

SEC
Access
Certification

26 August
2015

Content

1	Change Log	6
1.1	License	9
1.2	INTRODUCTION	10
1.3	Requirement Clauses	11
1.4	Reference model	12
1.5	Test	13
2	SEC - Common	14
2.1	MAC Flooding	14
2.2	Multicast injection	16
2.3	Spanning-Tree	18
2.4	Other protocols	20
2.5	Control Plane Protection	22
2.6	Network Protection	27
2.7	Network Equipment Management	29
2.8	Further Study	30
3	SEC - IPv4	31
3.1	IPv4 overview	31
3.2	ARP	35
3.3	ICMPv4	38
3.4	IPv4 Source Address Spoofing	40
3.5	IGMP	42
3.6	Multicast group overload	43
3.7	Protocol protection	43
3.8	Forced Forwarding	50
3.9	Secure Boot, Access Switch	52
3.10	VRRP / HSRP spoofing	53
3.11	Routing Protocols	55
4	SEC - IPv6	56
4.1	IPv6 overview	56
4.2	IPv6 Whitelist Database	57
4.3	IPv6 source address spoofing	60
4.4	Neighbor Solicitation / Advertisement	61
4.5	Neighbor Unreachability Detection	63
4.6	ICMPv6 Redirect Messages	65
4.7	Duplicate Address Detection - DAD	67

4.8	Router Advertisement attack.....	69
4.9	IPv6 Routing Header	71
4.10	MLD/MLDv2	72
4.11	IPv6 Multicast Group Overload	73
4.12	ND cache	74
4.13	Protocol Protection	76
4.14	Forced Forwarding	79
4.15	Secure Boot, Access Switch	80
4.16	VRRPv6 / HSRPv6	81
5	<i>Sample configurations</i>	82
5.1	Sample 1, L2 in access, client traffic exchanged on L2. shared VLAN/L2 for customers	82
5.2	Sample 2, L2 in access, client traffic exchanged on L3, shared VLAN/L2 for customers	84
5.3	Sample 3, L2 in access, client traffic exchanged on L3, separate VLAN/L2 per customer port	86
5.4	Sample 4, L3 in Access, client traffic exchanged on L3, separate L2 per customer port, no VLANs	88
6	<i>IPv6 Standards</i>	90
6.1	General	90
6.2	Address Selection.....	91
6.3	DHCP	91
6.4	DNS	91
6.5	Multicast.....	91
6.6	Transition	92
6.7	CPE.....	92

Requirements, Common

<i>Limit number of MACs per customer port.....</i>	<i>15</i>
<i>Report when the number of MACs per customer port has been reached</i>	<i>15</i>
<i>Filter unwanted sending of IPv4 and IPv6 multicast from customer ports.</i>	<i>17</i>
<i>Select what multicast groups can be received on customer ports.....</i>	<i>17</i>
<i>Disable spanning-tree on customer ports</i>	<i>19</i>
<i>Disable unsupported protocols on customer ports.....</i>	<i>21</i>
<i>Protect the control plane from excessive traffic.....</i>	<i>23</i>
<i>Avoid starvation of packet buffers due to excessive traffic.....</i>	<i>25</i>
<i>Drop fragmented IPv4 and IPv6 packets before reaching the control plane.</i>	<i>25</i>
<i>Limit broadcast, multicast and unicast flooding traffic, per port.....</i>	<i>27</i>
<i>Detect loops on customer ports, if loop detected disable the port.</i>	<i>28</i>
<i>If a loop is detected, log the event.....</i>	<i>28</i>

Requirements, IPv4

<i>Prevent users from communicating with equipment management.</i>	<i>29</i>
<i>Log all assigned dynamic addresses</i>	<i>31</i>
<i>DHCPv4 requests tagged with customer port identification.....</i>	<i>33</i>
<i>Drop packet or replace customer specified port identification if present in DHCPv4 requests</i>	<i>33</i>
<i>Filter out ARP Poisoning packets</i>	<i>37</i>
<i>Filter out ICMP redirect packets.....</i>	<i>39</i>
<i>Pass on valid ICMP traffic.....</i>	<i>39</i>
<i>Filter out spoofed IPv4 packets in the access node.</i>	<i>41</i>
<i>Track when spoofed IPv4 packets are received</i>	<i>41</i>
<i>Do not use IGMPv1</i>	<i>42</i>
<i>Use the IGMPv2/IGMPv3 messages to learn about interesting multicast receivers.....</i>	<i>42</i>
<i>Limit number of concurrent multicast groups a customer port can join using the IGMP protocol.....</i>	<i>43</i>
<i>Filter out all DHCP server traffic coming from customer ports.....</i>	<i>45</i>
<i>Track invalid DHCPv4 server packets.....</i>	<i>45</i>
<i>Limit the number of IPv4 addresses that can be assigned to a customer port.....</i>	<i>46</i>
<i>IPv4 UPnP (SSDP), LLNMR, mDNS and Bonjour must be blocked between customer ports.</i>	<i>48</i>
<i>Drop packets with too small fragment offset</i>	<i>49</i>
<i>Filter out IPv6 when IPv6 isn't intended to be used.....</i>	<i>49</i>
<i>Drop packets between customer ports in the same L2 broadcast domain.....</i>	<i>51</i>
<i>When Access Switch boots, only accept DHCPv4 / BOOTP packets from uplinks.....</i>	<i>52</i>
<i>VRRP/HSRP from customer ports should be blocked.....</i>	<i>54</i>

