

## ***Business and Technical Cases Build for Data Center Consolidation, Modernization, Centralization***

*Transcript of a sponsored BriefingsDirect podcast on how data center consolidation and modernization helps enterprises reduce cost, cut labor, slash energy use, and become more agile.*

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

**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 how [data-center consolidation](#) and [modernization of IT systems](#) helps enterprises reduce cost, cut labor, slash energy use, and become more agile.



We'll look at the business and technical cases for reducing the numbers of enterprise data centers. Infrastructure advancements, standardization, performance density, and [network services efficiencies](#) are all allowing for [bigger and fewer data centers](#) that can carry more of the total IT requirements load.

These [strategically architected](#) and located facilities offer the ability to seek out best long-term outcomes for both performance and cost -- a very attractive combination nowadays. But, to gain the big payoffs from fewer, bigger, better data centers, the essential list of user expectations for performance and IT requirements for reliability need to be maintained and even improved.

[Network services](#) and Internet [performance management](#) need to be brought to bear, along with the latest data-center advancements to produce the full desired effect of topnotch applications and data delivery to enterprises, consumers, partners, and employees.

Here to help us better understand how to get the best of all worlds -- that is high performance and lower total cost from data center consolidation -- we're joined by our panel. Please join me in welcoming [James Staten](#), Principal Analyst at [Forrester Research](#). Welcome, James.

**James Staten:** Thanks for having me.

**Gardner:** We're also joined by [Andy Rubinson](#), Senior Product Marketing Manager at [Akamai Technologies](#). Welcome, Andy.

**Andy Rubinson:** Thank you, Dana. I'm looking forward to it.

**Gardner:** And, [Tom Winston](#), Vice President of Global Technical Operations at [Phase Forward](#), a provider of integrated data management solutions for clinical trials and drug safety, based in Waltham, Mass. Welcome, Tom.

**Tom Winston:** Hi, Dana. Thanks very much.

**Gardner:** Let me start off with James. Let's look at the general rationale for data-center modernization and consolidation. What are the business, technical, and productivity rationales for doing this?

### *Data-center sprawl*

**Staten:** There is [a variety of them](#), and they typically come down to cost. Oftentimes, the biggest reason to do this is because you've got sprawl in the data center. You're running out of power, you're running out of the ability to cool any more equipment, and you are running out of the ability to add new servers, as your business demands them.



If there are new applications the business wants to roll out, and you can't bring them to market, that's a significant problem. This is something the organizations have been facing for quite some time.

As a result, if they can start consolidating, they can start moving some of these workloads onto fewer systems. This allows them to reduce the amount of equipment they have to manage and the number of software licenses they have to maintain and lower their support costs. In the data center overall, they can [lower their energy costs](#), while reducing some of the cooling required and getting rid of some of those power drops.

**Gardner:** James, isn't this sort of the equivalent of Moore's Law, but instead of at silicon clock-speed level, it's at a higher infrastructure abstraction? Are we [virtualizing our way](#) into a new [Moore's Law](#) era?

**Staten:** Potentially. We've always had this gap between how much performance a new CPU or a new server could provide and how much performance an application could take advantage of. It's partly a factor of how we have designed applications. More importantly, it's a factor of the fact that we, as human beings, can only consume so much at so fast a rate.

Most applications actually end up consuming on average only 15-20 percent of the server. If that's the case, you've got an awful lot of headroom to put other applications on there.

We were [isolating applications on their own physical systems](#), so that they would be protected from any faults or problems with other applications that might be on the same system and take them down. Virtualization is the primary isolating technology that allows us to do that.

**Gardner:** I suppose there are some other IT industry types of effects here. In the past, we would have had entirely different platforms and technologies to support different types of applications, networks, storage, or telecommunications. It seems as if more of what we consider to be technical services can be supported by [a common infrastructure](#). Is that also at work here?

## *Unique opportunity*

**Staten:** That's mostly happening as well. The exception to that rule is definitely applications that just can't possibly get enough compute power or enough contiguous compute power. That creates the opportunity for unique products in the market.



