

## ***Harnessing cloud technologies -- Many of the Technical Underpinnings Are Already in the Data Center Today***

*Transcript of a BriefingsDirect Podcast examining how enterprises are increasingly focused on delivery and consumption of cloud based infrastructure and services.*

[Listen](#) to the podcast. [Download](#) the podcast. Find it on [iTunes/iPod](#) and [Podcast.com](#). Download the transcript. Learn more. Sponsor: [Hewlett Packard](#).

Offer: Get a free copy of Cloud for Dummies courtesy of Hewlett Packard by going to: [www.hp.com/go/cloudpodcastoffer](http://www.hp.com/go/cloudpodcastoffer)

**Dana Gardner:** Hi, this is [Dana Gardner](#), principal analyst at [Interarbor Solutions](#), and you're listening to BriefingsDirect.



Today, we present a sponsored podcast discussion on quickly harnessing the technical benefits of [cloud computing](#) approaches. We will examine how enterprises are increasingly focused on delivery and consumption of cloud-based infrastructure and services.

But, we'll look at how many of the technical underpinnings of cloud are available now for organizations to leverage in their in-house [data centers](#), whether it's moving to highly scalable servers and storage, deeper use of virtualization technologies, improved management and automation for elastic compute provisioning, or service management and [governance](#) expertise. Much of what makes the cloud tick is already being used inside of many data centers today.

We expect that the way the clouds are built will be expanding into more and more enterprises overtime. The goal is gaining the efficiency, control and business benefits of an [everything-as-a-service](#) approach, without the downside and risks.

The interest in cloud adoption is being fueled by economics, energy concerns, skills, shortages, and complexity. But, getting the best paybacks from cloud efforts early and often and by bringing them on premises, can help prevent missing the rewards of cloud models later by being unprepared or inexperienced now.

Here to help us better understand how to make the most of cloud technologies, are four experts from [Hewlett-Packard \(HP\)](#). Please join me in welcoming [Pete Brey](#), worldwide marketing manager for HP [StorageWorks](#) group. Welcome to the show, Pete.

**Pete Brey:** Thank you.

**Gardner:** We're also joined by [Ed Turkel](#), manager of business development for HP [Scalable Computing and Infrastructure](#). Welcome, Ed.

**Ed Turkel:** Thank you.

**Gardner:** We are also joined by [Tim Van Ash](#), director of [software as a service \(SaaS\)](#) products in the HP Software and Solutions group. Welcome, Tim.

**Tim Van Ash:** Hi, Dana.

**Gardner:** And also [Gary Thome](#), chief strategist for infrastructure software and blades at HP. Welcome to the show, Gary.

**Gary Thome:** Thank you very much.

**Gardner:** Ed, let me take the first question to you. HP has been a supplier of the picks and shovels, if you will, to cloud service providers for many years. As we're starting to take these technologies to the enterprise for their requirements around scale, lower cost, flexibility and efficiency, what are we talking about when we discuss cloud? What comprises cloud for these enterprises and how they are adapting to it?

### *What do we mean by cloud?*

**Turkel:** The first thing is when we talk about cloud, what do we mean? What is our definition of cloud? We like to talk about the cloud as a means by which global-class, highly scalable, and flexible services can be delivered and consumed over the Internet on an as needed and pay-per-use business model. This enables new access, new capabilities, and new connections.



As a scalable computing and infrastructure organization, we have been selling to the major cloud providers for the last few years, we have been seeing this major trend towards a style of scale-out computing that they are then delivering as a service to their customers.

This is causing some significant trends simply in the way that they internally themselves deploy IT. First of all, they are building very large environments. So, when we talk about scale-out, we're talking about extreme scale-out. We are talking about numbers of servers, not in the tens or hundreds, but in the thousands, and often tens of thousands, within a single computing environment. We are talking about volumes of storage that go well beyond petabytes of storage, again in a single environment.

That creates challenges in the data centers in which they are deploying in terms of an almost pathological focus on power and cooling, because if you're putting together an environment that

large, every penny-per-kilowatt has a huge impact on the [return on investment \(ROI\)](#) for those environments.

What we've seen happen over the years, as the cloud providers themselves have been building out their environments, is that their customers are looking at the cloud providers and thinking to themselves, "If these guys can do it, if these guys can get some great benefits on reduced costs, on improved power efficiency, on increased agility, through their computing environments, why can't I?"



