ESnet

ENERGY SCIENCES NETWORK

LHCONE P2P Discussion

John MacAuley, ESnet

Lawrence Berkeley National Laboratory

LHCOPN/LHCONE Meeting

Ann Arbor, MI, USA

September 15-16, 2014

OSCARS Update

OSCARS Update

- OSCARS provides Ultimate Provider (uPA) functionality exposing ESnet network resources to the NSI community (nsi-bridge).
- Fully supports the NSI CS v2 provider specification (revision 117).
- Supports mutually authenticated TLS with certificate DN access control as per NSI security profile.
- New topology generation tool takes existing OSCARS NM topology, converts, and publishes into NSI Document Distribution Service (NSI DDS).
- OSCARS team provides other interoperable NSI tools such as reusable NSI protocol libraries, and a CLI for testing or scripting use.

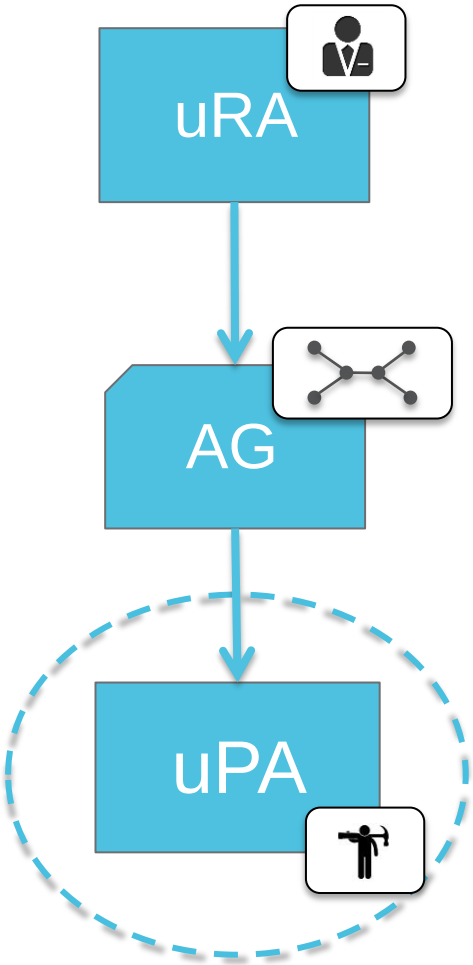

ESnet

Aggregator Update

- The “nsi-safnari” aggregator is jointly developed by ESnet and SURFnet.
- It is a 100% pure aggregator NSA with no associated uPA capabilities
  - Provides networks with their own uPA an aggregator option.
  - Allows deployment of project specific aggregators.
- Currently deployed version of aggregator supports:
  - NSI CS v2 provider and requester specification (revision 117).
  - Mutually authenticated TLS with certificate DN access control as per NSI security profile.
  - Both chain and tree signaling options.
  - Draft NSI Document Distribution Service (DDS).
- Deployed in three sites with more on the way!