More and more applications are [being broken down into modules](#), and, much like the [web services](#) and [web applications](#) that we see today, they're broken into tiers. Individual logic runs on its own engine, and all of that can be spread across some more monetized, consistent infrastructure. We are learning these lessons from the dot-coms of the world and now the [cloud-computing](#) providers of the world, and [applying them to the enterprise](#).

**Gardner:** I've heard quite a few numbers across a very wide spectrum about the types of payoffs that you can get from consolidating and modernizing your infrastructure and your data centers. Are there any rules of thumb that are typical types of paybacks, either in some sort of a technical or economic metric?

**Staten:** There's a wide range of choices from the fact that the benefits come from how bad off you are when you begin and how dramatically you consolidate. On average, across all the enterprises we have spoken to, you can realistically expect to see about a 20 percent cost reduction from doing this. But, as you said, if you've got 5,000 servers, and they're all running at 5 percent utilization, there are big gains to be had.

**Gardner:** The economic payoff today, of course, is most important. I suppose there is a twofold effect as well. If you're facing a capacity issue and you're thinking about spending \$40 or \$50 million for an additional data center, and if you can reduce the need to do that or postpone it, you're saving on capital costs. At the same time, you could, perhaps [through better utilization](#), reduce your operating costs as well.

**Staten:** Absolutely. One of the biggest [benefits you get from virtualization](#) is flexibility. It's so much easier to patch a workload and simply keep it running, while you are doing that. Move it to another system, but apply the patch, make sure the patch worked, deploy a clone, and then turn off the old version.

That's much more powerful, and it gives a lot more flexibility to the IT shop to maintain higher [service-level agreements \(SLAs\)](#), to keep the business up and running, to roll out new things faster, and be able to roll them back more easily.

**Gardner:** Andy Rubinson, this certainly sounds like a no-brainer: Get better performance for less money and postpone large capital expenditures. What are some of the risks that could come into play while we are starting to look at this whole picture? I'm interested in what's holding people back.

**Rubinson:** I focus mainly on [delivery over the Internet](#). There are definitely some challenges, if you're talking about using the Internet with your data center infrastructure -- things like performance latency, availability challenges from cable cuts, and things of that nature, as well as [security threats](#) on the Internet.



It's thinking about how can you do this, how can you deliver to a global user base with your data center, without having to necessarily build out data centers internationally, and to be able to do that from a consolidated standpoint.

**Gardner:** So, for those organizations that are not just going to be focused on employees, or, if they are, that they are a global organization, they need to be thinking the most [wide area network \(WAN\)](#) possible. Right?

**Rubinson:** Absolutely.

**Gardner:** Let's go to our practitioner, Tom Winston. Tom, what sort of effects were you dealing with [at Phase Forward](#), when you were looking at planning and strategy around data center location, capacity, and utilization?

### *Early adopter*

**Winston:** Well, we were in a somewhat different position, in that we were actually an early adopter of virtualization technology, and certainly had seen the benefits of using that to help contain our data-center sprawl. But, we were also growing extremely rapidly.



When I joined the organization, it had two different data centers -- one on the East Coast and one on the West Coast. We were facing the challenge of potentially having to expand into a European data center, and even potentially a Pacific Rim data center.

By continuing to expand our virtualization efforts, as well as to leverage some of the technologies that Andy just mentioned as far as, Internet acceleration, via some of the Akamai technologies, we were able to forego that data center expansion. In fact, we were able to consolidate our data center to one East Coast data center, which is now our primary hosting center for all of our applications.

So, it had a very significant impact for us by being able to leverage both that WAN acceleration, as well as virtualization, within our own four walls of the data center. [Editor's note: WAN here and in subsequent uses refers to public wide area networks and not private.]

**Gardner:** Tom, just for the edification of our listeners, tell us a little bit about Phase Forward. Where are your users and where do your applications need to go.

**Winston:** We run [electronic data capture \(EDC\)](#) software, and [pharmacovigilance](#) software for the largest pharmaceutical and clinical device makers in the world. They are truly global organizations in nature. So, we have users throughout the world, with more and more heavy population coming out of the Asia Pacific area.