So, we're starting to see enterprise customers who are looking at the cloud providers themselves as sort of a best of breed kind of IT environment and they're starting to look at how can they emulate this within their own environments. Thus they are saying, "Why can't we do that? How can we buy the same environments? What else should we build out to be able to get those kinds of advantages?"

**Gardner:** Now, we talked about some of the economic impetus around this. Tim Van Ash, there is also, of course, the simultaneous trend in the business around converting IT to an IT service or a managed service organization. As someone who has been dealing with SaaS for sometime, does moving a cloud technology to the enterprise, dovetail well with this whole notion of a service provider role for IT?

**Van Ash:** When you look at becoming a service provider, technology is a key part of it, architecting yourself to be able to support the service levels around delivering a service, as opposed to some of the more traditional ways that we saw IT evolve. Then, applications were added to the environment, and the environment was expanded, but it wasn't necessarily architected around the application in all cases.



Another thing is that, when they move to a service provider role, it's as much about how they structure their organization to be able to deliver those services. That means being able to not only have the sort of operational teams that are running and supporting the application, but also have the customer-facing sides, who are managing the business relationships, whether they would be internal or external customers, and actually starting to run it as if it were a business. So, what is the profit and loss statement for a particular service?

**Gardner:** I suppose that when you need to run it on a profit-and-loss basis, that every bit of efficiency counts, which is a little different from the previous models, right?

### ***Not just a cost model***

**Van Ash:** It is, and it's also about realizing that it's not just a cost model, but it is very much a business model. That means you need to be actively out there recruiting new customers. You need to be out there marketing yourself. And, that's one area that IT traditionally has been quite poor at -- recognizing how to structure themselves to deliver as a business.

The technology is really one of the key enablers that come into that and, more importantly, enables you to get scale and standardization across the board, because one of the issues that IT has traditionally faced is that often architecture is forced on them, based on the application selection by the business.

When you start to move into cloud environments, which feature, in many cases, high levels of [virtualization](#), you start to decouple those layers, as the service provider has a much stronger control over what the architecture looks like across the different layers of the stack. This is really one of the areas where cloud is hoping to accelerate this process enormously.

**Gardner:** Another unfortunate reality today is the lack of dollars. Discretionary spending has pretty much evaporated in many organizations. So for enterprises to move towards these cloud technologies, I would think it needs to be a very rapid return.

Let me take this to Pete Brey. Storage, of course, is a very high-cost area. I would think that moving to the cloud on the storage level might be a strong economic story, at least in terms of ROI.

**Brey:** Absolutely, and that is indeed one of the key things that we are looking at in HP StorageWorks, developing and delivering to market new classes of scale-out storage. Now, not only do you have your scale-out compute environments, you need to also pay attention to the storage piece of the equation and delivering the platforms. The storage platforms need not only to scale to the degree that we talk about into the [petabyte](#) ranges, but they also need to be very simple and easy to use, which will drive down your total cost of ownership and will drive down your administrative costs.



They also deliver a fundamentally new level of affordability that we have never really seen before in the storage marketplace in particular. So these combinations of things, scalability, manageability, ease of use and overall affordability, are driving what I consider almost a revolution in the storage marketplace these days. We're working on a lot of different things in the StorageWorks group at HP to deliver on all four of those capabilities.

**Gardner:** I've heard in many places recently that folks refer to [business intelligence \(BI\)](#) as the “killer application” for cloud. I would think that those petabyte-scale [warehouses](#) are a key focus for you. Is that the case?

**Brey:** Absolutely, that's the case. That's one of the prime application areas that we hear, as we talk to different customers, but that's not the only area. We see explosive data growth across the wide range of market segments. This includes everything from the traditional Web-based service providers to the communications, media, and entertainment industries, where they move towards higher and higher definition formats.

## ***Explosion in content***

It's driving this explosion to the medical field, where new innovations are happening in that particular space that are also driving an explosion in content. So, it's all of these factors coming together, and people are demanding new levels of scalability and affordability that are driving these types of storage platforms to support cloud environments.

**Gardner:** Gary Thome, is there a similar story, when it comes to the infrastructure that supports these cloud fabric and service fabrics? Is there an ROI story here as well?

