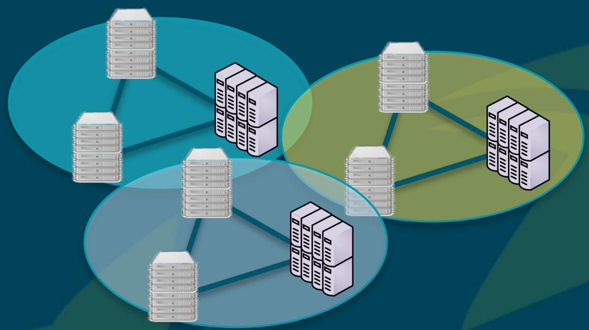

GEANT and “Programmable” WAN Services

Jerry Sobieski (NORDUnet)

GEC 19

Atlanta, Mar 18, 2014

Short Answer:

- **NSI (!)**

- **Network Services Interface**
  - Standard OGF
  - Explicitly deals with inter-Domain provisioning
  - Leaves intra-domain provisioning to local admin

- **Deploying within the production R&E Networks**
  - GEANT, NORDUnet, SURFnet, Esnet, StarLight, NetherLight, Internet2, ...

Key Network Components

- Transport Links
  - Transparent pipes from point A to point Z
  - There is no research value for the pipe... It simple moves data
- Where research,experimentation, forwarding, etc is required...
  - → You break the pipe and reconnect to a intermediate processing node
  - A [virtual] switch, router, VM, OpenFLow fabric...
- This model preserves the pipe as the virtualized resource that delivers raw information between Nodes.
  - The model allows the transport service to abstract and hide the underlying transport services.

A Programmable WAN

- The "WAN"
  - The "Wide Area Network" – presumably differentiable from the LAN or MAN...
  - **Why is WAN special?? – Multi-domain, core infrastructure, serving 10^6+ users, ...**
- Programming the WAN
  - The Programming is in the routers and switches
  - The transport links (between forwarding elements) remain transparent pipes – arguably always will.
- This also implies that there is some other devices that are not "network"...
  - ***Assert: All components are network components – we only really care how they connect and who owns them...***

Visibility and Authority

- Programming the WAN is [arguably] a multi-domain process
  - (otherwise it's a intra-domain issue and no different than programming your lab or campus.)
- What is a "domain"?
  - Technology boundary – homogeneous protocol region
  - Service boundary – A region that is served by a common service
  - **Administrative boundary – authorization policy**
- MD provisioning implies that you have visibility to some map of the MD environment – topology.
  - The topology information is non-verifiable... A domain can announce whatever it wants (!) – or only what it wants.

Questions/thoughts

- Virtualization...
  - Networks are not flat... Services overlayed on other services...over other service... Over yet other services...
  - Networks are not homogeneous...layer 0,1,2,3,4,5/eth/SDH/OTN/IP/MPLS/
  - Networks are not public (!) – we really know *very little* about their state or engineering...and this trend will accelerate.
- Security and Privacy ... → Trust
  - Making the behaviour of the network (WAN or otherwise) dynamic poses increased risk that the wrong agents will program it...and/or they will program it wrong.
  - How do we insure that information remains secure? (This is relevant to all applications!)

Whats required?

- A transport provisioning service with global reach
  - NSI – Network Services Interface v2.0 standard
  - Pathfinding -> Global Topology
  - Topology representation/architecture/distribution/security/ verification and semantics is still very immature
- A service that provides computational elements (including switching and forwarding nodes)
  - Must be virtualized if we want to be dynamic
- We need *virtual* resources –
  - No one will give you control of their actual infrastructure
  - You don't want other "you" messing with your testbed
  - Virtual allows us to jointly manage expensive resources as we need them

SDN in the Inter-domain

- In order to program the WAN..
  - The agent must know the topology among the fabrics (Topology visibility again!)
  - The agent must have rights to modify the fabrics (authority again)
- These rights are implicit in the data center or the campus...*but not in the WAN*
- There must be some way to announce a service topology and a means to learn a service topology available to your agent...
  - A domain need not announce their entire topology (summarization)
  - ...or even their real topology... (abstraction)

Virtualization

- Service providers (that provide programmable computational platforms) can provide *virtualized* elements
  - Virtualization bounds the capabilities of the resource.
  - OpenStack, VMware, Parallels, ... All can provide VMs.
- Other services can provide the transparent transport links between the [virtual] computational elements
  - NSI services: OSCARS, AutoBahn, GLAMBDA{A,K}, OpenNSA, etc.
  - These service provide a single integrated contiguous global connectivity...i.e. they can link computational elements located anywhere [in any NSI domain]

What does GEANT Provide?

- SA3 Bandwidth on Demand
  - Ethernet framed Point to Point Connections
  - Transparent “pipes” between GEANT NREN partners
  - NSI standard
    - *Inter-domain provisioning protocol*
    - *Dynamic, scheduled, guaranteed*
    - *AutoBAHN implementation*
  - International service reach (EU to NA, SA, APAC,...)
- Wave Services
  - Low layer transport across/using the EU R&E infrastructure
  - 10Gbps, (40,100 possible...)
- GEANT NREN partners
  - Regional national R&E infrastructure to reach GEANT
  - Academic partners

The GEANT Testbeds Service
From Innovation to Infrastructure

- Network Innovation requires testing to prove out...
  - Testing in live networks can have unintended effects on non-combatants. Other users and network providers don't like being **crash test dummies.**
  - "Production" environments have the required scale but are *highly* risk averse.

How do we evolve innovations from concept to production with minimal risk to infrastructure, services, and applications already in place providing on-going stable and reliable services?

Networking R&D Laboratories

- The network research community needs "Laboratories" to test novel concepts ...
  - Constructed from stable underlying infrastructure
  - Allow high risk experiments to be carried out...
  - Yet prevent unexpected or errant behaviour from interfering with production services or other research activities
  - Provide reliable and effective work environment for the researcher
  - Enable a broad range of innovation – i.e. technology agnostic
  - Agile: Ability to rapidly prototype new ideas or integrate new results
  - Scalable: Ability to construct large scale test environments
  - These laboratories must be able to duplicate real world scenarios such that research results are useful and valid

GN3+SA2 Testbeds as a Service

GÉANT

- **SA2 key service capabilities:**
  - **Dynamic "Packet" Testbeds** – *dynamically allocated, virtual networks* provisioned over production transport and switching infrastructure with a pan-European footprint.
    - *Under control of the researcher*
    - *Insulated to prevent collateral damage*
    - *Flexible user defined network resources, can morph as necessary*
    - *Extensible support for novel hardware*
  - **"Dark Fiber" Testbeds** –photonic testbeds over dark/dim fiber along long haul routes between a limited set of major EU metro areas.
    - *Virtualization of these resources is hard...but we'll see...*
  - **"GOFF"** – **a prototype OF testbed** originally fielded by GN3-JRA2
    - *Software emulated OpenFlow switching (OVS)*
    - *Bridge service as TaaS ramps up...*

- **SA2 is a GEANT *Production Service***
  - The test beds it creates are expected to be reliable and consistently available.
  - Which means the SA2 support processes must be stable and secure
  - This integrated "multi-species" virtualization represents new technology and continues to evolve in the community ... There continues to be many research efforts, and many emerging frameworks and service models...

