

Article development led by **ACM** **queue**
queue.acm.org

The age of cloud computing has begun. How can companies take advantage of the new opportunities it provides?

BY MACHE CREEGER

CTO Roundtable: Cloud Computing

MANY PEOPLE READING about cloud computing in the trade journals will think it's a panacea for all their IT problems—it is not. In this CTO Roundtable discussion we hope to give practitioners useful advice on how to evaluate cloud computing for their organizations. Our focus will be on the SMB (small-to medium-size business) IT managers who are underfunded, overworked, and have lots of assets tied up in out-of-date hardware and software. To what extent can cloud computing solve their problems? With the help of five current thought leaders in this quickly evolving field, we offer some answers to that question. We explore some of the basic principles

behind cloud computing and highlight some of the key issues and opportunities that arise when computing moves from in-house to the cloud. Our sincere thanks to all who participated in the roundtable, and to the ACM Professions Board for making this event possible.

Participants

Werner Vogels is the CTO of Amazon.com, responsible for both e-commerce operations and Web services. Prior to working for Amazon he was a research scientist at Cornell University, studying large, reliable systems.

Greg Olsen is the CTO and Founder of Coghead, a platform-as-a-service (PaaS) vendor on both sides of the cloud equation. Coghead sells cloud-based computing services as an alternative to desktop or client/server platforms and is also a consumer of cloud services. The company built its entire service on top of Amazon's Elastic Compute Cloud (EC2), Elastic Block Storage (EBS), and Simple Storage Service (S3). Previously, Olsen founded Extricity, a company that provided business-to-business integration.

Lew Tucker is CTO of cloud computing at Sun Microsystems. In the 1980s he worked on the Connection Machine, a massively parallel supercomputer that sparked his interest in very large-scale computing. He spent 10 years at Sun as VP of Internet services running Sun's popular Web sites. Tucker left Sun to go to Salesforce.com, where he created AppExchange (<http://www.salesforce.com/appexchange/>), and afterward went to a start-up called Radar Networks. Recently he returned to Sun to lead its initiative in cloud computing.

Greg Badros is senior engineering director at Google, where he has worked for six years. Before that he was chief architect at Infospace and Go2Net. He earned his Ph.D. in constraint algorithms and user experiences from the University of Washington.


Geir Ramleth is CIO of Bechtel, where he provides cloud services for internal company use. Prior to his current



Mache Creeger



Lew Tucker



Steve Bourne



The Group



Geir Ramleth



Werner Vogels



Greg Badros



Greg Olsen

PHOTOGRAPHS BY MARJAN SADOUGHTI

job, Ramleth started a company inside Bechtel called Genuity, which was an early ISP and hosting company. Genuity was later sold to GTE.

Steve Bourne is CTO at El Dorado Ventures, where he helps assess venture-capital investment opportunities. Prior to El Dorado, Bourne worked in software engineering management at Cisco, Sun, DEC, and Silicon Graphics. He is a past president of ACM and chairs both the ACM Professions Board and the ACM Queue Editorial Board.

Moderator

Mache Creeger is principal of Emergent Technology Associates, where he provides marketing and business development enterprise infrastructure consulting for large and small technology companies. Beginning his career as a research computer scientist, Creeger has held marketing and business development roles at MIPS, Sun, Sony, and InstallShield, as well as various startups. He is an ACM columnist and moderator and head wrangler of the ACM CTO Roundtable series.

CREEGER: Let's begin the discussion with a general question and then dig down into some of the deeper issues. How would you define cloud computing?

TUCKER: Cloud computing is not so much a definition of a single term as a trend in service delivery taking place today. It's the movement of application services onto the Internet and the increased use of the Internet to access a wide variety of services traditionally originating from within a company's data center.

BADROS: There are two parts to it. The first is about just getting the computation cycles outside of your walled garden and being able to avoid building data centers on your premises.

But there's a second aspect that is equally important. It is about the data being in the cloud and about the people living their lives up there in a way that facilitates both easy information exchange and easy data analysis.

The great search tools available today are a direct result of easy access to data because the Web is already in the cloud. As more and more user data is stored in the cloud, there is a huge opportunity that transcends just computation being off-premises because there



LEW TUCKER

Cloud computing is not so much a definition of a single term as a trend in service delivery. It's the movement of application services onto the Internet and the increased use of the Internet to access a variety of services traditionally originating from within a company's data center.



is a relatively high-bandwidth connection to all those bits.

TUCKER: Tim O'Reilly's definition of Web 2.0 was that the value of data significantly increases when a larger community of people contributes. Greg [Badros]'s characterization complements that nicely.

