22-09-2009

Deliverable DJ2.2.1: Specification of enhancements and developments for the AutoBAHN system

Deliverable DJ2.2.1

Contractual Date:	31-07-2009
Actual Date:	22-10-2009
Grant Agreement No.:	238875
Activity:	JRA2
Task Item:	T2
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	PIONIER
Document Code:	GN3-09-040
Authors:	M. Balcerkiewicz (PSNC), Eduard Escalona (UEssex), Vaggelis Kapoulas (GRNET), Gina Kramer (DANTE), R. Krzywania (PSNC), Rasmus Lund (NORDUNET), Jacek Lukasik (PSNC), Andrew Mackarel (HEAnet), Brian Barch Mortensen (NORDUNET), Vassilis Papapanagiotou (GRNET), Victor Reijs (HEAnet), Afrodite Sevasti (GRNET), Kostas Stamos (GRNET)

Abstract

The document describes the functionality extensions and further development plan of the AutoBAHN system, designed and deployed during GN2. The new functionality will involve support of more network transport technologies, extensive usage of AAI, accounting functionality, and resiliency options. The new features will help AutoBAHN to develop and improve to a production level service aiming at deployment of the BoD services in European NRENs.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

This deliverable provides an overview of the enhancements and developments planned for the Automated Bandwidth Allocation across Heterogeneous Networks (AutoBAHN) system during the GN3 project. This includes a review of

- Transport technologies that AutoBAHN may support for dynamic circuit provisioning (see *Domain Network Technology Support in AutoBAHN* on page 6).
- Open issues in AutoBAHN's multi-domain functions (see *Inter-Domain Stitching Issues* on page 26).
- How security and accounting may be enhanced (see *The Lookup Service* on page 44 and *Authentication, Authorisation and Accounting (AAA)* on page 51).
- Collaborative experiences (see *Inter-Systems Protocols and Collaboration* on page 60).

The conclusion summarises the challenges that the next iteration of AutoBAHN faces and provides time and resource estimates for the different tasks.

1 Introduction

The Automated Bandwidth Allocation across Heterogeneous Networks (AutoBAHN) system, which was piloted during the GN2 project, provides an interface for setting up dynamic circuits over global research and education network infrastructures. Dynamic circuit provisioning is an important factor in developing the next-generation GEANT network, using transport technologies to offer new services in addition to IP-based services [AutoBAHN].

This deliverable provides an overview of the enhancements and developments planned for the AutoBAHN system during the GN3 project. As this deliverable is issued at the start of the project, some enhancements are defined in general, providing several options to be chosen during the project lifetime. The detailed specification of the enhancements and its implementation is a significant research effort, which is undertaken by Task 2 (Hybrid Network Provisioning) of the JRA2 activity (Multi-Domain Network Service Research).

This document comprises:

- An overview of what was achieved during GN2 and how work in GN3 can continue and progress from what has already been accomplished.

- A review of different transport technologies that are under investigation by the AutoBAHN group, as AutoBAHN's capabilities are planned to be expanded to support these technologies for dynamic circuit provisioning (see *Domain Network Technology Support in AutoBAHN* on page 6).

- A look at open issues in the multi-domain functions of AutoBAHN, such as technology stitching or inter-domain links discovery (see *Inter-Domain Stitching Issues* on page 26).

- A look at security and accounting ideas. The next version of AutoBAHN is expected to support the Lookup service, and to be fully integrated with eduGAIN services (see *The Lookup Service* on page 44 and *Authentication, Authorisation and Accounting (AAA)* on page 51).

- A description of the experiences on working with other groups and activities. The directions AutoBAHN should take in the future to be interoperable with other similar systems are discussed (see *Inter-Systems Protocols and Collaboration* on page 60).

- A summary of the challenges that face the next iteration of AutoBAHN. It also gives estimates of the time and resources needed to complete each task (see *Implementation Time Lines and Conclusions* on page 62).

This document assumes the reader is familiar with the concept and architecture of AutoBAHN system. The deliverables “DJ3.3.1: Definition of Bandwidth on Demand Framework and General Architecture” [DJ3.3.1] and “DJ3.3.4: Functional Specification and Design of a Generic Domain-centric Bandwidth on Demand Service Manager” [DJ3.3.4] serve as a good starting point to these subjects.

2 A Retrospective Look at GN2 Results

The work performed by JRA2 T2 in GN3 is a follow-up of the efforts made during the GN2 project. AutoBAHN is a result of research that was conducted to provide a proof of concept with minimum functionality and support for specific network technologies popular in NRENs. This goal was achieved, and in the GN3 project the AutoBAHN work is divided into SA2 and JRA2 activities, with the first one being responsible for introducing a production class service, based on research results of GN2 and the GN3 JRA2 activity.

At the end of the GN2 project, AutoBAHN was supporting three transport technologies, which are Ethernet, SDH, and MPLS. The deployment was performed in the following NRENs:

- GÉANT2 network – using the Alcatel ASN interface to mange the Synchronous Digital Hierarchy (SDH) network.
- HEAnet – using BlueNET as a Network Management System (NMS) to manage HEAnet's Multi-Protocol Label Switching (MPLS) enabled network.
- GRNET – using ASNTool as an NMS to manage the GRNET MPLS and Ethernet (VLANs) networks.
- PIONIER – configuring MPLS circuits in the Ethernet network.
- GARR – configuring Ethernet circuits (VLANs).
- CARNet – configuring Ethernet circuits (VLANs).

Except in HEAnet's case, AutoBAHN was operating in a testbed environment, isolated from the production network. In HEAnet, the BlueNET tool configured the production network, serving as an NMS to the AutoBAHN Technology Proxy. At the end of the GN2 project several issues remained unsolved, providing input for the definition of GN3's JRA2 T2. The implemented Authentication and Authorisation Infrastructure (AAI) mechanisms were limited to authenticate end users via a Graphical User Interface (GUI) at circuit request submission and tracking. The components of the AutoBAHN system and inter-domain communication were secured with simple X.509 certificates, not utilising the advantages of [eduGAIN]. Any authorisation or policy mechanisms, as well as accounting tools and procedures were missing. These are subject of the JRA2 T2 research, which will result in specifications and prototypes ready to be translated into production class services in the SA2 activity.

Within the lifetime of the GN2 project the stitching framework was defined, responsible for specifying and agreeing technical attributes for inter-domain circuits definitions. Although this framework is a key point for inter-domain links management by automated system, only a minor part of it was implemented in AutoBAHN at the end of GN2 project. GN3 is about to extend the stitching framework with new logic and attributes, and also to ensure that AutoBAHN implements proper mechanisms for automated links configuration.

Although the GN2 project has come to an end, international and inter-initiatives forums continue their collaborative work on generic definition and standardisation of inter-domain protocols. Since the start of GN3, the AutoBAHN team has been following the results of this, and will contribute to further work in this area with gained knowledge and experience.

New transport technologies are considered to have potential to be included in the AutoBAHN cloud. This applies to emerging technologies, like Optical Transport Network (OTN) or Carrier Grade Ethernet, and to technologies already known and extensively used, like Multi-Protocol Label Switching (MPLS) and Generalised Multi-Protocol Label Switching (GMPLS). GN3's JRA2 T2 is also open to interaction with "last mile" solutions, allowing extension of L1/L2 circuits to networks to be outside of AutoBAHN management, like L3 campus networks, etc.

3 Domain Network Technology Support in AutoBAHN

3.1 GMPLS

AutoBAHN can interact with a Generalised Multi-Protocol Label Switching (GMPLS) network using an overlay or a peer model [RFC3945]. In an overlay model, the level of visibility of the network is limited in terms of topology or resource usage, so different mechanisms should be implemented to support advanced functionalities over the GMPLS. However, a peer model would allow AutoBAHN to be part of the GMPLS Signalling Communication Network (SCN) and share the complete resource information by implementing an instance of the routing and signalling protocols at the AutoBAHN/GMPLS proxy.

This section raises several issues that need to be taken into account to enable the control of GMPLS networks under the AutoBAHN umbrella. Also, some considerations will be made focusing in the AutoBAHN-GMPLS proxy implementation.

3.1.1 Topology Abstraction

GMPLS uses a dynamic link-state IGP routing protocol to flood information about nodes and links in the network. This information is gathered by all nodes to build what is called the Traffic Engineering Database (TED). This database provides the topology information that includes state parameters such as availability or bandwidth utilisation. In GMPLS, each involved node deploys an instance of the routing protocol to exchange routing information. For instance, OSPF-TE uses Link State Advertisements (LSA) for this purpose.

In an Overlay model, topology information has to be manually configured in AutoBAHN and updated accordingly upon any topology change. Two possibilities arise:

- Edge + Border

  AutoBAHN is configured with the edge and border points of the GMPLS network. This way the network will be seen as a cloud with ingress and egress points. In this case, the GMPLS cloud is assumed to be fully meshed.

- Full topology

  AutoBAHN is configured with the full topology of the GMPLS network so it will have a complete view of all nodes and links.

In the first case, only data types for edge and border points of the GMPLS network have to be translated to the data model in AutoBAHN. In the second case, the full GMPLS topology has to be translated, as well as element relationships and network events. Consequently, improvements in the data modelling in AutoBAHN are envisioned.

In a Peer model, AutoBAHN, as part of the GMPLS SCN, has complete knowledge of the network topology and network changes are also dynamically updated. This information is stored in a global TED that is available in the AutoBAHN/GMPLS proxy.

The system architecture depends on the model chosen for implementation. Figure 3.1 shows a preliminary design of both models. The overlay model implements a simple standard interface (i.e. OIF UNI) towards the GMPLS control plane, but the simplicity of this interface entails further enhancements in the AutoBAHN system to support advance functionalities to a certain level. However, the peer model complexity relies on the deployment of the GMPLS protocol controllers within the AutoBAHN/GMPLS proxy, but in this case all native GMPLS supported functionalities can be used. In comparison to the peer case, implementing the overlay friendly TP requires much less work. At this stage of the project, both options are investigated in detail to provide a final solution, keeping in mind that the choice will influence a large part of the JRA2 T2 task. At the time of writing the overlay model appears to have more advantages, especially due to the amount of work required to create such a TP. The final decision will be taken at the beginning of November 2009.

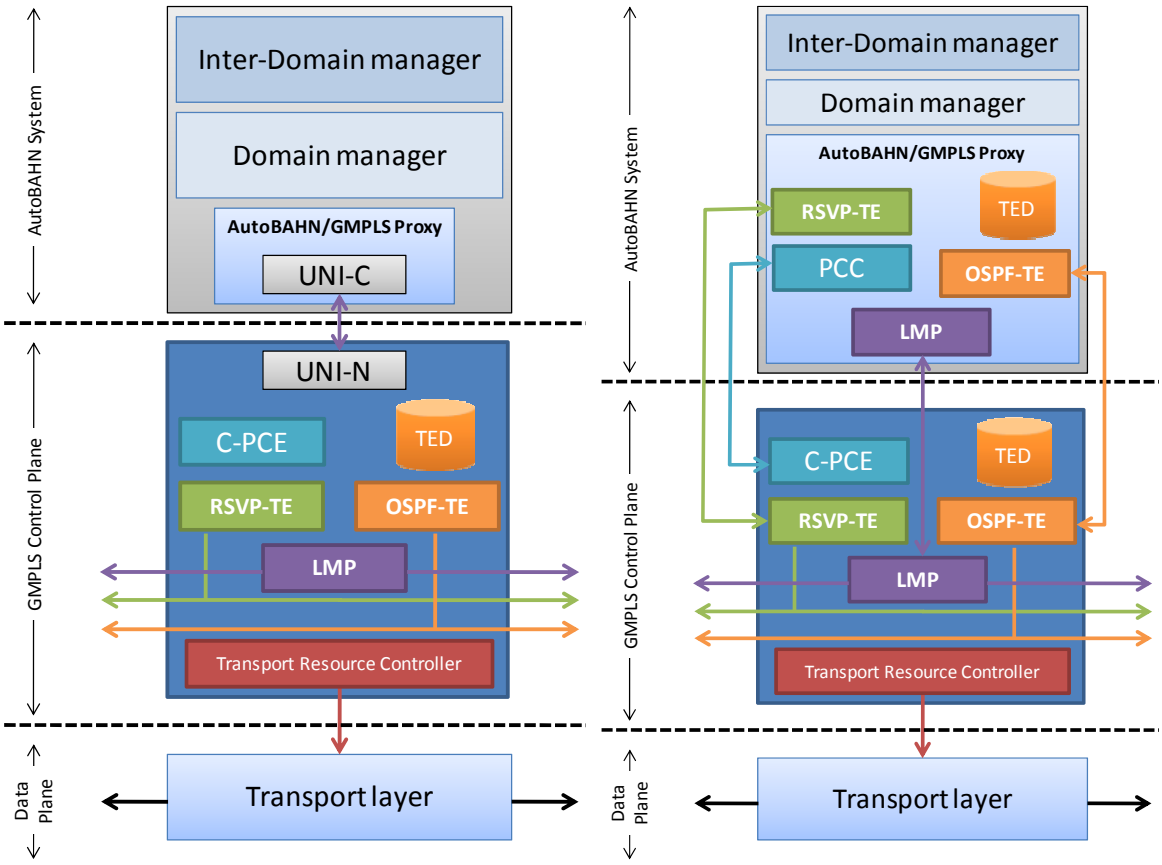

Figure 3.1: Overlay model architecture (left). Peer model architecture (right).

3.1.2 Intra-Domain Pathfinding

Path finding in an Automatically Switched Optical Network (ASON) is performed by the routing controller. In the Internet Engineering Task Force (IETF) a Path Computation Element (PCE) is responsible for implementing a routing (or RWA) algorithm and calculating the best route based on some input parameters (such as source, destination, bandwidth or other constraints [RFC4655]). The request arrives from the Path Computation Client (PCC), which is typically implemented in the Connection Controller (CC) and then processed at the PCE operating these parameters on the TED. PCC and PCE communication is performed using a standard protocol, PCEP.

One or more PCEs may be deployed in a GMPLS network but all instances need to access an updated TED with global network resource knowledge, which is fed by the IGP protocol. Depending on the architecture and the level of information available on each PCE, one or more PCEs may be involved in the path computation. The location of the PCEs should be known by its associated PCCs and this may or may not be in the head-end of the path, since the computation may need other PCEs to resolve parts of the path (e.g. loose hops).

In an Overlay model, path finding is carried out in a PCE within the GMPLS cloud, since OIF UNI does not allow the signalling of an explicit route. Thus, a connection request will just specify source and destination Transport Network Addresses (TNAs).

In a peer model, the AutoBAHN/GMPLS proxy could implement a PCC to communicate with the PCE available in the GMPLS network, or use the TED to perform the path finding itself. When the route is calculated, the proxy forwards the signalling request by means of the RSVP protocol, including the calculated path in the Explicit Route Object (ERO).

It is important to note that the technology proxy subtask will just focus on single technology domain management, so the support of inter-domain path computation between GMPLS domains is out of the scope of this task.

3.1.3 Domain Level Resiliency

The concept of network resilience refers to the actions taken to maintain the correct functioning of the network upon the detection of a failure. In GMPLS networks, this is achieved using recovery mechanisms that consist of localisation and isolation of the fault, followed by a repair and reconfiguration using survivability mechanisms (protection and restoration). In protection, the recovery paths are already reserved and provisioned before the failure occurs. In restoration, the recovery path is established after the failure detection.

GMPLS SCN can be deployed in-fibre or out-of-fibre. Failure in in-fibre control channels can be detected with the same mechanisms that the transport layer offers (e.g. loss of light). A transport plane failure in a network with an out-of-band control plane can be detected and repaired using other mechanisms, since the signalling communication continues upon the failure. The GMPLS Link Management Protocol (LMP) is intended to maintain node resources and provide other functionalities (such as control channel connectivity maintenance, link property correlation, link verification and fault isolation), as its implementation is of critical importance.

Restoration schemes have better capacity efficiency compared to protection schemes, but the implementation of a mesh restoration scheme is quite complex and requires sophisticated algorithms. Recovery can be classified typically by the segment to be recovered: span-level recovery, segment-level recovery and path-level recovery. Path restoration is better in terms of network resource usage, but is prone to experience more traffic loss until a new path is established. Also end-to-end restorations can be pre-planned to be dynamically allocated. However, restoration schemes are always slower than any kind of protection scheme, for which the 50ms bound can be guaranteed in nearly any kind of meshed network.

SCN recovery typically relies on the neighbour liveliness detection by the routing protocol instance running in the DCN. In case of OSPF IGP, this is done by exchanging “Hello” messages between adjacent nodes. The mechanism is the same applied in IP networks and implies a loss of CP communication for periods in the range of 30 seconds. This would be unacceptable in a GMPLS environment, in which the occurrence of a failure in the SCN could derive from a Transport Plane failure. Thus, in that timeframe the Control Plane could be in the maximum of its activity for connection recovery.

For this purpose, GMPLS Control Plane takes the benefit of the LMP functionality for Control Channel maintenance. LMP senses the status of its configured CCs and upon the occurrence of a CC down event, it associates a new CC (if any) to the affected links, thus negotiating what other CC(s) to use when sending CP traffic to an adjacent node.

Since SCN recovery is internal for a proper GMPLS operation, it is out of the scope of this task.

The recovery procedures supported by GMPLS via failure indication messages and switchover messages are described in Table 3.1.

Span protection	Unidirectional 1+1 dedicated protection
	Bi-directional 1+1 dedicated protection
	Shared M:N protection
Path protection	Unidirectional 1+1 protection
	Bi-directional 1+1 protection
	1:1 protection with extra traffic
	1:N (N>1) protection with extra traffic
Path restoration	Pre-planned rerouting with extra traffic
	Shared mesh restoration

Table 3.1: GMPLS recovery procedures.

For the interoperation of AutoBAHN and GMPLS, GMPLS data plane recovery procedures can be enabled in a peer model by using the corresponding RSVP-TE objects in the GMPLS UNI interface. In an Overlay model, if using OIF UNI, these procedures cannot be triggered or requested, thus, recovery mechanisms should be implemented in AutoBAHN.

3.1.4 Resource Management

Resource management in GMPLS is done by the control plane itself, which is responsible for selecting and allocating the best resources upon a request (with some constraints). Network resources can be discovered using OSPF-TE, but these resources can be correlated and verified within a GMPLS controlled network using LMP capabilities. GMPLS creates label switched paths (LSP) to be used as dynamic data-links, to deliver traffic between two adjacent nodes in a particular network layer.

In the same network static data links can also coexist, however, they are outside of the control of GMPLS. These resources are advertised in the form of TE Links, which are logical grouping of network resources to ease routing operations. Whereas paths are calculated based on TE Links, LSPs are provisioned over the actual data links forming those TE Links. Regarding the actual physical resources, GMPLS needs to develop specific interfaces depending on the transport devices to be configured. These interfaces map abstract configuration messages into specific commands that the device is able to understand using the available interface communication (SNMP, TL1, RS-232, etc.).

Resources in GMPLS are typically reserved and allocated at the same time that they are requested. Even if some studies have been carried out to support Advance Reservations in GMPLS networks [Phosphorus-G2MPLS] [ESCAL-COMNET], the GMPLS standards only consider bandwidth on demand requests (also known as immediate reservations). Other initiatives have tested and proved the feasibility and convenience of adding advance reservation capabilities to GMPLS [Harmony-AR], by means of creating a Thin Network Resource Provisioning System (Thin-NRPS) that implements the OIF UNI client (UNI-C) for GMPLS. The fact that the Thin-NRPS operates with the OIF UNI forces the interworking with AutoBAHN to the Overlay model, but it is a solid starting point to be considered.

When interfacing AutoBAHN and GMPLS, the resource management level is also dependent on the model chosen:

- Overlay

  In this model, GMPLS is seen as a cloud, so AutoBAHN has no direct access to network resources. Communication is done using the OIF UNI (switched connections), specifying the source and destination TNAs, bandwidth or directionality. Resources can be monitored and recorded if full topology information is configured in AutoBAHN. A PCC may be implemented in the AutoBAHN/GMPLS proxy to ask the PCE available in the network for a route between the TNAs specified in the OIF UNI request. This way, considering that AutoBAHN is the only “user” of the GMPLS network, the returned route will be the same and AutoBAHN will be able to keep track of the used resources at all times for enhanced resource management.

- Peer

  In this model, the AutoBAHN/GMPLS proxy is part of the GMPLS network. Resources are controlled directly from the proxy using the signalling protocol for connections establishment.

With Advance Reservations, since they are not supported by GMPLS, two implementation strategies can be considered:

- Advance reservations are implemented within AutoBAHN, using an enhanced Domain Manager or Technology Proxy.

