Architectural Overview for a New SURA Regional Infrastructure

A white paper developed by the SURA RII Architecture Working Group (AWG)

SURA gratefully acknowledges the contribution of both intellect and time from members of the SURA RII Architecture Working Group towards the development of this paper.

SURA RII AWG Members:

- Steve Corbato, Internet2
- Gary Crane, SURA
- James Deaton, OneNet
- Tim Deeves, Tulane University
- Jed Diem, Tulane University
- Thomas Durkin, Geo-Matrix
- Larry Flournoy, Texas A&M University
- Doyle Friskney, University of Kentucky
- Jeff Fritz, West Virginia University
- Zane Gray, University of Oklahoma
- Chris Griffin, University of Florida
- Daniel Grim, University of Delaware
- Philippe Hanset, University of Tennessee
- Rene Hatem, CANARIE
- Ron Hutchins, Georgia Institute of Technology
- Bill Johnson, OneNet
- Mark Johnson, North Carolina R&E Network
- Michael Krugman, Boston University
-
- Dewitt Lattimer, University of Tennessee
  Paul Love, Internet2
- John Nichols, Virginia Tech
- Ken Orgill, West Virginia University
- David Pokorney, University of Florida
- Ana Preston, University of Tennessee
- Mike Rackely, Mississippi State University
- Sandra Redman, NASA
- Dave Reese, 4CNet
- Stewart Seruya, University of Miami
- Jerry Sobieski, Mid-Atlantic Crossroads
- Troy Travis, University of South Carolina
- Ryan Vaughn, University of Florida
- John Watters, University of Alabama
- Roy Whitney, Jefferson Lab
- Bill Wing, Oak Ridge National Lab
- Mary Fran Yafchak, SURA
- Robert Zimmerman, University of Arkansas

An Architectural Overview for a New SURA Regional Infrastructure

Executive Summary (Reminder to develop one once the rest of the draft is complete.)

Introduction

The set of universities that constitute the Southeastern Universities Research Association (SURA) have decided to explore the feasibility and opportunities associated with constructing a Regional Information Infrastructure (RII) within the 16 state footprint that SURA represents. A broad interpretation of the architectural overview for this infrastructure would include a study of the issues associated with constructing a comprehensive networking infrastructure as well as a common middleware and applications roadmap. This document currently addresses only the networking issues.

It can be argued that the networking issues are strongly, if not entirely, dependent upon the higher layer requirements imposed by middleware and applications. However, it can also be argued that the various layers of the network stack – which extend all the way up into the core of applications – should in fact be largely independent of one another, solely reflecting a set of standard interfaces. In reality, of course, things are a blend of both. In addition, application characteristics and, therefore, the requirements for the infrastructure are expected to change dramatically during the lifetime of this project, partly due to the new capabilities that this network will provide and partly due to totally independent influences.

This document will approach the network architectural issues using the practical experiences acquired by the members of the SURA RII Architecture Working Group (AWG), who have built and supported both regional and national networks in support of higher education and evolved these networks amid changing environments. The AWG has developed this architectural overview to describe one potential strategy for creating the desired capabilities. It constitutes a broad vision where traditional limitations, e.g. bandwidth, are no longer an issue constraining the applications developer, and lays out a plan for constructing and managing the infrastructure as an inter-institutional shared resource.

Influential Trends & Context

Any coordinated effort by the SURA member universities to create a long term telecommunications infrastructure should recognize not only what the members current and currently anticipated needs are but should also recognize the difficulty in predicting what capabilities and requirements such an infrastructure would have to meet over time. This infrastructure will also need to track with similar changes over time occurring in relevant national and international networks.

The AWG believes that any new national/international academic network infrastructure will fundamentally be a very high performance backbone interconnecting regional networks that in turn will interconnect local institutions and campuses. This is certainly the current model of the Internet2 infrastructure and seems to be evolving well. However, the advent of optical wave division multiplexing, advances in long-haul and ultra-long-haul optical technology, gigabit Ethernet and 10gig Ethernet, and advanced network services such as lambda switching and routing, all conspire to make additional details of the network architecture of the future somewhat difficult to predict. These technologies - and others such as high-speed wireless, proliferation of mobile and nomadic computing devices and the explosion in trans-oceanic fiber capacity - will have a broad and significant impact on networking in the future. For instance, cheap and flexible bandwidth on demand may allow full or partial mesh interconnections among network components. While this may not affect the selection of the end-to-end bearer service, e.g. IP, it will have a definite impact on the engineering and architecture of the core network components, meaning fewer routers with more interior paths between them, primarily at the edge of administrative boundaries, and more sophisticated core optical transport capabilities.

