Title:

Enhancing E2Emon monitoring with NML topology exchange

Authors:
Freek Dijkstra, Sander Boele, Roman Łapacz, Ronald van der Pol

Affiliations:
Freek Dijkstra, Sander Boele, Ronald van der Pol:
SARA, Science Park 121, 1098 XG, Amsterdam, The Netherlands
Email: freek.dijkstra@sara.nl, sanderb@sara.nl, rvdp@sara.nl
Roman Łapacz:
PSNC, ul. Noskowskiego 12/14, 61-704 Poznań, Poland
Email: romradz@man.poznan.pl

Keywords:
Topology description, monitoring, multi-domain, multi-layer

Technical Paper
We intend to publish a full paper after acceptance of the extended abstract.

Abstract

E2Emon is a PerfSONAR-compatible service for monitoring the link status of end-to-end circuits that span multiple domains. The software is currently in use by the LHCOPN [E2Emon-LHCOPN]. E2Emon suffers from multiple limitations, most of which are related to the way topology information is represented and exchanged. One of the most notable limitations is that it is unable to relate monitoring information from multiple layers. In order to overcome this and other limitations, a proof-of-concept is under construction, which overhauls the topology exchange, and separates the monitoring service and topology service from each other.

The Network Markup Language working group (NML-WG) of the Open Grid Forum was chartered in 2007 to make a standard for topology exchange. While the NML-WG has not produced a ratified schema, it has progressed to a state where the ideas can be applied to working code. Our proof-of-concept for cross-layer E2E monitoring will be based on NML, creating an early implementation. A standardised topology service can be shared amongst different applications, including path finding, provisioning, monitoring, and visualisation. The advantage of this approach is that all tools will have a coherent view of the network, making integration between the different tools easier.

Architecture

The perfSONAR suite uses a data description and set protocols, standardised in respectively the Network Measurement (NM) and Network Measurement and

Control (NMC) working groups in the Open Grid Forum. In the E2Emon architecture, each domain contains E2E Measurement Points that collect data of the monitored links. The monitoring information of different domains is gathered by a E2E Monitoring System that stitches together the segments and provides a status overview for the end-to-end link. In the current architecture, the Measurement Points provide both status information of each port (up, down, alarm, administrative status, etc.) as well as topology information for each port (the name of the E2E link it is part of, the name of its node, and the name of the peering node).

This architecture has the following limitations.

1. It is not easy to cope with missing or wrong information
2. Logic is ambiguous in the unlikely event that an E2E path passes the same node twice
3. Topology Service is integrated and can not be shared with other services (like AutoBAHN)
4. Unable to relate monitoring information across layers, or relate two links with different names in series
5. Unable to relate monitoring information of two links with different names in series
6. Unable to relate monitoring information of two parallel links (e.g. for a 1+1 protected path)
7. Limited to point-to-point connections (no VLAN support)
8. Topology information is stored in a static configuration file, no historic topology information is kept (although historic monitoring information is available).

The first two limitations can be fixed by relating a port to its peering ports instead of to its peering node, as it the case now. The third limitation will be solved by separating the monitoring service from the topology services, as shown in figure 1 and figure 2.

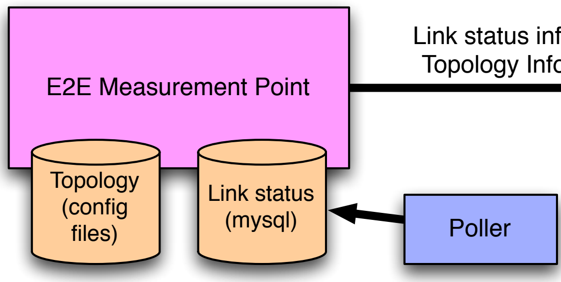

Figure 1. Previous E2Emon architecture.