- A new mediator for GMPLS is incorporated in the AutoBAHN architecture using Thin-NRPS as described in [Harmony-AR].

In any case, an entity within the architecture of AutoBAHN should keep track of the GMPLS resource calendar. In case of the peer model this can be achieved if AutoBAHN is the only “user” of the network; otherwise, an approximate approach should be used and resource availability cannot be 100% guaranteed which is not the best solution for AutoBAHN-GMPLS internetworking stability and resilience.

3.2 MPLS

Multi-Protocol Label Switching (MPLS) was developed as a packet-switching technology and has become popular in core IP networks. MPLS does not replace widely used IP-based routing, but works with existing and future routing techniques. MPLS adds new capabilities to services offered by IP networks, including Traffic Engineering (optimisation of traffic flow in network), Quality of Service (establishment of paths with guaranteed bandwidth) and Virtual Private Network (VPN, connecting private networks over the Internet).

Due to the features offered, the Traffic Engineering extension of MPLS is more relevant to the AutoBAHN service model and thus will be examined further in this section.

MPLS forwarding relies on labels rather than prefixes to route packets through the network. The labels are distributed over the MPLS domain using a label-distribution protocol. MPLS routers are called Label Switching Routers (LSR) because they operate on labels rather than on IP prefixes when forwarding packets.

The concatenation of the installed labels in the different LSRs is called a Label Switched Patch (LSP). An LSP is set up between the ingress LSR and the egress LSR. These edge LSRs are also called Label Edge Routers (LER). Packets that belong to a certain FEC are then mapped on an LSP. Determining the FEC of a packet is only necessary in the ingress of the LSP. Because LSPs are based on FEC-to-label binding, an LSP is a unidirectional path.

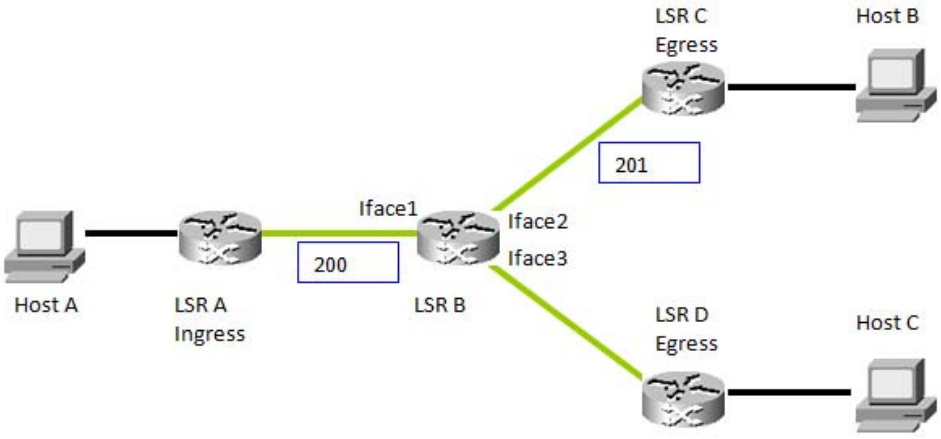

Incoming Label	Next Hop Label Forwarding Entry	Outcoming Label	Outcoming Interface
100	LSR A	101	Iface1
200	LSR C	201	Iface2
300	LSR D	301	Iface3

Figure 3.2: Label tagging in MPLS networks.

3.2.1 Topology Discovery

The creation of LSPs requires the forwarding tables at each LSR to be populated with the mappings `{incoming interface, label value}` to `{outgoing interface, label value}` as part of the label distribution process. The MPLS architecture does not require a single protocol for the distribution of labels between LSRs. In fact it specifically allows multiple different label distribution protocols for use in different scenarios, including:

- LDP
- CR-LDP
- RSVP-TE
- BGP4
- OSPF

Several different approaches to label distribution can be used depending on the requirements of the hardware that forms the MPLS network, and the administrative policies used on the network. The AutoBAHN system will need to retrieve the IP layer topology information and should be able to access the LSP information for its inter-domain pathfinding and request handling functions.

3.2.2 Intra-Domain Pathfinding

Pathfinding in MPLS networks allows nodes to choose routes based on traffic engineering constraints. It is assumed that information (like availability or bandwidth) can be distributed by the IGP throughout the area. MPLS Traffic Engineering can be used to find a suitable path that is:

- Intra-area
  The whole path is contained in a single area.
- Inter-area
  The path can lead through multiple areas belonging to one carrier.
- Inter-carrier
  Where the path runs across multiple carriers.

AutoBAHN will be focusing on Intra-area path selection. It is assumed that AutoBAHN can use information propagated by the IGP to obtain MPLS specific view of the underlying topology.

It is straightforward to choose a path for a LSP if the ingress and egress point are within the same area. The IGP floods information about the area's nodes and links to all nodes within this area. The IGP may spread the following information about each link:

- Available bandwidth.
- Administrative cost.
- Delay.
- Other properties of the link.

When LSPs are created or removed, traffic engineering properties associated with affected links (such as available bandwidth) will also change. The IGP continuously updates the link state information, thus allowing each node to build up a database of every link in its area. If a node receives a request to set up a LSP that must meet certain set of traffic engineering constraints, then this database can be used to calculate a complete route to the specified egress point for the LSP. The ingress point can also calculate a backup route at the same time to provide additional reliability for the requested link. The pathfinding process is therefore undertaken by the MPLS-TE signalling, and thus the AutoBAHN system is expected to rely on MPLS-TE for this. However, the AutoBAHN system will have to trigger LSP requests in order to discover intra-domain paths that are eligible for a user's request.

3.2.3 Domain Level Resiliency

Detecting and reporting failures in MPLS networks is realised in different ways:

- Power loss

  Usually means electrical outage and can be recognised by the line card attached to the device. The error message is propagated to the upstream and downstream node.

- Loss of light

  Optical link failure that can be discovered by the downstream node. Additional mechanisms are used to notify the upstream node about the failure.

- Link Management Protocol

  Link Management Protocol has been defined to monitor and discover links.

- Keep-alive

  Keep-alive is a mechanism where neighbour nodes poll each other at some constant interval to check availability of the neighbours. This method is considered slow in terms of failure detection.

- Topology updates

  Routing protocols within Autonomous Systems propagate current state of the network, including any broken links. This method is slower compared to hardware failure detection.

- Signalling notifications

  Signalling protocols specify mechanisms to report LSP creation failures as well as provide possibilities of notifying upstream nodes.

- Crankback

  Crankback extends properties of error and notification messages. With the underlying cause of the failure, additional information is provided such a broken link identifier.

3.2.4 Protection Switching

In protection switching, data flow is switched from a failed LSP (primary LSP) to a backup, pre-provisioned LSP. Usually switching takes place at the ingress node, although it is possible to mark another point on the LSP.

Several models describe protection switching in situations when the primary LSP is backed up by another pre-provisioned LSP:

- 1+1

  The backup LSP is already provisioned and ready for use. In this mode, data flows through primary and backup LSP at the same time. In case of failure, only the egress node is required to read data from the backup link instead of the primary LSP. This changeover could be triggered by a notification message or by comparison of the signal integrity on both LSPs. This resiliency model offers the fastest way of recovery but at the price of double network resources being used.

- 1:1

  The backup LSP is already provisioned and ready for use. This time, data is only flowing through the primary link, where the backup remains idle. When failure is detected, the ingress must be notified, so it can switch the data flow. It is possible that the backup link can be used for low priority traffic so that when the primary LSP fails it can be discarded.

- 1:N

  The backup LSP is already provisioned and ready for use. In this mode the backup link supplies multiple primary LSPs. If one of the primary LSPs fails, flow is switched to the backup LSP and the rest of LSPs are left without protection. In practice, failure of more primary LSPs is very unlikely and this type of protection is often considered sufficient.

3.2.5 Local Switching

The communication of signalling information in MPLS uses IP. When a failure occurs in the network, an LSR upstream of the failure can attempt to re-signal the LSP. LSP signalling relies on IP routing, and therefore can take advantage of the fact that the routing table may be updated with new routes to the downstream nodes. However, this may take many seconds, and will not necessarily result in a new route being available.

Data is forwarded based on the {incoming interface, incoming label} to {outgoing interface, outgoing label} mappings, and not information in the IP routing table. Therefore, updates to the IP routing table do not affect the data flows. Data paths can only change once a new LSP has been signalled and devices on the LSP programmed with the new label mappings.

3.2.6 Fast Re-route

Fast re-route is a process where MPLS data can be directed around a link failure without the need to perform any signalling at the time that the failure is detected. Unlike protection switching, the repair point is the point of failure detection. Consequently there is no requirement to propagate the error to the repair point using the signalling protocol.

Most fast re-route protection schemes rely on pre-signalled backup resources. When the failure is reported to the repair point, it simply updates the programming of its switch so that data that was previously sent out of one interface with one label is sent out of a different interface with another label.

3.2.7 Link Protection

The simplest form of fast re-route is link protection. An LSP tunnel is set up through the network to provide a backup for a vulnerable physical link. The LSP provides a parallel virtual link. The capacity of the backup LSP should, of course, be sufficient to carry the protected LSPs. If all LSPs on a link are to be protected, then the capacity should equal the bandwidth of the protected link. This can potentially lead to a huge amount of backup bandwidth being required, especially if multiple links must be protected in this way.

Note that not all LSPs using a link need to be protected by the same backup LSP, or even at all. By leaving some LSPs over the link unprotected, the backup bandwidth requirement can be reduced.

3.2.8 Node Protection

Link protection only handles the case where a single link between two LSRs has failed. However, it is also possible that an entire LSR will fail.

When the failure is detected, the traffic is re-routed down the tunnel and data continues to flow. Using additions to RSVP-TE, the initial setup of the protected LSP reports the labels in use on each link as part of the Record Route object. This object contains a list of LSR IDs and labels describing each hop. It is passed upstream during LSP establishment (on the Resv message) so that every node on the path knows the labels used on every link.

Once this information has been passed back upstream, each LSR can determine the correct labels to use in the label stack when it re-routes an LSP after failure.

3.2.9 Summary of Protection Methods

Each protection method has its advantages and disadvantages. Table 3.2 summarises protection features:

Protection method	Network resources	Recovery time	Configuration
Local protection.	Recovered LSP takes same amount of resources.	Tends to be slow.	Signalling remains the same. New path may not be optimal in terms of length.
Protection switching.	Backup LSP is pre-provisioned. In some cases may be used by other primary LSPs.	Usually information about failure must be delivered to the ingress point.	Signalling remains the same.
Fast Re-route link protection.	Backup link is required for each protected link.	Recovery process starts as soon as failure is detected.	Link backups have to be preconfigured.
Fast Re-route node protection.	Backup node is required for each protected node.	Recovery process starts as soon as failure is detected.	Node backups have to be preconfigured.

Table 3.2: Summary of protection methods.

The AutoBAHN system will have to leverage the MPLS protection mechanisms to provide protected circuits over an MPLS-TE domain.

3.2.10 Resource Management

It is important to know what functionality of MPLS is available for management purposes. From AutoBAHN's point of view an MPLS device must provide the following functionalities:

- Access to the underlying topology.
  AutoBAHN's internal view of topology must be synchronised with the actual MPLS network.
- Create LSP.
  AutoBAHN must be able to specify a reservation path, with additional parameters such as required bandwidth, delay and resiliency mode.

- Remove LSP.

  AutoBAHN must be able to turn down previously created path.

- Query available LSPs.

  AutoBAHN must be able to list of all paths that have been created.

3.2.11 High Level Design

Depending on the network model, the LSRs in the MPLS may have different information about the current status of the entire network. The network model can be:

- Overlay

  Each LSR holds only IP/MPLS layer information. In this model, the MPLS network will be seen as a cloud by AutoBAHN.

- Peer

  Each LSR holds information about the topology and status of physical links. If this model is chosen, AutoBAHN will have a full view of the underlying network topology.

3.3 Carrier Grade Ethernet

Ethernet technology provides high bandwidth at a relatively low cost, easy installation, and capability of point-to-point and point-to-multipoint operation. These properties made Ethernet emerge as a WAN protocol even in provider backbones.

Carrier Grade Ethernet is defined by MEF as “a ubiquitous, standardised, carrier-class Service and Network defined by five attributes that distinguish it from familiar LAN based Ethernet” [MEF]. These attributes are:

- Standardised Services
- Scalability
- Reliability
- Service Management
- Quality of Service

The basic technologies defined by IEEE that provide essential support for backbone deployment of Ethernet are:

- Link Layer Discovery Protocol – IEEE802.1AB.
- Provider Bridges / Q-in-Q – IEEE802.1ad.
- Provider Backbone Bridging (PBB) / Mac-in-Mac - IEEE802.1ah.
- Provider Backbone Bridge Traffic Engineering (PBB-TE) - IEEE802.1Qay.

- Ethernet Operations, Administration and Maintenance (OAM) – IEEE802.1ag (also published as ITU Y.1731).

LLDP has little effect on AutoBAHN, as its main purpose is to discover the elements and the topology of an Ethernet-based network. The discovered topology will have to be maintained by AutoBAHN, managed and abstracted. Already at this point, LLDP is used by cNIS for automated discovery of native Ethernet topologies.

Q-in-Q and Mac-in-Mac are essentially encapsulating technologies that enable customer VLANs (or entire MAC addresses) to be transparently handled by the backbone network, achieving scalability and manageability.

802.1Qay, whose standardisation and ratification within the IEEE is currently being finalised, is the amendment of 802.1. This is probably most relevant to the AutoBAHN service model as it allows explicitly selected traffic-engineered paths within Provider Backbone Bridge Networks. It is expected that functionality adhering to 802.1Qay will be offering explicitly routed, bandwidth guaranteed paths over a PBB core, so the corresponding AutoBAHN technology proxy will have to be developed accordingly.

802.1ag provides operation, administration and maintenance functions. These are not directly relevant to the dynamic circuit reservation and provisioning functionality required by AutoBAHN. However, it is expected that such functions will be essential for the monitoring of the performance and troubleshooting of AutoBAHN-provisioned circuits, and thus exploited accordingly by advanced releases of the AutoBAHN CGE technology proxy.

Carrier Ethernet definition by the MEF specifies two service types (E-Line and E-LAN), which can be supported in port mode (as EPL and EP-LAN) or VLAN-multiplexed mode at the UNI (EVPL and EVP-LAN). In MEF 6.1 an additional service type (E-Tree) has been defined, with corresponding EP-Tree and EVP-Tree modes. Services are implemented using EVC (Ethernet Virtual Connections).

Since AutoBAHN deals with point-to-point connections, the focus is on the support of E-Line service types; EPL (EWS in Cisco terminology) and EVPL (ERS in Cisco terminology) services.

The level of implementation of the AutoBAHN Carrier Ethernet technology proxy will depend on the underlying implementations of the aforementioned standards. It is expected that PBB/PBB-TE functionality will, in early phases, be supported only at the management plane. Therefore the specification of a proxy will be heavily reliant on individual implementations rather than a standardised control plane.

Preliminary research has shown that part of the necessary functionality (such as intra-domain pathfinding) will be undertaken by the CGE management plane. In such cases, the AutoBAHN TP proxy will be accommodating the information flow between the management plane and the AutoBAHN system, and the AutoBAHN DM will have to suspend its pathfinding functionality. Some efforts to provide GMPLS control of Ethernet PBB-TE are underway in the IETF ([GMPLS-PBBTE01]) and will be examined based on their acceptance by the industry.

At the management level, Ethernet OAM functions will be combined with monitoring and diagnostic information functionality in later stages. In more detail:

- Topology abstraction

  It is expected that the AutoBAHN TP will be interacting with the CGE management plane in overlay mode, so that topology information will be already condensed when reaching the AutoBAHN DM module. Abstraction to a technology-agnostic topology will be undertaken by the AutoBAHN DM module. However, at the technology-specific level, for inter-domain stitching purposes, the requirement will be to ensure that cNIS can sufficiently model Q-in-Q and Mac-in-Mac capabilities, where Ethernet VLAN and MAC spaces respectively are stacked in layers.

- Topology discovery

  Although not a core function of the AutoBAHN system but an essential one for its operation, topology discovery is envisaged to be possible by exploiting protocols such as IEEE 802.1ab (LLDP) Ethernet OAM functions. The level of support of topology discovery functionality by CGE management systems is to be investigated.

- Intra-domain pathfinding

  In PBB, and especially in PBB-TE specifications, Ethernet self-learning protocols (such as MAC learning) and self-control protocols (such as STP) are disabled, and the management plane is responsible for configuring frame forwarding on the data plane. Thus, intra-domain pathfinding is expected to be driven by the management functionality offered. The AutoBAHN technology proxy will have to enable information flow on intra-domain paths between the management plane and the AutoBAHN DM. Again, control plane options for intra-domain pathfinding seem to be less possible in a short time frame.

- Domain-level resiliency-protection

  CGE Ethernet (and more specifically PBB-TE) cater for resiliency by supporting protection paths for configured service instances, based on 802.1ag (CFM) signalling a working path failure with failover times in the order of 50ms. This functionality will have to be exploited by the AutoBAHN CGE technology proxy, when requests for protected circuits are submitted to the AutoBAHN system.

- Resources management

  Advance reservation management is currently not supported in CGE solutions, although in some cases it is an area of interest. Therefore, advance reservations will be based on the calendar functions of the AutoBAHN DM. Signalling of circuits at the moment of enabling of advanced reservation will be based on messaging between the DM, the CGE technology proxy and the underlying management plane.

3.4 OTN

With the intention of providing solid management features to support the high bandwidth capabilities of DWDM networks, the Optical Transport Network (OTN) standard facilitates the benefits of SDH/SONET management, while additionally bringing other advantages to network operators like transport of transparent mix of traffic protocols:

- Backward compatibility of protocols.
- Enhanced maintenance capabilities.
- Forward error correction (FEC).

Although the OTN standard is available and has been accepted in the industry for some time, the functionalities and features that it provides have not yet been fully utilised within the European NREN communities. There are not many different client signal interfaces available in vendor portfolios. Therefore, the client traffic is typically transported on a line terminal basis, and thus any multiplexing or cross connections are done on the higher layers. Furthermore, there are currently no management capabilities at pure optical level on the OTN.

3.4.1 OTN Standards

The two most basic defining standards of OTN are G.709 and G.872. G.709 defines the interfaces and the rates, while G.872 describes the architecture of OTN.

Currently there are no standards defined for inter-domain communication, and also vendors are taking slightly different directions in implementing their products. This proprietary approach limits the inter-domain compatibilities, and therefore it's not possible to connect optical transport sections (OTS) from two separate domains without considering connection parameters (for example, line speeds or modulation schemes). As a result of this, the flexibility and transparency of the OTN technology is somewhat unexploited.

Therefore the development of a "domain network technology support" for enhancement of the AutoBAHN system will begin with studies on which operators or vendors are in fact employing the multiplexing hierarchy of OTN today, and perhaps use any knowledge they might have with the other functionalities it brings (such as enhanced E2E management, or APS protection mechanisms).

Based on the above studies, (as an enhancement to the AutoBAHN system), an OTN topology abstraction tool shall be developed. As there are different approaches to implementing this, a vendor-independent study of possible solutions shall be performed. It must be focused on finding optimal abstraction algorithms and the best suited topology representation for OTN. Also it must be investigated whether the existing SDH/SONET topology abstraction developments from AutoBAHN system can be reused in order to take advantage of the substantial similarities between SDH/SONET and OTN.

In addition, it will be determined whether the intra-domain routing and path finding functionalities in SDH/SONET management platforms can be reused in the development of the OTN topology abstraction. It must also be decided whether intra-domain path finding should be performed by the management locally in the domain or by the AutoBAHN domain manager.

This could be implemented if the AutoBAHN domain manager is only administrating border nodes at the Inter Domain Interface (IrDI). And the local management system is administrating the Intra Domain Interfaces (IaDI). The principle is shown in Figure 3.3.

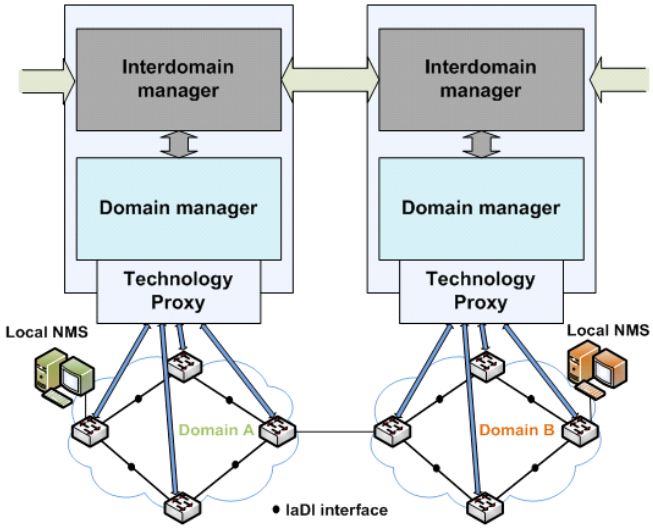

Figure 3.3: Local management of IaDI interfaces.

