Insight Series

EDUCATION LEADERS ON...

HOW INSTITUTIONS ARE MEETING THE COMING BANDWIDTH TSUNAMI

A massive growth in network bandwidth requirements in higher education is underway, driven by factors ranging from e-learning to research needs to rapidly growing residential networks. High-performance managed optical fiber networks are in use already by governments, commercial carriers, large enterprises, and many higher education and research institutions. As the bandwidth wave sweeps forward, how can you best position your institution's network to serve your users? *Campus Technology* interviewed a panel of experts in education technology networks, asking them to share their strategies for meeting future network demand.

CAMPUS TECHNOLOGY | ciena Level(3)

Participants

Charlie McMahon,
Tulane University

Paul Schopis,
OARNet

Mark Johnson,
MCNC

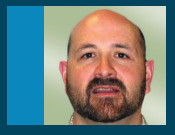

Lonnie Leger,
Louisiana State University and LONI

Steve Perry,
University of New Mexico

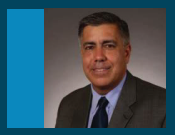

Jim Archuleta,
Director of Research and Education Initiatives, Ciena Government Solutions

Craig Cupach,
Director of Sales, Level 3 Communications

Scott Pohlman,
Senior Account Director, Level 3 Communications

1

At your institution (or the institutions you work with), what is driving the massive growth in bandwidth requirements – is it research, e-learning needs, residential networks, or addition of K-12 and government entities to the RON networks? If it's a combination of factors, which is the most pressing need, and why?

McMahon: Like most campuses, Tulane sees bandwidth demand from residential students, teaching and research increase every year. The increased demand from our users notwithstanding, the most important force driving massive bandwidth increases in higher education is economic. The cost of creating, processing, and storing data has plummeted, and so has the cost of provisioning bandwidth. Because costs have fallen so quickly and dramatically, many schools with 10G backbones are opting to skip 40G technologies and instead invest in 100G backbones.

Schopis: It's a combination of factors. First, in current research models, big data is king. Projects such as the Large Hadron Collider (LHC) in high energy physics, genomics, and climatology all depend on very large data sets. Moving the data around for processing requires very high bandwidth capacity.

Secondly, particularly at residential campuses, student expectations for online access are increasing dramatically. Most students arrive with at least three network-enabled devices and more times than not, students are no longer watching TV on a physical TV. They tend to use online access for their recreational content as well as their academic pursuits. K-12 is also gravitating toward online resources. The Ohio Department of Education funded ITCs and large urban districts upgrades from one-gigabit connections to ten-gigabit connections.

Last but not least, the state is embracing e-resources for the citizenry with the expressed goal of making it easier to use state resources online.

Johnson: MCNC has been operating NCREN for over 30 years. We have 30 years of traffic data and we are thus pretty comfortable predicting how much traffic will grow. At the same time, we've also learned over time that it is very difficult to predict exactly what will drive that growth. We didn't predict the music industry moving online, or YouTube or Netflix.

We do know a few things that are contributing to increase traffic: All science is becoming a digital enterprise; education is increasingly dependent on digital resources; healthcare is moving online with electronic health records, all diagnostic imaging systems are digital, electronic communications are common now with insurers, pharmacies, and others. Also, cities are embracing the Internet of Things to instrument and automate utilities. At the same time, more people in

> "THE COST OF CREATING, PROCESSING, AND STORING DATA HAS PLUMMETED, AND SO HAS THE COST OF PROVISIONING BANDWIDTH."
>
> —Charlie McMahon, Tulane University

all of those fields have computers, more of those computers are connected to the Internet, and the connections are becoming faster.

Leger: For us, it's a combination of factors. In the area of research, it's really about the flow of big data. Researchers need to move massive data sets to and from our computing center, back to their home institutions, or to another storage system for post-processing, then back to us.

From the e-learning side, more and more data is going into the cloud through applications-as-a-service or software-as-a-service. Also, there are more and more students who have no concept of life without wi-fi — they've always had a cell phone. Their high expectations for both connectivity and content is definitely driving bandwidth demand.

Perry: At UNM, drivers are as follows: Research, e-learning, and planning for the addition of K-12 and government entities on the RON. Our most pressing driver is primarily research, because our research community has levied requirements to provide high-speed connectivity for computational data sets, with both local data transfers and the desire to work with other higher education institutions throughout the US. We have received National Science Foundation (NSF) funding to build the local high speed "ScienceDMZ" network, and have received a separate NSF grant (in partnership with New Mexico State University and New Mexico Tech) to set up regional planning workshops to assist the New Mexico comprehensive schools.

