

06.08.08

Deliverable DJ.2.1.1,5: Revised GÉANT2 Security Recommendation and Policy



Deliverable D.J.2.1.1,5

Contractual Date: 30/06/08
Actual Date: 06/08/08
Contract Number: 511082
Instrument type: Integrated Infrastructure Initiative (I3)
Activity: JRA2
Work Item: 1
Nature of Deliverable: R (Report)
Dissemination Level: PU (Public)
Lead Partner: DANTE
Document Code: GN2-08-170

Authors: M. Molina (DANTE), M. Mogensen (DANTE), D. Kalogeras (GRNET), J. Mohacsi (HUNGARNET), H. Nussbacher (IUCC), M. Garcia (DANTE), R. Sabatino (DANTE), M. Wright (DANTE), G. Kramer (DANTE), W. Routly (DANTE)

Abstract

This document is an update to the 4th Security Recommendation and Policy deliverable, "DJ2.1.1,4: Revised GÉANT2 Security Recommendation and Policy [DJ2.1.1,4]". The purpose of this document is to provide insight into how a network such as GÉANT2 is secured and to provide a checklist of best practices. It is hoped that this will assist NREN operators, their connected end sites and any other parties that may find it useful in securing their own networks.

Table of Contents

0	Executive Summary	iv
1	Introduction	1
	1.1 Security Policies	2
2	Securing the network infrastructure	4
	2.1 Layer 3 Services	4
	2.1.1 General	4
	2.1.2 Best Practices	5
	2.2 Layer 2 Services	19
	2.2.1 General	19
	2.2.2 Best practices	20
	2.3 Layer 1 Services	22
	2.3.1 General	22
	2.3.2 Best Practices	23
	2.4 Physical Security	27
3	Securing End-to-End Network Services	28
4	Securing other Service Delivery Equipment	30
	4.1 Workstations	30
5	Protecting End Customers	32
	5.1 Netflow-Based Anomaly Detection	32
	5.2 Darknets	33
	5.2.1 Creating a Darknet	34
	5.2.2 Using Darknet Data	35
	5.3 Black Holes	36
6	Conclusion	37
7	References	38
8	Acronyms	40

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

Table of Figures

Figure 2.1: Router view of forwarding, control and management plane.	7
Figure 2.2: uRPF strict operation.	18
Figure 2.3: Martini Layer 2 circuit – Atlas project.	20
Figure 2.4: SDH switch management.	24
Figure 2.5: Functional diagram of the OSC.	25
Figure 2.6: Spectrum Analysis showing the OSC at 1544 nm.	26
Figure 5.1: A Darknet implementation.	35

0 Executive Summary

This document is an update to the deliverable “DJ2.1.1,4: Revised GÉANT2 Security Recommendation and Policy [DJ2.1.1,4]”, which was submitted in September 2007.

The purpose of this document is to provide insight into how a network such as GÉANT2 is secured and to provide a checklist of best practices. It is hoped that this will assist NREN operators, their connected end sites and any other parties that may find it useful in securing their own networks.

The previous deliverable focused on possible security policies that would be implemented for the expected end-to-end services available through GÉANT2. In comparison, this update focuses on updating the security policies for new services that are now operational in GÉANT2.

The recommendations presented in this document are based on experience gained during the deployment and continued operation of end-to-end services, and cover the following areas:

- Network infrastructure.
- End-to-end services.
- Service delivery equipment.
- End customer protection.

1 Introduction

This document is an update to the deliverable “DJ2.1.1,4: Revised GÉANT2 Security Recommendation and Policy [DJ2.1.1,4]”, which was submitted in September 2007.

The purpose of this document is to provide insight into how a network such as GÉANT2 is secured and to provide a checklist of best practices. It is hoped that this will assist NREN operators, their connected end sites and any other parties that may find it useful in securing their own networks.

Based on experience gained over the last 10 years in securing the Internet Protocol (IP) network and the IP based services, no changes to securing network services since previous revisions are foreseen.

Considering the new services introduced by GÉANT2 (10Gbps multi-domain point-to-point or end-to-end services were made available from July 2006 onwards, and Gigabit Ethernet services using the Alcatel Metro Core Connect (MCC) switching platform were delivered from October 2006 onwards), there are no known current security issues, and therefore no recommendations are made beyond what is described as best practice for Layer 1 services.

The structure of this document is designed to cover all aspects of the GÉANT2 project, from the basic Layer 1 infrastructure, to the routed IP network and to service delivery equipment such as workstation access. Although the document is centred around the core network, possible methods to be used within the core to assist in the protection of end customers equipment and networks are also documented.

As the security policies described are based on a technical analysis of the security requirements of each service, these policies are followed by the relevant technical analysis. The security policies are divided into:

- Network infrastructure.
- End-to-end services.
- Service delivery equipment.
- End customer protection.

The policies focus on the demarcation points of the services to define responsibility at the administrative domain border. Procedures for dealing with incidents are defined. The existing policies for Layer 3, Layer 2 and Layer 1 services are documented under the network infrastructure heading.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

The best practice documentation and the technical analysis of the security requirements for GÉANT2 services result from the analysis of the management, control and forwarding plane.

1.1 Security Policies

The aim for security policies is to be vendor-independent as far as possible.

The security policies for new services are based on a technical analysis of security requirements. The technical analysis and best practice documentation is maintained in the Joint Research Activity (JRA) 2 Document Management section of the GÉANT2 Wiki [WIKI].

In the following, the services offered by GÉANT2 are divided using the Open Systems Interconnection (OSI) model [OSI], 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	IPv4, IPv6, Multicast	Routers, Workstations
Layer 2	Layer 2 circuit	Routers, Switches
Layer 1	Lambda, Synchronous Digital Hierarchy (SDH) circuit, Point to point Ethernet	SDH, Wavelength-Division Multiplexing (WDM)

Table 1.1: Services per OSI model layer

In the following sections, a security recommendation and policy for services is 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 are summarised in *Layer 3 Services* on page 4.
- Layer 2 services were partly covered in the initial deliverable, but are covered separately in *Layer 2 Services* on page 19, as the security requirements are significantly different than for Layer 3 services.
- Layer 1 services were previously presented as an initial security recommendation and policy. During the operational phase these initial security recommendations proved useful and as such are included here as best practice in *Layer 1 Services* on page 22.
- Physical security is introduced as an additional consideration as this is often an overlooked area when securing networks and associated equipment.
- An additional category has been added to cover service delivery equipment not currently covered by the layers 1-3 policies. Currently this documents only workstation security measures.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

- Possible methods to assist in the protection of end customers is included. While this is not essential to securing the network core it is an important aspect of providing a secure environment.

An end-to-end service within GÉANT2 is expected to use at least one of the three OSI layers and to implement the other aspects as appropriate. The best practice is described in the relevant section for each layer. However, it is still necessary to define a security recommendation for end-to-end services to have a common framework for discussion, especially for service handover at the administrative domain border.

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

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

2 Securing the network infrastructure

2.1 Layer 3 Services

2.1.1 General

Layer 3 services have a clearly defined domain boundary which is typically structured as follows:

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

This structure has the benefit of making the responsibilities and demarcations clear, so that each organisation is responsible for their own domain. This is possible because Layer 3 services have a clear boundary between administrative domains. However, Layer 3 services have a significant overlap between the management plane, control plane and forwarding plane, which makes the service vulnerable. Fortunately, there are several good practices that greatly reduce the risk. The following discussion is focused on routers; workstation security is covered in *Workstations* on page 30.

The best practice section presents 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 Message Digest Algorithm 5 (MD5) authentication for Border Gateway Protocol (BGP) peerings.

To protect the forwarding plane, the following pro-active security measures have been implemented on the border between GÉANT2 and National Research and Education Networks (NRENs):

- uRPF (Unicast Reverse Path Forwarding) 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 Réseaux IP Européens (RIPE) Autonomous System (AS) macro of the NREN. The reason for doing two checks is that legitimate traffic could fail the first check. For further details see *Unicast Reverse Path Forwarding (uRPF)* on page 17.

- All packets entering GÉANT2 with Bogon source addresses are cross-referenced against an up-to-date list of Bogon addresses [BOGON], and discarded if they match. It is important to keep the Bogon list up-to-date to prevent legitimate traffic from newly assigned IP ranges from being discarded. The GÉANT2 list is kept up-to-date by subscribing to the Bogon announce list [LIST]. It is believed that as many as 70% of brute force Denial of Service (DoS) attacks use Bogon source addresses [CYM]. These filters protect all NRENs from such attacks from other NRENs and the wider Internet.
- On all GÉANT2 NREN access ports, filters that control bandwidth for certain protocols/ports have been implemented to control the traffic that would be destined for GÉANT2 equipment (any equipment allocated with a prefix from within the GÉANT2 /19 core address space). The purpose of this is to protect the GÉANT2 infrastructure and related services from DoS attacks, and to prevent the scanning activity that is usually a precursor to hacking activity. Production traffic from NRENs is not affected in any way.

Re-active security measures have also been implemented:

- The GÉANT2 Network Operations Centre (NOC) can filter NREN traffic on request. If needed, the GÉANT2 NOC can use router logs to 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.

