

Expanding the Community

Serving Community Anchor Institutions for BTOP, U.S. UCAN and beyond

John Moore, MCNC
Bill Owens, NYSERNet

October 25, 2010

Introduction

This paper is presented to the NTAC in order to provide a framework for discussion about the technical, operational and policy attributes required to support several new constituencies mandated by the NTIA BTOP program. We approached this task by first illustrating the traditional differences between R&E networks and commercial carriers, then examining the opportunities and challenges presented by the new environment and the ways BTOP-funded organizations may need to change their R&E networks and operations in order to satisfy the requirements of the BTOP program.

There have been numerous regional BTOP awards, in addition to the national award to a team led by Internet2. The nature of the relationship between the national and regional R&E networks is such that it would be difficult to discuss the requirements these new constituents place on the national project (U.S. UCAN) without also looking at how the regional networks will need to respond. Thus, this paper discusses both Internet2 and the regional networks.

Landscape

Research and Education networks in the US typically cite the same set of distinguishing features to separate themselves from the 'commercial' or 'commodity' Internet. R&E networks support advanced protocols like IP multicast and IPv6. Connections to the network are restricted to members of the community even if, as with many regional networks, the services include commercial Internet access in addition to research and education traffic. They have more direct routing between endpoints, which reduces end-to-end latency, and their minimalistic design and ample capacity reduce jitter. Commercial networks, driven by different economics, must utilize their facilities more fully and wait until backbone and peering links are near-full before upgrading them; R&E networks typically plan to have at least half of their link bandwidth unused, based on average utilization. This excess capacity allows the use of experimental protocols that make unusual bandwidth demands or don't tolerate congestion.

Those are only the surface differences, however. The design and operation of the R&E infrastructure is distinct as well, driven by the technical requirements for the provisioning of those special capabilities, and as a natural reaction to the different environment of the R&E world.

Motivation

The R&E networking community in the US has been cited by the federal government as a very successful model that could serve as a template for serving a specific set of important constituencies, first described in the FCC's National Broadband Plan as Community Anchor Institutions. Many of the regional networks, along with Internet2, have been recipients of NTIA

grants to extend their infrastructures to support the entire spectrum of CAIs.

The NTIA BTOP program clearly defines this set of community anchors: it is comprised of community colleges, schools, libraries, medical and healthcare providers, community support organizations, and public safety entities. In the BTOP2 NOFA, the NTIA gave specific priority to projects that serve community colleges and public safety entities.

Many regional networks already serve a subset of these CAIs, and have been doing so for some time. They have evolved their service offerings to meet the needs of these new groups. Programs such as Internet2's Sponsored Education Group Participants (SEGP) program and the FCC's Rural Health Care Pilot Program (RHCPP) encourage these non-traditional constituents to participate in national R&E networking, through the regional network serving their area.

Those regional networks that are primary recipients of BTOP grants are committed to serving CAIs, based on the details of their individual proposals. Others that are sub-recipients, or have not received BTOP funding, will likely have varying degrees of interest in expanding their constituency to support some or all of the CAIs in their region.

Attributes

The following sections discuss technical, operational and policy attributes that we feel are important to be considered as the community evolves to support the broader mandate specified in the National Broadband Plan and embodied in the U.S. UCAN project. The list of attributes is not exhaustive, but reflects our thinking at this point in time (October, 2010).

Bandwidth and service quality

The original distinction between R&E and commercial networks was their size - the R&E path, all the way from the campus to the backbone, was typically much larger and less heavily utilized than in the commercial world, and as a result the performance for all applications was improved. Those users with significant demands for sustained high bandwidth, low latency, low jitter or large headroom had been unable to meet their needs with commercial Internet connections, but were able to make the applications work across the research networks.

Those among the new CAIs with similar applications will find the same benefits; perhaps more importantly, they will be able to deploy applications like medical imaging, high-quality conferencing and telepresence, database sharing and real-time applications that were never available to them via commercial Internet connections. The growth in participation brought on by the BTOP funding at regional and national levels will challenge R&E network operators to continue to maintain this technical advantage and distinction from the commercial providers.

