Monitoring End to End Bandwidth on Demand Circuits over Ethernet Infrastructure

C.Bouras[1,2], A.Gkamas[1] and K.Stamos[1,2]

[1]Research Academic Computer Technology Institute, N. Kazantzaki Str, University Campus 26500, Rio Greece
[2]Computer Engineering and Informatics Dept., University of Patras
e-mail: bouras@cti.gr, gkamas@cti.gr, stamos@cti.gr

Abstract

AutoBAHN (Automated Bandwidth Allocation across Heterogeneous Networks) has specified and is now prototyping a Bandwidth on Demand (BoD) service intended to operate in a multi-domain environment using heterogeneous transmission technologies. The AutoBAHN system aims at providing a guaranteed capacity, connection-oriented service between two end points. This paper highlights the architecture of the monitoring system for the AutoBAHN service, the implementations challenges and the decisions taken, focusing on monitoring end to end BoD reservations over ethernet infrastructure. In addition this paper presents the architecture and the components of the AutoBAHN monitoring system for ethernet infrastructures.

Keywords

Bandwidth on Demand, Quality of Service, Network Monitoring, Ethernet Monitoring

1. Introduction

The GN2 European project (GEANT2 Web site, 2007) is a research project funded by the European Union and Europe's NRENs (National Research and Education Networks). The purpose of GN2 is to build the seventh generation of the pan-European research and education network, which connects universities, institutions and other research and educational organizations around Europe and interconnects them to the rest of the Internet using high-speed backbone connections.

Within the GN2 project, a number of joint research activities (JRAs) are being pursued. It has been understood that in the research and academic environment, there are applications and research fields (such as radioastronomy, high-energy physics and general Grid applications) with strict demands for the provisioning of guaranteed and dedicated capacity. For this reason, the GN2 project is developing the AutoBAHN (Bandwidth on Demand – BoD) Joint Research Activity 3. The AutoBAHN service aims at providing a guaranteed capacity, connection-oriented service between two end points. Due to the intrinsic multi-domain environment of research and education networking in Europe, JRA3 decided to focus most of its initial implementation effort in the inter-domain communication, as it is the part of the system which can rely less on known, proven and existing standards. An integral part of the AutoBAHN service is the monitoring of the provisioned end-to-end connections. We have therefore undertaken the design and implementation of the monitoring module for the AutoBAHN system, and this is the focus of this paper.

The overall architecture of the AutoBAHN system, its goal and the network mechanisms it employs are thoroughly presented in Campanella et al. (2006). This paper highlights the architecture of the monitoring system for the AutoBAHN service, the implementations challenges and the decisions taken, focusing on monitoring end to end BoD reservations over ethernet infrastructure. The rest of the paper is structured as follows:

Section 2 presents the general architecture of the AutoBAHN system, at its current development status. Section 3 discusses related work, focused on whether and how relevant projects have dealt with the monitoring issue. Section 4 presents the general architecture of the monitoring module, while section 5 gives the implementations details for the parts comprising the monitoring module. Finally, section 6 concludes the paper and presents future fields of related study.

2. GN2 / AutoBAHN BoD System

The GN2 Joint Research Activity 3 (JRA3) or AutoBAHN (Automated Bandwidth Allocation across Heterogeneous Networks) has specified and is now prototyping a Bandwidth on Demand (BoD) service intended to operate in a multi-domain environment using heterogeneous transmission technologies.

The GN2 JRA3 BoD system is composed of the following modules (as displayed in Figure 1): Domain Manager (DM), Inter-domain Manager (IDM), Technology Proxy (and Resource Modeling) module, AAI module, Policy module, Pathfinder (inter-domain and intra-domain). It also contains two ancillary modules: the Information Storage System and the Location service.

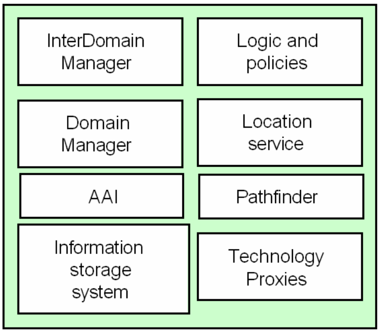

Figure 1: JRA3 BoD system Modules

The Inter-domain Manager (IDM) is the BoD module responsible for receiving BoD service requests from a user, application or another domain and is responsible for approval and instantiation of these requests.

The main function of the Domain Manager (DM) is to instantiate BoD instances within a single domain. The DM has a detailed knowledge of the topology of its domain. It participates in the inter-domain pathfinding process by examining the feasibility of providing an end-to-end path within its local domain. It contacts the Technology Proxy module to request the configuration of the BoD service instance.