VOGELS: It's not just data. I also believe that clouds are a platform for general computation and/or services. While telcos are moving their platforms into clouds for cost-effectiveness, they also see opportunities to become a public garden platform. In this scenario, people can run services that either extend the telco's services or operate independently. If, for example, you want to build an application that has click-to-call or a new set of algorithms such as noise detection in conference calls, then you can run those services connecting to the telco's platform. The key is having execution access to a common platform.

Because we have a shared platform, we can do lots of new things with data, but I believe we can do new things with services as well.

TUCKER: I see it as three layers: SaaS (software-as-a-service), which delivers applications such as Google Apps and Salesforce.com; PaaS (platform-as-a-service), which provides foundational elements for developing new applications; and IaaS (infrastructure-as-a-service), which is what Amazon has led with, showing that infrastructure can also be accessed through the cloud. I believe it is in this infrastructure layer—in which we've virtualized the base components of compute and storage, delivering them over the Internet—where we have seen the fundamental breakthrough over the past two years.

VOGELS: Understanding cloud computing requires a look at its precursors, such as SaaS before it became this platform-like environment; SOA (service-oriented architecture); virtualization (not just CPU virtualization but virtualization in general); and massively scalable distributed computing.

These were technologies that we needed to understand fully before cloud computing became viable. We needed to be able to provide these services at scale, in a reliable manner, in a way that only academics thought about 10 years ago. Building on this foundation, we

have now turned these precursors into the commercial practice of cloud computing.

TUCKER: A handful of companies, such as Amazon, Google, and Yahoo, demonstrated the advantage of very, very large scale by building specialized architectures just to support a single application. We have started to see the rest of the world react and say, “Why can’t we do that?”

BADROS: While I agree that the emergence of the massive scale of these companies plays a critical part, I also think that the development of client-side technology such as HTML, CSS, AJAX, and broadband connectivity is very important.

CREEGER: What about virtualization? It provides an encapsulation of application and operating system in a nice, neat, clean ABI (application binary interface). You could take this object and put it on your own premises-based hardware or execute it on whatever platform you choose. Virtualization makes execution platforms generic by not requiring the integration of all those horrible loose ends between the application and the operating system every time you want to move to a new machine. All that is required for a virtualized application/operating-system pair to execute on a new platform is for that platform to support the VM (virtual machine) runtime.

TUCKER: An important shift has been to use basic HTTP, in the form of REST (representational state transfer) APIs (http://en.wikipedia.org/wiki/Representational_State_Transfer), as an easier-to-use SOA framework. Everything that made services hard before, such as CORBA (<http://www.omg.org/getting-started/corbafaq.htm>) or IDL (http://en.wikipedia.org/wiki/Interface_description_language), went away when we said, “Let’s do it all over HTTP.”

BOURNE: Let’s be practical. What are the economics of clouds? What’s the CapEx (capital expenditure) and what is the OpEx (operational expenditure)? At the end of the year, did I spend more or less?

VOGELS: CapEx forces you to make massive investments. In the past, you had some measure of control over your customers; these days your customers have control over you. They know what to choose and have perfect informa-

tion. So if you build products today as an enterprise, but also as a young business, you have no idea whether you’re going to be successful or not. The less investment you have to make upfront, the better.

OLSEN: What inspired me about the cloud was that I could start a company and not buy any servers, phones, or software licenses. We were dedicated to using cloud services from day one. We started our company relying solely on services for email and the Internet and went from there to putting our source control on as a service. I wrote an article titled “Going Bedouin” where I expressed these views in more detail (<http://webworkerdaily.com/2006/09/04/going-bedouin/>).

BADROS: Clouds are clearly a huge win to get started with a business or product offering. At Google, we see internal people using the GAE (Google App Engine; <http://code.google.com/appengine/>) as a means of deploying something very quickly before they worry about scaling it on our base infrastructure. People do this because it is so much faster to get going, even inside Google where you have lots of infrastructure available.

Today’s developer has a decision to make: after I am a success, am I going to switch off of this initial platform? That’s the trade-off. Once it’s obvious that something like an Amazon S3 is able to outperform the best that the vast majority of companies can ever deploy, then it’s obvious you should just work entirely within the cloud. In this way you never have to suffer the replacement CapEx for the initial infrastructure.

VOGELS: For many customers, using our cloud products requires new expertise. You are no longer looking for a typical system administrator. If you have a large company, you’re looking for someone with the specific expertise to support 50,000 internal customers. Using the cloud, you no longer have to deal with things at the physical level. You no longer need to have people running around the data center replacing disks all day.