2.1.2 Best Practices

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 is 2 million packets per second. This is not the case for the types of routers known as software routers. This type of router is present in a lot of campus networks and some NREN networks, and handles connectivity demands ranging from multiple Mbps to Synchronous Transport Module (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 tends to be uncommon due to widespread implementation of ingress filtering (such as Best Common Practice 38 [BCP38]).

Another method to achieve a forwarding plane attack is improper handling and processing of ICMP and other control packets, which can lead to a commonly exploitable phenomenon known as "ping of death". It is therefore important to consider how to detect and respond to such forwarding plane attacks in the networks' core. In some cases, forwarding plane attacks can be associated with a short time increase of flows destined to the same subnet or host. However, forwarding plane attacks look like normal network traffic in the core of the network, which means they are quite difficult to detect.

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

In the following section, each type of equipment is analysed and recommendations for a best practice for securing them are presented. For each equipment type, how to secure the management, the control and the forwarding plane is discussed.

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 that 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 or compromised.
- 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 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 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. This typically occurs 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. However, DDoS attacks should be detected and stopped in the interest of the affected end users and their equipment.

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.

2.1.2.1 Routers

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

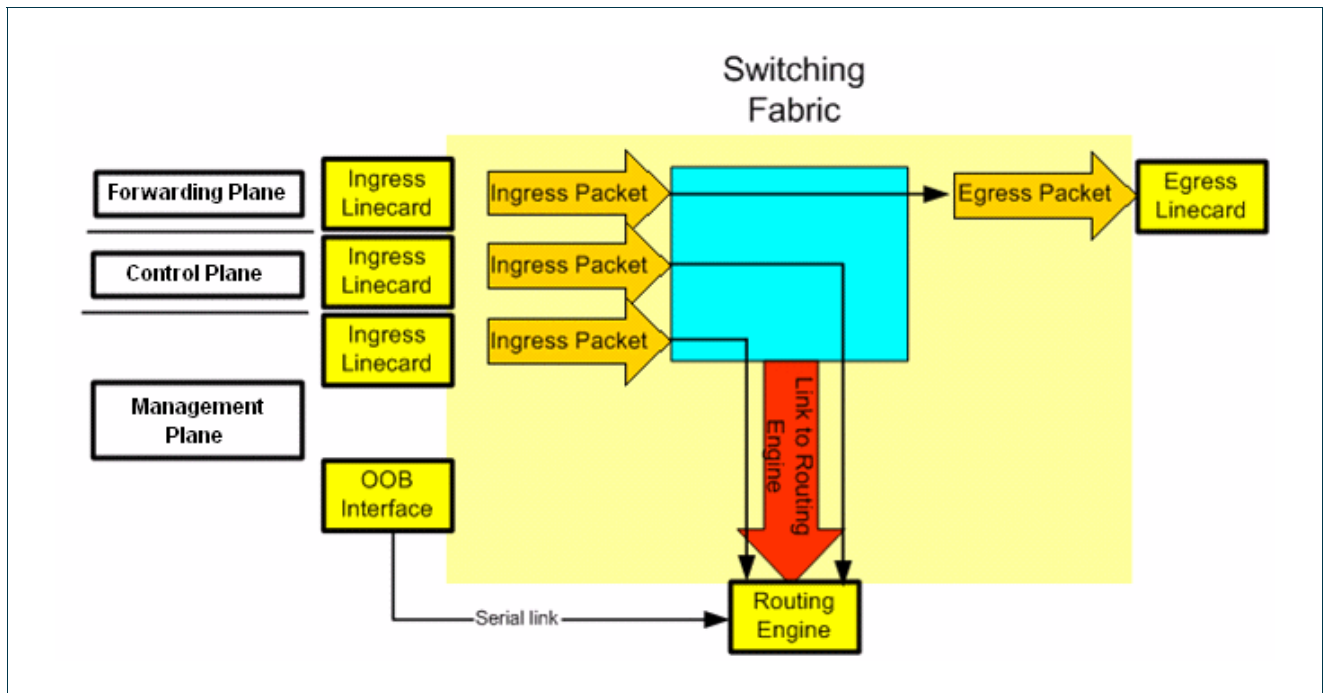


Figure 2.1: Router view of forwarding, control and management plane.

- The management plane gives user access to the routing engines (one active, one standby). The engines are 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 via an internal link between the switching fabric and the routing engine. For an analysis of the protocols used, see *Table 1.1: Services per OSI model layer* on page 2.

Protocol traffic to the routing engines can be rate-limited to reduce the impact of a DDoS attack on the control plane, and this has to be implemented with great care. However, the control plane can still be attacked by a DDoS assault 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, simply 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 into the control plane, which eventually causes 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. As in the case of 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 either the filtering features or the firewall or access control features of the core IP routers. This is because data can be filtered and logged at line rate with great granularity. Another option would be to use a separate/non-routable address space to address management interfaces.

In this way, the management plane is used 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 only see packets addressed to the routing engine itself, while 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, both the control plane and management plane can be disrupted. Traffic to the routing engine can be rate-limited, but this should be implemented with care, and protocol authentication (with MD5 for example) should be used as much 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 Local Area Network (LAN) equipment, workstations first of all.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

2.1.2.2 Protecting the Management Plane of Routers

There are many ways of improving the security of the management plane of routers.

In general only authorised users should be able to access the router, and only a 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 Virtual Local Area Networks (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 of each router that belongs to the same administrative domain. ACL lists are very good for limiting the exposure of the router. Secure Shell (SSH) is used on top of the ACLs to authenticate users.

Routers run a large number of services where some services use more than one protocol. For an analysis of each protocol for its attack surface and recommendations for securing them, see *Table 1.1: Services per OSI model layer* on page 2.

Many built-in services in some routers software are not needed in an ISP backbone environment. These features (for example, echo server, HTTP server) should be turned off in the default configuration.

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

- IP redirection on interfaces should be switched off.
- IP directed broadcast should be switched off on all interfaces, otherwise “Smurf”-type attacks [SMURF] can be conducted against the management plane.
- Proxy Attribute Release Policy (ARP) is usually not needed in a backbone environment. Proxy ARP, as defined in Request for Comments (RFC) 1072 [RFC1072], is used by the router to help hosts with no routing capability to determine the Media Access Control (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 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 not 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, if they connect to the router, that only authorised users are permitted to connect. They should be given (for example) an official warning and told to contact the helpdesk. Information about the system and the provided services should not be provided.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

Passwords and Simple Network Management Protocol (SNMP) community strings kept in any stored configuration should be encrypted, preferably in a non-reversible form for non-administrative users. This can prevent attacks against the routers' management plane.

A timeout should be implemented on all management interfaces, as each open management session uses precious resources of the routing engine. This will also minimise the risk of an operator leaving their terminal logged into the router.

Secure authentication to access the management plane should be implemented. SSH must be used 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, Authorisation and Accounting (AAA) 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 passwords in local configuration files, and all access is logged on 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, that are typed into the router. The AAA infrastructure usually allows router command auditing.

In addition, it is recommended to enable the logging of system messages to an external syslog server for later analysis. The syslog server should be synchronised from the same time source as the routers to allow comparison of logs with peers and/or NRENs. Ideally the time source should be in Universal Time Coordinated (UTC). The collected logs should be analysed in order to detect unwanted or malicious activities.

A proper boot strategy and configuration management should be maintained to ensure that the networking devices always boot an untainted operating system image and use the correct configuration files. Before uploading a new image in the networking device, it must be checked to ensure it is coming from an authentic source and has not been tampered with since the original shipment or download.

2.1.2.3 *Protecting the Control Plane of Routers*

Control plane traffic should be filtered based on source address as much as possible, so that only protocol data from expected sources is accepted. Filtering in-bound traffic on all ingress network points is recommended, 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 (RFC 1918 [RFC1918] private IP address space, for example).

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

In IPv4 it is easier to filter out (deny) a packet originated from Bogon routes, while in IPv6 it is easier to allow legitimate packets. At the time of writing, the allocation of IPv6 addresses by Internet Assigned Numbers Authority (IANA) is limited so this method of filtering is easily implemented. As the IPv6 address space grows this method will not scale, however, it is expected that, by that point, there will be a readily available Bogon list as there is for IPv4.

Current IPv6 unicast allocations can be viewed online [IANA1] [IANA2]. Special assignments are mentioned in RFC 5156 [RFC5156]. It is also worth considering IPv6 BGP filtering recommendations [BGP].

In summary, all special purpose prefixes that should not be seen on the Internet should be blocked, all allocated addresses should then be accepted and everything else denied. It is important that the list of assigned unicast addresses is checked regularly, as otherwise legitimate newly assigned addresses may be blocked inadvertently.

Priority	Rule
1	Deny RFC 5156 [RFC5156] addresses.
2	Allow assigned unicast addresses.
3	Deny everything else.

Table 2.1: Bogon Filtering Firewall Rules in IPv6.

