

12.09.06

# Deliverable DJ2.1.1,3: Revised GÉANT2 Security Recommendation and Policy



## Deliverable DJ2.1.1,3

Contractual Date:	31/08/06
Actual Date:	12/09/06
Contract Number:	511082
Instrument type:	Integrated Infrastructure Initiative (I3)
Activity:	JRA2
Work Item:	1 (One)
Nature of Deliverable:	R - Report
Dissemination Level	RE - Restricted
Lead Partner	DANTE
Document Code	GN2-06-205v3

**Authors:** M. Mogensen (DANTE), D. Kalogeras (GRNET), J. Mohacsi (HUNGARNET), H. Nussbacher (IUCG). M. Garcia (DANTE), R. Sabatino (DANTE)

## Abstract

This document is an update to the 2nd Security Recommendation and Policy deliverable, DJ2.1.1,2.

# Table of Contents

0	Executive Summary	iv
1	Introduction	1
2	Policies	2
3	Layer 3 Services	4
4	Layer 2 Services	6
5	Layer 1 Services	8
6	End-to-end Services	10
7	Conclusion	12
8	References	13
9	List of Acronyms	14
10	Appendix A – Best Practice Documentation	16
10.1	Layer 3 services	16
10.2	Overview of layer 3 services	17
10.3	Securing Layer 3 equipment	17
10.3.1	Routers	18
10.3.2	Protecting the management plane of routers	20
10.3.3	Protecting control plane of routers	21
10.3.4	Protecting the forwarding plane of the router	25
10.4	Layer 2 services	27
10.5	Layer 1 services	28
10.5.1	Data Communications Network	29
10.5.2	SDH equipment	29
10.5.3	WDM equipment	30

## Table of Figures

<b>Figure 5.1:</b> Router view of data, control and forwarding plane	19
<b>Figure 5.2:</b> uRPF strict operation	26
<b>Figure 5.3:</b> Loose uRPF	27
<b>Figure 5.4:</b> Martini layer 2 circuit – Atlas project	27
<b>Figure 5.5:</b> SHD switch management	30
<b>Figure 5.6:</b> Functional diagram of OSC	31
<b>Figure 5.7:</b> Spectrum Analysis showing OSC	32

## 0 Executive Summary

This deliverable presents an update to the 2nd revision (DJ2.1.1,2) which was submitted in May 2006.

There is ample experience in securing the IP network and the IP based services, gained over the last 10 years and therefore we do not foresee any changes since revision 2. As for the new services introduced by GÉANT2, 10Gbps multi-domain point to point (or end to end) services are being made available from July 2006 onwards, whilst Gigabit Ethernet services using the switching platform (Alcatel MCC) will be delivered from October 2006 onwards. Hence, to date there is no working experience with these services to use as a basis to update the security recommendations and policies, as set out in revision 2. As a result, the remainder of this document is mostly unchanged from the second revision.

It is expected that working experience with these services will be available in time for fourth revision of the deliverable.

The deliverable is organised in a way that the policies and the technical analysis have been separated in order to remove technical contents from the policies as far as possible. The policies are the main part of the document and the technical analysis with the best practice documentation is attached in Annex 1. The security policies are based on a technical analysis of the security requirements of the services so Annex 1 provides further technical documentation to support the policies.

The security policies are divided into layer 1, layer 2 and layer 3 services where end-to-end services are covered separately because they can span more than one layer. The policies focus on the demarcation points of the services in order to define responsibility at the administrative domain border. Procedures for dealing with incidents are defined. The existing policies for layer 3 and layer 2 services are documented and a first set of provisional set of policies for layer 1 services and end-to-end services are defined.

Annex 1 contains best practice documentation and a technical analysis of the security requirements for GÉANT2 services. The attack surface of each type of equipment is analysed in order to present a best practice recommendation for securing each type of equipment. The technical analysis is based on an analysis of the management plane, control plane and forwarding plane.

This document will be updated in DJ2.1.1,4 Updated Security Recommendation and Policy.

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

# 1 Introduction

This document is an update to deliverable DJ2.1.1,2 Initial Security Recommendation and Policy which was submitted January 2005.

The purpose of the update is to revise the recommendations and policies defined in the initial deliverable. The initial deliverable focused on the security policies that had been implemented at the time of writing and it presented the expectations for GÉANT2.

In comparison, this update focuses on defining a security policy for the services that are new in GÉANT2 compared to GÉANT. At the time of writing, not all services planned for GÉANT2 are yet operational so the policies presented here are provisional policies that are expected to be revised with operational experience and if the service requirements change. The revised policies are planned to be presented in the next deliverable update.

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

## 2 Policies

The security policies are based on a technical analysis of security requirements for the new services. The technical analysis and the best practice documentation have been separated from the policies and are included in Annex 1. This is primarily to focus this document on the security policy but also to make it simple to maintain the best practice documentation so it can be updated with collaborators. To facilitate this process, the technical analysis and best practice documentation is maintained in the JRA2 Document Management section of the GÉANT2 Wiki [GN2-JRA2-Wiki]. For some new services, there is no established good practice yet. In that case, the security recommendation is based on a technical analysis. The aim is to be vendor independent as far as possible.

In the following, the services offered by GÉANT2 are divided using the OSI model [OSI7498]. The OSI model is used here because it provides a layered model for defining how to secure services. Most services clearly belong to a specific layer in the OSI model as shown in the table below:

OSI model layer	Services	Type of equipment
Layer 3	IP, IPv6, Multicast, QoS	Routers, Workstations
Layer 2	Layer 2 circuit	Routers, Switches
Layer 1	Lambda, SDH circuit, Point to point Ethernet	SDH, WDM

**Table 1:** Services per OSI model layer

In the following, a security recommendation and policy for services will be presented. Compared to the initial security recommendation and policy, the status is:

- Layer 3 services were extensively covered in the initial deliverable and the results will be summarised here.
- Layer 2 services were partly covered in the initial deliverable but they will be covered separately as the security requirements are significantly different than layer 3 services

- Layer 1 services were only covered as expectations in the initial deliverable because at the time of writing it had not been decided how they would be implemented. This is now clear, so an initial security recommendation and policy is presented here.

End to end services can span more than one layer if they are provided by more than one type of equipment. As an example, an end-to-end service between user A in NREN A and user B in NREN B could use IP (layer 3) for some part of the path, a layer 2 VPN (layer 2) for a part of the path and WDM (layer 1) for a part of the path. How end-to-end services will be provided in practice will be clearer once these services become operational. However, it is still necessary to define a security recommendation for end-to-end services to have a common framework, especially for service handover at the administrative domain border.

The security recommendation and policy presented here is based on an analysis of the management plane, control plane and forwarding plane of each type of service as defined in the description of work.

### 3 Layer 3 Services

Layer 3 services have a clearly defined domain boundary which is typically as illustrated below.

Site <-> Regional Network <-> NREN <-> GÉANT2 <-> NREN <->Regional Network <-> Site

On the one hand this has the benefit of making the responsibilities and demarcations clear so each organisation is responsible for their own domain. This is possible because layer 3 services have a clear boundary between administrative domains. On the other hand, layer 3 services have a significant overlap between the management plane, control plane and forwarding plane which makes the service vulnerable but there are several good practices which greatly reduce the risk. The following is focused on routers – workstation security is covered later in this chapter.

The technical annex present specific recommendations for securing control plane protocols needed to provide layer 3 services. Some security measures are mandatory for connecting to GÉANT2, for example the use of MD5 authentication for BGP peerings.

