# E2E Service Verification Architecture

## An architectural proposal for defining, engineering, and verifying Performance Guaranteed services

Jerry Sobieski

NORDUnet

GLIF 2011

Hong Kong  24 Feb, 2011

# Context and Motivation

- The R&E (GLIF) community has been exploring lightpath services for almost a decade now..
  - We are now recognizing that the practical useability of lightpaths is dependent upon a combination of technologies that are necessary to deliver a "lightpath _service_":
    - E.g. Standardized and ubiquitous user interfaces, on-demand and book-ahead scheduling, inter-domain reservation and provisioning processes, topology distribution and intelligent automated path computation, security and authorization,...

- A key feature common to these new services is the "performance guarantee" (PG)
  - i.e. the network _guarantees_ a certain performance level
  - Might be transport capacity, or availability, or some other characteristic,...or some combination there of.

- As dynamic global layer 0/1/2 connection services emerge, "*performance*" on these services means something different than conventional best-effort IP services:
    - These are not IP services so you cannot assume IP verification models will be appropriate (e.g. perhaps an ethernet VLAN is being requested...)
    - Not all "performance verification" is about packet loss or congestion (perhaps the service is required at a specific time, or via specific route, or with certain protection features...)
    - Virtualization and mult-protocol layering hide physical layer topology
    - Contemporary security, privacy,and scaling realities make detailed global E2E network information gathering difficult if not impossible.

# A New Notion of "Service"

- Just as service *delivery* of performance guaranteed services requires new service paradigms, service *verification* of these new services requires a new approach; different assumptions...
  - We need a fresh notion of "service architecture" that better addresses the global E2E issues of Performance Guaranteed services
- These questions have been debated in many forums of recent years –
  - Especially GLIF R&E community,
  - In commercial forums and consortia...
  - And quite recently, and in great detail, within the OGF NSI Working Group
- This talk will pose an architectural approach that tries to integrate these ideas into a formalized set of design prinicples.