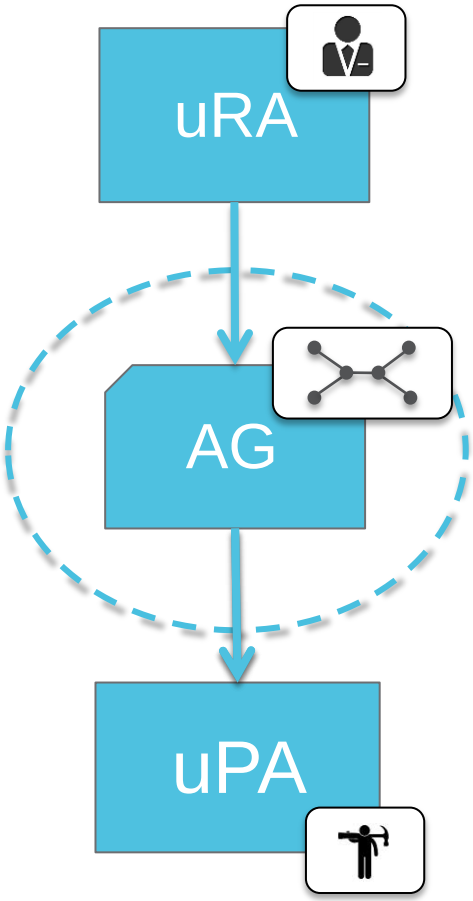

ESnet

Reference Deployment

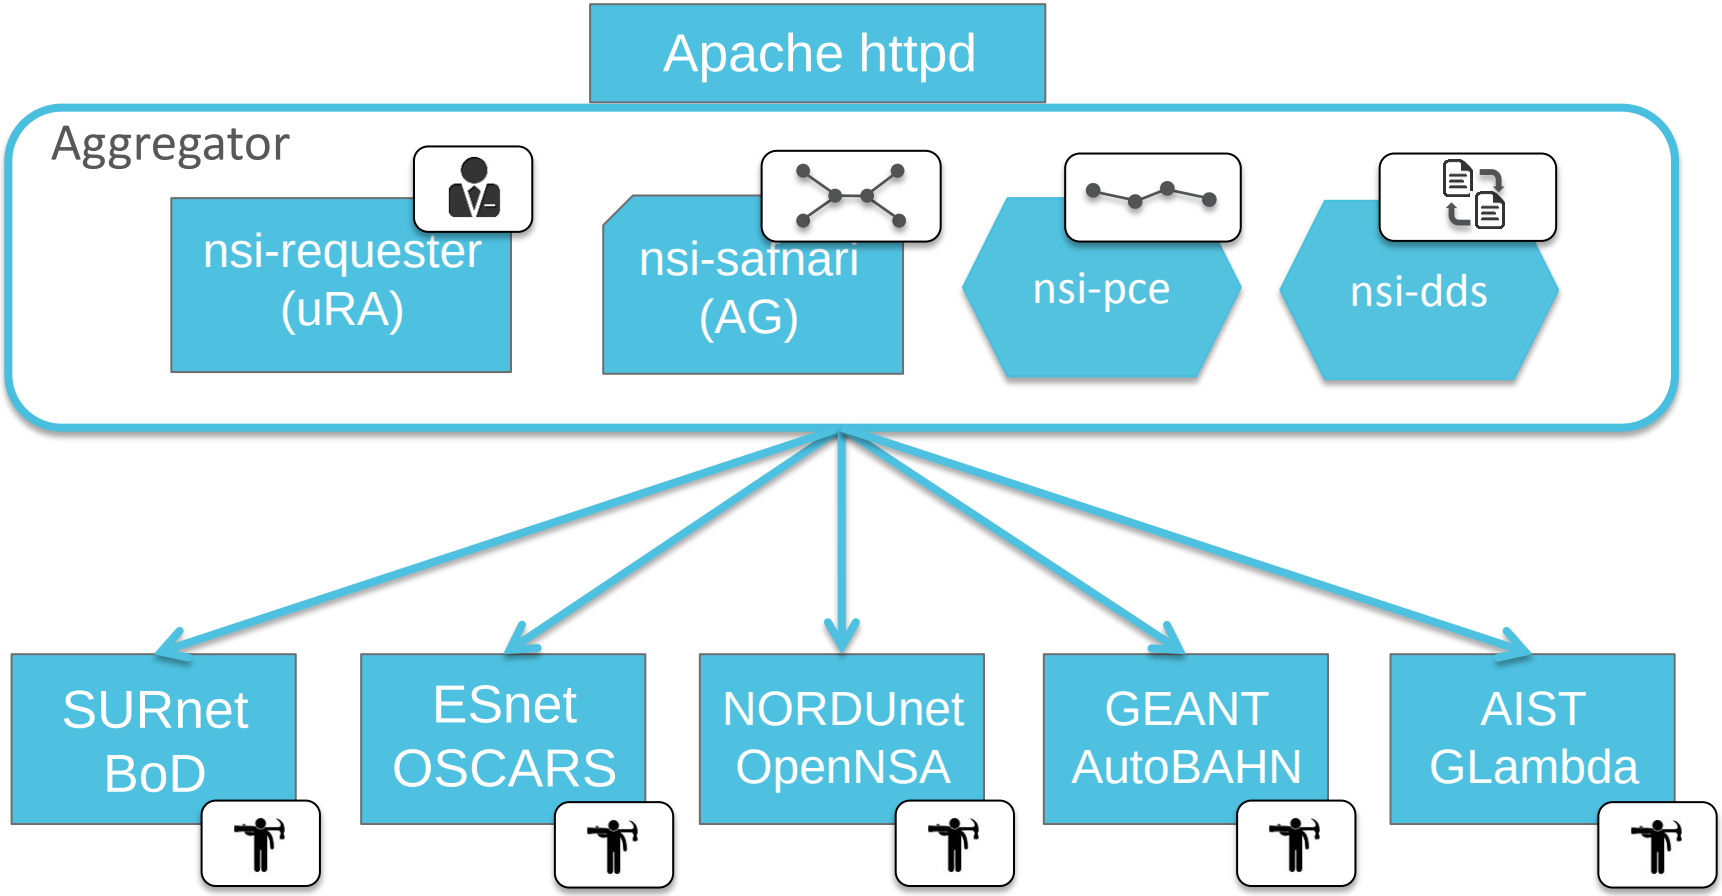

Software available @ https://github.com/BandwidthOnDemand

ESnet

Building an AutomatedGOLE backbone

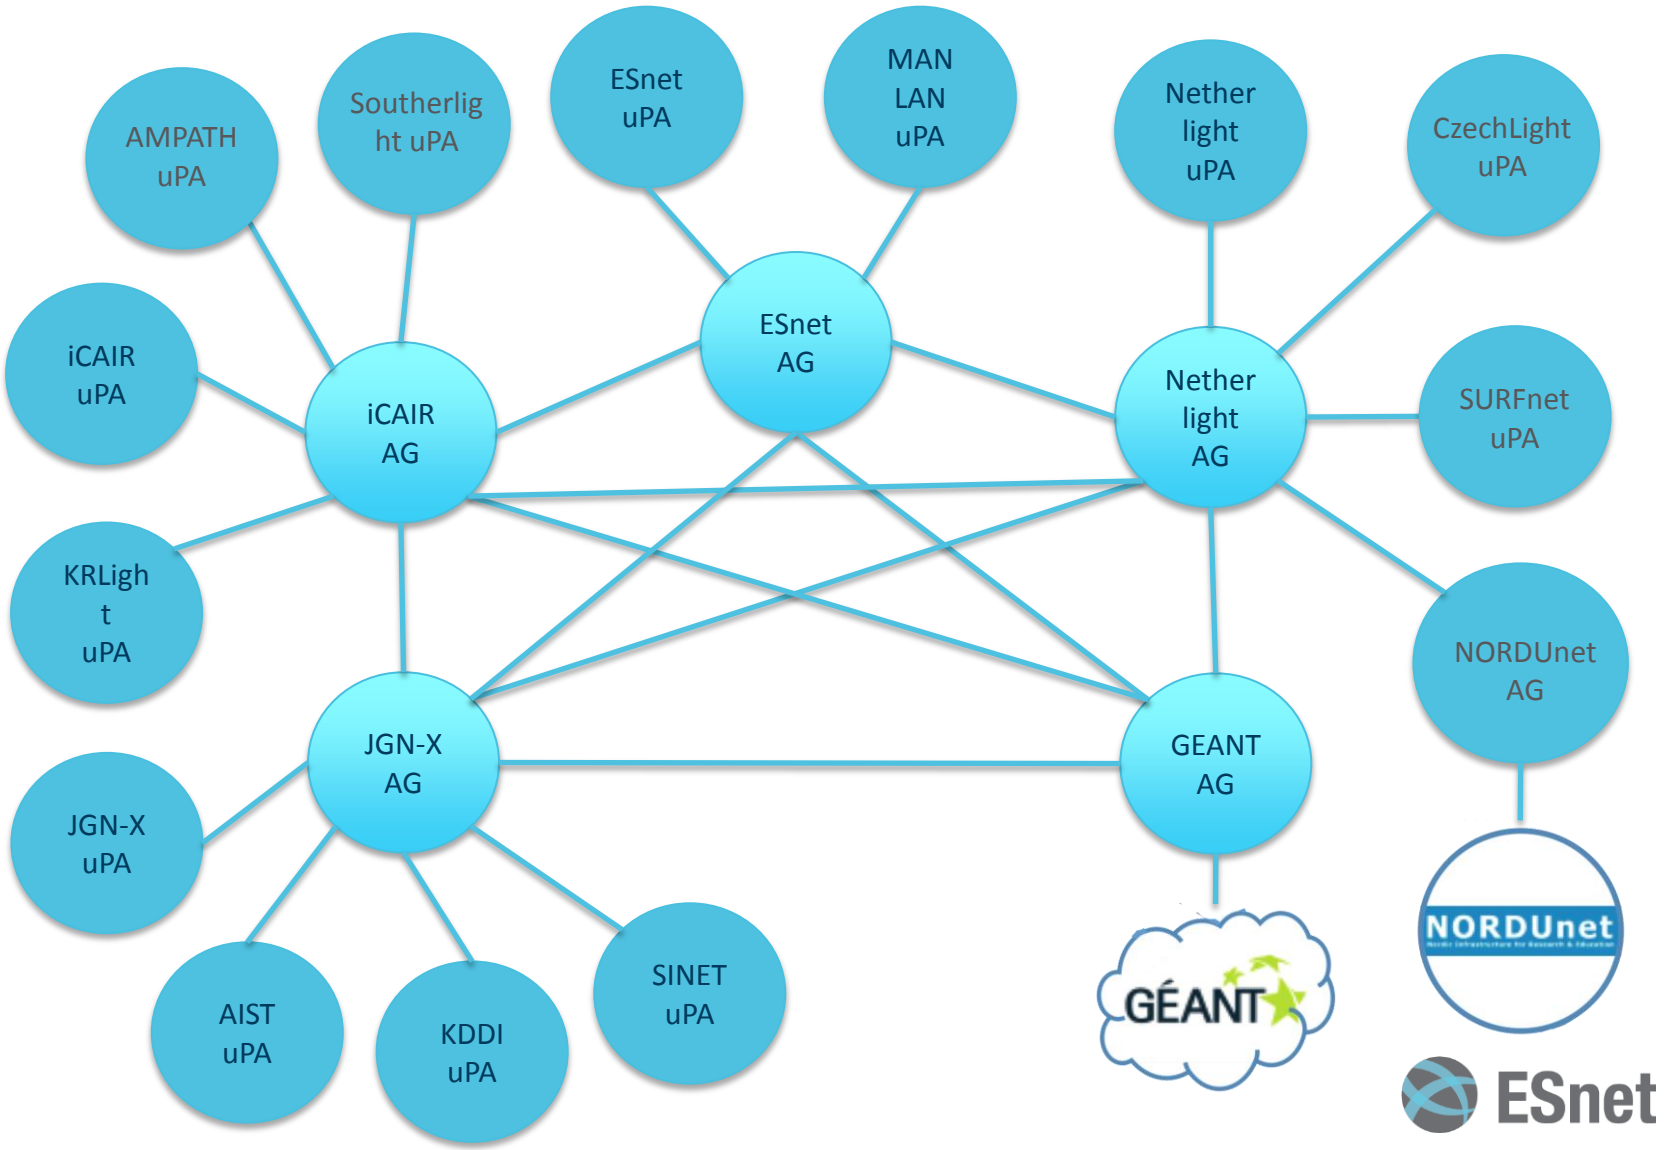

ESnet's NSI Task Overview

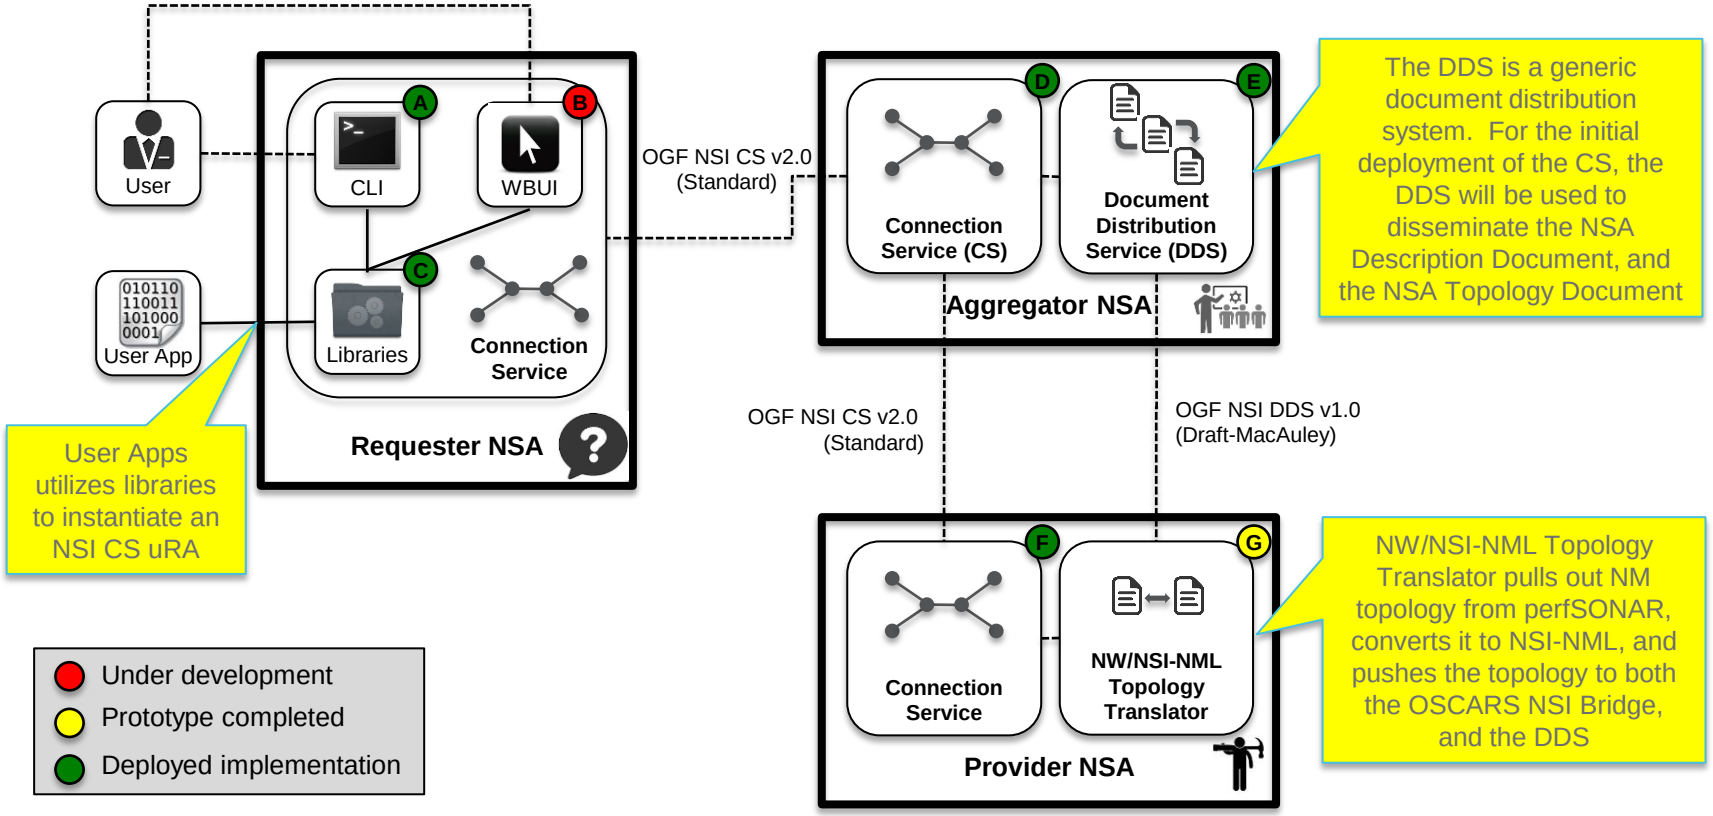

A. (uRA) CLI component is part of NSI CLI code base (Available*: https://oscars.es.net/repos/trunk/oscars/nsi-cli) [ESnet:Vangelis]
B. (uRA) WBUI is proposed to be MEICAN (Targeted 4Q2014) [RNP:Lisandro, ESnet:Vangelis]
C. (uRA) NSI CS Client Libraries is part of NSI CLI code base (Available*: https://oscars.es.net/repos/oscars/trunk/nsi-client) [ESnet:Vangelis]