To protect the forwarding plane, the following pro-active security measures have been implemented on the border between GÉANT2 and NRENs:

- uRPF (Unicast Reverse Path Forwarding) which ensures that only traffic using legitimate IP addresses is allowed into the GÉANT2 core. To avoid dropping legitimate traffic, traffic is first checked against the routing table to see if it has been received on the interface that has the best route back to the source. If this fails, NREN traffic is checked against the RIPE AS macro of the NREN. The reason for doing two checks is that legitimate traffic could fail the first check.
- In addition to uRPF, all packets entering GÉANT2 with bogon source addresses [Bogon] are discarded. It is believed that as many as 75% of brute force DoS attacks use bogon source addresses; these filters protect all NRENs from such attacks from other NRENs and the wider Internet.
- On all GÉANT2 NREN access ports policer limits have been implemented to control the traffic that would be destined for the GÉANT2 equipment (any equipment allocated with a prefix from within the GÉANT2 /19 core address space). The purpose of this is twofold: firstly to protect the GÉANT2 infrastructure and related services from DoS attacks and secondly to prevent the scanning activity that is usually a precursor to hacking activity. Production traffic from NRENs is not affected in any manner.

Re-active security measures have also been implemented:

- The GÉANT2 NOC can filter NREN traffic on request. If needed, the GÉANT2 NOC can determine the sources of an attack and can inform NRENs when an attack has stopped. To ensure only legitimate filtering request are implemented, an authentication procedure has been implemented that NRENs have to follow.

Co-located Project Workstations are discussed here because of an increasing demand to place workstations in the core, typically for project use. So far, measurement workstations have been deployed by SA3 and JRA1 in the GÉANT2 core and in NREN networks.

Workstation security relies on keeping two separate sets of components secure – the operating system itself and the software components running on top of it. The skill set required of the people maintaining these two sets of components is substantially different, and it is therefore anticipated that they will generally be different individuals rather than one individual with joint responsibilities. This is a point to note, as it is important that the person responsible for a component has enough knowledge of how that component works and where possible vulnerabilities are likely to occur, in order to act appropriately.

- To ensure that a workstation is being correctly maintained, it is vital that there is a single point of contact for each workstation responsible for the maintenance of these components. This person is then in charge either of maintaining the machine themselves, or delegating responsibility onto other people who are better positioned to respond.
- The activity leader is ultimately in charge of workstation security, even if the responsibility for the work itself is delegated. Any requests for changes to network security should be approved by the activity leader.

## 4 Layer 2 Services

The only layer 2 service deployed via GÉANT2 is layer 2 VPNs, more specifically Martini layer 2 circuits [Draft Martini], which provides point to point circuit emulation using MPLS. Layer 2 circuits have been extensively deployed to provide support for projects with high bandwidth demands and for projects that require to be directly connected at layer3, i.e., distributed testbed.

A current best practice for securing layer2 circuits has been implemented. The technical analysis and the best practice rules are documented in the technical annex. Layer 2 circuits do not have inter-domain model which makes it necessary to run intra-domain protocols (RSVP and LDP) across domain borders which is not what the protocols were designed for.

The following policies apply to layer 2 circuits:

- The LSPs (Label Switched Path) will be configured on the GÉANT2 routers and the NREN routers. Those LSPs will be stitched on the GÉANT2 routers using Juniper CCC functionality to build an end-to-end tunnel. The LSPs coming from NREN routers are not allowed to transit GÉANT2, meaning that they cannot have both ingress point and egress point outside GÉANT2 without DANTE's knowledge.

This policy ensures that any unauthorised LSPs can be identified quickly and action can be taken to remove them.

- The layer 2 circuit has to be set up with RSVP signalled LSPs. This is to allow traffic engineering and to allow LSPs to be established across a domain boundary.

Due to the high bandwidth requirements of layer 2 circuits, it is usually necessary use traffic engineering to route the LSPs on specific paths which requires RSVP. In addition, to be able to run LSPs across a domain border (two different traffic engineering domain) RSVP is required.

- The NRENs involved have to configure the bandwidth and priority values as well as the strict and loose hops requested by DANTE. If any other LSP parameters are used, DANTE will need to have knowledge of it and will apply the necessary configuration changes. GÉANT2 will have the option of limiting the available bandwidth in the access to ensure that an LSP that has ingress point outside GÉANT2 does not reserve more bandwidth than necessary.

RSVP has no inter-domain model which means any router in the RSVP domain can freely reserve available bandwidth in the domain which could make it impossible to establish new LSPs. To avoid this risk, the reservable bandwidth on an NREN access interface can be lowered. Note that this limitation only limits the RSVP bandwidth reservation request for establishing the LSP and does not restrict the effective bandwidth of the layer 2 circuit.

## 5 Layer 1 Services

Layer 1 services are a new type of service for GÉANT2. It is important to stress that these services are not operational yet and the recommendation and policy presented here is a first version which will be modified when more operational experience has been gained.

Layer 1 services can be provided by two types of equipment:

1. SDH (Synchronous Digital Hierarchy) equipment which can provision SDH circuits and point to point Ethernet circuits.
2. WDM (Wavelength division multiplexing) equipment which can provision dedicated wavelengths (lambdas)

Layer 1 services are used to provide not only NREN connections – the majority of the IP trunks in the GÉANT2 core are provisioned using GÉANT2 WDM equipment at the time of writing.

The WDM and SDH equipment in the GÉANT2 core are managed by a centralised NMS (Network Management System). It is essential to secure the NMS from attacks because all higher layer services could be disrupted if management access to the NMS was lost. The NMS servers are protected by firewall filters and the user logins are audited. Only DANTE, GÉANT2 NOC (Network Operations Centre) and the vendor have direct access to the NMS. For monitoring purposes, alarms are exported using SNMP and the results are available to the service users via a proxy server. This avoids the need to allow direct access to the NMS servers.

The network used for communication between the equipment and the NMS is called a DCN (Data Communication Network). For GÉANT2, the DCN carries both IP and OSI traffic. The DCN for the GÉANT2 core is analysed in the Technical Annex. A DCN is typically proprietary but the functionality between vendors is very similar. The DCN does not run on any client interfaces which is a significant difference to Layer 2 and Layer 3 services because it separates the forwarding plane (client interface) from the management plane (DCN). This means user data on a user client interface cannot disrupt the DCN. On a client SDH interface, DCC signalling should be disabled.

It is still essential to secure the management interfaces on the network elements. It is recommended to use private IP addresses or OSI addresses for the management interfaces if this is feasible because these should not be routable across a domain border.

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

SDH and WDM equipment does not normally use a control plane because the NMS is centralised. For a GMPLS/G.ASON [G.8080] deployment, the network elements will create a distributed control plane and take over some functions from the management plane. The new control plane will only run via the DCN so user data cannot disrupt the DCN. In case UNI [UNI] is deployed to allow user provisioning of services, a separate client interface is required for signalling. Restrictions can be set on which resources a user can provision via UNI. If UNI will be deployed, recommendation on how to secure a UNI interface will be produced.

## 6 End-to-end Services

It is expected that GÉANT2 end-to-end services will be point-to-point services between two end sites in two different countries. To provision and operate an end-to-end service, several organisations need to be involved. The demarcation points between each organisation need to be clearly defined to make the responsibility of the service clear.

An AUP (acceptable use policy) for end-to-end services needs to be defined. The use of an end to end service could be a security problem in the sense that it is being used for transiting data which is a usage policy matter more than a matter of protecting the service against disruption.