Traditionally, the cost of network links has been bandwidth based and a function of the distance between the termination points for each link. Since most of the termination points were only served by an Incumbent Local Exchange Carrier (ILEC) or had very little competition, the pricing of such services were/are "market based", i.e. as much as the market could bear. This put universities in a particularly untenable position in that they needed to push the network and applications technologies but were severely constrained by the high cost of doing so. Further, the costs of such services were – and still are – distance sensitive, placing an inordinate burden on institutions that were not located in or near major metropolitan areas that had developed into telecommunications hubs.

The development of optical technology and the presence of more competition have resulted in a current downward trend in the cost of wide area bandwidth. This, combined with the current market conditions of excess capacity in the long haul markets, could remove this historical limitation – bandwidth would be taken off the table as a constraint. However, in order to build a new cost-based model, the higher education community will need to gain a greater level of control over the network structure and may have to build its own network. Doing so could potentially provide significant financial savings over time, but perhaps more importantly, it will afford the higher education community the ability to evolve the network infrastructure to meet their own needs. The higher education community needs this flexibility to insert new technology on their own schedule, to upgrade or migrate to new services as they define them and are capable of deploying them.

One caveat to this scenario is that we don't know what will happen in the commercial market place. Would a precipitous decline in the cost of bandwidth obviate the need for the SURA schools to build their own network? The cost of adding a 10 gigabit/second network link appears to be approaching the $100K mark. However, the hidden costs include the wide-area engineering that is required to support such a "cheap" link: the distances that such an interface card can drive, intermediate equipment that will likely be required (WDM), the cost and availability of fiber capable of supporting such high speed connections, etc. These are very real and potentially expensive hidden costs that need to be factored into the "build or buy" decision.

The RII AWG anticipates a major shift to all optical networks, with the core initially utilizing the high capacity multi-wavelength optical technologies but with these capabilities becoming more broadly deployed, and more pervasive and sophisticated over time. The implications of this expected trend is that access to actual fiber to support this technology becomes crucial. The fiber plant and the engineering, particularly in the "long haul" segments, will be the most difficult and likely the most expensive aspects of the project, and few networking personnel within the university community have ever undertaken such an engineering task.

The strategy used for developing the SURA RII network architecture will also be influenced by the general approach that SURA is taking to identify potential commercial partners for the project. The general process of the SURA RII is to identify potential commercial partners before attempting to define a network design or business plan. This opportunistic approach to identifying potential partners/providers of transport services provides for wider choice and encourages creative thinking about network partners and design. It also means that network design and implementation are likely to be iterative processes and based on the partnerships as they develop. This does not diminish the need for an overall strategy - in fact it heightens the need.

Design Principles & Strategy

Given the conditions and challenges listed above, the AWG is proposing specific design principles and strategy to be used in guiding the creation and evolution of a new networking infrastructure for the SURA region.

The AWG recommends that the following assertions be adopted as the primary design principles for the SURA RII:

Possible additional design principle from 7/16 RII WG call; need to discuss and address. Regarding the issues of the inherent reliability and multiple diverse connections from a given campus. Many of our larger campuses maintain multiple local loops to multiple ISPs specifically for the added insurance of diverse local loops. A single, albeit, high bandwidth local loop connecting a given institution to the RII would forfeit some of this insurance.

Design Principles:

1) **Bearer Service -**

a) **The Primary service will be an Internet Protocol (IP) bearer service.** IP is currently the only practical means of delivering application traffic end-to-end. It is likely to remain so. However, the advent of wireless and mobile computing associated with the proliferation of inexpensive network enabled personal devices is driving a strong push for IP version 6. Any “advanced” network going into service in or after CY 2002 should strongly consider IP v6 as the native layer3 protocol. Ipv4 should be considered transitional only.

b) **Optical Service - There will have to be an active optical transport layer.** The Optical transport layer will consist of all components required to deliver a data stream between two points on the network including amplification and regeneration of signals if required. This layer will support inputs based on the traditional telco SONET hierarchy as well as others such as 1Gbps and 10Gbps Ethernet. The optical layer will include the ability to perform dynamic switching of transport streams.

c) **Some telecommunications research efforts will require separate fiber infrastructure.** Some types of research will be able to use separate optical paths (e.g. lambdas) along the “production” optical network. But very low layer efforts will necessarily require separate fiber in order to insulate upper layers from destructive testing. Such “research” oriented infrastructure need not be as extensive as the “production” infrastructure. However, an extensive fiber plant that shadowed the production infrastructure would ease migration to new technologies by providing the capability to roll new production services out on the [newer] research fiber, then replace the old fiber with newer yet research fiber.