# The Fundamental Transaction

- The "Fundamental Transaction" for PG network services:
  - The user requests, in advance, a specific level of service of the network
  - The network has an opportunity to verify that it can indeed meet the user's criteria, and arrange to do so...
  - If the network can meet the requested service criteria, it will respond with a confirmed commitment to the user.
  - If the network cannot meet the service requested, it then has an opportunity and responsibility to reject the request.
- A correlary of the Fundamental Transaction:
  - Both the requester and the provider should be able to *independently* determine if the service provided meets the requested constraints.
  - Guaranteed performance in a modern global internetwork cannot rely on trust that a delivered service meets spec -> service performance must be <u>measurable and verifiable</u>.
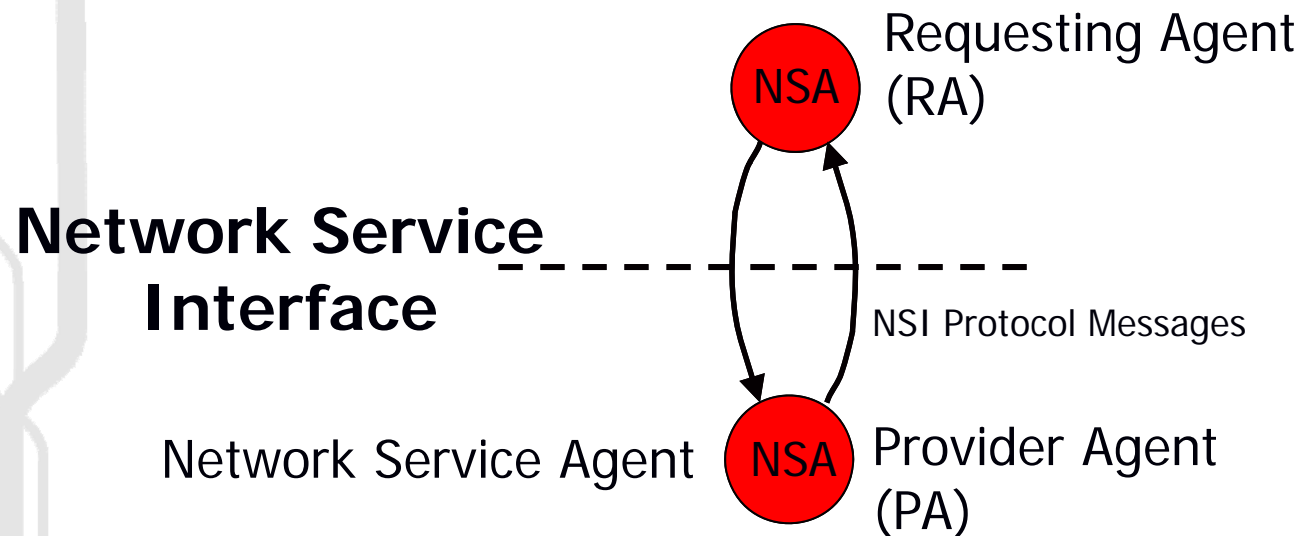
- The sacrosanct requirementof the FT:
  - Once the network commits to providing a service, it is irrevocably responsible for meeting that commitment.
  - A failure to meet the obligation – _for any reason_ – constitutes a service outage.
  - But the network is accountable to the constraints as formally presented by the request
- The concomitant responsibility of the user then is:
  - To be exact and complete in specifying the necessary service constraints, as there is no guarantee implied for any other aspect of the service
  - And be prepared to assume the cost of the resources allocated and dedicatd tofulfilling those constraints.
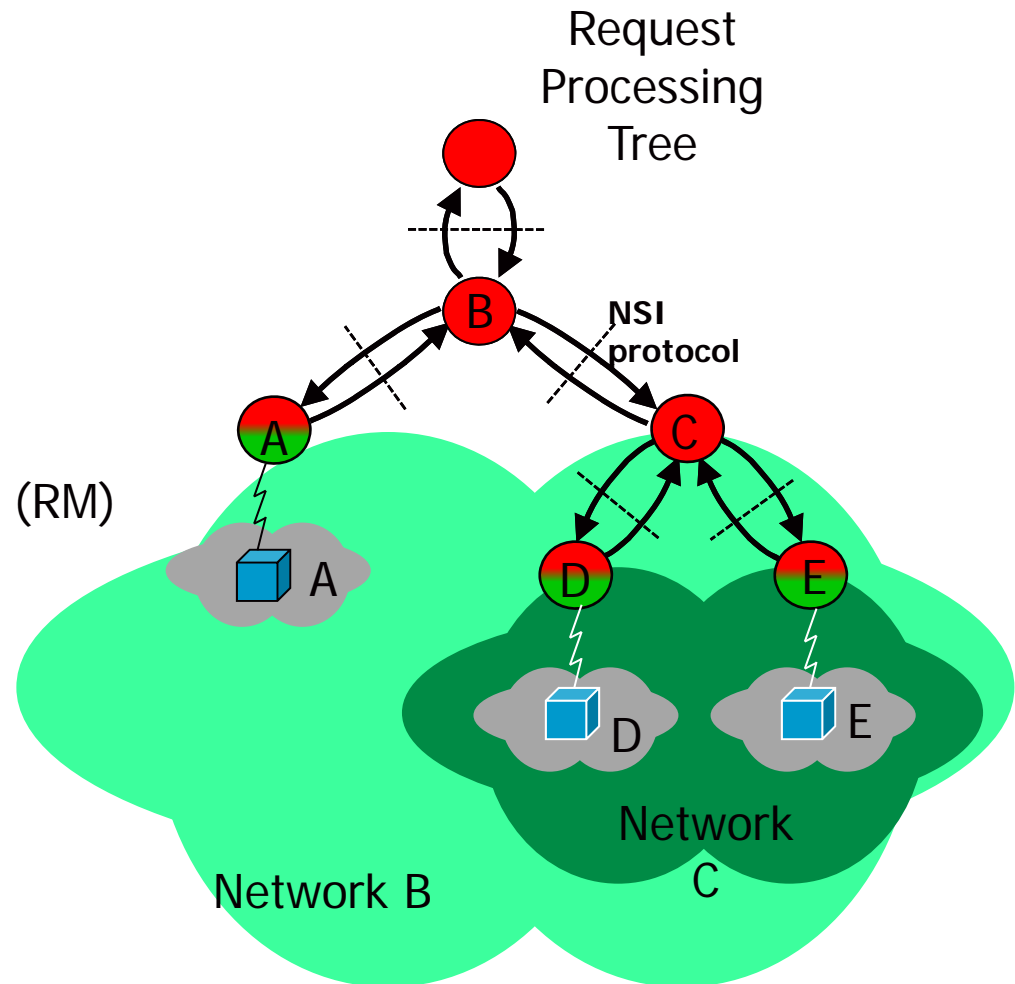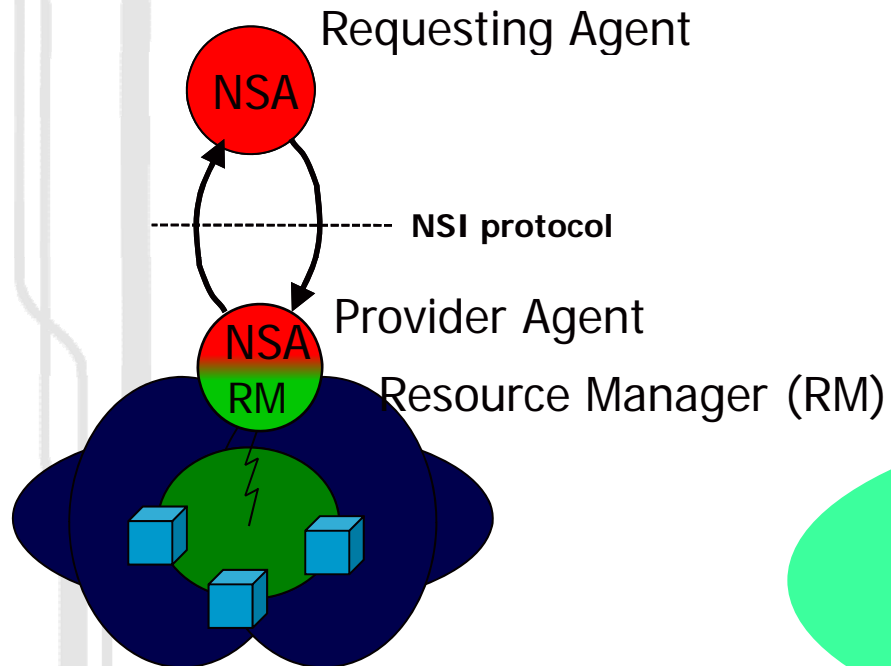
- The FT also *objectively* defines the service expectations:
  - If the constraints are satisfied, the connection is Good,
  - If the constraints are *not* satisfied, the connection is Bad.
- The fundamental transaction established "delegated responsibility"...
  - Delegated responsibility can be used by the network to subdivide the user's request into smaller pieces and delegate each piece similarly to other service providers
- Delegation hides complexity and provides scalability
  - Each network can function as an opaque autonomous system – it's internal structure and processes are private
- Likewise, delegation enables federation
  - "Networks of networks" can be composed that reflect common service preferences or shared resources...