Protect routing protocols from malicious users..... 55

Requirements, IPv6

Identification info must be added to the DHCPv6 snooped packets..... 57

IPv6 whitelist database..... 58

Filter out spoofed IPv6 packets in the access switch..... 60

Track when spoofed IPv6 packets are received..... 60

Filter out ND/NS messages with unknown source IPv6 addresses..... 62

Neighbor Unreachability Detection [NUD, RFC4861] filtering..... 63

Filter out incorrect ICMPv6 redirect messages on customer ports..... 66

Drop invalid replies to IPv6 Duplicate Address Detection packets..... 68

Block Received Router Advertisement packets on customer ports..... 70

Drop all received packets on customer ports that contains a type 0 Routing Header..... 71

Use MLD/MLDv2 messages to learn about interesting multicast receivers..... 72

Protect against filling up ND cache..... 75

DHCPv6 must be blocked between subscriber ports..... 76

IPv6 UPnP, LLNMR, mDNS and Bonjour must be blocked between subscriber ports..... 77

Drop IPv6 packets with too small fragment offset..... 78

Drop packets between customer ports in the same L2 broadcast domain..... 79

VRRPv6/HSRPv6 or similar protocol must be protected against malicious users..... 81

1 Change Log

Version 2015-08-26, Torbjörn Eklöv

Number	Change
4.2.5	Clarified between DHCPv6 and SLAAC

Version 2015-05-12, Torbjörn Eklöv

Number	Change
4.13.3	Marked red
3.7.5	Filter out IPv6

Version 2014-05-30, Torbjörn Eklöv

Number	Change
3.4.3	Switch -> Node
4.13.12.1	From address to address / prefix

Version 2014-05-19, Torbjörn Eklöv

Number	Change
4.2.5	Removed per vlan req

Version 2013-05-27, Marcus Jonsson

Number	Change
Many	Added details on how tests should be performed for Common and IPv4
2.7	Added new requirement, "Network Equipment Management"
3.2	Requirement 3.2, "IPv4 Whitelist Database" was removed. It overlapped with some other requirements.
3.7.3	Changed name from "IPv4 UPnP" to "IPv4 Multicast Discovery"

Version 2013-01-31, Torbjörn Eklöv

Number	Change
3.8.3.2	Added mDNS, LLMNR and Bonjour to Upnp rule
Many	Marked minimum level for SEC testing IPv4 and IPv6 with red
4.13.3.1	Added mDNS, LLMNR and Bonjour to Upnp rule
Some	Changed Must to Should

Version 2012-08-27, Torbjörn Eklöv

Number	Change
3.6.2	Must to should
3.8.1.3	Added must and should
3.8.3.3	Should to Must
3.11.3	Must bot not tested
4.1.3	Empty ->Must
4.2.5	Should to must
4.3.1,4.4.3,4.5.3,4.6.3	Empty -> Must
4.6.4	Added Force Forward
4.7.3	Empty -> Must
4.7.4	Added Force Forward
4.8.3,4.9.3	Empty -> Must

4.10.2	Must to Should
4.12.3	Should to Must
4.13.1.1	Empty -> Must

Version 2012-04-04, Torbjörn Eklöv

Number	Change
4.2	Typo, changed IP4 to IPv6.
5.3	Added: "Identification info must be added to the DHCPv6 snooped packets, with option 18 and 37 to identify customer ports" in IPv6 section

Version 2012-03-21, Anders Löwinger

Number	Change
1.1	Changed license from CC BY-SA to CC BY. Added copyright notice.
1.2	Changed the link to the English http://secureenduserconnection.se/
4.5.1, 4.5.2	Neighbor Unreachability Detection. Added picture and description
4.9.1, 4.9.2	IPv6 Routing Header. Added picture and description
4.12	ND Cache. Added picture and description
5.1, 5.2, 5.3, 5.4	Sample Configurations. Added better description on how each scenario is designed.
All tables	Fixed broken borders
Document	Changed to .xlsx format

Thanks to Henrik Nordström Consulting for review and feedback.

