

Network Service Interface Signaling and Path Finding

Status of This Document

This document provides information to the Grid community on signaling and pathfinding in the NSI protocol. Grid Working Document – Informational (GWD-I). This document obsoletes GFD.217.

Copyright Notice

Copyright © Open Grid Forum (2012-2017). Some Rights Reserved. Distribution is unlimited.

Abstract

This document provides a discussion of the signaling and path finding models supported by the NSI protocol.

Contents

Abstract	1
Contents	1
1 Introduction	1
2 Signaling and Path Finding Models.....	2
2.1 Tree-Based Signaling and Path Finding Model	2
2.1.1 Tree-Based Hop-by-Hop Routing	4
2.1.2 Tree-Based Source Routing	5
2.2 Chain-Based Signaling and Path Finding Model	6
2.2.1 Chain-Based Hop-by-Hop Routing	7
2.2.2 Chain-Based Source Routing	8
3 Conclusion	9
4 Security Considerations	9
5 Contributors.....	9
6 Glossary	9
7 Intellectual Property Statement.....	11
8 Disclaimer	11
9 Full Copyright Notice.....	11
10 References.....	12

1 Introduction

The NSI Service Framework defines a Transport Plane and a Service Plane. The connectivity services are built using the resources available in the Transport Plane. NSI messages are exchanged on the Service Plane. The NSI messages follow a workflow that has been designed to be independent of the Transport Plane connectivity.

The NSI framework provides the flexibility to support NSI message workflows as simple as a chain or as complex as multi-level trees. The fundamentals of NSI that support this flexibility can be found in the OGF NSI Framework document [1]. To provide clarity on the various scenarios for signaling and path finding, the following sections outline specific workflows and their requirements.

2 Signaling and Path Finding Models

The model for signaling and path finding has been driven by the following architectural requirements that were established early in the inception of the NSI working group. These requirements have been driven by the need to support a flexible Service Plane that is topologically decoupled from the Transport Plane. Recent developments in Software Defined Networking (SDN) have confirmed the importance of decoupling of the Service and Transport Planes.

1. The Network Service Agent (NSA) Service Plane topology does not need to be congruent with Transport Plane topology.
2. Not all NSAs peer with each other, i.e. it cannot be assumed that there is a full mesh of reachability between all NSAs. Pair-wise peering between NSAs dictate the Service Plane topology. NSA Service Plane inter-connectivity will be guided by security and administration considerations and NOT exclusively Transport Plane considerations. Therefore, a requester NSA may not have a direct Service Plane peering to the complete set of Provider Agents (PAs) involved in a reservation.
3. Users may request a reservation from an NSA (i.e. an Aggregator NSA (AG)) that is not directly managing resources in the Transport Plane. Based on item #2 above, this AG may not have direct Service Plane peering with all the NSAs involved in the reservation request.
4. Users may request reservations between endpoints that are not in their Network, or the Network of their NSA. This implies the user request may not originate from the NSA managing the source end of the Connection.

These all imply that a reservation request in a tree workflow may pass through many AGs on its way to instantiating Connection segments on the children uPAs. If the AG at the top of the request tree is to perform its duties, then it will need to understand the Service Plane topology. This implies a knowledge of all relevant AGs within the Network, all the relevant uPAs and of course, the applicable Transport Plane topology. Now for the most important bit: this top-level AG *can only get all this information through its direct peers.*

The following sections distinguish between the tree and chain workflows even though a chain workflow can be considered to be a degenerate case of a tree workflow. The reason for this is to highlight certain concessions that can be made when considering the chain workflow independently.

This following sections also distinguish between source-based and hop-by-hop workflows. In the source-based case, the first AG takes responsibility for calculating the full path for the whole Connection. In the hop-by-hop case, each AG can choose to perform only sufficient path computation to determine where to forward the request. By supporting both modes, the NSI CS allows the AGs to decide how much or little path finding they chose to perform. The hop-by-hop mode allows a light-weight PCE to be implemented that only knows where to forward path request. Alternatively, an AG may choose to implement a PCE that is aware of the whole service domain.

