

Network Transformation: Taking a Second Look at Network Architectures as we Tackle New Demands in the Enterprise

Transcript of a BriefingsDirect Podcast examining how data-center transformation requires a new and convergent look at network architecture.

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

Offer: For more insight into best practices for transforming your data center to increase agility, lower costs and improve business continuity and availability, download three valuable Data-Center Transformation whitepapers from HP by going to: www.hp.com/go/dctpodcastwhitepapers

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 reevaluating [network architectures](#) in light of newer and evolving demands. Most enterprise networks are the result of a patchwork effect of bringing in equipment as needed over the years to fight the fire of the day, with little emphasis on strategy and the anticipation of future requirements.



Nowadays, we see that network requirements have, and are, shifting, as IT departments adopt improvements such as [virtualization](#), [software as a service \(SaaS\)](#), [cloud computing](#), and [service-oriented architecture \(SOA\)](#).

The network loads and demands continue to shift under the weight of Web-facing applications and services, security and regulatory compliance, [governance](#), ever-greater data sets, and global area service distribution and related performance management.

It doesn't make sense to embark upon a data-center transformation journey without a strong emphasis on network transformation as well. Indeed, the two ought to be brought together converging to an increasing degree over time.

Here to help explain the evolving role of network transformation and to rationalize the strategic approach to planning and specifying present and future enterprise networks, is Lin Nease, director of Emerging Technologies, [HP ProCurve](#). Welcome to the show, Lin.

Lin Nease: Thank you.

Gardner: We're also joined by [John Bennett](#), worldwide director, [Data Center Transformation Solutions](#) at HP. Hello, John.

John Bennett: Hi, Dana.

Gardner: And, [Mike Thessen](#), practice principal, Network Infrastructure Solutions Practice in the HP Network Solutions Group. Welcome to the show, Mike.

Mike Thessen: Thank you. Hello, everyone.

Gardner: John Bennett, let's start with you. Tell me a little bit about the typical enterprise network as it's evolved, and how does that affect data-center transformation?

Helping customers

Bennett: Let's start by reminding people what data-center transformation is about. Data-center transformation is really about helping customers build out a next-generation data center, an adaptive infrastructure, that is designed to not only meet the current business needs, but to lay the foundation for the plans and strategies of the organization going forward.



In many cases, the IT infrastructure, including the facilities, the servers, the network, and storage environments can actually be a hindrance to investing more in business services and having the agility and flexibility that people want to have, and will need to have, in increasingly competitive environments.

When we talk about that, very typically we talk a lot about facilities, servers, and storage. For many people, the networking environment is ubiquitous. It's there. But, what we discover, when we lift the covers, is that you have an environment that may be taking lots of resources to manage and keep up-to-date.

You may have an environment that is not capable of moving network connections as quickly as servers, applications, and storage devices need and you want them to in order to meet your agility objectives.

We also find an environment that can be a cabling nightmare, because it has grown in a kind of an organic way over time. So, in looking at data center strategy and looking at data-center transformation, we have to make sure that the whole data-center architecture, including the network infrastructure both inside the data center and the services it provides to the organization, are really aligned to meet those goals and objectives.

It becomes increasingly important to have, as we continue to experience incredible explosions in storage and data volumes, with the new types of information, and with the historical information that's maintained.

The networking infrastructure becomes key, as an integration fabric, not just between users in business services, but also between the infrastructure devices in the data center itself.

That's why we need to look at network transformation to make sure that the networking environment itself is aligned to the strategies of the data center, that the data center infrastructure is architected to support those goals, and that you transform what you have and what you have grown historically over decades into what hopefully will be a "lean, mean, fighting machine."

Gardner: Lin Nease, from the perspective of an architect, is there a lag, or perhaps a disconnect, between the trajectory and evolution of networks and where the entire data center has been moving towards?

Multiple constituencies

Nease: Absolutely. The network has basically evolved as a result of the emergence of the Internet and all forms of communications that share the network as a system. The server side of the network, where applications are hosted, is only one dimension that tugs at the network design in terms of requirements.



You find that the needs of any particular corner of the enterprise network can easily be lost on the network, because the network, as a whole, is designed for multiple constituencies, and those constituencies have created a lot of situations and requirements that are in themselves special cases.

In the data center, in particular, we've seen the emergence of a formalized virtualization layer now coming about and many, many server connections that are no longer physical. The history of networking says that I can take advantage of the fact that I have this concept of a link or a port that is one-to-one with a particular service.