Although the R&E community has largely abandoned IP QoS, there may be instances where its use is desirable. Beyond the traditional QoS model, however, R&E networks may provide differentiated services, for example by allowing research traffic to consume bandwidth in preference to commercial peering traffic. New communities may have similar requirements: public safety agencies may request absolute priority, or healthcare consortia may need specific amounts of guaranteed bandwidth for critical patient care applications.

To the extent that CAI requirements dictate consistent end-to-end service quality across one

or more regional networks and U.S. UCAN, the community will need to work together closely to agree on reasonable standards for implementing and managing such services. Much of the work already invested in dedicated services for the research community will be a useful starting point. Indeed, developing and managing such a service, spanning multiple backbones managed by different organizations, could provide a key distinguishing feature for U.S. UCAN. This would have the unfortunate side effect of reducing the status of R&E traffic from primary to at best co-equal with other high-priority traffic, or potentially pushing it to a lower priority level. Some objection from the existing R&E community might reasonably be expected if such an arrangement is proposed.

Service characteristics

The current offering from Internet2 to the regional networks consists of several separate services: R&E IP, commercial peering (CPS) and dedicated circuits. This arrangement is generally flexible enough to satisfy the needs of both parties without unduly burdening either with a confusing and operationally challenging configuration. The offering can be provided on one physical interface, or provisioned across two physical interfaces which can be in separate locations.

Some regional networks choose to keep each service separate when they provision their downstream customers, particularly the R&E IP and commercial peering services. Other regional networks provide a simple blended IP service to their smaller constituents, consisting of R&E IP, commodity IP, and CPS. Those same regional networks will also provide research institutions with (for example) a blended IP service *and* access to dedicated circuits. The permutations are numerous.

As new constituents are added, and the relationships between participants in the U.S. UCAN community become more complex (e.g. arrangements other than the traditional national-regional-campus hierarchy), the challenge will be to maintain service flexibility and allow for customization where needed. It is likely that the regional networks will continue to provide customized services to the new CAIs they serve, while also continuing to rely on a flexible upstream service from U.S. UCAN. In rare cases where U.S. UCAN provides service directly to a CAI, it is likely that a simple blended IP service will be the optimal solution.

Defining specific new offerings will require more significant engagement with the new CAIs at a technical level. As those discussions occur, it is important to keep in mind the overarching goal is to provide a service that has significant value over and above standard commodity Internet service. In many cases that service will have the characteristics outlined in the previous section that the R&E community has long been enjoying; in other cases the value proposition will be defined by the unique needs of the CAI in question.

Scaling and reach

R&E networks optimize for few customers with large connections; commercial operators plan for many customers with smaller connections. An R&E customer typically connects directly to a backbone router, which minimizes latency and jitter, and makes the network simpler to administer by reducing the number of 'moving parts' in the design. Commercial networks must employ multi-level connectivity to meet their scaling demands. Home and small business customers connect to aggregation devices (appropriate to the technology - DSL, PON, or cable), which connect to provider-edge routers; larger customers also connect there. Those routers connect to the backbone, and will frequently have multiple uplinks to diverse backbone

sites for redundancy and traffic tuning.

The two national R&E backbones taken together have a total of eighteen backbone router nodes, but with overlap they cover just eleven cities. Adding transport locations, they reach a total of 27 cities. Independent regional networks (those not funded by state government) tend to be similarly sparse. Commercial ISPs not only overlap with each other, their networks are much denser. As an illustration of that density, though underserved areas remain across the country, the FCC's National Broadband Plan reports that 95% of Americans have access to high-speed Internet access.

A major component of regional BTOP proposals is a large expansion of the regional operating footprint. BTOP 1 focused on pushing backbone middle-mile services to underserved areas. BTOP 2 continued that effort, and added the additional requirement of connecting CAIs directly to the middle-mile backbone where possible.

In most cases, the addition of a large number of new users connecting to a significantly expanded footprint will cause the regional to rethink their architecture. The extent to which existing architectures need to change is not yet clear, since there are still many unknowns regarding the size and scope of the new CAI population. As we ponder the implications of a large increase in the number of constituents who use our networks, one recent phenomenon in the R&E community may serve as a constructive example: the addition of K-12 schools to regional networks that previously served mostly higher education.

A number of regional networks have connected K-12 schools to their IP backbones. In many cases, the physical connection is between the local school district office and the backbone; the existing district network aggregates all the school connectivity within that district. Leased lines from a local carrier are often used to connect between the district office and the backbone, and to knit the schools and together district-wide.