2.1 Tree-Based Signaling and Path Finding Model

To support the NSI tree-based workflow (see Figure 1), the following must be supported:

1. A Network Service request may be initiated at any AG in the service region, and this AG does not have to be associated with any Network resources involved in the creation of the service.
2. A Network Service request can be propagated to as many uPAs as there are sub-connection segments. The AG receiving the request does not need to fully resolve the path of the Connection, but can delegate the task further down the tree to other AGs.

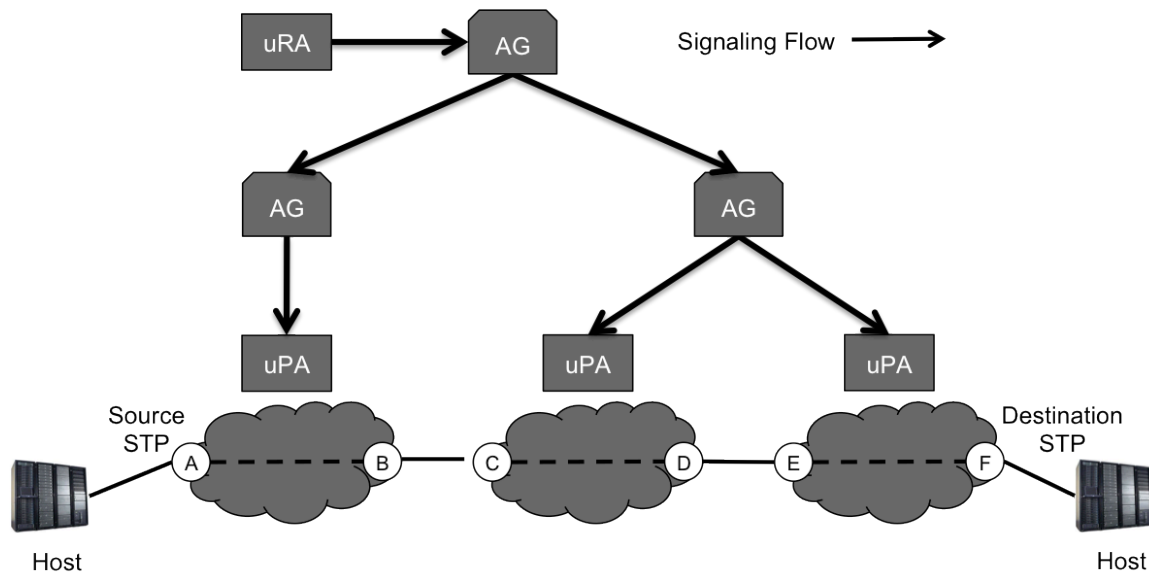


Figure 1. Tree-based signaling flow

In order to support the above two premises, the following requirements need to be met:

1. The Aggregator NSA needs a view of as much of the Network topology as is needed to perform advanced "intelligent" routing decisions. At a minimum, the Aggregator NSA needs to know about the relevant Networks, the Transport Plane peering ports (to derive SDP), the services offered, associated service domains (Switching Services), and any adaptations within the Networks.
2. An AG may hide Network details of child Networks if desired, and advertise a summarized NSI Topology as needed. The key with this model is that external NSAs need not know the details of these internal child NSAs.
3. An AG is restricted to communicating with only direct peer NSAs based on administrative policies, however, *it needs access to **all** relevant NSA Description Documents and NSI Topology documents to perform tree based routing.* In addition to peer communications, the NSA Description Documents and NSI Topology documents are needed to determine the NSA managing Networks and the Service Plane topology to determine routing paths of requests in the connected graph. The NSI Topology Documents from each Network are then used to build NSI Topology for the service region for routing of the Connection. The security, verification, and distribution of the NSA Description Documents and NSI Topology documents are outside the scope of this document and are addressed in separate NSI working group documents [3].