D. (AG) NSI CS Aggregator is SAFNARI (Available: https://github.com/BandwidthOnDemand) [SURFnet:Hans, ESnet:John]
E. (AG) NSI DDS (Available: https://github.com/BandwidthOnDemand/nsi-dds) [ESnet:John]
F. (uPA) NSI CS uPA is OSCARS v0.6 running NSI Bridge (Available: https://code.google.com/p/oscars-idc/) [ESnet:Andy, Vangelis]
G. (uPA) NW/NSI-NML Topology Translator (Target 3Q2014) [ESnet:Sowmya]

*Requires https://oscars.es.net/repos/oscars/trunk/nsi-soap

NSI Update

NSI Status

- NSI Connection Service v2.0 has standard status as OGF document GFD-R-P.212 (May 13, 2014).
- Group is actively building 2.0 errata based on developer feedback.
- GLIF NSI Implementation Taskforce has been addressing deployment issues and providing feedback and new contributions into the NSI Working Group for standardization.
- NSI-WG meets September 22nd to discuss details surrounding:
  - Policies impacting path finding functions.
  - Topology service proposals to address both chain and tree-based solutions.

ESnet

NSI Values

- Complete decoupling of signaling plane from the data plane.
- Deployment of Network Service Agents with no network associations.
- Ability to perform centralized path finding with a complete view of the inter-domain topology.
- Facilitate advanced network resource workflows for network aware applications.
- Support for both tree and chain based signaling as a deployment option.