Figure 3.3 illustrates the situation where routing and path finding functions are handled by the local NMS by managing the Intra domain interfaces, and have the domain manager handling the nodes at the border of the domain and their Inter domain interfaces (IrDI), shown in Figure 3.4.

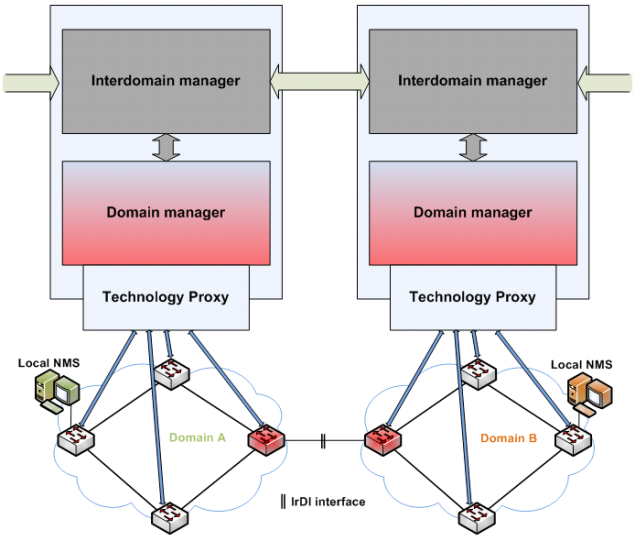

Figure 3.4:. DM handling the IrDI interface.

Alternatively, all routing and path finding functions could be handled by the domain manager, (both IrDI and IaDI interfaces). In this case, no routing and path finding are done by the local NMS. (Figure 3.5).

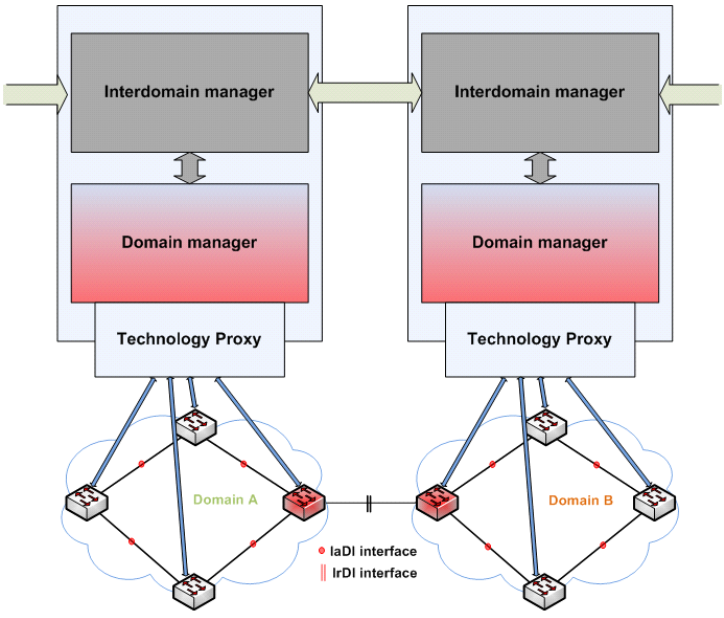

Figure 3.5:. DM handling the IrDI interface.

Another issue is to make use of the existing protection technologies of OTN and determine whether it can be integrated into the AutoBAHN system and so add inter-domain resilience. Therefore, a study must be performed to determine whether these functions should be performed by local the NMS or the AutoBAHN system.

Part of determining the topology abstraction methods should be to investigate and select which interfaces and protocols should be used for the communication between the DM and the network. (e.g. TL1 or ISN).

The above pre studies and decisions will lead to a more detailed plan for the actual design of the OTN domain network topology support.

4 Inter-Domain Stitching Issues

4.1 Technology Stitching

Technology stitching is required because NRENs implement different network technologies. Stitching provides a homogeneous method to interconnect domains. Stitching technology refers to the elements involved in the junction or interconnection between two different domains. These domains may or may not use the same network technology. Certain parameters need to be exchanged to help determine the best or most successful way to establish the required connection.

The focus on stitching during GN2 was on explicitly describing the process with sufficient detail that will allow stitching to be automatically configured (as described in the deliverable "DJ3.5.3: Report on Testing of Technology Switching" [DJ3.5.3]). Technologies addressed were: Ethernet, SDH/SONET, L2 MPLS VLL, IP MPLS QoS and PIP (Premium IP) services.

The proposed developments in GN3 focus on new technologies that are starting to be deployed in NRENs (such as Carrier Ethernet, OTN, Dynamic Lambda Provisioning, and GMPLS). Data representation models may be produced in the most suitable format structure to help Technology Proxy support these technologies.

In the GN2 version of the data representation model, a consistent structural formatting style of **Attribute**, **Parameter Names** and *attribute values* was used. This format and naming structure may be revised by Task 2 of the GN3 JRA2 activity, as a result of agreements from standardisation work in the NREN community.

The stitching framework will only identify the parameters and those associated attributes that allow stitching between peering domains using the same technology. A peering domain is the next domain where a certain protocol layer is terminated. New technology stitching attributes and parameters will be identified for these technologies for terminating domains that will use or be linked by these technologies as Peering Domains. The domain types are:

- Termination domain

  Where the path begins or ends for a Bandwidth on Demand link. At least two termination domains must exist in each path. A termination domain can be a single workstation or a full-blown network (with demand for multiple paths). Normally a termination domain would be called a user domain. This is an essential domain for the stitching framework.

- Linking domain

  A linking domain has known interfaces/protocols/processes and is one of the domains that make up the path between two termination domains. All linking domains together will form the established AutoBAHN path. This is another essential domain for the stitching framework.

It is important to bear in mind that a linking domain is autonomous, and is not aware of the internal details of its neighbouring domains (except their external connections for which to peer with). For stitching, some information about the linking domain is essential to create a path between the two user domains. The Linking domain has to ensure that its interfaces work with their peering interfaces, and also that the routing/switching is done properly inside its own domain: Stitching needs to ensure that the participating linking domains will configure their internal routing/switching, and adapt in such a way that the two logical interfaces are linked. This stitching must occur at each logical interface level (a grouping of parameters relevant to that technology interface) which is applicable for that domain. Internal domain technology implementations have to be catered for by the local domain manager, which may be a specific type of technology proxy.

Technology stitching of a circuit is achieved using the following systems:

- The Inter-Domain Manager (IDM) determines the best path between domains and relays requests either in-band or out-of-band.

- The Domain Manager (DM) configures local network domain equipment (for instance through the provisioning system). The DM evaluates all its own domain and technology specific parameters

The IDM handles inter-domain communication while the DM makes sure that all domain specific information is translated in an abstracted representation. The IDM can send messages to its direct neighbours only, so that only adjacent domains can communicate. This implies a chain communication model, which is enforced by the AutoBAHN application logic, Stitching framework and reservation processing which are all components of the IDM.

A data model has been created to map out parameters, attributes, peering relationships with each domain and how it may stitch linking domains together to form a path. In Figure 4.1, the user representing the termination domains asks for a path, through its home domain. The home domain uses technology abstraction and pathfinding (two other components of the IDM) to calculate the possible paths between the termination domains. For each path a feasibility check will be done, based on the parameter values, to check if actually a path can be made. This is essentially a constraint analysis based on reservation scheduling matching and technology stitching carried out in each IDM. A required path, by definition, should go from the source to the destination domain.

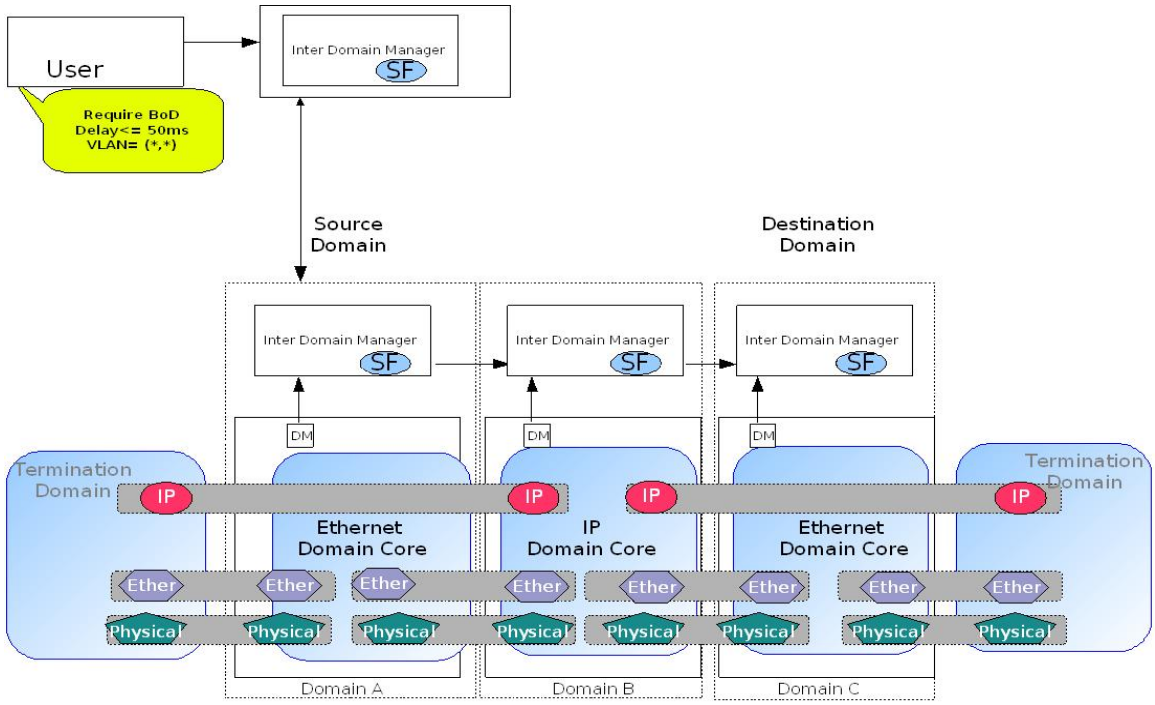

Figure 4.1: Initial BoD reservation.

If the DM of domain A has available resources to perform the reservation, the IDM of domain A forwards possible parameter values to domain B's IDM. Domain B performs a local resources check too, and defines its own constraints for the reservation. If it is possible to accept the reservation's constraints and technology, and stitch a path to the C domain, the request and constraints are forwarded to domain C's IDM.

The user, who wants to create a BoD path between the two termination domains, can identify specific requirements as a part of the BoD request. In this example, the user specified a Delay value lower then 50 msec for the whole path and does not care which VLAN ID will be issued on both sides (see Figure 4.2). The BoD request from the user will be forwarded by the source domain's IDM (domain A) towards the destination domain's IDM (domain C). This information will be used in the parameter constraint analysis for reservation scheduling and technology stitching in the destination domain's IDM.

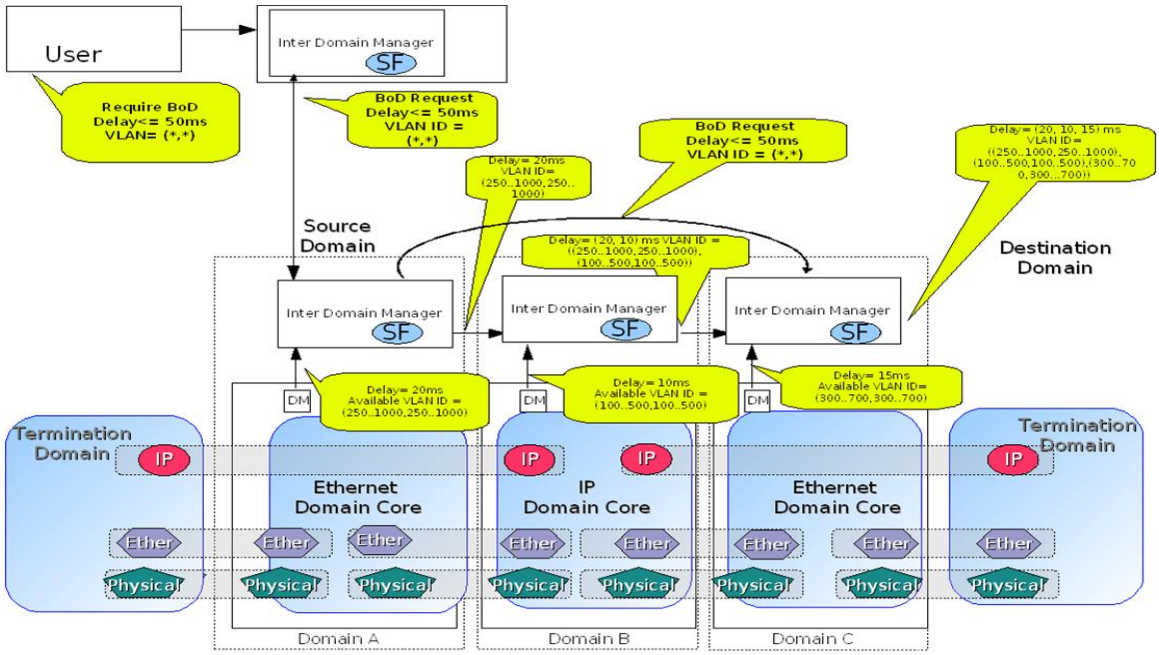

Figure 4.2: Initial request forwarded to Neighbouring domains (parameter concatenation).

Domain C's DM also performs a local resources check to define its local domain parameter values. All parameters are now available to the IDM of the destination domain (domain C), as shown in Figure 4.3. The IDM will do a full parameter/constraint analysis for stitching and reservation scheduling. If this evaluation still allows multiple parameter values for the path, a random choice will be made by the IDM of the source domain.

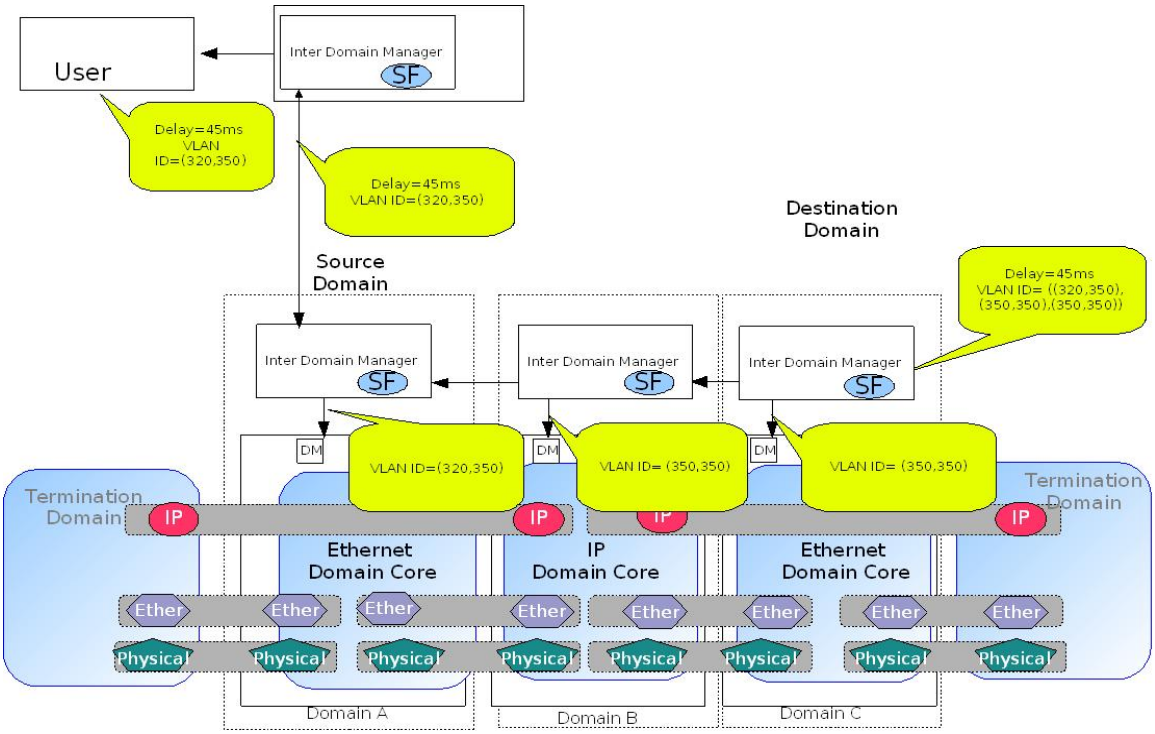

Figure 4.3: Agreeing the Reservation schedule.

The parameter/constraint analysis in this example looks at the total ***Delay*** along the path; it must not be higher than requested by the user and this condition is met. As all domains should agree available ***VLANIDs*** for Ethernet domains, the common random values (like 320 and 350) are selected for the Termination domains from the possible set (300...500). It has been possible to stitch the parameters and domains to complete the path.

Since all the parameters now have values, and all domains have resources to schedule matching the ***Schedule*** parameter values, the preparation for the final schedule can now be started (Figure 4.4). Domain C's DM books the local domain resources and schedules the reservation for the specified time to execution. The response now requires that the requisite parameters are forwarded to domain B and then to domain A. All domains must schedule the reservation with specified parameters defined by the destination domain. The user is now informed. Should any domain not have been able to stitch together the appropriate resources, a suitable error message would have to be returned to the source domain and the user.

Extra parameters and attributes can quickly be added to the model. In the example above, ***Delay*** (maximum value specified by the user) is identified as being a critical requirement.

Peering dependencies, such as sources of intervention to setup the path and how domains can influence parameter attribute values are mapped by the current data model to allow for the future development of an automated database solution to be implemented as part of GN3 T2.

The stitching framework and data model approach shows which functions will have to be implemented in the logic of resource scheduling and stitching to evaluate this parameter. The current version of the stitching framework specified that parameters are best viewed as an instantiation of the object PARAMETER, which has certain attributes and methods attached, to help it align with object-oriented data modelling tools.

PARAMETER as an object has the following attributes:

- **PathDomain**

  The Terminating or Linking domain that the PARAMETER belongs to.

- **Involved**

  Shows what part of the domain is involved in determining the value of a parameter (e.g. (***InterfaceSpeed*** is determined by the egress/ingress interface itself), while the whole domain (section) will assign a value to a parameter in case (e.g. contact info for the ***NOCInformation***).

- **Name**

  The name of the parameter. For example, ***Bandwidth***. Each **Name** plus **Type** have to be unique in the data model.

- **Data**

  The actual value of the parameter. For example, the ***Bandwidth*** is *10000000*.

- **Dimension**

  Specifies the units of the **Data** attribute. This can be a string, a number, vector, etc. For example, e.g. ***Bandwidth*** is in `bit/s`.

- **Type**

  Indicates which kind of logical interface this parameter belongs to. This clusters parameters together, helping to generate categories for parameters for different logical interfaces and domain cores types. For instance ***MACAddress*** is categorised with a **Type** of *ethernet*.

- **Dependency**

  A parameter's **Data** can have a **Dependency** based on the value of another parameter (such as *full/half* ***Duplex*** in ethernet, which depends on ***AutoNegotiate*** and ***InterfaceSpeed***).

- **Logic**

  Logic defines how the **Data** value of a parameter is determined. For example, some values are dependent on the peering domain (like ***IPAddress***es), some parameters need to be *contiguous* (the same) in a domain (e.g. ***VLANID***s values sometimes need to be *contiguous*). In other cases, for **PatchPanel** the values have to be *non-contiguous* (different). Some parameter's **Data** value may need to be determined by all domains in a *path* (such as ***Bandwidth*** and ***MDS***).

- **Method**

  The logic method used to compare parameter values between or across peering domains (defined by the **Logic** attribute value). e.g. ***IPAddress***es must be different, while ***IPSubnetMask*** must be the same. In some cases the Method of logic is a complex function. For example, to determine the ***Bandwidth*** value a domain needs to know the ***InterfaceSpeed***, ***VLANTagging***, etc.

- **Intervention**

  Intervention types, for example *human* action (physical, like ***Cables***, etc.), *remote* configuration (typing in ***VLANID***s, ***IPAddress***es, configure routing/switching etc.) or automated processes (***AutoNegotiate*** in ethernet and ***DHCP*** in IP, etc.) record the type of mechanism used to check and match parameters required for stitching at a particular logical level.

- **Default**

  If a peer of a parameter is not specified, a default value can be assigned to this missing parameter in the other peering domain. In that situation, any feature check can implement a total check on values selected and how eligible are they for the situation. For example, if a domain does not specify any ***VLANID***s, another domain can use the default range (which is e.g. 0:4096).