It is possible that some end-to-end services will be provisioned using only one type of equipment but end-to-end services can also be hybrids in the sense that they use more than one type of equipment between the two end points. As an example an end to end service could be provided using IP routing, layer 2 circuit and WDM wavelength depending on what equipment is available between the two end sites.

To secure an end-to-end service requires clear demarcation points for each part of the service. For the countries without a GÉANT2 PoP (Point of Presence), the demarcation point is the NREN router interface. For a country that has a GÉANT2 PoP, the demarcation points to the local NREN are:

- For IP services – the NREN GÉANT2 router interface (and backup router interface if in use)
- For layer2 circuits – the NREN GÉANT2 backup router interface (or primary router interface if not in use). A separate interface can be used with specific permission from DANTE Operations.
- For SDH or WDM services – an ODF position in the GÉANT2 PoP specified by DANTE Operations.

The following is a proposed security recommendation for end-to-end services.

- A demarcation point needs to be clearly agreed between two organisations that connect to each other with the purpose of providing an end-to-end service.
- Each organisation involved is responsible for securing and monitoring the end-to-end service up to its demarcation points.

- Each organisation involved may be requested to give monitoring access to their equipment. If no direct monitoring access is possible, access to monitoring data can be given via a proxy server.
- An end-to-end service is seen as a point-to-point service between the two end sites. The data transported by the end-to-end service can only be transmitted by the end sites. The data send via an end-to-end service is the responsibility of the end site that transmits the data. The end sites have the responsibility of avoiding any breach of acceptable use policies.

## 7 Conclusion

The updated security recommendation and policy presented here will be updated again in D.J.2.1.1,3 when end-to-end services will be in general deployment and operational experience will have been gained.

If new service requirements are defined, JRA2-WI1 will analyse the requirements and present security recommendation which should lead to new security policies. Naturally, this is an ongoing process where security policies are modified and disseminated when they change.

The security policies are disseminated as part of the “NREN Operational Procedures” which are available on the GÉANT2 website [NREN OPS]. Any changes are announced to the GÉANT2 APM list.

The technical annex is meant to be a shared resource that NRENS and other GÉANT2 service users can access and help to maintain.

## 8 References

[Bogon] <http://www.cymru.com/Bogons/>

[Draft Martini] <http://www.ietf.org/internet-drafts/draft-martini-l2circuit-encap-mpls-10.txt>

[G.8080] ITU-T G.8080/Y.1304 - <http://www.itu.int/itudoc/itu-t/aap/sg15aap/history/g8080/g8080.html>

[GN2-04-148] M. Mogensen, M. Garcia, D. Kalogeras – Initial Security Recommendation, January 2004.

[GN2-JRA2-Wiki] <http://wiki.GÉANT2.net> (login is required)

[OSI7498] [http://standards.iso.org/ittf/PubliclyAvailableStandards/s020269\\_ISO\\_IEC\\_7498-1\\_1994\(E\).zip](http://standards.iso.org/ittf/PubliclyAvailableStandards/s020269_ISO_IEC_7498-1_1994(E).zip)

[NREN OPS] <http://intranet.GÉANT2.net/server/show/nav.922>

[UNI] <http://www.oiforum.com/public/impagreements.html#UNI>

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

## 9 List of Acronyms

APM	- Access Port Manager
CLNS	- ConnectionLess Network Service
DCN	- Data Communication Network
DoS	- Denial of Service
ISO	- International Organization for Standardization
LDP	- Label Distribution Protocol
LSP	- Label Switched Path
MPLS	- Multi Protocol Label Switching
NMS	- Network Management System
NOC	- Network Operations Centre
ODF	- Optical Distribution Frame
OSI	- Open System Interconnection
POP	- Point of Presence
QoS	- Quality of Service
RSVP	- Resource Reservation Protocol
SDH	- Synchronous Digital Hierarchy
URPF	- Unicast Reverse Path Forwarding

- VPN - Virtual Private Network
- WDM - Wavelength Division Multiplexing

## 10 Appendix A – Best Practice Documentation

GÉANT2 has a significant focus on end-to-end services compared to GÉANT. These end-to-end services take different forms and will be developed into the various activities in GÉANT2. Securing the equipments individually is indispensable. However, securing each and every piece of equipment is not enough for a secure research network, rather, “securing” is a continuous process to always improve security to match the provided services.

Some new, and as yet unprovided, end-to-end services are expected to make use of the type of equipment which is already deployed or being deployed on GÉANT and NREN networks. This will be:

- End-to-end and multi-domain services such as end-to-end QoS and bandwidth on demand. Those services will require applications deployed in a distributed way and that need to be accessed across multiple domains.

Besides new equipment, IP routers will continue to deliver the same services as on GÉANT and NREN networks and need to be secured as well. The majority of the security measures for protecting equipment and services are implemented using the filtering or so called firewall features of the core IP routers, because data can be filtered and logged at line rate with great granularity.

### 10.1 Layer 3 services

This technical annex presents recommendations services in GÉANT2 and NREN networks. Layer 3, 2 and 1 equipment will be analysed in order to present best practice for securing the management plane, control plane and forwarding plane. The aim is to be vendor independent as far as possible.

## 10.2 Overview of layer 3 services

The forwarding plane in GÉANT2 is provided by the GÉANT2 backbone routers and the NREN routers. It is common that those routers can handle traffic at line rates. This means that a router does not degrade its performance no matter how many packets are switched through the switching fabric. A typical performance for line rate for Gigabit Ethernet connection for 64 byte packets would be 2 million packets per second. This is not the case for routers which are called software routers. This type of router is present in a lot of campus networks and some part of the NREN networks; it handles connectivity demands ranging from multiple of Mbps to STM-1, STM-4 and gigabit speeds.

A forwarding plane attack could appear near the "edge of the network" by exploiting a Distributed Denial of Service attack vector. This is sometimes achieved by spoofing the source addresses but nowadays this tends to be less frequent – since widespread implementation of ingress filtering ala BCP38. Another method to achieve a forwarding plane attack would be improper handling and processing of ICMP and other control packets which can lead to a commonly exploitable phenomenon known as "ping of death". This means it is important to consider if it is possible to detect such forwarding plane attacks in the core of networks and how to respond. In some cases forwarding plane attacks can be associated with a short time increase of flows destined to the same subnet or host. Forwarding plane attacks look like normal network traffic in the core of the network which means it is quite difficult to detect.

The fact that core networks cannot easily detect forwarding plane attacks and are not exposed to those attacks does not mean that NRENs should stay indifferent to such potential incidents. It is obvious that that the attackers would like to exploit the GÉANT2 infrastructure to deploy attacks.