The total K-12 population can be quite large, in the millions of potential end users, but the relatively small number of districts (tens or hundreds, up to just above one thousand, per state) helps to limit the number of direct backbone connections. This preserves the overall backbone connection architecture, a small number of large participants connecting to a sparse backbone.

Regionals that have won BTOP awards will expand their footprints in two major ways that will effect K-12. First, they will increase the number of PoP locations. Second, they will be replacing some (but likely not all) of the carrier circuits with direct lateral fiber between the backbone and K-12 district offices. In most cases the configuration of the intra-district networks will not be affected by BTOP builds (though perhaps by other changes, such as the new E-Rate rules).

In general, BTOP builds will allow regional networks to expand the number of PoP locations at which their service is available, and increase the number of downstream connectors. The connectors themselves, however, will continue to provide aggregation of their users (school districts, hospitals, community colleges, etc.) and the BTOP-financed expansion will not necessarily cause the regional networks to add a proper aggregation layer. Last-mile providers will continue to function as aggregators, where needed.

Middle-mile services: open access and pricing

Over the last several years, as regional and national R&E networks began to construct facilities-

based backbones, it became possible for them to offer transport-layer connectivity to their communities. These services may include DWDM lambdas, multiplexed Ethernet over SONET circuits, or Ethernet VLANs. Generally similar services are also available from some commercial providers, but with the important difference that R&E circuits are typically priced at well below market rates.

These services may be of great interest to some of the new CAIs. For example, public safety agencies building next-generation wireless networks may want Ethernet backhaul from local towers, or a health consortium might wish to build a mesh of paths between facilities with particularly demanding requirements for bandwidth, privacy or security.

It is important that U.S. UCAN and the regional networks consider the implications of providing such services to a broader community. By accepting federal funding, BTOP middle-mile awardees must comply with interconnection and non-discrimination requirements and, according to the BTOP NOFA, all awardees must:

... offer interconnection, where technically feasible without exceeding current or reasonably anticipated capacity limitations, on reasonable rates and terms to be negotiated with requesting parties. This includes both the ability to connect to the public Internet and physical interconnection for the exchange of traffic.

It is likely that BTOP awardees within the R&E community have different interpretations of the NTIA's open access requirements, and significant discussion may be needed before consensus can be reached on a harmonized view of open access.

Internet2 has stated that the open access requirement for U.S. UCAN can be satisfied with CPS: companies wishing to interconnect can do so on a settlement-free basis by complying with existing peering rules that Internet2 has established. Access to the IP network serving new and existing CAIs, as well as circuit services, will be limited to users that fit the definition of a CAI.

Regional networks could take a more expansive view of open access by providing dark fiber and circuit services to CAIs and last-mile providers alike. Individual regional networks will likely make this decision based on a number of factors that might include the local political climate, local market conditions, the extent to which the organization is chartered to foster economic development, etc.

The BTOP NOFA also includes this requirement:

Applicants should demonstrate that this pricing is competitive and affordable to their target markets. However, pricing levels also should be reasonable and market-based, so as to maximize the efficient use of Federal grant funds.

Open access requirements and pricing are intimately linked by market realities.

Many facilities-based R&E networks, both regional and national, make circuit services available to their members at the cost of equipment, by covering other costs on the basis of membership fees, IP connections or other revenue sources. These services are therefore only available to members, but are offered at well below market pricing as a benefit of membership and network participation.

Prior to BTOP funding, most R&E network providers had only enough fiber assets to deploy a single optical path on a given route. Thus, there is not a significant amount of experience within the community with being a fiber provider (apart from the brokering service, FiberCo). Today, however, many regional networks who are building new BTOP-funded fiber routes will likely have a significant surplus of fiber on popular routes.

In both cases (circuits and dark fiber), the community must contend with a limited resource that needs to be managed properly. Unless the price of service includes enough margin to cover upgrade costs, pricing that is substantially below market rates may result in resource exhaustion in short order.

Risk mitigation strategies that the community might consider are tiered pricing for circuit services, based on a specific class of constituents investment in the base (non-BTOP funded) infrastructure, or set-asides for specific community participants with respect to dark fiber, with the remaining assets available for last-mile providers.