That is no longer the case. What we're seeing with virtualization is challenging the current design of the network. That is one of the requirements that are tugging at a change or provoking a change in overall enterprise network design.

Gardner: Mike Thessen, from a systems integrator problem set, have there been different constituencies, perhaps even entirely different agendas, at work in how networks are put together and then how the data center requirements are coming around?

Thessen: Sure. From the integrator perspective, we get involved with clients' real problems and real requirements. People listen to the press and they are trying to do the right thing. What we are

finding is that, many times, they get distracted by that and lose sight of the fact that what they're really trying to do is provide access to applications within their data center to their user base.

We try to bring them back around, when we work with them on a consulting basis, to not so much focus on products, but on what they are trying to achieve overall at a very high level -- just to get things started.

Gardner: Is there a new philosophy perhaps that needs to be brought to the table, Mike, around planning for network, data center, and storage requirements in tandem? This seems to have been something that happened on a fairly linear basis in the past. How do we get that to be simultaneous, or is that the right way to go?



Thessen: In my mind, you are really talking about collaboration. Data-center networking certainly cannot happen in a vacuum. In years past, you were effectively just providing [local area network \(LAN\)](#) and [wide area network \(WAN\)](#) connectivity. Servers were on the network, and they got facilities from the network to transport their data over to the users.

Now, everything is becoming converged over this network -- "everything" being data storage, and telephony. So, it's requiring more towers inside of corporate IT to come together to truly understand how this system is going to work together.

Gardner: Lin Nease, is this about services orientation? Do some of the same methods and best practices and architectural approaches from the application side come to bear on the network features and functions as well?

The only way out

Nease: Absolutely. In fact, that's the only way out. With the new complexity that has emerged, and the fact that traditional designs can no longer rely on physical barriers to implement policies, we have reached a point, where we need an architecture for the network that builds in explicit concepts of policy decisions and policy enforcement. They're not always in the same place, and it's not always intuitive where a policy should be enforced or decided upon.



As a result of that, the only way out is to regard the network itself as a service that provides connectivity between stations -- call them logical servers, call them users, or call them applications. In fact, that very layering alone has forced us to think through the concept of offering the network as a service.

So, service orientation is crucial and it will allow those who build infrastructure to no longer be forced into an ad-hoc situation, where they build infrastructure in an extremely fluid manner with respect to how applications are being designed, but rather to a much more formal presentation of

what they are doing as a server. That presentation becomes the designed target for both sides, and, as I said, that's probably the only way out with the complexity that's emerged.

Just to give you an example, look at a virtual server today and some of the new technologies that are being proposed like [Single Root I/O virtualization](#) combined with virtual switching, combined with edge, blades, switching, and standalone switches. You could have seven or eight queues that separate an application from just the core of the network. That's far more than in the past. That complexity is going to cause applications to break. So, this mentality is probably the only way out.

Gardner: I'd like to drill down, if we could, on virtualization. There are several different layers and levels of this course. We're starting to even hear more about [desktop virtualization](#), PC-over-IP, and different approaches to bring the essence of an operating system environment to the user, but without them really having the actual compute power locally.

Let's go back to John Bennett. John, tell us a little bit about the different dimensions of virtualization, and how that has an impact on this complexity issue in the network?

Bennett: Virtualization is a major theme. [Ann Livermore](#) in her keynote at [VMworld](#) challenged people to think about moving from virtualizing servers to virtualizing their infrastructure, and even go beyond that. Server virtualization has just been the starting point for people moving to a shared infrastructure.

In parallel with that, we see an increasing drive and demand for virtualizing storage to have it both be more efficiently and effectively used inside the data center environment, but also to service and support the virtualized business services running in virtualized servers. That, in turn, carries into the networking fabric of making sure that you can manage the network connections on the fly, as Lin talked about.

Then, it reaches outside of the data center. Desktop virtualization is seen as a major opportunity to not only provide for better control of desktop devices, but also better security and better protection of end-user data. Now, you are taking an environment that used to run locally in an office with just data connections back to the data center to now you have an environment which depends upon the data center for all of the services that are being provisioned on a virtualized desktop. So, you have that complexity taking place as well.

Virtualization is not only becoming pervasive, but clearly the networking fabric itself is going to be key to delivering high quality business services in that environment.