We have a very large, diverse user base that is accessing our applications 24x7x365, and, as a result, we have performance needs all the time for all of our users.

In an age where, as James mentioned, people are expecting things to be moving extremely quickly and always available, it's very important for us to be able to provide that application all the time, and to perform at a very high level.

One of the things James mentioned from an IT perspective is being able to manage that virtual stack. Another thing that [virtualization](#) allows us to do is to provide that stack and to improve performance very quickly. We can add additional compute resources into that virtual environment very quickly to scale to the needs that our users may have.

**Gardner:** James Staten, back to you. Based on Tom's perspective of the combination of that virtualization and the elasticity that he gets from his data center, and the ability to locate it flexibly, thanks to some network optimization and reliability issues, how important is it for companies now, when they think about data center consolidation, to be flexible in terms of where they can locate?

### *All over the place*

**Staten:** It's important that they recognize that their users are no longer all in the same headquarters. Their users are all over the place. Whether they are an internal employee, a customer, or a business partner, they need to get access to those applications, and they have a performance expectation that's been set by the Internet. They expect whatever applications they are interacting with will have that sort of local feel.

That's what you have to be careful about in your planning of consolidation. You can consolidate branch offices. You can consolidate down to fewer data centers. In doing so, you gain a lot of operational efficiencies, but you can potentially sacrifice performance.

You have to take the lessons that have been learned by the people who set the performance bar, the providers of Internet-based services, and ask, "How can I optimize the WAN? How can I push out content? How can I leverage solutions and networks that have this kind of intelligence to allow me to deliver that same performance level?" That's really the key thing that you have to keep in mind. Consolidation is great, but it can't be at the sacrifice of the user experience.

**Gardner:** When you find the means to deliver that user experience, that frees you up to then place your data centers strategically based on things like skills or energy availability or tax breaks, and so forth. Isn't that yet another economic incentive here?

**Staten:** You want to have fewer data centers, but they have to be in the right location, and the right location has to be optimized for a variety of factors. It has to be optimized for where the appropriate skill sets are, just as you described. It has to be optimized for the geographic constraints that you may be under.

You may be doing business in a country in which all of the citizen information of the people who live in that country must reside in that country. If that's the case, you don't necessarily have to own a data center there, but you absolutely have to have a presence there.

**Gardner:** Andy, back to you. What are some of the pros and cons for this Internet delivery of these applications? I suppose you have to rearchitect, in order to take advantage of this as well.

**Rubinson:** There are two main areas from the positives, the benefits, and that's the cost efficiency of delivering over the Internet, as well as the responsiveness. From the cost perspective, we're able to eliminate unnecessary hardware. We're able to take some of that load off of the servers, and [do the work in the cloud](#), which also helps reduce them.

### *A lot of cost efficiencies*

There are a lot of cost efficiencies that we get, even as you look to Tom's statement about being able to actually eliminate a data center and avoid having to build out a new data center. Those are all huge areas, where it can help to use the Internet, rather than having to build out your own infrastructure.

Also, in terms of responsiveness, by using the Internet, you can deploy a lot more quickly. As Tom explained, it's being able to reach the users across the globe, while still consolidating those infrastructures and be able to do that effectively.

This is really important, as we have seen more and more users that are going outside of the corporate WANs. People are connecting to suppliers, to partners, to customers, and to all sorts of things now. So, the private WANs that many people are delivering their apps over are now really not effective in reaching those people.

**Gardner:** As James said earlier, we've got different workloads and different types of applications. Help me understand what Akamai can do. Do you just accelerate a web app, or is there a bit more in your quiver in terms of dealing with different types of loads of media, content, application types?

**Rubinson:** There are a variety of things that we are able to deliver over the Internet. It includes both web- and [IP](#)-based applications. Whether it's [HTTP](#), [HTTPS](#), or anything that's over [TCP/IP](#), we're able to accelerate.

We also do streaming. One of the things to consider here is that we actually have a global network of servers that kind of makes up the cloud or is an overlay to the cloud. That is helping to not only deliver the content more quickly, but also uses some caching technology and other

things that make it more efficient. It allows us to give that same type of performance, availability, and security that you would get from having a private WAN, but doing it over the much less expensive Internet.