**NORDUnet**
Nordic infrastructure for Research & Education

- The OGF Network Service Interface WG – spawned from GLIF activities – has described a service interface architecture that supports the Fundamental Transaction.

- The Network Service Interface (NSI) Framework describes a set of interactions between a Requesting Agent (the user) and the Provider Agent (a network).

  - The NSI Connection Service protocol (NSI-CS) begins with a ReserveRequest primitive that constitutes the fundamental transaction.

**NORDUnet**
Nordic infrastructure for Research & Education

**Network Service Interface**

Network Service Agent

NSA — Requesting Agent (RA)

NSA — Provider Agent (PA)

NSI Protocol Messages

# Delegation in the NSI Model



Requesting Agent

NSA

NSI protocol

Provider Agent

NSA
RM

Resource Manager (RM)

Request Processing Tree

B

NSI protocol

A

C

D

E

A

D

E

Network B

Network C

The user application

**NORDUnet**
Nordic infrastructure for Research & Education



Ingress
Service Termination Point
"A"

Egress
Service Termination Point
"Z"

Access section
Ingress Framing

Transport section
Transport framing

Access section
Egress Framing

- The User (RA) specifies connection constraints (ostensibly externally measurable) for the access portion of the service instance
- The Network (PA) decides how to fulfil those constraints across the transport section.

- The FT relies on clearly specifying any and all constraints associated with a connection request:

```
ReserveRequest{
    Orig="//NTNU/CloudCluster-p1";
    Dest="//KeioUniversity/Vizstation3";
    Capacity=1 Gbps;
    StartTime=2011/3/1 00:00:00 EST;
    EndTime=2011/3/1 18:00:00 EST;
    FrameLossRate=10^(-8);
    BurstSize=1 Gbps;
    Auth="Abcdef";
    Policy="prefer domain=[NORDUnet,GEANT,JGN2Plus]
        include domain=MANLAN, exclude=domain USLHCNET"
}
```

- "Service Definitions" describe the scope and range of the service offered by a particular network.

- The Service Definition abstracts the network specific service capabilities away from the processes and functions that implement the [NSI] Connection Service

  - This allows the service implementations in different network domains to be defined and tailored as a process separately from the software and service architecture that delivers these services.

- Each network defines their respective service offering (SD)

- A group of networks can get together and define a *common* service definition

**NORDUnet**
Nordic infrastructure for Research & Education

- The Service Definition specifies all relevant service parameters associated with a particular network service offering
  - A provider constructs the SD docment listing the service parameters they are willing/able to support
  - Each parameter has a range of allowed values and a default where appropriate
  - The Service Definition acts as a template for service requests; User specified values overlay those found in the SD base.
  - Thus all service requests are fully specified