The specific requirement that an NSA would only directly peer with a subset of NSAs, and not be able to directly communicate with all NSAs is due to administrative limitations and scaling. Provisioning a full mesh of trust relationships between all NSAs is considered prohibitive, and anonymous retrieval of the NSI documents have negative security implications. This results in the need for Service Plane topology so that trust relationships between NSAs can be discovered, and NSI description information can be distributed to all NSAs within the service region.

2.1.1 Tree-Based Hop-by-Hop Routing

In a tree-based hop-by-hop routing workflow, the (root) AG receiving the request from the uRA may perform “partial” path finding and segment the initial request into smaller components that it can pass on to other NSAs to resolve. Figure 2 shows an example of such a workflow with the associated Path Computation Engines (PCEs):




Figure 2. Example of a tree-based source routing workflow

1. The uRA makes a reservation request to the root Aggregator AG_0 on behalf of the user to Connect STP A to STP F.
2. AG_0 receives the request (A,F), and after validation, performs path finding to determine a viable path for the requested end-points A and F.
3. PCE_0 uses NSI Topology over the whole service region to determine a viable Transport Plane path for the request, returning two Connection segments (A,B)(C,F).
4. AG_0 then sends the Connection request for segment (A,B) to AG_1 which represents the only signaling path to uPA_1 .
5. AG_1 receives the request (A,B), and after validation, performs path finding to determine which local STP should be involved in the Connection reservation.
6. PCE_1 determines that both the source and destination STP are managed by uPA_1 and returns the Connection segment (A,B).
7. AG_1 then makes a reservation request to the uPA_1 for the local Connection segment (A,B).
8. AG_0 then sends the Connection request for segment (C,F) to AG_2 which represents the signaling path to uPA_2 and uPA_3 .
9. AG_2 receives the request (C,F), and after validation, performs path finding to determine a viable path for the requested end-points C and F.
10. PCE_2 uses NSI Topology over its service region to determine a viable Transport Plane path for the request, returning two Connection segments (C,D)(E,F).

11. AG₂ then makes a reservation request to the uPA₂ for the local Connection segment (C,D).
12. To complete the end-to-end Connection, AG₂ makes a reservation request to the uPA₃ for the local Connection segment (E,F).

NB: It should be noted that since Connection segments (A,B) and (C,F) are distinct, steps 4-7 and 8-12 can be done in parallel. This can be generalized to anywhere an AG has multiple children

2.1.2 Tree-Based Source Routing

In a tree-based source routing workflow, it is typical for the (root) AG receiving the request from the uRA to perform path finding and segmentation using NSI Topology for the whole service region and route the requests to the appropriate NSAs accordingly. Figure 3 shows an example of such a workflow with the associated Path Computation Engines (PCEs):