RAMLETH: You can get your smartest guys to work on what matters most rather than having them work on mundane stuff. That’s a huge benefit.

CREEGER: What about the people who need to run a flat-load, basic accounts receivable package? Once they get their

software and hardware in place and get their operational process down, it’s pretty straightforward and they can amortize the capital expenditure over a very long time period.

TUCKER: Every three years they’ve got to upgrade the software and the hardware.

RAMLETH: We spent \$5 million last year on an upgrade that did nothing for our business processes or end users. The software vendors told us that if we did not upgrade, they would stop supporting us.

OLSEN: I always wondered why we think software is so different from anything else. If a restaurant was growing its own food, slaughtering its own animals, generating its own power, collecting rainwater, and processing its own sewage, we would all think they were idiots for not using ready-made services. For a long time people built their own stack from the ground up, or ran their own servers because they could. Viewing the state of our industry, any student of economics will tell you that you have to start layering.

VOGELS: There are restaurants that do not buy their own herbs; they grow them on-site. They would argue that it contributes to the quality of the end product. They will never generate their own electricity, however, because that will not produce better food.

OLSEN: Realistically, however, software is really extreme in terms of how many people are doing undifferentiated tasks, on their own, at all kinds of levels. Look at the auto industry: there are many tiers of subcontractors, each providing specialized services and products. We just haven’t evolved to that same level of efficiency.

RAMLETH: We have dramatically reduced our data-center capital expenditures as a direct result of virtualization, allowing us to reuse our capital many more times than we ever could before. Before we started our effort, the average server utilization in our global server park was 2.3%. Going to virtualization has increased it to between 60% and 80%.

When we started, the core side of our central data centers, not including peripheral things, ran 35,000 square feet. Today, the equivalent of those 35,000 square feet is now operating in less than 1,000 square feet. We are utilizing

our hardware in very different ways than we could ever do before. The lesson we learned is that a very big part of building these public and private clouds is to be sure that you can get utilization factors significantly better than traditional company operations.

VOGELS: If you run your services inside the company, privately, utilization becomes an issue. It amortizes your costs over a number of cycles. If you run services outside, on a public service, it is no longer an issue for you.

RAMLETH: We are operating hundreds of servers that are processing data for projects that no longer exist and are no longer generating revenue. We do this because there may be a time and place where we would need this information, such as in a warranty situation.

Amazon taught us that we can move these programs from our data center to EC2, get them operational, capture that image, and then shut it down. At this point we have incurred very minimal costs. When conditions arise that require the execution of one of those programs, we can do it. By using Amazon EC2, we can transform what used to be a fixed cost of allocating a dedicated in-house server—regardless of whether we need the information—to a variable cost that is incurred only when the business case requires it.

The cost savings of using Amazon is quite compelling. A basic server, operating internally, that sits and does nothing costs us about \$800 to a \$1,000 per month to run. We can go to Amazon and get charged only for what we use, at a rate of 10 cents to 15 cents an hour.

TUCKER: This is the promise of utility computing. Users will be able to move their applications and their platforms off-site, and they will have more choices. There will be many different kinds of cloud service providers and, ultimately, opportunities for arbitrage. We are moving to a scenario where it will not matter where things execute, and where choosing an execution platform will be based on a number of different characteristics such as cost, security, performance, reliability, and brand awareness.

The great thing is that self-service has now moved into the provisioning of virtualized compute, storage, and networking resources. Without even talking to anybody at Amazon, you can use

its service with just a credit card. Enterprise customers are looking at their internal customers the same way. If the marketing department wants to run a new kind of application, traditionally you had to get the IT department to agree to help you build and deploy that application. Now IT departments are able to say, “You’ve got your own developers over in your area. If they want to develop and run this, fine, go ahead. Here are the policies for infrastructure services.”

BADROS: One of the key benefits is that not only is it easier to get going at start up, but also there is no discontinuity as things grow. It’s never the case that you are debating internally whether you should buy that extra server, invest in a more sophisticated infrastructure, or be able to scale to that second machine.

TUCKER: We need to be a little careful. Not all applications scale easily. While there is a whole class of applications that have very easy scaling characteristics, others do not. Databases are part of this class unless you are using something that has been set up to scale, such as Amazon’s SimpleDB (<http://aws.amazon.com/simpledb/>). If you’re running your own database, unless it has been designed to be scalable, don’t count on it happening.

CREEGER: How does that poor person sitting at a small- to mid-cap company make a decision to invest in clouds? What is he going to do next quarter or next year when the CEO comes in and says, “I read this thing in the *Wall Street Journal* stating that all the smart companies are going to cloud computing.” How is this guy going to respond?