ESnet

Chain-based signaling model

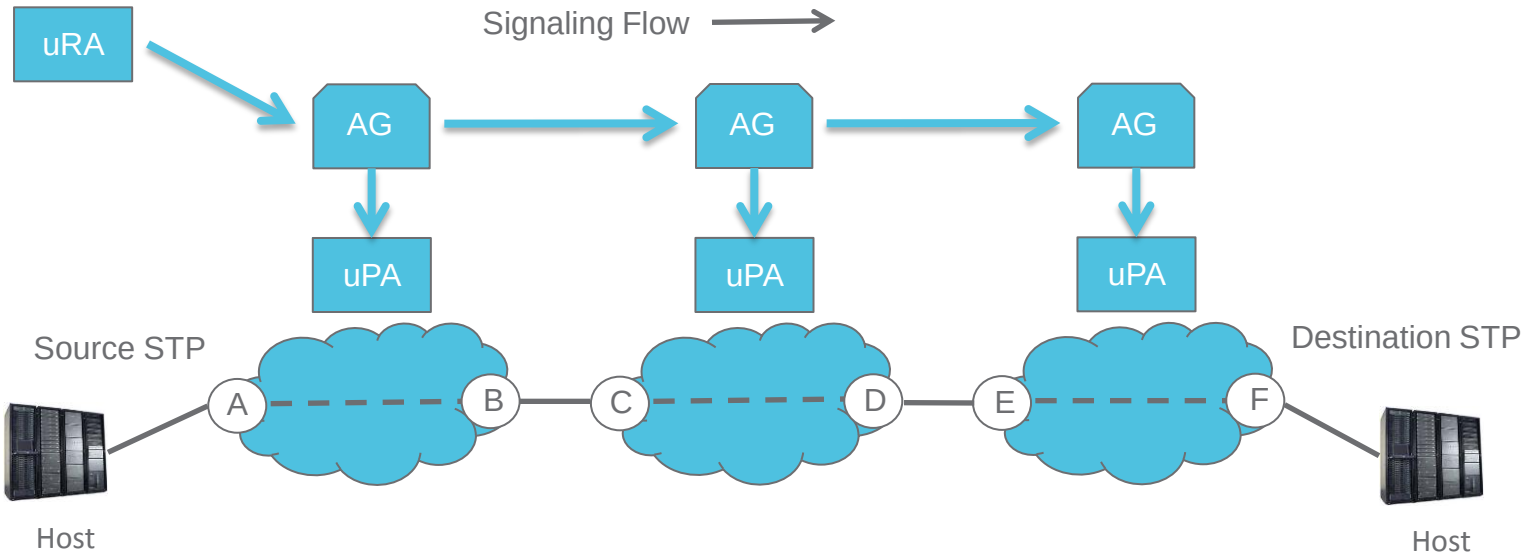

Every NSA associated with network resources must be an Aggregator capable of propagating a reservation request to the local uPA component and at most one adjacent (child) NSA associated with the next connection segment in the data path.

ESnet

Enhanced chain-based signaling model

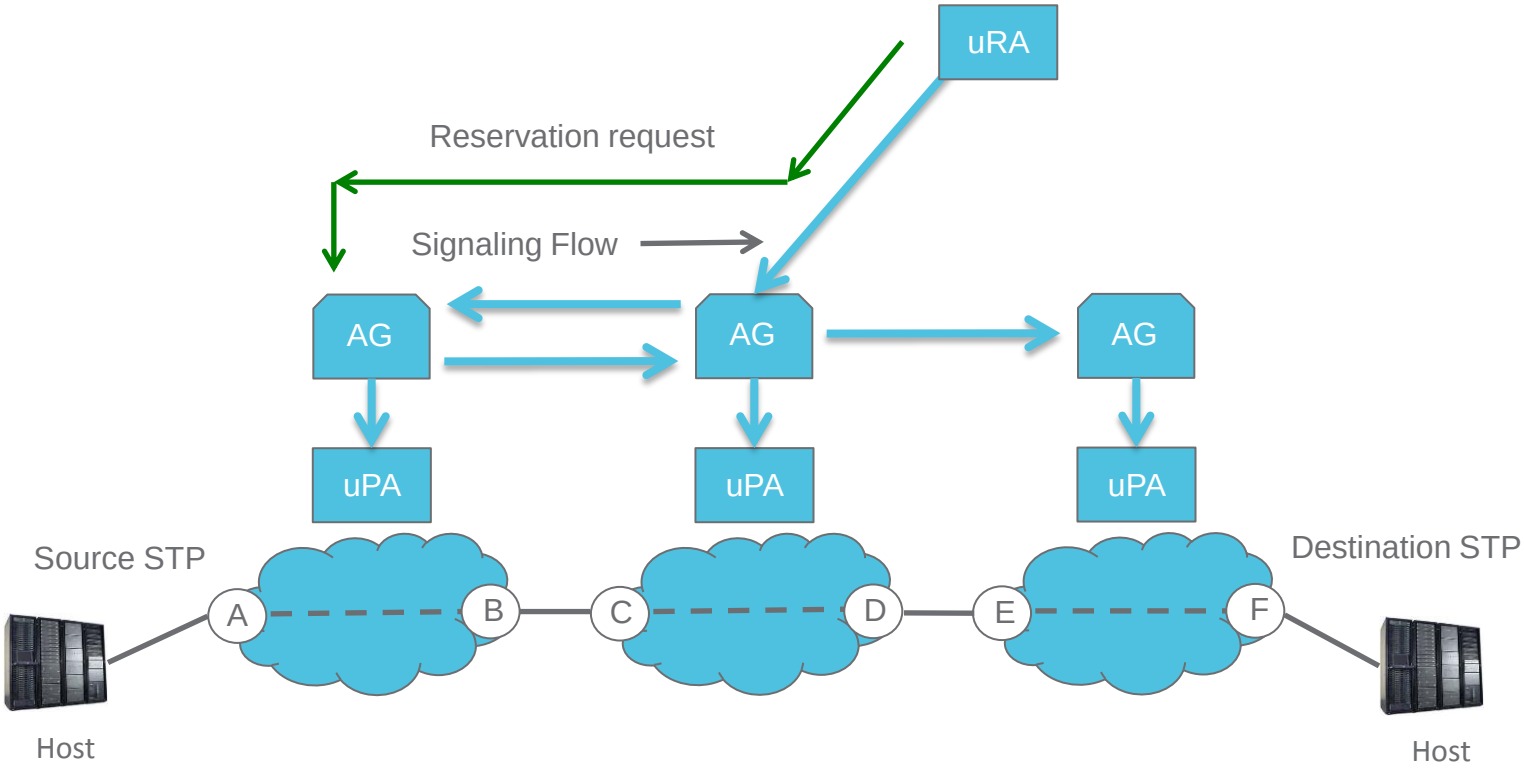

Control plane will “forward” a reservation request to the head-end node before beginning reservation of data plane segments..