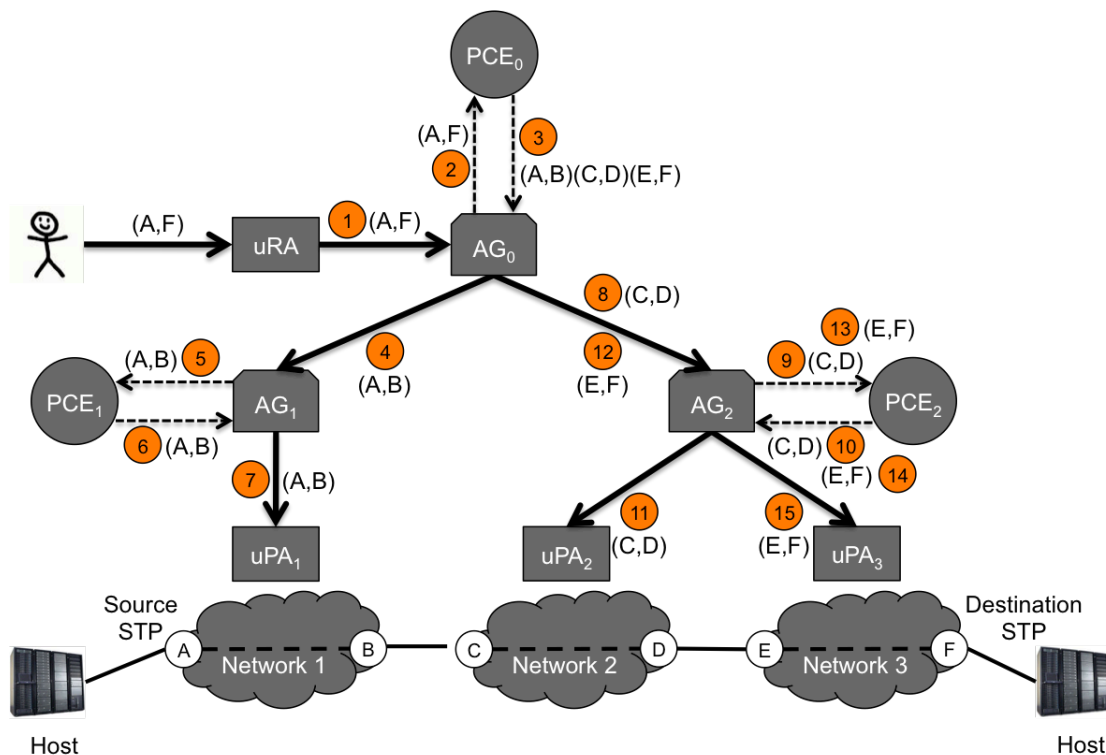


Figure 3. Example of a tree-based source routing workflow

1. The uRA makes a reservation request to the root Aggregator AG₀ on behalf of the user to Connect STP A to STP F.
2. AG₀ receives the request (A,F), and after validation, performs path finding to determine an end-to-end path through the various networks.
3. PCE₀ uses NSI Topology over the whole service region to determine an explicit Transport Plane path for the request, returning three Connection segments (A,B)(C,D)(E,F).
4. AG₀ then sends the Connection request for segment (A,B) to AG₁ which represents the only signaling path to uPA₁.
5. AG₁ receives the request (A,B), and after validation, performs path finding to determine which local STP should be involved in the Connection reservation.
6. PCE₁ determines that both the source and destination STP are managed by uPA₁ and returns the Connection segment (A,B).

7. AG₁ then makes a reservation request to the uPA₁ for the local Connection segment (A,B).
8. With uPA₂ having only a trust relationship with AG₂, AG₀ sends the request to uPA₂ for Connection segment (C,D) via AG₂.
9. AG₂ receives the request (C,D), and after validation, performs path finding to determine which local STP should be involved in the Connection reservation.
10. PCE₂ determines that both the source and destination STP are managed by uPA₂ and returns the Connection segment (C,D).
11. AG₂ then makes a reservation request to the uPA₂ for the local Connection segment (C,D).
12. To complete the end-to-end Connection, AG₀ sends the request for Connection segment (E,F) to AG₂ as it is the only signaling path to uPA₃.
13. AG₂ receives the request (E,F), and after validation, performs path finding to determine which local STP should be involved in the Connection reservation.
14. PCE₂ determines that both the source and destination STP are managed by uPA₃ and returns the Connection segment (E,F).
13. AG₂ then makes a reservation request to the uPA₃ for the local Connection segment (E,F).

2.2 Chain-Based Signaling and Path Finding Model

Chain signaling with the NSI CS protocol requires every NSA to be an AG capable of propagating a reservation request to the local uPA component (associated with local Network resources) and at most one adjacent (child) NSA associated with the next Connection segment in the data path. In general, a Connection request issued by a uRA should be to the AG associated with the head-end STP of the service, as the service request will only be signaled in one direction, from the NSA associated with source STP through to the NSA associated with the destination STP. If a uRA issues a Connection request to an arbitrary AG, the AG must first route the request to the AG associated with the head-end STP of the service. It is important to note that if the request was initially issued to an intermediate AG (i.e. an AG in the middle of the chain), the intermediate AG will maintain two records of the reservation, the first to track the routing of the request message to the head-end AG, and the second to manage the request for resources and the forwarding of the message to the downstream AG in the chain.