2) Physical Infrastructure -

a) **A corollary to an active optical transport layer is the need to provision fiber on a much broader scale to all - or nearly all - of the participating institutions regardless of their geographic location.** The optical layer of the network will not fill the expected requirements if it is only accessible from the core network nodes. The optical layer must extend all the way to the campus and into the campus. This implies that a large number of traditionally “fiber challenged” institutions should be part of the plan to provide robust fiber routes to their institutions, and a plan for managing that resource to benefit the entire community. This is a subtle but critical design criteria and should be explored carefully and thoroughly.

b) **The low layer infrastructure should extend down to conduit wherever possible.** This will allow the most flexibility in the out years to replace today’s initially deployed fiber with new forms of fiber which will undoubtedly be required to support the emerging optical technologies. This will also allow the insertion of additional or different (e.g. marine grade) fiber to meet un-anticipated requirements. If conduit is not available, or is overly expensive, then the effort should pursue fiber IRUs for at least 15 years.

c) **Given the likely direction of the technologies that will be deployed within the RII, the possibilities for a logical infrastructure that permits a much greater density of interconnections exists.** We may not be able to predict the timing and/or shape of particular logical configurations that may be possible and desirable in the future. When designing the physical infrastructure, this potential for increased density of connections should be considered.

3) **Business Model -**
   a) **The capital costs of infrastructure should be amortized over 5 years.** In order to stay on the leading edge, the fiber being deployed should be reviewed and may need to be refreshed every 5 years or so. If in fact this fiber can be cost justified over 5 years, then any additional productive use of this fiber will be an added bonus. Optics and electronics equipment are also likely to need replacement over five years. Conduit investments might be reasonably amortized over a longer period.
   b) **The network should leverage the SURA University community to create new confluence points for both the RII network as well as commercial service offering points at those universities.** The universities could become a new set of telecom Points of Presence. This would benefit the universities by providing easy access to competing commercial services. The carriers get a stable anchor tenant that justifies the investment. The cost to the universities is minimal, essentially space and power. However, there is an emerging collocation market that could provide a significant financial revenue flow to offset services the university may require.
   c) **Restrictive use policies should be avoided.** The advantages of an open access policy for this new network should be considered. This would essentially allow any entity to buy transport services on this new network to supply services to their customers. This could provide incentive to multiple public and private parties to invest in last mile solutions and dramatically increase the reach of this initiative. This will only be possible if the infrastructure that is used to support this new network is available to us without restrictions.
   d) **Advanced network equity for all SURA's members.** Many of SURA's member institutions are located in non-metropolitan areas where there is little competition for broadband access services. Many of these areas lack access to fiber optic infrastructure and represent areas of smaller aggregate demand for network services, thereby making access to high bandwidth network services costly and difficult, if not impossible, to obtain. Every attempt should be made to leverage this project to lower these hurdles.

4) **Management -**
   a) **The network must have a clear point of demarcation in order to bound management and operational responsibilities.** Regardless of the organizational structure that is developed to manage this project it is important that network boundaries are clear to all participating organizations involved. While the final network configuration may be comprised of several hierarchical layers, each must have clearly defined boundaries with clearly identified responsible parties.

b) **The network should peer with local and/or regional Internet service providers.** This initiative should leverage the magnitude of the Universities in their community to entice the prevailing Internet Service Providers in a region to peer with this network at no cost. Each byte sent to a local [no cost] provider, is one less byte sent out the default route to a high cost commodity transport service.

Strategy:

In conjunction with the above design principles, the AWG recommends that the following strategy be used to develop the final network design and evolve the management and operational plans in support of that design.

The regional infrastructure and the partnerships that will enable and sustain it should be addressed at multiple levels. With respect to planning, implementation and even funding, the levels to be considered are regional, sub-regional, and local.

Regional
The SURA RII is intended to benefit the SURA region as a whole and will move forward most effectively if "top down" coordination knits together contributing activities across all levels. This coordination could be provided by or through SURA and should ensure that there is equal opportunity for participation and an equal voice for all participants. Coordination will be necessary to move the RII forward as a cohesive set of compatible interconnections on the network level. It will also be needed when designing effective overall management in the areas of: contracts, service level agreements, network operations, user support, development/deployment of higher level services, and continued strategic planning.

Sub-Regional / Local
Sub-regional distinctions within the RII allow for exploitation of opportunities or resources that are bound to particular locales, states, or multi-state groups. The actual necessity for and boundaries of desirable sub-regional and local RII components will become more clear as SURA moves further through the partner identification and investigation process. Specific activities and directions to be taken at the sub-regional level can be initiated (motivated and informed) and enabled at the regional level. Such activity could be supported or partially supported using regional resources, e.g., the availability of SURA-sponsored sub-regional leads to work on targeted efforts.