The table below focuses on the control protocols running on the GÉANT2 core routers. The term “external interface” refers to an interface facing another network. The term “internal interface” refers to an interface facing another router within 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 Time To Live (TTL) Security Check Mechanism (RFC 3682 [RFC3682]) for BGP, which introduces a lightweight security mechanism to protect external Border Gateway Protocol (eBGP) peering sessions from CPU utilisation-based attacks using forged IP packets.</p> <p>It is also useful to filter BGP packets by source/destination addresses, so that 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 or another Regional Internet Registry (RIR) database in order to only accept the routes that this peer should be advertising.</p> <p>A maximum number of prefixes should be configured per peer, whereby reaching a</p>

	<p>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 Routing Working Group's Recommendations on Route-flap Damping [RIPE].</p> <p>The private AS numbers should be removed from the ASPATH.</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 [RFC3682]) for MSDP, which introduces a lightweight security mechanism to protect external MSDP (eMSDP) peering sessions from CPU utilisation-based attacks using forged IP packets.</p> <p>MPDP packets should be filtered by source/destination addresses.</p> <p>Rate-limiting source announcements is highly recommended while simply rate-limiting all MSDP traffic is not useful, as it effectively lowers the amount of traffic an attacker has to send before packet loss occurs. The following methods of applying rate limits offer a better solution, and used together they provide good protection against unwanted source announcements:</p> <p>For per-source rate limiting a good setting is to allow a maximum of 1000 active streams per source. If a source has more than 1000 active streams (which is unlikely), this can be explicitly specified.</p> <p>Per-peer limits work in a similar way to the BGP max prefix limits and limit the number of Source Announcements (SAs) received from an MSDP peer.</p> <p>The final rate limiting option is per-instance. This method is better suited to protecting the router from being overloaded rather than policing the network. It simply sets a limit on the number of SAs in a routers routing instance.</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>

<p>IPv4 PIM</p>	<p>Only packets from the physical interfaces of peers should be accepted (but there is currently no authentication check available for Protocol-Independent Multicast, PIM).</p> <p>No external parties should use the internal rendezvous points for registering sources or joining multicast groups, and Bootstrap Router (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> <p>A more wide-ranging implementation is to block all PIM joins for groups that do not belong on the Internet. This can be done so that the PIM joins still transit the network but any join to the specified groups is denied.</p>
<p>IPv6 PIM</p>	<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 an embedded Rendezvous Point (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>
<p>IS-IS</p>	<p>International Standards Organization (ISO) Connectionless Network Service (CLNS) packets should not be accepted on any external interfaces of the ISP backbone network, so that Intermediate System-to-Intermediate System (IS-IS) adjacencies cannot be established with external networks. IS-IS is based on ISO, not IP, so the external attack surface is drastically reduced compared to IP protocols.</p> <p>IS-IS can be further protected by using MD5 Hash Message Authentication Code (HMAC) authentication for areas/domains. The MD5 HMAC digest allows authentication at the IS-IS routing protocol level, which prevents unauthorised routing messages from being injected into the network routing domain.</p>

OSPFv2	<p>Open Shortest Path First (OSPF) v2 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 should 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 unauthorised routing messages 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 should be further protected by using Internet Protocol Security (IPSEC) HMAC authentication for areas/domains.</p>
RIP and RIPng	<p>Routing Information Protocol (RIP) and Routing Information Protocol next generation (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>SNMP traps should be configured for all the authentication failures, even for SNMP authentication failure.</p> <p>SNMPv3 features that can be used are role-based access control and encrypted message exchange.</p>
IGMP	<p>Internet Group Management Protocol (IGMP) should only be enabled on internal interfaces connected to workstations and not on external interfaces.</p>

MLD	Multicast Listener Discovery (MLD) should only be enabled on internal interfaces connected to workstations and not on external interfaces.
NTP	<p>Network Time Protocol (NTP) packets should only be allowed from known NTP servers that the routers are set up to communicate with. If the server is external, NTP packets should be authenticated.</p> <p>Network elements should be synchronised to preferably two or more reliable NTP sources to keep the logs and all the operations punctual.</p>
HTTP	Hypertext Transfer Protocol (HTTP) or Hypertext Transfer Protocol Secure (HTTPS) server on the routers should not be allowed.

Table 2.2: Recommendations for securing possible protocols encountered on the control plane.

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

Handling IP Header Options

The headers of IP packets in v4 and v6 have space reserved for several options to be set. Usually there is no reason to discard or ignore such option bits. Notable exceptions are:

- A header option that allows the source of a packet to specify the route to be taken to reach the destination, bypassing the decision of any routers along the path. For IPv4 this feature is known as “source-routing” and almost universally disabled as it is generally thought to be problematic.
- When version 6 of the protocol was devised, the problems of source-routing were forgotten and the source-routing capability was included. This is known as the “Type 0 Routing Header” flaw. The flaw was first identified in 2001 but was not widely known until it was demonstrated in 2007 [FLAW].

Once the exploit was demonstrated, vendors released various fixes. In some cases this involved discarding any packet with the Type 0 Routing Header extension, a solution that had to be specifically added to firewalls on the routers. At the time of writing, most, if not all, IPv6 stacks ignore the Type 0 Routing Header, so discarding the packet is no longer necessary. Further information from Juniper is available online [JUNIPER].

There are also steps to deprecate the Type 0 Routing Header function. More information can be found in the related RFC 5095 [RFC5095].

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

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 recommended to rate-limit the informational type ICMP messages (e.g. icmp echo, icmp reply) that a router will send out, although this affects ping responsiveness.

Whilst it may be a common practice to filter all ICMP packets, it is not recommended as ICMP is used for much more than just echo requests and blocking these control packets can impair the day-to-day working of a network [ICMP].

Filtering of ICMP packets in IPv6 should be considered more carefully [RFC4890]. As its name suggests, Internet Control Message Protocol for IPv6 (RFC 2463 [RFC2463]) is the control and foundation protocol for the operation of IPv6, not an auxiliary protocol that can easily be omitted. Our recommendations are:

- Enable link-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-Maximum Transmission Unit (MTU) discovery to operate properly.
- You must generate ICMPv6-packet-too-big messages properly if your MTU is different 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.

The following table summarises the ICMPv6 recommendations:

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.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

Packet too big	Important for PATH MTU discovery.
MLD messages	Required for Multicast operations.

Table 2.3: ICMPv6 recommendations

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

Source Address

It is recommended 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, as a loopback interface is virtual and always working independent of the state of the physical interfaces.

2.1.2.4 Protecting the Forwarding Plane of Routers

Some routers require explicit enabling for faster packet forwarding. In the backbone network this is essential, so this must be switched on. This is also a requirement for generating netflow records.

The most important tool to protect an Internet Service Provider's (ISP) resources and its customers' network is ingress and egress filtering, as according to BCP 38/ RFC 2827 [RFC2827]. These rules allow enforcing policy, as well as reducing the risk of being the network chosen by hackers to launch an attack on other networks.

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

uRPF is presented as a proactive countermeasure because it can be used to block attack traffic which uses address space that is either not allocated for general Internet use or assigned to a different network than the network the packets are received from. uRPF should be used with care to avoid dropping any legitimate traffic.

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

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), for every incoming IP packet a router checks 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. There are two uRPF modes, strict and loose. Strict will drop any packet which arrives on an interface that is not the best route back to the source. Loose only drops packets if there is no known route to the source. Figure 3.2 shows the strict mode implemented with a fail filter:

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

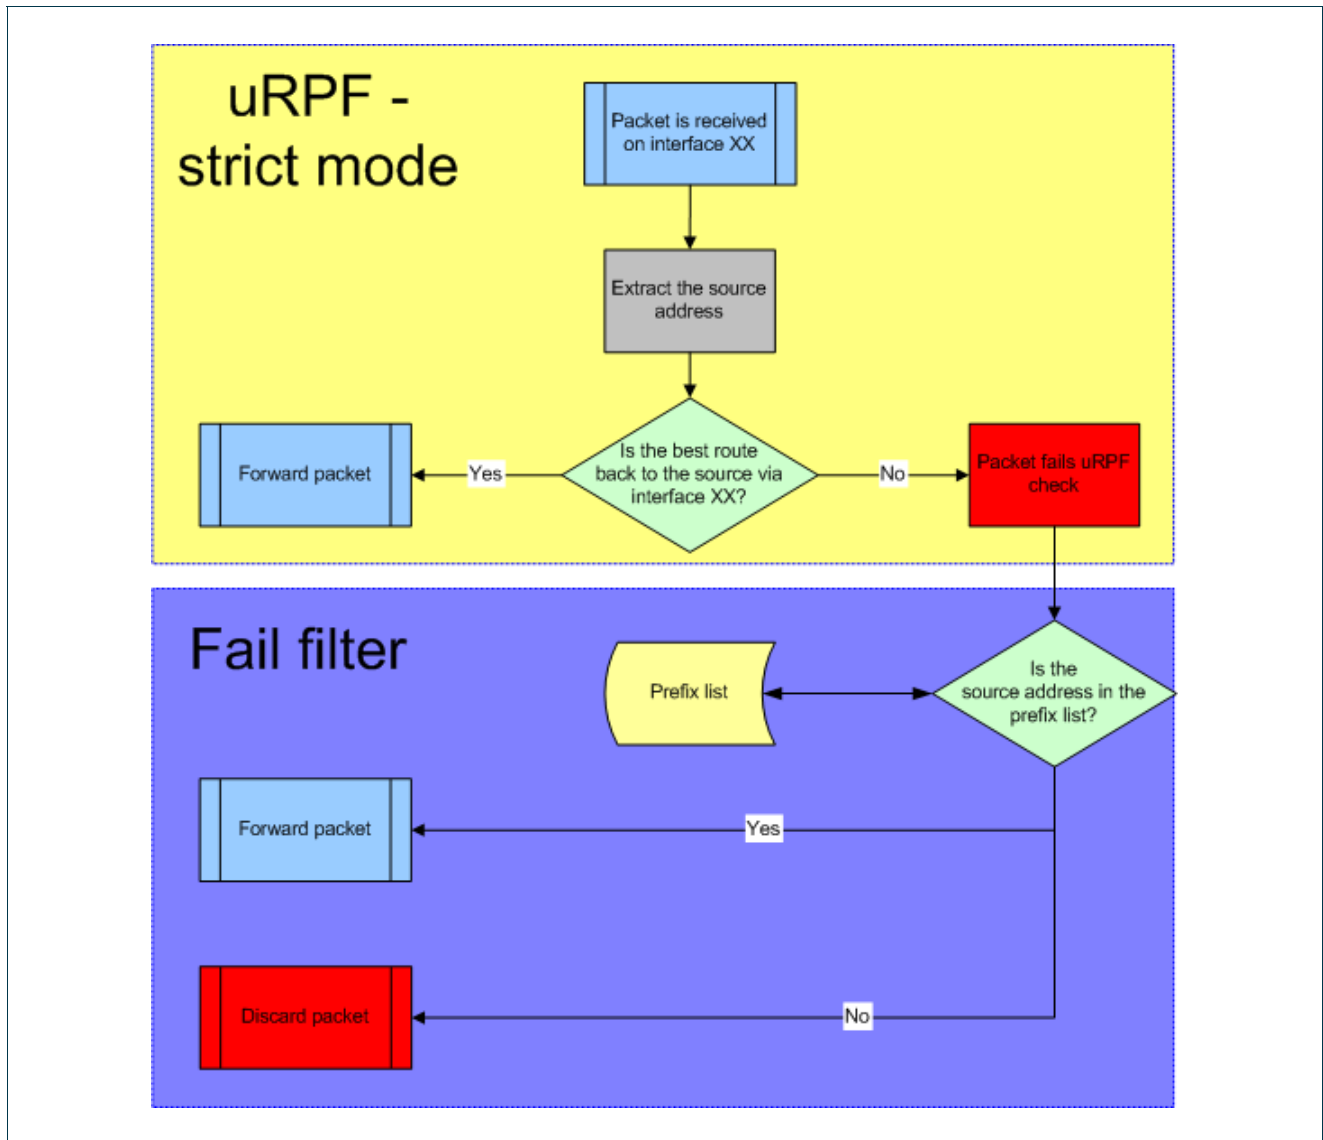


Figure 2.2: uRPF strict operation.

The fail filter in strict mode adds an extra check to packets that would otherwise be discarded. It works by querying a pre-defined list of prefixes. If the source address is within this prefix list then the packet is forwarded.

It is good practice to use the same prefix list to filter BGP announcements and check uRPF. This prefix can often be automatically built using a peers RIPE AS macro.

It is recommended, that NRENs run strict uRPF on interfaces facing sites, to ensure that sites only send traffic that has a source address which corresponds 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 GÉANT2.

2.2 Layer 2 Services

2.2.1 General

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 Layer 3 (distributed test bed).

The only Layer 2 services deployed via GÉANT2 are Layer 2 Virtual Private Networks (VPNs). These have two elements:

- Martini Layer 2 circuits [MARTINI], which provide point-to-point circuit emulation using Multiprotocol Label Switching (MPLS).
- Ethernet VLAN transport using Alcatel MCCs and SDH circuits

Ethernet VLAN transport using the MCCs is sufficiently well secured due to the management and control traffic being transmitted separately from any user data within the VLAN. The following is therefore aimed mainly at securing the MPLS services.

Layer 2 circuits do not have an inter-domain model. This 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 Circuit Cross-Connect (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 must be set up with Resource Reservation Protocol (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 to use traffic engineering to route the LSPs on specific paths that require RSVP. In addition, RSVP is required to run LSPs across a domain border (two different traffic engineering domains).

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

- The NRENs involved have to configure the bandwidth and priority values as well as the strict and loose hops as 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. DANTE 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 that any router in the RSVP domain can freely reserve available bandwidth in the domain. This could make establishing new LSPs impossible. To avoid this risk, the amount of bandwidth that can be reserved on an NREN access interface can be lowered. Note that this only limits the RSVP bandwidth reservation request for establishing the LSP, and does not restrict the effective bandwidth of the layer 2 circuit.

2.2.2 Best practices

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 Figure 3.3:

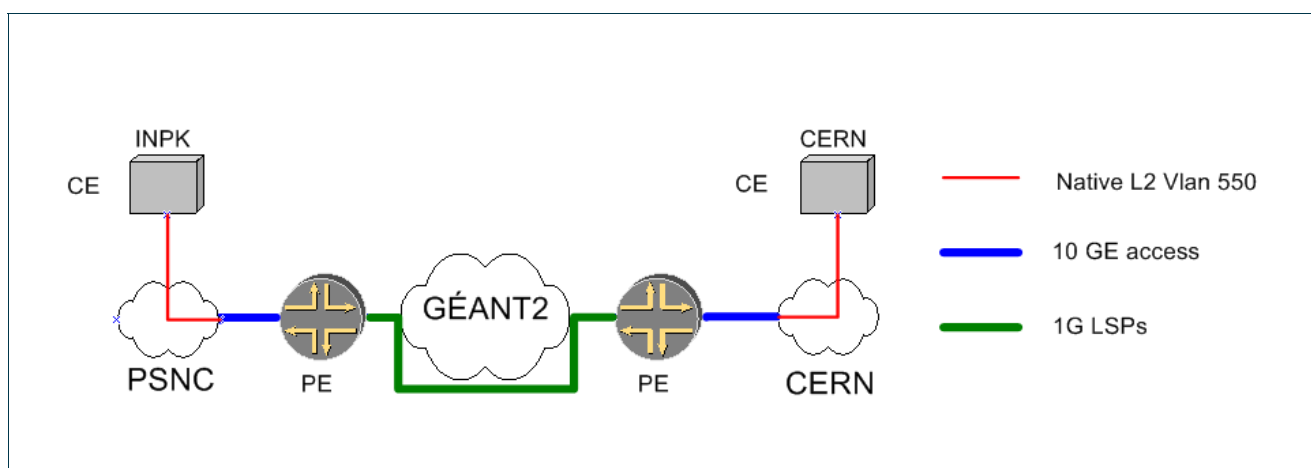


Figure 2.3: Martini Layer 2 circuit – Atlas project.

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. To make this work, the following protocols are needed:

- Label Distribution Protocol (LDP) signalling is used between the loopback interfaces of the two Provider Edge (PE) routers.
- RSVP signalling is used to set up the LSPs. RSVP is used to enable traffic engineering.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

- 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; therefore the inner label is the VC (Virtual Circuit) label and the outer label is for the LSP.

The following table focuses on the control protocols running on the GÉANT2 core routers, and NREN and campus routers. Martini Layer 2 circuits are provisioned using the following protocols:

Protocol	Recommendation for securing
RSVP	RSVP should only be enabled on external interfaces 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). Since the contents of RSVP cannot be filtered, RSVP should be used with care.
LDP	<p>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.</p> <p>It is recommended to use Generalised TTL Security Check mechanism (RFC 3682 [RFC3682]) for LDP which introduces a lightweight security mechanism to protect external LDP sessions from CPU utilisation-based attacks using forged IP packets.</p>
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 2.4: Recommendations for securing layer 2 protocols.

2.3 Layer 1 Services

2.3.1 General

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

- SDH (Synchronous Digital Hierarchy) switching equipment, which can provision SDH circuits and point-to-point Ethernet circuits over SDH wavelengths.
- DWDM (Dense Wavelength Division Multiplexing) equipment, which can provision long haul dedicated wavelengths (lambdas).

DWDM services form the very basis of the network by providing the majority of the IP trunks between routers, point-to-point STM-64/10GE wavelengths for NRENs, and the STM-64 trunks to support the SDH network. The DWDM and SDH equipment in the GÉANT2 core are managed by a centralised NMS (Network Management System).

The NMS system consists of two geographically diverse systems configured as hot failover, so should one system fail the other can resume control with no loss of service. It is essential to secure the NMS against attacks, because all higher layer services could be disrupted if management access to the NMS is compromised. Only DANTE, GÉANT2 NOC (Network Operations Centre) and the vendor have direct access to the NMS. This is accomplished by using firewall filters on both the local router and the NMS system. Also, user logins are audited. The centralised NMS is replicated on a backup system using hot failover. Both systems are geographically diverse and have redundant connections to the main GÉANT2 network. Should one NMS fail for any reason, the other NMS will take control.

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. For an analysis of the DCN for the GÉANT2 core, see *Data Communications Network* on page 23. A DCN is typically proprietary, but the functionality between vendors is very similar. The DCN does not run on any client interface, 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, Data Communication Channel (DCC) signalling should be disabled (see *SDH Equipment* on page 23).

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

SDH and DWDM equipment does not normally use a control plane, as the NMS is centralised. For a GMPLS/G.ASON [G8080] deployment, the network elements create a distributed control plane and take over some functions from the management plane. The new control plane only runs via the DCN, so that 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 is deployed, recommendation on how to secure a UNI interface will be produced.