A Brief Dive into the Internals:

The TaaS Architecture treats all [testbed] networks as *graphs*

Testbed "Alpha" Description

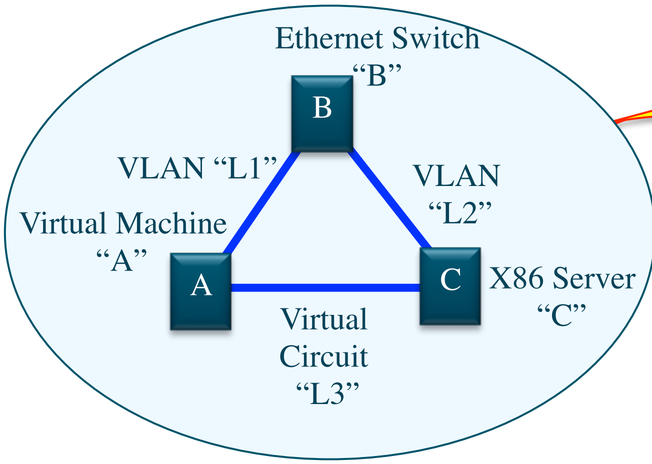

Internally, TaaS represents all testbed components as "**virtual resources**" with data flow Ports. User specified Port adjacency relations define the testbed topology.