Gardner: Back to Mike Thessen. Tell me a bit about this from the integration perspective. How do these virtualization complexities need to be considered as folks move towards a network transformation of some sort?

Understanding requirements

Thessen: Virtualization, from the network perspective, really centers on several aspects. First, from the system and application perspective, we have to understand the requirements of how blade server interconnectivity is going to be achieved, how things like dynamic movement of hypervisors will be managed, and basically how much [Layer 2](#) adjacency is required in the network.

While Layer 2 is expanding in the data center, it really needs to be contained such that it's limited to what is required within a pod, cell, module, or whatever the term is that a client may use to define a span of the data center.

We don't want to allow Layer 2 domains to expand across the entire data center or be unlimited between data centers. We want to contain this Layer 2 environment. While it's getting bigger we don't want to have the attitude that we'll just allow it to go everywhere. There will still be issues with that large span of Layer 2 Ethernet connectivity, and from a manageability perspective it gets very complex.

Second, there is a trend to utilize network device virtualization to eliminate the need for things like [spanning tree](#), the redundant default gateway mechanisms, and so forth. Those are different ways to use technology to expand the Layer 2 domains, but limit the risk associated with that.

Third, this tends to utilize device and routing control-plane virtualization for logical separation of external-facing applications, especially in things like financial industries, and so forth.

The test and development and QA environments are extremely important, especially as things become more virtualized. Things really have to be tested. We really promote having our clients spend the extra money to have that lab always available to test things in. Then, naturally, how do we do that more cheaply and less expensively?

You can use certain virtualization techniques in the networking hardware to separate those environments in a logical manner, as opposed to having to buy completely separate networks to do your testing.

The fourth thing is that networks need to be prepared for the convergence of the communication paths for data and storage connectivity inside the data center. That's the whole conversion -- enhance, Ethernet, [Fiber Channel](#) over Ethernet. That's the newest leg of the virtualization aspect of the data center.

Gardner: Of course, nowadays, the IT and other departments, and telecom folks are all under pressure to cut cost. So, if we are going to transform networks, not only do we have to look at complexity and bringing up the support of additional requirements, but folks are looking for efficiencies as well.

Lin Nease, what about network transformation? It can allow our requirements to be met, but perhaps, at the same time, find efficiencies, higher utilization, and cut overall cost.

Accounting for use cases

Nease: It's important to account for all the use cases that are critical to the enterprise, and it's possible to design networks that have what I will call a most common denominator. What we've found with HP's own huge data center consolidation was that ruthless standardization was the key to cutting cost.

The way I cut cost is I don't have an artificial metric like server utilization, CPU utilization or network utilization. I have a simple metric of budget. And from budget, comes all other mechanisms for optimization.

A most common denominator network is another way of saying, "I can build a substrate of the data center network at least for this point in time, call it a pod, call it a cell, call it a unit of modularity. I can put it in place and it will solve every use case I care about -- everything from the high balance, low latency requirements of middleware clustering or database clustering all the way down to the mundane."

I can cross server tiers for example with relatively low bandwidth requirements. But I can do that all from a very high performance substrate. If I have one design, I have only one thing I need to manage. Operationally, it solves multiple problems at once. But, rather than being purpose built for each application, now the network is built once as a standard. As a result of the fact that it solves all the problems, I can change the nature of what my change review-boards look like, for example.

If I want to go in and put in new servers, I don't have to worry about including someone from a particular department in the decision, because I know the network works for all the use cases I care about.

Gardner: Lin, as we try to get that overall perspective of a solution approach, do we also find ourselves able to cut energy use? That seems to be an important part of a lot of transformation initiatives as well.

Nease: Oddly enough, on one side of the coin, I just talked about ruthless standardization. The simplest way to drive down the cost of my process is to overkill. On the other hand, now we have the concept that overkill is bad, because it consumes a lot of energy. Well, here's our way out. Here's the degree of freedom we have.

Thermal management is probably the biggest hitter in terms of energy savings and consumption, going forward. The key for networking in particular is, number one, to enable higher utilization

of servers. That's the most direct way of saving on energy. Then, secondly, to make sure that the thermal design of the data center is optimized.

It's quite possible to have a completely independent architectural approach to logical topology, versus the approach to physics. In this case, when I say, physics, I mean how I support the hot aisle, cold aisle, the ventilation, and the energy pressure drops.