## 10.3 Securing Layer 3 equipment

In the following section, each type of equipment will be analysed and recommendations for a suggested best practice for securing them will be presented. For each type of equipment it will be discussed how to secure the management plane, the control plane and the forwarding plane. In general, the three planes can be defined as follows:

- The management plane refers to the management features of the equipment where configuration changes can be applied and the equipment is monitored. In general, the management plane needs to be secured so only authorised users can access the management functions. Access is either in-band (on the same interface used for forwarding data) or out of band (OOB – on an interface not used for data forwarding). A security breach on the management plane can cause all services to be disabled.
- The control plane refers to the features that control the operation of services. The control plane is distributed and runs between equipment using protocol signalling in order to maintain a service and to deal with fault conditions. To secure the control plane, each service and their

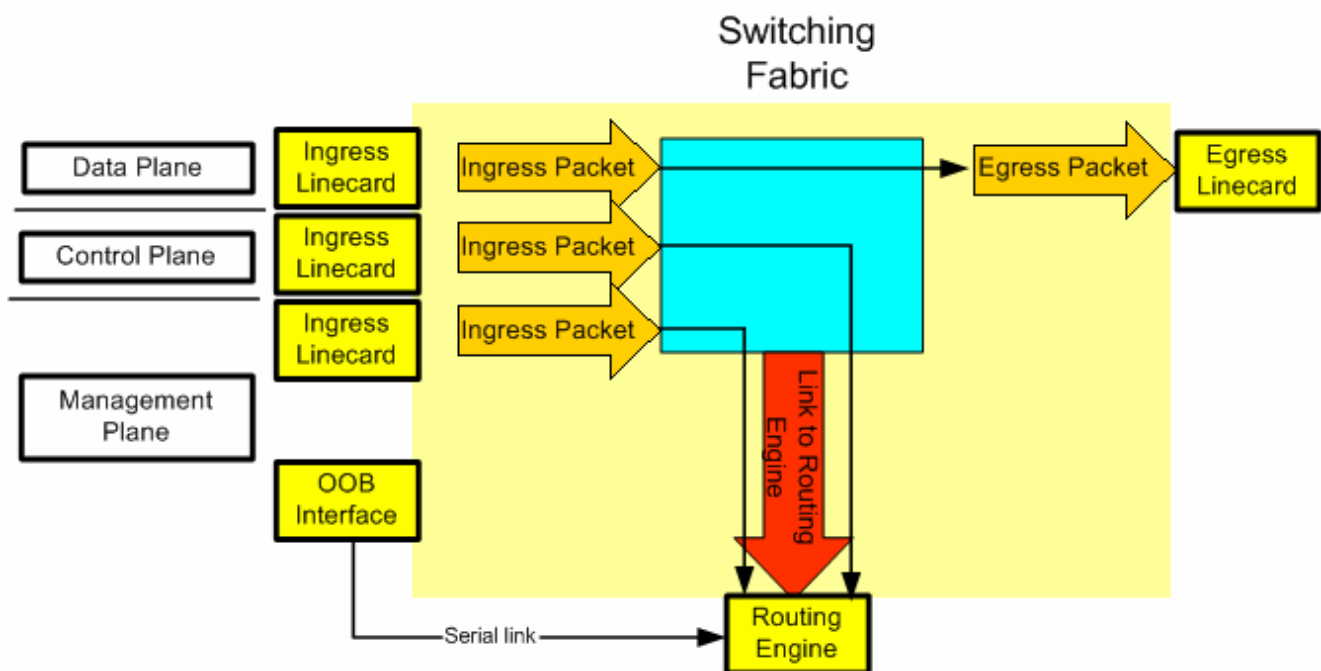
respective protocols need to be analysed individually for their vulnerabilities. An attack on the control plane can potentially disable services temporarily until the attack can be filtered.

- The forwarding plane refers to the features that move data from one interface to another. Forwarding is usually based on forwarding tables (maintained by the control plane) and can be done either in hardware (specially designed forwarding modules) or software (CPU processing). The forwarding plane can be attacked by overloading equipment with more traffic that it can forward, typically in the form of a distributed denial of service attack (DDoS) which can lead to loss of data. In a high capacity backbone network a DDoS attack may not have any operational impact on the backbone itself, if the attack is not large enough to disrupt the forwarding plane, but DDoS attacks should be detected and stopped in the interest of the affected end users.

In order to minimise the impact of attacks, equipment vendors try to separate the management plane, control plane and forwarding plane as much as possible, for example by using different physical or logical interfaces for each plane. However, an attack on the management plane can still put both the control plane and forwarding plane out of service, and an attack on the control plane can cause the forwarding plane to stop forwarding data for one or more protocols.

### 10.3.1 Routers

IP backbone routers have an essential role in providing services because they have both rich functionality for providing services and high forwarding performance. In a core router, the three planes are typically implemented as shown in the diagram below:



Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

**Figure 10.1:** Router view of data, control and forwarding plane

- The management plane gives user access into the routing engines (one active, one standby) that is accessed either in-band via loopback or via a virtual terminal interface (Telnet, SSH, etc) or via a serial interface (OOB).
- The control plane is used by the router to operate the services running on the router so that each protocol can communicate with the corresponding protocol on other routers. The control plane traffic accesses the routing engines in-band therefore packets destined to the routing engine are forwarded to the routing engine via an internal link between the switching fabric and the routing engine. The protocols used will be analysed individually below. Protocol traffic to the routing engines can be rate-limited to reduce the impact of a DDoS attack on the control plane though this has to be implemented with great care. However, the control plane can still be attacked by a DDoS attack if the control plane packets between two routers are lost due to congestion caused by such an attack. The more bandwidth available, the less likely it is that a DDoS attack can cause control plane packets to be lost due to congestion because it is more difficult for an attacker to cause congestion.

In addition to DDoS attacks, the common control plane attacks are based on injecting control information to the control plane, that eventually causes wrong or erroneous responses from the routing engine.

- The forwarding plane is implemented in hardware in the form of switching fabric modules (usually redundant) each of which holds a copy of the router's forwarding table. The active routing engine uses the routing tables to build forwarding tables, which are copied to the switching fabric modules when any changes are detected. Forwarding plane traffic, which is not management plane traffic or control plane traffic, is never sent to the routing engine because the forwarding tables are sufficient to switch packets from one interface to another. In this way, the routing engine does relatively little packet processing, and data forwarding is not dependent on the processing power of the routing engine.

The forwarding plane can be attacked by congesting an interface which will cause packet loss if some traffic cannot be forwarded. Like for the control plane, the more bandwidth available, the less likely it is that a DDoS attack can cause packet loss.

The majority of the security measures for protecting equipment and services at the management plane are implemented using the filtering or so called firewall or access control features of the core IP routers, because data can be filtered and logged at line rate with great granularity.

Another option would be to use separate/non-routable address space to address management interfaces.

In this way, the management plane is used by users to manage the router whereas the control plane is used for communication between routers to operate services. Usually, the management plane and control plane are separated from the forwarding plane, for example using an internal fast link. This means that the management

plane and the control plane will only see packets addressed to the routing engine itself– all other packets are handled by the forwarding plane. In this way, the control and management plane are difficult to attack because few packets are processed by the routing engines. However, the routing engines should still be protected because any packets destined for the routing engines needs to be processed. If the link between the forwarding engines and the routing engines is congested, control plane and management plane can be disrupted. Traffic to the routing engine can be rate-limited, though this should be implemented with care, and protocol authentication (with MD5 for example) should be used as far as possible.

Besides this, the interface which gives access to the routing engines can have even more specific filters to secure the management access. The router filter can also be used to protect LAN equipment, first of all workstations. This will be covered in more detail in the workstation section later in this document.

### 10.3.2 Protecting the management plane of routers

There are many ways of improving security of the management plane of routers. In general only authorised users should be able to access the router and limited set of these users should be able to access the management functions. To improve access security, control can be achieved by using infrastructure access control lists (ACLs), infrastructure firewalls or infrastructure VLANs – i.e. separate VLANs for accessing the management interface of the routers. The infrastructure ACLs rely on a protected/un-routable IP address that is used on the management interface of the router. Security is dependent on preventing access to this IP address space at the forwarding plane on each router that belongs to the same administrative control.

Routers run a large number of services where some services use more than one protocol. Below, each protocol is analysed for its attack surface and recommendations are made for securing them.

Many built in services in some router software are not needed in an ISP backbone environment. These features (e.g. echo server, HTTP server) should be turned off in your default configuration.