OLSEN: First, adopt a philosophy of buy first, build second—even at the basic level of “I’m going to start a company, I need IT services.” Do I look to hire engineers and buy equipment or do I assume that there’s some outside service that might meet my needs? To me that’s half of it. I’m going to assume that services meeting my needs are already available or are going to evolve over time. I take a philosophy that says, “I’m all about my core business. I buy only infrastructure that directly supports my unique contributions to the marketplace.”

RAMLETH: I agree with you, if you are rational. However, you’re dealing with humans, and they are often not ratio-

nal. When a CEO goes down to his IT manager’s office and asks, “How are we utilizing cloud computing?”, the first thing that manager asks is, “What will this mean to me?” The biggest obstacle to change at our company was our own IT guys trying to protect their jobs. The change we have done at Bechtel has been 20% technology and 80% managing the change.

I believe an important part of your value proposition should be to explain to both the decision maker as well as the user how this tool enhances their professional futures. If it does not, those folks are going to be your obstacles.

TUCKER: There are certainly different approaches for different businesses at different points in their life cycles. A start-up has a certain set of needs. I completely agree with Greg [Olsen] to look for all the services that you can purchase before you think of building it yourself.

Animoto is a new company that makes movies out of photographs synced with music. It started with 50 instances running on Amazon. They launched it on Facebook and had very high success. In a matter of three days they went to 3,500 instances.

Can you imagine going to your IT department and saying, “We’re running on 50 servers today, and in two to three days we want to go to 3,500 servers”? It just would not have been possible.

CREEGER: So, for the zero- to a million-miles-an-hour overnight business plan that is stalled because of up-front CapEx costs, cloud computing is going to be your answer.

What other types of criteria can we give to people to evaluate how effective their internal IT infrastructure is in supporting business goals?

VOGELS: There are many first steps that corporations take into this world. Engineers can start by experimenting with these services, using them for small projects and comparing cost savings. I find that many of the first steps that enterprises take are just something small, easy, simple, and cost effective.

The *New York Times* scanned images covering a 60-year period in history and wanted to place them online. These guys moved four terabytes into S3, ran all the stuff on a Sunday, spent \$25, and got the product done.

BADROS: Replacing existing organi-

zation structure or IT functionality is harder in larger companies. Often you have a better chance of success if you introduce something that provides new value, perhaps by enabling a new type of collaboration, rather than replacing or modifying existing functionality. In this way you can avoid the risk of encountering resistance resulting from complexity or politics. In today's tougher economic times, you may also want to make your proposal more compelling by showing that operational TCO (total cost of ownership) can be significantly lowered when using a cloud.

OLSEN: The assumption that it's central IT making decisions about other technologies is wrong. Cloud computing has become successful not because a whole bunch of central IT groups proclaimed that cloud computing is good. Cloud computing has become popular from grassroots acceptance, from IT decisions made by small businesses, new providers, or at the departmental level. Cloud computing is coming into IT only at the end of this. My company does not sell to CIOs. We don't even try.

CREEGER: That's fine, but there are CIOs who will have to provide plans after their CEOs read that one can realize massive savings with cloud computing.

BOURNE: So who should pay attention to cloud computing?

OLSEN: I'm either a consumer of information technology needs: I need applications, I need storage; or I'm a producer: I'm somebody who's going to provide a service. Both of those audiences need to know what they can build from and how they can sell what they have. To me, it's not primarily about central IT. Central IT is an important constituent, but all these little system integrators, consultants, little ISVs, VARs—these are the folks who actually deploy computation on a broad scale to businesses and people. Any person who is in that space, either as a producer or a consumer of IT, needs to understand how to use cloud services.

BADROS: To me, the value proposition of cloud computing is so broad that the beauty of it is you can sell to almost anybody in the organization. Different aspects of the solution appeal to different sets of folks. Depending on whom I'm talking to, the story is different in order to let them see how it's going to be better for them.



WERNER VOGELS

If you run your services inside the company, privately, utilization becomes an issue. It amortizes your costs over a number of cycles. If you run services outside, on a public service, it is no longer an issue for you.



The individual who has been using consumer email and Google Calendar is excited about having the home experience at work and about the rich search capabilities and collaboration of Calendar. We see people using docs and spreadsheets to manage their wedding on the docs collaboration suite. Then when they are doing a similar type of project at work, they don't understand why they are stuck in early 1990s-style thinking with a set of applications that don't talk to one another. For that person, the collaboration story is the value proposition.

If an enlightened CIO comes to us and is wondering how this thing helps his organization, then cost of ownership, ease of scaling, and simplicity of starting new geographically distributed offices are really rich selling points.

