models that may help cultural memory organizations work together for their mutual benefit. This guide is written with a broad audience in mind that includes librarians, archivists, scholars, curators, technologists, lawyers, and administrators. Readers may use this guide to gain both a philosophical and practical understanding of the emerging field of distributed digital preservation, including how to establish or join a network.

http://www.metaarchive.org/GDDP

2010 Horizon Report

Copyright: EDUCAUSE

From the Abstract:

The annual Horizon Report is a collaborative effort between the EDUCAUSE Learning Initiative (ELI) and the New Media Consortium (NMC). Each year, the report identifies and describes six areas of emerging technology likely to have a significant impact on teaching, learning, or creative expression in higher education within three adoption horizons: a year or less, two to three years, and four to five years. The areas of emerging technology cited for 2010 are: Time to adoption, One Year or Less: Mobile Computing and Open Content.

http://www.educause.edu/ELI/2010HorizonReport/195400

Managing External Innovation: The Case of Platform Extensions

Copyright: Nadia Noori

From the Abstract:

In recent years, high-technology industry witnessed the emergence of a product system where a focal company provides core components (a platform) and external parties provide complements. Increasing user demands for new and customized features product require companies to seek external sources of innovation. To align external innovation with company goals and protect the platform integrity, companies need to take a more active role in coordinating participation of external parties and controlling the quality of innovation outcome. This research examines strategies adopted by platform owners to manage the quality of platform complements developed by external parties. The research is focused around software platforms and one type of complement: platform extensions.

http://www.carleton.ca/tim/theses/2009/Noori2009.pdf

The Total Growth of Open Source

Copyright: Amit Deshpande, Dirk Riehle

From the Abstract:

Software development is undergoing a major change away from a fully closed software process towards a process that incorporates open source software in products and services. Just how significant is that change? To answer this question we need to look at the overall growth of open source as well as its growth rate. In this paper, we quantitatively analyze the growth of more than 5000 active and popular open source software projects. We show that the total amount of source code as well as the total number of open source projects is growing at an exponential rate. Previous research showed linear and quadratic growth in lines of source code of individual open source projects. Our work shows that open source is expanding into new domains and applications at an exponential rate.

http://dirkriehle.com/publications/2008/the-total-growth-of-open-source/

Top Threats to Cloud Computing

Copyright: Cloud Security Alliance

From the Abstract:

The purpose of this document, Top Threats to Cloud Computing, is to provide needed context to assist organizations in making educated risk management decisions regarding their cloud adoption strategies. In essence, this threat research document should be seen as a companion to Security Guidance for Critical Areas in Cloud Computing. As the first deliverable in the CSA's Cloud Threat Initiative, the "Top Threats" document will be updated regularly to reflect expert consensus on the probable threats which customers should be concerned about.

http://www.cloudsecurityalliance.org/topthreats.html

NEWSBYTES

March 8

Edmonton Takes Five Steps Towards Open Government

Edmonton, AB

The City of Edmonton is the fourth municipality in Canada moving ahead with open government initiatives rooted in the IT department. Five major announcements were made by Chris Moore, chief information officer of the City of Edmonton, at the Open City Workshop held in Edmonton, Alta. on Saturday.

http://itworldcanada.com/news/edmonton-takes-five-steps-towards-open-government/140161

April 13

DemoCamp Guelph

Guelph, ON

Rob Hyndman, a Toronto lawyer who focuses specifically on tech business, will be our invited speaker. Rob will be speaking about the main legal mistakes that startups make, and how to avoid them.

http://guestlistapp.com/events/15906

April 24

Open Data Hackfest

Ottawa, ON

We're planning a civic hacking fest happening in April, 2010 and you're invited. Ottawa is getting on board with open data, and we want apps that showcase how fantastic that is.

http://opendataottawa.ca/

May 4-5

BCNET/Canarie Conference

Vancouver, BC

The 2010 Conference theme, New Ideas for Shared Solutions, will explore innovative ideas for sharing technology services to solve common problems. Interact with and learn from mistakes and successes, discuss real-life situations, speak to panels of academic and industry technology experts and join conversations for viable solutions.

http://bc.net/2010-conference/

May 6

Startup Camp Montreal

Montreal, QC

This is the sixth edition of Startup Camp Montreal, an event dedicated to everything Startup. It is the forum where early stage companies, investors and onlookers alike can share information and validate pitches live. Join us as we learn from each other about the ins and outs of pitching a start up. This event is not just for Montrealers, all are welcome. Space is limited, so you need to register soon.

http://scmtl6.wikidot.com/

May 6-7

Future Play

Vancouver, BC

Future Play 2010 is the place to explore groundbreaking ideas alongside some of the world's most interesting games researchers and developers. Peer-reviewed researchers from universities around the world will be presenting their findings.

http://www.futureplay.org/

May 7-8

Northern Voice

Vancouver, BC

Northern Voice is a two-day, non-profit personal blogging and social media conference held at the UBC main campus.