**Gardner:** You're looking at specifics of an application in terms of what's going to be delivered at frequent levels versus more infrequent levels, and you can cache the data and gain the efficiency with that local data store. Is that how it works?

**Rubinson:** A lot of folks think about Akamai as being a [content delivery network \(CDN\)](#), and that's true. There is caching that we are doing. But, the other key area where we have benefit is through the delivery of dynamic data. By [optimizing the cloud](#), we're able to [speed the delivery of information](#) from the origin as well. That's where it's benefiting folks like Tom, where he is able to not only cache information, but the information that is dynamic, that needs to get back from the data center, goes more quickly.

**Gardner:** Let's check in with Tom. How has that worked out for you? What sort of applications do you use with wide area optimization, and what's been your experience?

### *Flagship application*

**Winston:** Our primary application, our flagship application, is a product called [InForm](#), which is the main EDC product that our customers use across the Internet. It's accelerated using Akamai technology, and almost 100 percent of our content is dynamic. It has worked extremely well.

Prior to our deployment of Akamai, we had a number of concerns from a performance standpoint. As James mentioned, as you begin to virtualize, you also have to be very conscious of the potential performance hits. Certainly, one of the areas that we were constrained with was performance around the globe.

We had users in China who, due to the amount of traffic that had to traverse the globe, were not happy with the performance of the application. Specifically, we brought in Akamai to start with a very targeted group of users and to be able to accelerate for them the application in that region.

It literally cut the problem right out. It solved it almost immediately. At that point, we then began to spread the rest of that application acceleration product across the rest of our domains, and to continue to use that throughout the product set.

It was extremely successful for us and helped solve performance issues that our end users were having. I think some of the comments that James made are very important. We do live in a world where everybody expects every application across the Internet to perform like Google. You want to search and you expect it to be back in seconds. If it's not, people tend to be unhappy with the performance of the application.

In our application, it's a much more complex application. A lot more is going on behind the scenes -- database calls, whatever it may be. Having an application perform to the level of a

Google is something that our end users expect, even though obviously it's a much different application in what it's attempting to solve and what it's attempting to do. So, the benefits that we were able to get from the acceleration servers were very critical for us.

**Rubinson:** Just to add to that, we recently commissioned a study with Forrester, looking at what is that tolerance threshold [for a page to load]. In the past it had been that people had tolerance for about four seconds. As of this latest study, it's down to two seconds. That's for [business to consumer \(B2C\)](#) users. What we have seen is that the [business-to-business \(B2B\)](#) users are even more intolerant of waiting for things.

It really has gotten to a point where you need that immediate delivery in order to drive the usage of the tools that are out there.

**Gardner:** I suppose that's just human nature. Our expectations keep going up. They usually don't go down.

**Rubinson:** True.

**Gardner:** Back to you, Tom. Tell me a little bit more about this application. Is this a [rich Internet application \(RIA\)](#)? Is this strictly a web interface? Tell us a little bit more about what the technical challenge was in terms of making folks in China get the same experience as those on the East Coast, who were a mile away from your data center.

### *Everything is dynamic*

**Winston:** The application is one that has a web front-end, but all the information is being sent back to an [Oracle](#) database on the back-end. Literally, every button click that you make is making some type of database query or some type of database call, as I mentioned, with almost zero static content. Everything is dynamic.

There is a heavy amount of data that has to go back and forth between the end user and the application. As a result, prior to acceleration, that was very challenging when you were trying to go halfway around the globe. It was almost immediate for us to see the benefits by being able to hop onto the Akamai Global Network and to cut out a number of the steps across the Internet that we had to traverse from one point to our data center.

**Gardner:** So, it was clearly an important business metric, getting your far-flung customers happy with their response times. How did that however translate back when you reverse engineered from the experience to what your requirement would be within that data center? Was there sort of a meeting of the minds between what you now understand the network is capable of, with what then you had to deliver through your actual servers and infrastructure?

I guess I'm looking for an efficiency metric or response in terms of what the consolidation benefit was.