## NORDUnet
**Nordic infrastructure for Research & Education**
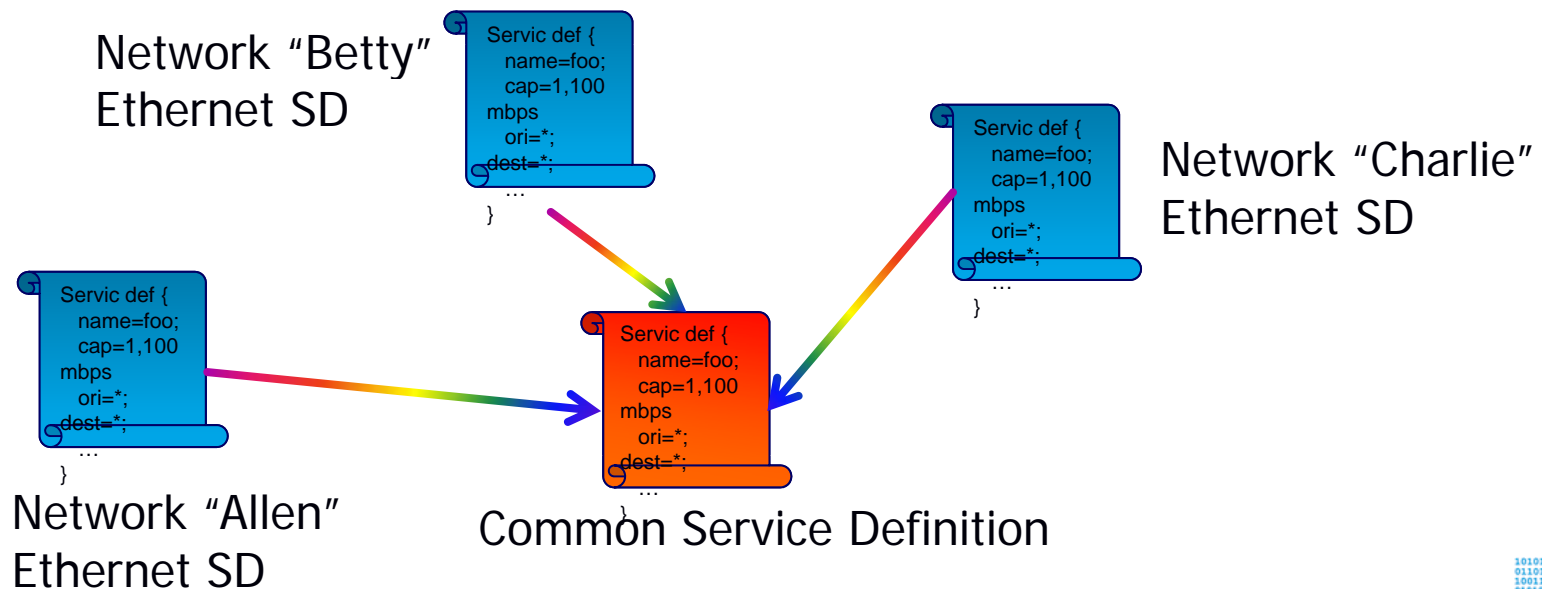
- Service Definitions are defined separately from the service protocols

```
Service Defintion{
    Name="GLIF-Ethernet";
    Guarranteed {
        Orig:=[NSI-DirectorySvc(<p>)];
        Dest:=[NSI-DirectorySvc(<p>)];
        Capacity:=[Range(1,1000,1)] Mbps;
        Framing:=802.1ad |802.1, default=802.1;
        StartTime:= DateTime(<p>);
        EndTime:= DateTime(<p>) | Dur(<p>);
        Auth:=[EduGain(<p>)];
    }
    Prefered {
        Policy:= default="none";
    }
}
```
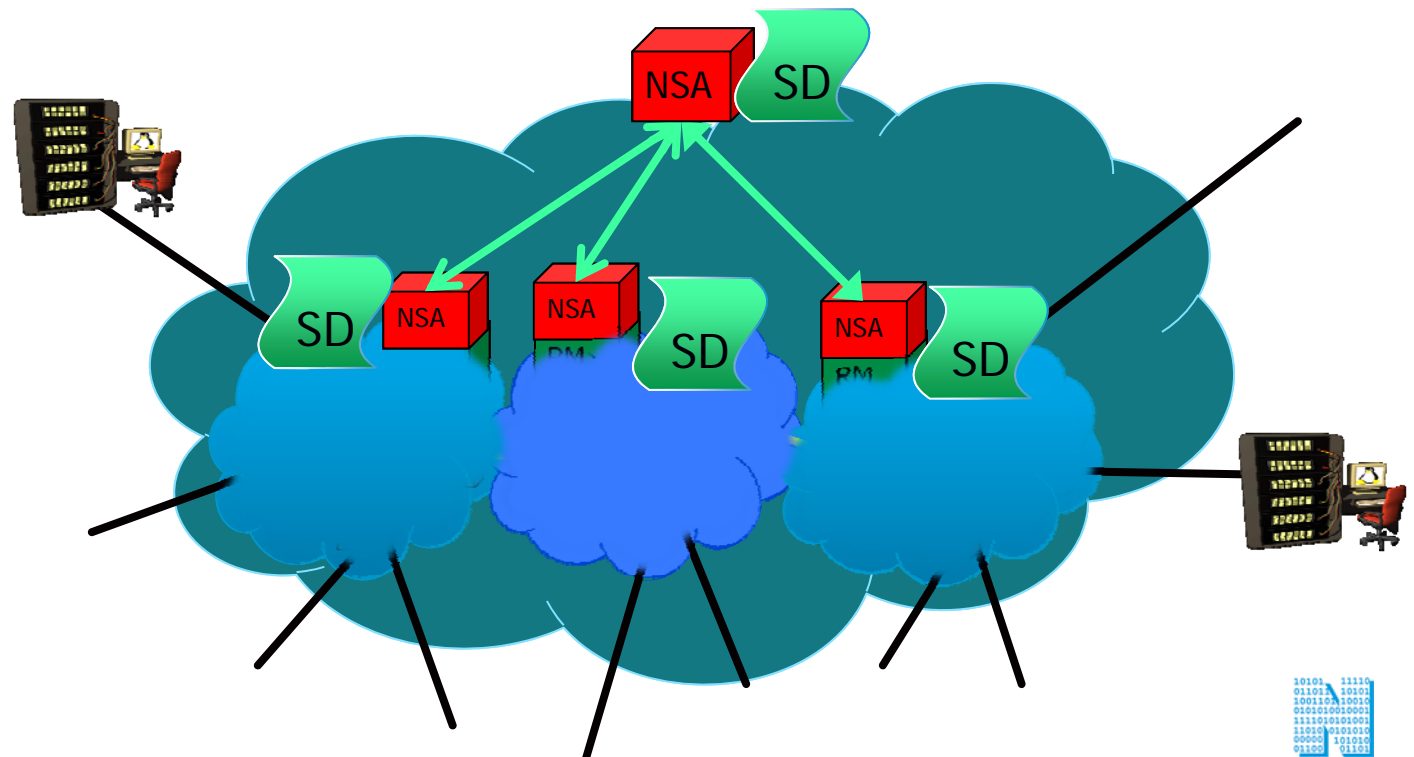
**NORDUnet**
Nordic infrastructure for Research & Education