The importance of coming to a community consensus on these issues should be emphasized: *the new business of wholesaling middle-mile services will present regional networks and U.S. UCAN with significant challenges.*

Interactions with other BTOP awardees

Not every region of the country is covered by a BTOP award, and some existing regional networks are not BTOP recipients; even those that are may not be able to serve every institution in their footprint in an efficient and cost-competitive manner. As a result, some CAIs will not be directly connected to regional middle-mile BTOP awardees. In some cases, last-mile providers will bid to serve those CAIs directly. Those last-mile providers may request middle-mile services from regional BTOP awardees that would allow them to participate in U.S. UCAN. In such cases, the regional (and U.S. UCAN) must be able to provide a service to the CAI through the last-mile provider that is equivalent to the service provided to a directly-connected CAI.

Since U.S. UCAN is considered a “national middle-mile” facility, it is possible that some commercial providers (last-mile or middle-mile) may wish to serve CAIs by getting services directly from U.S. UCAN. In geographies already covered by existing regional networks, it is important that the regional have the right of first refusal in serving the CAI in question. If the regional does not choose to serve that CAI directly, other providers may be acceptable as alternative connectors to U.S. UCAN, but the other providers should augment rather than compete with the regional. Open communication and coordination between all parties involved is key in such a situation.

Advanced services

Universities were for many years the only significant users of IPv6 in the United States, and the R&E networks were the only infrastructure that supported it. They have more recently been joined by federal government agencies operating under IPv6 deployment mandates, and by a few commercial content providers (notably Google). Regional and national backbones have years more experience in building and operating IPv6-capable networks. As deployment spreads, this expertise will be invaluable - particularly for those CAIs whose internal networking capabilities are more limited than the traditional R&E community. One likely outcome of the expansion into new constituencies is a greater involvement in address assignment and related

operational activities (address space management, WHOIS, reverse DNS, and complaint handling).

Native IP multicast support has also been a hallmark of the R&E networks at all levels, and although usage has never been extremely high, the applications that do take advantage of this capability find it essential. Few commercial ISPs are willing to support inter-network multicast at all, and although some networks use it internally for converged video services they do not make it available to their customers. New communities will be able to take advantage of multicast for the first time in the UCAN environment, leveraging the ability to connect large numbers of like institutions to a set of networks that uniformly support multicast capabilities. As these applications are deployed, regional networks will have increased involvement in education, planning, deployment, and especially troubleshooting.

There could be a reaction against the deployment of advanced services, given that there is always a tension between the network's ability to support them and its operational complexity. With every additional service layered on top of the basic IP network, there is a larger potential set of bugs, with the prospect of affecting both basic and advanced services; a greater chance that otherwise benign changes to the network will break the advanced service and necessitate additional, potentially disruptive maintenance; and an increased risk of simple mistakes, given the intricacy of the network's configuration. All of these issues reduce network stability. Multicast seems particularly vulnerable to this phenomenon given the relatively low uptake in the existing R&E community, its high rate of operational problems and the inherent risk of abuse; a new community that sees no need for multicast might push for its removal from the network. Indeed, these concerns are frequently cited by commercial ISPs as their reasoning for not deploying advanced services, and particularly for not matching the R&E community's pace of development.

Support

An often overlooked aspect of this expansion of user base and services offered is the change in the support model that each regional will need to consider. An increase in the number of served endpoints, many with less sophisticated local support, may require a higher level of centralized support or a transition to a well-coordinated distributed support model.

The National Broadband Plan, referencing the brief presented by Internet2 and SHLB that proposed U.S. UCAN, recommends supporting and assisting CAIs in obtaining and utilizing broadband connectivity. The plan notes that:

(m)any community institutions lack the institutional resources to undertake the many tasks necessary to maximize their utilization of broadband. Facilitating collaboration on network design and how best to utilize applications to meet public needs could result in lower costs and a far more efficient and effective utilization of broadband by these institutions.

Internet2 currently provides some services that could help fulfill this requirement in a high-level sense. Open source tools developed with Internet2 resources such as Shibboleth, perfSONAR, etc. provide reasonably well-supported applications that can make life easier for those CAIs that have the ability to help themselves.