For the short term, the RII could progress at the regional level using a continuance or close variation of the current SURA staff/working group(s) arrangement for management and design. For the longer term, it is likely that additional staff or formal dedication of current staff will be necessary. As the regional group continues to investigate potential partners who may be able to provide various portions of the infrastructure, they should also identify sub-regional and local target areas and partners with specific interests in those. As areas and activities emerge that are sub-regional in nature, these can be bounded as "sub-regional components" with regional support extended to them for purposes of coordination. A similar approach would work for local areas/activities but it is also likely that more localized efforts will be most dependent on local motivation, information gathering, and negotiation.

A regional infrastructure that truly illuminates a new model for advanced networking is not likely to develop quickly or all at once. The network will need to be built in stages with perspective and foresight so that the stages blend into a smooth evolution. This evolutionary approach also creates an environment where the benefits of the developing infrastructure for various constituencies can be further explored and, if appropriate, can influence the direction and/or timing of the network build. The SURA region and the SURA membership are both large and diverse and the human and programmatic connections between the SURA community and other communities within the region are numerous.

The main focus of the RII is to realize a new model of networking for the South. Given that, the first and most important criteria for deciding when, where, and how to build is whether or not a particular direction/component embodies new technology and/or pricing structures in support of that goal. A second valuable evaluation criterion (valuable for community building; valuable for procuring funding) is to whom and how a particular direction/component provides benefit. The AWG encourages further and ongoing examination of "benefit to community" aspects of the RII in conjunction with the selection and development of the physical infrastructure.

The examination of community benefit and the selection of partners and technologies to implement the RII should proceed hand-in-hand. As promising partners or combinations of partners capable of connecting across the region are identified, the coverage that each partner or combination can provide should be examined to determine what gaps in coverage, if any, would remain. The goal of the SURA RII is to dramatically reduce the cost per megabit for network transport for all SURA members, with complete coverage of the region in as short a term as possible and sustainable for the long term. The need to initiate specific sub-regional and/or local efforts that engage targeted sub-regional or local providers will be driven by the gaps that would need to be filled. The final selection of network partners and design should remain cognizant of these things. In addition, the AWG recommends that SURA establish a partners' forum or other ongoing means of communication between SURA, RII participants, and RII partners in order to support coordinated and collaborative development of long term strategies.

Strawman Network Design

Additional consideration for Network Design, from 7/16 RII WG call; need to discuss and address - Regarding the issues of the inherent reliability and multiple diverse connections from a given campus. Many of our larger campuses maintain multiple local loops to multiple ISPs specifically for the added insurance of diverse local loops. A single albeit, high bandwidth local loop connecting a given institution to the RII would forfeit some of this insurance.

The effort to architect the physical design of an optical network infrastructure for the SURA region, i.e. identify the "correct" locations for core nodes, choose fiber routes, vendors, etc., will have to be broken into two aspects:

First, a working map of the core backbone should be established. This map would be intended only as a tool to initiate a discussion regarding a new physical network infrastructure and to allow a first cut at the degree of cost and complexity for the network backbone. The core map should be reviewed in conjunction with the discussions that are currently taking place with regional fiber providers. It may become apparent that the physical fiber infrastructure that is currently in place in many of the major metropolitan areas within the SURA footprint may dictate the placement of core nodes.

Second, but perhaps more important, an in-depth study should be completed concerning fiber routes and "nesting", i.e. where fiber routes could be created to construct localized robust rings and/or meshes that will reach as far into the set of regional campuses as possible. The objective is to minimize fiber miles yet maximize the degree of fiber connected institutions. The ability to reach possibly many remote or relatively small campuses that are not on major fiber routes will be key to creating a successful optical network for the SURA region.

Figure 1, below depicts the first of two example networks. The "Terrestrial RII" network map was constructed by first identifying three key locations which have existing extensive telecommunications infrastructure. These locations were Washington D.C., Atlanta, and Houston. A core link was extended between these locations using route maps from two fiber providers. Additional core routes were then added to pass and pick up other "important" locations – i.e. major universities and/or research labs, additional access points (e.g. Miami), etc.

It should be noted that the priority on choosing routes was based on building a broad network that reached the most SURA institutions, **not** based upon cost or ease of access. This strategy results in many locations that traditionally have had great difficulty in acquiring cost effective services now having access to a major telecommunications route. In some such cases, routes may run "near" several institutions, but still miss them by many miles. By reaching fiber providers before the routes are locked in, we can influence such routes to run directly past or through these institutions. In these situations, the institutions can fast track permitting with local jurisdictions, provide regeneration or access pops for the long haul circuits, and create markets for service providers where there were none previously.