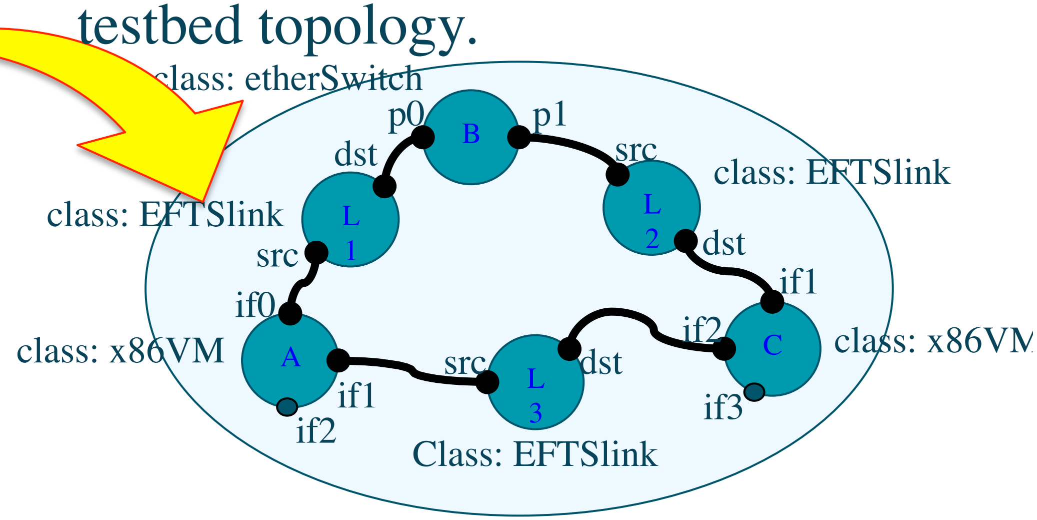

Data plane resource graph

TaaS Dynamic Testbed Provisioning - How it works

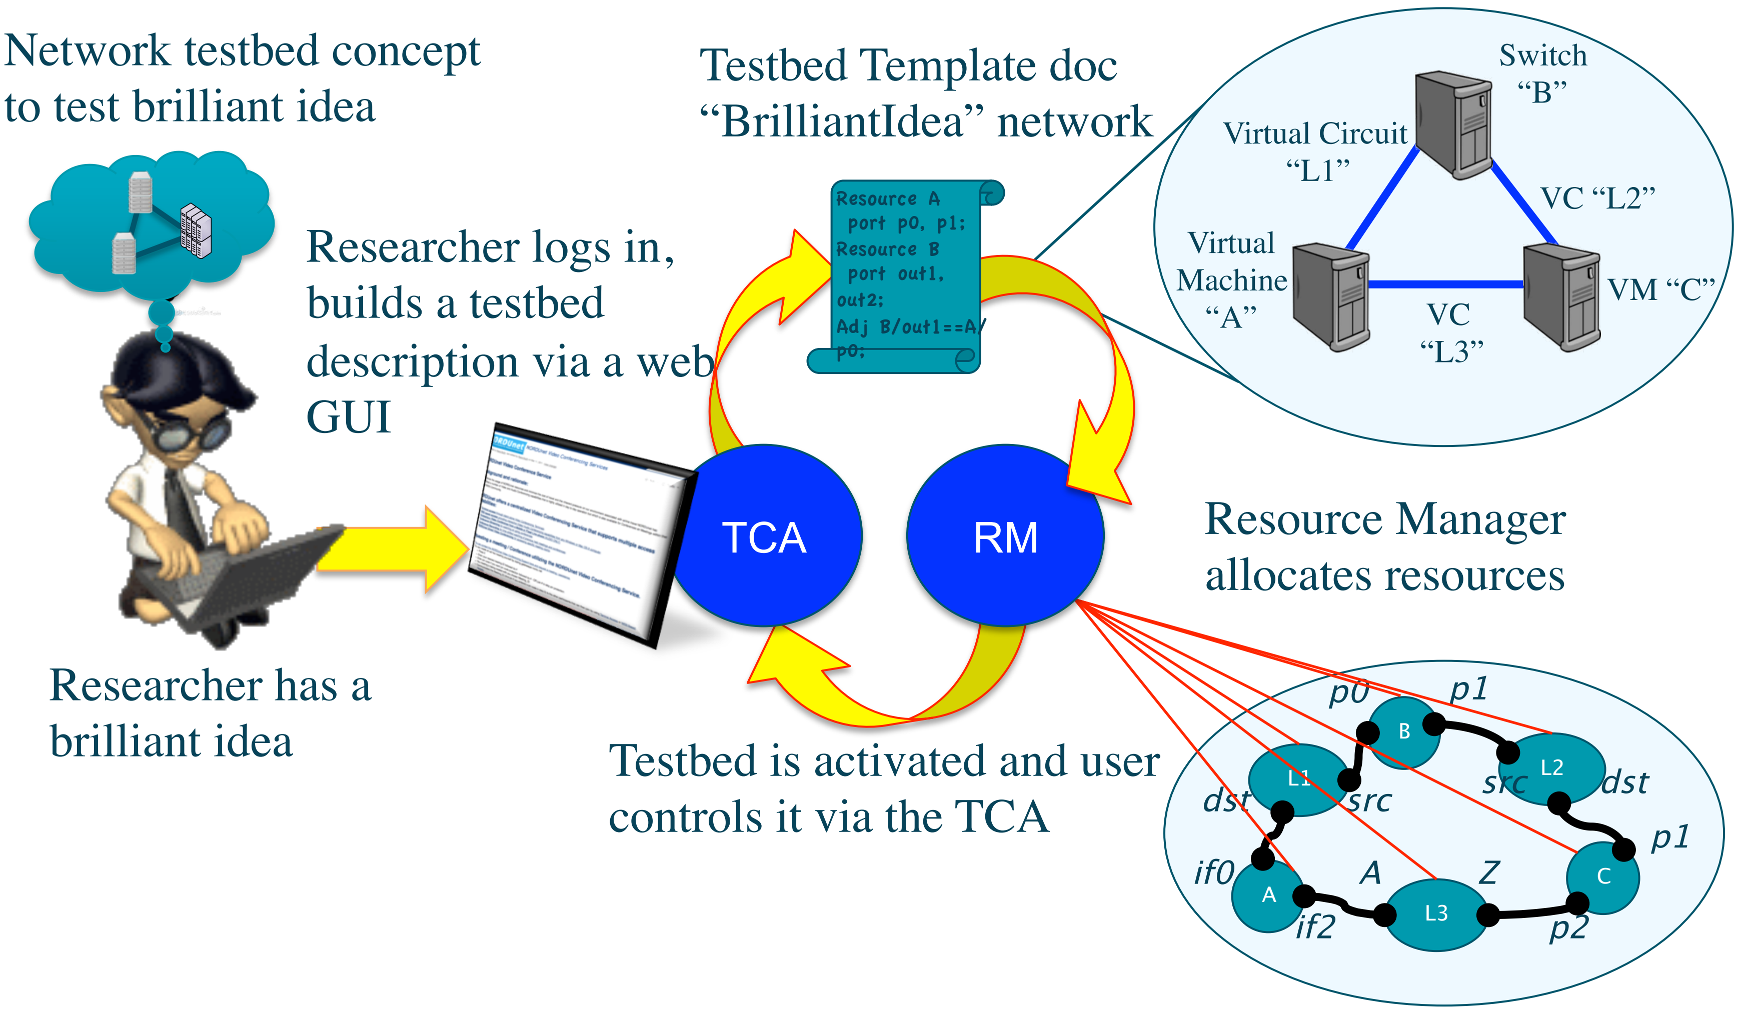

Testbed DSL Endgame ... (not to scale)

```
IMPORT foo.dsl

TYPE Foo {
    PORT p1 { attributes=transparent }
    PORT p2 { attributes=transparent }
    ATTRIBUTE memsz=4GB
    ATTRIBUTE location=? {
        CPH | AMS | BRA | LON }
    --- Children within the composite ---
    LinuxVM es1 { cpuSpeed=2.3GHz }
    LinuxVM es2 { cpuSpeed=2.0ghz }
    EftsVC  L1 { capacity=1000Mbps }
    ADJACENCY es1.p1, Foo.p1
    ADJACENCY es1.p2, L1.src
    ADJACENCY es2.p1, Foo.p2
    ADJACENCY es2.p2, L1.dst
    --- No methods defined by user ---
}

IMPORT testbed.dsl

TYPE Testbed {
    Foo Alpha, Beta {}
    ADJACENCY Alpha.p1, Beta.p2
    ADJACENCY Alpha.p2, Beta.p1
}

Testbed myTestbed1
```