2.3.2 Best Practices

2.3.2.1 Data Communications Network

SDH and DWDM equipment is managed by a centralised NMS. In the GÉANT2 core, the Alcatel NMS consists of 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, as it acts as an Application Layer Gateway (ALG).

The NMS needs a Data Communication Network (DCN) to be able to communicate with the network elements (NE). SDH and DWDM equipment is managed differently. This is analysed further in section 4.3.2.2 and 4.3.2.3.

The NMS to NE communication is difficult to disrupt because the NEs use proprietary commands and protocols that are transmitted in a separate plane to the user traffic. The NEs are configured only to accept packets from the specified NMS servers (the source address is checked). Although the source address of an incoming packet to GÉANT2 could be spoofed to that of the NMS, it would be dropped at the edge of the GÉANT2 network as part of the Bogon filter rule as the NEs all use private address ranges.

2.3.2.2 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 is carried via the SDH 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) are managed using either OSI between Q3 interfaces. IP access is also used for downloading software updates. OSI traffic can be tunnelled over IP, also if required, IP can be tunnelled over OSI (CLNS). Both the OSI and IP traffic use Link Access Protocol — D Channel (LAPD) as the Layer 2 encapsulating protocol. Dynamic routing protocols are used for both IP and OSI to make all network elements aware of the network topology, so that in the event of an outage on an SDH trunk the DCN traffic is re-routed. For IP, OSPF is used and for OSI, IS-IS is used.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

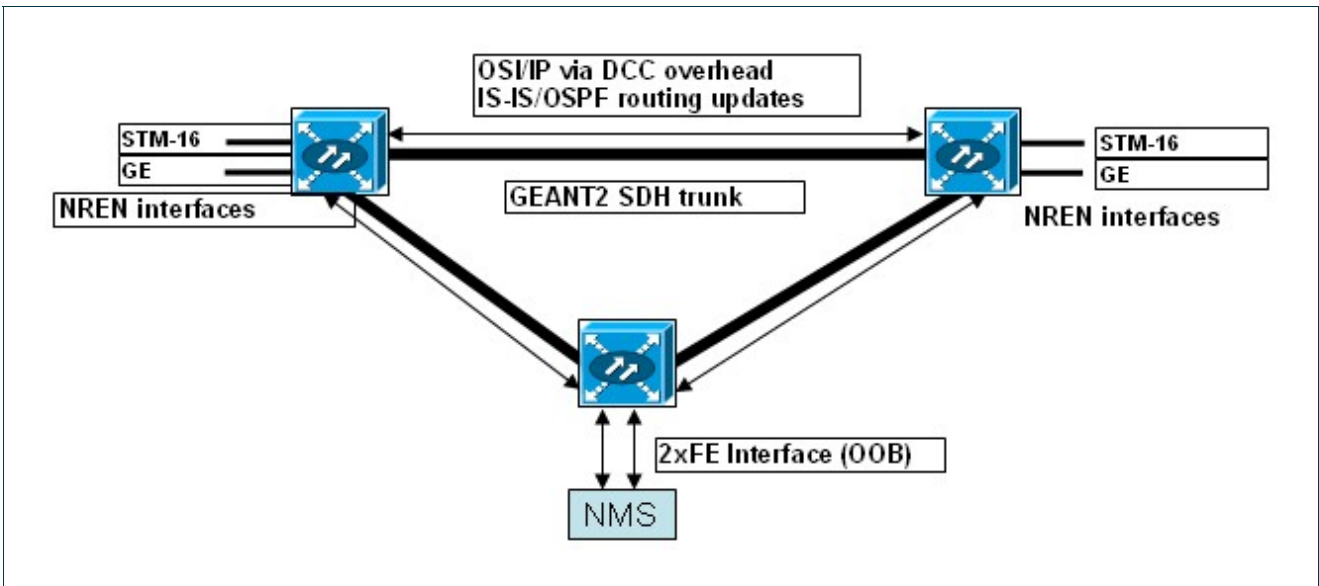


Figure 2.4: SDH switch management.

The 1678MCC can be managed entirely in-band (an MCC in the same Point of Presence, 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 no in-band channels are available.

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 are discarded on ingress. OSI packets are not normally routable across a domain border so OSI packets are discarded on ingress. For GÉANT2, all OSI packets from external peers are discarded.

If an NREN uses an SDH switch to connect to a GÉANT2 SDH switch, DCC signalling should be disabled on that interface. The GÉANT2 SDH switch ignores the DCC overhead bytes from a client interface, so it is not affected, but could cause unnecessary confusion.

2.3.2.3 DWDM Equipment

DWDM equipment is managed using a combination of in-band management (optical supervisory channel, OSC) and out-of-band management (Ethernet interface). The DWDM 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 in Figure 2.5.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

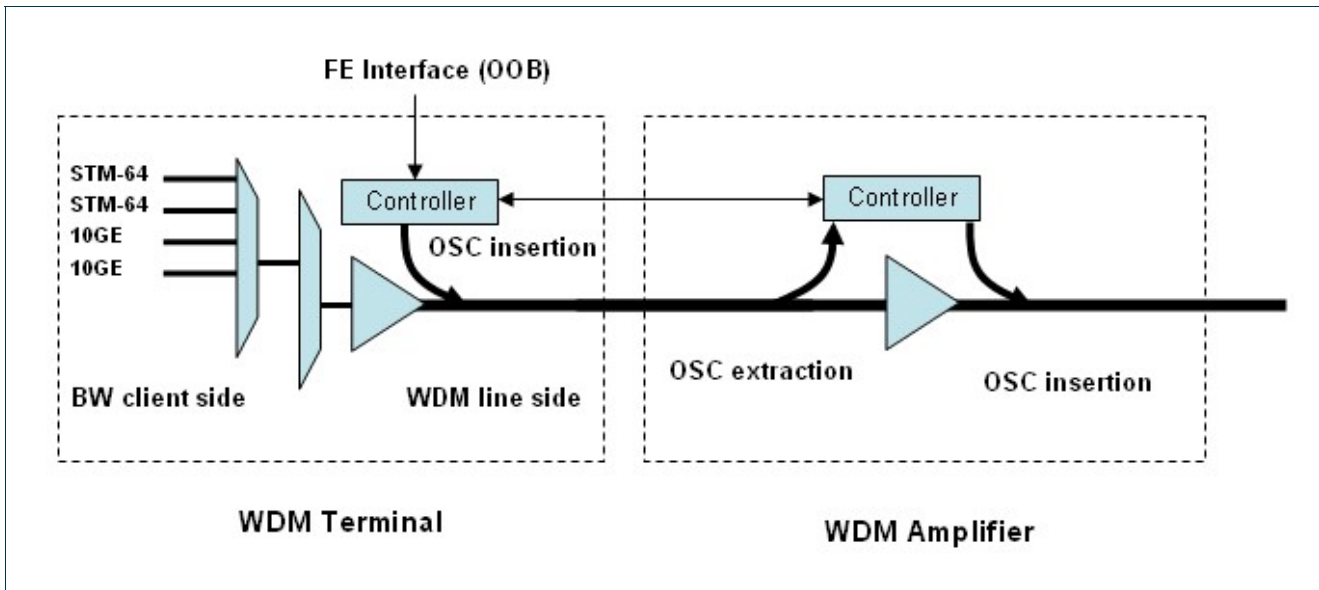


Figure 2.5: Functional diagram of the OSC.

The OSC is generated on a DWDM terminal and electrically regenerated on each WDM in-line amplifier (ILA) until it is terminated on the remote DWDM terminal at the end of the route. For GÉANT2, the OSC provides a 4 Mbps channel. The OSC is inserted and extracted using optical filters. In this way, the controller board of each DWDM network element is reachable in-fibre. The DWDM terminal works as a gateway to an ILA. This 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 DWDM equipment. The diagram below illustrates how the OSC uses a wavelength different to the client signals.