ESnet

Tree-based signaling model

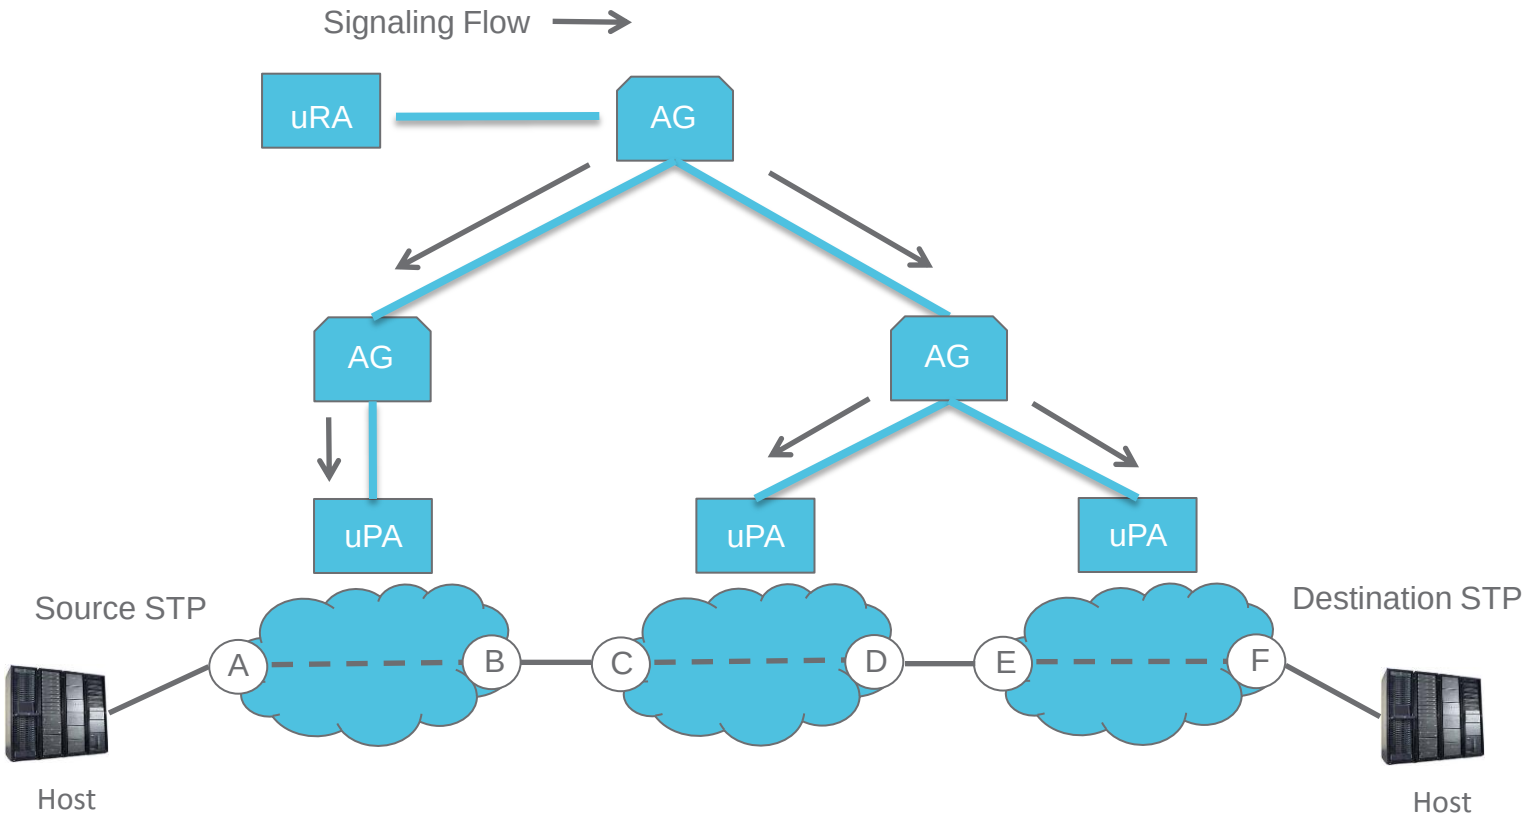

An Aggregator involved in a connection reservation does not have to be associated with any network resources involved in creation of that service. A uRA can issue a service request to an Aggregator NSA anywhere in the network if authorized to do so, and the NSI CS protocol with handle creating the reservation.

ESnet

NSI implementation testing

ESnet

Status

- Base integration testing of NSI CSv2 protocol (r117) and client authenticated TLS is completed for implementations in North America, Europe, and Scandinavia.
- Other implementations are coming on-line daily
  - AMpath, SouthernLight, and RNP are being tested now.
  - Waiting for implementations in Asia to be made available for testing.
- What next?
  - Designers are busy fixing bugs in implementations, specifically around a clean and consistent connection lifecycle.
  - Consistent and useful error feedback for reservation failures.
  - Path finding feedback to guide reservation retries (using error feedback).
  - Individual implementation and control plane stress testing.

ESnet

AutomatedGOLE dynamic circuit addressing

Test hosts provisioned on allocated VLAN an IP address ranges for connectivity testing using ping and iperf.