The Technology Proxy module performs the translation of requests received by the DM (in abstract network language) into vendor- or equipment-specific configurations. The BoD architecture foresees a set of proxy modules, each one independent of the others, which can be modularly attached to the BoD system. The proxy module may configure the network via an existing Network Management System (NMS) or act as a GMPLS (Mannie, 2004) agent; consequently it does not act directly upon the network, but relies on an intermediate control layer.

The Policy module contains all the rules and policies which are available for use by other modules when inspecting and elaborating on a request. The rules are collected in a single module for easy maintenance, modification and to enforce coherence.

The Pathfinder module contains the algorithms and the logic to search for a path that satisfies each BoD reservation request according to specific sets of constraints, algorithms and policies. The module is composed of two main blocks devoted to inter-domain and intra-domain enquiries. The search returns a list of candidate paths using constraint-based shortest path algorithms.

The Information Storage System and the Location Service function mainly offer support to the other modules. The Information Storage System is responsible for providing storage, archival and database functionalities for data explicitly relevant to the BoD system, while the Location Service locates the addresses of all type of services and modules.

If the local administrative domain does not provide basic network management services, such as a Network Management System (NMS), monitoring services and an Authorization and Authentication Infrastructure (AAI), the BoD architecture may implement the minimum required level of these functionalities.

The architecture recognizes the importance of an addressing and labeling scheme for all the components of the BoD system. The AutoBAHN system uses IPv6 addresses both as labels and addresses at the control plane.

In this paper, we are interested in the end to end monitoring module for ethernet-based technology infrastructure, which is part of the DM and is presented in detail in section 4.

3. Related Work

The AutoBAHN BoD system has been influenced by a number of other projects dealing with similar challenges for bandwidth on demand provisioning. In this section we present some of the most closely related ones, with an emphasis on their approach to monitoring the status of the provisioned resources.

The DRAGON project (Leung et al., 2006) is also conducting research and developing technologies to enable dynamic provisioning of network resources on an interdomain basis across heterogeneous network technologies. It mainly deals with GMPLS enabled domains. It currently provides web-based real-time monitoring information, based on cricket and nagios tools.

OSCARS/BRUW project (Guok, 2005) focuses on L3 MPLS QoS and adopts SNMP queries to the routers for monitoring LSP teardown and usage.

The UCLP community project and the related Argia commercial product enable users to control and manage network elements for the purposes of establishing End-to-End (E2E) lightpaths (Wu et al., 2003, Argia Web site, 2007). In the UCLP model, the users fully own and control the network resources.

The MUPBED project attempts to integrate and validate, in the context of user-driven large-scale testbeds, ASON/GMPLS technology and network solutions (Cavazzoni, 2007). MUPBED takes into account that for an optical network operated according to OTN, specific fields of the frame are reserved to user's monitoring of a connection. Moreover, with tandem connection monitoring (TCM, International Telecommunications Union standard G.707), a layered multi-operator monitoring can be carried out, allowing lightpath provisioning also in a multi-domain environment.

VIOLA (VIOLA Web site, 2007) is another related project for the development and test of software tools for the user-driven dynamical provision of bandwidth but does not focus on the development of tools for the monitoring of the provisioned resources. VIOLA is focused on the area of network technology and application development as part of a testbed environment.

4. Architecture of the Monitoring System

Before the description of the monitoring module architecture, some basic principles are presented for the concepts of End to End (E2E) Monitoring:

- Intradomain link: Intradomain link is a link between two nodes (routers or switches) that both belong to the same domain.

- Virtual link: Virtual link is the conceptual link which connects two border nodes (routers or switches) that both belong to the same domain. A virtual link is the result of the conjunction of many Intradomain links. Virtual links are the kind of information that a domain presents to the outside world about its internal state: it does not propagate detailed information about its internal links; instead it provides aggregated information about Edge to Edge connections represented by Virtual links.

- Interdomain link: Interdomain link is a link between two border nodes (routers or switches) of two different domains. Interdomain links are typically monitored by both domains they connect to, but the case where only one domain monitors an interdomain link can also be handled.

- E2E (End to End) link: E2E link is the conceptual link which connects two end points from different domains. An E2E link is the result of the conjunction of many Interdomain and Virtual links.

Each domain provides information about the links it monitors (these links can be either Virtual links or Interdomain links). The E2E Monitoring system computes the overall status of an E2E link by aggregating the status information from the involved domains.