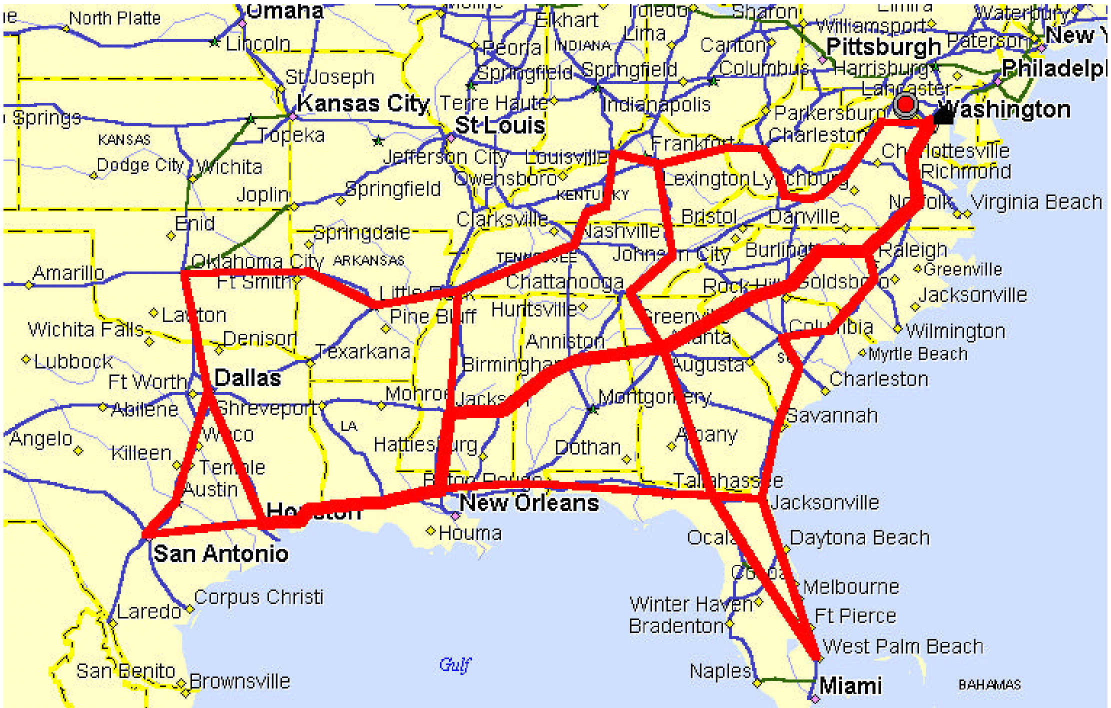

Figure 1: Terrestrial RII

The Terrestrial RII design consists of approximately 10,000 route miles of fiber. It contains four major IP peering points: Washington, Atlanta, Houston, and Miami. It also includes approximately 15 optical switching points, with as many as 50 optical add/drop muxs (OADMs) serving the constituent institutions. Implicit in such a configuration is the Intermediate Light Amplification sites (ILAs) that will be required every 60 km (~40 miles).

As another example of a potential network architecture, the “Surf & Turf RII” map, Figure 2 below, describes a network that combines both land based and marine fiber. Such a system has both advantages and disadvantages: offshore long haul routes are relatively inexpensive to put in place and landing points are flexible, but the inland points of interest may be more challenging in terms of reach and permitting and cost.

Figure 2: Surf & Turf RII

Where fiber is just too expensive, cannot be deployed immediately, or for other reasons may be delayed, traditional commercial services may be employed to connect an institution into the net. This will sacrifice to some degree that institution's ability to participate fully in an optical network, but this should not be too onerous to smaller institutions, and commercial services may in fact begin to offer lambda services over the next few years.

Upon examining the maps of SURA institutions and other related facilities, one observes that many of these institutions are concentrated in areas served by natural affinities: common state network initiatives, Internet2 gigapops, metropolitan areas, and the like. So it might be important to develop and fund a set of regional initiatives that focus on putting fiber into the local campuses and homing that fiber infrastructure back to one or more regional nodes. These regional nodes would be the confluence points for the multi-state network aggregation of optical and/or IP traffic.

The optimal number of core nodes is an interesting issue and merits further discussion. How will we balance the choice of location and number of core nodes against the cost of campus access services and the need for costly amplifiers and regeneration equipment?

The active optical transport capability provides the network with new options for delivering the IP bearer services. Indeed, given the optical connectivity, the definition of a "core node" may change to reflect the long haul optical aggregation rather than the traditional IP aggregation. The lambdas would be provisioned in [possibly] a full mesh across the network, only touching down at the border routers.

In fact, the lambda provisioning would allow interesting new models of service provisioning in the network. For instance, coastal marine cable might be utilized to carry “long haul” traffic between land-based network nodes. In parallel, ocean based instrumentation platforms could access gigabit Ethernet lambdas to carry sensor data from the sensors to the supporting institution.