- Each network defines their own service offering (SD)
- A group of networks can then get together and define a "common service definition" CSD
- The CSD contains a set of agreed common service parameters, but each network can still augment this set and/or set specific ranges for the common parameters

Network "Betty"
Ethernet SD

```
Servic def {
    name=foo;
    cap=1,100
mbps
    ori=*;
    dest=*;
    ...
}
```

```
Servic def {
    name=foo;
    cap=1,100
mbps
    ori=*;
    dest=*;
    ...
}
```

Network "Charlie"
Ethernet SD

```
Servic def {
    name=foo;
    cap=1,100
mbps
    ori=*;
    dest=*;
    ...
}
```

```
Servic def {
    name=foo;
    cap=1,100
mbps
    ori=*;
    dest=*;
    ...
}
```
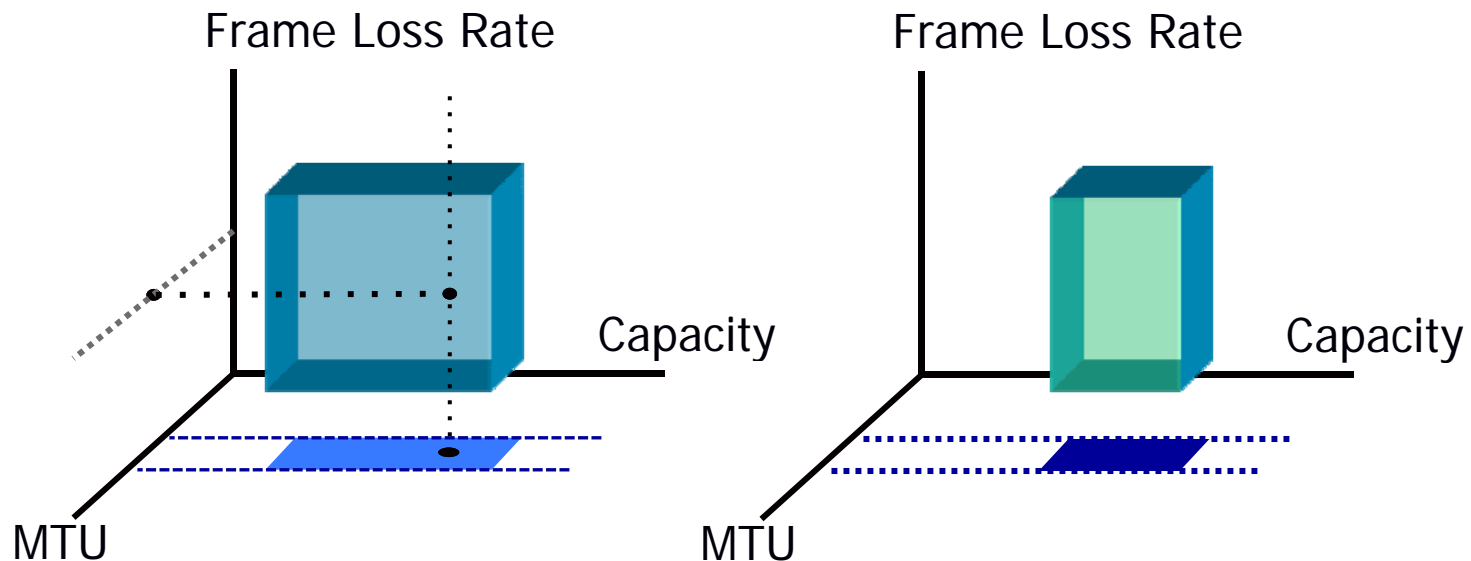
Network "Allen"
Ethernet SD

Common Service Definition

- Common Services are established when several networks agree on a common set of service parameters and values

# Service Definition

- A Service Definition can be conceived of as a multi-dimensional volume in n-space.
  - Each service parameter defines an axis
- Service Requests can similarly be conceived as a point within the same n-space
  - Service requests that project inside this n-dimensional service space can be provisioned.
  - Service requests that lie outside cannot be established.