If I can service a cooling aspect of the data center, it turns out that air conditioning and cooling account for more than half the energy consumption in a typical data center. So, by optimizing on the thermal front, I can still have a very simple network and I can separate the concern by topology architecture.

Gardner: We hear the term "convergence" batted around a lot, and just from our call today, I can tell that convergence really has multiple aspects and perhaps even multiple levels of convergence.

Back to you, John Bennett. In talking about our network transformation philosophy, are we converging storage and data with applications. Is it data center and network? Is it on-premises and off-premises or cloud? How can we set a taxonomy or get a handle on what we mean by convergence nowadays?

Better integration

Bennett: Fundamentally, convergence is about better integration across the technology stacks that help deliver business services. We're saying that we don't need separate, dedicated connections between servers for high availability from the connections that we use to the storage devices to have both a high-volume traffic and high-frequency traffic accesses to data for the business services or that we have for the network devices and the connections between them for the topology of the networking environment.

Rather, we are saying that today we can have one environment capable of supporting all of these needs, architected properly for particular customer's needs, and we bring into the environment separate communications infrastructures for voice.

So, we're really establishing, in effect, a common nervous system. Think about the data center and the organization as the human body. We're really building up the nervous system, connecting everything in the body effectively, both for high-volume needs and for high-frequency access needs.

Gardner: Mike Thessen, we've heard here about the need for convergence and managing complexity. We're hearing about a lot more services coming into play. It now sounds as if we are taking what people refer to as a utility or grid approach for data and applications and we are applying that now to networks. Is that how we should be thinking about this, instead of getting

bogged down in convergence? Is this really more of a cloud or fabric approach that includes network services?

Thessen: As someone said, a few minutes ago, at some level of the network the primary aspect needs to be utility. When you're talking about clouds, clouds can be not what people think about clouds from [Amazon](#) or wherever, but even clouds inside a client's own IT environment. So, it's possible to do something like replace the way they typically would do external client access to their data. These things come to mind especially in the financial industry.

The most important thing is really still the brutal standardization, as Lin said -- network modularity, logical separation, utilizing those virtualization techniques that I talked about a few minutes ago, and very well-defined communications flows for those applications. When things may not go right, when things break, or when there are performance issues, there is documentation there that defines who is talking to what.

Additionally, you need those communication flows especially in these SaaS or cloud-computing, or convergence environments to truly secure those environments appropriately. Without understanding who is talking to whom, how applications communicate, and how applications get access to other IT services, such as directory services and so forth, it's really difficult to secure them appropriately.

We haven't talked about WAN very much. We've been focused on the data center. But, data centers are more or less useless, without people being able to access them through some sort of wide-area facility.

We focus a lot on determining how these new applications are going to communicate over the WAN by doing dependency mapping of the applications and by doing transaction profiling of the applications from the network perspective. We identify, not only how much bandwidth required by transaction and how many users they're going to be hitting at any given point in time, but also how latency is going to affect the end-user experience.

If you move everything into the cloud, which implies virtual centralization, the users are now more separated from that. So, you really have to pay close attention to how latency is going to affect these new environments.

Gardner: It certainly seems to be an awful lot to chew, to bite off, and factor when we move towards network transformation. I wonder what some of the common mistakes are that people make as they approach a certain path, a crawl-walk-run, or a methodological, or an architectural overview. What might they do that perhaps prevents them from getting to where they want to be? Let's start with you Lin. What are some common mistakes people get into as they start to move towards network transformation?

The most common mistake

Nease: This one is near and dear to my heart, being an evangelist for the networking business. Too often people are compelled by a technology approach to rethink how they are doing networking. IT professionals will hear the overtures of various vendors saying, "This is the next greatest technology. It will maybe enable you to do all sorts of new things." Then, people waste a lot of time focusing on the technology enablement, without actually starting with what the heck they're trying to enable in the first place.

Unfortunately, I think this is the most common mistake by far. I'll give you an example of how the inevitability of some technology trend will probably lead people far less down a path of great optimization than what they are thinking, and this is in convergence of storage traffic with data traffic.

There is a big difference between storage traffic and voice traffic. Voice traffic is limited by the perceptions of human beings. It will never require more bandwidth for a phone call, and probably less than it does today, as technology evolves. It's very easy to incorporate that into a common plumbing.