Again, these illustrations are intended only as tools to encourage discussion among ourselves and with potential commercial partners. The fact that SURA is not approaching this project using the traditional RFP process means that a final physical network infrastructure will be highly dependent on the network assets that can be delivered by our choice of commercial partners. It is fully expected that the process of finalizing a region-wide physical network architecture will require many iterations. It is also anticipated that it is highly unlikely that we will be able to identify a set of partners that will be able to provide complete coverage of the SURA footprint. Areas where gaps are identified in regional coverage will be targets for sub-regional or local investigations and may require the development of additional, more local, partnerships or direct investments in the construction of fiber segments by affected member institutions.

Working with our consultant, Geo/Matrix, the AWG will develop a recommendation as to the “sweet spot” in terms of network parameters (conduit vs. fiber vs. lambdas, where the routes should go, what services are provided and when) at any (or every) point in the SURA footprint. The recommendation will be the result of in-depth discussions with numerous potential partners, cost considerations, organizational support and structure, priorities, and timelines - and, ultimately, "best guesses" as to what the future holds and how best to take advantage of emerging optical networking technologies.

Management and Operations

Recommendations regarding the governance of a new SURA regional infrastructure are not within the scope of this paper. However, the management and governance of a new SURA regional infrastructure will define many of the assumptions about the network – its purpose, its scope, its financing, etc – which will have a very real impact on its architecture. The SURA RII AWG recommends that the issues regarding the governance of this new SURA regional network be appropriately analyzed in a separate white paper devoted to governance issues.

The network must adopt some level of operational service goals. Whether this is “three 9s” or “five 9s” can be argued. But whatever the target level, there must be a positive and active program to monitor and characterize the network performance and to discover network outages.

There must also be a clear and well-defined delineation of operational responsibility for the network, i.e.; the core infrastructure must have a well-defined boundary and administrative authority. There will be multiple levels of network aggregation – some historical in nature, some born out of financial or technical pragmatism. Wherever these lines of responsibility are drawn, the ability to support a dependable and reliable network from the top down should be the prime consideration.

This does not rule out a more distributed model of network support, e.g. a confederacy of regional networks that have a common transit policy serving their participants. In theory, a confederation of regional networks offering transit to each other and agreeing to a common routing policy could be made to work. In practice, however, such models tend to balkanize the network and do not address the distribution of costs necessary to serve the entire confederation equally. For example, some regions are “better connected” than others, imposing excessive burden on the remote university (such as land grant schools).

Further, since the network will be attempting to create an active and interoperable optical networking capability, there must be a common coordinated engineering process to insure a broad level of compatibility and extensibility into the future. This engineering process must be driven by a common, coordinated, strategic planning process to align the network design with future application and usage goals.

Funding Issues

Creating a network of the type described in this paper will result in significant costs to SURA and SURA member institutions. To date the SURA Regional Infrastructure Initiative has been funded from the SURA operating budget, SURA member contributions of staff time and travel and the SURA IT Fund, which was generated from the sale of SURANet. These funds have been adequate to seed this activity to date and should be adequate to complete the planning and design phases of this initiative. However, the actual construction of a network of the scope outlined in this paper will require a combination of creative corporate partnerships, participating institutional funding support, continued support from the SURA IT Fund and possibly funding from state and/or federal sources.

It is anticipated that a network of the type described in this paper would be capable of supporting all of the current network traffic at each of the participating institutions and have the capacity to support significant enhancements to the connectivity at these institutions. It is anticipated that by redirecting the revenue streams that pay for this existing connectivity the participating institutions should be able to realize a significant increase in the bandwidth available to their institutions. In other words, it is believed that by forging creative corporate partnerships and using new technologies that allow for significant cost reductions in wide area transport, the network conceived in this white paper can be funded for the same cost as the aggregate current expenditures by SURA member institutions for lower levels of service. If one considers the potential to also roll the current aggregate expenditures for voice services onto this new network infrastructure, the potential for reducing the overall expenditures for communications by SURA member institutions while supporting this new infrastructure are greatly enhanced.

A simple examination of Abilene connectivity costs within the SURA region, arguably the lowest cost network services available to higher education on a dollar per megabit basis, reveals that there are currently 22 separate connections to Abilene within the SURA region (18 at OC3 and 4 at OC12). This represents a cost of $3.06 million per year just in Abilene connection fees. Local loop fees for these connections, if estimated at an average of $80,000/year/connection, amount to $1.76 million. Consolidating these connections into 4 OC48 connections distributed across the SURA region could result in estimated savings in excess of $3 million per year. Potential savings from similarly aggregated connectivity to commodity ISP services would likely far exceed this level.

The AWG recommends that a detailed survey of network connectivity costs be conducted by SURA to better quantify the potential savings from region-wide aggregated network services.