Cupach: It certainly is a combination across those things. You have students entering campus with several network connected devices. The broadband networks on campus become their entertainment, whether it's an xBox, Netflix, Hulu, or Pandora, there is a concentration of bandwidth consumed from the entertainment perspective. As kids break for the summer, we see IP utilization goes down on campus network.

Also online education is taking off, both from remote or strictly on-line students. On-line education is also becoming part of curriculum. One school we work with, recently mentioned they have a goal to double enrollment over next two years by adding everybody online. It's more cost effective and their potential addressable market for students around the globe becomes larger. Research and collaboration also drives bandwidth. We are also seeing large bandwidth demands from teaching hospitals.

Archuleta: At this point science is global endeavor. Any university has needs and requirements to connect with national labs to drive progress and divide the challenge. Networks are the face of this strategy.

Pohlman: Two of our largest customers operate large backbones. Their traffic is all about research. They don't have any student traffic, it's all research traffic and it grows by factors of ten every 4 to 7 months. They see that continuing and starting to accelerate. And it's pure science data.

"THERE ARE MORE AND MORE STUDENTS WHO HAVE NO CONCEPT OF LIFE WITHOUT WI-FI — THEY'VE ALWAYS HAD A CELL PHONE. THEIR HIGH EXPECTATIONS FOR BOTH CONNECTIVITY AND CONTENT IS DEFINITELY DRIVING BANDWIDTH DEMAND."

— Lonnie Leger, Louisiana State University & LONI

2

Beyond network growth drivers today, what additional current and emerging applications do you see driving network needs at your institution in the future? What about outside your institution, in higher education in general?

McMahon: As I mentioned in the previous answer, the falling cost of provisioning bandwidth has made 100G affordable for many institutions. With the exception of a few academic researchers who deal with very large data sets, these high-bandwidth networks exceed the needs of most applications and users. Hence, this is not a situation where applications are driving the need for speed. Rather, abundant, low-cost bandwidth will lead the way for new applications to emerge that take advantage of the high-bandwidth networks.

Schopis: Although it's too soon to tell for sure, a trend toward more automation and use of sensor networks seems to be emerging as a growth area. Allowing physical systems to respond to changing conditions proactively and in coordination with the rest of campus seems to be a desire of large research institutions. Perhaps it's currently an overused term, but the notion of the Internet of Things (IoT) seems intuitively correct.

Johnson: At MCNC, we are not a campus, of course, but a regional network supporting education, healthcare, government, and public safety. Right now, healthcare is the fastest growing sector in terms of new connections, and K-12 education is growing traffic fastest. We think public safety will soon account for significant traffic as next-generation 911 systems and the FirstNet network for first responders come online. The Internet of Things (IoT) will soon dominate the number of connected devices across all sectors. Moore's law will ensure that the traffic-generating capacity of each device will grow rapidly over time. IoT also has implications for what we think of as "mobile" devices. Many IoT devices may be fixed in location, but connect wirelessly.

Leger: Everything we read says that higher education is changing drastically as the culture and the audience changes. So how does a public institution like ours get funded appropriately to meet expectations?

Customer demographics are indeed changing; they are much different than they were 25 years ago. But beyond that, there are other the drivers. One huge thing I see driving network speed is the personalization of the user experience within the network. New users will be growing up in an age where the user is not going to be just an individual among the masses. The way a user interacts with the network is going

> "ABUNDANT, LOW-COST BANDWIDTH WILL LEAD THE WAY FOR NEW APPLICATIONS TO EMERGE THAT TAKE ADVANTAGE OF THE HIGH-BANDWIDTH NETWORKS."
>
> —Charlie McMahon, Tulane University

to be much more personal. The network is going to be much more personal, more contextual, more about where you are, what you're doing, and how you're doing it. It's going to put the user in a software-definable experience. It might be like Netflix, which learns your behavior and suggests what you should be watching next. We're going to go to more and more of those concepts in personalizing the experience of the network user.

Perry: I'd have to say cloud-based applications, in response to both questions. For instance, UNM migrated to O365 and saw a marked increase in commodity Internet-based traffic. We are also looking at hybrid on-premises/cloud computing and storage solutions, in order to determine their feasibility. As we examine the pros and cons, we are looking for a balanced approach. We are seeking vendors that can offer a robust solution that can meet a cost model beneficial to the university and any of our customers and partners.

Archuleta: There's certainly a national trend toward adopting the cloud model. Universities today are exploring hybrid cloud models where certain data may be close by or on campus and other data in the cloud. Many applications are well on their way to being entirely virtualized. That drives bandwidth and performance requirements.