Some IP features are useful for campus networks but harmful in ISP backbone:

- IP redirection on interfaces should be switched off
- IP directed broadcast should be switched off on all interfaces, otherwise “Smurf” type attacks can be conducted against the management plane.
- Proxy ARP is usually not needed in a backbone environment. Proxy ARP, as defined in RFC1072 is used by the router to help hosts with no routing capability to determine the MAC addresses of hosts on other networks or subnets. Relying on proxy ARP in an Internet backbone router potentially carrying a huge number of MAC addresses could potentially be problematic for the router’s performance.

Proprietary protocols (like Cisco Discovery Protocol - CDP) used for discovering and managing the networks are useful in a small environment, but there is not much sense to run it on an ISP's backbone. It is strongly recommended that CDP be disabled on all public-facing interfaces.

Management plane traffic should be filtered based on source address as much as possible, so that only protocol data from expected sources is accepted.

It is important to warn users that if they connect to the router that only authorised users are permitted to connect. It is wise to inform users with an official warning and to contact the helpdesk, but at the same time, not to reveal too much information about the system and the provided services.

Keep passwords and SNMP community strings in any stored configuration encrypted – preferably in a non-reversible form for non-administrative users. This can prevent attacks against the routers control plane.

Implement a timeout on the management interface: It is recommended to implement a timeout on all management interfaces since each open management session uses precious resources of the routing engine – if it is not used for a long time – free the resources. It will also minimize the risk that the operator leaves his/her terminal logged into the router.

Implement secure authentication to access the management plane. It is strongly advised to use SSH instead of telnet to access the management plane of the router, since SSH is less susceptible to eavesdropping.

From a management point of view it is wise to implement a central Authentication and Authorisation infrastructure, where all authentication and authorisation information is stored. The advantage of the central scheme is that some routers do not support secure storing of password in local configuration files and all access are logged to the AAA server. It is of course important that communication between the routers and AAA infrastructure is properly secured.

It is also recommended to track all commands or a limited set of commands typed into the router. The AAA infrastructure usually allows router command auditing.

It is also recommended to enable logging the system messages to an external syslog server for later perusal.

### 10.3.3 Protecting control plane of routers

Control plane traffic should be filtered based on source address as much as is possible, so only protocol data from expected sources is accepted.

It is recommended to filter in-bound traffic on all ingress network points so that any traffic with a destination IP address matching the core address space is dropped unless it is specifically permitted. In the out-bound direction, all control plane traffic should be allowed. Bogon and Martians filtering should be applied on external interfaces. "Bogon traffic" is traffic with a source address that belongs to an address range that has not been

allocated for use. “Martian traffic” is traffic with a source address that belongs to an address range that is either reserved or for special use (RFC1918 private IP address space for example).

In IPv4 it is easier to filter out (deny) packet originated from bogon routes, while in IPv6 it is easier to allow legitimate packets as shown in the table below:

Rule	Meaning
deny 2001:db8::/32 any	Filter out documentation prefixes
allow 2001::/16 any	Allow RIR allocated prefixes 1
allow 2003::/16 any	Allow RIR allocated prefixes 2
allow 2002::/16 any	Allow 6to4 relay prefix
allow 3ffe::/16 any	Allow 6Bone prefixes - deprecated after 6th June 2006
deny any any	Deny everything else

**Table 10.1:** Bogon Filtering Firewall Rules in IPv6

The table below focuses on the control protocols running on the GÉANT2 core routers and NREN and campus routers. The term “external interface” refers to an interface facing another network. The term “internal interface” refers to an interface facing another router in the same network.

Protocol	Recommendation for securing
BGP	<p>It is recommended that BGP (Border Gateway Protocol) peers use MD5 authentication based on a shared secret.</p> <p>It is recommended to use Generalised TTL Security Check mechanism (RFC 3682) for BGP, which introduces a lightweight security mechanism to protect external Border Gateway Protocol (eBGP) peering sessions from CPU utilization-based attacks using forged IP packets.</p> <p>It is also useful to filter BGP packets by source/destination addresses so only specific peers are able to send BGP packets and all other BGP packets are dropped.</p> <p>In addition, a specific prefix list per BGP peer should be derived from the RIPE database in order to only accept the routes that that peer should be advertising.</p> <p>A maximum number of prefixes should be configured per peer, whereby reaching a certain limit would alert the operator or would reset the BGP session to prevent routing table blow-up.</p> <p>BGP route flap damping is not recommended. See the RIPE-378 document (May 10, 2006)</p>
MSDP	<p>The MSDP (Multicast Source Discovery Protocol) MD5 password authentication feature should be used to protect the TCP connection between two MSDP peers.</p> <p>It is recommended to use Generalised TTL Security Check mechanism (RFC 3682) for MSDP, which introduces a lightweight security mechanism to protect external MSDP (eMSDP) peering sessions from CPU utilization-based attacks using forged IP packets.</p> <p>MPDP packets should be filtered by source/destination addresses.</p>

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

	<p>Additionally, a maximum limit of source announcements should be set per peer and per router.</p> <p>A prefix-list could filter group announcements, to avoid groups that should never be announced from external peers (for example the GÉANT2 AS multicast space as well as multicast bogons).</p>
IPv4 PIM	<p>Only packets from the physical interfaces of peers should be accepted (but there is currently no authentication check available for PIM).</p> <p>No external parties should use the internal rendezvous points for registering sources or joining multicast groups, and BSR packets should be filtered at the edge.</p> <p>The only PIM packets that should be accepted on external interfaces are the PIM join messages.</p>
IPv6 PIM	<p>Only packets from the physical interfaces of peers should be accepted (but there is currently no authentication check available for PIM).</p> <p>External parties might use the internal rendezvous points for registering sources or joining multicast groups – especially if embedded RP is used, but BSR packets can be filtered at the edge.</p> <p>The only PIM packets that should be accepted on external interfaces are PIM join messages.</p> <p>The PIM join messages can be filtered based on the scope accepted by the IPv6 PIM domain.</p>
IS-IS	<p>ISO CLNS packets should not be accepted on any external interfaces of the ISP backbone network, so that IS-IS adjacencies cannot be established with external networks. IS-IS is ISO based and not IP based, so the external attack surface is quite drastically reduced compared to IP protocols.</p> <p>IS-IS can be further protected by using MD5 HMAC authentication for areas/domains. The MD5 HMAC digest allows authentication at the IS-IS routing protocol level, which prevents unauthorized routing message from being injected into the network routing domain.</p>
OSPFv2	<p>OSPFv2 packets should not be accepted on any external interfaces of the ISP backbone network, so that OSPFv2 adjacencies cannot be established with external networks</p> <p>OSPFv2 can be further protected by using MD5 HMAC authentication for areas/domains. The MD5 HMAC digest allows authentication at the OSPFv2 routing protocol level, which prevents unauthorized routing message from being injected into the network routing domain.</p>
OSPFv3	<p>OSPFv3 packets should not be accepted on any external interfaces of the ISP backbone network, so that OSPFv3 adjacencies cannot be established with external networks</p> <p>OSPFv3 can be further protected by using IPSEC HMAC authentication for areas/domains.</p>
RIP and RIPng	<p>RIP and RIPng (IPv6 RIP) should not be used since they are inherently insecure and unscalable.</p>
SSH	<p>SSH should be filtered based on source address to allow management access only for specific users known to access the service from those specific source addresses.</p>
SNMP	<p>It is recommended to use only read-only communities – and allocate them on a per usage/project basis.</p> <p>SNMP should be filtered based on source address.</p> <p>It is wise to configure SNMP traps for all the authentication failures –even for SNMP</p>

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

	authentication failure. It is best to use SNMPv3 features – role based access control and encrypted message exchange.
IGMP	IGMP should only be enabled on internal interfaces connected to workstations and not on external interfaces.
MLD	Like for IGMP, MLD should only be enabled on internal interfaces connected to workstations and not on external interfaces.
NTP	NTP packets should only be allowed from known NTP servers that the routers are set up to communicate with. If the server is external then NTP packets should be authenticated.
HTTP	HTTP server on the routers should not be allowed

**Table 10.2:** Recommendations for securing layer 3 protocols

It is recommended that control plane traffic is rate limited to reduce the risk that the routing engines can be attacked by a DDoS attack.