AutoGOLE VLAN and IP addressing

Please see this list for L2 and L3 addressing for the AutoGOLE environment.

Layer 2: VLANs 1779-1799 will be used for the Automated GOLE on all interconnecting links.

Layer 3: The AutoGOLE uses prefix 10.250.xx.yy/24 where xx are the AutoGOLE VLAN numbers 79-97 (abbreviation of full AutoGOLE VLAN range: 1779-1799) and yy is the GOLE/Network Identifier, this is assigned as below:

	IP prefix	Domain	Hostname	STPs
GOLEs & Networks	10.250.xx.1	NetherLight		
	10.250.xx.2	StarLight		
	10.250.xx.3	MAN LAN		
	10.250.xx.4	CzechLight		
	10.250.xx.5	NORDUnet		
	10.250.xx.6	CERN		
	10.250.xx.7	UvA		
	10.250.xx.8	PSNC		
	10.250.xx.9	JGN2		

ESnet

Control plane testing

Control Plane: Testing implementation of two-way TLS and NSI-CSv2.0R117

		TO							
		AutoBAHN	*DynamicKL*	*G-lambda/A (PA)*	*G-lambda/K*	*OpenNSA*	*OSCARS*	*SURFnet BoD*	nsi-safnari
FROM	*AutoBAHN*	X							
	DynamicKL		X				X		
	G-lambda/A			X			X		
	G-lambda/K				X		X		
	OpenNSA					X	X		
	OSCARS		X	X	X	X	X	X	
	SURFnet BoD						X		
	nsi-safnari								
	nsi-requester								
	G-lambda/A (Aggr)								

	NSI-CSv2-R117 & Two-way TLS OK
	Two-way TLS OK
	NSI-CSv2-R117 OK
	Started testing
	Found errors/On-Hold
	Not started testing

ESnet

Data plane testing

Data Plane: Testing implementation of two-way TLS and NSI-CSv2.0R117

(!) Please see Result table below

FROM \ TO	AIST	AMPATH	CESNET	ESnet	GEANT	GRnet	Internet2	JGN-X	KDDI	KISTI	MANLAN	NetherLight	NORDUnet	PSNC	RNP	SINET	StarLight	SouthernLight	SURFnet	UvA	Funet
AIST	X																				
AMPATH		X																			
CESNET			X																		
ESnet				X																	
GEANT					X																
GRnet						X															
Internet2							X														
JGN-X								X													
KDDI									X												
KISTI										X											
MANLAN											X										
NetherLight												X									
NORDUnet													X								
PSNC														X							
RNP															X						
SINET																X					
StarLight																	X				
SouthernLight																		X			
SURFnet																			X		
UvA												1781 w	1781 worked						1781 Could	X	
Funet																					

	Data plane fully working (ping succesful)
	Not sure if data plane is working (control plane works, no ping yet)
	Started testing
	Found errors/On-Hold

ESnet

Authentication and Authorization

A Control Plane of Trust

- The control plane is built up through a series of NSA peering agreements (may or may not follow data plane peering).
- Client authenticated TLS is used to authenticate peer NSA, as well as ensure the integrity and confidentiality of the messages traveling through the control plane.
- Control plane security is based on transitive trust: I trust my neighbors and the neighbors they trust.