**Thome:** Definitely. Very much so. Certainly, when customers are thinking about going to a cloud infrastructure or shared-service model, they really want to look at how they are going to get a payback from that. They're looking at how they can get applications up and running much faster and also how they can do it with less effort and less time. They can redirect administrative time or people time from just simply getting the basic operations, getting the applications up and running, getting the infrastructure up and running for the applications, to doing more innovative things instead.



Customers are looking for those things, as well as the cloud model, a shared-services platform, to be able to get higher utilization out of the equipment. So, they definitely look for those kinds of ROI.

**Gardner:** Ed Turkel, is there a different sales approach in the enterprise? Someone mentioned earlier that so much of IT has followed on from the applications, but when we think about the architecture of a cloud, we are really thinking about an abstraction of infrastructure that applications can be deployed to and we can get provisions and better efficiency out of. Do you have to go to these enterprises at a different level to sell this? What is the difference between selling to an enterprise and a service provider?

**Turkel:** It's definitely selling in a different model. First of all, the approach to selling is much more of a holistic view of the IT environment and selling a broader solution, than simply going in and selling a server with some storage and so on for a particular application. It tends to touch a broader view of IT, of the data center, and so on.

As was discussed in some of the other comments a moment ago, it has to look at working with the CIO or senior staff within the enterprise IT infrastructure, looking fundamentally at how they change their model of how they deliver their own IT service to their internal customers.

Rather than just providing a platform for an application, they are looking at how they provide an entire service to their customer base by delivering IT as a service. It's fundamentally a different

business model for them, even inside their own organizations. So absolutely, it's a completely different way of selling.

**Gardner:** Pete Brey, how does this notion of architecture sale, rather than a technology sale, affect the storage business?

### *Profound effects*

**Brey:** It has very profound effects in terms of the end-to-end application that the customer is using and understanding the unique requirements of those applications and how that gets driven down into the technology that supports those requirements. So, it's a fundamental shift in the way we think about it and the solutions that we deliver from a storage standpoint into the marketplace.

**Gardner:** Tim Van Ash, management, of course, is a crucial part of this. But, we're going to be managing, many of us analysts predict, across heterogeneous environments of on-premises, delivered cloud services, traditional legacy services and applications, and then the third-party, outside applications.

As enterprises consider these technologies, it seems to me important to consider how you would manage them not just on their own, but in the context of a larger cloud ecology.

**Van Ash:** The thing that we're seeing from our customers is how they extend enterprise control in the cloud, because cloud has the potential to be the new silo in the overall architecture. As you said, in a heterogeneous environment, you potentially have multiple cloud providers. In fact, you almost certainly will have a multi-sourced environment.

So, how do you extend the capabilities, the control, and the governance across your enterprise in the cloud to ensure that you are delivering the most agile and the most cost effective solution, whether it would be in-house or leveraging cloud to accelerate those values?

What we're seeing from customers is a demand for existing enterprise tools to expand their role and to manage both private cloud and public cloud technologies. One of the big steps that HP has taken this year is enabling both of the services. The Software-as-a-Service Group delivers IT management as a service, which can manage both your private cloud capabilities and your public cloud capabilities, and manage the security performance and service-level aspects around both your internal and your external consumption.

**Gardner:** Ed Turkel, when we think about taking these technologies from what had been a service provider environment into enterprises, I think the requirements on service providers are often higher than enterprises are accustomed to, in terms of availability and reliability. Is this proving a benefit that they recognize? What's the transition, in terms of the management and requirements around performance?

**Turkel:** In those environments, the way that they look at management of the environment, the resilience or reliability of individual servers, storage, and so on, is done a little differently, partially because of the scale of the environments that they are creating.

If you look at many of the cloud providers, what they've done is they've implemented a great deal of resilience in their application environment, in a sense, moving the issues of resiliency away from the hardware and more into software. When you look at an environment that is as large as what they are doing, it's somewhat natural to expect that components of that environment will fail at some level of frequency. If you have tens of thousands of servers, or tens of thousands of disk drives, some number will fail on a somewhat regular basis.

### ***Resiliency capabilities***

So, their software infrastructure has to be able to deal with that. Many of the very largest of the cloud providers have implemented resiliency capabilities into their software infrastructure to allow for that. It fundamentally changes things, because of the nature of the scale of the environment. It also changes the way that we work with those same folks in terms of how we provide things like technical services and break-fix services into those environments.