### 10.3.3.1 Handling ICMP messages

The ICMP destination unreachable messages are key messages used to determine the state of the network. ICMP unreachable messages are responses sent by a router/host/switch whenever the destination host address is unreachable, the specific protocol is unreachable, or the destination networks are not listed in the forwarding table. ICMP unreachable messages are a normal function of the TCP/IP protocol, but can be exploited to overload network devices. Therefore it is recommended to rate limit the ICMP unreachable messages or to prevent them from being generated.

It is also wise to rate limit the informational type ICMP messages (e.g. icmp echo, icmp reply) a router will send out, although this will affect ping responsiveness.

It is a very common (although questionable) practice to filter completely the ICMP messages in IPv4. This is no longer possible with IPv6. As the name that it stands for suggests, Internet Control Message Protocol for IPv6 (RFC 2463) is the control and foundation protocol for the operation of IPv6, not an auxiliary protocol that can be easily omitted. Our recommendation is the following:

- Enable links scoped ICMPv6-Neighbour-Solicitation and Neighbour-Advertisement (Type 135 and 136) for the Neighbour Discovery function to operate properly and ICMPv6-Router-Solicitation and Router-Advertisement (Type 133 and 134) if the Stateless Address Auto configuration function is used.
- You must enable incoming ICMPv6-packet-too-big messages (Type 2) as answers to outgoing IPv6 packets for the Path-MTU-discovery to operate properly.
- You must generate ICMPv6-packet-too-big messages properly if your MTU is different

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

anywhere within your network from the MTU on the link between you and your provider. Be prepared to forward ICMPv6-packet-too-big messages on the firewalls and routers.

To summarise the ICMPv6 recommendations in the following table [19]:

ICMPv6	Usage
Echo request/reply	Debugging
Destination unreachable	Debugging – better indicators
TTL Exceeded	Error report
Parameter problem	Error report
NS/NA	Important for IPv6 operation - except if you use Static ND entry
RS/RA	For Stateless Address Auto configuration
Packet too big	Important for PATH MTU discovery
MLD messages	Required for Multicast operations

**Table 10.3:** ICMPv6 recommendations

Note: Each IPv6 specific feature is marked with Blue, and each required feature marked with Red.

### 10.3.3.2 Source address

It is wise to configure an IP source interface for each subsystem in the router. This address is important, since it can be used without ambiguity in all firewall rules. It is recommended to use a loopback address as a source address since a loopback interface is always working independently of the physical interfaces.

## 10.3.4 Protecting the forwarding plane of the router

Some routers require explicit enabling for faster packet forwarding. In the backbone network it is essential so it has to be switched on – this is also a requirement for generating netflow records.

The most important tool to protect an ISP's resources and its customers' network is ingress and egress filtering as per BCP 38/ RFC2827. These rules allow enforcing policy, as well as reducing the risk of being the network chosen by the hackers to launch an attack on other networks.

There are several ways to implement BCP 38: using firewall filters, router ACLs or by using uRPF (unicast Reverse Path Forwarding) checking.

uRPF is presented as a proactive countermeasure because it can be used to block attack traffic that uses address space that is either not allocated for general Internet use or traffic belonging to a different network than

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

the network the packets are received from. uRPF should be used with care so as to avoid that any legitimate traffic is dropped.

It is recommended that the tools and services from JRA2/WI-2, when available, be used for detecting security incidents before they affect any service.

### 10.3.4.1 Unicast Reverse Path Forwarding (uRPF)

uRPF is a popular technique among major routing vendors for mitigation of forwarding plane attacks, where the attacker uses spoofed source IP addresses. The scheme works as follows. Assuming the existence of a Forwarding Information Base - FIB, a router checks for every incoming IP packet: the source address field, and the interface from where this packet was received. If this association has an absolute match in the FIB, the packet is forwarded, otherwise it is rejected. This is shown in Figure 5.6




Figure 10.2: uRPF strict operation

In contrast to the above mentioned technique which is also known as strict uRPF, there is a slight modification called loose uRPF. According to this flavour of uRPF there is no need for a strict match in the FIB. The necessary criterion is the existence of a route in the FIB. The loose operation is shown in Figure 5.3.

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3




Figure 10.3: Loose uRPF

It is recommended that NRENs run strict uRPF on interfaces facing sites to ensure that sites only send traffic that has a source address corresponding to their allocated address space. uRPF can also be implemented between NRENs and GÉANT2 in a loose manner similar to the uRPF check implemented in GÉANT.

## 10.4 Layer 2 services

Martini layer 2 circuits have been deployed since the GÉANT project. The specific implementation depends on the requirements and the equipment available. A typical example is shown in the Figure below:

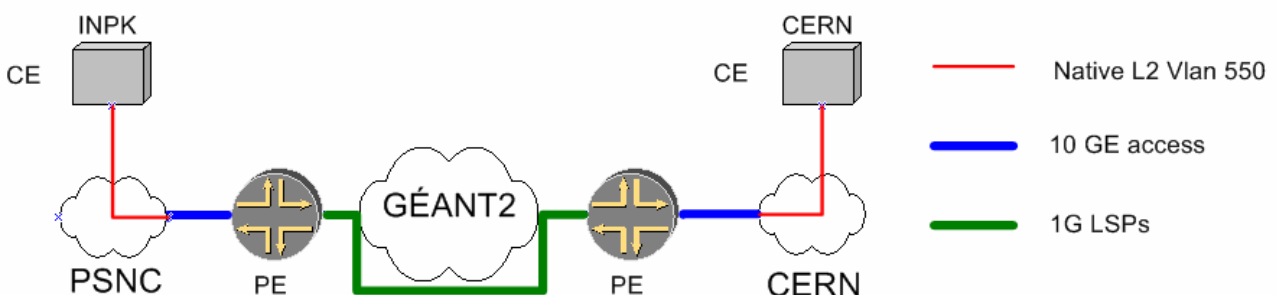


Figure 10.4: Martini layer 2 circuit – Atlas project

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

The layer 2 circuit connects two sites in two different NRENs. Each end site will see the remote site as directly connected at layer 2 via a VLAN. In order to make this work, the following protocols are needed:

- LDP signalling is used between the loopback interfaces of the two PE routers.
- RSVP signalling is used to set up the LSPs. RSVP is used to enable traffic engineering.
- MPLS is required to forward traffic. Here, MPLS is simply the label switching functionality which enabled the router to forward packets with an MPLS header. Label stacking is used so the inner label is the VC (virtual circuit) label and the outer label is for the LSP.

The table below focuses on the control protocols running on the GÉANT2 core routers and NREN and campus routers.