Providing real value to the CAIs in this area, however, will require the equivalent of an on-site consulting service for networks (and perhaps network-enabled applications) that can assess

local CAI needs and make recommendations for how best to exploit the service being offered by the regional and U.S. UCAN. Some regional networks that have both a service provider and community convener role may find this approach particularly attractive. Although this goes far beyond the service provided by most regional networks, there is at least one models that is currently being successfully employed.

MCNC, the regional network in North Carolina, has a Client Network Engineering group whose mission is to provide network consulting services to K-12 and Community College customers in order for them to maximize the value they get from NCREN. The small CNE team provides network and security assessments, as well as troubleshooting and mitigation advice. Funding is provided centrally on a cost recovery basis.

Most regional networks depend on a limited number of technical staff working normal business hours, with on-call or outsourced call center services for after-hours support. Despite the significant increase in operational costs, manned 24/7 help desk and network operations centers will likely be required to support mission critical, high-reliability services for health care, public safety, and other entities that deal with life and death situations. Reliance on the network outside of business hours will constrain the planning of maintenance work and upgrades, and require that service parts and technicians be available on short notice.

Flexibility

Like most businesses, commercial ISPs strive for efficiency and lowered costs by standardizing their service offerings. They must also protect the operational integrity of their networks by enforcing demarcation points between customer and provider, and preventing customer network problems, misconfigurations or misbehavior from affecting the provider network. R&E networks, with much smaller customer bases, can afford to customize to a much greater degree, and to offer services that commercial ISPs would consider too risky or too likely to reveal details of the ISP's network and operations.

This customization capability will be attractive to individual communities within the new classes of anchor institutions, each with their own needs, and each accustomed to different services and network types (for example, the dedicated networks often used by public safety entities today). The ability to provide tailored service offerings will attract participation, ease the transition from existing services to the new infrastructure and help to fulfill the promise of the BTOP program.

At the same time the growth of participation will make customization more difficult, and the need to maintain the other characteristics of the R&E environment while increasing the depth of support, adding resiliency and controlling costs will tend to push R&E operators more in the direction of commercial ISPs and away from extensively customized offerings. This is likely to result in a loss of some of the flexibility currently seen in the R&E networks at all levels, though the regional networks will feel the most pressure given their smaller staff sizes and budgets. Resource constraints will force the network operators to prioritize their customization opportunities and eliminate those that are too costly or risky, or which benefit only a minority of their users.

Resiliency

Most Internet users have extremely low tolerance for failures of their connections, even at home - and especially if their ISP is also their television and telephone provider. Although many single points of failure remain at the edges of most ISP networks, their designs attempt to minimize the

effects of any failure, and to avoid outages entirely in the core of the network through designed-in redundancy and resiliency. Major customers can and do pay for redundant facilities, or purchase connectivity from multiple independent sources.

Those institutions with separated commercial Internet and R&E service typically treat their research connections as important, but not essential, and are much more tolerant of service outages because the commercial Internet offers a lower-performance but acceptable substitute during the outage. They almost always have only a single connection to an R&E provider and no interest in a second one. As a result, R&E network operators invest in bandwidth capacity and advanced features rather than redundancy. Their support staff are centralized and limited in number, with on-call rotations or outsourced services taking the place of 24x7 manned network operations centers. Field services are provided by contractors, and equipment is maintained under vendor support contracts instead of spare parts inventory and redundant installations.

Other regional networks have opted for redundant connections to the Internet2 backbone, in some cases by partnering with nearby connectors to share the service and split the cost. As existing customers have expressed a need for higher-availability services, these regional networks are responding by expanding their presence into areas that they have not previously served. SOX, for example, is expanding beyond Atlanta into Nashville, Dallas and Chicago to provide their customers with redundant commodity peering service.

New constituents such as health care and public safety may be expected to raise the bar still higher. In this case, redundancy in both the regional and national networks becomes more important, to guard against outages at the customer site. At the regional level, redundant connections to the customer site will be required. The difficult part is providing the degree of diversity to make those redundant connections useful. In cases where regional networks rely on partners to jointly serve health care or public safety, the separate partner infrastructure may provide sufficient diversity.