- **Proposed**

  In case there are multiple possible solutions for the path, the value for **Data** for a particular parameter maybe proposed by the stitching engine.

Several attribute values have fixed value settings (which do not change per technology domain, like **Dimension**, **Type**, **Involved**, **Method** and **Default**) for a particular type of parameter. Other parameter attributes might change their value depending on the technology domain's properties.

Logic functions can perform any necessary functional checks within a domain. Examples are:

- In some domains a parameter may require that the other domains (like peering domains) take part in some of the decision making to determine the **Data** value of a parameter. (e.g. *VLAN IDs* entered in ***VLANID***s sometimes need to be the same (contiguous) on both interfaces of a technology domain).

- Some parameters' data values may need to be determined and reached by consensus by all domains in a *path* (such as ***Bandwidth*** and Maximum Datagram Size ***MDS***).

A prototype model has been created to show the computational logic and interactions, between the parameters as they are entered or calculated across linking domains, trying to establish a Bandwidth on Demand Path. The value of the model is that it can be expanded or reduced in size by adding or reducing the number of parameters per logical interface. It can help designers model the interaction of particular parameters and their logical attributes so that they can simplify what is the level of stitching that must be implemented between the Inter-Domain Managers and the Pathfinding Engine.

It is expected that this data modelling structure will change to align with work being carried out by other groups (such as Network Service Interface (NIS-WG)) when that data structure has been agreed. It is forecasted that most of the conversion exercise will take the form of nomenclature changes and that most relationships attributes will remain quite similar.

The stitching framework model currently handles L2 MPLS VLL and also IP QoS-enabled domains for providing a premium IP service. For Carrier Grade Ethernet the current data model can quickly be modified to handle some of the extra features required, such as stitching functions to support the Mac-in-Mac addressing - IEEE802.1ah for the Provider Backbone Bridging (PBB) option.

There are difficulties faced in stitching some of the additional technologies, such as GMPLS-enabled domains. The intra-domain control plane only keeps track of LSPs operational status, routing paths, and traffic engineering and it has no full exposure to the actual peering interfaces involved in stitching the domains together. Interworking in the GMPLS control plane is very complicated, as different suites of protocols are used in separate networks such as routing, private network-to-network interface versus OSPF–TE used.

GMPLS switching can be packet-based, TDM-based, wavelength-based, waveband-based, or fibre-based. Therefore, different stitching technology data models may have to be constructed depending on the particular technology used. GMPLS can operate at peer-to-peer level or as an overlay technology similar to AutoBAHN. Stitching these technologies at their peering interfaces and controlling that this peering is appropriate and allowed becomes a greater challenge for the IDM and DM interactions between linking domains.

4.2 Inter-Domain Links Discovery

The AutoBAHN system needs information about the topology of the network it uses. There are two views of the topology available in the system: the technology specific intra-domain topology and abstract inter-domain topology. Abstract topology is computed in the conversion process from the technology specific one. It hides some details (such as VLAN identifiers, STM link types, MTU, and so on). Usually the network administrators have the detailed information about the network topology of the domain they manage. Information about the topology must be inserted manually into the Domain Manager. A user can enter it directly to the database using a cNIS web-based front-end. The intra-domain information for a domain should contain technology-specific data describing the devices inside the domain as well as information on edge interfaces, links to adjacent domains and remote interfaces.

The problem is to match the edge interface from the local domain with the correct interface belonging to a neighbouring domain. When this information is completed in both domains, the same link is correctly represented at both sides, allowing the identification of a common abstract identifier during the topology abstraction process. Pairing the interfaces belonging to the different domains could be performed in several ways. In the current implementation of the system the manual approach is used.

4.2.1 Manual Discovery

This concept assumes that the administrators of adjacent domains communicate with each other through a not AutoBAHN specific channel (such as email, phone or an instant messenger). The information that interests an administrator is the public name of the interface that is connected to our domain. The problem becomes complicated when there is more than one link between a pair of domains. Then administrators should check which interface is connected to particular remote interface using some kind of connectivity experiment.

When the matching process is completed, both administrators update their intra-domain topology databases with the correct port identifier, meaning that the topology information is complete. Domain manager can now proceed with the abstraction process and inter-domain topology will be calculated properly.

4.2.2 Semi-Automatic Discovery

The concept of automating the process of inter-domain links discovery relies on registering the edge interfaces of a domain in a repository. Using this repository, other domains can search for the matching entry. The key is the pair of domain identifiers. During the lifetime of the GN3 project the functionality of the Lookup service in AutoBAHN will be expanded. One of the probable cases is the use of the Lookup module as a repository for storing the information on the edge interfaces. Note that this solution requires the domain identifier to be unique for the whole AutoBAHN instance.

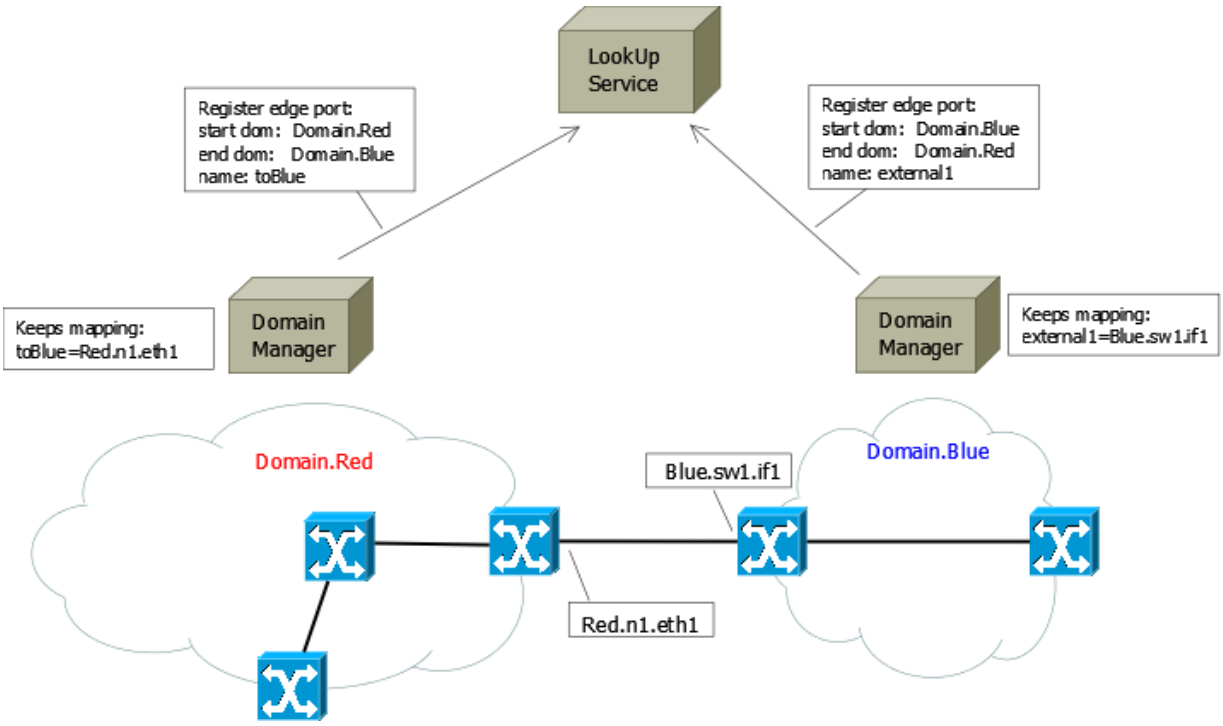

Figure 4.4: Usage of Lookup Service to register edge ports.

Figure 4.4 shows the process of registering the public name of the edge interfaces with the pair of domains between which it terminates a connection. Note that the registered identifier by the domains is not an actual label used in the domains (this is private information). Domain manager keeps the mapping between public and real name.

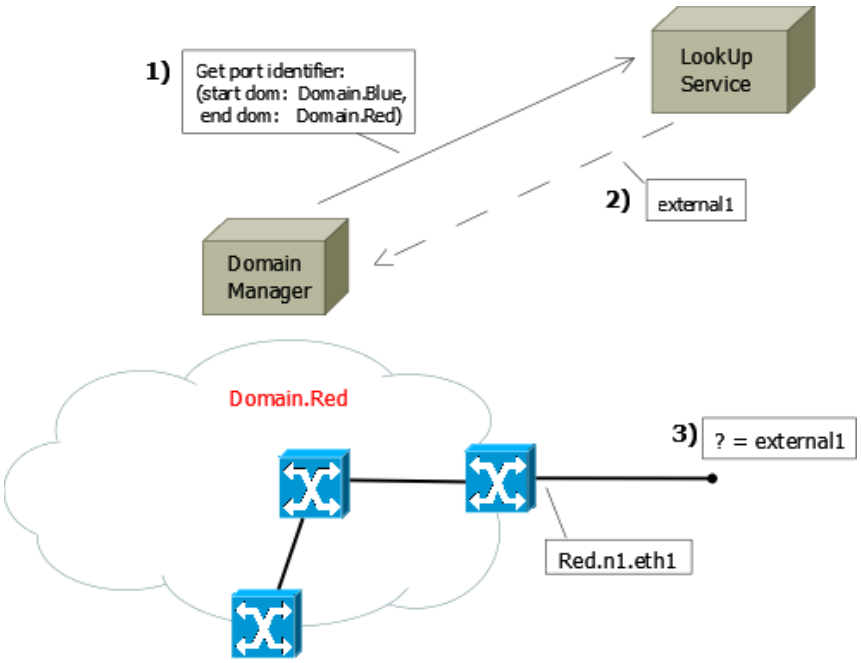

Figure 4.5: Getting information about the external interface connected to our domain.

Figure 4.5 shows the process of getting the information from the Lookup Service about the remote interfaces that are used to connect two particular domains (Domain.Blue and Domain.Red). An analogical process is performed in Domain.Blue. After this operation both domains have full view of the topology and are able to perform abstraction properly. When there is only a single link between two domains, the process is automated and no manual intervention is required. Fetching the information from LS can be implemented both in the Domain Manager and directly in cNIS (which allows you to view the remote interface in the GUI when inserting the topology). However, this introduces a dependency on an external source of data, and this is not envisioned in the forthcoming releases of cNIS.

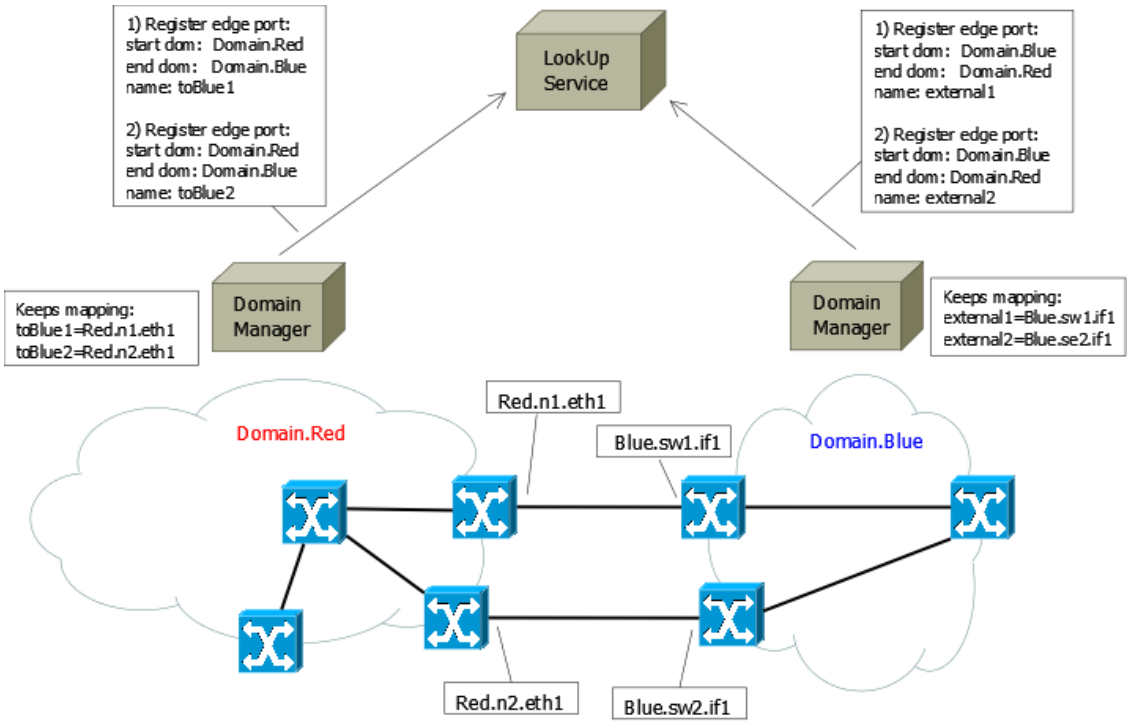

Figure 4.6: Registering edge ports in case of multi-homing.

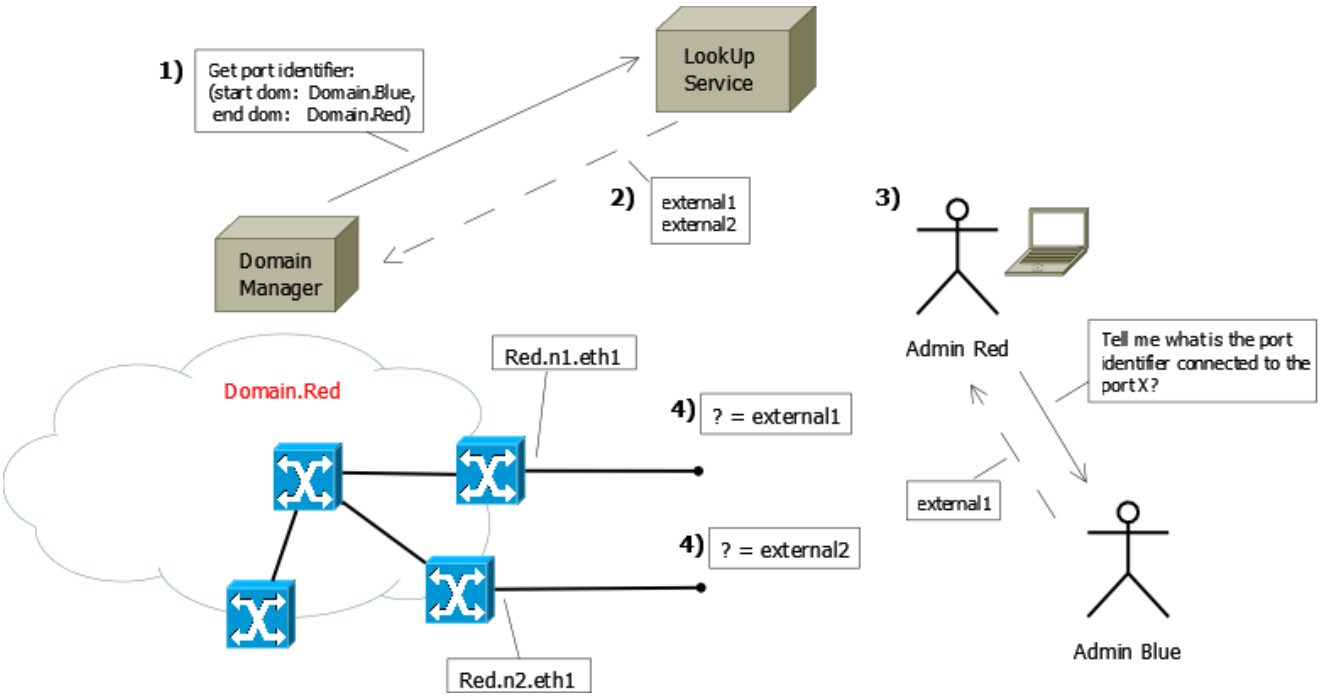

Figure 4.7: Getting identifiers from LS in case of multi-homing.

The situation is far more complicated when there are multiple links between two domains (multi-homing). Obviously the pair of domain identifiers is not sufficient to cover this situation. In this case more than one edge interface identifier was returned. This requires an administrator's intervention to determine which remote interface is connected to which local interface. This solution can only provide a list of possible remote interfaces and the administrator should do the correct matching.

4.2.3 Automatic Discovery

There are some approaches to the automatic topology discovery inside an administrative domain proven to be working, though discovering inter-domain links is often limited by security issues. Some proprietary discovery protocols (like the Cisco Discovery Protocol (CDP)) rely on browsing the SNMP databases of network devices to match certain attributes to verify that there is a connection between them. This approach is limited to the CDP enabled devices, and authorisation for reading the SNMP database of a device from an external domain is required. However, if the CDP is properly configured and the traffic is not blocked on the edge of a domain, some information about an external device (such as IP address and MAC address) could be available in the MIB database of the local device.

The work here will focus on investigating existing functionality in inter-domain link discovery at the control and management planes for different technologies, and how that can be exploited by the AutoBAHN system for automatic discovery of inter-domain peerings for its internal representation.

4.3 Resiliency

The AutoBAHN design includes support for the resiliency mechanism, but the specification of those mechanisms was not defined exactly during GN2. The work of JRA2 T2 will investigate the possibilities of introducing resiliency features at inter- and intra-domain levels, according to user requirements and technologies available.

The resiliency or circuits protection mechanisms can be described at two levels: inter- and intra-domain. The general concept of resiliency in a multidomain AutoBAHN system is created based on the combination of both possibilities. There will be several options that can be selected by user for global circuit resiliency:

- Level 0 Full Diverse Path Protection + Intra-domain Paths Protections
- Level 1 Partial Diverse Path Protection + Intra-domain Paths Protection
- Level 2 Full Diverse Path Protection
- Level 3 Partial Diverse Path Protection
- Level 4 Intra-domain Paths Protections
- Level 5 Best Effort Protection
- Level 6 No protection

The user interface at the moment of request creation should allow the user to select a required level of protection at inter-domain level. The decision on which intra-domain protection mechanisms to use is the responsibility of each single domain.

4.3.1 Level 0 Protection

Level 0 protection should be able to give full protection of the link in case of any failure in the network. The requirement is to find two independent inter-domain paths (source and destination domains will be the same, as the end-points are attached to them). Having two paths, which affects different chain of domains, prevents failure at inter-domain links and within each particular domain. Despite that, each single domain enables local path protection, according to its technology features. The example case is shown in Figure 4.8.

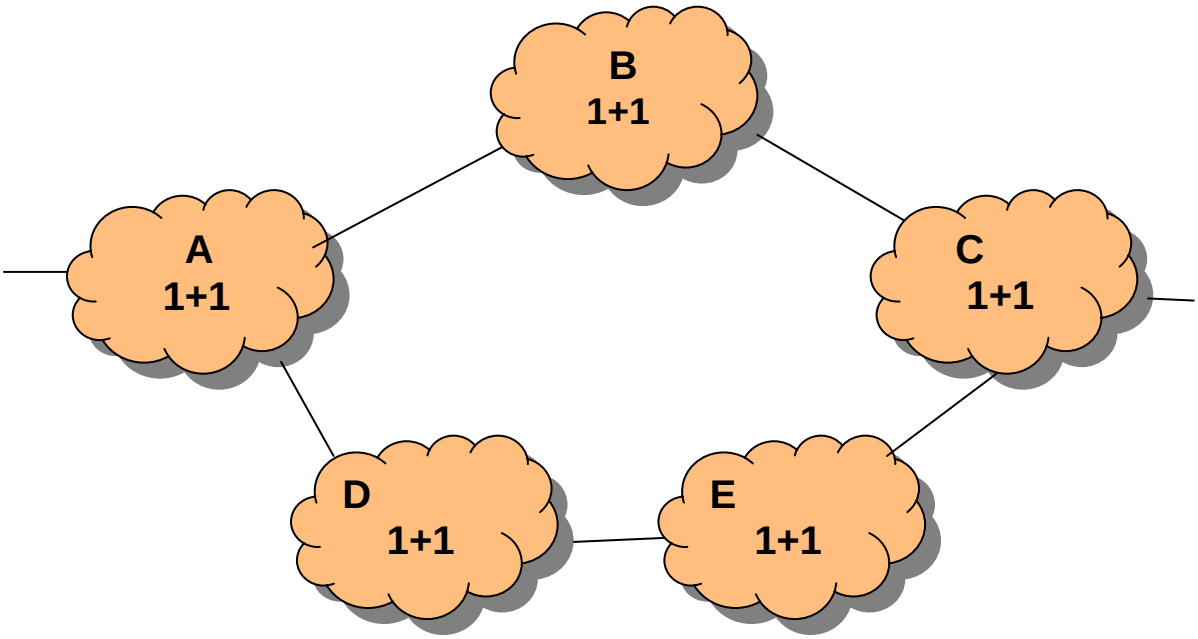

Figure 4.8: Level 0 Full Diverse Path Protection + Intra-domain Paths Protections.