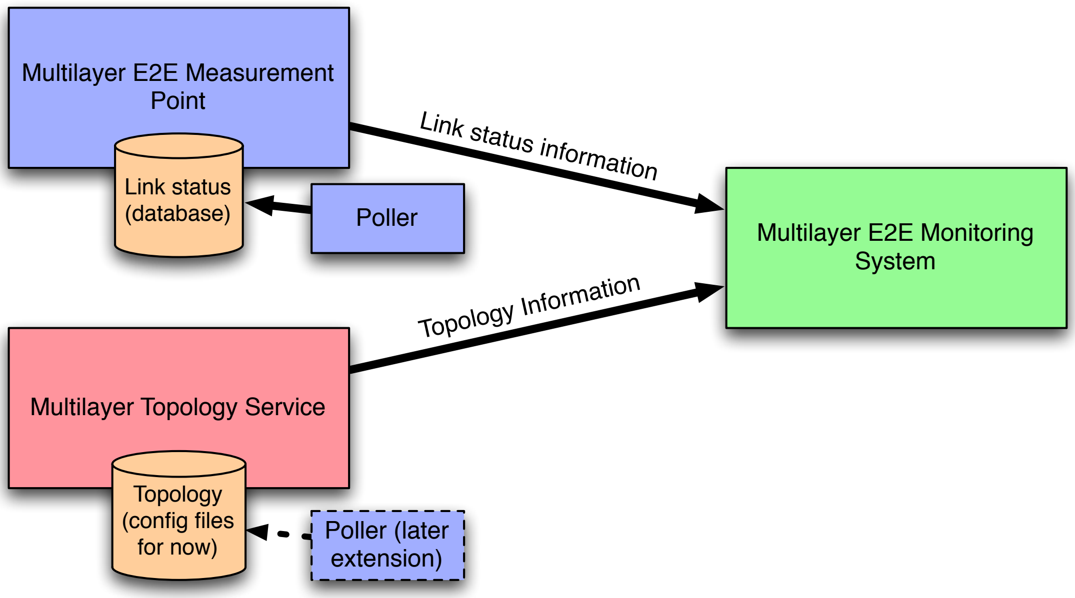

Figure 2. New E2Emon architecture.

It is expected that points 4 to 7 can be solved by using the port and link relations as defined by the NML-WG, and this proof-of-concept serves as a test case.

Status of NML standardisation

The NML working group was chartered in early 2007 to define a network topology description. Since that that time, volunteers with backgrounds in UNIS (PerfSONAR), cNIS, NDL, VXDL and other backgrounds [GFD.165] have worked on coming to a joint vision.

By 2009, the group defined a common terminology that was acceptable by all [NML-SCHEMA]. In 2010, the group made steady progress in face-to-face meetings. While we do not claim that the schema is stable yet, the OGF in October 2010 marked an important change. For the first time, over half of the slides contained examples, rather than discussions about terminology.

This is an indication that NML is getting ready for early implementations. This does not mean that the group is finished. The schema and syntax are still preliminary and in flux. In fact, the hardest part is still ahead of us. The group needs to make sure that what has been said on paper is translated into running code with actual applications, and verify that the terminology, syntax and logic is complete, unambiguous and -dare we say- actually useful in practice. For that, input from coders and implementers will become more and more useful for the NML working group.

The network measurement (NM) and network service interface (NSI) working groups, which respectively standardise monitoring and provisioning protocols agreed to use the NML model as the underlying topology model. The proof-of-concept described in this paper is one of the first adopters, which will provide feedback to the working group. The topology service in the perfSONAR-PS software suite is another early implementation, which supports a combination of topology schemas as developed in the DICE collaboration, and the NM and NML working groups [PSPS-TOPO]. It should be observed that both implementations cannot be truly NML-compatible for the simple reason that the NML schema is still in flux.

Early Results

Early results indicate that NML is able to relate different links together (eliminating limitations 4-7 in the second section.), although some minor changes are required. For example, it was possible in the preliminary NML syntax to describe a path as a sequence of consecutive link segments, but it was not possible to say that a single segment is part of a larger path, without naming the other segment. This was an easy fix, and is exactly the kind of feedback required by the NML working group.