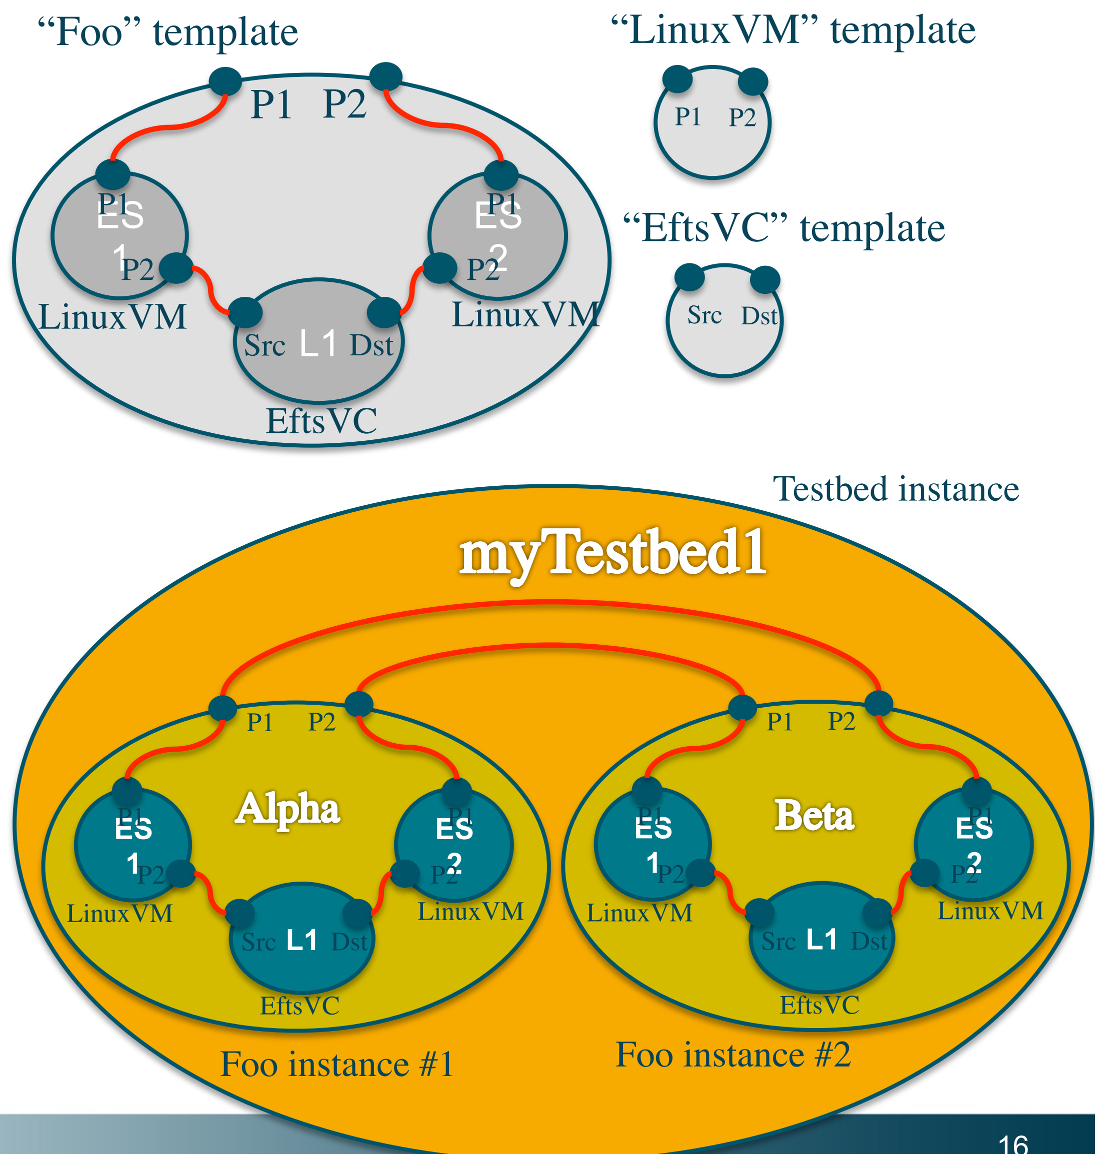

GEANT "Testbeds as a Service" Virtualization, Management, and Control Layers

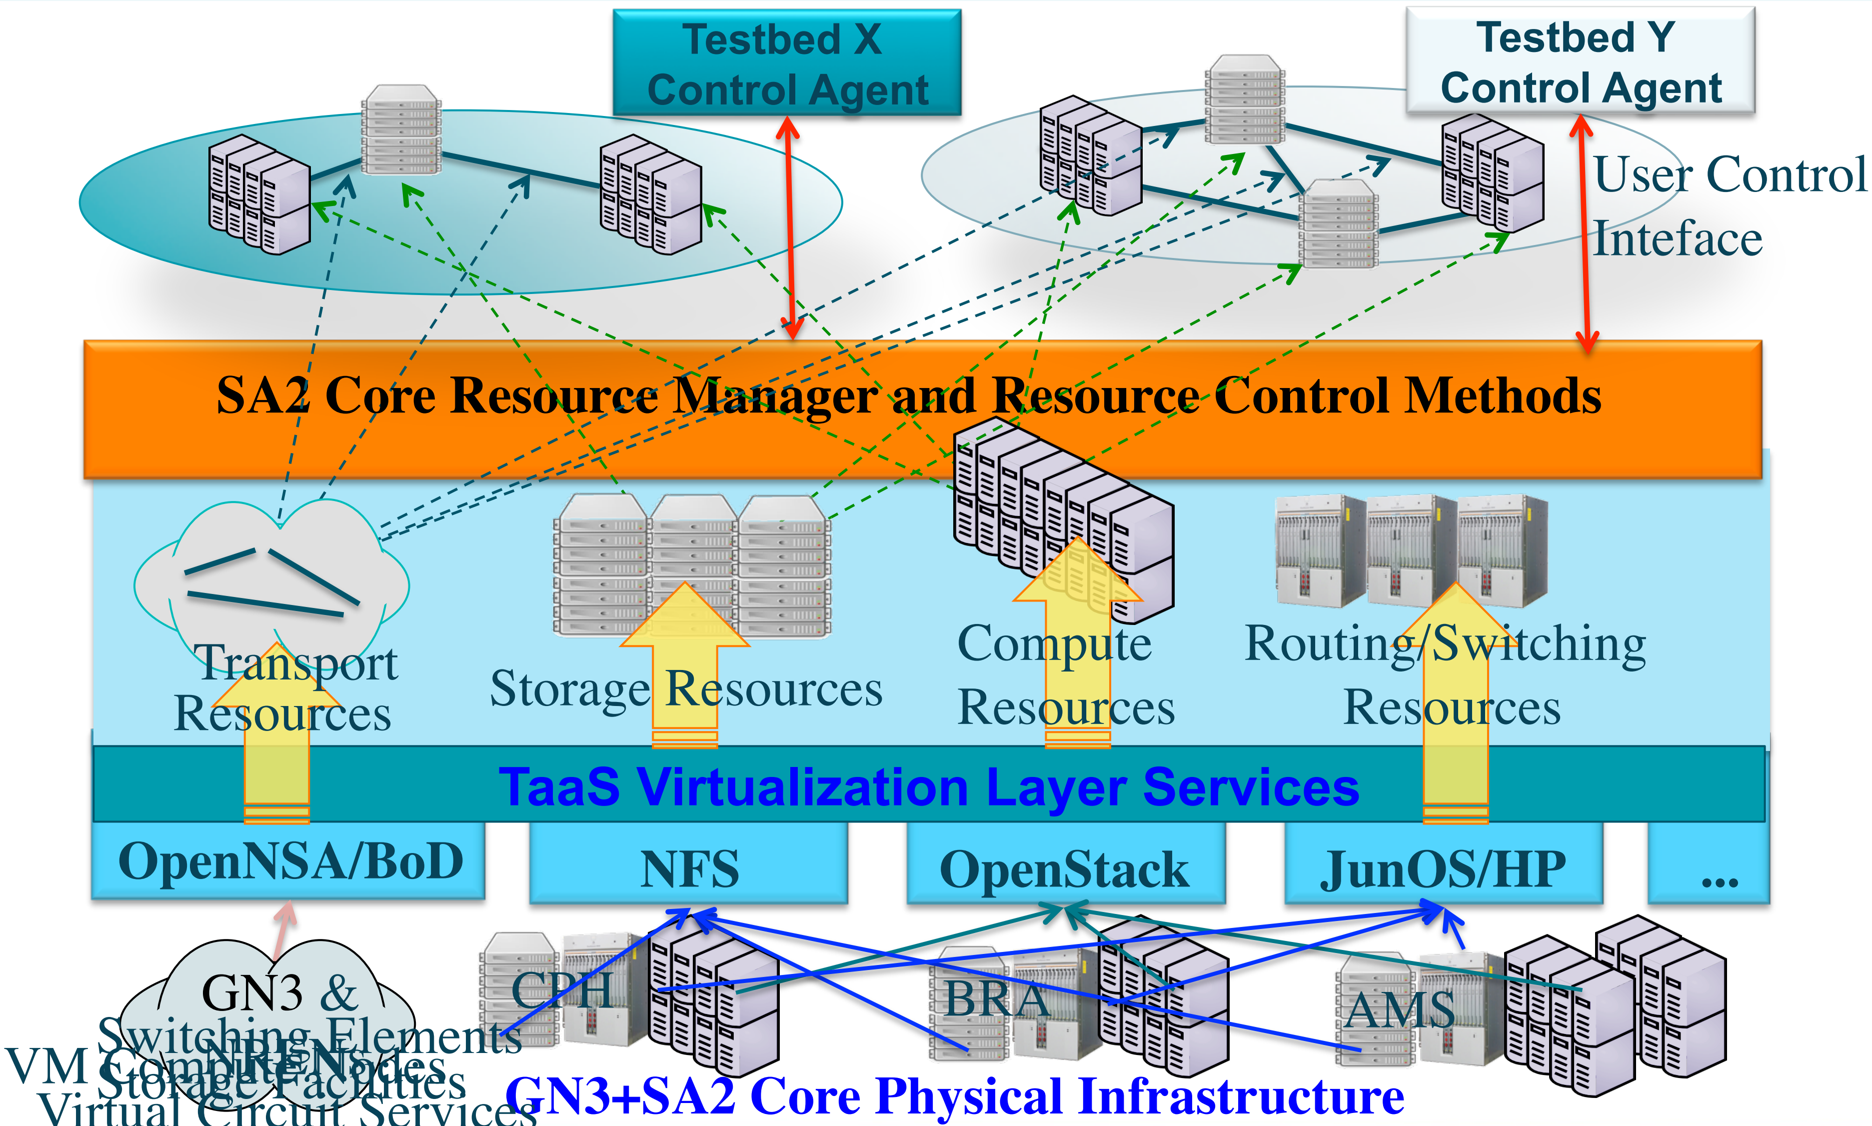

TaaS Deployment Plan
(as of Jan 2014)