The requirement for having Level 0 protection is to have a topology where two completely independent paths (not involving the same domains nor inter-domain links) could be created between particular two end-points. The primary circuit (green one) is created through domains A, B, and C, where each domain provides 1+1 protection. In case of a circuit failure in any of the domains, a domain will put a user traffic via protection intra-domain path, without affecting the inter-domain path. In case of inter-domain link failure, the whole global path will be rerouted to an alternative path (red one). The mechanisms for this rerouting are subject for further investigation within JRA2 T2's work. The alternative path can be instantiated in parallel to the primary one (this will increase cost, but reduce switching time) or the resources can be only booked when the circuit is instantiated in case of failure (decreases costs, but increases switching time). The choice may be either static, decided by AutoBAHN, or the user may have an influence on that choice during circuit request.

4.3.2 Level 1 Protection

The Level 1 protection is similar to Level 0 protection, except the domains and links along reservation paths on primary and protection path does not need to be different in all cases. This may apply in topologies where it is not possible to avoid configuring circuit through some domains. For example, in Figure 4.9 the primary path is defined as A, B, C, and D. The protection path is unable to avoid domain B, as there are no alternatives. Thus domain B will also be present on the protection path. However, the domain C and its inter-domain links can be avoided through domain E. Each domain on both paths is providing additional intra-domain protection regarding supported technologies (1+1 in this case).

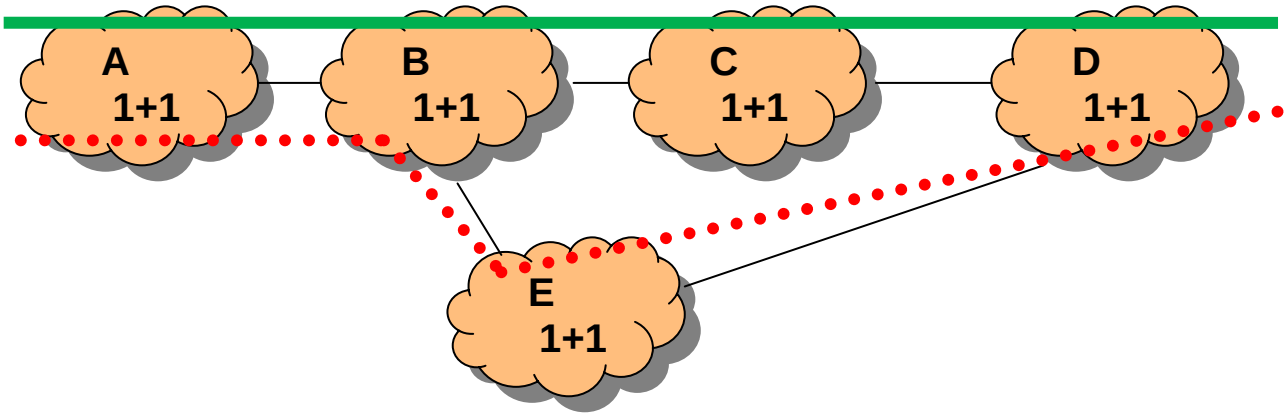

Figure 4.9: Level 1 Partial Diverse Path Protection + Intra-domain Paths Protection.

4.3.3 Level 2 Protection

Level 2 protection is similar to Level 0 protection. It is provided to a user when all domains along the primary reservation path can be replaced with alternative domains, so that two paths (primary and protection path) have only two elements in common (source and destination domain). The domains do not provide any protection mechanism at intra-domain level, thus any failure on an inter-domain link or within a domain always causes the rerouting of the whole inter-domain path. The situation is shown in Figure 4.10.

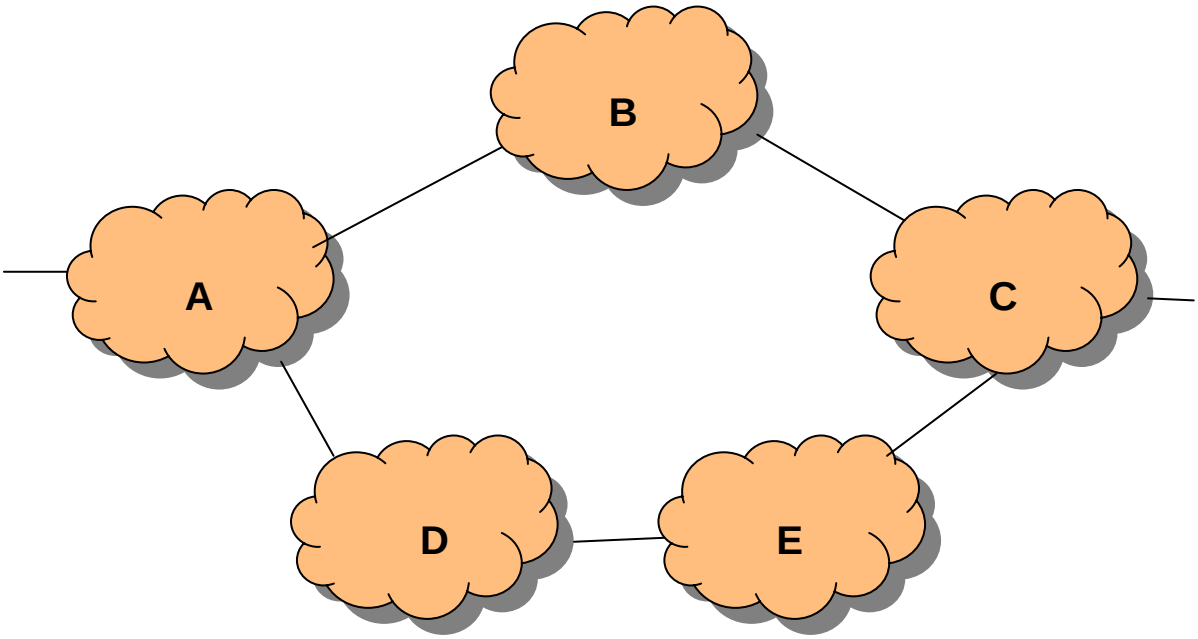

Figure 4.10: Level 2 Full Diverse Path Protection.

4.3.4 Level 3 Protection

Level 3 protection is similar to Level 1 protection. It may be provided when a particular domain on the protection path cannot be avoided, and intra-domain protection mechanisms are disabled. The path will be rerouted when a failure occurs (on the inter-domain links or within particular domain). However if the failure occurs within a domain that is common to both the primary and the protection path, the service will not be provided to the user. Figure 4.11 shows an example of Level 3 protection.

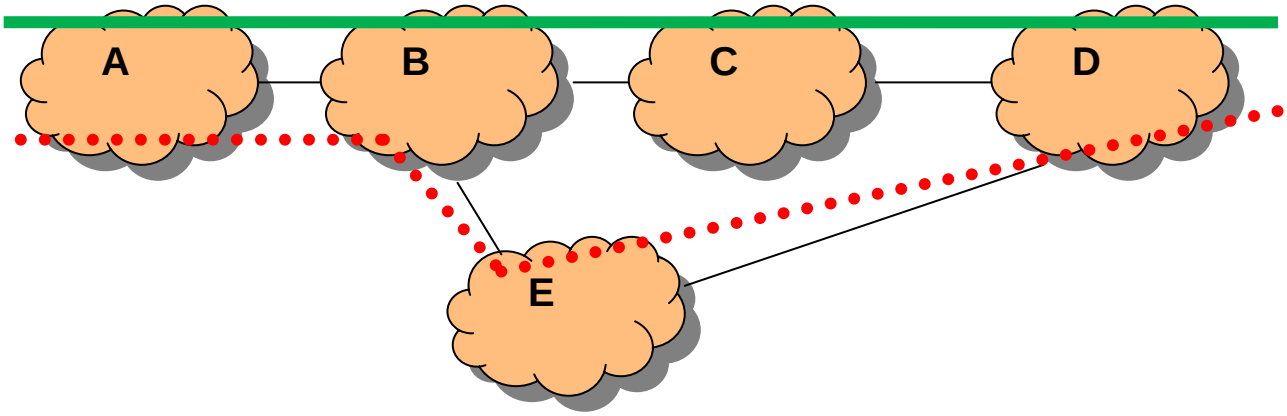

Figure 4.11: Level 3 Partial Diverse Path Protection.

4.3.5 Level 4 Protection

Level 4 protection relies on intra-domain protection mechanisms, typical for topologies within a domain. A user cannot select a particular mechanism, as it is the domains responsibility to choose an optimal protection type. The inter-domain links are not protected at all, therefore in case of failure at this level the transport service will become unavailable. An example of Level 4 protection is shown in Figure 4.12.

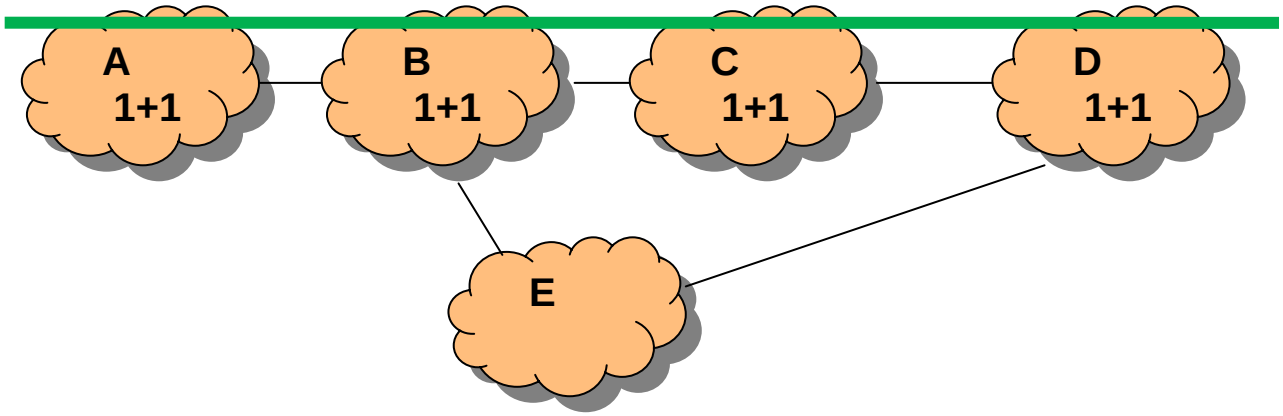

Figure 4.12: Level 4 Intra-domain Paths Protections.

4.3.6 Level 5 Protection

Level 5 protection does not provision a protection path at request processing time. A user has only the primary path available and configured, and no other resources are booked. In case of failure, AutoBAHN will start to reroute the circuit according to the current network condition and scheduled reservations. If an alternative path can be found, it is immediately configured and provided to the user. In case a path cannot be found, the service will become unavailable. This involves the process of pathfinding, resource negotiations and circuit configuration, thus the actions are the same as for new request processing. The process is time consuming and does not guarantee any success.

4.3.7 Level 6 Protection

Level 6 protection means that there is no circuit protection at all. Only one circuit is provided to the user, and in any case of failure, the service will not be provided to the user.

4.3.8 Mixed Protection Levels

It is possible to have mixed protection levels in case some options are not possible in particular domains. Figure 4.13 gives an example. The path is passing domains A, B, C, and D, but domain C does not support any protection for this reservation. The path is then not fully protected at domain C, and failure at this level will require rerouting of the whole circuit to a protection path. Mixed levels are the subject of further investigation, as they are special cases of defined protection levels. Separate scenarios and system behaviour need to be defined for such cases. JRA2 T2 will investigate this problem. However the results may be conceptual only, without prototypes or a service-level solution.

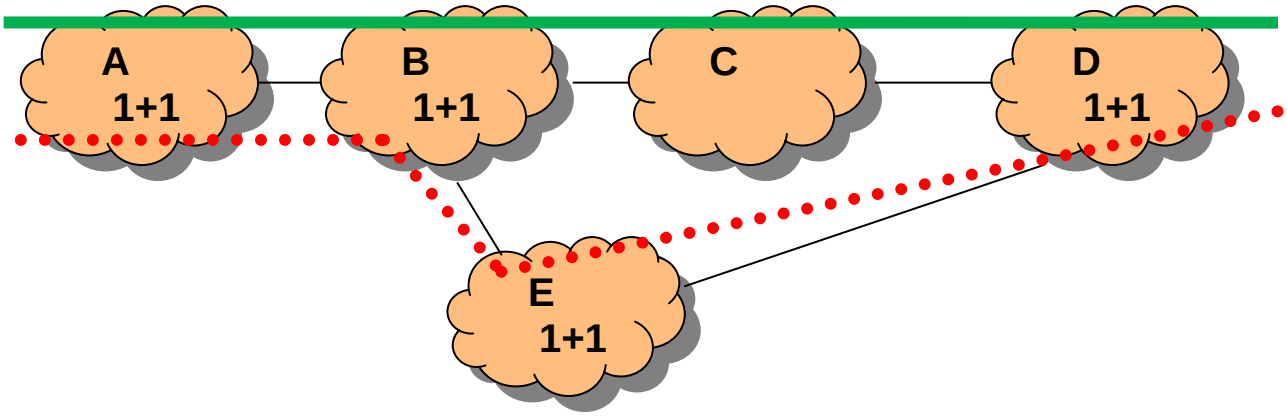

Figure 4.13: Mixed Protection Levels.

4.3.9 Additional Ideas

After defining a global path for any resiliency level, a user may be able to select a resiliency level for any element on the path, according to available options. For example, during the resources negotiations phase, each domain creates a constraint attribute describing possible resiliency options. At the destination domain, where constraints attributes need to be agreed, a user interaction could be possible to select from available resiliency options. For the topology shown in Figure 4.14, an AutoBAHN system has proposed the path marked as red.

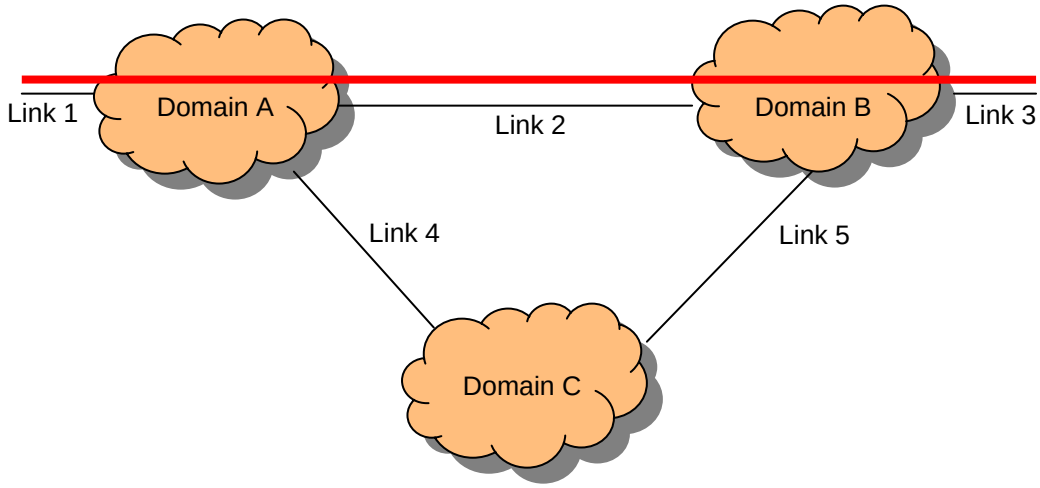

Figure 4.14: Custom Protection Levels.

A user can then select a resiliency level for the following elements:

- Domain A
- Link 2
- Domain B

Domain A allows the user to select a link protection for the GMPLS network of 1:1 or 1+1. Domain B only allows the user to select link protection for the GMPLS network 1+1. For Link 2 a user can choose between not protected or diverse path. A user, having a dedicated GUI, may select 1:1 in the Domain A, 1+1 in the Domain B, and diverse path option for the Link2. This will cause Domain A and B to be configured according to user selection, and the AutoBAHN system will find an alternative path avoiding the Link 2. In this case the path will go through Domain C.

As new domains are involved in the protected path, a user may again have an opportunity to select the protection level for Domain C. The process of the interactive building of protected circuits requires additional investigation due to the required GUI and modifications to the resource negotiation mechanism. This process will be part of the work of JRA2 T2. However, the results of it may be conceptual only, not supported by prototype or service-level solution. This functionality may be important for advanced users, aware of network protection mechanisms, expecting high levels of resiliency for data transfers. Such users need to be defined first in order to define process requirements and detailed functionality.

5 The Lookup Service

The AutoBAHN system has been designed and built as a distributed system and the service it provides is based on the collaborative effort of many domains. Users interface with the system by requiring reservations from one endpoint to another endpoint at a different domain. Each domain has its own internal identifiers for its data plane, which are not disseminated to other domains nor are generally known by the users of the AutoBAHN system.

All domains participating in the AutoBAHN architecture exchange their topologies in an abstracted form (using a different, AutoBAHN-specific identifier namespace), which is therefore universal among AutoBAHN instances but also unfamiliar to end users of the system.

Furthermore, each AutoBAHN instance comprises several modules and communication between these AutoBAHN modules takes place using web services technology. Each web service needs to discover the others. Therefore, the two broad categories of requirements for the Lookup Service with respect to an AutoBAHN instance deployed at a specific domain can be differentiated as “external” and “internal” requirements, as explained below.

5.1 External Requirements

The following are requirements for the user-friendly operation of the AutoBAHN architecture and for the discovery of neighbouring AutoBAHN instances.

- Translation of identifiers of end ports in the abstracted topology to a friendly form for end users.

  When the user wants to make a reservation, it is necessary to present a user-friendly list of endpoints (the start and end port of a requested reservation). Since the AutoBAHN identifiers for the endpoints are unfamiliar to the users (they are automatically produced during the abstraction process) the user-friendly identifiers must allow the user to understand if the service is available in their location and in the remote location they would like to connect to, by providing a combination of physical location and organisation information about the displayed endpoints. The role of the Lookup Service is to have a mapping of this user-friendly list of endpoints to the identifiers used by the system for the same endpoints, so that the AutoBAHN system instances can dynamically retrieve them.

- Location of IDM modules for neighbouring domains.

  To perform an inter-domain reservation, the IDM modules of the corresponding domains have to communicate (currently in a chain-like fashion). The unique identifier of an AutoBAHN domain is defined by the web service address of its specified IDM, as the identifier is in the form of a URL, containing the DNS name of the host running the IDM system. This approach implicitly uses the existing DNS service but requires manual insertion of the URLs. The requirement from the Lookup Service in this case is that each IDM should be able to retrieve the location (URLs) of the IDM modules of neighbouring domains, using only its knowledge of the neighbouring domain names.

- Interdomain links identification.

  As detailed in *Inter-Domain Links Discovery* on page 33, AutoBAHN faces the problem of matching an edge interface from one domain with the correct interface belonging to a neighbouring domain in order to properly identify interdomain links. A requirement from the Lookup Service is to enable the transition from the (currently used) manual approach to a semi-automated one (see *Semi-Automatic Discovery* on page 34). The Lookup Service will function as the repository where edge interfaces will be registered.

  The records in the Lookup Service need to contain the pair of domains that the interdomain link joins, and an identifier for the edge interface. For example, a record could be (domA, domB. port-x1), meaning that at domA there is an edge interface named port-x1 that is connected to the neighbouring domain domB. Then domB can query the lookup service for all edge interfaces from domA connected to domB, and properly match the information with its own edge ports (multi-homing cases still require manual intervention though).

Based on the above, the Lookup System entries will have to contain the following entry types:

- End port identifier / user-friendly name / (optional) domain name: E.g. (port1, Institute A at City X main connection). Optionally, this record may also contain the domain name that registered this endpoint; e.g. (port1, Institute A at City X main connection, domainA). This information is redundant (since end port identifiers are unique), but may be useful for the user interface, so that it can effortlessly present the available endpoints grouped by domain.
- Domain name / IDM instance URL: E.g. (DomainA, http://idm.domaina.net:8080/autobahn/interdomain).
- Start domain / end domain / public edge port identifier : E.g. (domA, domB. port-x1).

To meet these requirements, the interface to/from the Lookup Service should offer the following functionalities:

- Registration of end port identifier / user-friendly name mapping.

  Each local administrator is capable of identifying the physical location and organisation information about an end port in the domain-specific topology of his administrative area. By inputting this information to the topology-specific part of the database, the Topology Abstraction module of the AutoBAHN system can associate the user-friendly information to the abstracted topology identifiers and register it to the Lookup Service.

- Update or removal of end port identifier / user-friendly name mapping.

  When there is a topology update, a former endpoint may be updated or cease to exist. The Topology Abstraction module propagates this information to the Lookup Service, which correspondingly updates or removes the affected record.

