

Network Services Interface

gateway for future network services

Authors and Affiliations:

<p>Radosław Krzywania <radek.krzywania@man.poznan.pl> Poznan Supercomputing and Networking Center, Noskowskiego 12/14, 61-704 Poznan, Poland</p>	<p>Tomohiro Kudoh <t.kudoh@aist.go.jp> National Institute of Advanced Industrial Science and Technology (AIST) Tsukuba-Central 2, 1-1-1 Umezono Tsukuba Ibaraki 305-8568 Japan</p>
<p>Joan A. García-Espín <joan.antoni.garcia@i2cat.net> I2CAT C\ Gran Capità 2-4 Edifici Nexus I, 2ª planta, desp. 203 08034 Barcelona Spain</p>	<p>John MacAuley <john.macauley@surfnet.nl> SURFnet Radboudkwartier 273 3511 CK Utrecht The Netherlands</p>
<p>Chin Guok <chin@esnet.net>, Lawrence Berkeley National Laboratory, Energy Sciences Network, 1 Cyclotron Road Mail stop 50A-3111 Berkeley, CA 94720 USA</p>	<p>Guy Roberts <guy.roberts@dante.org.uk> DANTE City House 126-130 Hills Rd Cambridge CB2 1PQ England</p>
<p>Inder Monga <imonga@en.net> Lawrence Berkeley National Laboratory, Energy Sciences Network, 1 Cyclotron Road Mail stop 50A-3111 Berkeley, CA 94720 USA</p>	<p>Jerry Sobieski <jerry@nordu.net> NORDUnet A/S Kastruplundgade 22 DK-2770 Kastrup DENMARK</p>
<p>Jeroen van der Ham <vdham@uva.nl> University of Amsterdam Science Park 904, 1098 XH Amsterdam, The Netherlands</p>	

Keywords

Network, protocol, connections, services, global, multi-domain

Abstract

The Network Services Interface (NSI) [1] was created as a result of collaborative development by network and application engineers primarily associated with the Research and Education (R&E) community. The objective was to deliver network infrastructures as a service to both novice and expert end users. The first step was to design a protocol which could enable the automated creation of multi-domain heterogeneous network circuits and offer it as a “*Connection Service*” at global scale.

The NSI Connection Services (NSI-CS) [2] concept was formalized into a protocol specification in August 2011 under the umbrella of the Open Grid Forum (OGF) NSI-WG. Within three months, seven distinct implementations of the NSI protocol demonstrated interoperability at a global scale by negotiating resources and provisioning circuits across twelve domains. The first protocol interoperability demonstration took place at the GLIF Technical meeting in September 2011 in Rio de Janeiro (Brazil), followed by provisioned circuits over a global topology at the Future Internet Week in October 2011 in Poznan (Poland), and at the SuperComputing’11 event in November 2011 in Seattle (WA, USA). The current list of tools for network provisioning, which are capable of using the NSI protocol, includes AutoBAHN (GÉANT), DynamicKL (KISTI), G-lambda (AIST, NICT, KDDI Labs), OpenDRAC (CESNET, SURFnet, UvA), OpenNSA (NORDUnet, UvA, CANARIE/Northwestern University), and OSCARS (ESnet, Internet2, RNP). These protocol implementations are currently prototypes based on draft-standard documents, however there are plans for future operational deployment once the final standard is published.

The NSI is not just a protocol, but also a framework for services and protocols development. It consists of an abstracted set of principles and relations that form the basis of the protocol, which can be used to build a variety of architectures and services. The first specified service was the Connection Service that delivers end-to-end provisioning features. Other services are being explored, including topology exchange, performance verification, fault localization and remediation, and provisioning of general IT resources. The NSI-WG is an open working group, comprised of a well-rounded group of contributors with a wide variety of applicable skillsets. The open involvement ensures that the NSI framework is developed within a multi-dimension environment.

Currently defined draft-standard version of the NSI CS is stable and functional, making it easy to migrate into a production deployment over networking infrastructures around the globe. It has a well-specified reservation state machine, that assures that processing is performed consistently and correctly according to the NSI CS rules and system policies. The current state machine has eleven states and supports the following operations; i) accepting a request, ii) reserving resources, iii) provisioning a connection, iv) releasing a connection after a pre-defined time, and v) terminating a reservation at any time. Reservations can be requested to begin immediately or to be scheduled for a future time, based on availability of network resources. Once confirmed, these reserved resources are made available for end users use at the appointed time.