To the CEO, it may be the fact that the IT department doesn't need to be as large as it is. The CEO is often scratching his head asking why he is spending 20% of his people budget just so the rest of his people can get their email. So, it really depends on the audience to understand what the best value proposition is. The beauty of cloud computing is there is a story for everyone—it's that compelling.

CREEGER: Does cloud computing enable new types of functionality that were not feasible under more traditional IT architectures?

VOGELS: In the past, I always thought that you could not build data warehouses out of general components. It's highly specialized, and I thought being really fine-grained precluded you from doing scatter-gather of lots of data operations. I think MapReduce (<http://labs.google.com/papers/mapreduce-osdi04.pdf>) has shown us that brute force works, and while it's not the most efficient approach, it allows you to get the job done in a very simple way.

A number of small companies now provide data warehousing as a service. The data movement is a little more inefficient than it used to be, but they're getting access to much smarter, much easier-to-use computational components.

It turns out that we have many customers who do not need a data warehouse 24-hours-a-day. They need it two hours a week. In the worst case, they're willing to spend a bit more on computational resources just to get these two

hours. They are still ahead on cost, given the alternative of having to purchase the hardware outright and build it up to support a peak load.

CREEGER: So, the analogy would be to analyze the cost of either purchasing a car or taking taxis to meet personal transportation needs?

VOGELS: Engineers are not well trained to think about end-to-end cost. MapReduce and other examples have shown us that the end-to-end picture of cost looks very different from what you would normally expect. We have to learn to think about the whole package—at storage, computation, and what the application needs to do—and really reason about what the axis of scale and cost really is.

CREEGER: I'd like to go around the room once and give some final recommendations to the folks who are struggling to try to make sense of all this.

RAMLETH: This is not a technology game but a change-management game. The goal is to get people to understand that it is not dangerous to think this way. We have three rules:

- ▶ Think about what you can do that can benefit service delivery in aggregate; don't focus on the small subcomponents that can lead to suboptimal solutions.

- ▶ Don't think about how you're going to distribute your costs before you start any effort. Make sure that internal charging mechanisms (allocations) are not obstacles for change and progress.

- ▶ Don't think about and design future organization changes. Base decisions on organizational benefit and not on increased power to you as a manager or to your organization.

If you think about these three things, it's amazing what an organization can actually do.

BADROS: The beauty of what we're talking about is that it's so easy to try. You don't need a big budget or approvals to get started. The fact that you can do this so simply enables innovation that would be unavailable if you needed to purchase a big piece of hardware ahead of time.

TUCKER: As services move into the Internet, they become easier and more cost effective. This also means a shift in power in IT away from those who control capital resources to the users and developers who use self-service to pro-



GREG OLSEN

Cloud computing presents a compelling opportunity for consumers of information technology and producers of information services.



vision their own applications. When FedEx went online, people were taken out of the support loop and customers could find their package status information themselves whenever it was needed. You can now apply the same principle to the provisioning of computing resources. A developer can have a server provisioned to run an application without having to contact a human. That cuts the most costly aspect of computing out of the equation.

OLSEN: Cloud computing presents a compelling opportunity for consumers of information technology and producers of information services. Application builders should take advantage of existing functionality they can buy as opposed to the past practice of building their own and focus their resources on the unique capability they alone can deliver. Consumers of information technology have got to rethink where they look for functionality. If they don't adapt their service delivery models, then they will quickly become obsolete.

CREEGER: Reducing cost and enabling overall agility are what I believe you all are trying to say. Cloud computing has the potential for removing business friction to make more services possible and to do so much more easily, with less risk and capital outlay. I think that is as good a summary as any for something as transformative as cloud computing. Thank you all very much for your time, talent, and wisdom. □

Related articles on queue.acm.org

For the complete version of this CTO Roundtable discussion, visit

<http://queue.acm.org/detail.cfm?id=1551646>

Describing the Elephant

Ian Foster and Steven Tuecke

<http://queue.acm.org/detail.cfm?id=1080874>

Enterprise Software as Service

Dean Jacobs

<http://queue.acm.org/detail.cfm?id=1080875>

CTO Roundtable: Virtualization

Mache Creeger (moderator)

<http://queue.acm.org/detail.cfm?id=1400229>

Mache Creeger (mache@creeger.com) is a technology industry veteran based in Silicon Valley. Along with being a columnist for *ACM Queue*, he is the principal of Emergent Technology Associates, marketing and business development consultants to technology companies worldwide.

© 2009 ACM 0001-0782/09/0800 \$10.00