Frame Loss Rate

MTU

Capacity

**NORDUnet**
Nordic infrastructure for Research & Education

- Since peering networks may have slightly different service definitions, there is a need to "compare" service requests E2E for compatibility:



- Individual service requests can be projected onto this intersection to ascertain path viability

# Service Verification

- A confirmed service instance with guaranteed performance criteria should be independently verifiable:
  - i.e. the user should not have to take the word of the provider that the service meets spec; (provider based assurances have an inherent conflict of interest)
  - Required service constraints should be *detectable* by the user.  If the constraints were indeed met, then they can be detected and measured.

- A user agent must be able to measure the "as built" service characteristics and verify that all the constraints were met
  - For pure performance and schedule constraints, conventional iperf style measurements at the end points will suffice
  - But for policy constraints, independent verification is not obvious as to how it would be accomplished...

- The provider should similarly be able to verify independently (from the user) that the "as built" performance meets spec
  - For pure performance and scheduling aspects, the PA should be able to verify that the information flow presented by the user at the ingress is within profile and that this flow is being presented at the egress intact.
    - Why? E.g. Tcp windowing problems are examples of end system issues that masquerade as network performance problems... The pro-active network is able to refute or coraborate user claims of service outages.
- But policy verification is much more insidious...
  - How can a user independently verify policy constraints?
    - E.g. "include MANLAN" or "do not transit Libya" ...
  - Access to as-built and other circuit information will be necessary to confirm policy cpnstraints

**NORDUnet**
Nordic infrastructure for Research & Education

Ingress
Service Termination Point
"A"

Egress
Service Termination Point
"Z"

Transport section

The Network measures performance between
these two end points from inside the transport section

The User measures performance
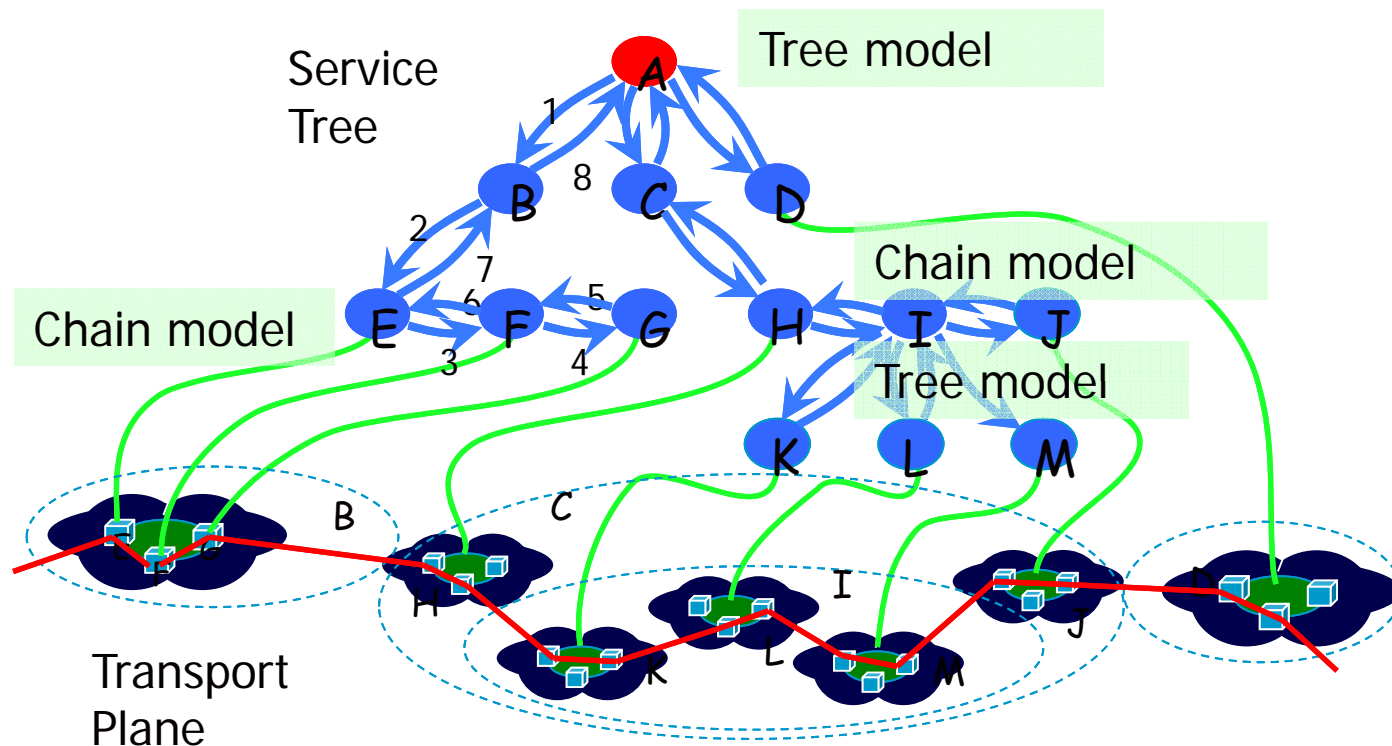between these two end points from outside