A more fundamental result that hopefully comes out of the proof-of-concept is what exactly a topology service should encompass. A topology service may contain the following types of information:

- Topology: information on links between nodes or between domains
- Configuration: information on the dynamic state within a node or network
- Capability: information on available data transport services within a node or network

It is evident that the topology information should be part of a topology service, as it is used by most applications (path finding, visualisation and monitoring). This is less evident for configuration (which is used by provisioning and monitoring) or capability (which is only used for path finding).

In our proof-of-concept, we have decided to place the configuration information in the topology service, but it may well be decided later that the monitoring service is more appropriate, or that this is even a separate service.

Acknowledgements

This work is performed in the GÉANT (GN3) project, partially funded by the EC in the 7th framework programme.

The authors are much indebted to Jeroen van der Ham of the University of Amsterdam and Jason Zurawski of Internet2 for their detailed proofreading and significant contributions to this document.

We would like to thank the contibutors of the NML and NMC working groups, the perfSONAR-MDM and perfSONAR-PS projects, and the participants of JRA2 Task 3 in the GÉANT project. Jeff Boote, Eric Boyd, Aaron Brown, Jakub Gutkowski, Jacek Łukasik, Martin Swany, and Jason Zurawski have provided valuable feedback and input in recent discussion on the perfSONAR architecture.

Finally, we would like to thank the original authors of the E2Emon system for creating a system we can base our work on.

References

[E2EMON-LHCOPN] Fausto Vetter, Domenico Vicinanza, perfSONAR Multi-Domain Monitoring Service Deployment and Support: The LHC-OPN Use

Case, TNC June 2010, http://www.geant.net/Media_Centre/Media_Library/Media%20Library/perfsonar%20Deployment%20and%20Support.pdf

[GFD.165] Paola Grosso, et. al, Network Topology Descriptions in Hybrid Networks, GFD-I.165, Open Grid Forum, March 2010, http://www.ogf.org/documents/GFD.165.pdf

[NML-SCHEMA] NML schema, draft version July 2009, http://forge.gridforum.org/sf/go/doc15481

[PSPS-TOPO] PerfSONAR-PS Topology Service version 3.2, Internet2, November 2010, http://psps.perfsonar.net/topology/

Vitae

Freek Dijkstra received his M.Sc. in applied physics from the Utrecht University in 2002, and received his Ph.D. in Computer Science at the University of Amsterdam in June 2009 on the topic of "Path finding in multi-layer transport networks". Dijkstra is currently working as network researcher at SARA, were he develops standards and software for use in innovative networks. His interests are topology descriptions and monitoring of multi-layer and multi-domain networks. He is co-chairing the Network Markup Language working group in the Open Grid Forum (OGF).

Sander Boele received his M.Sc. in Cognitive Psychology from the Vrije Universiteit in Amsterdam in 2005. After he worked in the IT department of various companies, he joined SARA as a network engineer and innovator. His main interests lie in network monitoring and building innovative networks.

Roman Łapacz is a network specialist at Poznań Supercomputing and Networking Center (Poland). He graduated from Poznań University of Technology in 2000 and completed postgraduate management studies (Poznań University of Economics and Poznań School of Banking). He also studied at Université Paris-Dauphine within the Socrates programme. The current main work area is the network monitoring systems. Roman has been deeply involved in the perfSONAR project from the beginning of the initiative. Now he leads Joint Research Activity 2 Task 3 group in the GN3 project and works on new innovative extensions for the multi-domain monitoring perfSONAR framework.

Ronald van der Pol is a network researcher at SARA. He has been working in the field of Education and Research Networks for more than twenty years. His former employers include the VU University in Amsterdam, SURFnet and NLnet Labs. At SURFnet he specialized on IPv6 and is co-author of several RFCs. In recent years his focus is on optical networks, especially hybrid network monitoring, end-to-end performance and network access for demanding applications. He holds masters degrees in both Physics and Computer Science.