All NSAs in the chain request must contain a Transport Plane Connection segment associated with the reservation request. An NSA in the chain must have a peering relationship with all NSAs that it is connected to at the data plane. This implies that the Service Plane MUST be congruent with the Transport Plane. Figure 4. shows this basic chain signaling flow and NSA components involved in the deployment.




Figure 4. Chain-based signaling flow

The NSI CS 2.0 architecture currently supports two path-finding models for chain-based deployments: Hop-by-hop routing, and source routing. These are discussed in the next two sections.

2.2.1 Chain-Based Hop-by-Hop Routing

Hop-by-hop routing is a chain-based solution that makes a localized routing decision at each NSA along the service path while using the NSI Topology to guide next hop decisions. Figure 5 shows the following hop-by-hop routing workflow:

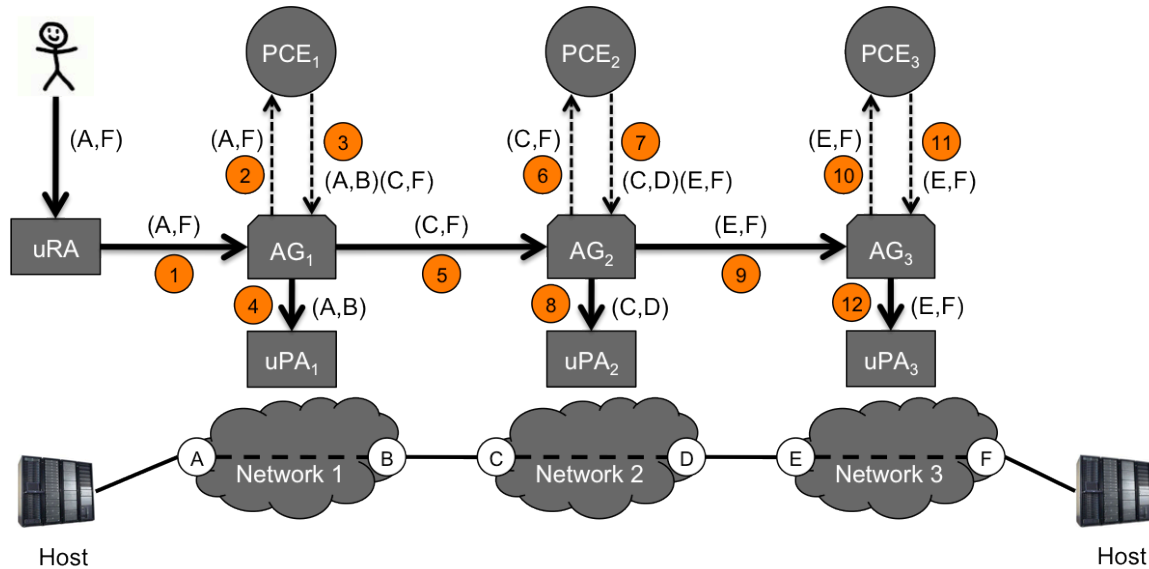


Figure 5. Example of a chain signaling workflow using hop-by-hop routing

1. The uRA make a reservation request to the head-end Aggregator AG₁ on behalf of the user to interconnect STP A to STP F. This request is made to the head-end Aggregator NSA associated with the source STP in Network 1.
2. AG₁ receives the request (A,F), and after validation, performs path finding to determine which local STP should be involved in the Connection reservation, and therefore, the peer NSA that will be next in the chain.
3. The PCE₁ uses NSI Topology over the whole service region to determine a loose path for the request, returning the local Connection segment (A,B), and the remaining path segment in the form of two STPs: the ingress STP in the adjacent Network 2 and the original destination STP (C,F).
4. AG₁ makes a reservation request to the local uPA₁ for the local Connection segment (A,B).
5. AG₁ then sends a new reservation request to the next NSA (Network 2) in the chain for the remaining Connection segment (C,F).