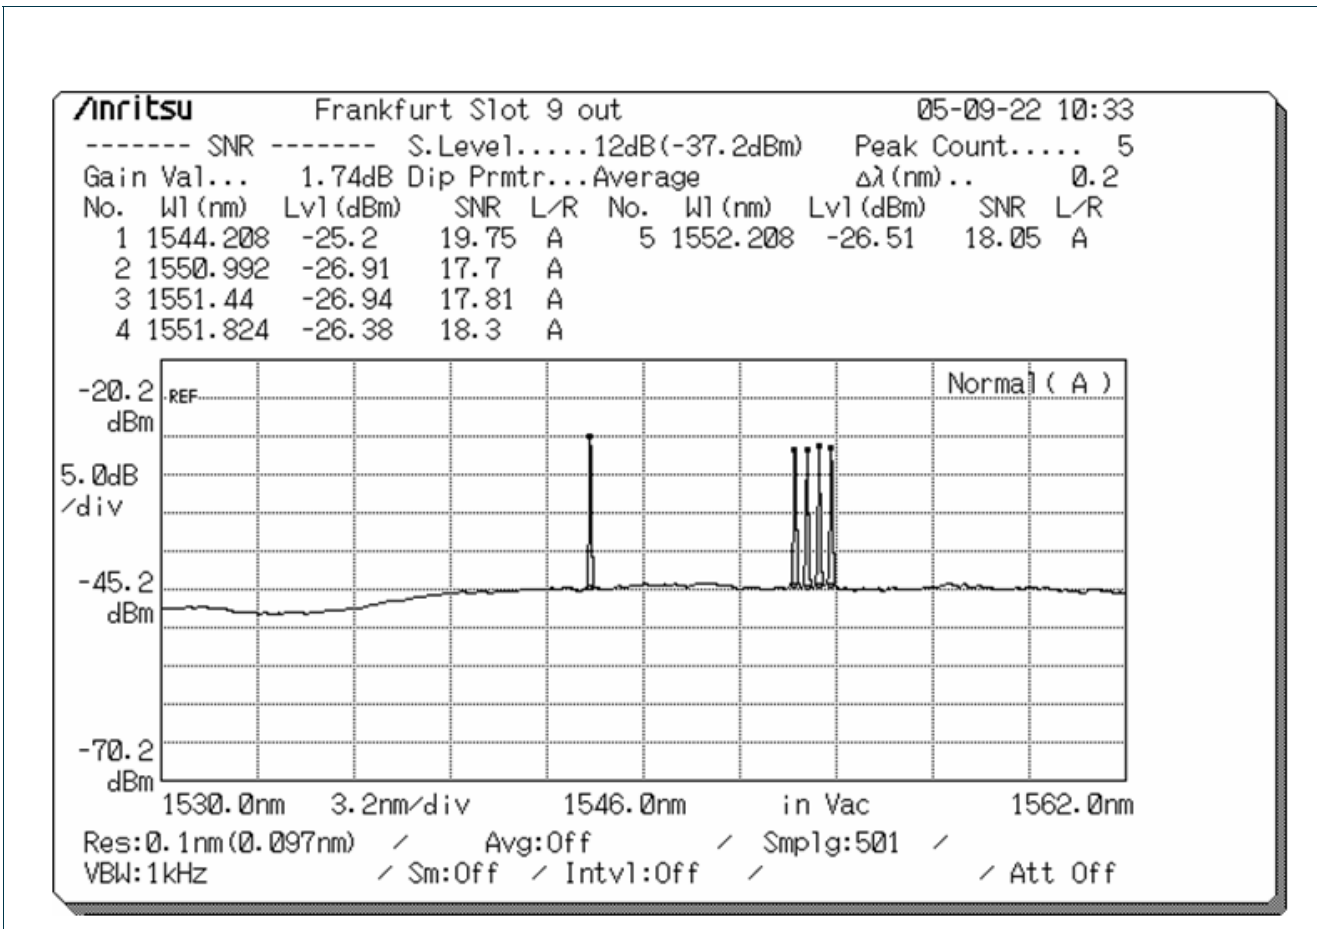


Figure 2.6: Spectrum Analysis showing the OSC at 1544 nm.

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

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

2.4 Physical Security

The physical security of networking hardware must not be overlooked. It stands to reason that a denial of service attack simply has to disrupt network services to be classed as successful. Usually this is done with network traffic, but the same effect can be achieved by unplugging circuits from hardware or disconnecting power. In many cases a physical attack is more damaging than an attack within the network. Network attacks can be investigated and measures put in place to mitigate attacks. However, if a line card is removed from a router, a replacement must be sourced. This is not only expensive but can entail lead times of several weeks.

With the flagship models of some router vendors costing around €1,000,000, equipment is very valuable and attractive to thieves.

Security is easier to establish in large carrier houses that house core network equipment. These usually have 24/7 security guard cover, CCTV, swipe cards to gain entry and locked cages. Racks themselves often have either a key or combination lock. For smaller installations, the same method can be used. Equipment can be secured in locked cages or rooms.

Access rights should be defined in advance with the landlord. These include a list of people allowed to request access and also remote work requests. Every time an access is requested the person requesting it should be authenticated to prevent misuse. Keeping the list of permitted people as small as possible is also good practice. This should include remote access, as an attacker does not have to be physically present to cause harm. If they can request that the landlord perform work such as disconnecting the power to a router then they will still have accomplished their goal of denying network services.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

3 Securing End-to-End Network 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 may be a security problem in that it could be used for transiting data which is in breach of a usage policy.

It is possible that some end-to-end services will be provided using only one type of equipment, but these services can also be hybrids in 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 DWDM 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 GÉANT2 router interface at which the NREN connects (and backup router interface if in use)
- For Layer 2 circuits: The GÉANT2 backup router interface, at which the NREN connects (or primary router interface if not in use). A separate interface can be used with specific permission from DANTE Operations.
- For SDH or DWDM services: An Optical Distribution Frame (ODF) / patch panel 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.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

- Each organisation involved is responsible for securing and monitoring the end-to-end service up to its demarcation points.
- Each organisation involved is encouraged to use perfSONAR [PS] measurement points, or provide access to their monitoring equipment. If no direct monitoring access is possible, access to monitoring data can be given via a proxy server. How to restrict the access to perfSONAR measurement points making use of the eduGAIN [DJ5.2.2,2] infrastructure is described in the deliverable “DJ1.2.4: PerfSONAR AA Service Specification [DJ1.2.4]”.
- 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 viewed by the end sites. The data sent 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.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

4 Securing other Service Delivery Equipment

4.1 Workstations

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 Service Activity (SA) 3 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. As the skill set required to maintain these two sets of components is substantially different, it is anticipated that two different persons rather than one with joint responsibilities will carry out the maintenance. 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 of delegating responsibility to other people who are better able 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.

To secure these workstations, a framework is in place where each workstation is placed into one of the following categories which determine the level of access to the GÉANT2 core and the level of access from the wider Internet. Each category has a specific range of IP addresses and a default firewall rule configuration.

- Internal

Internal is the default and allows the workstations full access to the GÉANT2 core. All traffic going to the internal address range from outside of the GÉANT2 network is dropped at the edge of the network.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

These workstations are often used to monitor the state of the network by communicating with network elements. A proxy system is used to display this information to external clients.

If a workstation needs outside connectivity, its interface is re-assigned an external address, either protected or not. All traffic to the external address ranges is allowed by the border routers and the access is controlled on the local router that the workstation is attached to.

- External

For External connectivity the only filters applied on the routers are to block any traffic from the workstation to the GÉANT2 address space. Any traffic to or from an NREN or upstream provider is permitted. It is the responsibility of the workstation user to secure their workstation. However, should the workstation be compromised and used as a tool in an attack, filters are put in place on the routers to prevent any hostile traffic.

- Protected External

Protected External combines parts of the Internal and External categories. The workstation is placed in an externally accessible address range, but stringent filters are put in place on the local router, and it is advised to also put these on the workstation. The filters have a default rule to discard almost all traffic in and out of the VLAN the workstation is attached to. As well as the discard term there are other terms included in the default template to allow access for DANTE and the GÉANT2 NOC, and also ICMP traffic for ping and traceroute applications. The ICMP traffic is subject to a rate limit to prevent a ping flood type DoS attack from affecting the workstation.

The basic Protected External template is then modified to cater for the needs of the project the workstation is assigned to. This usually includes permitting SSH traffic to and from a known group of hosts to allow developer access, and often specific ports that a server is listening on. These terms are always as specific as possible and specify the port number and IP address(es).

The IP address ranges are assigned to separate 802.1Q VLANs on the router's Gigabit Ethernet (GE) interface. This allows firewall filters to be set per VLAN and therefore per address range. The GE interface from the router is connected to a switch. This switch must be configured to allow access from a specific port to a specific VLAN. By default the switch ports are not configured with any VLAN tags, so if a rogue machine was to be plugged into the switch, it would not be able to communicate with any other machines.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

5 Protecting End Customers

Anomaly detection tools deployed in core networks are very useful to enhance end-customer protection, and complement edge defences that campus area network administrators or end users should (but not necessarily do) have in place. Examples of anomalies that can be effectively detected are:

- Hosts in the community being the target of scanning activity, DoS or DDoS attacks.
- Hosts participating in scanning or sourcing DoS attack traffic.
- Hosts connecting to suspicious sites (e.g. known botnet Command & Control servers), indicating that the host is potentially infected with malware.
- Generic indications of worm spreads and scanning activities.

5.1 Netflow-Based Anomaly Detection

Ongoing DoS or DDoS attacks on hosts of the GÉANT2 (GN2) community can be effectively detected with Netflow analysis tools such as Netflow Sensor (NfSen). If the attack is coming from an easily identifiable set of hosts, it is possible to deploy ACLs in the routers of the network to blacklist these sources and avoid a continued or recurring attack. Also, trying to get in touch with the administrators of the attacking end sites might help (note that we are not addressing what security practices administrators should follow after such alerts, like patching OSs, installing anti-viruses and firewalls, and so on).

However, more frequently attacks come from a distributed multitude of sources. This makes it more difficult to defend against recurring attacks, since the sources are likely to be infected themselves and under the control of a bot commander. Still, if this can be communicated to some cooperative administrators, it can be possible to identify the communication channel used to command the hosts and thus block further occurrences of the attack. Another degree of uncertainty is introduced by spoofing, as most attacks come from spoofed sources. However, this should not happen if uRPF is implemented at network boundaries.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

The detection of scanning activities is also possible with NetFlow analysis tools. Both hosts that are the target of a scan or scanners of other hosts can be evidenced. When either one or the other end is found within the GN2 community, informing the end administrators should help in mitigating the consequences of the scan. The scan itself is not producing any harm but it may be preliminary to a DoS or DDoS attack on an open service on a system (when the systems receives the scan), or it may be an indication that a system is infected by a worm or a Trojan and is trying to automatically “recruit” other victims to spread the worm or Trojan. Another effective way to identify end hosts potentially under the control of a bot commander is to look (always through Netflow data) for the connection to known botnet Command & Control servers which should be enough to trigger an analysis on the end system to look for possible Trojan software presence.