You start looking at technical service from a different viewpoint. You don't send a field service engineer into those environments every time a component fails. You do it more on a scheduled basis or, in many instances, some of those customers do their own maintenance and simply maintain a parts depot within their environment to get replacement parts. Again, it's fundamentally different because of the scale that they are operating at.

**Gardner:** Well, what's interesting to me is that we can take what is an expectation and requirement in a business-to-consumer environment, like a service provider deals with, and can apply that now to a business-to-business type of applications or requirements, but you couldn't do vice-versa.

**Turkel:** No, I think it does go somewhat in both directions. Enterprise IT environments, as they are consolidating their environments into a single large infrastructure, rather than the application silos we touched on a little bit earlier, they are dealing with some of the same issues of scale. The way that they service and the way that they design the environment has to be somewhat similar to those cloud providers.

But then, they are delivering all of that as a service to their customer. So, as you say, it becomes more of a business-to-consumer way of delivering their services rather than, as you suggested, the business-to-business model, or a less direct non-service oriented approach to doing it.

**Gardner:** Let's look at some examples of where HP has brought some of these technologies into enterprises and what some of the paybacks have been, I don't know whether you can name

companies, but maybe a use-case scenario. Pete Brey, can you provide an example on what some of the paybacks have been?

**Brey:** Absolutely. In fact, there is a very notable example that we announced this past summer, a partnership that we've developed with [DreamWorks Animation](#). DreamWorks is using HP storage to host their animation environments, and this would be an example of an enterprise building up a cloud-based environment.

They have multiple locations. When they're working on a film, they have animators spread across geographic boundaries, across countries and continents. They have a need to virtualize those environments into an enterprise cloud-like setting for their animation environments. They are building this solution, as we speak, using HP components, HP servers, HP storage, and software to link it altogether.

For them, it's really a great opportunity to evolve their infrastructure to meet some of the new requirements that they have around high-definition content and also around rapidly increasing their productivity, in terms of the number of films that they can turn out in a given amount of time. In the not-too-distant past, they were able to produce two, three, maybe four films a year, where now they have been able to double that.

The technology that HP has been able to provide to them has helped them significantly in achieving those levels of productivity. So, it's really an exciting relationship with DreamWorks. And, they are very excited to be working with us too, helping us drive our own cloud strategies around things like key-based storage archive systems, some really new and innovative features that are going to make storage and compute environments even simpler to use in these cloud environments.

**Gardner:** Gary, what about some of your products and strategies for applying to enterprises? Is there a Matrix story in terms of examples of undergirding cloud-type environments?

### ***Cloud-like experience***

**Thome:** Yes, very much so. [BladeSystem Matrix](#) is designed to help customers, provide a cloud-like experience for their enterprise applications.

For many enterprises, unlike the cloud that Ed was talking about earlier where they are able to put things like the resilience and scalability into the software, many enterprises don't own all their applications, and there are a variety of different applications on a variety of different operating systems.

So, they really need a more flexible platform that gives them an abstraction between the applications and the hardware itself. Products like BladeSystem Matrix, with technologies such

as our [Insight Orchestration](#) and our [Virtual Connect](#) technology, allows customers to get that abstraction.

They can turn on applications very quickly, and then be able to scale them up and scale them down very quickly as well, without having to rely on specialized software to do it. The servers themselves are doing it.

We've got one company, [Micros-Fidelio](#) which itself is a service provider in the hospitality industry. They have a need to be able to stand up applications very quickly for their customers. Technology, such as Insight Orchestration, gives them the capability to do that very quickly.

**Gardner:** Tim Van Ash, do you have any examples of the use of these technologies in the enterprise environments?

**Van Ash:** From HP Software's perspective, this has been a core business of ours for some time and there are numerous examples. One of the most exciting examples that I have seen recently has been taking the enterprise technology around provisioning of both physical and virtual servers in a self-service and a dynamic fashion and taking it to the service provider.

[Verizon](#) recently announced one of their cloud offerings, which is Compute as a Service, and that's all based on the business service automation technology that was developed for the enterprise.

It was developed to provide data-center automation, providing provisioning and dynamic provisioning to physical and logical servers, networks, storage, and tying it altogether through run book automation, through what we call Operations Orchestration.

Verizon has taken that technology and used that to build a cloud service that they are now delivering to their customers. So, we're seeing service providers adopting some of the existing enterprise technology, and really taking it in a new direction.