- Registration of domain name / IDM instance URL mapping.

  Upon initialisation, each IDM instance should register this information to the Lookup Service.

- Update of domain name / IDM instance URL mapping.

  An IDM instance may be relocated. In that case, upon initialisation of the new IDM instance, it registers its information to the Lookup Service. The Lookup Service should be able to identify that there has been a previous entry for the same domain and update it accordingly.

- Removal of domain name / IDM instance URL mapping.

  Upon normal shutdown, an IDM instance should notify the Lookup Service and the Lookup Service may remove the corresponding record. This functionality may be optional, since IDM instances have internal mechanisms for handling error situations in their inter-domain communication.

- Registration of start domain / end domain / public edge port identifier mapping.

  Each AutoBAHN instance registers with the Lookup Service one record per edge interface, which contains the public name of the edge interface and the names of the domains it connects. If there are no multiple links connecting the same domains, this information can be used to automatically identify inter-domain links (see *Semi-Automatic Discovery* on page 34).

- Update or removal of start domain / end domain / public edge port identifier mapping.

  When there is a topology update, an inter-domain link may be updated or cease to exist. The Topology Abstraction module propagates this information to the Lookup Service, which correspondingly updates or removes the affected record.

Note that the above specifications rely on the assumption that domain names are unique.

The nature of external requirements lends itself to a centralised implementation solution, with a universally accessible repository from which all user interface module instances can retrieve user-friendly port mappings, Topology Abstraction module instances can retrieve inter-domain links identifiers and IDM module instances can retrieve the location of neighbouring IDM instances.

In case AutoBAHN deployment is widespread to the level where the centralised solution is no longer efficient, a more structured approach would be required (probably similar to perfSONAR's hierarchical Lookup Service). In this case, there are several layers of lookup servers (in the case of perfSONAR 2 levels, global and home lookup servers), and lower layer servers perform a summarisation of their information before transmitting it to the global Lookup Service (which functions similarly to a root DNS server). This process significantly complicates the update and retrieval of information, so during the project the related trade-offs will be evaluated.

5.2 Internal Requirements

The following are requirements for the internal operation, flexibility and ease of installation of a single AutoBAHN instance.

- Location of certain AutoBAHN services inside a domain (DM, TP, and possibly other standalone modules such as Topology Abstraction and Resource Reservation Calendar).

  This functionality might be useful because of the modular nature of the AutoBAHN system. In particular, each domain contains several standalone modules, which communicate with each other via a well-defined WS interface.

- Availability status of registered modules.

  It is useful for the Lookup Service to be able to monitor the availability status of the registered modules. The Lookup Service keeps additional information about each registered service, in the form of a single up/down flag, indicating whether the module is alive or not. This functionality can be provided either with a push (where the components notify the Lookup Service of their availability) or a pull mechanism (where the Lookup Service queries the modules for availability). Since modules may go down and not have the capability to inform the Lookup Service, it seems preferable to implement a pull mechanism, combined with the possibility for the modules to register and de-register themselves from the Lookup Service.

  The preferred solution seems to be a keep-alive mechanism, where each module is responsible for periodically sending a simple message to the Lookup Service indicating that it is alive. The Lookup Service periodically refreshes its records, and marks any service modules that have not sent keep-alive message for some while as down. All error handling in case a module is down is handled by the other modules that want to communicate with that module.

Based on the above, the lookup system entries will have to look as follows:

- Module name / Web Services URL / availability status

To meet these requirements, the interface to/from the Lookup Service should offer the following functionalities:

- AutoBAHN module registers itself in the Lookup Service.
- AutoBAHN module de-registers itself from the Lookup Service.
- AutoBAHN module retrieves the point of access to other modules.

5.3 Implementation Options

Currently AutoBAHN does not fulfil the external lookup requirements, while it uses a manual approach for meeting the internal lookup requirements. The short-term predicted scale of AutoBAHN architecture deployment indicates that the current approach for meeting the internal requirements is adequate for the foreseeable future.

The external requirements of the Lookup Service for AutoBAHN do not seem to justify a very complicated solution and therefore a gradual approach such as deploying a simple centralised repository seems to be the most suitable option at this stage.

However, an important demand from the Lookup Service is high reliability, as the operation of the AutoBAHN system will depend on it and unavailability of the service will cause significant degradation or complete unavailability of the AutoBAHN system. Robustness of the system can be achieved by the combination of two approaches, one implemented by the Lookup Service and one implemented by the modules acting as clients to the Lookup Service:

- Replication

  For increased resiliency the main lookup server could be supported by one or more backup servers that should be kept synchronised with the main one. The location of the backup servers can be registered in the main server. Upon first retrieval of information from the main server, the AutoBAHN modules can keep the backup addresses in a local cache and utilise them in case of failure of the main lookup server. Backup addresses can even be stored in AutoBAHN modules using external intervention (e.g. administrator configuration), in order to avoid service breakdown even if the main server remains unavailable for long periods. The replication requirement means that the Lookup Service should be aware of the backup instances and relay any new information to them as well, so that the synchronisation of the backup instances is transparent to the clients of the Lookup Service.

  An example of replication for the Lookup Service is shown in Figure 5.1. Each lookup server keeps a list of all the other lookup servers. Upon receiving an information update, it relays it to all lookup servers in its list. When a new lookup server is initialised, it has to know at least one already existing lookup server. It announces itself to the existing lookup server (so the existing lookup server is responsible for flooding this information) and it retrieves from the existing lookup server the whole repository.

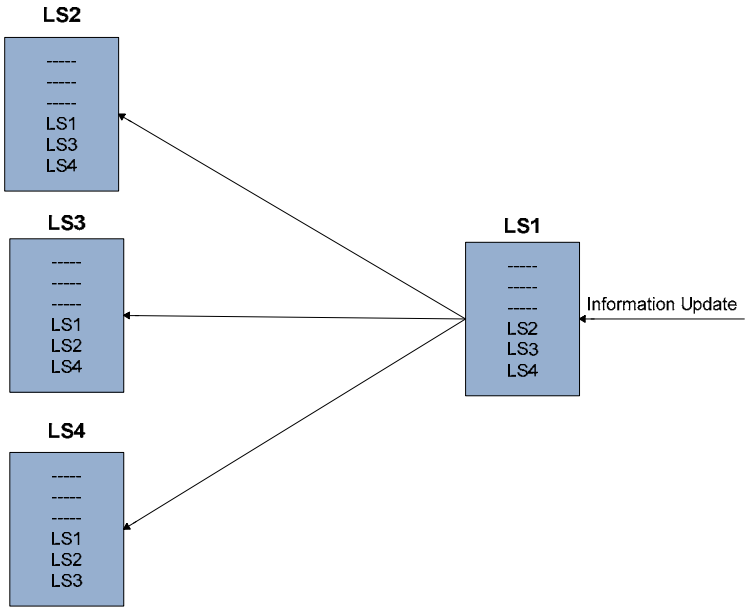

Figure 5.1: Replication for the Lookup Service example.

- Fallback mechanisms

  In case all lookup servers are unavailable (for example if they are hosted at the same network and access to that network is unavailable), fallback mechanisms can be used by having some AutoBAHN modules (for example the IDM) keep a local cache of the information learnt from the Lookup Service. This feature is also beneficial for reducing traffic and request load to the Lookup Service. The local cache can either have en expiration date for its records or the modules can query the Lookup Service when the resources indicated in the local cache are unavailable.

  To meet the internal lookup requirements (if judged to be within the scope of the project), several approaches are available:

  - UDDI is an existing mechanism for web services discovery. The purpose of a UDDI registry is to offer a standard-based mechanism for listing, classifying and managing Web Services. UDDI has evolved from focusing on a universal registry to focusing on federated control. The core data stored in UDDI registries as of UDDI v3 are:
    - The description of the business function of a service.
    - Information about the organisation responsible for the service.
    - The technical details of the service.
    - Relationships among entities.
    - Requests to track changes to a list of entities.
    - Various other attributes such as taxonomy and digital signatures.

According to the UDDI v3 categorisation of UDDI registries, the internal requirements of the Lookup Service match to the “Corporate/Private” registry type (the most restricted one), where no data is shared with other registries (in different domains). The advantage of using UDDI is that existing frameworks such as the Java-based jUDDI can be used for implementing the registries. The trade-off is reduced flexibility as the data has to be made to conform to the UDDI specifications. UDDI v3 supports extension through a derivation mechanism provided by XML Schema, so this feature could be used for data that does not fit in the existing UDDI information model.

- A DNS-like distributed mechanism, similar to the one developed by JRA1 during the GN2 project (perfSONAR). In that case, the Lookup Service allows every independent service to be a visible part of the system. The perfSONAR Lookup Service mechanism is hierarchical, using global and home Lookup Services. The home Lookup Service (hLS) in each domain registers with any global Lookup Service (gLS) that is maintained by a perfSONAR partner organisation. Each gLS is regularly synchronised with all other gLS instances to ensure they have an up-to-date global view of which web services are available. Client applications (usually visualisation tools) query a gLS instance for the address of an hLS that might store detailed information about these services. An important difference between the perfSONAR requirements that led to this approach and AutoBAHN is that the perfSONAR Lookup Service should be able to provide information about multiple modules across different domains. In contrast, AutoBAHN modules (with the exception of the IDM and the User Interface to it) are only accessible within their domain, negating the need for an elaborate hierarchical structure.

It is planned that in any case, any suggested solution will have to be integrated with all the tools supported by SA2 and not just the AutoBAHN system.

6 Authentication, Authorisation and Accounting (AAA)

6.1 AAI

AutoBAHN is recommended to use the existing eduGAIN infrastructure for authentication and authorisation. The eduGAIN infrastructure consists of Federation Peering Points (FPP), the modules used by federations to upload data to the eduGAIN Metadata Service (MDS), and Bridging elements (BE), which serve as a means of establishing appropriate trust links among federation components and user applications and of adapting syntax, semantics and procedures used by the participating federations. The eduGAIN trust is maintained among these elements. Bridging elements serve also from transforming the trust matrix problem from NxM problem to one-to-one mapping from eduGAIN trust into the corresponding internal federation trust. Current eduGAIN decisions may lead to the deprecation of Bridging Elements as their functionality can be replaced by standardised Security Assertion Markup Language (SAML) 2.0 communication.

eduGAIN has specified four profiles that support corresponding use cases:

- Web SSO profile

  Used for the Single-Sign On (SSO) use case, where a user is authenticated using the credentials from his home domain.

- Automated Client (AC) profile

  Used by services that run without direct user interaction and is therefore based on X.509 certificates instead of user credentials.

- User behind a Client (UbC) profile

  Applicable to WS clients that run under the direct control of a human.

- Client in Web Container (WE) profile

  Applicable to clients that are accessed by end users through a web container, and the users are authenticated to the client through the Web SSO procedure.

In the framework of the GN2 project, several steps have been made in the direction of integrating AutoBAHN with the existing eduGAIN infrastructure and AAI mechanisms. The purpose of the AAI task in the current project is to continue the existing process, re-evaluate any decisions if deemed necessary, and undertake work in the areas that have not been yet initialised.

6.1.1 Current AAI Status in AutoBAHN

The current version of the AutoBAHN system provides an interface to perform user authentication and authorisation, however, only a limited implementation has taken place. In particular:

- User authentication is performed at the web-based user interface used for managing the system and requesting resource reservations. The UI follows the Web SSO profile. However, the eduGAIN filter necessary to implement the profile has been unstable, and stable AutoBAHN operation requires deactivating eduGAIN-based authentication.

- Trusted communication between AutoBAHN modules (inter-domain IDM-IDM communication and intra-domain communication) has not been implemented, although planned to follow the AC profile.

- Authorisation mechanisms have also not been yet implemented.

6.1.2 User Authentication

When a user wants to make a reservation in a resource, eduGAIN SSO infrastructure will be used for authentication and authorisation purposes as described below.

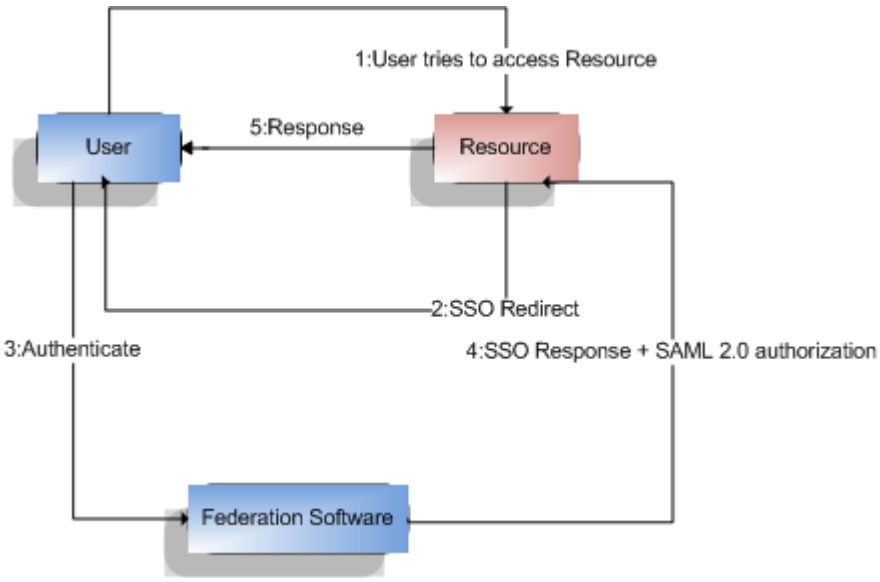

Figure 6.1: Message flow when a user wants to make a reservation.

In principle, if a user tries to make a reservation directly, the resource redirects the user to the Single Sign-On service of their federation. Then the user is authenticated through the federation software, which sends the SSO response and SAML 2.0 authorisation back to the resource. The response contains both authentication and authorisation information as SAML 2.0 attributes. Finally, the resource checks the SSO response and SAML 2.0 attributes and responds to the user appropriately about their reservation request. The proposed attributes transmitted are detailed in *Authorisation and Related eduGAIN Attributes* on page 57.

Below is presented the detailed procedure in the context of the AutoBAHN system, as specified during the GN2 project.

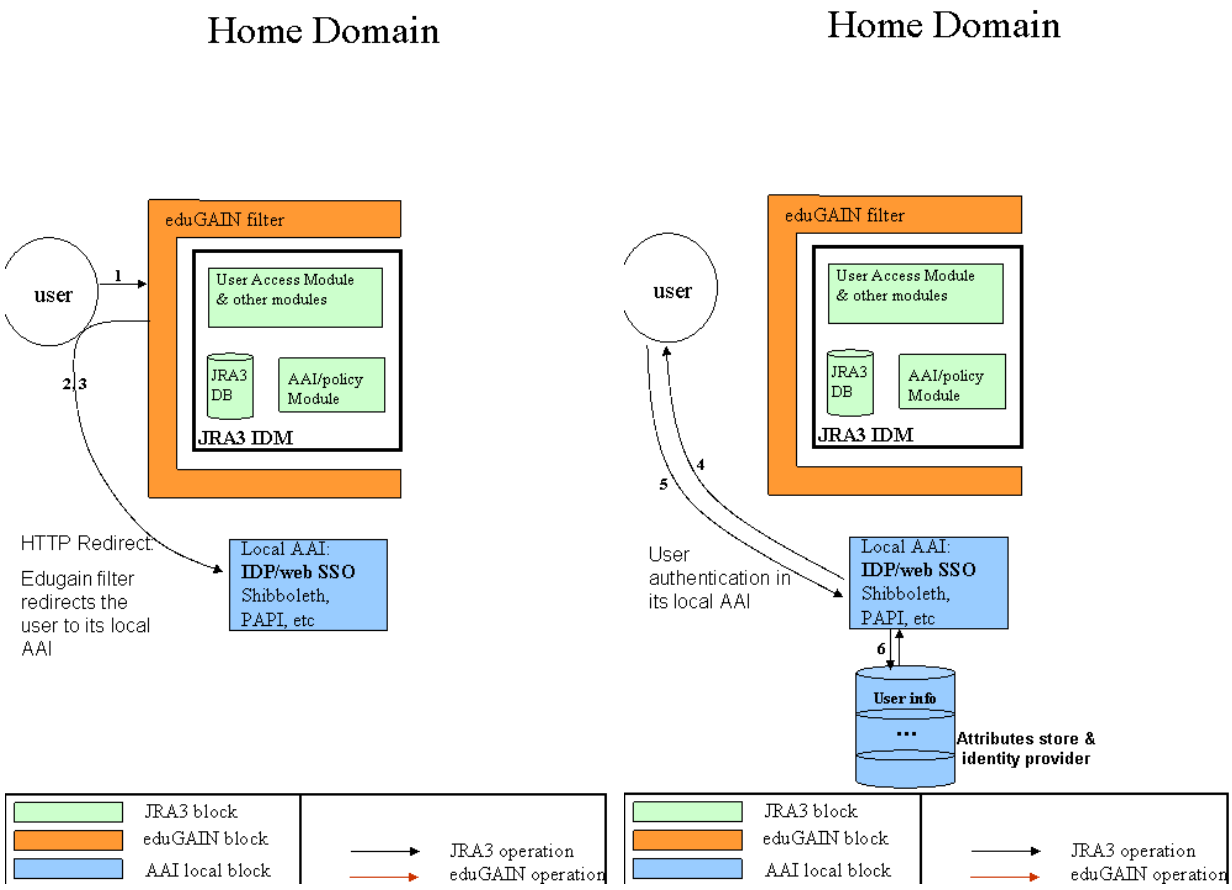

Figure 6.2: Attribute authentication 1.

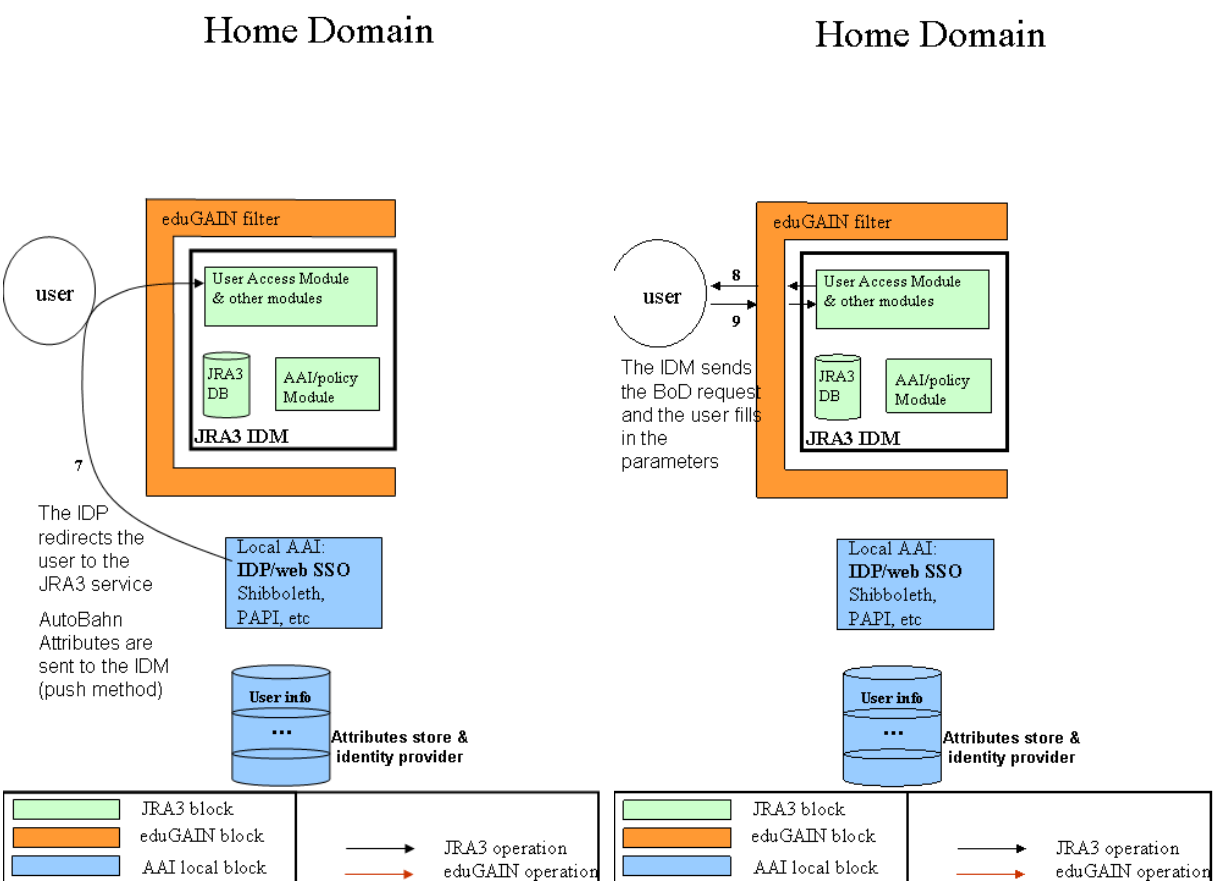

Figure 6.3: Attribute authentication 2.