Netflow data analysis can be used for detection of the anomalies described above. Also in core network where Netflow collection is coming from sampled packets with small sampling rates (e.g. 1/100 or 1/1000) these activities can be easily identified. What cannot be guaranteed, especially in cases where scans are for specific ports and just involve the exchange of a few packets, is that all the hosts being scanned are identified. It may well happen that some of them will not have even one single sampled packet.

It should be noted that there are various versions of Netflow records. The most common is version 5, which provides the basic information. If the hardware supports it, it is recommended to use version 9 which allows the recording of MPLS labels and also IPv6 addresses and ports. A network carrying IPv6 traffic should consider it a necessity to use version 9, as without it network operators are left blind to any attacks based on IPv6.

5.2 Darknets

The Netflow-based anomaly detection described in *Netflow-Based Anomaly Detection* on page 32 has the advantage that it can immediately pinpoint the host(s) involved. Its drawback, however, is that false positives (legitimate traffic mistaken for suspicious activity) are inevitable. The deployment of Darknets reverses this situation. The term “Darknet” has the following definitions:

- A private virtual network where users can only connect to a group of individuals that they trust. In this sense Darknets can be seen as covert communication networks. This definition is often used to refer to file-sharing networks.
- Any routed network space that does not have any networking devices, but serves as a virtual black hole in which data that flows into this segment is logged for analysis. The objective behind this kind of setup is that as there are no devices within this address space, no traffic should be destined for it. This raises an interesting point as it implies that traffic entering this network could be of a malicious nature.

The advantage of using Darknets is that they are relatively easy to set up and do not require costly equipment to work. They are useful tools for detecting suspicious activities without having to filter the false positives that are prevalent on conventional traffic segments. As there is no legitimate traffic on the Darknet segment, all traffic is worthy of investigation. Attackers often hide their activities within the conventional stream of network traffic, making it more difficult to detect their attacks. The disadvantage of using a Darknet is that, if its use is known, it may be bombarded with meaningless traffic, rendering it ineffective.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

Tracking compromised machines in large networks can be extremely difficult. Very often the solutions in play are not able to scale to large ISP scale networks, where the solution is to move the traffic rather than investigate a hard to diagnose single machine.

Team Cymru describes the Darknet as “a portion of routed, allocated IP space in which no active services or servers reside” [CYMRU]. The term dark refers to the fact that these networks have no services or anything within them.

The only device that may reside within this segment is the server that collects all the packets sent to this network. These netflows/packets can then be used for analytical and forensic purposes. As previously mentioned, no legitimate traffic should enter the network, although misconfiguration may result in certain traffic getting through. The majority of the traffic that enters the network is therefore a result of malware. The purpose of malware is to scan for vulnerable devices on the segment. The scanning action, therefore, stands out and can be analysed as an activity on the segment.

A Darknet can have multiple uses. It can, for example, be used as flow collectors, packet sniffers and IDS devices without the problem of false positives. The objective of a Darknet is to create awareness and to ease the mitigation of threats to the network by making bad traffic and its source easily identifiable.

5.2.1 Creating a Darknet

Creating a Darknet requires address space, for example, a /24 for IPv4 or /64 in case of IPv6. By allocating more address space to the Darknet, its visibility to outside traffic can be increased. As the Darknet is a segment of routed space, Bogon address space must not be used because these prefixes are often excluded from scanning activities. To protect the device from malicious threats, flow collector and management device addresses should not be in the same address space as the Darknet segment. As the purpose of the Darknet is to detect malicious traffic, no legitimate traffic should be routed through this segment and poison the results.

Figure 3.5 shows a possible implementation of a Darknet, displaying the logical implementation of the Darknet segment.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

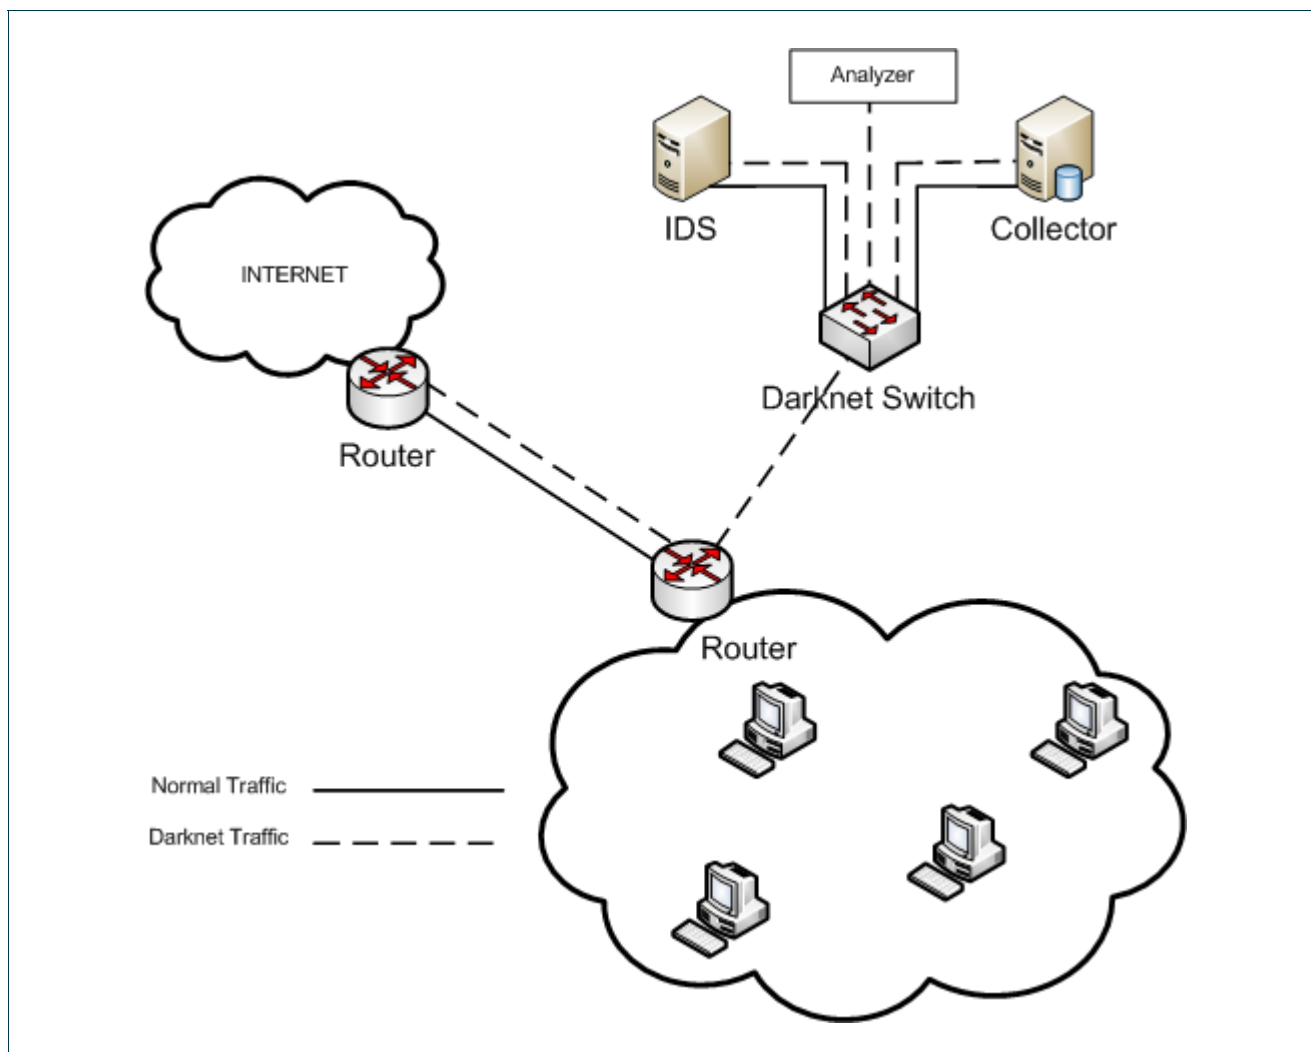


Figure 5.1: A Darknet implementation.

5.2.2 Using Darknet Data

There are many possible uses for the data collected from a Darknet implementation:

- Various CERTs publish advisories on exploits, listing the ports used in these attacks. This information can be compared to the tcpdumps from the segment to check if any of the attacks are targeting your network. You can then give this information to customers within your network, providing them with source IP addresses and scan timestamps, so they can clean the infected devices and create filters on firewalls and routers.
- The information gathered from the Darknet can be used to generate reports, advisories for within the local security community and for internal statistical purposes.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

- Team Cymru notes that a Darknet can also be used as a "garbage meter" to determine the volume of bad traffic on the network [CYMRU].
- From a management perspective, Darknets provide a proactive approach toward security as they enable detection of malicious activity at an early stage, and prevent malware from spreading its payload or spamming.

For these reasons, Darknets are a very powerful security tool that CERTs can use to spot scanning without the need to have complex analysis equipment. The Darknet setup can be extended to include an Intrusion Detection System such as SNORT [SNORT].

