GNA – Global Network Architecture
Web: http://gna-re.net/

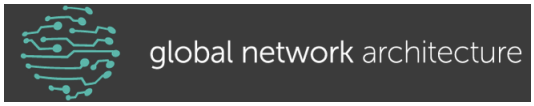

Document name:	**Describing a Reference Architecture for a Global R&E Network**
Author(s):	Dale Finkelson, Dave Reese
Contributor(s):	GNA Technical Group
Date:	January 2017
Version:	v1.0

Describing a Reference Architecture for a Global R&E Network

Table of Contents

Executive Summary

This document describes the components of the Global Network Architecture (GNA). It outlines technical and operational considerations at a high level. More in-depth information on specific issues can be found in separate documents on the GNA website. This is designed as an introduction to the GNA Technical Working Group set of documents.

Introduction

As long as Research and Education (R&E) networks exist, they will work to connect to each other on a worldwide scale. Global Exchange Points (GXPs) focused on serving the needs of R&E networks in facilitating these connections have been around since the early 1990s. Internationally, the Global Lambda Integrated Facility (GLIF) consortium was established as a collaborative partnership of networks and GXPs to focus on the needs of R&E that were not being met by the commercial exchanges.

Data-intensive science applications have always been major drivers of technology because they encounter technical barriers long before these limitations are experienced within other application domains. Currently, multiple science domains are motivating the creation of new advanced communication services to meet the requirements of the most data-intensive applications in the world, including those in high-energy physics, astronomy, computational astrophysics, genomics, computational chemistry, and others.

Applications are increasingly requiring capabilities that support the high-volume, long-duration streams with guaranteed assurances of high-quality communication services based on a dynamically changing infrastructure. Increased capacity integrated with capabilities for network programmability is a requirement not only for aggregate traffic but also for individual data streams, including direct control of those individual streams.

However, the connections between R&E networks are typically “one-off” specialty circuits dedicated to a specific purpose or short-term project and/or demonstration. The general consensus was that if R&E networks could work together under a common reference architecture, the reach of these global networks for research and education could be broader, easier to use and more robust.

Therefore, in 2013, the members of the Global R&E Network CEO Forum (CEO Forum) initiated a project to define a reference architecture for global R&E networking. Since then participation in GNA has broadened beyond the CEO Forum, and the GNA Technical Working Group is now autonomous. Any R&E network with an international focus that is willing to actively participate and contribute resources to this collective work is encouraged to join.

The aim of this project is to define a reference architecture that R&E networks and funding agencies around the globe can use to (further) align their investments to serve R&E on a global scale.

In order to define this reference architecture it is important to understand what the components of this reference architecture will be and how they will be described. The purpose of this document is to provide an overview of the components of that work as well as pointers to more specific documentation. Other documents referred to in this paper can be found at gna-re.net under the 'Resources' tab.

A Global Network Architecture

To fully develop a Global Network Architecture it is important to look at what is intended by use of the term "architecture" in this context.

Fundamental to this discussion is the understanding that the GNA is not a network; it is a *reference architecture* that can be used by organizations or consortia that are investing in global infrastructure to ensure their infrastructure is consistent with the GNA view of global connectivity.

It is a goal of the GNA that the collective set of resources that are built in conformance to this reference architecture will enable the seamless delivery of services across network boundaries worldwide.

At one fairly straightforward level, the GNA is described as a set of interconnecting circuits (international and national) between GXPs where participants' networks interconnect. For this to be an effective infrastructure there would need to be sufficient interconnects to ensure resilient, physically redundant and operationally stable services with well-documented SLAs.

A desired outcome of this activity is the ability to deliver services, across network boundaries, in a way that current interconnects are unable to deliver. This activity aims to enable new capabilities that in turn create new possibilities for interaction between R&E organizations.

Another goal of the GNA is for a participating organization to be able to deliver services to its users, at any place and at any time, with as close to the same characteristics as local delivery of that service. Some differences may be unavoidable (latency being the obvious example).

Ingredients of a Global Network Architecture

The Global Network Architecture is not a unified traditional approach where the main issues are physical design and cost. As individual components are brought into service by participating organizations, the degree to which those components are a useful part of the global infrastructure can be determined by how closely they conform to the reference architecture. Rather than a unified network design, the GNA is a multi-faceted design.