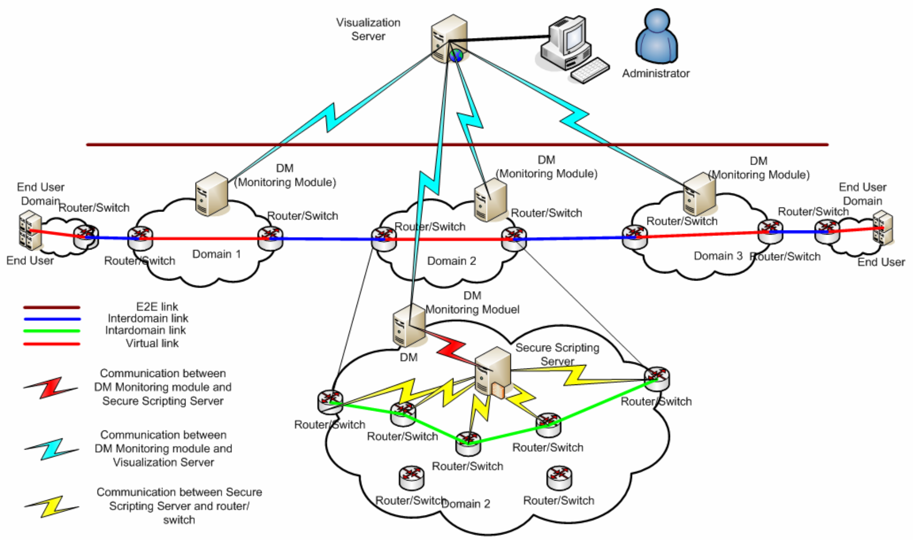

Figure 2: Ethernet Module Architecture

Figure 2 presents the architecture of the E2E monitoring system. It consists of the following modules:

- DM monitoring module: The monitoring module of the DM is responsible for monitoring the status of both Virtual and Interdomain links and provide this information to the Visualisation server.

- Secure scripting server: This module is responsible for the communication with the network devices of the domain in order to check the status of the physical links.

- Visualisation server: The visualisation server collects monitoring information from the DMs (through the corresponding monitoring module) of the domains involved in the E2E monitored link and presents monitoring information to the administrator through a graphical interface.

During the monitoring of an end-to-end Bandwidth on Demand circuit over ethernet Infrastructure the following steps take place:

The IDMs and DMs of the involved domains cooperate in order to establish an end-to-end Bandwidth on Demand circuit over the existing infrastructure. The results of this establishment will be an end-to-end circuit in the abstracted topology (presented in more detail in section 5.1) of the involved domains. This end-to-end Bandwidth on Demand circuit or E2E link for shortness is the input for the E2E monitoring system which is responsible to monitor the status of the E2E link. The E2E monitoring system extracts from the abstracted topology the domains and the corresponding interdomain and virtual links which comprise the E2E link. In repeated time intervals the E2E monitoring system polls the DM monitoring modules of the involved domains and aggregates all the status information (as detailed in section 5.3). Each DM monitoring module reports aggregate status about the links that it monitors to the E2E monitoring system. In order for the DM monitoring module to check the status of an interdomain link or an intradomain link, it has to communicate (via web services) with the Secure Scripting Server. This communication for security reasons in based on SSH. The Secure Scripting Server communicates with the corresponding routers / switches in order to check the status of the physical link that corresponds to an intradomain or interdomain link (the process is detailed in section 5.2). An interdomain link can be monitored by either one of the domains connected by this interdomain link or by both connected domains. In the latter case the E2E monitoring system is responsible for combining the monitoring information from both involved domains into a unified status for the link. The E2E monitoring system computes the overall status of an E2E link by aggregating the status information of the involved domains and presents the overall status of an E2E link through the visualization server. The communication between the visualization server and the DM monitoring system is based on web services (detailed in section 5.4).

The E2E monitoring system in its current version provides information for two monitoring parameters: The **Operational State**, which is derived from the operational state of the involved physical devices, and the **Administrative State**, which reflects the management processes performed by the domains (detailed in section 5.3).

5. Implementation Details

This paragraph provides details in various aspects regarding the E2E monitoring system for monitoring of an end-to-end Bandwidth on Demand circuit over ethernet infrastructure.

5.1. Topology Abstraction

A fundamental part of the implementation for the monitoring module is the topology abstraction process. Because the AutoBAHN system has to interoperate between multiple domains with different technologies, its main focus has been on the development of the Interdomain Managers (IDM), which are technology-agnostic and for that reason, they only recognize abstract links.

A typical request to the monitoring module as described in section 6 will request the status of an end-to-end provisioned path between two sites that can belong to different administrative domains (NRENs).