The U.S. UCAN proposal has addressed this in a similar manner by providing two separate IP networks (red and blue), managed by Internet2 and NLR. If two networks are indeed both contributing resources to the U.S. UCAN service, then a thorough understanding of the level of optical layer and fiber diversity that underlies these IP services is required before high reliability services can be provisioned and sold (this same caveat applies to regional high-reliability services).

Health care and public safety customers may require SLAs from the regional networks. The prevailing view of the R&E community is that these contractual commitments don't actually improve reliability, they are merely bets that provide some financial and perceptual incentives for providers to make the right architectural and operational choices. In order for the R&E community to distinguish itself in the area of reliability, it will be important to leverage a key aspect of traditional R&E service: transparency.

Transparency

The technical details of a commercial ISP's network are carefully kept secrets, since knowledge of topology, equipment, configuration, connectivity and network utilization could be used against them by a competitor. The same secrecy applies when a customer tries to work with the ISP to resolve a problem, asks for information about maintenance work, or requests a post-mortem report after an outage. On the other hand, R&E networks, with their smaller customer

base, strong hierarchy of customer relationships (backbone-regional-campus-enduser) and requirement to provide advanced services, have a history of providing much more open and responsive customer support services, including visibility into the network's operations and configuration that is unheard of in the commercial world.

Institutions with sufficient technical capability are able to use this transparency to better manage their use of the network and to understand both its capabilities and limitations. They can access real-time information about the state of the network (utilization, errors, outages), can test their network performance to known-good measurement points, and in many cases can troubleshoot network problems on their own, using tools similar to those available at the network operations center. The combination of applications that are sensitive to network degradation and the ability of users to see the details of the network performance helps keep the network free of problems like packet loss from congestion, errors, and misconfigurations.

The extensive transparency that has been part of the culture of R&E networks at every level will continue to be an important distinction. However, the value of this window into the network will be lower for those connectors with lower levels of in-house network expertise. Additionally, it is possible that some new communities (particularly public safety) may be concerned with the openness of the R&E networks, preferring that the details of the network not be known to anyone other than the provider.

Regulatory environment

R&E networks have long enjoyed freedom from most of the regulations imposed on commercial carriers and Internet service providers. To the extent that they operate private networks with restricted membership and purpose, R&E providers can avoid entanglement with tariffs, USF, DMCA, CALEA and other requirements that commercial entities must deal with. By the same token, when regional and national networks begin to offer services that make them appear similar to commercial providers, they begin to become subject to those regulations. Recipients of BTOP funding, in particular, are specifically required by the program to:

comply with all applicable Federal and State communications laws and regulations as applicable, including, for example, the Communications Act of 1934. . . Telecommunications Act of 1996. . . and the Communications Assistance for Law Enforcement Act (CALEA).

The effort required to meet these regulations will vary widely between networks, depending amongst other things on technology, topology and operational structure. Requirements for lawful intercept are currently in the news with speculation that the rules and requirements may change in the near future, unfortunately making this a moving target.

The FCC's E-Rate program is also in the midst of change, generally seen as favorable to the new community anchor institutions; since regional networks are now more likely to be able to provide services that are reimbursable under E-Rate, they will need to become Service Providers, and more fundamentally be conversant with the rules and regulations of the program so as to correctly advise new institutions when they have questions about service eligibility.

The national R&E networking organizations have been strongly in favor of network neutrality, commenting in ongoing policy discussions and attempting to influence vendors to maintain neutral networks. Network neutrality has been raised as one justification for the national R&E networks' development of large-scale commercial peering networks (TransitRail and CPS),

under the theory that if major ISPs began to favor certain content providers, the users of TR/CPS would still be able to access all content on an equal, neutral basis (at least, whatever content is reachable over the peering connections). This attitude can be expected to extend to the future U.S. UCAN network and to regional networks connecting the new CAI communities, and will be a selling point if other networks fail to remain neutral. It is also likely that any new regulations regarding neutrality will extend to U.S. UCAN and any networks receiving BTOP funding, and it is possible that such rules will impact the ability of the networks to provide QoS-differentiated services for different communities.

The relevance to regional networks and U.S. UCAN of regulations that involve data privacy (HIPPA, FERPA, etc.) and require content inspection and filtering (CIPA, etc.) will need to be revisited in light of their importance to the new constituencies.