**Gardner:** What's interesting, along the lines of what Ed Turkel was saying, is that this is a two-way street where you can apply underlying cloud fabric. That's a fascinating observation -- that is to say, between the types of technologies we would expect in an enterprise IT environment and the types that we would expect in a service provider environment.

### ***Significant changes***

**Van Ash:** While we are seeing some significant changes in both the economics model and the scale, in many ways, cloud is really building on a series of innovations that we have been seeing for some time, as IT move towards more of the utility type model around this.

It's utility, both in terms of being able to take a power cord and plug it into the socket, but also utility in the sense that you are enabling customers to do many of the things that, once upon a time, would require them to open a ticket and have teams of people manually working on their activities in the background. Now, they can do this in a self-service fashion that really ties all these processes together in an automated way.

So, while cloud is currently going in a very exciting direction, it really represents an evolution of many of the technologies that we at HP have focused on now for the last 20 years.

**Gardner:** It sounds almost as if cloud computing, as a vision, is providing somewhat of a unifying theory around many of the different aspects of computing and technology development over the past decades. A unifying theory is something, of course, has been elusive in the realm of physics.

Okay, Ed Turkel, on this notion of an example, do you have any of the use-case scenarios or actual companies that you could offer in terms of this trend?

**Turkel:** Well, we're somewhat challenged in being able to talk about some of the leading cloud providers that we're actually selling to, because virtually every one of them will not allow us to talk about them for the fundamental reason that their IT infrastructure is part of their unique value add and part of their value proposition to their own customers. So, it is very competitive within each of those environments. They tend not to let us mention them by name.

But, if you look across the set of customers that we talk to, for example, we have one that's a leading email house. Another is a leading social networking company, and so on. I can't name names and I can't tell you exactly how they're using our systems, but some of those environments are again very, very large.

We're also seeing some interesting crossover from another part of our market that has been very traditionally a scale-out market. That's the high-performance computing or technical computing market, where we are seeing a number of large sites that have been delivering technical computing as a service to their customers for some time, way back when they called it time sharing. Then, it became utility computing or grid, and so on.

Now, they're more and more delivering their services via cloud models. In fact, they're working very closely with us on a joint-research endeavor that we have between HP Labs, Yahoo, and Intel called the cloud Testbed, more recently called the [Open Cirrus Project](#).

### ***Model is expanding***

It's where some of our largest HPC customers are implementing their scale-out environments as cloud services where they are offering high performance computing environments as a service to enterprise customers, to academic customers, and so on, over the Internet using that same cloud

model. We're seeing this model expanding, and beyond just those big cloud providers into some of those traditional HPC environments.

**Gardner:** I'm afraid we'll have to leave it there. We've been discussing how technologies that have supported cloud, utility, and service provider infrastructure for years, are beginning to work their way into enterprises under the category of cloud computing but giving them some technical underpinnings for new business models, approaches, and efficiencies.

To help us discuss this, we've been joined by Pete Brey, worldwide marketing manager for HP StorageWorks group. I appreciate your input, Pete.

**Brey:** Thank you.

**Gardner:** We were also joined by Ed Turkel, manager of business development for HP Scalable Computing and Infrastructure. Thanks.

**Turkel:** Thank you.

**Gardner:** And, Tim Van Ash, director of SaaS products at the HP Software and Solutions group. Thank you, Tim.

**Van Ash:** Thanks very much, Dana.

**Gardner:** And also, Gary Thome, chief strategist for infrastructure software and blades. Thank you, Gary.

**Thome:** Thanks for the time.

**Gardner:** This is Dana Gardner, principal analyst at Interarbor Solutions. You have been listening to a sponsored BriefingsDirect podcast. Thanks for listening and come back next time.

[Listen](#) to the podcast. [Download](#) the podcast. Find it on [iTunes/iPod](#) and [Podcast.com](#). Download the transcript. Learn more. Sponsor: [Hewlett Packard](#).

Offer: Get a free copy of Cloud for Dummies courtesy of Hewlett Packard by going to: [www.hp.com/go/cloudpodcastoffer](http://www.hp.com/go/cloudpodcastoffer)

Transcript of a BriefingsDirect Podcast examining how enterprises are increasingly focused on delivery and consumption of cloud based infrastructure and services. Copyright Interarbor Solutions, LLC, 2005-2009. All rights reserved.