Storage, on the other hand, is directly tied to server performance. Storage is going to continuously grow in terms of requirements. There is so much focus on replacing the technologies that we don't like that we forget about what we're trying to enable from an application perspective.

How do I have applications that are deployed on infrastructure that follow the potential energy of business requirement changes, rather than first focusing on how the plumbing works? That's the biggest mistake people make. What they'll find is that they'll look at a lot of these technologies. Two years from now, you'll see networks that look quite a bit like they do today, because the focus has not been on enabling what it is that people are trying to actually accomplish in their business.

Gardner: Mike Thessen, the same question. Are there some common mistakes, from your vantage point, as a systems integrator, that folks fall prey to as they move into network transformation mode?

Thessen: Lin focused on picking a technology. If you take that a step farther, lots of times our clients are hell-bent on picking a specific product or a specific vendor, prior to actually defining the requirements.

I have a couple of sayings that a bill of materials is not a design and it takes a lot of effort to get a PowerPoint presentation into something that you can actually implement. We often see clients coming to us and they've already got a bill of materials. Then, they want you to back into an architecture and a design an implementation strategy.

As I said earlier, we take a different approach. We prefer to get in earlier, and really strategize with the client -- what are we trying to do and what are we trying to come from and get to.

Testing is required

I mentioned this earlier. Sufficient testing of this new technology in a dedicated lab environment is absolutely required, whether you are talking about how the applications are going to work, or just making sure that you can get the network components working together properly with all the new features and functions that you might want to implement. It's absolutely key, especially for the data center environments, to have that test.

Sometimes, we also see that the need for one gigabyte, or lower, speed transports are being forgotten about. Everybody is all wrapped up around 10 gig, 10 gig, 10 gig -- they've got to have it everywhere.

We typically recommend a mix of 1 gig and 10 gig, based on the requirements for the existing servers. Ten gig coming from blade is absolutely right on target, but do we really need 10 gig for every server on the network? Probably not, at this time.

What we need to look at is what is the real performance of these services and when the next technology refresh cycle for these servers is going to occur. Possibly, if we do our modularity and standardization process right, we can rev servers and network gears simultaneously within those modules, and really keep our cost down, as opposed to piling ten gig on everything right now.

A lot of times, our clients also forget about the fact that a lot of the management interfaces on some equipment don't even support one gig yet. So, there is this mix of the technologies and products that need to come together and they really need to be thought through, before you go out and buy the bill of materials prior to having all your requirements in check.

Gardner: Lin Nease, how about some examples of where folks in the field have embarked on a network transformation, and taken into consideration some of these issues that we have been discussing today? What have they found? What are some of the paybacks? Any examples of success or perhaps things to avoid?

Nease: I'll start with the biggest success, and it's very close to home -- Hewlett-Packard's own data center consolidation -- 85 data centers down to six. The thinking on the network, which was very consistent thinking overall, was simplicity. If you were desperate and you had to actually have a low budget, what types of actions would you take? If you don't take those actions, you should justify the benefits of doing something different than what you would do well, if you are desperate.

It wasn't desperation, but rather sheer cost savings, incredible cost savings, that HP got out of IT. They actually deployed 5,000 of our ProCurve devices, for example, in a network that was deliberately kept extraordinarily simple. As we talked about earlier, we ruthlessly standardized the whole [VLAN](#) topology. The approach to layering was kept very, very simple.

The law of large numbers said that they didn't actually have to build an extremely complex network to get big gains, and there is a lot more behind that than you might think. In the process of doing that in the network, HP saved well over \$90 million in the deployment of just six data centers on network gear alone.

Gardner: Mike Thessen, do any examples come to mind in terms of how this should be done in the field?

Right on track

Thessen: We're working in all kinds of environments all the time, from financials, to manufacturing, to retail. The things we are covering here today are all right on track with what our clients are asking for: "How do we implement virtualization? I want to consolidate my voice and my data infrastructures. I want to be prepared for any level of convergence within the data center from a storage and data perspective, but I want to do it in the right way."

We had several instances where we put in some of the converged infrastructures just from a future-proofing perspective, because the client's timeline was right. In other cases, we talked the client out of going down that path and doing what I typically call more of a discrete network non-converged, if you will, data, and fiber channel over Ethernet topology, just because they were standardizing more on blades.