As the NSI CS by itself is not a provisioning system, but rather a stateful protocol for resource scheduling and provisioning, it interfaces with independent provisioning tools that are deployed within each individual network domain in order to manage the specific hardware resources. These local resources management tools are extensively used by the NSI CS protocol for verification of local resource availability, path-finding, reserving, and provisioning of resources. The NSI protocol augments this support to provide the ability to deliver inter-domain services and extend service

coverage on a global dimension. By involving more and more deployments in addition to recent demonstration partners, the reach of the NSI is continually growing. Despite the fact that the current environment was built for a proof of concept, there are already end users developing against the offered service, showing the potential and interest in unified network services delivered at the largest possible scale.

The OGF NSI-WG activity is leveraging the critical mass to push the work forward. The demonstrations have shown the usability of the protocol as a proof of concept for global network services development. Moving forward, the NSI-WG is now focused on releasing a new stable version of the NSI CS protocol that will be the base for operational deployment. Most of the engaged NRENs and organizations have expressed interests in providing a NSI CS capable peering for their infrastructures that support approximate timescale and resource commitments. In 2012, it is expected that there will be a series of NSI showcase events to expand the NSI cloud in order to reach more infrastructures and end users.

The NSI CS protocol readiness for operational activities has been proven, however, there are still important functionality required in order to deliver the highest reliability and level of the service. The NSI-WG has created a list of features which are subject to investigations for further protocol releases. The list involves issues learned from the past three demonstrations of the NSI protocol, as well as from experiences and requirements of the NSI contributors and potential users.

The two most crucial issues are the Authentication Authorization Infrastructure (AAI) and Topology exchange functionality. Security is of utmost importance in any contemporary network service or protocol in order to prevent intruder attacks, unauthorized control, or resource abuse. The NSI CS service primitives are authenticated and authorized at all network service boundaries. However, in a multi-domain environment, just as each network may incorporate different technologies, likewise AA requirements may vary from network to network. While the NSI framework provides the architecture for effective policy enforcement, the specifics of those policies, and how a common shared security profile is defined and projected among a group of globally distributed autonomous networks remains a challenge. A best practice recommendation for initial security profile is forthcoming based on experience with the 2011 pilot demonstrations.

The second aspect that needs additional consideration is topology exchange, which will provide mechanisms to support dynamic network service management at global scale. The NSI path finding functions rely on topology information in order to make efficient use of infrastructure resources. To be truly scalable at a global scope, topology management must be a distributed process, with a comprehensive ontology and standard semantics, representation, and common functional primitives enabling a continuous dynamic exchange of topology information among NSI networks. The current NSI Topology model is a very basic topology model that provides network and endpoint identification, adjacencies, and specification of the automated agent responsible for each network. Due to the small working pilot project over the last year, the experimental infrastructure topology has been managed manually and centrally as the NSI CS protocol was refined. Now, the issue of Topology management must be revisited to provide a more scalable and robust NSI Distributed Topology eXchange protocol (NSI DToX). Once both the AAI and Topology distribution issues have been addressed and a consensus approach decided, the protocol will become a fully functional proposition for a wide variety of service deployments in either research or production environments.

The list of future features are not limited to these two previously mentioned issues, but also include monitoring, accounting, service reliability and robustness to name a few. Developing all these functionalities will not be trivial but it provides a systematic roadmap for the NSI-WG and the evolutionary path for the Network Services Interface Framework.

References

[1] NSI Architecture Document v1.0, Guy Roberts, Tomohiro Kudoh, Inder Monga, and Jerry Sobieski, OGF, 2011

[2] NSI Connection Service Protocol v1.0, Guy Roberts, Tomohiro Kudoh, Inder Monga, and Jerry Sobieski, OGF, 2011

Author Biographies

Radosław Krzywania received the M.Sc. degree in Computer Science – Software Engineering from the Poznan University of Technology in 2003. He is working in Poznan Supercomputing and Networking Center as a network engineer. He is responsible for research task of GÉANT3 for develop new functionality of AutoBAHN Bandwidth on Demand system, and is an activity leader for network infrastructure in FEDERICA project. He is also interested in resources virtualization, efficient network utilization an management, as well as running, diving, and harp playing

Joan A. García-Espín is a research project manager at the Distributed Applications and Networks Area of the i2CAT Foundation in Barcelona, Spain. He received his MSc degree from the Technical University of Catalonia (UPC) in 2007 for a thesis on design and implementation of TE-enabled, DiffServ-aware MPLS networks providing end-to-end QoS. He is currently the work package leader for the design and implementation of the Logical Infrastructure Composition Layer in the EU-FP7 GEYSERS project. He supported the design and implementation of the bandwidth on demand tool named Harmony in the EU FP6 Phosphorus project. He was also one of the contributors to the Network Services Framework and Interface recommendation from the Open Grid Forum.