Martini Layer 2 circuit are provisioned using the protocols in the table below:

Protocol	Recommendation for securing
RSVP	RSVP should be enabled on external interfaces only if it is required for a specific project, i.e. for signalling LSPs for an end-to-end layer 2 VPN. If so, RSVP messages should be accepted from specific peers only (using the physical interface addresses for source and destination addresses) but since the contents of RSVP cannot be filtered, RSVP should be used with care.
LDP	Like RSVP, LDP should only be enabled on external interfaces if it is required for a specific project, i.e. for the control plane signalling for an end-to-end layer 2 or layer 3 VPN (in case of layer 3 VPN Carrier supporting Carrier setup). In that case, LDP packets should be filtered by source and destination addresses. It is recommended to use Generalised TTL Security Check mechanism (RFC 3682) for LDP which introduces a lightweight security mechanism to protect external LDP sessions from CPU utilization-based attacks using forged IP packets.
MPLS	MPLS is not in itself a protocol but is used for providing MPLS services via label switching. MPLS packets should normally only be accepted on internal interfaces but can be allowed on external interfaces when particular projects need end-to-end layer 2 VPNs or layer 3 VPN (in case of layer 3 VPN Carrier supporting Carrier setup).

**Table 10.4:** Recommendations for securing layer 2 protocols

## 10.5 Layer 1 services

Note: The following is a work in progress and focuses on the GÉANT2 core SDH and WDM equipment (Alcatel) though the technologies of the other vendors are similar.

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3

### 10.5.1 Data Communications Network

SDH and WDM equipment is managed by a centralised NMS (which can be a redundant set of servers). An NMS is typically proprietary. In the GÉANT2 core, the Alcatel NMS is the 1353NM (element manager) and 1354RM (service management). In addition, the 1359IOO is used to export alarms from the NMS to a workstation. The use of the IOO prevents the need for any external application to connect directly to the NMS servers.

The NMS needs a Data Communication Network (DCN) to be able to communicate with the network elements (NE). SDH and WDM equipment is managed differently which is analysed further below.

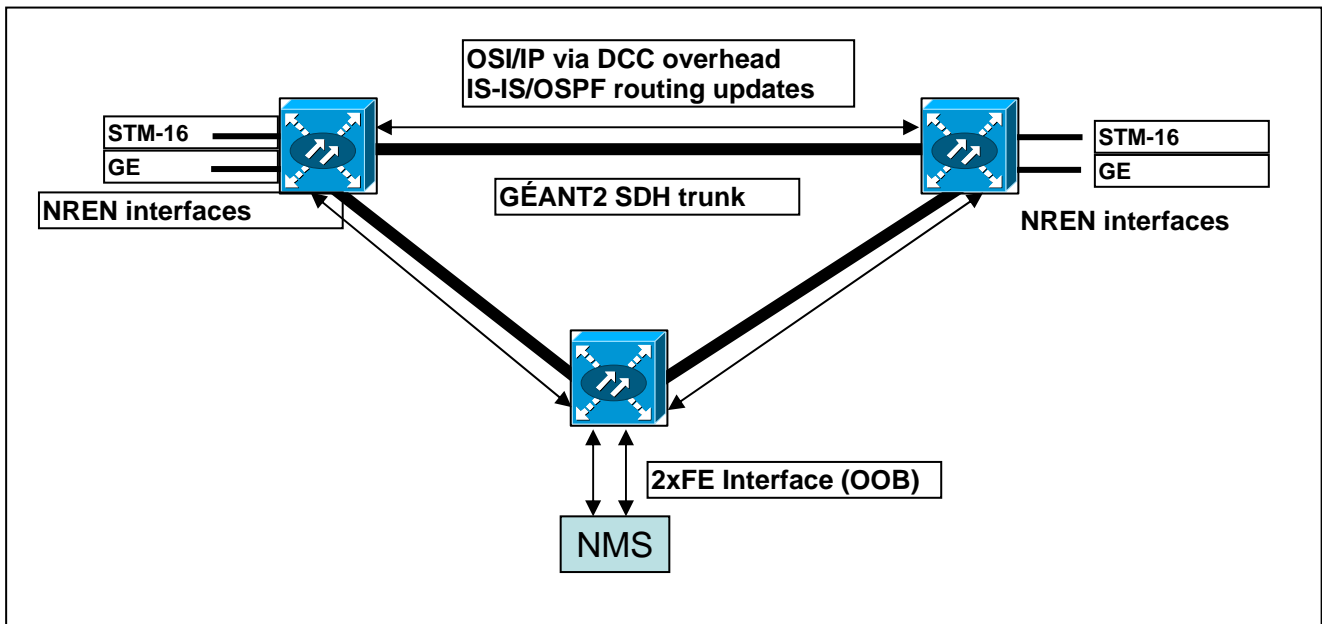
The NMS to NE communication is difficult to disrupt because the NEs only accept TL1 commands via CORBA interface. The NEs are configured only to accept packets from the specified NMS servers (the source address is checked).

### 10.5.2 SDH equipment

SDH equipment can be managed in-band via the DCC channel or out-of-band via an Ethernet interface on the switch controller (which is typically redundant).

For SDH, the DCC channel are carried via the regeneration section overhead bytes (D1,D2,D3 bytes). This provides a low speed management channel (512 kbps for the GÉANT2 SDH equipment) between two neighbouring SDH switches.

The GÉANT2 SDH switches (Alcatel 1678MCC) can be managed using either OSI only or a combination of IP and OSI (CLNS). Both the OSI and IP traffic is using LADP as the layer 2 encapsulating protocols. Dynamic routing protocols are used for both IP and OSI to make all network elements aware of the network topology so DCN traffic cannot be routed in case of outages on the SDH trunks. For IP, OSPF is used and for OSI IS-IS is used.



**Figure 10.5:** SHD switch management

The 1678MCC can be managed entirely in-band (where only switch in the same PoP as the NMS is used as a gateway). Out-of-band management can be enabled to ensure that the NMS servers can still manage the SDH switch even if all in-band channels are unavailable.

On SDH circuits, user data (payload) has no access to the DCC channel which means that the in-band channels cannot be affected by a DDoS attack. Since no external access is required, it is recommended to use private IP address for the management interfaces – for GÉANT2, all packets with private IP addresses will be discarded on ingress. OSI packets are not normally routable across a domain border so OSI packets will be discarded on ingress – for GÉANT2, all OSI packets from external peers are discarded.

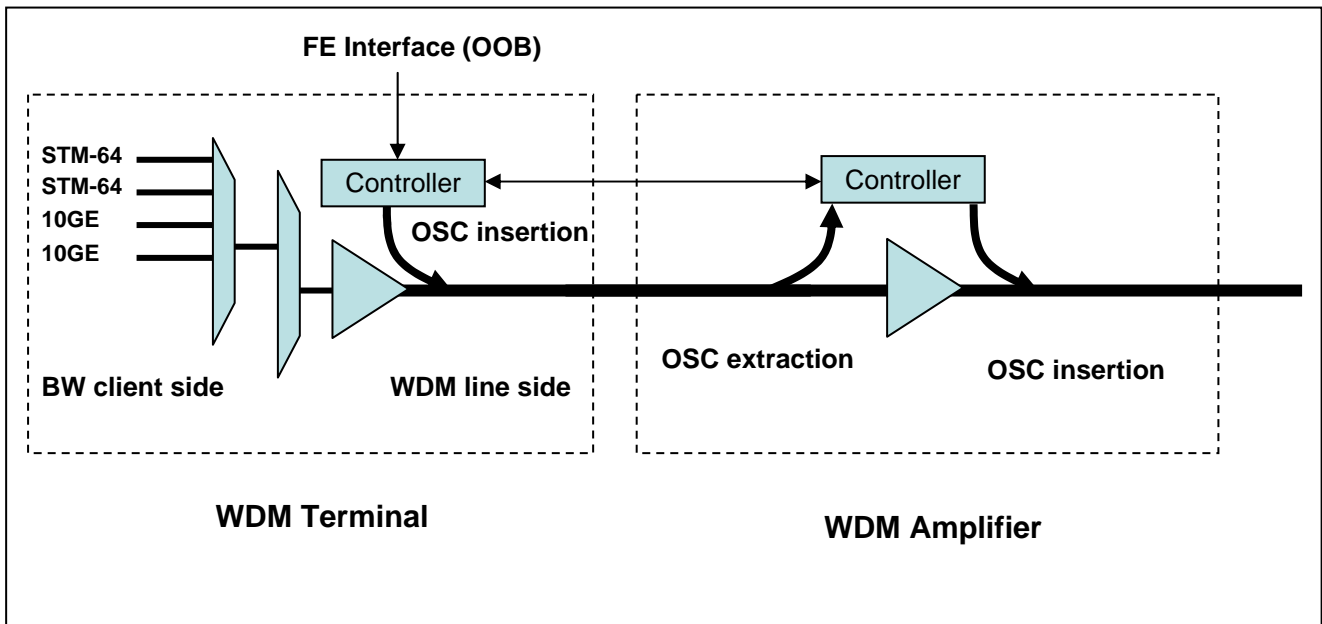
If a NREN use an SDH switch to connect to a GÉANT2 SDH switch, DCC signalling should be disabled on that interfaces. The GÉANT2 SDH switch will ignore the DCC overhead bytes from a client interface so it will not be affected, but it would cause unnecessary confusion.