6. AG₂ receives the request (C,F), and after validation, performs path finding to determine which local STP should be involved in the Connection reservation, and therefore, the peer NSA that will be next in the chain.
7. PCE₂ also uses the NSI Topology over the whole service region to determine a loose path for the request, returning the local Connection segment (C,D), and the remaining path segment in the form of two STPs: the ingress STP in adjacent Network 3 and the original destination STP (E,F).
8. AG₂ makes a reservation request to the local uPA₂ for the local Connection 3 segment (C,D).

9. AG₂ then sends a new reservation request to the next NSA (Network 3) in the chain for the remaining Connection segment (E,F).
10. AG₃ receives the request (E,F), and after validation, performs path finding to determine which local STP should be involved in the Connection reservation.
11. PCE₃ determines that both the source and destination STP are within the local Network and returns the local Connection segment (E,F). The chain is complete and no further segments are returned.
12. AG₃ makes a reservation request to the local uPA₃ for the local Connection segment (E,F).

2.2.2 Chain-Based Source Routing

Source routing is a chain-based solution where the head-end Aggregator NSA (or uRA) uses The NSI topology for the whole service region to partially or completely specify a path the Connection will take through the network. This detailed path information is passed from NSA to NSA along signaling path using an Explicit Route Object (ERO) within the service request. Each NSA is bound by the ERO to follow the path segments specified during its own path finding activities. If an NSA along the path cannot meet the constraints specified in the ERO, the reservation request is rejected and an error is returned to the head-end Aggregator (or uRA) that can attempt an alternative path. Figure 6 shows the following source routing workflow:

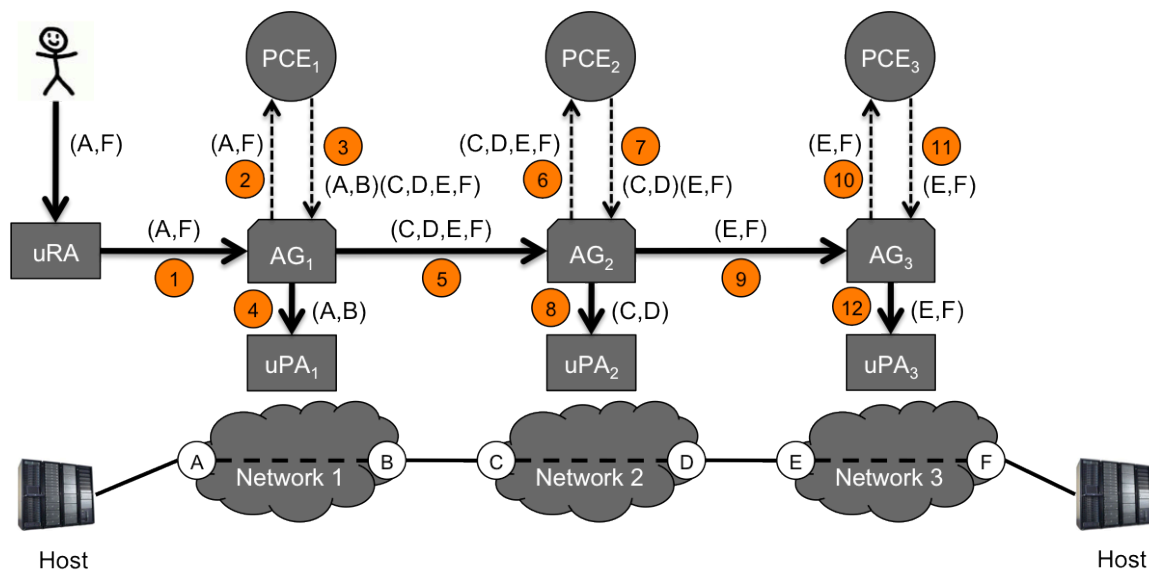


Figure 6. Example of a chain signaling workflow using source routing

1. The uRA make a reservation request to the head-end Aggregator AG₁ on behalf of the user to interconnect STP A to STP F. This request is made to the head-end Aggregator NSA associated with the source STP in Network 1.
2. AG₁ receives the request (A,F), and after validation, performs path finding to determine an end-to-end path through the network.
3. The PCE₁ uses NSI Topology over the whole service region to determine an explicit path for the request, returning the local Connection segment (A,B), and the remaining path segment in the form of a set of STP in other Network to include in the service (C,D,E,F). The egress STP B from Network 1 is connected to the ingress STP C in the adjacent Network 2, and therefore, the peer NSA that will be next in the chain.
4. AG₁ makes a reservation request to the local uPA₁ for the local Connection segment (A,B).

5. AG₁ then sends a new reservation request to AG₂ in the chain for the remaining Connection segments (C,D,E,F).
6. AG₂ receives the request (C,D,E,F), and after validation, performs path finding to determine which local STP should be involved in the Connection reservation, and therefore, the peer NSA that will be next in the chain.
7. PCE₂ uses NSI Topology over the whole service region to determine path for the request, returning the local Connection segment (C,D), and the remaining path segment (E,F).
8. AG₂ makes a reservation request to the local uPA₂ for the local Connection segment (C,D).
9. AG₂ then sends a new reservation request to AG₃ in the chain for the remaining Connection segment (E,F).
10. AG₃ receives the request (E,F), and after validation, performs path finding to determine which local STP should be involved in the Connection reservation.
11. PCE₃ determines that both the source and destination STP are within the local domain, so returns the local Connection segment (E,F). The chain is complete and no further segments are returned.
12. AG₃ makes a reservation request to the local uPA₃ for the local Connection segment (E,F).

3 Conclusion

The above sections have outlined several rudimentary examples of how an end-to-end path can be segmented and corresponding requests routed to the appropriate NSAs. It does not however dictate the algorithm that the path finder uses to determine the “optimal” path and segments, which is beyond the scope of this document. It should be noted however that in all the schemes discussed above, it is necessary to apply the path computation against the NSI Topology over the whole service region topology and not simply partial fragments of topology.

4 Security Considerations

Security considerations are dealt with in Open Grid forum GWD-R draft-trompert-gwdi-nsi-aa-v04, NSI Authentication and Authorization [3].

No additional security issues have been raised.

5 Contributors

Chin Guok, ESnet
Tomohiro Kudoh, AIST
John MacAuley, SURFnet
Guy Roberts, GEANT
Henrik Thostrup Jensen, NORDUnet
Hans Trompert, SURFnet

6 Glossary

The following terms are defined in NSI:

Aggregator (AG)	The Aggregator is an NSA that has more than one child NSA, and has the responsibility of aggregating the responses from each child NSA.
-----------------	---