Transport framing
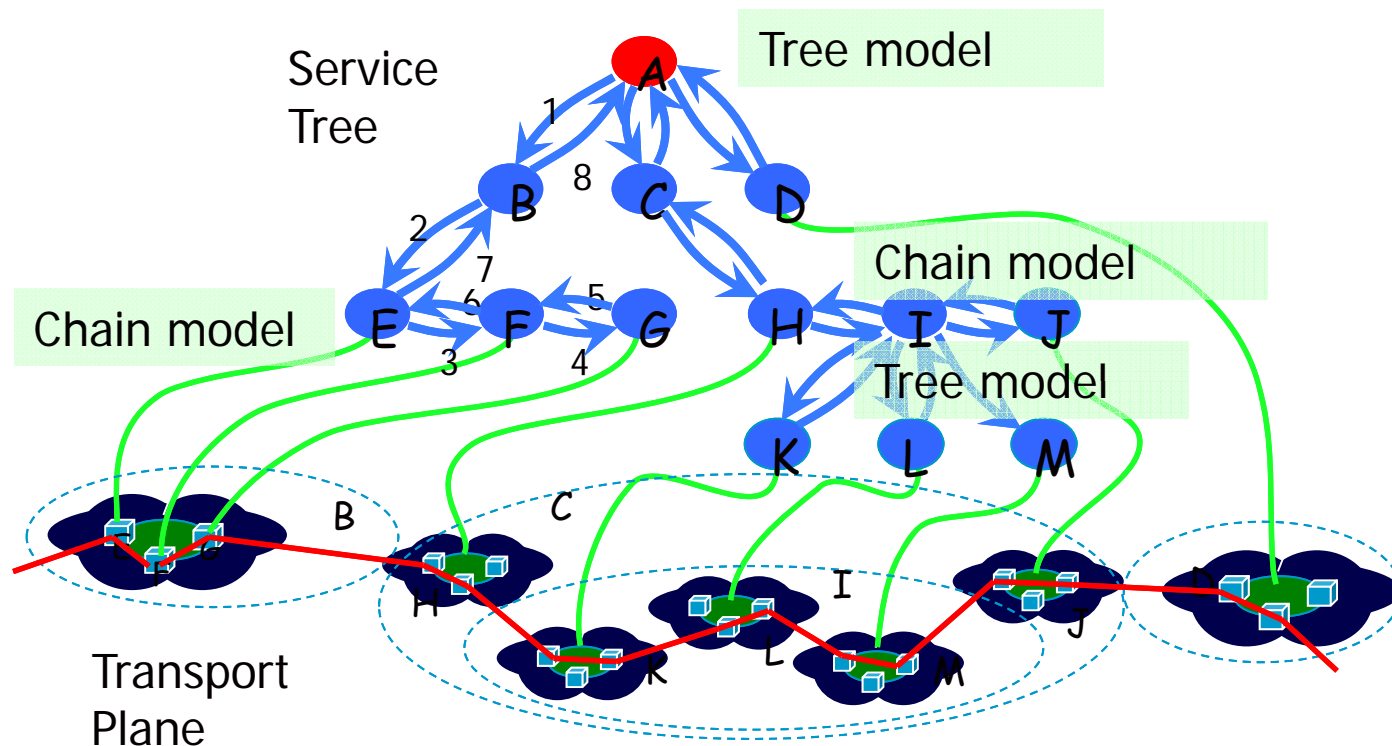
# Verification Processes

- Circuit services move information from ingress to egress.
    - The end points are key for performance and availability constraints
- A circuit oriented performance verification and AFU process should be congruent with the service architecture-
    - I.e. Performance verification should correspond to the service requests.
    - By construction, and asserting delegation, we can identify which service instance is not meeting spec.

- The service reservation process creates a "service tree".
  - The path and authorization structure is inherenent in the NSAs that comprise the tree for any given request