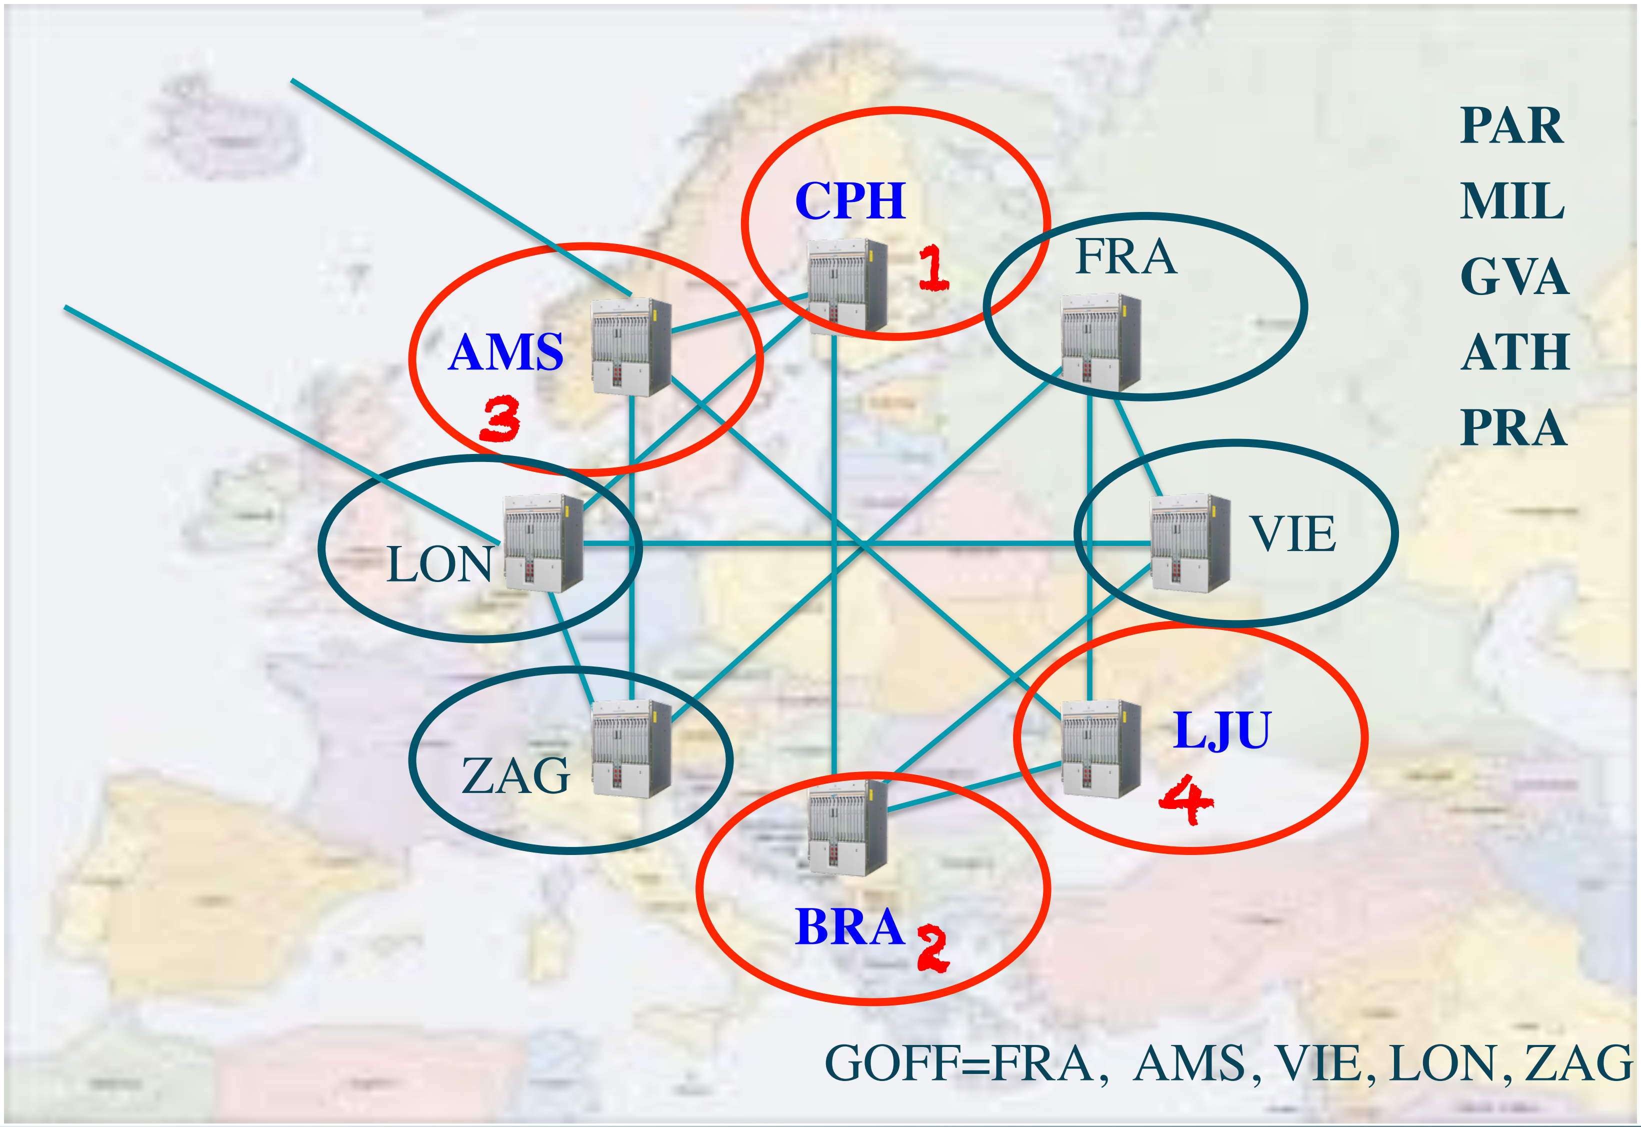

TaaS initial multi-domain interconnection concept

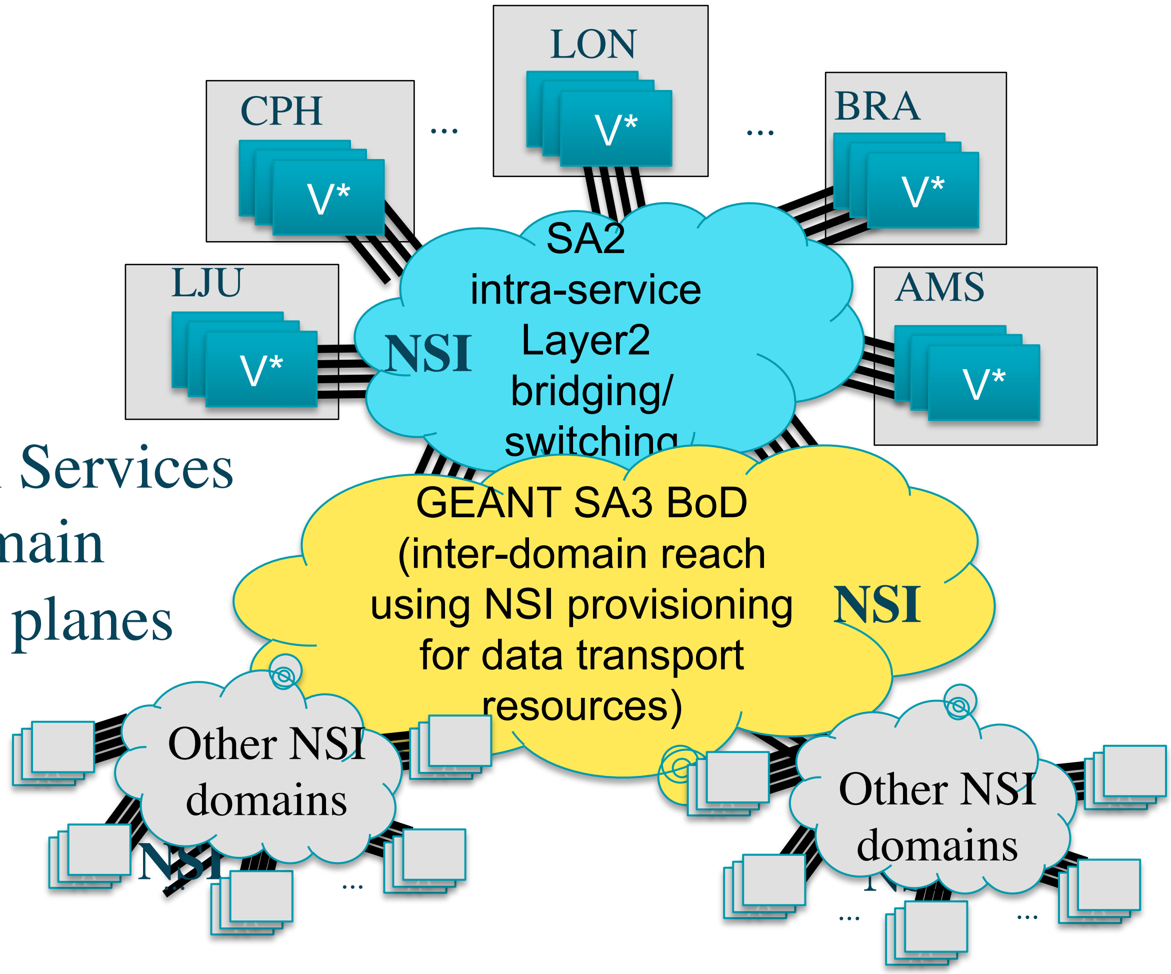

NSI Connection Services create multi-domain transparent data planes

The Automated GOLE International Testbed Fabric

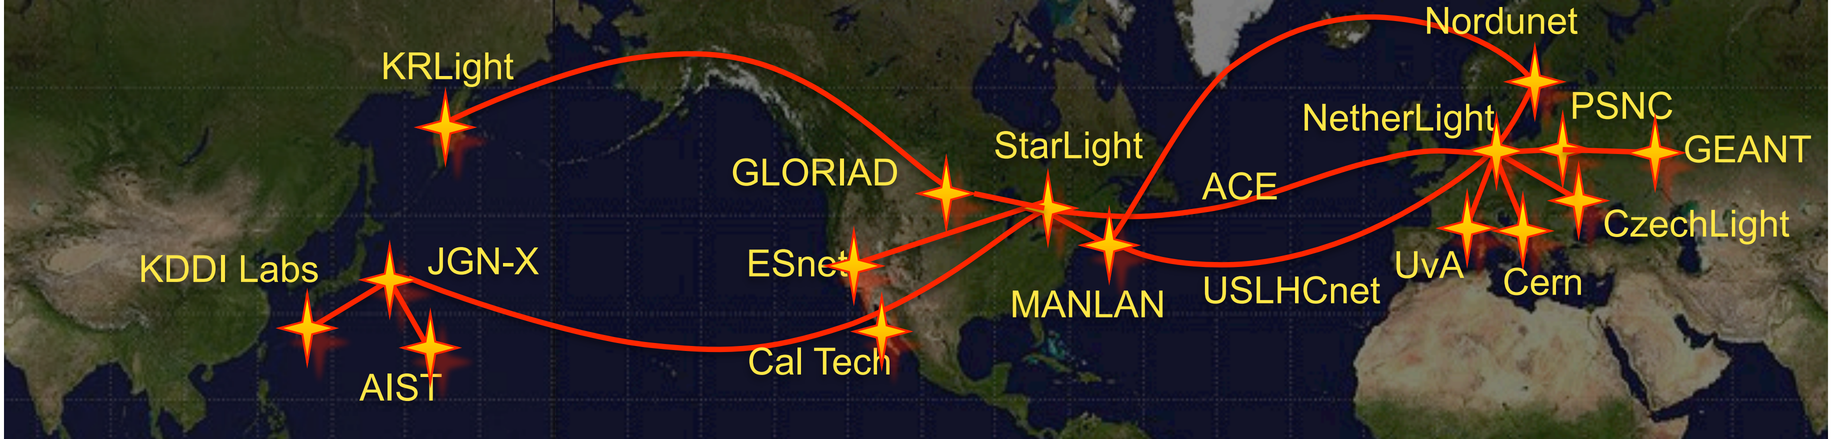

- The GLIF Automated GOLE Testbed Facility The GLIF Automated GOLE Pilot was initiated in 2010 to provide a global fabric of GLIF Open Lightpath Exchanges (GOLEs) for the express purpose of maturing the dynamic provisioning software (NSI), demonstrating the value of GOLEs to emerging network service models, and to develop a set of BCP for these services.

TaaS Multi-domain Testbeds (Phase 2)