Networking Research

The basic goal of the SURA RII is to create a network architecture that can address the day-to-day needs of the higher education community well into the future. The implication is that while the network would support research, it would not primarily be a “research network” or test-bed for network developments. The network would necessarily be an “early adopter” infrastructure, but the fundamental intent is to provide a flexible yet stable network.

There is however a secondary need to support the research efforts of the telecommunications sciences. The technologies that are developed and proven in concept in experimental settings need a place where they can be deployed and tested on a wider scale. These technologies include new transmission technologies as well as experimental signaling protocols and applications. Such experiments may be destructive to various layers of the network. For instance, advanced WDM experiments can cause crosstalk across wavelengths thereby making such experiments incompatible with production wavelengths. Such experiments will need access to their own fiber. On the other hand, some experiments with new higher-level protocols may cause a network link to fail. While this may be destructive at the IP layer if it does not interact with the optical layer, these experiments could be performed using "production" lambdas to create an "experimental [higher layer] network".

Without access to a separate fiber infrastructure, potentially destructive research uses of the network must be carefully reviewed on a case-by-case basis to ensure that the more operational uses of the network are not disrupted. This review process will determine the degree of isolation such experiments should receive. It seems obvious (though may not be so) that separate fiber infrastructure for research would be an expensive prospect. However, supporting basic network technology research is an important aspect of maintaining a leading edge fiber based regional infrastructure. In order to roll out new production services and technologies, an ability to test such technologies in order to gauge their maturity is an absolute requirement. Doing so over long distances with real world issues is the most credible means for doing this.

The AWG recommends that the SURA Regional Infrastructure Initiative provision facilities for "destructive" testing at the lowest level in the network hierarchy that is economically practical. This includes the possibility of creating a dedicated fiber infrastructure that could be used to prove the viability of new transport technologies that could then be migrated to a more production environment, as they become stable. An early deployment of the SURA RII project could even be built as a research network with the goal of testing new technologies that are targeted for the production version of the RII. For example, the SURA network research community could be engaged to work with potential corporate partners to trial and develop 10 Gigabit Ethernet technologies for use in a wide area network.

One other important possibility would be for SURA to fund a special initial research network to address the actual issues associated with the RII purpose. I.e., if RII is to build a network, it better have first hand experience and knowledge to rely on as it does so. Some of this can be bought via outsourcing various aspects, but some of the technology is so new and evolving so quickly that there are no credible outside firms that can provide such a service. In this latter case, SURA should insure that it has from among its members, adequate institutional know-how to execute on this initiative. Possible suggestions along these lines would be to continue supporting metro area optical initiatives such as MAX or SOX, but to explore adding a long haul component to this approach such as a marine based backbone, or a similar fiber project between say Miami and Washington. Such an effort need not physically touch many institutions along the path since its purpose is research and engineering. But it could provide a great deal of crucial experience in deploying new technologies.

Production Services

Commodity Internet

Access to the commodity Internet has become a critical service required by all universities. The number of faculty, staff, and students who routinely use the commodity Internet is approaching 100% and their available access speeds are increasing. Typical access for faculty and staff is 100Mbps Ethernet and students often have 10Mbps Ethernet to their dormitory rooms. Each computer is becoming faster and the applications used require more bandwidth. The end result is a seemingly unending requirement for more Internet bandwidth and very rapidly escalating costs. Challenges to SURA institutions include the high cost of local access to sufficient bandwidth to meet campus needs and, in many rural areas, the lack of a competitive environment for ISP services resulting in large disparities in the cost of commodity service across the SURA region.

The SURA RII plans to attack the rising cost of commodity services on two fronts. It will provide IP transport to common aggregation points where ISPs will deliver services.

Common aggregation points will allow institutions in remote areas to leverage access to reach multiple service providers. Resulting aggregation and competition will keep prices reasonable. The RII will also support and enable peering, caching, and replication services that will reduce the growth of bandwidth needs. Existence of the RII will mean that traffic between SURA institutions will not need to travel over expensive commodity links.

Maximum benefit from aggregating commodity service comes from buying in bulk. In order to take advantage of bulk buying and fairly distribute the cost of bandwidth it will be necessary to develop an approach for allocating costs associated with the bandwidth.

A side effect of enhanced competition and access to multiple carriers will be that institutions that previously could only afford one ISP because of high transport costs will be able to use multiple ISPs and experience enhanced reliability of their service.

The AWG recommends that as corporate partnerships are developed for the basic transport for the RII, SURA should actively pursue aggregated purchase and/or peering relationships with potential partners. It is particularly important to establish peering with ISPs offering high-speed residential services such as DSL and cable modem access.