Jeroen van der Ham Jeroen van der Ham received his MSc in Artificial Intelligence from Utrecht University in 2002, his MSc in System and Network Engineering in 2004, and his PhD in 2010 at the University of Amsterdam on the topic of "Semantic descriptions of complex computer networks". He is currently working as a researcher at the System and Network Engineering research group at the University of Amsterdam. His research interests are in semantic descriptions of multi-layer and multi-domain networks and (virtualised) resources, as well as associated algorithms and architectures. He is currently involved in various EU-funded projects, including NOVI.

Chin Guok joined ESnet in 1997 as a network engineer, focusing primarily on network statistics. He was a core engineer in the testing and production deployment of MPLS and QoS (Scavenger Service) within ESnet. He is the technical lead of the ESnet On-Demand Secure Circuits and Advanced Reservation System (OSCARS) project, which enables end users to provision guaranteed bandwidth virtual circuits within ESnet. He also serves as a co-chair of the Open Grid Forum On-Demand Infrastructure Service Provisioning Working Group.

Inder Monga is developing new ways to advance networking services for collaborative and distributed science by leading research and services within ESnet. He also serves as the co-chair of the Network Services Interface working group in the Open Grid Forum. Monga's research interests include network virtualization, network energy efficiency, grid/cloud computing and sensor networking. He currently holds 10 patents and has over 15 years of industry and research experience in telecommunications and data networking at Wellfleet Communications, Bay Networks, and Nortel. He earned his undergraduate degree in electrical/electronics engineering from Indian Institute of Technology in Kanpur, India, before graduate studies in Boston University's EECS Department.

Tomohiro Kudoh received his Ph.D. degree from Keio University in Japan in 1992. He joined National Institute of Advanced Industrial Science and Technology (AIST) in 2002. He currently serves as the group leader of the Grid Infraware Research Group of Information Technology Research Institute, AIST. In the past few years his research has focused on network as a Grid infrastructure. His recent work also includes the G-lambda project which target is to define an interface to manage network as a Grid resource. He is a co-chair of the OGF NSI working group.

John MacAuley received a M.Sc. degree in Computer Science from The University of Western Ontario in 1996 while working as a software designer in the optical networking division of Bell-Northern Research. He had a long career at Nortel performing varying architecture roles within the company ranging from network management software architecture to protocol design and standardization. In 2005 while at Nortel he developed the Dynamic Resource Allocation Controller (DRAC) for dynamic reservation and provisioning of optical, SONET/SDH, and Ethernet services, which later that year was deployed as a service within the SURFnet network. He joined SURFnet in 2009 as a consultant to continue his work as technical lead on the newly open sourced OpenDRAC project.

Guy Roberts received his BEng degree from RMIT University in Australia in 1991 and completed a PhD on integrated semiconductor optical amplifiers for fast packet switching at the University of Cambridge in 2007. His experience in the telecommunications sector began with the rollout of SDH into Telstra's transmission network. In 1995 he moved to Fujitsu as product architect for the multiservice access platform FSX2000 and later relocated to the UK where he was involved in the development of the FDX --- Fujitsu's xDSL access platform. In 2006 Guy joined DANTE where he works with the network engineering and planning team, he is also co---chair of the OGF NSI working group.

Jerry Sobieski joined NORDUnet (Copenhagen, DK) in 2008 as the Director for International Research Initiatives. He received his degree in Computer Science from the University of Houston in 1985. He headed the Laboratory for Parallel Computing at the University of Maryland Institute for Advanced Computer Studies from 1990 to 1997, and then Joined Internet2 to construct the Abilene network in the US from 1998-1999. Mr. Sobieski served as Director of Engineering and then Director of Research at the Mid-Atlantic Crossroads, the regional network in Washington DC, from 2000 to 2008. Mr. Sobieski has been actively involved in protocol development for multi-layer connection oriented services for over 20 years, and is an active participant in the GLIF and the OGF NSI WG. He currently resides in Washington DC and represents the Nordic R&E community in a wide array of advanced networking topics in Europe, the Americas, and the Asia/Pacific region.