Towards an Information Architecture for the GNA

Version 0.6

Jon Dugan, Inder Monga, Ed Balas

Introduction

Science and commercial clouds are moving towards large scale data with global reach. Facilitating this trend requires the community of network service providers to find ways to provide an improved end to end experience for users of our collective infrastructure. In particular, this means finding ways to more effectively orchestrate provisioning, improve global situational awareness of operational issues and coordination across administrative boundaries to build the cross domain services needed to facilitate use of global clouds or enable worldwide scientific collaboration.

Improving situational awareness, orchestration and coordination require a consistent definition of what data needs to be shared for a given task. Global collaborations will not succeed if they are incapable of coordinating on a global scale. This implies that their planning processes need to be linked using a common terminology and interoperable technology. As a basis for collaboration, operators need to have a common understanding of what data needs to be shared, a model to facilitate reasoning about the data and policies for sharing policies this data. In this document, we will explore architectural options for sharing this data in a decentralized way between the various Global Network Architecture (GNA) participants.

General Characteristics And Scope

Given that there are approximately 240 countries in the world and some countries have more than one NREN, it seems that the GNA Information Architecture should be designed to scale to something on the order of 1000 participants.

Wherever possible this effort will take advantage of already existing technology solutions instead of reinventing the wheel. The means looking into the feasibility of drawing from existing work including, but not limited to perfSONAR, NSI and InCommon/Shibboleth. By using solutions that are already available we can move forward more quickly. Given that this is an architecture and not a design we also want to keep a clear separation at the conceptual layer between what kinds of information we need to exchange and the specific solutions used to implement them.

What Information To Exchange

There are four broad categories of information that need to be shared.

Documentation	Provides canonical data about what is available and how it is ideally configured.	Services, capabilities, topology
Operations	Provides information about the operational state of the network. Monitoring data may be used here to create issues or notifications.	Issue tracking, maintenance calendar, notifications
Measurement	Measurements of various parts of the system. Usually this is a time series of such measurements.	Interface counters (traffic & error), CPU load
Monitoring	Monitoring of parts of the system to determine their operations status. This may be as simple as watching the up/down status of a device or it may be more complex such as anomaly detection which is derived from the measurement data.	System status (up/down), service status, anomaly detection, thresholding

It is neither practical nor desirable for humans to be in the loop of exchanging this data, as a result it is important that all data is made available in machine readable format. Whenever possible the specifics of the machine readable formats should be adopted from other widely used standards. For example, perfSONAR provides several formats for exchanging various types of measurement data and NSI has developed a topology model that may be relevant here.

TODO: Add specific list of information needed?

Canonical vs. Actual

When talking about the network we often switch between describing the actual configuration or state of the network with the canonical state of the network. It is important to be clear about which is being discussed at any given time. The **canonical** state describes how we have decided that the network should be. It is often phrased as canonical, the source-of-truth or as a policy. The **actual** state describes how the network actually is and is often phrased as the current configuration or what is derived from monitoring.

It is possible for the canonical state to differ from the actual during outage or maintenance scenarios. Longer term decision making (such as advanced reservations/calendaring) are usually made on canonical state.

Information Exchange

Each participant has their own internal systems for tracking the information that will be need to be shared. We propose an architecture where the is a GNA Information Adaptor service, run by each participant, that provides the ability to retrieve information from internal systems and makes that information available to other GNA participants in a common format. See Figure 1 for a strawman of this architecture.

Exchange Model

Two approaches have been discussed. The first is a peer to peer model where each GNA participant creates a peer relationship with every other GNA participant for which they are interested in exchanging information. The other model that has been proposed is that an identical global representation of the GNA information is provided to every participant. Figure 1 is based on the peer to peer model.

It will be important to decide on a minimal useful subset of each kind of data to share. For example, for trouble tickets it is clearly necessary to share the ticket number and some summary of the contents, however it may not be desirable (or perhaps permissible) to share other details in the ticket. The information adaptor will provide the capability for each participant to decide what view of their internal information is made available to a given GNA participant.

There has been some discussion of whether a single global view of the entire GNA is provided to each participant or if participants can choose which resources they will share with each specific participant. A single unified view has the advantage of providing a way to do things like path finding across all parts of the GNA. However, it seems limiting to force all participants to share all resources with all participants. Instead it seems that the ability to share some resources with a subset of the participants provides significantly more flexibility.