**Winston:** As I mentioned, we had already consolidated from a virtualization standpoint within the four walls of the data center. So, we were continuing to expand in that footprint. But, what it allowed us to do was forego having to put a data center in the Pacific Rim or put a data center in Europe to put the application closer to the end user.

**Gardner:** Let's look to the future a little bit. James, when people think nowadays about cloud computing, that's a very nebulous discussion and topic set. It seems as if what we're talking about here is that more enterprises are [going to have to themselves start behaving like what people think of as a cloud](#).

**Staten:** Yes, to a degree. There is obviously a positive aspect of cloud and one that can potentially be a negative.

Operating like a cloud is really operating in this more homogeneous, virtualized, abstracted world that we call server virtualization in most enterprises. You want to operate in this mode, so that you can be flexible and you can put applications where they need to be and so forth.

But, one of the things that cloud computing does not deliver is that if you run it in the cloud, you are not suddenly in all geographies. You are just in a shared data center somewhere in the United States or somewhere in your geography. If you want to be global, you still have to be global in the same sense that you were previously.

### *Cloud not a magic pill*

**Rubinson:** Absolutely. Just putting yourself in the cloud doesn't mean that you're not going to have the same type of latency issues, delivering over the Internet. It's the same thing with availability in trying to reach folks who are far away from that hosted data center. So, the cloud isn't necessarily the answer. It's not a pill that you can take to fix that issue.

**Gardner:** Andy, I don't think you can mention names, but you are not only accelerating the experience for end users of enterprise applications like a Phase Forward. You're also providing similar services for at least several of the major cloud providers.

**Rubinson:** It really is anybody who is using the cloud for delivery. Whether it's a high-tech, a pharma company, or even a hosting provider in the cloud, they've all seen the value of ensuring that their end users are having a positive experience, especially folks like [software-as-a-service \(SaaS\)](#) providers.

We've had a lot of interest from SaaS companies that want to ensure that they are not only able to give a positive user experience, but even from a sales perspective, being able to demonstrate their software in other locations and other regions is very valuable.

**Gardner:** Now, James, when a commercial cloud provider provides an SLA to their customers, they need to meet it, but they also need to keep their costs as low as possible. More and more enterprises are trying to behave like service providers themselves, whether it's through [ITIL](#)

adoption, IT shared services or [service-oriented architecture \(SOA\)](#). Over time, we're certainly seeing movement toward a provider-supplier, consumer-subscription relationship of some kind.

If we can use this acceleration and the ability to use the network for that requirement of performance to a certain degree, doesn't this then free up the folks who have to meet those SLAs in terms of what they need to provide? I'm getting back to this whole consolidation issue.

**Staten:** To some degree. Obviously, by using the best practices that we've adopted to have blazing fast websites and applying them to make sure that all of your applications, consumed by everyone, are still blazing fast means that you don't have to reinvent the wheel. Those practices work for your website. You just apply them to more areas.

If you're applying practices you already know, then you can free up your staff to do other things to modernize the infrastructure, such as deploying ITIL more widely than you have so far. You can make sure that you apply virtualization to a larger percentage of your infrastructure and then deal with the next big issue that we see in consolidation, which is [virtual machine \(VM\) sprawl](#).

### *Can get out of control*

This is where you are allowing your enterprise customers, whether they are enterprise architects, developers, or business units to deploy new VMs much more quickly. Virtualization allows you to do that, but you can quickly get out of control with too many VMs to manage.

Dealing with that issue is what is front and center for a lot of enterprise IT professionals right now. If they haven't applied the best practices or performance to their application sets and to their consolidation practices, that's one more thing on their plate that they need to deal with.

**Gardner:** So, this also can relate to something that many of us are forecasting. Not much of it happening yet, but it's this notion of a hybrid approach to cloud and sourcing, where you might use your data center up to a certain utilization, and under certain conditions, where there is a spike in demand, you could just offload that to a third-party Cloud provider.

If you're assured from the WAN services that the experience is going to be the same, regardless of the sourcing, they are perhaps going to be more likely to pursue such a hybrid approach. Is that fair to say, James?