**NORDUnet**
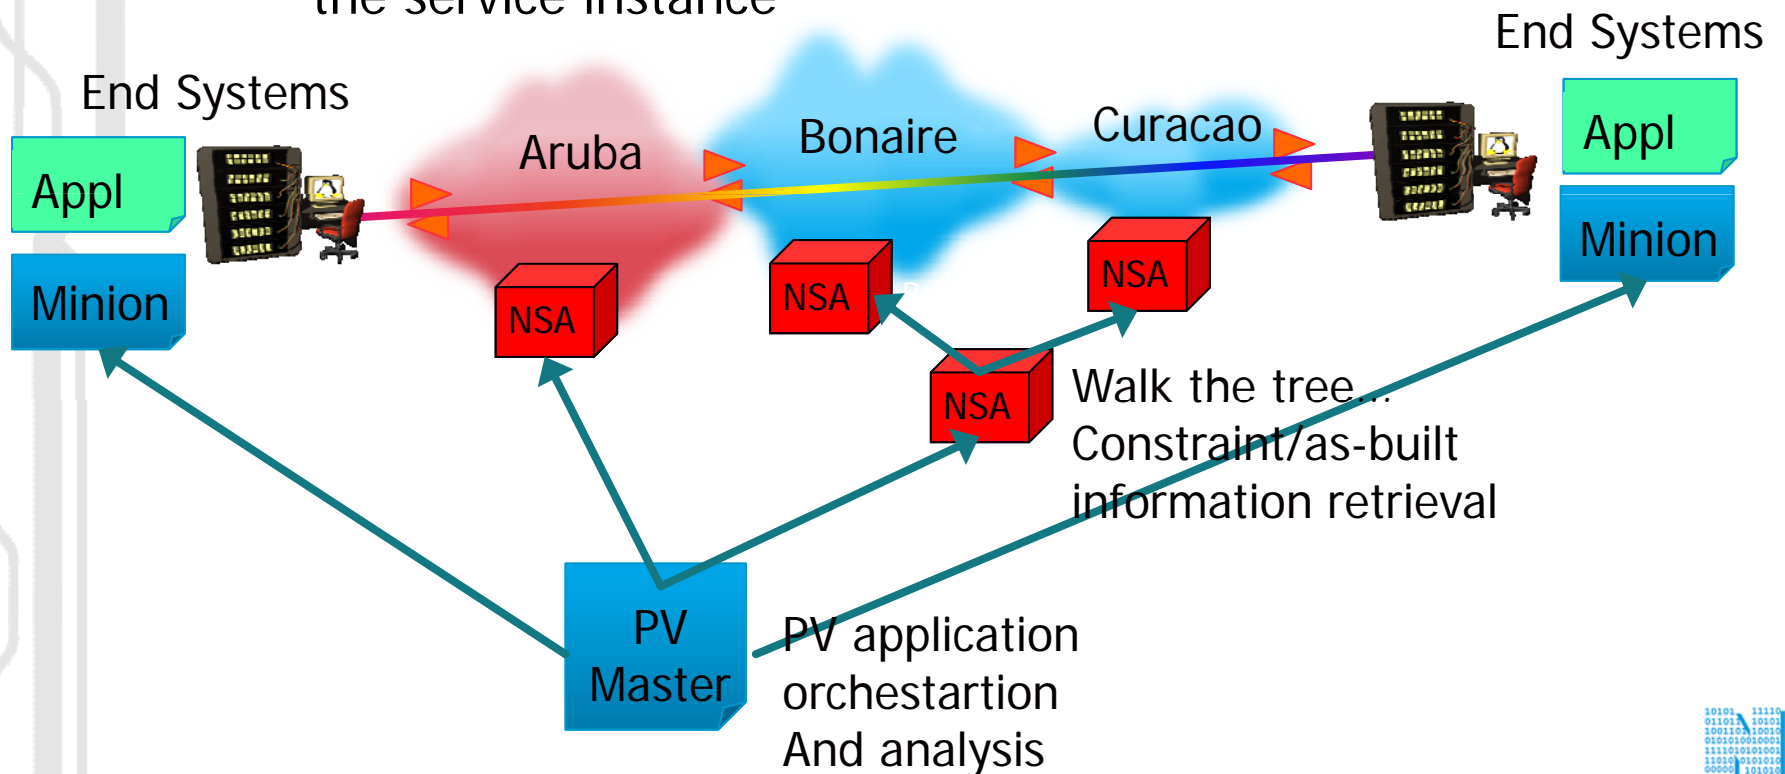Nordic infrastructure for Research & Education

- By walking the service tree, an AFU agent could [theoretically] reconstruct, under authorized access, the delegation breakdown and data plane path of a service instance.

Service Tree

Tree model

Chain model

Chain model

Tree model

Transport Plane

**NORDUnet**
Nordic infrastructure for Research & Education

- Since the service instances will be integral components of distributed applications, it stands to reason that the PV/AFU process should also be a similar distributed application

  - The PV/AFU application dynamically assume a congruency with the service instance



End Systems

End Systems

Appl

Minion

Aruba

Bonaire

Curacao

Appl

Minion

NSA

NSA

NSA

NSA

Walk the tree...
Constraint/as-built
information retrieval

PV Master

PV application
orchestartion
And analysis

# Summary

- The PV/AFU process must be adapted to reflect the nature of the services under test, and the environment in which they will be used.

- Independent verification of service delivery is critical to future application requirements

- Delegation of responsibility for PG services allows us to decompose the problem to divide and conquer

- New protocols and new models of integrating the PV function into the user virtual environment must be designed and developed

- We must assume a secure and autonomous network model for these PV and AFU processes – no more free access to sensitive information. But access should still be available to authorized agents.

**NORDUnet**
Nordic infrastructure for Research & Education

- We need a conprehensive approach that considers how performance verification processes is intrinsically linked to and parallel with the delivery of performance guaranteed services.

- Such a formalized architecture will be much more effective approach to performance verifcation and fault localiztion.


- Lets engage on this...