Cupach: We're also seeing a lot of hybrid cloud models. Another is the Internet of Things. Everything is becoming a connected device. There may be different vending machines with an IP interface that provides data for finance or IP security cameras. There are more devices connected to the Internet driving different types of applications. Even more essential than network growth, is how to secure that network. Network security is a huge topic right now. There's an amount of risk protecting students, protecting data, and intellectual property that might be a target for bad actors.

Pohlman: There are two ways big networks are dealing with security. Some are turning to third parties to assist in monitoring & protecting their network, while others have just bought bigger pipes, however that doesn't ultimately address the route cause. Being able to protect your network both at the edge and at the core is a best practice.

"MANY APPLICATIONS ARE WELL ON THEIR WAY TO BEING ENTIRELY VIRTUALIZED. THAT DRIVES BANDWIDTH AND PERFORMANCE REQUIREMENTS."

—Jim Archuleta, Director of Research and Education Initiatives, Ciena Government Solutions

3 What specific benefits will 100G bandwidth offer for your user community?

McMahon: A 100G network will remove one more barrier to collaboration and cooperation. Specifically, researchers who rely on transferring large data sets will realize the benefit of faster networks immediately. Eventually, all users will benefit from increased network speeds. And, as I said earlier, 100G networks will open the door for new applications that can mine and analyze enormous amounts of data that will enrich all our lives.

Schopis: Benefits have been several-fold. First, OARnet members have seen an uptick in the number of cyberinfrastructure grants awarded. At the national level we are doing quite well. Being able to cite our investment in infrastructure appears to play well at funding agencies.

Secondly, having a 100 Gbps network allows us to turn up services that are logically separate but share the same transport platform, in a much more timely manner. We are able to employ network virtualization techniques to meet all of the needs of our community, yet control costs. The other benefit is that we are dealing in such large bandwidth, we are able to use super aggregation, if you will permit the term, to drive cost per unit down.

Johnson: As with every technological step in capacity, we will see specialized research users adopt 100G first, along with links aggregating traffic from many sources. As the cost of 100G parts drop over time, 100G will spread into other parts of the network.

Leger: For the most part, I think we'll see benefits on the backbone. Again, 100G speeds will allow our researchers the capacity to transmit very large data sets quickly, file sizes that are terabits and larger. Researchers have an expectation of connectivity that was commonplace 15 years ago, and we may be able to supply that with gigabit connectivity on the backbone. We're trying to extrapolate in the 100G arena, and handle those data sets appropriately. So that's probably where our most specific benefit will be.

Perry: A few of the biggest benefits of 100G are: Research computing and data transfer; enhanced mass user video streaming and distance learning capability both on our main campus, and between our main and branch campuses; and delivery of e-learning capabilities.

> "100G NETWORKS WILL OPEN THE DOOR FOR NEW APPLICATIONS THAT CAN MINE AND ANALYZE ENORMOUS AMOUNTS OF DATA THAT WILL ENRICH ALL OUR LIVES."
>
> —Charlie McMahon, Tulane University

Cupach: A bigger pipe certainly removes the network barrier for researchers to do various projects. Think about it from the other side, though. The university may want to throttle that bandwidth. That gives them some control over how much bandwidth they would allow for various users. It gives universities a path to easily scale bandwidth for users in a cost predictable manner, allows users and or applications to access network features that might not be available otherwise.

Archuleta: At this point, 100G is becoming pervasive. That helps drive converged architectures and better price points.

Pohlman: It potentially lets you simplify the network from a design and architecture perspective. When you have a lot of bandwidth, you just overprovision for different applications because you have that "limitless" bandwidth at your disposal.

4 What are the key build-versus-buy considerations you face in growing your college or university's network?

McMahon: The main considerations when choosing whether to build or buy include long-term cost and network flexibility. Higher education has a long history of building and operating high-bandwidth networks at costs that are much lower than the same bandwidth and services could be purchased from traditional Internet providers. By choosing to build, an institution can design the network to specific institutional needs and can more easily upgrade equipment or add new services as institutional needs change.

Schopis: It really comes down to cost/benefit. We found it necessary to build our own network because of two factors. First, our performance requirements are usually ahead of where vendors are able to deliver. For example, when we lit our 100Gbps backbone 100G was not commercially available. That is only the most recent example, and perhaps a little trivial. One of the motivations for an R & E network is to provide a testing ground for emerging technology, if we are successful, it will be incorporated into the commercial vendor offerings. What we have done historically is applied research, i.e. it should work and theoretically it will, but it has not been tested to scale. Back in 2004 when we lit the original dark fiber network we were the first networks to use an MPLS/Logicalrouter architecture. We tested it thoroughly in the lab and were pretty certain it would work at scale.