**Staten:** This is a really good point that you're bringing up. We wrote about this in a report we called 'Hollow Out The MOOSE'. MOOSE is [Forrester's term](#) for the Maintenance and Ongoing Operations, Systems, and Equipment, which is basically everything you are running in your data center that hasn't yet been deployed up to this point.

The challenge most enterprises have is that MOOSE consumes 70 or 80 percent of their entire budget, leaving very little for new innovation and other things. They see things like cloud and they say, "This is great. I'll just move this stuff to the cloud, and suddenly it will save me money."

No. The real answer is that you need to choose the right type of solution for the right problem. We call this Strategic Rightsourcing, which says to take the things that others do better than you and have others do them, but know economically whether that's a positive tradeoff for you or not. It doesn't necessarily have to be cash positive, but it has to be an opportunity to be cost positive.

In the case of cloud computing, if I have something that I have to run myself, it's very unique to how I design it, and it's really best that I run it in my data center, you're not saving money by putting that in the cloud.

If it's an application that has a lot of elasticity, and you want it to have the ability to be on two virtual machines during the evening, and scale up to as many as 50 during the day, and then shrink back down to 2, that's an ideal use of cloud, because cloud is all about temporary capacity being turned on.

A lot of people think that it's about performance, and it's not. Sure, load balancing and the ability to spawn new VMs increases the performance of your application, but performance is experienced by the person at the end of the wire, and that's what has to be optimized. That's why those types of networks are still very valuable.

**Gardner:** Tom Winston, is this vision of this hybrid and the use of cloud for ameliorating spikes and therefore reducing your total cost appealing to you?

### *Has to be right*

**Winston:** It is, but I couldn't agree more with what James just said. It has to be for the right situation. Certainly, we've started to look at some of our applications, potentially using them in a cloud environment, but right now our critical application, the one that I mentioned earlier, is something that we have to manage. It's a very complex environment. We manage it and we need to hold it very close to the vest.

People have the idea that, "Gee, if I put it in the cloud, my life just got a lot easier." I actually think the reverse might be true, because if you put it into the cloud, you lose some control that you have when it's inside your four walls.

Now, you lose the ability to be able to provide the level of service you want for your customers. Cloud needs to be for the right application and for the right situation, as James mentioned. I really couldn't agree more with that.

**Gardner:** So, the cloud is not the right hammer for all nails, but for when that nail is correct, that hybrid model can perhaps be quite a economic benefit. Andy, at Akamai, are you guys looking at that hybrid model, and is there something there that your services might foster?

**Rubinson:** This is really something that we are agnostic about. Whether it's in a data center owned by the customer or whether it's in a hosted facility, we are all about the means of delivery. It's delivering applications, websites, and so forth over the public Internet.

It's something we're able to do, if there are facilities that are being used for, say, disaster recovery, where it's the hybrid scenario that you are describing. For Akamai, it's really about how we're able to accelerate that. How we are able to optimize the routing and the other protocols on the Internet to make that get from wherever it's hosted to a global set of end users.

We don't care about where they are. They don't have to be on the corporate, private WANs. It's really about that global reach and giving the levels of performance to actually provide an SLA. Tell me who else out there provides an SLA for delivery over the Internet. Akamai does.

**Gardner:** Well, we'll have to leave it there. We've been discussing how data center consolidation and modernization can help enterprises cut costs, reduce labor, slash their energy use, and become more agile, but also keeping in mind the requirements about the performance across wide area networks.

We've been joined by James Staten, he is a Principal Analyst at Forrester Research. Thank you, James.

**Staten:** Thank you.

**Gardner:** We were also joined by Andy Rubinson, Senior Product Marketing Manager at Akamai Technologies. Thank you, Andy.

**Rubinson:** Thank you very much.

**Gardner:** Also, I really appreciate your input Tom Winston, Vice President of Global Technical Operations at Phase Forward.

**Winston:** Dana, thanks very much. Thanks for having me.

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

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

*Transcript of a sponsored BriefingsDirect podcast on how data center consolidation and modernization helps enterprises reduce cost, cut labor, slash energy use, and become more agile. Copyright Interarbor Solutions, LLC, 2005-2009. All rights reserved.*