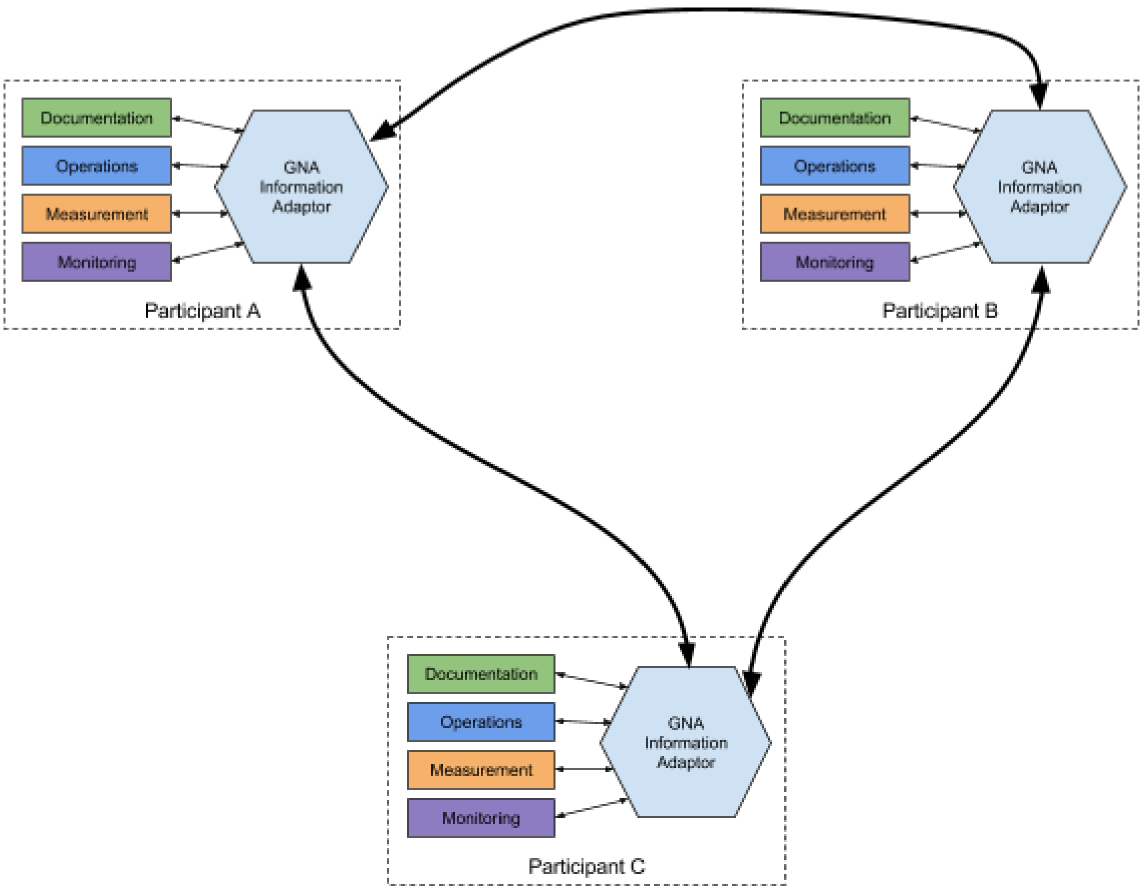

Figure 1 – Strawman Architecture

Path Finder

In order to explore this space, we seek information and participation from NRENs that will be interested in implementing a path finder to test out the principles presented in this document. Since implementing the full Information Architecture is a major undertaking, we propose to start with a small, well defined part of the overall vision. One option is presented below, but other options to be discussed within the larger GNA group.

Proposed Path Finder: Outage and Maintenance Notifications

The exchange of outage and maintenance events is a critical part of this kind of collaboration. The current best practice seems to be to send emails between NOCs in an undefined way. The core of this path finder is to use the problem of exchanging outage and maintenance notifications as an example of how we might exchange many kinds of data eventually. From a strategic point of view if we can succeed in this well scoped area, it will open many more possibilities for exchanging other kinds of data and enabling the larger vision of the GNA. From a technical point of view the investment here might be quite high, however that investment is justified by the strategic gains.

In collaborations such as ANA and AutomatedGOLE this underspecified exchange of information is a significant problem. In the case of ANA there are organizations depending on others to provide services and there is an agreement that notifications will be provided but the content and semantics of those notifications is no well defined. In a similar way, AutomatedGOLE would greatly benefit from notifications with clear semantics and naming so that problem areas can be automatically routed around. This also has an additional benefit of reducing the burden on NOC staff to translate emails into actionable data.