The overall GNA activity is described by a multilayer model:

- GNA Vision
- Funding
- Policies and Principles
- Organizational Framework
- Technical Architecture Planning

This document will not address all of the layers in this model. Other documents will describe the vision, potential funding models, policy considerations and organization. This document is focused on the Technical Architecture Planning. The areas relevant to this discussion that need to be considered include:

- Technical and Operations:
    - Acceptable Use Policy (AUP) considerations
    - Changing dynamics of circuit costs
    - Higher bandwidth (upgrades) and new technologies
    - Changing realities of where R&E needs to connect to do its work
    - Data privacy
    - Federated operations models
    - Network services beyond best effort IP
    - Production quality GXPs able to meet the service requirements of its participants.
    - Location of Content required by the participants
    - End-to-end SLAs
    - The Commons

Governance and Policy Components

The reference architecture describes a shared, cooperative activity. As such it is important that there be adequate governance models for each component and that there be a methodology for interaction between potential governance groups in the overall infrastructure. It is expected that future directions regarding the funding, policies and/or governance of the GNA will be discussed in a separate working group that is in the process of being formed. Watch the website gna-re.net for updates regarding the GNA and associated working groups.

Global Network Architecture GXPs

A key component of the reference architecture are the GXPs where the services of the GNA are delivered to the organizations participating in the GNA.

An Exchange Point must allow any R&E organization (assuming fair and appropriate payments) to establish a connection and may not force or restrict the ability of that organization to establish services with the other participants of the Exchange Point.

Separate documentation of all the capabilities of an Exchange Point is important and should be made available publicly. Templates for this are discussed in the Exchange Point documentation.

End-to-End SLAs or SLSs

The ability to deliver end-to-end SLAs or SLSs is critical to the success of the participating R&E networks. The vision of the GNA is seamless end-to-end service delivery, whether it is local or global, in a manner not otherwise achievable.

Along with specifying the characteristics of the services delivered, other parameters within the SLA should be specified. Among these would be:

- delivery time of a new service;
- time it takes to get feedback on the time needed to set up the service;
- cost to turn up a new service;
- performance characteristics; and
- repair times (such as MTTR, MTBF).

GXPs and interconnecting circuit owners must make available information about SLA's and SLS's.

Acceptable Use Policy

The GXPs in the GNA reference architecture are a critical component to the open and free exchange of traffic among the participants at an Exchange Point. All traffic that is originating or terminating inside an R&E participant's network at an Exchange Point must be able to traverse the exchange without limits or restrictions.

This does not imply that individual participant organizations cannot have or impose AUP restrictions within their networks. They are free to define their own policies as needed.

Interconnecting circuits that are provisioned under this reference architecture may restrict the type and/or volume of traffic that is allowed to transit the circuit.

All interconnecting circuit owners must make their policies publicly available in order to allow GNA participants to route traffic appropriately.

Economics and Cost Sharing of Circuits

By combining and aligning investments, R&E networks across the globe are able to create a more powerful and future proof set of interconnects of their national, regional and international efforts, for the benefit of Research & Education.

As the cost of these interconnects drops over time, the building of more diverse infrastructures will become both possible and desirable.

It is also expected there will be another significant step up in bandwidth requirements, as flows of well over 10 Gb become the norm. This growth will be seen especially in the big sciences and in cloud access.

Data Privacy

All participating organizations must have well established and documented policies on the collection and distribution of traffic data for everything that traverses their portion of the infrastructure. It is highly encouraged that these policies be made available online. More detail on privacy considerations can be found in a separate document.

Architecture Components: Content Delivery and Cloud Services

Content delivery and cloud-based services are currently in high demand globally throughout R&E. Therefore it is highly likely that content and cloud management strategies will enter into the GNA as it becomes fundamental to the activities of R&E. While the GNA working groups may not be directly dealing with this, work on storage and computing access might spin off from this activity.

One possibility is that GXPs should be open to the possibility of hosting content servers when that is feasible as well as providing access to content services that might otherwise only be available via commercial Internet services.

Given the increasing globalization of scientific instruments, and the desire to partner across borders on a global scale, having a comprehensive global connectivity strategy will likely be required for campuses and researchers across the globe.

Connectivity meeting the requirements of this reference architecture will allow that strategy to develop.

Architectural Components: Operational Requirements

There are also significant Operational elements to consider in this architecture. The aim is to enable participants to use well-defined service delivery models that are provisioned over a stable and secure infrastructure. Operational elements are described in a separate document.

Network Services Beyond Best Effort IP

The requirements for quality of service and other circuit based connectivity models will continue to drive the need to offer services beyond best effort IP. Best effort IP will still be a valuable and useful resource, but may not be adequate for the delivery of these advanced service types. Network services supported under the GNA will be described in a separate document.

The Commons

A critical piece of this architecture is providing a common space that is available to the community at large, independent of their participation in providing circuits or other components. Understanding the Commons, how it is used, its potential benefits as well as potential abuses will be explored in much greater depth in another document.

# # # # #