Secondly, it is simply cost. When we built the original dark fiber network, we predicted we would hit the break-even point in seven years – we hit it in five. As I mentioned earlier, by aggregating and making bandwidth plentiful, we have been able to drive per-unit costs down.

Johnson: It used to be that it made more sense to build computing capacity and buy network capacity. Today the opposite is true. Services like Amazon Web Services and Microsoft Azure are much more cost-effective than running private systems. On the other hand, once a network reaches the point where wavelength-division multiplexing (WDM) systems are required, then it makes sense to build your own.

Leger: We have to consider the longer-term use of the capital investment that we might be undertaking. In a sense, we are a niche market – not much different in a capital investment strategy sense than some ISPs and telcos.

Also, sometimes the things that we want to do just aren't available in the market to buy. Sometimes it's too costly to do a build scenario; sometimes it isn't. So one

> "WHEN WE BUILT THE ORIGINAL DARK FIBER NETWORK, WE PREDICTED WE WOULD HIT THE BREAK-EVEN POINT IN SEVEN YEARS — WE HIT IT IN FIVE."
> —Paul Schopis, OARNet

consideration is availability, and a second one is what restrictions exist even if what we need is available. In higher education, we tend to be very philanthropic with our investments – community enrichment, community engagement and so forth. We're not talking about strictly a vertical market and a commercial container. Some services we might otherwise consider have difficulty with the far-reaching impact of our mission as an educational institution.

And, as with any decision, we also have to consider speed to market. Can I deliver this in a relatively rapid manner? That also determines our build-versus-buy decision. I tell many vendors that we are a unique vertical market. We are not like others, and we appreciate their consideration as such. So that also goes into our specific consideration of key technologies.

Perry: We have to consider both available capital funding vs. leasing costs (Capex vs. Opex), and the return on investment. For example, by investing in the "last mile" build, we were able to use a franchise agreement to connect a branch campus. That investment will eliminate about $8,000 a month in carrier costs once the fiber is lit. The ROI was a fairly easy decision. When the upfront costs are higher, the decision becomes much tougher to justify.

Another consideration along the same lines is funding and staffing to maintain any owned network assets.

Cupach: A lot of customers look to build because of the ownership model. Being able to own the asset has value, both from a predictability and financial perspective. You have to make sure to factor in not only the cost of the asset, but also what its going to take to operate and manage, and stay current with security patches. Do you have the skill set to run that and is that your core competency as a university? That's the beauty of a managed optical fiber network solution. It lets users control the network and delivers a differentiated IT experience.

Archuleta: And the industry has customized an approach to fit the situations of every university, where certain things are offered as a service others as an asset.

5

How are you funding network upgrades? (For example, are you using research grants, NSF grants, ERATE, or a combination?) How was the decision made to fund the latest upgrade? Briefly, what challenges have you faced in finding funding?

McMahon: Two NSF grants provided the funding for Tulane's research network; Tulane's production network upgrade is funded centrally by the university. I think a centrally funded network is a good model but it puts the network at risk when times get tough. For example, the Tulane network was scheduled for an upgrade when Katrina hit New Orleans. The financially tough times that followed delayed the network upgrade another six years.

Schopis: Our funding model at OARnet is straightforward. The state provides OARnet with an operating subsidy and capital subsidy. The balance of the cost is made up through cost recovery and user fees.

Johnson: We build network upgrades into our rates so that we are not scrambling for capital at the last minute. Every piece of equipment has an expected lifetime and we plan for replacement at the end of that lifetime. We also do our best to extend the life of our gear beyond its nominal life by doing things like redeploying devices elsewhere in our network as they are upgraded.

Leger: Oh, there are challenges, very much so. I can't speak for the private post-secondary arena, but in public post-secondary, there's been an evolution in politics, which drives funding. Whether it's coming from a federal program like the National Science Foundation, or the federal E-Rate program... it's definitely a challenge. With the federal government, whether the Department of Education or the Department of Commerce, we're often talking about 50-year-old-plus legislative acts by Congress, so they are constantly under pressure for policy changes. On the public side, funding can come from the local economy – the taxation policy – and with any given state and or national policy, that has a huge impact.

Given all of that, it really comes down to having a core group of like-minded stakeholders in the public domain that maintains a necessary level of financial commitment. Without that sort of comprehensive and cohesive ecosystem approach to the problem, we would definitely falter. We'd be in a state of chaos around the work we do collectively. Also, the public looks to us for comprehensive uses in post-secondary education. So our ability to present a uniform front – both in policies and programs, and in funding strategies – is critical to keeping the funding progression going in an ever-changing technological landscape.