They had already more or less achieved a lot of their cabling reductions, because of the mechanisms they were using within the blades environments. So, they were able to leverage existing cabling infrastructures and didn't have to add anymore. Still, they took advantage of a lot of the cost-reduction features simply by implementing a different computing platform as opposed to going down to the full converged type data center network and storage environment.

It all depends on what the requirements of the client are. As an integrator, that's our first step -- what are the requirements -- and then matching technology and products to it.

Gardner: John Bennett, we are working towards network transformation. We're certainly seeing a lot of data-center transformation, bringing the two together in a cohesive, organized, perhaps harmonious way -- maybe that's wishful thinking. How do you get started on that? How do you bring these things together, and what are some of the initial steps you expect to see from people to do this successfully?

Bennett: Harmonious is actually something you can expect, and Lin has made note of the HP example here several times. That's definitely a nice outcome to have, and a possible outcome to have.

How do you get started? If you have the capabilities in-house to really do your own networking, architectures, and business process analysis, then take a step back and revisit what you've done in the past. Take a hard look at your networking environment and look at where you need to have the business in the organization running in three to five years. Basically, do the analysis to architect and build it out yourself. Then, make use of current generation tools and capabilities, as you do that.

Clearly, if you are virtualizing your infrastructure and moving in those directions if you want to automate a great deal of the data center environment in the business services you are running, you need to have a clean networking infrastructure as much as you have a clean storage and server environment. The ruthless standardization is foundational for doing that.

If you don't have that experience, if I were a customer, I'd be calling for help, because networking is one of the areas that is most challenging for me personally. Take advantage of a system integrator like HP and their capabilities.

Mapping your strategy

We have people who can come in and not tell you what to do, but work with you to map your business strategy to your IT and your data center strategies, and then look at what you should do over time, in order to change from what you have been to what you can be.

So, if you are self-capable, take a strategic look at it. If you want to take advantage of the experience -- and we have been doing networking since it was created -- take advantage of the experts, someone like HP, and really take that fresh look, and then the implementation and plans after that all follow, but focus on the strategy and the architecture first.

Gardner: Mike Thessen, any rules of thumb that you fall back to, when folks come to you and ask how to get started? What are the first things we need to start doing or thinking about?

Thessen: What we focus on is really developing a good strategy first. Then, we define the requirements that go along with business strategy, perform analysis work against the current situation and the future state requirements, and then develop the solutions specific for the client's particular situation, utilizing perhaps a mix of products and technologies.

One thing to note here is that HP makes networking product. We make great blade products. We have more or less everything a client would need, if it fits their solution. From our perspective in the Network Solutions Group, we know that HP solutions aren't going to fit in every case. So, we

are still one of the largest Cisco worldwide gold partners. We have a vast array of other partnerships in the network space to bring together the right solution for our clients.

Gardner: The last word today goes to you, Lin Nease. Tell me a bit about what your opening salvo is when folks come to you and say, "Wow, an awful lot to think about. How do we put this into a chunk that we can get started on?"

Nease: The advice to the network architect is to look at the portfolio of applications that you are trying to enable. Don't look at the data-center network the same way you look at the enterprise network. It is different. It is specialized. Consider strongly those unique special cases that you handle today as exceptions. Think of how you would handle them more as a mainstream provider.

Be honest with which applications are going in what ways and what demands will be on the network in the future. Also, look to simplify. Always assume that your first step as an architect is to figure out how to simplify what you are trying to accomplish in your data center network design.

Gardner: Very good. I want to thank you all for joining us. We have been on a sponsored podcast discussion today on transforming network architectures in anticipation of evolving demand.

I want to thank our panel. We've been joined by Lin Nease, director of Emerging Technologies, HP ProCurve. Thank you, sir.

Nease: Thank you very much.

Gardner: John Bennett, worldwide director, Data-Center Transformation Solutions at HP.

Bennett: Thank you, Dana.

Gardner: And, Mike Thessen, practice principal, Network Infrastructure Solutions, practice in the HP Network Solutions Group. Great to have you with us, Mike.

Thessen: Thank you, everyone.

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.

Offer: For more insight into best practices for transforming your data center to increase agility, lower costs and improve business continuity and availability, download three valuable Data-Center Transformation whitepapers from HP by going to: www.hp.com/go/dctpodcastwhitepapers

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

Transcript of a BriefingsDirect Podcast examining how data-center transformation requires a new and convergent look at network architecture. Copyright Interarbor Solutions, LLC, 2005-2009. All rights reserved.