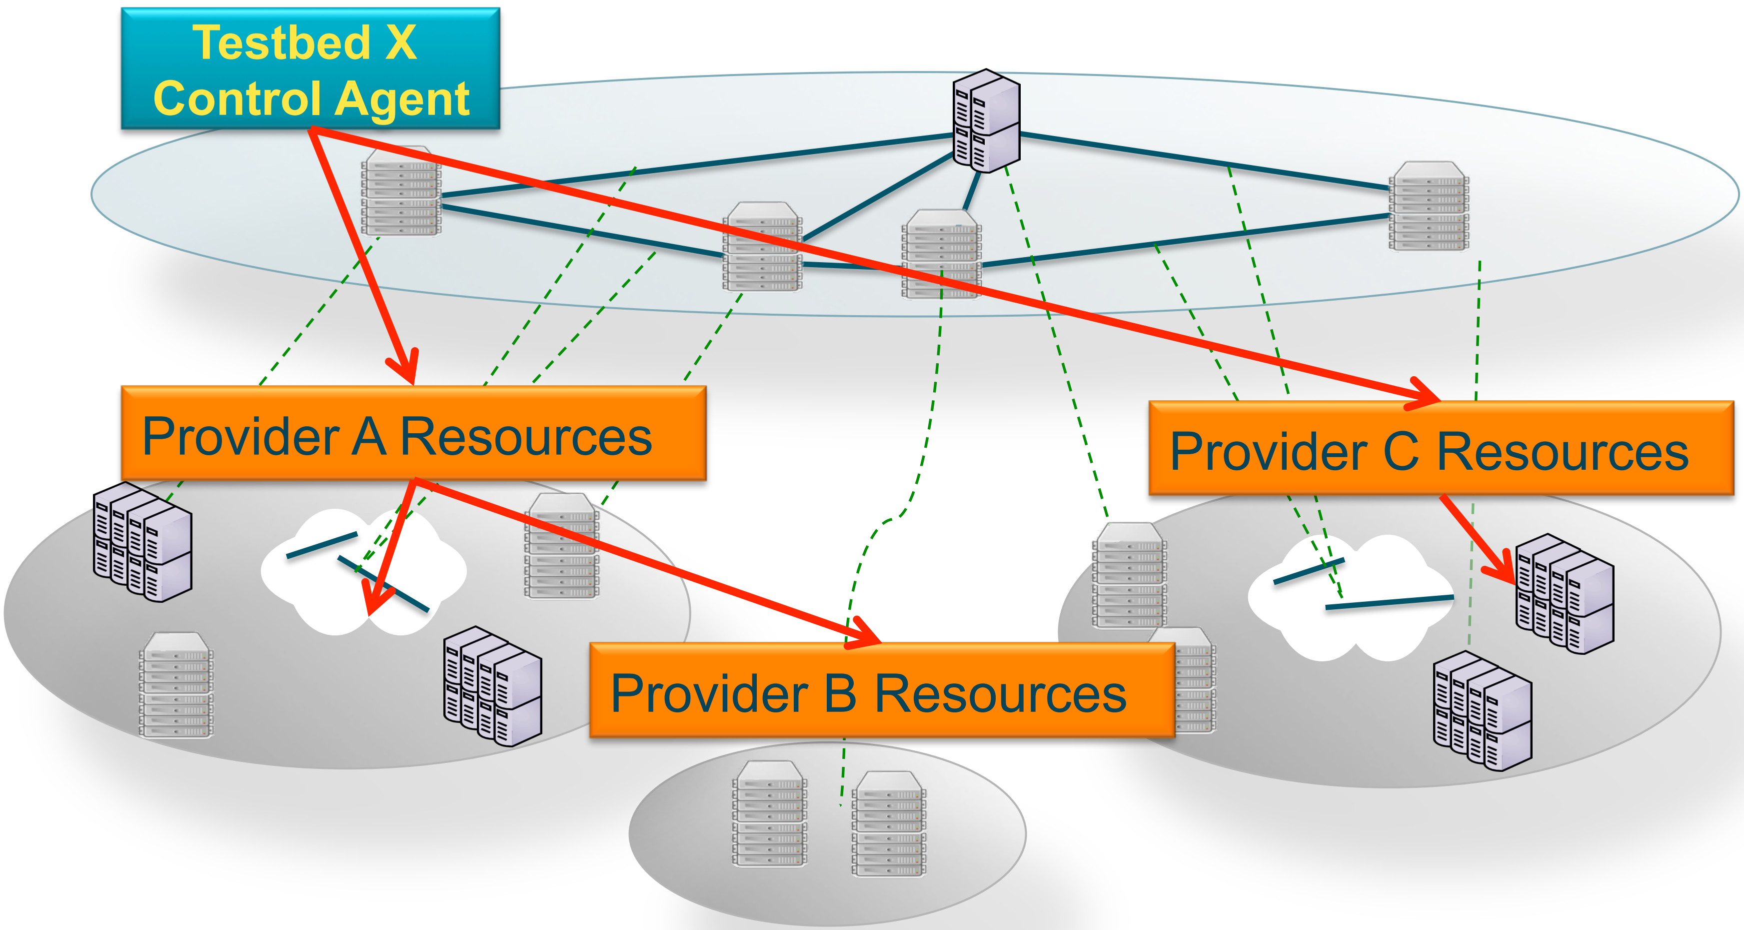

Globally interoperating virtualized services domains establish a globally distributed user controlled [SDN] VNE domain...

NSI inter-Domain Transit Services

GÉANT

NSI Enabled SDN Controller

NSI enabled opaque transit domains

No change to the internal SDN architecture or operation

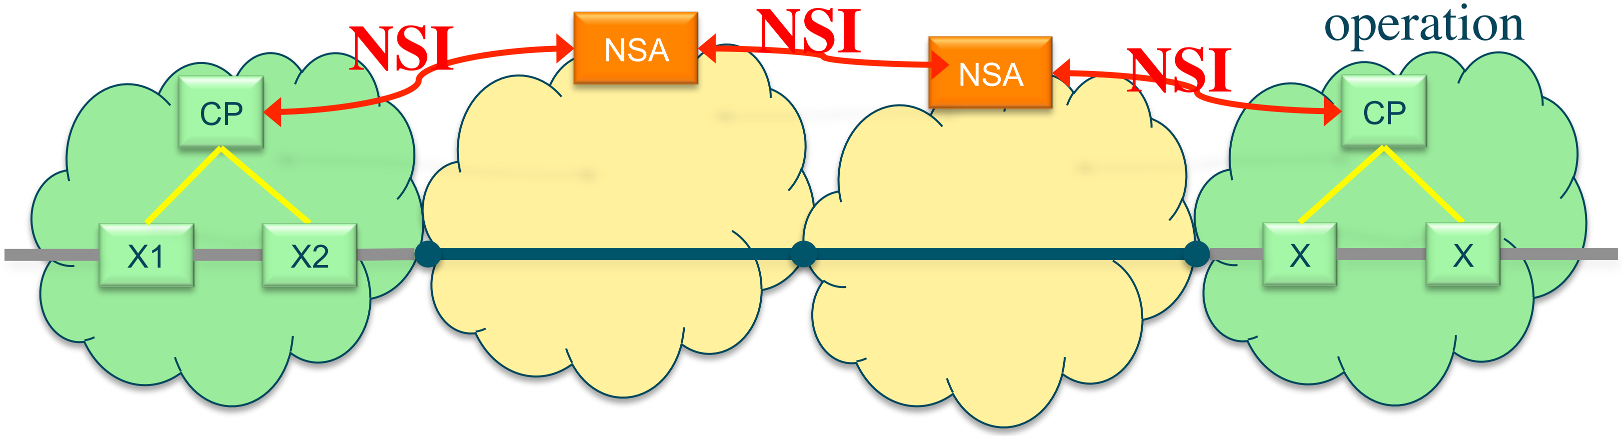

- The inter-domain transport model:
  - NSI-CS is used to establish transparent conduits between SDN islands.
  - These NSI Connections establish virtual SDN adjacencies over/thru intervening multi-layer transit networks
  - The NSI domains are transparent to the SDN application - they just provide the basic atomic network function of data transport.