### 10.5.3 WDM equipment

WDM equipment is managed using a combination of in-band management (optical supervisory channel – OSC) and out-of-band management (Ethernet interface). The WDM equipment deployed in the GÉANT2 core is the Alcatel 1626LM platform.

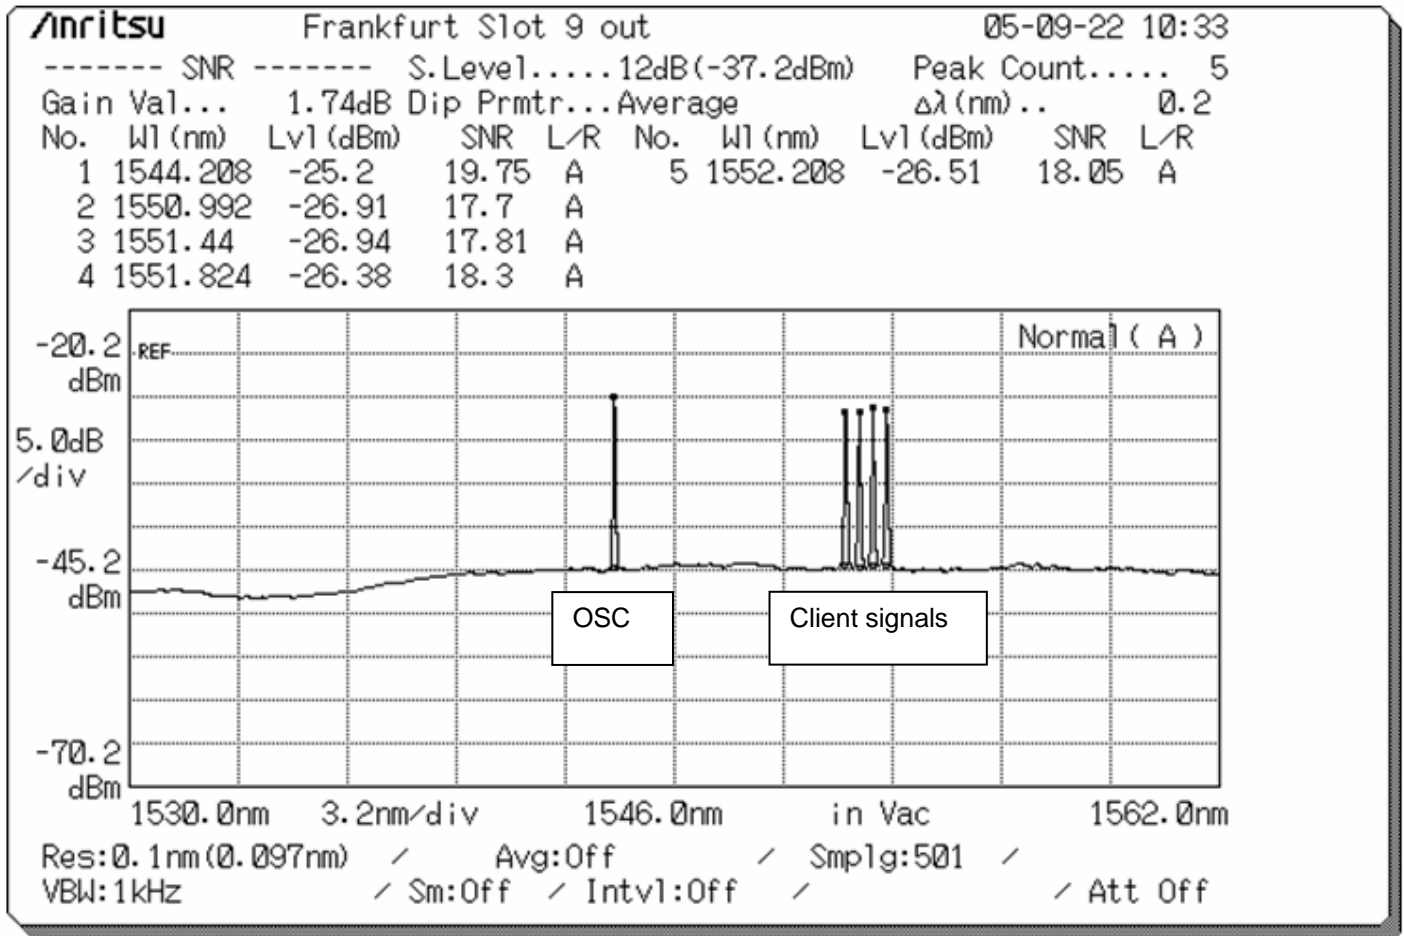
For in-band management, the OSC is inserted as a dedicated wavelength as shown on the diagram below:

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3



**Figure 10.6:** Functional diagram of OSC

The OSC is generated on a WDM terminal and electrically regenerated on each WDM in line amplifier (ILA) until it is terminated on the remote WDM terminal at the end of the route. For GÉANT2, the OSC provides 4 mbps channel. The OSC is inserted and extracted using optical filters. In this way, the controller board of each WDM network element is reachable in-fibre. The WDM terminal works as a gateway to an ILA. This design means that a client interface (or BW – black and white interface) has no access to the OSC. Between two terminals, the client signals are processed only as an optical signal so the user data has no effect on the WDM equipment. The diagram below illustrates how the OSC uses a wavelength different from the client signals.



**Figure 10.7:** Spectrum Analysis showing OSC

Out of band management is available via an Ethernet interface on the terminals which gives access to the controller.

The WDM equipment deployed in the GÉANT2 core is managed entirely with OSI (CNLS). The routing topology is maintained with IS-IS so any topology changes (i.e. a fibre cut) causes traffic to be rerouted. The OSI network is one flat network where each WDM network element is an OSI router will full topology knowledge. The use of OSI only means that no packets from outside the GÉANT2 core will be sent to the WDM equipment because OSI traffic is discarded on ingress to GÉANT2.

Project:	GN2
Deliverable Number:	DJ2.1.1,3
Date of Issue:	12/09/06
EC Contract No.:	511082
Document Code:	GN2-06-205v3