1. The user (through a web browser) tries to access the AutoBAHN service (the web-based User Interface) of the starting point of the required reservation.
2. The eduGAIN filter intercepts the request and sends to the web browser an http redirection to an IdP (or GIdP). In order for this redirection to take place, eduGAIN has implemented a WFAYF (Which Federation Are You From) service, which is based on the eduGAIN MDS and allows the user to select the appropriate IdP for further processing.
3. The user's web browser sends an http request to the IdP server.
4. The IdP server sends to the web browser a page to authenticate the user.
5. The user sends their credentials (login + password, certificate, etc) to the IdP server.
6. The IdP authenticates the user using the credentials and the local database (ldap, etc). The user attributes concerning AutoBAHN are also retrieved.
7. The IdP server redirects the web browser to the AutoBAHN service. The local AAI also sends the AutoBAHN attributes to the IDM. The IDM stores these attributes.
8. The IDM sends the BoD request page.
9. The user fills in the page and sends it to the IDM. From then on, the reservation request procedure is initiated by the IDM.

The purpose of further development in this case will be to elevate, in coordination with eduGAIN, the current implementation to production level. Furthermore, the possibility of establishing Identity Providers specifically for the purposes of authenticating AutoBAHN users will be investigated.

6.1.3 Trusted Communications between AutoBAHN Modules

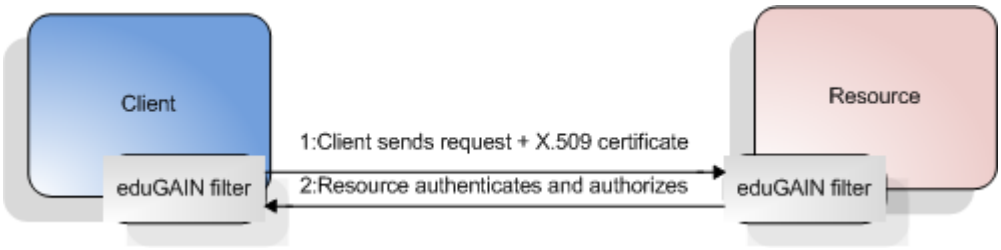

Figure 6.4: Message flow when an automated client wants to make a reservation.

In principle, when the client module wants to communicate with another module (the resource), it sends its request to the required resource along with its X.509 certificate through its eduGAIN filter. The eduGAIN filter of the resource authenticates the client by validating its certificate. The certificate contains identification information that allows the resource to authenticate only designated clients.

The detailed procedure in the context of the AutoBAHN system for the trusted communication between AutoBAHN modules is as follows:

1. The AutoBAHN module that wants to communicate (client) must have a certificate, so no interaction for credentials is needed. The X.509 certificate is issued by a Certificate Authority (CA) subordinated to one of the eduGAIN roots of trust [DJ524].
2. The client module sends its request and the certificate to the resource.
3. The resource module performs trust validation by checking that the whole trust path of the certificate correctly resolves to the root(s) of trust defined by eduGAIN.
4. The resource checks that the client module is allowed to access it.
5. The resource provides the requested answer to the client module.

The purpose of further development in this case will be to use the eduGAIN trust fabric (composed of a hierarchy of Certification Authorities) in order to make the trusted communication between AutoBAHN modules possible.

6.1.4 Multi-Domain User Authorisation

After a user has been authenticated and is able to submit a resource reservation request, an authorisation procedure should take place that determines, according to the specified policies, whether this specific user should be able to reserve the resources. This decision has to be taken in every domain along the reservation path, based on user attributes that have to be transmitted with the reservation request and mapped to the policies implemented by each domain.

Below is presented the detailed procedure in the context of the AutoBAHN system, as specified during the GN2 project for the multi-domain authorisation procedure.

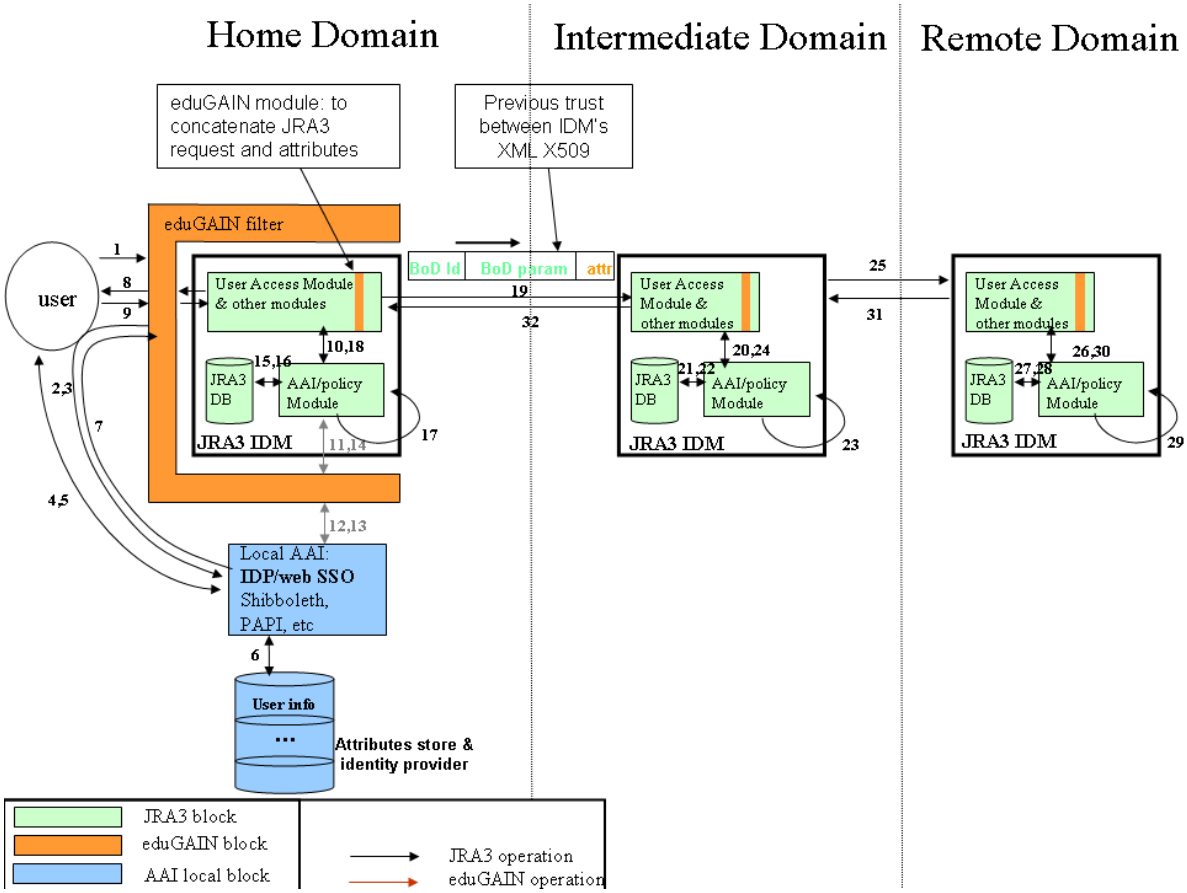

Figure 6.5: Multi-domain authorisation.

Steps 1-9 are the user authentication procedure. Steps 10-18 are the possible authorisation procedure within the start domain of the reservation. When the authorisation policy module examines a reservation request, it might wish to retrieve additional attributes, in which case it should obtain them through the eduGAIN filter. During the course of the project, the desirability of retrieving additional attributes will be evaluated (as opposed to acquiring all needed attributes with the initial user authentication).

When the reservation request has been authorised in its Home Domain and the IDM wants to propagate further down the selected reservation path, it has to send the request to the next domain. The attributes are sent in the same request. An eduGAIN module is planned to be used to concatenate these attributes in the AutoBAHN request (XML).

Upon arrival at the next domain, the possible authorisation procedure is repeated there (steps 20-24) and at every subsequent domain (steps 26-30).

The purpose of further development in this case will be to evaluate the desired complexity for the authorisation policy and develop a prototype implementation accordingly.

Issues that have not yet been addressed will be investigated, such as the involvement of the destination endpoint in the authorisation procedure, and the possible differentiation of authorisation policies among domains. For example, different domains may allow different capacity reservation limits for requested circuits depending on the role of the user, his organisation etc. Concerning the classification of users there are several different options:

- Each reservation made by an authenticated and authorised user is credited to the user individually.
- Each reservation made by an authenticated and authorised user is credited to the user's home domain, and counts against an aggregate limit for all users from the same domain.
- Each reservation made by an authenticated and authorised user is credited to the total number of reservations and counts against an aggregate limit for all AutoBAHN users.

It is possible that each domain chooses its own policy regarding the classification of users, or that over time policies change. The authorisation procedure should therefore be able to handle all of the above possibilities.

Furthermore, some sort of granularity in terms of authorisation flexibility is required, so that for example users can perform a subset of the available actions through the AutoBAHN management interface (e.g. monitor service, use service, and administrate service). The structure of the subset of allowable actions can also be defined by each domain.

6.1.5 Authorisation and Related eduGAIN Attributes

The authorisation granularity and the classification of users will be supported by the type of attributes transmitted by the eduGAIN infrastructure (the identity providers of the users) during authentication. The AutoBAHN system can then transfer these attributes to the policy module to implement the specified authorisation policies.

When a service uses eduGAIN for authentication, it first redirects the user to the eduGAIN mechanism of their IdP. The user enters their credentials and then the eduGAIN mechanism communicates with the service (typically using SAML) to make the authentication possible.

The attributes that should be provided by the IdP (or a separate attribute provider) and which will be propagated along the reservation path to be mapped to policies such as the ones described above are:

- Name/Email

  A unique ID of the user wanting to make a reservation. This could be either the name or the email of the user, or a combination of both.

- Organisation

  The organisation/domain/federation of which the user is a member.

- Project Membership

  This attribute should contain a specified value (e.g. AUTOBAHN) that demonstrates that this user is an authorised AutoBAHN user.

- Project Role

  This attribute offers granularity in terms of the subset of available actions that the user is allowed to perform, and can contain values such as Administrator, Developer, User, etc.

The above attributes are going to be named in compliance with LDAP schemas for academic and research environments, such as eduPerson/MACE-Dir and SCHAC, with attributes for describing individuals and facilitating inter-institutional data exchange. For example, when expressed in a SAML assertion in accordance with the MACE-Dir SAML Attribute Profiles, the first attribute could be urn:mace:dir:attribute-def:mail, while according to SCHAC specifications the rest of the attributes could be expressed accordingly as urn:mace:terena.org:schac:schacHomeOrganisation, urn:mace:terena.org:schac:schacProjectMembership and urn:mace:terena.org:schac:schacProjectSpecificRol. It is planned that the same schemas used by current eduGAIN services (based on MACE-Dir and SCHAC specifications) are also going to be applied here.

The possibility of modifying the list of transmitted attributes is going to be investigated according to the final formulation of the authorisation policies.

6.2 Accounting

Accounting of network resources is important in a multi-domain network environment were the offered service relies on some kind of effort from the participating domains. The accounting is typically agreed upon in a service contract supplemented with a Service Level Agreement (SLA) that determines which parameters to account for, and the necessary service level to deliver in order to account for the service.

Traditionally accounting in, for example, Plain Old Telephony Systems (POTS) is done by traffic volume, for example, Erlang hours. However, in a multi-domain environment the accounting parameters will be agreed upon between the participating network operators. Therefore the accounting system parameters should be adaptable to the operator demands.

Accounting in the network can be done in a number of different ways. In one approach accounting should start/stop metering when the requested service is first/last available for the end user who requested the service. The results of this approach are that some domains will actually set up the requested resource before the resource is available to the user. If the resource is a guaranteed service (e.g. wavelength or guaranteed bandwidth), it will take resources from the operator that are not accounted for.

Another approach is to start/stop service metering when the user requires the service (or wants to terminate it). This approach has the disadvantage that the user may end up paying for time when the service is not available (imagine an environment where the signalling has to traverse many IDMs to be established).

Whichever approach is chosen, it is important that the chosen approach is well documented and described in the SLA between the participating domains. It is evident that both approaches have pros and cons for the user and the participating domains, depending on the number of service request and the distributing of these among the domains.

For the sake of implementing an accounting system that is simple and scalable, metering of resources could be done in each domain when a path is established in that particular domain. Figure 6.6 shows a resource reservation sequence in an environment with three different domains. Here metering of the resources should be started/stopped whenever signalling between the Inter Domain Manager and the Domain Manager takes place. In Figure 6.6 this is steps 2, 5 and 7. Again this approach is not entirely fair to the user, but this can be compensated by, for example, lowering the average rate that has to be paid by the user (since the system may be simpler to operate and implement).

The AutoBAHN system accounting enhancements will be based on a specification of the accounting logic and data model together with the principles of accounting principles in each domain.

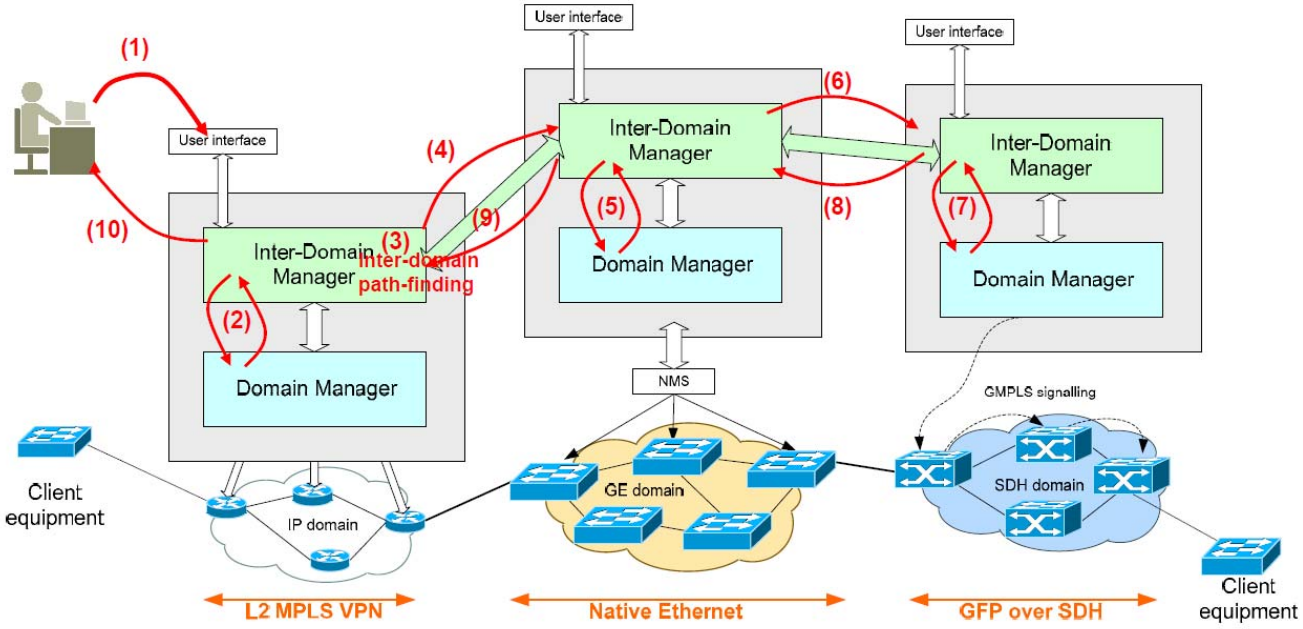

Figure 6.6: Signalling sequence for resource reservation.

7 Inter-Systems Protocols and Collaboration

AutoBAHN established a number of different collaboration opportunities with other BoD initiatives during the GN2 project lifetime. A close collaboration with Internet2 and ESnet organisations was established under the DICE control plane group, enabling AutoBAHN to create circuit reservations with end points located in the Internet2/ESnet domains in the United States of America, utilising the respective DRAGON/OSCARS systems.

A new protocol was defined to exchange information between BoD systems via intelligent proxy instances. This protocol was defined as the common substrate of the inter-domain reservation protocols used in AutoBAHN and DRAGON/OSCARS systems. As a follow-up for this work, AutoBAHN representatives have been involved in the OGF Network Service Interface Working Group (NIS-WG).

Supporting NSI-WG will be one of the priorities in standardisation efforts in JRA2 T2. The experience gained in AutoBAHN and collaboration with DRAGON/OSCARS will allow having significant influence on the interface and data models for provisioning network services design in NSI-WG. The official description of NSI-WG states:

> "The main purpose of the NSI WG is to facilitate interoperation between Grid users, applications and network infrastructures spanning different service domains, via the development of abstract messaging and protocols.
>
> The NSI WG must provide a general and open definition independent of implementation of provisioning systems (e.g., Grid and network). It should be sufficiently flexible, modular and scalable to facilitate future enhancements. The NSI WG recommendation will allow any user and network service to interoperate by using a common naming and message definition.
>
> The NSI WG will also focus on identifying existing standardisation activities/documents, understand their relevance and specify the relationships with regards to NSI (e.g., OGF (NM-WG, NML-WG) IETF, OIF)." [NSIWG]

The emphasis will be placed on the definition of resource reservation protocol, including advance reservation features. As OGF is a place for exchange of requirements and contributions between large number of NRENs, institutions and vendors, a valuable feedback is expected allowing development of practical and functional protocols. The NSI-WG activity is strictly focused on resources negotiations and reservations and is not dealing with pathfinding algorithms or AAI. The protocols for inter-domain topology and resources information exchange are also not the highest priority. However, a set of requirements will be defined. NSI-WG is collaborating with other OGF work groups to assure compatibility with other standardisation works and improve the quality of the final results.

The AutoBAHN activity in JRA2 T2 will not only be focused on participating to the NSI-WG, but also on implementing the interface support and associated protocol stack deployed there as part of the AutoBAHN interoperability work. This will enable immediate collaboration with other services/tools implementing the NSI interface, and also contribute to validating it and stimulating its development and adoption. Internet2/ESnet are expected to be partners in this effort.

Since the GN3 project started, the collaboration within the DICE control plane group has been re-established. It is planned to participate in the design of the inter-domain reservation protocols and continuously exchange experience in the area of inter-domain reservation procedures. The first steps will be focused on definition of particular additions to the current functionality, including AAI features, topology exchange and resources negotiations. Also, integration of services will be a part of the research agenda (e.g. global interaction of AutoBAHN and perfSONAR). A testbed infrastructure is expected to be created. While European NRENs will be deploying the AutoBAHN system for production service, the production level interoperation with Internet2/ESnet will also be considered. The continuous collaboration will require participation in meetings and periodical video-conferences.

In the background, JRA2 T2 will also investigate other developments on inter-domain resource reservation, especially under the umbrellas of IETF but also GLIF, G-Lambda and the outcomes produced in previous projects and initiatives (such as the EU-FP6 PHOSPHORUS project with its Harmony system). Within the task's lifetime, it will be impossible to fulfil compatibility requirements for all other BoD solutions, thus a prioritisation is needed, according to the GEANT-NREN and user community requirements. The main scope is to keep track of the relevant developments in the commercial world, and improve the operational aspects and adoption rates of AutoBAHN, rather than implementing a message proxy for each potentially collaborating BoD system. An attempt will be made to introduce compatibility of alternative BoD systems with AutoBAHN through research results of DICE and NSI-WG initiatives.

8 Implementation Time Lines and Conclusions

The JRA2 T2 task is defined for two years duration. The timelines of particular sub-tasks are defined on the GANTT chart in Figure 8.1:

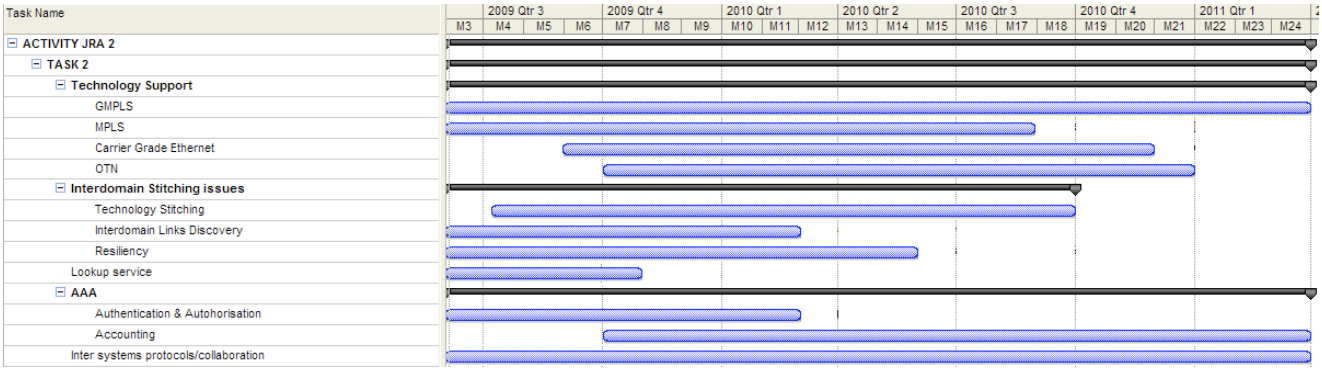

Figure 8.1: JRA2 Task 2 timelines.