This seems to be a clear place where we could reap benefits for the community and work towards understanding the the issues that might arise in the implementation of the full information architecture. This also is a chance to leverage existing collaborations on ANA and AutoGOLE.

This path finder would work on defining what kind of information needs to be carried in these notifications and a format and mechanism for exchanging this information. It is key that the definition and semantics of the information be defined separately from the format and mechanisms of exchange so that the model can be reused as new formats and/or mechanisms replace current mechanisms. This description is intentionally high level as it is the work of this pathfinder to work out the details.

Questions to consider

What’s the the scale we need to solve for?

- The larger the scale the more design / complexity
- Smaller scale can be simpler
- Do we plan to support O(10), O(100) or O(1000000) participants?
    - Consensus in Chicago was O(1000) -- there are 240 countries in the world, some of which have more than one NREN

How can we reuse things that already work?

- Let’s not spend years re-inventing wheels.
- Can we adopt Is shibboleth / in-common for authentication? Is there a story from that community regarding authorization?
- Is there a way to build on perfSONAR work for sharing measurement data?
- Is it possible to use NSI definitions for exchanging topology information?

Data sharing model

- View
    - Everyone can seen everything (all resource announcements)
    - Pairwise decisions of what to share (ala BGP peering)
    - Everything for some classes of data / pairwise for other?
- What kinds of data do we need a global view for?
    - How to support pathfinding across multiple networks?
- Can we use pointers into perfSONAR data for example?
    - Instead of sending back some new format of data, stick with the perfSONAR format
    - We want to delegate the implementation.
- Can we reuse NSI topology?
- NSI Document Distribution Service as mechanism for sharing of data?
- Filtered / partial view
    - Especially important for issue tracking -- may not be willing/able to share all details

Discovery:

- Peer to peer between the GNA information Adaptors -- bring up a peering relationship with all of the participants you want to talk to
- Special DNS entry, gna-services.example.com
- perfSONAR lookup service?
- NSI Document Distribution Service?

AuthN/AuthZ

- Is shibboleth / in-common possible?
- public key exchange -- can manually add access to a adaptor given a public key

How do we make this work with a non-GNA provider or a non-compliant network (like a cloud provider)?

- Do we need a set of specifications and APIs that we can expose to broad community and ask for compliance?
- Is there a translator module that needs to be built?
- How do we compare the canonical to the actual?

What needs to be added to the document?

In addition to answering the questions above, the document should be enhanced with some notional exchanges, or adopt current norms to exchange information. Comments on the real-time nature of information sharing should also be discussed - what information is shared near-real time, what changes slowly.

Questions

- This document can contain just the architecture and high level discussion or go into technical details of standard format and information exchange? Whats most appropriate?
- How can ANA be pathfinder for this document?

Appendix

GNA Services

The current list of GNA services is:

1. Multi-point service
2. Overlay services (like LHCONE): L2 / L3 / MDVPN / Content peering
3. Point-to-point Anywhere-in-the-World
4. Monitoring & Trouble tickets (flowing back along provisioning lines across multiple networks / verification)
5. Tie into Compute & Tie into Storage (as part of the L3 overlay /towards CDN service & caching)
6. Slicing (of OSS)
7. Encryption & Certification on transport layer (P2P / P2MP /Overlay)
8. Security Services / Forensics / Boundary DDoS mitigation

This document uses the list of services above as it’s basis but strives to provide a level of generality which should allow it to extend beyond this list.

Notes / Scratchpad

This will be removed when the document is complete.

What should the information architecture contain:

I think at a first level we should provide a conceptual model of what types of information any GNA Network Operator is likely to posses and the at least common use cases where that data is shared. One of the goals for this architecture should be to provide just enough conceptual detail to allow Operations to create policies and guidelines that encourage appropriate levels of sharing. A second goal would be to place requirements on service definitions themselves as to what information requirements they should specify in their service definitions. There was also some discussion at the meeting about who drives resolution and sharing of data in particular notifications, i think a third goal of the architecture would be to define how these information flows are controlled.