http://2010.northernvoice.ca/

May 10-11

Canada 3.0

Stratford, ON

Help make Canada the world leader in digital media. Start establishing relationships with other visionaries, strategists, and entrepreneurs on the Canada 3.0 Community.

http://www.canada30.ca/

May 11-14

BSDCan

Ottawa, ON

BSDCan, a BSD conference held in Ottawa, Canada, has quickly established itself as the technical conference for people working on and with 4.4BSD based operating systems and related projects. The organizers have found a fantastic formula that appeals to a wide range of people from extreme novices to advanced developers.

http://www.bsdcan.org/2010/

May 12

Atlantic Entreneurship Expo

Charlottetown, PEI

An unbeatable lineup of keynote speakers, a tradeshow, and a networking lunch make up the full-day event. Hundreds of entrepreneurs come together to learn, grow, and share their success. Business relationships are made, new skills are learned, and deals are signed.

http://www.atlanticexpo.ca/cities/charlottetown

May 18-21

PGCon

Ottawa, ON

PGCon is an annual conference for users and developers of PostgreSQL, a leading relational database, which just happens to be open source. PGCon is the place to meet, discuss, build relationships, learn valuable insights, and generally chat about the work you are doing with PostgreSQL. If you want to learn why so many people are moving to PostgreSQL, PGCon will be the place to find out why. Whether you are a casual user or you've been working with PostgreSQL for years, PGCon will have something for you.

http://www.pgcon.org/2010/


MASSIVE INNOVATION
Lead to Win
Exciting new businesses that are in Lead to Win now!
selectstart
CheckinWiz
media
ucreatemedia.com
catwalk networks
SCTi
EclipseSource
Castel RESEARCH
zeebumobile
Videotelephony Inc.
NOVABRAIN TECHNOLOGIES Inc.
data bridge
A NEXIA SERVICES COMPANY
EnTeraSec security
Accentway
Triple Through Sustainability
SMART ROTOR SYSTEMS
SiliconPly
EventuSix
the speed of success
cornerportal
MARKTEK
Clearview
NewEraHealth
MarkOff Security

The goal of the Open Source Business Resource is to provide quality and insightful content regarding the issues relevant to the development and commercialization of open source assets. We believe the best way to achieve this goal is through the contributions and feedback from experts within the business and open source communities.

OSBR readers are looking for practical ideas they can apply within their own organizations. They also appreciate a thorough exploration of the issues and emerging trends surrounding the business of open source. If you are considering contributing an article, start by asking yourself:

1. Does my research or experience provide any new insights or perspectives?

2. Do I often find myself having to explain this topic when I meet people as they are unaware of its relevance?

3. Do I believe that I could have saved myself time, money, and frustration if someone had explained to me the issues surrounding this topic?

4. Am I constantly correcting misconceptions regarding this topic?

5. Am I considered to be an expert in this field? For example, do I present my research or experience at conferences?

If your answer is "yes" to any of these questions, your topic is probably of interest to OSBR readers.

When writing your article, keep the following points in mind:

1. Thoroughly examine the topic; don't leave the reader wishing for more.

2. Know your central theme and stick to it.

3. Demonstrate your depth of understanding for the topic, and that you have considered its benefits, possible outcomes, and applicability.

4. Write in third-person formal style.

These guidelines should assist in the process of translating your expertise into a focused article which adds to the knowledgable resources available through the OSBR.

Upcoming Editorial Themes

May 2010:	Communications Enablement
June 2010:	Growing Business
July 2010:	Go To Market
August 2010:	Interdisciplinary Lessons

Formatting Guidelines:

All contributions are to be submitted in .txt or .rtf format.

Indicate if your submission has been previously published elsewhere.

Do not send articles shorter than 1500 words or longer than 3000 words.

Begin with a thought-provoking quotation that matches the spirit of the article. Research the source of your quotation in order to provide proper attribution.

Include a 2-3 paragraph abstract that provides the key messages you will be presenting in the article.

Any quotations or references within the article text need attribution. The URL to an online reference is preferred; where no online reference exists, include the name of the person and the full title of the article or book containing the referenced text. If the reference is from a personal communication, ensure that you have permission to use the quote and include a comment to that effect.

Provide a 2-3 paragraph conclusion that summarizes the article's main points and leaves the reader with the most important messages.

If this is your first article, include a 75-150 word biography.

If there are any additional texts that would be of interest to readers, include their full title and location URL.

Include 5 keywords for the article's metadata to assist search engines in finding your article.

The Talent First Network program is funded in part by the Government of Ontario.

The Technology Innovation Management (TIM) program is a master's program for experienced engineers. It is offered by Carleton University's Department of Systems and Computer Engineering. The TIM program offers both a thesis based degree (M.A.Sc.) and a project based degree (M.Eng.). The M.Eng is offered real-time worldwide. To apply, please go to http://www.carleton.ca/tim/sub/apply.html.