At the interdomain level, each domain is treated as a semi-black box, where information only for “virtual” intradomain links is provided. These virtual intradomain links represent an intradomain path possibly consisting of multiple actual intradomain links. The topology abstraction module is responsible for maintaining the association between the full internal knowledge of the domain topology and the abstracted view that is propagated to the outside world. The monitoring module has to use the topology abstraction process in order to identify which actual links have to be queried about their status in order to reply to an end-to-end path status request.

For an ethernet technology domain in particular, the topology abstraction process is implemented in Java and works as follows: First, a specialized java class is used in order to map a virtual link to a list of intradomain links. It has to be noted here that a virtual link may have multiple mappings of actual intradomain paths. At this point, the implementation of the monitoring module simply selects the first mapping available for a virtual link. The second step is for each intradomain link to retrieve the appropriate ethernet technology links, and in turn their starting and ending interfaces.

From that point on, the technology-specific monitoring scripts take over, as will be described in the next section.

5.2. Status of each Ethernet Monitored Link

The DM Monitoring module in order to check the status on ethernet monitored link (either interdomain or intradomain link) communicates via a web service interface with the secure scripting server. The secure scripting server is using a read only account and a number of scripts in order to communicate with the corresponding routers/switches and check the status of the physical link in which an intradomain or interdomain link is based.

The monitoring scripts used, require as input, some means of identifying the circuit to be monitored. In order to find the correspondence between the interdomain / intradomain links and physical links information from the topology abstraction is used.

The monitoring scripts are implemented in perl programming language and they use appropriate libraries in order to communicate with the command line interface of the network devices and issue commands to network devices and collect responses from network devices.

The usage of secure scripting server ensures that the communication with the network devices of the domain will be done under specific security requirements which include secure and encrypted communication and usage of read only accounts for accessing the network devices.

The use of a standardized web service interface for accessing the secure scripting server means that a domain that does not wish to install and use the secure scripting server is free to do so. It can integrate its own monitoring solution for the physical links it controls with the monitoring module we have developed, simply by making its own monitoring solution adhere to the standardized web interface.

5.3. End to End Status Aggregation

As we have already mentioned the E2E monitoring system in its current version provides information for the monitoring parameters defined as **Operational State** and **Administrative State**.

The supported values for the Operational State are: **Unknown**, which means that the Domain could not acquire information about operational state (value 0), **Up** - Link is up (value 1), **Degraded** - Link is up, but has reduced performance (value 2), and **Down** - Link is down (value 3).

The supported values for the Administrative State are: **Unknown**, which means that the Domain could not acquire information about administrative state (value 0), **NormalOperation** - No administrative work is performed (value 1), **Maintenance** - Planned maintenance activity in progress (value 2), **TroubleShooting** - Trouble shooting is in progress (value 3), and **UnderRepair** - Repair process is in progress (value 4)

The domains have to provide both Operational and Administrative States of all links monitored by this domain to the E2E monitoring system. The monitoring system then aggregates the states of the monitored links to determine the state of E2E links.

To determine the status of an E2E link, the worst state value (i.e. having the highest value) of all monitored links (interdomain and virtual links) for this particular E2E link is used. This rule applies for both Operational and Administrative States.

5.4. Visualization

In order to present the retrieved information to the user, a web-based visualization approach has been chosen. The current visualization module is based on the visualization developed for another GN2 activity, namely JRA4, and can be seen at Figure 3. The monitoring module provides the infrastructure for easily extending or replacing the visualization module.

The monitoring data is periodically polled by the Visualization module. The information on the monitored items is retrieved by the database where it has been stored as described in the preceding sections. Figure 4 presents the overall information flow, from the low-level technology-specific monitoring module towards the visualization client that presents it to the user.

The visualization server we have currently implemented is a SOAP server, while the visualization client is a SOAP client. Implementing that part of the information flow with the SOAP protocol enables us to easily decouple the visualization server from the client and produce alternate clients that may display the information in a different fashion without having to deal with the rest of the monitoring system. SOAP communication takes place using XML messages with specific format and therefore and client adhering to that specification can interpret the communicate data.