- Monitoring
    - Alerts
- Measurements
    - Performance characteristics such as interface counters for traffic, errors, etc.
- Topology
    - External View (similar to NSI model)
- Capabilities (Services supported, Service Descriptions)
- Maintenance Events/Outages?
-
-

List of Services

Jon's list of pieces:

• Capability announcement
• what is contributed to the commons?
• Global service ID / handle• Reservation• Provisioning• Measurement• End user service portal• Verification
– Likely built on provisioning and measurement, but somewhat distinct
• Event notification
– Outages, service degradation, planned maintenance
• AAA

– It’s a tar pit but necessary
- Policy

– Local decision about who is trusted to use the shared capabilities?

Ed’s diagram:

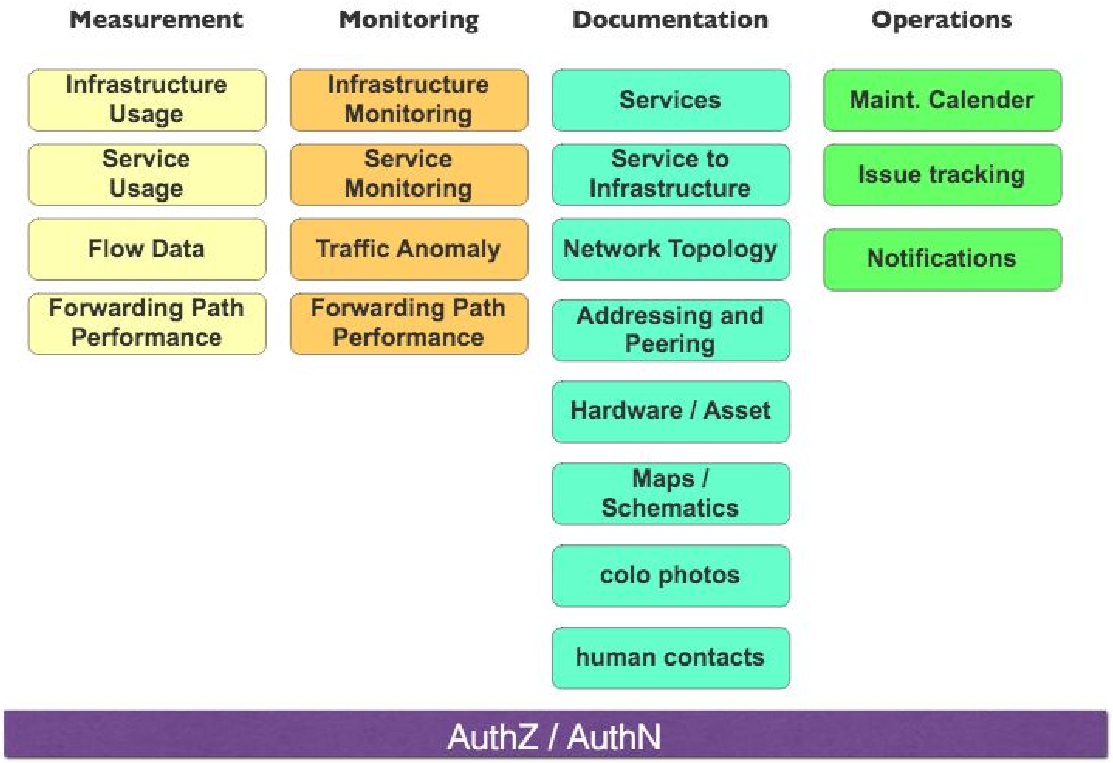

Measurement	Monitoring	Documentation	Operations
Infrastructure Usage	Infrastructure Monitoring	Services	Maint. Calender
Service Usage	Service Monitoring	Service to Infrastructure	Issue tracking
Flow Data	Traffic Anomaly	Network Topology	Notifications
Forwarding Path Performance	Forwarding Path Performance	Addressing and Peering	
		Hardware / Asset	
		Maps / Schematics	
		colo photos	
		human contacts	

AuthZ / AuthN

SURFNet

Data to share
- Service (de)commission details for all NOCs in service path (and back-up path)

– Service global ID, service parameters (e.g. type, guaranteed bandwidth or best-effort, endpoints, duration, addressing, user, SLA, policy, SDN controller)
and all characteristics to set it up
- NOC / user contact details per service• Reporting per service: usage, errors, outages (tickets) / monitoring• Reporting maintenance through tickets• Verification per service (does it ping? BW OK?)

Operations

• Portal for GNA NOCs

– To create/remove and test services

• testing done through measurement points at GNA Open Exchanges

– To view service reports / measurements– To lookup contact details for other domains, links, users, …– To review tickets?

• Automation of service setup and teardown

– Needs a network controller, topology description and per domain authorization– Needs collaboration on software development

●