NSI Across SDN Transit Domains

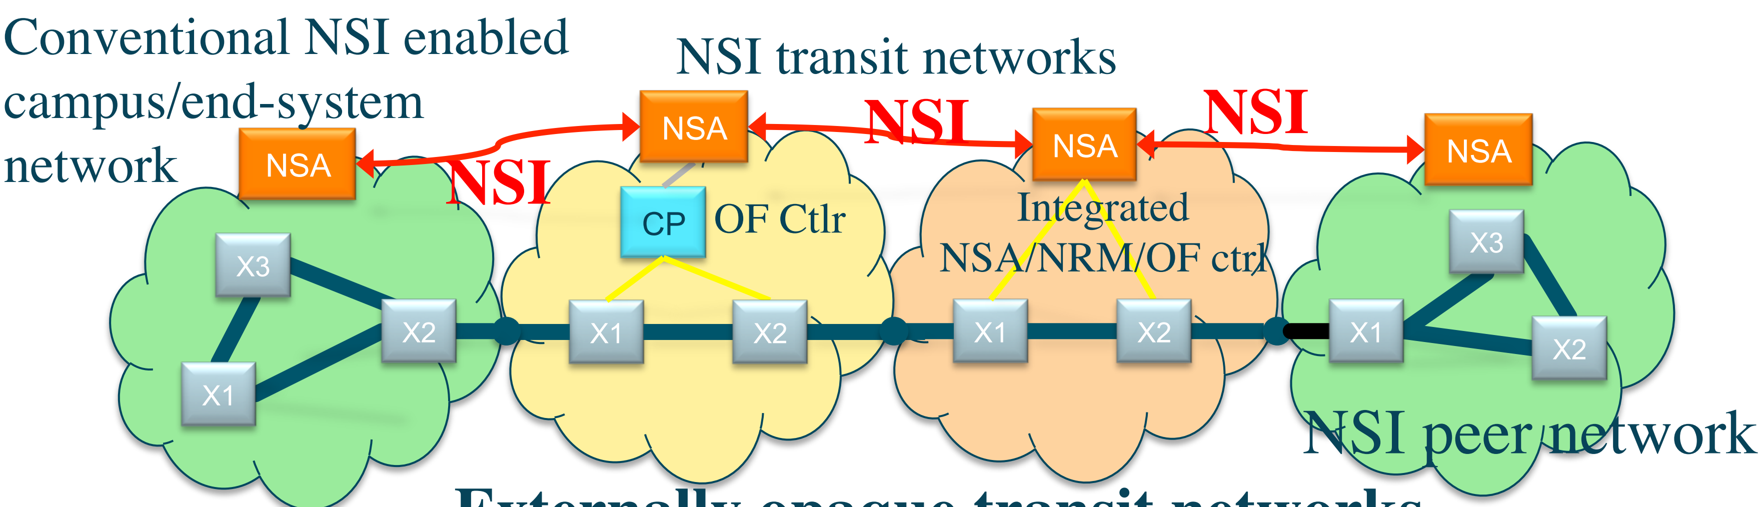

Externally opaque transit networks use SDN technologies internally

- The technology agnostic NSI-CS is used to provision Connection requests across an autonomous transit environment (that just happens to utilize SDN internally).
- This model allows NSI-CS to offer SDN flow spaces via the n-tuple STP mapping.
  - Enables “flow space defined” transport link connections.
  - It is consistent with NSI abstractions (domain boundaries, STPs, etc.)

NSI and Virtualized Network Environments

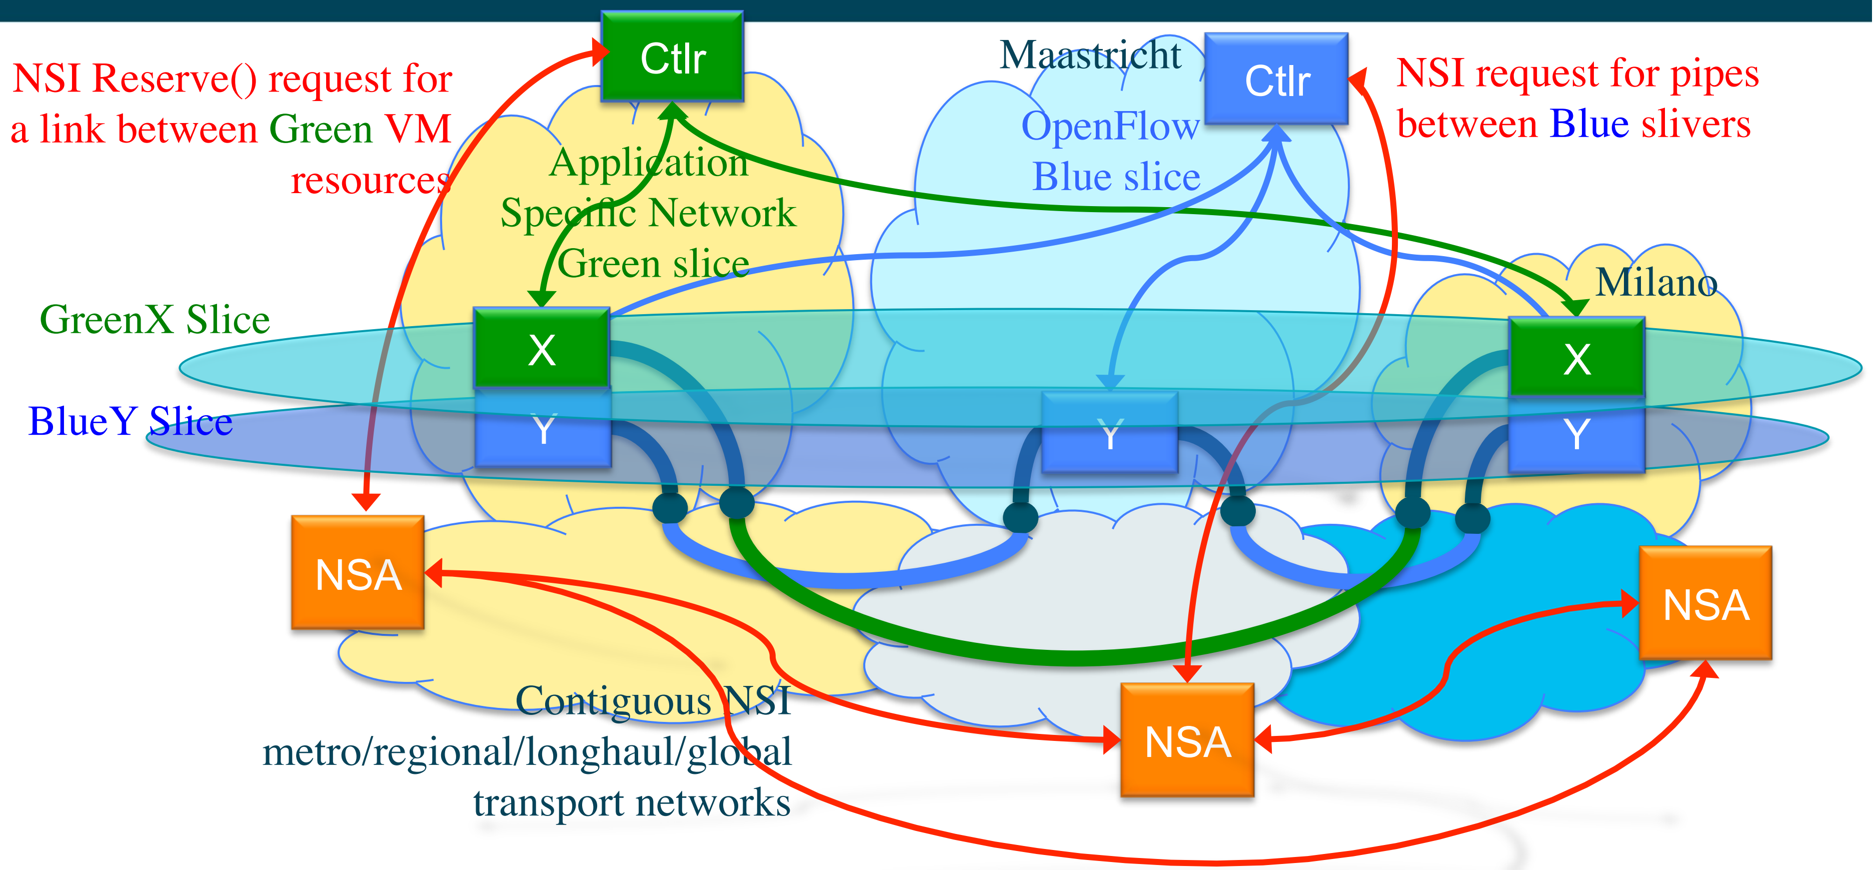

- Dynamic and collaborating Resource Managers construct global slices using NSI to provision transport links across the inter-domain global infrastructure between Cloud/VM/SDN resources...
- GN3+ SA2 “Testbeds as a Service” is deploying use this model.

Inter-Domain & Multi-Domain

- Testbeds (virtual network environments/slices) must extend/scale globally - yet preserve security, insulation, control, privacy, etc
- GEANT SA2-T4 is exploring a strategy
  - EU deployment within NRENs and Campus service implementations
  - Inter-domain interoperation with similar projects in other international regions
    - *US/NA (Internet2 AL2S, GENI)*
    - *SA*
    - *APAC*
  - Common service model, common inter-domain architecture, consensus protocol(s), ...

Questions?