Support for advanced applications and research

The network must support the research activities of faculty that demand advanced network capabilities. Some of the activities that are especially challenging are transfers of very large data sets, remote control of scientific instruments, and computing grid infrastructure. These advanced capabilities will manifest themselves as particular quality of service (QoS) requirements. These will range from latency limits from instrument control to bandwidth guarantees for very large file transfers. Each requirement translates into a specific service level agreement (SLA). Aggressive implementation of the type of features described above differentiates this network from typical commodity services. Some of the research requirements may in fact be best implemented by provisioning of parallel paths isolating traffic from the rest of the network.

Some network capabilities demanded by research requirements may become more ubiquitous service offerings. An example of this migration is the high quality video conferencing over IP. Originally a requirement for supporting some research some production videoconferencing is using the same technology now. We expect more advanced service requirements to be driven by research demands. Some services like voice over IP will be driven by the need to reduce the cost of traditional means of delivery.

Many advanced applications require new or enhanced middleware services in order to be fully realized. Middleware needs to be coordinated among SURA members and some of these services may have to be offered centrally in order to be effective. As the RII is developed SURA needs to participate in the Internet2 middleware working group and the Grid Forum standards process.

Of obvious and related importance is the relationship between the SURA RII and other relevant national and international networks and connection points (e.g., Abliene, ESNet, STARTAP). SURA members currently enjoy and depend on access to many such R&E networks to support the external research and collaborative work of their faculty and students. The SURA RII is intended to provide the level of service required to support advanced applications and research for member institutions; this service must be in the form of appropriate egress and peering arrangements in addition to advanced network capabilities.

Coordination/Cooperation with Other Initiatives

The SURA RII should attempt to leverage the work of other similar initiatives that are either in progress or are being planned in other parts of the US and Canada. Several projects are currently underway and should be tracked closely and wherever possible should be consulted and coordinated with as we progress toward implementing a new regional optical infrastructure. These initiatives include:

a) CANARIE's CANet3 and CANet4 Projects (http://www.canarie.ca/advnet/canet3.html)
b) CENIC's ONI Project (http://www.cenic.org/ONI.html)
c) The UCAID Network of the Future - next generation Abilene (http://www.ucaid.edu)
d) Other optical networking initiatives within the SURA region (e.g., MAX, NCREN3)

One method for SURA to maintain close involvement with the various other regional and national optical networking initiatives is to continue their involvement in TheQuilt (http://www.thequilt.net).

Recommendations

The AWG recommends that the architecture and engineering effort proceed to at least the point of developing a cost model for deploying and supporting this network. This process needs to include the identification and prioritization of the key points of interest on the network. Finding the means to interconnect these points will be an iterative process and some Points of Interest will clearly be easier to connect than others are.

Rather than focus primarily on points that can readily be connected, the AWG recommends that the list of difficult-to-connect points be examined carefully for aggregate need. This need can then be leveraged with carriers whose routes are yet undecided, to encourage and enable these carriers to build in areas that will be of benefit to the SURA RII.

The AWG recommends that the issues regarding the governance of this new SURA regional network be appropriately analyzed in a separate white paper devoted to governance issues. This second white paper should be completed in conjunction with the finalization of the architectural overview and should concentrate on analyzing the options for centralized management and NOC

The AWG recommends that plans for initial deployments be developed with the intent of proving the basic technical, managerial, governance and partnership concepts that the RII is based on. Further it is recommended that the RII Working Group thoroughly investigate any already identified "first build" scenarios as well as any others that emerge to ensure a balance of objectives and benefits across early deployment.

The AWG recommends that a formal structure of some type be established very early on to track developments and changes over time that are occurring external to the RII but in relevant national and international networks. This awareness will need to be brought to bear on both short and long term planning and operations.

The AWG recommends that a formal structure of some type be established very early on to track developments and changes over time that are occurring in the area of network technology. This awareness will need to be brought to bear on both short and long term planning and operations.

The AWG recommends that RII planning, implementation and management follow the design principles and strategy that are presented in this white paper. However, it should be remembered that these principles and strategy have been developed within a current context. As the RII proceeds, the guiding design principles and strategy should be reviewed and updated as necessary.

To sustain the RII beyond this initial period of partner and project definition, the AWG recommends that the major funding source shift from the SURA operating budget to a combination of the following: creative corporate partnerships, participating institutional funding support, continued support from the SURA IT Fund, possible funding from state and/or federal sources.

The AWG recommends that a detailed survey of network connectivity costs if SURA member institutions be conducted by SURA to better quantify the potential savings from region-wide aggregated network services.

The AWG recommends that the RII provision facilities for “destructive” testing at the lowest level in the network hierarchy that is economically practical. This includes the possibility of creating a dedicated fiber infrastructure that could be used to prove the viability of new transport technologies that could then be migrated to a more production environment, as they become stable.

The AWG recommends that as corporate partnerships are developed for the basic transport for the RII, SURA should actively pursue aggregated purchase and/or peering relationships for commodity ISP services with potential partners.