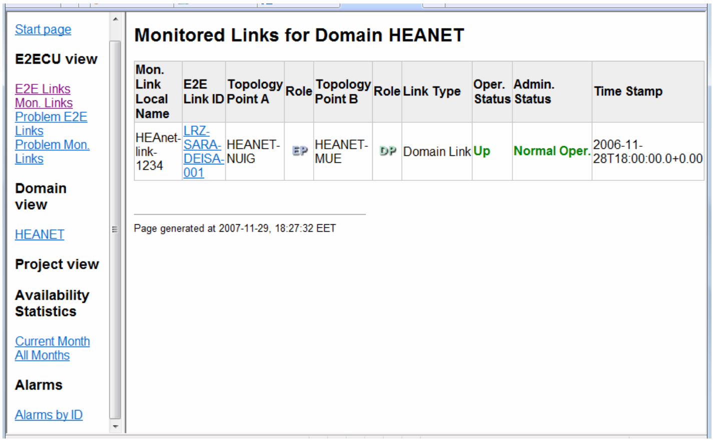

Figure 3: Visualization client GUI interface

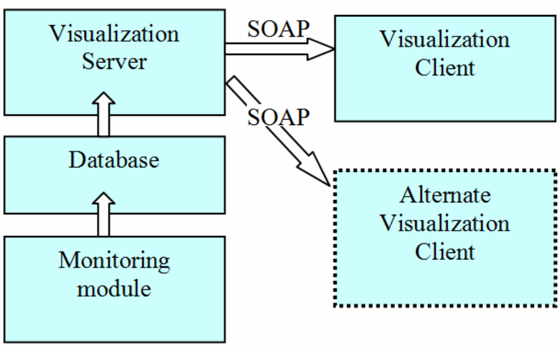

Figure 4: Information flow towards the visualization components of the monitoring system

The Visualization server has been implemented using the Apache Axis framework within an Apache Tomcat server installation. Specifically, it distributes the link status information using a "Message" style service as provided by the Axis framework, which enables us to manually customize the transmitted XML and to format it exactly as required by the E2E monitoring system request/response model.

6. Conclusion and Future Work

Designing and implementing a monitoring module in the context of an automated and heterogeneous interdomain Bandwidth on Demand system has proven to be a complex and challenging task. The main challenges we faced were the following: the introduction of a complex monitoring solution across multiple heterogeneous domains, which we managed by modularizing the system using several levels of abstraction have been adopted and by isolating compartments of the overall task. Each domain can therefore adopt the parts of the system it sees fit, and put in place its own custom components in the remaining places by obeying to the standardized interfaces we provide. Another challenge was the lack of centralized information, which we solved by compartmentalizing, abstracting and aggregating the relevant information as detailed in the technical descriptions. Current development and deployment of the monitoring module has focused on the ethernet technology which is more readily available and familiar. However the design concepts are absolutely valid for a heterogeneous technology environment, and we intend to validate this assumption by working on an extension for supporting the SDH/OTN technologies.

An important part of the future work is therefore to extend the technology specific part of this work beyond the ethernet realm. Also, the versatility of the decoupled visualization server and client pair can be used in order to develop enhanced visualization clients such as map-associated versions with geographical information factored in. In the immediate future, our plans involve the deployment of the Ethernet monitoring module to the HEANET and GRNET networks (Irish and Greek NREN respectively) which participate in the AutoBAHN activity, and the production of measurements and evaluation of the monitoring system's effectiveness.

7. Acknowledgments

The authors would like to acknowledge the contribution of their GN2 partners and especially the GN2-JRA3 group.

8. References

Campanella, M., Krzywania, R., Reijs, V., Sevasti, A., Stamos, K., Tziouvaras, C. and Wilson, D., (2006)"Bandwidth on Demand Services for European Research and Education Networks", *1st IEEE International Workshop on Bandwidth on Demand*, 27 Nov, San Francisco (USA).

Cavazzoni, C., (2007), "MUPBED Overview and Architecture", TERENA Networking Conference 2007, Copenhagen, Denmark, May.

GEANT2: The pan-European research and education network, http://www.geant2.net/

Guok, C., (2005), "ESnet On-Demand Secure Circuits and Advance Reservation System (OSCARS)", *Internet2 Joint Techs Workshop*, Salt Lake City, Utah, February 15.

Inocybe technologies Inc., Argia product, http://www.inocybe.ca/

Internet Engineering Task Force RFC 3945, (2004), "Generalized Multi-Protocol Label Switching (GMPLS) Architecture" E. Mannie, October.

Leung, F., Flidr, J., Tracy, C., Yang, Xi, Lehman, T., Jabbari, B., Riley, D. and Sobieski, J., (2006), "The DRAGON Project and Application Specific Topologies", *Broadnets 2006*, San Jose, California, USA, October 1-5.

VIOLA project - Vertically Integrated Optical Testbed for Large Applications in DFN, http://www.viola-testbed.de/

Wu, J., Campbell, S., Savoie, J.M., Zhang, H., Bochmann, G. v. and St.Arnaud, B., (2003), “User-managed end-to-end lightpath provisioning over CA*net 4”, *Proceedings of the National Fiber Optic Engineers Conference (NFOEC)*, Orlando, FL, USA, Sept 7-11, pp. 275-282.