The work on the GMPLS Technology Proxy (TP) is expected to be continued for the duration of JRA2 T2, as this technology offers a lot of functionality and requires detailed investigation. GMPLS-enabled networks are one of the priorities of the AutoBAHN deployments in the GN3 project. There are a number of decisions that need to be made to offer a functional and usable GMPLS TP.

MPLS TP was investigated during the GN2 project. This work is now extended to GN3, with an attempt to enable TE extensions and to improve scalability and efficiency of existing solutions (e.g. for Polish NREN PIONIER). This sub-task is expected to be terminated at M17 of the GN3 project.

Carrier Grade Ethernet is not currently widespread in NRENs and was not a subject of research during the GN2 project. This is also true for Optical Transport Network, for which an investigation of standards and available functionality is required. Those tasks are initiated in M6 and M7 and are expected to be terminated at M20 and M21. Due to potential limited accessibility of the testbeds or production deployment infrastructure, the sub-tasks may result in a rather theoretical study supported by simplified prototypes. Although a detailed specification may not be delivered for these technologies, the work done in JRA2 T2 will provide a good base for development of technology proxies in the future, as the technologies will become more popular.

During GN2, research was carried out on technology stitching for inter-domain reservation. It resulted in a good theoretical basis for work in the GN3 project. The stitching framework is expected to provide a well defined prototype for inter-domain resources negotiation procedures, especially for domains with different underlying technologies, taking also into consideration any progress in standardisation. The sub-task is expected to be terminated in M18. After that a prototype and its specification could be provided to the SA2 activity for further development.

The issue of inter-domain link discovery was experienced in GN2. Due to limited resources it was impossible to provide a satisfactory solution. In GN3 the investigation will continue and the best possible solution will be proposed. A prototype and specification for production class service will be delivered at the end of M11.

The intra-domain resiliency mechanisms will be studied independently for specific technologies, while working on technology proxies. In spite of that, resiliency mechanisms at inter-domain level will be investigated as a separate task. As a result, a specification for inter-domain path protection will be proposed. A prototype based on this specification is envisaged at M14.

To improve the quality, efficiency and maintenance of the AutoBAHN system, a Lookup Service will be specified. Its purpose will be to support external and internal AutoBAHN system functions, as well as to help administrators and users to interact with the provisioning environment. This task will result only in a requirements specification, which will be forwarded to SA2 T5 for implementation in M7.

In GN2 the collaboration with AAI was limited to user authentication while accessing the GUI. The authentication of IDMs and other software modules was based on X.509 certificates and did not require any external entities. In GN3 it is expected that AutoBAHN will utilise the AAI infrastructure for secure internal communications as well. Additionally, the authorisation mechanisms will be introduced, allowing more control over users and resources. The specification for this functionality should be provided by M11.

In the meantime, in M7, the work on accounting mechanisms will be initialised. The objective of this task is to provide a specification for the management of resources, tracking utilisation and user activities. The final results should be ready at M24.

During GN2 the AutoBAHN group was very active in international collaborations regarding BoD services and inter-domain resources provisioning. This work will be continued in GN3, with emphasis on standardisation efforts and also on enabling collaboration of services at a production level in global scale. This sub-task activity will last throughout JRA2 T2 and should be complete in M24.

References

[AutoBAHN]	Automated Bandwidth Allocation across Heterogeneous Networks http://www.geant2.net/server/show/nav.756
[DJ3.3.1]	M. Büchli, M. Campanella, G. Ivánszky, R. Krzywania, B. Peeters, D. Regvart, V. Rejis, L. Serrano, A. Sevasti, K. Stamos, C. Tziouvaras, D. Wilson, "Deliverable DJ.3.3.1: GÉANT2 Bandwidth on Demand Framework and General Architecture" http://www.geant2.net/upload/pdf/GN2-05-208v7_DJ3-3-1_GEANT2_Initial_Bandwidth_on_Demand_Framework_and_Architecture.pdf
[DJ3.3.4]	M. Campanella, R. Krzywania, A. Sevasti, S. Thomas, “Deliverable DJ3.3.4: Functional Specification and Design of a Generic Domain-centric Bandwidth on Demand Service Manager” http://www.geant2.net/upload/pdf/GN2-08-129-DS3-3-4_Functional_Specification_and_Design_of_Generic_Domain-centric_BoD_Service_Manager.pdf
[DJ3.5.3]	A. Escolano, A. Mackarel, D. Regvart, V. Reijs, G. Roberts, H. Popovski, "Deliverable DJ3.5.3: Report on Testing of Technology Stitching" http://www.geant2.net/upload/pdf/GN2-07-066v5-DJ3-5-3-Report_on_Testing_of_Technology_Stitching.pdf
[DJ524]	D. Lopez, J. Rauschenbach, A. Solberg, M. Stanica, S. Winter, C. Rodríguez, "eduGAIN Profiles and Implementation Guidelines" http://www.google.co.uk/url?sa=t&source=web&ct=res&cd=1&url=https%3A%2F%2Fmail.internet2.edu%2Fwws%2Farc%2Fshibboleth-dev%2F2008-02%2Fdoc00001.doc&ei=_Ni5SrLBGanajQfYz8X9BQ&rct=j&q=DJ.5.2.4+eduGAIN+Profiles+and+Implementation+Guidelines&usg=AFQjCNEMKP9wnyPqu-6mZe-2StJ1js6hUg
[eduGAIN]	http://www.edugain.org/
[ESCAL-COMNET]	"Advance Reservations for Service-Aware GMPLS-based Optical Networks", Elsevier Computer Networks, Vol. 52, Issue 10, Page(s):1938-1950, July 2008
[GMPLS-PBBTE01]	D. Fedyk, H. Shah, N. Bitar, A. Takacs, Generalized Multiprotocol Label Switching (GMPLS) control of Ethernet PBB-TE, draft-ietf-ccamp-gmpls-ethernet-pbb-te-02.txt, February 25, 2009
[GN3POwiki]	Project Office page on the GÉANT3 wiki http://wiki.geant2.net/bin/view/GEANT3/PO
[HARMONY-AR]	"Harmony - Advance Reservations In Heterogeneous Multi-domain Environments", IFIP/TC6 Networking 09, Aachen, Germany
[MEF]	http://www.pipenetworks.com/library/brochures/MEF.pdf
[NSIWG]	http://www.gridforum.org/gf/group_info/view.php?group=nsi-wg

[Phosphorus-G2MPLS] "Phosphorus Grid-enabled GMPLS Control Plane (G2MPLS): Architectures, Services and Interfaces", IEEE Communications Magazine, Special issue on Multi-Domain Optical Networks: Issues and Challenges, Vol. 46, Issue 6, Page(s):128 - 137, 2008

[RFC 2702] Requirements for Traffic Engineering Over MPLS http://www.rfc-editor.org/rfc/rfc2702.txt

[RFC 3945] Generalized Multi-Protocol Label Switching (GMPLS) Architecture http://www.rfc-editor.org/rfc/rfc3945.txt

[RFC 4655] A Path Computation Element (PCE)-Based Architecture http://www.rfc-editor.org/rfc/rfc4655.txt

Glossary

AAA	Authentication, Authorisation and Accounting. An architectural framework for configuring a set of three independent security functions in a consistent manner. It is designed to dynamically configure the type of authentication and authorisation on a per-user or per-service basis.
AAI	Authentication & Authorisation Infrastructure. An infrastructure that provides Authentication and Authorisation services. The minimum service component includes identity and privilege management with respect to user and resources.
AC	Automated Client. A profile for software that is not operated directly by humans when it engages in an authentication or authorisation interaction. This category includes elements such as daemons, autonomous servers, and programs subject to automatically scheduled execution.
AL	Activity Leader. The leader of a GN3 activity.
APS	Automatic Protection Switching. An idea for an automatic protection switching system that protects against link or node failure in a network. The system provides a fault recovery system that ensures that the data is properly routed to its destination in the event of a fault in a link.
ASON	Automatically Switched Optical Network. A next step in transport networks that allows for dynamic policy-driven control of an optical or SDH network based on signalling between a user and components of the network.
AutoBAHN	Automated Bandwidth Allocation across Heterogeneous Networks. A system that provides a user-friendly interface for instantiating dynamic circuits over global research and education (R&E) network infrastructures.
BE	Bridging Elements are always integrated within a participating federation and serve as a means of establishing appropriate trust links among federation components and user applications, and of adapting syntax, semantics and procedures used by the participating federations.
BoD	Bandwidth on Demand. A system that tries to deliver a guaranteed capacity, connection oriented service between two end user points.
CC	Connection Controller. A control plane component that services a single subnetwork and provides the abstract interfaces to other control plane components. The connection controller is responsible for coordinating the Link Resource Manager, Routing Controller and both peer and subordinate Connection Controllers.
CDP	Cisco Discovery Protocol. A protocol that enables access of the configuration of other directly connected equipment such as routers and switches. It runs mostly on Cisco equipment.
CGE	Carrier Grade Ethernet. Extends Ethernet from the local area network (LAN) to the wide area network (WAN), makingit possible to connect Ethernet LANs to service provider networks through the same Ethernet interface. It provides a transparent LAN service that connects LANs in different locations as if they were one network.
CP	Control Plane. Responsible for connection and resource management.

DB	Database. An integrated collection of logically related records that provides data for multiple use.
DCN	Data Communication Network. The control-plane in an optical network.
DM	Domain Manager. The AutoBAHN module of that implements the requested service in the form of a provisioned circuit within the local network domain. It deals with the technology-specific details of the particular domain, either by using existing provisioning and management tools or by directly operating upon the data plane. The DM provides all the domain-level functionality to the Interdomain Manager (IDM) and one of its tasks is to translate the technology-specific intra-domain topology to the abstract format used by the IDM.
DWDM	Dense Wavelength Division Multiplexing. An optical technology used to increase bandwidth over existing fiber optic links. It works by combining and transmitting multiple signals at the same time at different wavelengths on the same fiber. As a result, one fiber becomes multiple virtual fibers.
E2E	End-to-end. A principal design element of the Internet that allows network nodes to send packets to all other network nodes, without requiring intermediate network elements to maintain status information about the transmission.
ERO	Explicit Route Object. Identifies the route from head-end to tail-end.
EPL	Ethernet Private Line. A dedicated point-to-point connection from one user location to another with guaranteed bandwidth and payload transparency end-to-end.
ERS	Ethernet Relay Service. Enables multiple instances of service to be multiplexed onto a single user user-to-network (UNI), so that the UNI can belong to multiple ERSs.
EVC	Ethernet Virtual Connections. A logical relationship between Ethernet UNI in a provider-based Ethernet service.
EWS	Ethernet Wire Service. A point-to-point connection between a pair of sites. It is typically delivered over a shared switched infrastructure within the service provider network.
FEC	Forward Error Correction. A system of error control for data transmission.
FPP	Federation Peering Points serve as a means of publishing metadata about a federation through the MDS.
GMPLS	Generalised Multi-Protocol Label Switching. The ability to route a data transmission based on a wavelength of light that carries it. The routing device only analyses wavelengths (light frequencies) to make its forwarding choice. GMPLS extends MPLS to support optical networks.
IaDI	Intra Domain Interface. An interface managed by NMS.
IDM	Inter-Domain Manager. The AutoBAHN module responsible for processing each BoD service request and propagating the accepted request either locally to the DM or to another IDM in a neighbour domain. The service request is forwarded along a chain of domains until the service endpoint is reached.
IdP	Identity Provider. A service hosted by an organisation which publishes electronic identity information for user that have some relationship with the organisation.
IETF	Internet Engineering Task Force. A non-membership, open, voluntary standards organisation dedicated to identifying problems and opportunities in IP networks.
IGP	Interior Gateway Protocol. A category of routing protocols that run within a single area such as a local area network (LAN).
IP	Internet Protocol. A protocol used for communicating data across packet-switched internetworks.
IrDI	Inter Domain Interface. An interface between nodes at the border of the domain.
JRA	Joint Research Activity. A GN2 or GN3 research activity.
LAN	Local Area Network. A group of computers and associated network devices that share common communication lines in close proximity. Most local area networks are built with hardware such as Ethernet cables, network adapters and switches. Wireless LAN also exists.

LER Label Edge Routers. A device placed at the edge of an MPLS domain that uses routing information to assign labels to datagrams, and then forwards them to the MPLS domain.

LDAP Lightweight Directory Access Protocol. A software protocol that allows resources such as files and devices to be located in a network. LDAP is lightweight version of Directory Access Protocol (DAP) without security features.

LLDP Link Layer Discovery Protocol. A protocol used by network devices to announce their identity and capabilities on the local network. Information obtained with LLDP is kept in the device and can be obtained through the Simple Network Management Protocol (SNMP).

LMP Link Management Protocol. A protocol designed for easier configuration and management in optical networks devices. LMP offers automatic configuration of such devices, negotiation of capabilities and localisation of faults.

LS Lookup Service. Allows independent services to register and provide detailed information of their capabilities, and then make themselves visible to other services that are registered with the LS.

LSA Link State Advertisements. A basic communication form for the OSPF routing protocol. It spreads a router's local routing topology to all other local routers in the same area.

LSP Label Switched Path. A path that runs across an MPLS domain, set up by a signalling protocol such as LDP or RSVP-TE.

LSR Label Switching Router. A router located within an MPLS network. Its main task is to switch labels used to route packets.

MEF Metro Ethernet Forum. A global industry alliance that aims to accelerate the worldwide adoption of carrier class Ethernet networks and services.

MDS Metadata Service. A means of storing and providing metadata about eduGAIN interfaces, such as identity providers (IdP), Attribute Authorities (AA), Service Providers (SP) and others. It serves mainly to locate appropriate identity providers able to identity a certain entity in a given federation.

MPLS Multi-Protocol Label Switching. A highly scalable, protocol-agnostic, data-carrying mechanism. In an MPLS network, data packets are given labels. Packet forwarding choices are made according to the contents of this label instead of inspecting the whole packet.

MTU Maximum Transmission Unit. The size (in bytes) of the largest protocol data that can be forwarded. The MTU may be assigned by standards (i.e. Ethernet) or be negotiated at connect time.

NMS Network Management System. Both hardware and software used to monitor and administer a network.

NREN National Research and Education Network. An internet service provider focusing on the needs of the research and education communities within a country.

NSI Network Service Interface. The main goal of the NSI is to facilitate interoperation between Grid users, applications and network infrastructures that span different service domains, through the development of abstract messaging and protocols.

OAM Operations Administration and Maintenance. Commonly used term in the context of computer networks to describe the process, activities, tools and standards related to administering and maintaining any system.

OGF Open Grid Forum. An open community established to adopt applied distributed computing.

OIF Optical Internetworking Forum. Promotes the development and deployment of interoperable networking solutions through the creation of Implementation Agreements (IAs) for optical networking products and network processing elements.

OSPF Open Shortest Path First. A routing protocol used to determine the best path for routing based on distance between nodes and other constraints. OSPF is an interior gateway protocol (IGP) to work within an autonomous system.

OTN Optical Transport Network. A set of Optical Network Elements connected by fibre links that have transport, multiplexing, switching, management and survivability of optical channels capabilities.

OTS Optical Transport Sections consists of the optical multiplex section (OMS) along with an additional optical supervisory channel (OSC).

PBB Provider Backbone Bridging. A set of architecture and protocols for routing a user network over a provider's network.

PCC Path Computation Client. Any client application requesting a path computation to be performed by a PCE.

PCE Path Computation Element. An entity that computes a network path based on a network graph and applying computational constraints. A PCE might be a network node, network management system, or application that has knowledge of the network resources.

PCEP Path Computation Element Protocol. Defines communication details between PCCs and PCEs, and between cooperating PCEs.

PF Pathfinder. A module that implements an algorithm for finding all possible paths between nodes that make up a network. A path-finding algorithm can accept several constraints to return only paths that meet these constraints. For example, a node may be excluded from path-finding calculation in order to find paths that omit that node. Another constraint could be minimum bandwidth that searched path must provide.

PIP Premium IP. A service that offers network priority over other traffic. Premium IP traffic takes priority over all other services, such as Best Efforts (BE) and Less Than Best Efforts (LBE). During network overload Premium IP traffic gets a better and guaranteed level of network performance

POTS Plain Old Telephony Systems. Traditional telephone service for analog voice transmission.

QoS Quality of Service. Allocates different priorities to different applications, users or data flows, and guarantees some level of performance to data flow.

RSVP Resource Reservation Protocol. A protocol used by a host to request specific QoS from the network for particular data flows. RSVP is also used by routers to deliver QoS requests to all nodes along the path of the flow and to create and maintain requested services. RSVP messages are used to reserve resources in each node along the path.

RWA Routing and Wavelength Assignment. The main goal of RWA is to maximise the number of established optical connections. Each connection request is assigned a route and wavelength. Two connections requests can only share the same optical link if a different wavelength is used.

SAML Security Assertion Markup Language. An extensible and customizable XML-based framework for communicating user authentication, entitlement and attribute information between an identity provider (a producer of assertions) and a service provider (a consumer of assertions). SAML is authored by the Security Services Technical Committee of the Organization for the Advancement of Structured Information Standards (OASIS).

SCN Signalling Communication Network. Corresponds to the IP control plane topology enabling routing exchanges between GMPLS controllers (i.e. the routing adjacencies).

SLA Service Level Agreement. Part of a service contract where the level of service is formally defined. Most often, the SLA is used to refer to the contracted delivery time or performance. The SLA may specify the levels of availability, performance, operation and other attributes such as billing.

SDH Synchronous Digital Hierarchy. A group of fiber-optic transmission rates that can transport digital signals with different capacities over optical fiber or light-emitting diodes (LED).

SNMP Simple Network Management Protocol. A protocol that enables a management station to configure, monitor and receive trap(alarm) messages from network devices such as servers, printers, switches and routers on an Internet Protocol (IP) network.

SONET	Synchronous Optical NETworking. A group of fiber-optic transmission rates that can transport digital signals with different capacities over optical fiber or light-emitting diodes (LED). SONET is widely used in the United States and Canada.
T2	Task 2 (Hybrid Network Provisioning) of the GN3 JRA 2 activity (Multi-Domain Network Service Research).
TE	Traffic Engineering involves adapting the routing of traffic to the network conditions by dynamically analysing, predicting and regulating the behaviour of transmitted data in order to provide efficient use of network resources.
TED	Traffic Engineering Database. A repository for flooded topology link-state information learned from Interior Gateway Protocols and Traffic Engineering attributes that belong to the state or resources associated with each link.
TNA	Transport Network Address. The address of an entity within the transport network.
TP	Technology Proxy. A component of the AutoBAHN system that provides bidirectional access to different network technologies and Network Management Systems. In the BoD architecture, technology proxy creates an interface that abstracts a specific technology on a technology domain to the abstract format used by the BoD system. To achieve this, parameters that are used to specify the end-to-end BoD service are mapped onto technology-specific parameters.
UbC	User behind a Client. Applicable to web services clients that run independently of an HTTP (Web or application) server, but under the direct human control. These clients are standalone programs that can be freely installed at individual workstations or shared computers, and are operated under direct control of their users.
UNI	User to Network Interface. The service control interface between a client device and the transport network.
URN	Uniform Resource Name. A Uniform Resource Identifier (URI) that employs the urn scheme and does not imply availability of the identified resource.
VLAN	Virtual Local Area Network. A collection of hosts that operate within the same broadcast domain even though they are placed in different physical locations.
VLL	Virtual Leased Line. Provides Ethernet based point-to-point communication over IP/MPLS networks.
VPN	Virtual Private Network. A private network that uses a public network (mostly the Internet) to connect remote endpoints together. Instead of using dedicated or leased lines, the links between nodes of a virtual private network are formed over logical connections.
WE	Client in Web Container. This profile is applicable in cases where a certain software component (the client) is accessed by end users through a web container (i.e. an application server), and the clients acts on behalf of the end user when requesting services to other components.
WG	Working Group. A collaboration of researchers carrying out new research initiatives. The body of a Working group usually consists of experts on a given subject. Working groups are sometimes also called task groups or technical advisory groups.