"IT REALLY COMES DOWN TO HAVING A CORE GROUP OF LIKE-MINDED STAKEHOLDERS IN THE PUBLIC DOMAIN THAT MAINTAINS A NECESSARY LEVEL OF FINANCIAL COMMITMENT."

— Lonnie Leger, Louisiana State University & LONI

Our latest upgrade, to 100G, was a combination of several grants, and going back to my programmatic sponsor – the board of regents – I'd go back to them and communicate my need. I had the seed money from the grants, and then we petitioned the board and got them to match funding coming in at the institution level. With that approach... working together as an ecosystem, we were able to fund our latest 100G investment. Going forward, that's a good recipe for success. We partner with our institutions that want to achieve better connectivity, and mutually come up with some shared funding approach to the investment.

Perry: At UNM, we are funding network upgrades in a couple of ways. First, by using NSF grants -- we have received funding to build a ScienceDMZ via the CC*IIE program. Second, we use a "refresh fund" from our collaboration partners. This funding goes before the partners and we agree on the best effort to increase our capability.

Cupach: There are challenges, not only in higher education but across the board, to do more with less. They have to support more users, provide higher bandwidth, and they have to do it with the same amount of funding and protect against security threats.

Archuleta: A lot of our customers are small organizations with a large charter. They're looking for ways to own and operate a network that's failsafe and doesn't get them into hot water. We hope the network is considered an asset for the university—the offensive line of IT. We hope it's an enabler for executing on a wider scope of their mission and an extension of the campus to conduct broader reaching research.

6

SDN (software-defined networking) promises to enable orchestration of network infrastructure to create a virtualized pool of connectivity resources. What would a more automated and agile network mean to you? How would this change your ability to satisfy the needs of your user community?

McMahon: SDN is part of a trend that some refer to as SDX – Software Defined Everything. As we virtualize the entire stack – compute, storage, security, network – we make it possible to create applications that are aware of the available cyberinfrastructure. From there, it seems inevitable that developers will write autonomous applications that secure the compute, storage, security, and network resources needed to execute intended functions.

Schopis: SDN is another overloaded term these days – what it means to different vendors varies. The one item that is overarching in all of the conversations is the notion of orchestration. I think that is an idea whose time has come, and is a bit of a natural extension to our current provisioning model and tools. I was at a vendor event a year or so ago and a speaker asked the audience to raise a hand if you were a network person. He then quipped that "you have always been in the cloud business but didn't know that was what was called."

Similarly, SDN has actually been around for a while, it just wasn't as coherent an approach as it is now. It will take a little time for the standards or de facto standards to completely emerge, but efforts such an Open Daylight seem to be taking a balanced approach.

Johnson: MCNC is collaborating with Internet2 and our local universities and research centers on the use and operation of SDN infrastructure. We don't know yet what the most important benefit will be. SDN has the potential for disrupting the market for switches and routers by commoditizing some components, which may have huge economic benefits.

Leger: Among those who really understand SDN and what the acronym means, I predict lots of differing opinions. My feeling is, beyond network growth, as I said earlier, we're going to see personalization and individualization of the network as an experience. That's more my opinion as an individual than at an institutional level. In my opinion, it will allow the institution to create an individual perspective, to provide a privatized network experience... Based on my experience, that's what I'm anticipating SDN will enable.

> "SDN HAS ACTUALLY BEEN AROUND FOR A WHILE, IT JUST WASN'T AS COHERENT AN APPROACH AS IT IS NOW. IT WILL TAKE A LITTLE TIME FOR THE STANDARDS OR DE FACTO STANDARDS TO COMPLETELY EMERGE."
>
> —Paul Schopis, OARNet

Perry: In a nutshell, a more automated and agile network means much faster provisioning, and simpler and more robust management. By creating profiles and templates, and spinning virtual resources up in a matter of minutes, we can be very quick with our process. SDN will certainly change our ability to support our users and meet tighter SLAs. It will also allow us a user-based "self-provisioning" option that gives us a higher degree of confidence that a user will be successful without our intervention.

Archuleta: In some sense, that's the next step in driving down a lot of cost in network architecture—both in terms of the platform and operating networks. "Also, SDN powered applications will allow users to learn and collaborate in new ways."

Cupach: A big part of network DNA is being able to dynamically provision capacity and manage services. At the university level, that helps downstream with researchers. Being able to dynamically provision things and turn up and down bandwidth as needed gives that higher level of user experience. It reduces barriers and creates innovation.

Pohlman: You don't want users to have to think about bandwidth. Give them a model to use on demand, and they're much more apt to use it and be more creative.