Business model

Without exception, R&E networks in the US are run by not-for-profit or government-funded entities, and commercial Internet networks are for-profit. The differences between these three models are too many to list, and they appear at every level within the organizations: governance, financial rules, product design, service models, pricing, contracts, marketing, etc. And although the practice is disappearing at the national level, many regional networks still enforce acceptable use policies that determine the makeup of their membership and the uses those members can make of the network.

Existing business models of both the regional networks and U.S. UCAN will be stressed by the mix of new constituents and the commitments (and restrictions) imposed by NTIA-funded infrastructure. These organizations will be challenged to either stay true to their existing charter, modify it considerably (in conjunction with an expanded set of constituents with a vested interest), or consider forming separate entities (potentially for-profit) to deal with the new reality.

BTOP-funded networks that provide services to commercial last-mile networks or other for-profit customers may find that such activities are considered unrelated business income for tax purposes, and should consult with their financial advisors and tax experts.

Shared services

Commercial ISPs typically provide several add-on services at low or no cost: domain name service (caching resolvers and authoritative servers), address space assignment, and network utilization statistics. Those catering to the home and SOHO market often add e-mail boxes, web hosting and bundled double or triple play services (telephony and television). Often these services are made available through an individualized customer portal, to allow self-service for routine requests.

To date, the typical research network participant has had little need for these services, preferring to provide them internally. The new communities will likely expect to receive these and potentially other services from the regional network, requiring the creation of new back-office and portal capabilities. Different communities may expect different bundles of services, and usage of any particular offering may vary widely depending on the nature of the institutions being served. For example, the K-12 community requires security and content filtering services that would not appeal to most other institutions, but that they expect from their providers. Initiatives from the Internet2 community like federated identity management, high-performance distributed computing, logistical networking and cloud computing will be available to the new

community institutions for the first time, but their adoption will likely require the national and regional networks to take a role in education and support.

Summary and next steps

As the new CAIs grow to represent a larger fraction of the community, their different needs, desires and attitudes will increasingly influence the evolution of R&E networking at every level: governance, budgeting, support, future development and policymaking. The primary challenge facing the entire R&E community is simply stated: maintaining the distinctiveness of the networks and institutions that we have created, while expanding to serve new communities and meet the other requirements imposed by the opportunities we have sought out and won.

While the regional networks, national networks and U.S. UCAN make individual preparations, the ability of the community to move forward will depend on group coordination, harmonization and concerted effort on those issues that affect us as a whole.

To that end, we believe that the NTAC should consider the following questions:

How can we provide ample bandwidth for research and education traffic while allowing for the demands of new institutions and communities? Will this affect the design and operation of individual networks? What level of coordination and cooperation will be needed across networks, and how will this affect existing connections, the network links and the contracts or other arrangements that govern them?

BTOP awardees are already planning to increase the density and reach of their networks in order to accommodate new CAIs; will those regional networks not participating in BTOP also feel pressure to connect these new institutions? How can the relationship between the regionals and U.S. UCAN be structured to allow both to benefit from the participation of new communities?

Regional networks have different models for providing middle mile services; will meeting the BTOP requirements for open access to require changes to their pricing and service models? How will this affect the existing community's use of those services?

To what degree can the R&E networking community find commonality in their open access policies, and harmonize those policies? Can we agree on a single, consistent interpretation of the open access requirement? What other regulations, not previously applicable to R&E networks, will we be subject to?

What can we do to make advanced services, long a hallmark of the R&E network environment, attractive to new communities? How can we alleviate concern on the part of those communities who do not require them, and view the complexity that they add to the network as a drawback?

Newly connected institutions and communities will require new and different levels of support at every stage, from first engagement through the lifetime of their network connections. What can we learn from those networks that already have experience supporting some portion of the spectrum of institutions? How can we work, individually and together, to handle the changes in scale and depth of support that will be required?

How can we preserve our flexibility and openness to new ideas, technology and operational

models; use those characteristics to help the networks adapt to the new communities; and maintain flexibility as our networks and operations become more complex and subject to new demands for reliability and resiliency?

How can we provide new institutions with the tools and technical expertise to take advantage of the transparent operations model that many R&E networks have embraced? How will we maintain a balance between transparency and security of the infrastructure sufficient to satisfy new communities with different experiences and expectations?