Further reading on an NRENs experience of Darknets is available online [TELESCOPES].

5.3 Black Holes

In network terms, a black hole is something that silently discards incoming traffic. This can be both useful and dangerous.

Black holes are useful to mitigate the effects of a denial-of-service attack on the network. The nature of routing protocols is such that a router will almost always send traffic to the next hop that advertises the most specific prefix. By setting a more specific route in the routing table than the one currently known, you can pull traffic away from its usual path.

The disadvantage of black holes is that black holed traffic is difficult to detect, and the only way to see where traffic has disappeared is to compare traffic graphs.

Black holes can be implemented in a network using one of the following methods:

- You can advertise the target IP address as a /32 prefix in case of IPv4 and /128 in case of IPv6 from your own routers. This method has the benefit that you retain control of which traffic is black holed, but requires constant monitoring to determine when an attack has passed.
- You can set up a specific black hole router. Customers can peer with this router and announce to it /32 or /128 addresses that they wish to black hole. The routes are subsequently advertised within the AS. This draws all traffic destined for the customer /32 or /128 address to the black hole router. The router then discards all traffic that hits it. The black hole router must be placed on high bandwidth trunks or near the edge of the network, so that DoS traffic does not impact legitimate traffic. This method has the benefit that customers can make the decision whether to black hole their traffic or not. You must, however, ensure that the customer can only advertise their own address space, otherwise you would have created an efficient DoS tool.

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

6 Conclusion

If new service requirements are defined, JRA2-WI1 will analyse the requirements and present security recommendations that 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 [NRENOPS]. Any changes are announced to the GÉANT2 Access Port Manager (APM) list.

The best practice sections are meant to be a shared resource that NRENS and other GÉANT2 service users can access and help to maintain.

7 References

- [DJ1.2.4] M. Molina, A. Solberg, D. Lopez, J.M. Macias, J.W. Boote, "D.J.1.2.4: PerfSONAR AA Service Specification"
- [DJ2.1.1,4] M. Wright, M. Molina, M. Mogensen, D. Kalogeras, J. Mohacsi, H. Nussbacher, M. Garcia, R. Sabatino, "D.J.2.1.1,4: Revised GÉANT2 Security Recommendation and Policy"
- [DJ5.2.2,2] D.R. Lopez, R. Castro, B. Kerver, T. Lenggenhager, M. Linden, I. Melve, M. Milinovic, M. Molina, J. Rauschenbach, M. Stanica, K. Wierenga, S. Winter, H. Ziemek "D.J.5.2.2,2: GÉANT2 Authorisation and Authentication Infrastructure (AAI) Architecture – second edition"
- [BCP38] <http://tools.ietf.org/html/bcp38>
- [BGP] <http://www.space.net/~gert/RIPE/ipv6-filters.html>
- [BOGON] <http://www.cymru.com/Bogons/>
- [Bogon list] <http://puck.nether.net/mailman/listinfo/bogon-announce>
- [CYM] <http://www.team-cymru.org/documents/60Days.ppt>
- [CYMRU] <http://www.team-cymru.org/Services/darknets.html>
- [DARKNET] <http://www.cymru.com/Darknet/>
- [MARTINI] <http://www.ietf.org/internet-drafts/draft-martini-l2circuit-encap-mpls-10.txt>
- [FLAW] http://www.secdev.org/conf/IPv6_RH_security-csw07.pdf
- [G8080] 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.
- [IANA1] <http://www.iana.org/assignments/ipv6-unicast-address-assignments>
- [IANA2] <http://www.iana.org/assignments/ipv6-address-space>
- [ICMP] <http://www.cymru.com/Documents/icmp-messages.html>
- [JUNIPER] <http://www.juniper.net/alerts/viewalert.jsp?txtAlertNumber=PSN-2007-04-034>
- [NRENOPS] <http://intranet.GÉANT2.net/server/show/nav.922>
- [OSI] [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)
- [RIPE] <http://www.ripe.net/docs/ripe-378.html>
- [PS] <http://www.perfsonar.eu>
- [RFC1072] <http://tools.ietf.org/html/rfc1072>
- [RFC1918] <http://tools.ietf.org/html/rfc1918>
- [RFC2463] <http://tools.ietf.org/html/rfc2463>
- [RFC2827] <http://tools.ietf.org/html/rfc2827>
- [RFC3682] <http://tools.ietf.org/html/rfc3682>
- [RFC4890] <http://tools.ietf.org/html/rfc4890>
- [RFC5095] <http://tools.ietf.org/html/rfc5095>

[RFC5156]	http://tools.ietf.org/html/rfc5156
[SMURF]	http://en.wikipedia.org/wiki/Smurf_attack
[SNORT]	http://www.snort.org/
[UNI]	http://www.oiforum.com/public/impagreements.html#UNI
[TELESCOPES]	http://noc.ilan.net.il/research/telescope/
[WIKI]	http://wiki.geant2.net/bin/view/JRA2/Jra2WorkingArea#D2_1_1_v5

8 Acronyms

[AAA]	[Authentication, Authorisation and Accounting]
[ACL]	[Access Control List]
[ALG]	[Application Layer Gateway]
[APM]	[Access Port Manager]
[ARP]	[Attribute Release Policy]
[AS]	[Autonomous System]
[ASPATH]	[Autonomous System Path]
[AUP]	[Acceptable Use Policy]
[BCP]	[Best Common Practice]
[BGP]	[Border Gateway Protocol]
[BSR]	[Bootstrap Router]
[BW]	[Black and White interface]
[CCC]	[Circuit Cross-Connect]
[CCTV]	[Closed Circuit Television]
[CDP]	[Cisco Discovery Protocol]
[CLNS]	[Connectionless Network Service]
[CMISE]	[Common Management Information Service Element]
[CORBA]	[Common Object Request Broker Architecture]
[CPU]	[Central Processing Unit]
[DCC]	[Data Communication Channel]
[DCN]	[Data Communication Network]
[DDoS]	[Distributed Denial-of-Service]
[DoS]	[Denial of Service]
[DWDM]	[Dense Wavelength Division Multiplexing]
[eBGP]	[external Border Gateway Protocol]
[eMSDP]	[external Multicast Source Discovery Protocol]
[FIB]	[Forwarding Information Base]
[GE]	[Gigabit Ethernet]
[GN2]	[Geant2]
[HMAC]	[Hash Message Authentication Code]
[HTTP]	[Hypertext Transfer Protocol]
[HTTPS]	[Hypertext Transfer Protocol Secure]
[IANA]	[Internet Assigned Numbers Authority]

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170

[IGMP]	[Internet Group Management Protocol]
[ILA]	[In-Line Amplifier]
[IP]	[Internet Protocol]
[IPSEC]	[Internet Protocol Security]
[ISP]	[Internet Service Provider]
[ISO]	[International Standards Organization]
[IS-IS]	[Intermediate System-to-Intermediate System]
[JRA]	[Joint Research Activity]
[LAN]	[Local Area Network]
[LAPD]	[Link Access Protocol — D Channel]
[LDP]	[Label Distribution Protocol]
[LSP]	[Label Switched Path]
[MAC]	[Media Access Control]
[Mbps]	[Megabits per second]
[MCC]	[Metro Core Connect]
[MD5]	[Message Digest Algorithm 5]
[MLD]	[Multicast Listener Discovery]
[MPLS]	[Multiprotocol Label Switching]
[MSDP]	[Multicast Source Discovery Protocol]
[MTU]	[Maximum Transmission Unit]
[NE]	[Network Element]
[NfSen]	[Netflow Sensor]
[NMS]	[Network Management Station]
[NOC]	[Network Operations Centre]
[NREN]	[National Research and Education Network]
[NTP]	[Network Time Protocol]
[ODF]	[Optical Distribution Frame]
[OOB]	[Out Of Band]
[OSC]	[Optical Supervisory Channel]
[OSI]	[Open Systems Interconnection]
[OSPF]	[Open Shortest Path First]
[PE]	[Provider Edge]
[PIM]	[Protocol-Independent Multicast]
[PoP]	[Point of Presence]
[Q3]	[CMISE Q3 OPTICS-IM based management interface]
[RFC]	[Request for Comments]
[RIP]	[Routing Information Protocol]
[RIPE]	[Réseaux IP Européens]
[RIPng]	[Routing Information Protocol next generation]
[RIR]	[Regional Internet Registry]
[RP]	[Rendezvous Point]
[RSVP]	[Resource Reservation Protocol]
[SA]	[Service Activity]
[SAs]	[Source Announcements]
[SDH]	[Synchronous Digital Hierarchy]

[SNMP]	[Simple Network Management Protocol]
[SSH]	[Secure Shell]
[STM]	[Synchronous Transport Module]
[TL1]	[Transaction Language 1]
[TTL]	[Time To Live]
[uRPF]	[Unicast Reverse Path Forwarding]
[UTC]	[Universal Time Coordinated]
[VC]	[Virtual Circuit]
[VLAN]	[Virtual Local Area Network]
[VPN]	[Virtual Private Network]
[WDM]	[Wavelength-Division Multiplexing]

Project:	GN2
Deliverable Number:	DJ.2.1.1,5
Date of Issue:	06/08/08
EC Contract No.:	511082
Document Code:	GN2-08-170