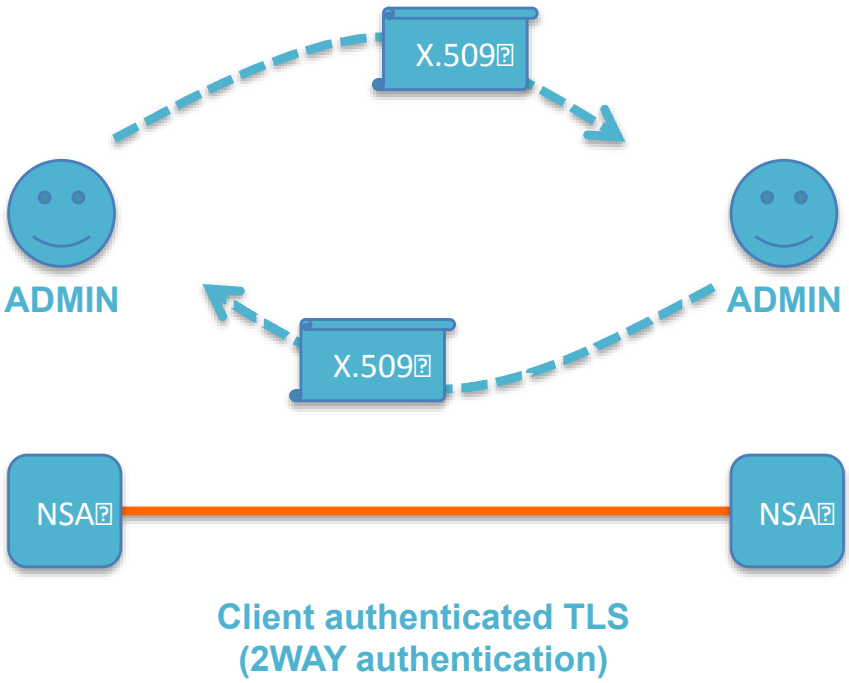

User access to control plane

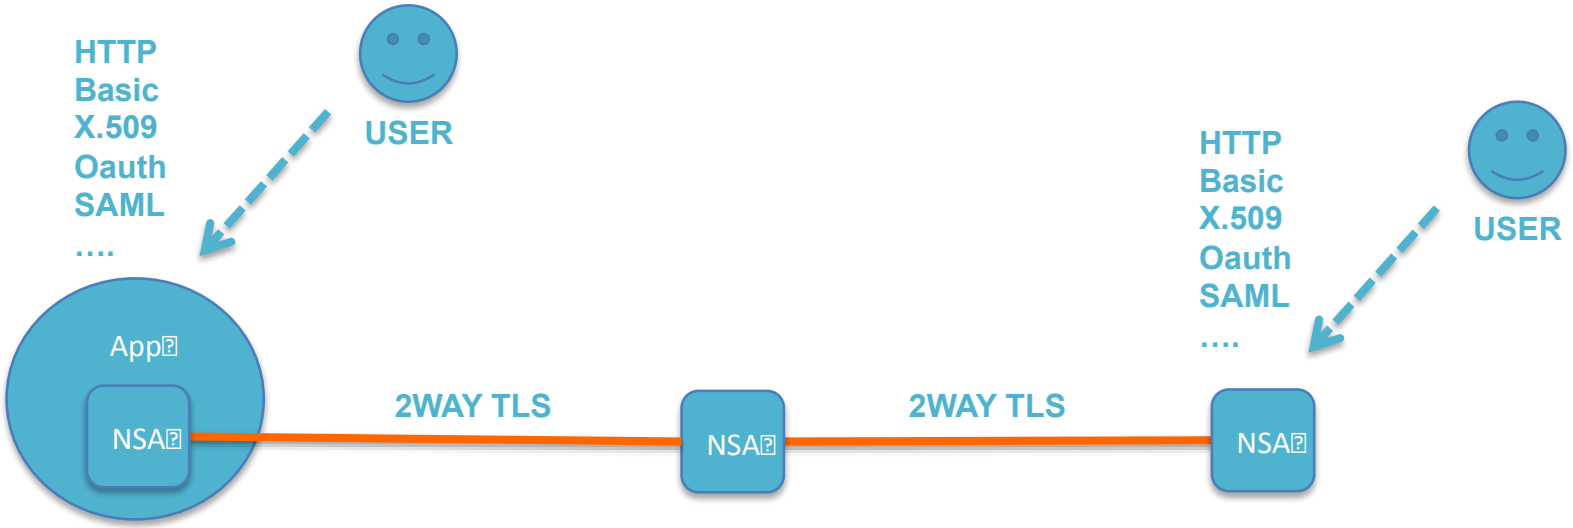

- uRA must authenticate originating user using locally defined authentication schemes.
- Identity information of the originating user can be added to the NSI message header and passed to peer NSA along with the reservation request.
- NSA along the reservation path can utilize identity information if needed for local policy enforcement.

Authorization policies

- Any NSA within the reservation path can do authorization on the request using local policies.
- Once in the reservation path, and NSA will receive all further messaging related to that reservation, applying additional policies as needed.
- Examples of authorization policies that could be made by an NSA include:
  - Transit policies - can the reservation transit the network?
  - Restricting STP access – can the reservation use the specified STP?
  - Bandwidth limits – is the reservation within allowable bandwidth limits.
  - Time-of-day access restrictions
  - Maximum number of reservations per day/week/year/…
  - Etc…

ESnet

Policy impacts on path finding

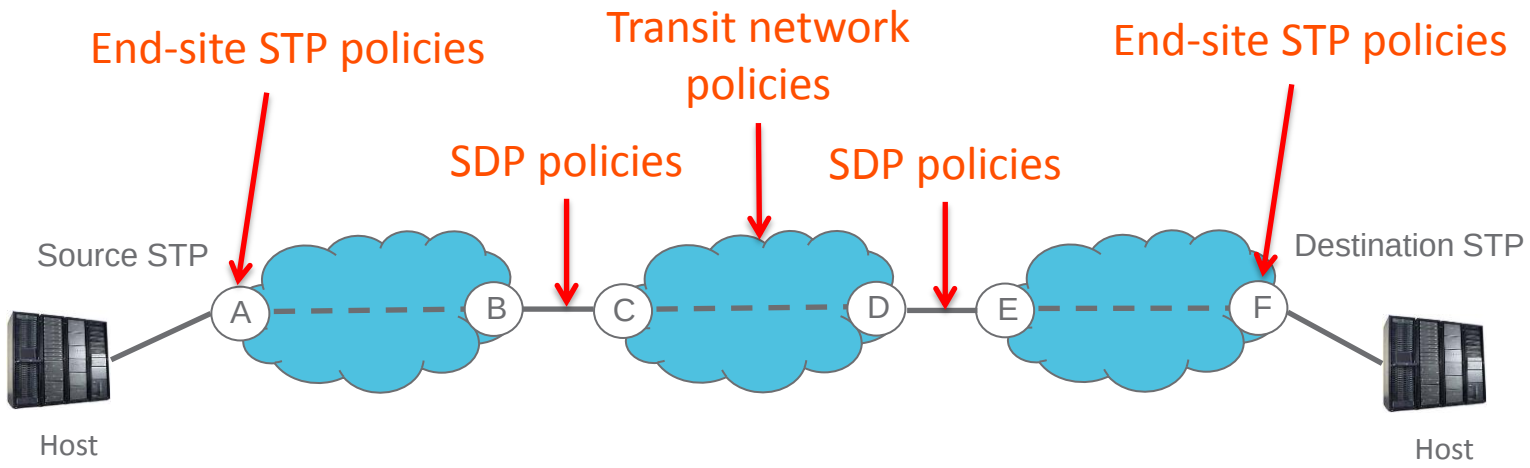

- Each uPA is ultimately responsible for enforcing any local policies on the reservation independent of any view of policy used by the path finder when choosing the path.
- End-site STP policies have minimal impact on path finding
  - An end-user is usually aware of any policies enforced on ports connected into their sites, and can provide the needed credentials.
- SDP and transit policies have a much greater impact as they are invisible to the end-user.

ESnet

Discussion

- Are their any LHCONE project specific policies?

ESnet