Connection	A Connection is an NSI construct that identifies the physical instance of a circuit in the Transport Plane. A Connection has a set of properties (for instance, Connection identifier, ingress and egress STPs, capacity, or start time). Connections can be either unidirectional or bidirectional.
Connection Service (CS)	The NSI Connection Service is a service that allows an RA to request and manage a Connection from a PA.
Connection Service Protocol	The Connection Service Protocol is the protocol that describes the messages and associated attributes that are exchanged between RA and PA.
NSA Description Document	The NSA Description Document is a document that provides information about the services and features supported by an NSA.
ERO	An Explicit Routing Object (ERO) is a parameter in a Connection request. It is an ordered list of STP constraints to be used by the inter-Network pathfinder.
Network	A Network is an Inter-Network topology object that describes a set of STPs with a Transfer Function between STPs.
Network Services	Network Services are the full set of services offered by an NSA. Each NSA will support one or more Network Services.
Network Service Agent (NSA)	The Network Service Agent is a concrete piece of software that sends and receives NSI Messages. The NSA includes a set of capabilities that allow Network Services to be delivered.
Network Services Framework (NSF)	The Network Services framework describes an NSI message-based platform capable of supporting a suite of Network Services such as the Connection Service and the Topology Service.
Network Service Interface (NSI)	The NSI is the interface between RAs and PAs. The NSI defines a set of interactions or transactions between these NSAs to realize a Network Service.
NSI Message	An NSI Message is a structured unit of data sent between an RA and a PA.
NSI Topology	The NSI Topology defines a standard ontology and a schema to describe network resources that are managed to create the NSI service. The NSI Topology as used by the NSI CS (and in future other NSI services) is described in: GWD-R-P: Network Service Interface Topology Representation [3].
Requester/Provider Agent (RA/PA)	An NSA acts in one of two possible roles relative to a particular instance of an NSI. When an NSA requests a service, it is called a Requester Agent (RA). When an NSA realizes a service, it is called a Provider Agent (PA). A particular NSA may act in different roles at different interfaces.
Service Demarcation Point (SDP)	Service Demarcation Points (SDPs) are NSI topology objects that identify a grouping of two Edge Points at the boundary between two Networks.
Service Termination Point (STP)	Service Termination Points (STPs) are NSI topology objects that identify the Edge Points of a Network in the intra-network topology.
Service Plane	The Service Plane is a plane in which services are requested and

	managed; these services include the Network Service. The Service Plane contains a set of Network Service Agents communicating using Network Service Interfaces.
Topology Distribution Description ServiceDocument	The NSI Topology distribution Service description document contains a description of the Network topology. This document is available to be distributed as a way of sharing topology information between trusted NSAs. is a service that allows the NSI topology to be exchanged between NSAs.
Transport Plane	The Transport Plane refers to the infrastructure that carries the physical instance of the Connection, e.g. the Ethernet switches that deliver the Connection. This is also commonly referred to elsewhere as the data plane.
Ultimate PA (uPA)	The ultimate PA is a Provider Agent that has an associated NRM.
Ultimate RA (uRA)	The Ultimate RA is a Requester Agent is the originator of a service request.

7 Intellectual Property Statement

The OGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the OGF Secretariat.

The OGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights, which may cover technology that may be required to practice this recommendation. Please address the information to the OGF Executive Director.

8 Disclaimer

This document and the information contained herein is provided on an "As Is" basis and the OGF disclaims all warranties, express or implied, including but not limited to any warranty that the use of the information herein will not infringe any rights or any implied warranties of merchantability or fitness for a particular purpose.

9 Full Copyright Notice

Copyright (C) Open Grid Forum (2012-2017). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included as references to the derived portions on all such copies and derivative works. The published OGF document from which such works are derived, however, may not be modified in any way, such as by removing the copyright notice or

references to the OGF or other organizations, except as needed for the purpose of developing new or updated OGF documents in conformance with the procedures defined in the OGF Document Process, or as required to translate it into languages other than English. OGF, with the approval of its board, may remove this restriction for inclusion of OGF document content for the purpose of producing standards in cooperation with other international standards bodies. The limited permissions granted above are perpetual and will not be revoked by the OGF or its successors or assignees.

10 References

1. OGF GFD-R-212, Network Service Interface Connection Service, v2.0
<http://www.ogf.org/documents/GFD.212.pdf>
2. Open Grid forum GFD-I-213, Network Services Framework v2.0,
<http://www.ogf.org/documents/GFD.213.pdf>
3. Open Grid forum GWD-R draft-trompert-gwdi-nsi-aa-publi-comment-v7, NSI Authentication and Authorization,
https://redmine.ogf.org/dmsf_files/13424?download=