Version 2012-03-19, Anders Löwinger

Number	Change
Multiple	Changed "Implementation" headings to "Implementation Example" to be clear that this is not requirements.
Multiple	Changed "address assignment" to "address configuration"
1.3, all requirements	Requirement clauses. Added Test field to each requirement
1.5	New section "Test". The test setup is described
2	Added a test description on all requirements
2.5.2.1	Avoid starvation of packet buffers. Added test description, this will not be part of SEC certification
2.6.2.2	Loop detect. Added implementation detail, could possibly use Ethernet Loopback 0x9000
3.1.2	IPv4 address configuration. Added some techniques that can share IPv4 address prefixes between L2 broadcast domains
3.2.3	DHCPv4 snooping. Added reference to IETF draft
3.4.2	Attack – ICMPv4 redirect. Changed picture, it was incorrect and had IPv6 addresses
3.6.3	IGMP snooping. Added clarification why RFC4541 exist and should be used.
4.2.3	DHCPv6 snooping. Added link to IETF draft, SAVI solution for DHCP
4.3.2	IPv6 source address spoofing. Added description of IETF SAVI and some links to the drafts.
4.8.4	Router Advertisement Attack. Added link to IETF draft "Implementation advice for Ipv6 router Advertisement Guard (RA-Guard)"
4.12.4	ND Cache. Added link to IETF draft "Operational Neighbor Discovery Problems"
4.13.4.2	IPv6 Fragmentation attack. Added link to IETF draft "Processing of IPv6 atomic fragments" and "Tiny fragments in IPv6"

Thanks to Jakob Schlyter @ Kirei AB for review and feedback.

Version 2011-11-25, Anders Löwinger, first draft

Number	Change
	First draft, for internal review

SEC Access Certification

1.1 License

SEC Access Certification. Revision 2012-08-26

© Anders Löwinger and Torbjörn Eklöv 2011-2015

The text and illustrations are protected under copyright law and is provided under the license Creative Commons License, Attribution CC BY 3.0, as described here

<http://creativecommons.org/licenses/by/3.0/>



You are free:

to Share — to copy, distribute and transmit the work

- to Remix — to adapt the work

Under the following conditions:

- Attribution — you must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).

Attribution should be done to the document and the authors.

- Anders Löwinger, Abundo AB, anders@abundo.se
- Torbjörn Eklöv, Interlan Gefle AB, torbjorn.eklov@interlan.se

Any derivative work of this document must remove the SSNF logotypes and graphical elements. They are protected by law and are not included by the Creative Commons License above.

1.2 INTRODUCTION

SKA (Säker Kund Anslutning) started 2005 with the purpose to ensure that the Swedish Broadband market doesn't contain incorrectly built and insecure networks. It was a fact that Broadband Networks was designed and built with insecure connections and security issues, which end-users could not control and fix. Sad to say, these types of insecure networks are still being built.

In the beginning SKA only certified hardware, but from 2009 a complete Broadband Network can be certified. It is an excellent method for a Broadband Network to show that they have good security and their end-users can feel safe when using the network.

After more than five years of SKA, the product is internationalized and renamed to SEC – Secure End user Connection. The main reason to change name and language is to avoid language problems with international vendors. The homepage <http://secureenduserconnection.se/> will only contain English information and the SEC documents will also be written in English.

This document describes the various issues, problems and possible solutions that exist when deploying L2 networks in the edge using the IPv4 and IPv6 protocol. Most of the issues cannot be mitigated by end-users; it is the broadband network that must create the protections.

The demands were primary written for networks with shared broadcast domains, but all the issues should be tested independent of what type of design/network topology is used (for example one VLAN per customer). In our experience there are many possible configuration errors in networks that can create security issues.

SEC focuses mainly on

- Man-In-The-Middle (MITM) attacks
 - The ability to eavesdrop and possibly change traffic without the customer being aware of it.
- Denial-of-Service (DoS) attacks.
 - The ability for one customer to affect the services of one or multiple other customers
- Abuse - Tracking of end-users (IP addresses)
 - The ability to identify a customer, if there has been some incorrect usage of the services.

SEC certification can be done on IPv4, IPv6 or both protocols.

All "Implementation example" sections are for reference only. How the SEC functional requirements are implemented is outside the scope of SEC, SEC does not want to enforce or judge different types of implementations against each other as long as they give the proper end-user protection.

There are other aspects than technical and functional, when deciding on how to implement SEC. Some aspects can be related to patents, licensing, proprietary functions, and interoperability. Due to this, SEC does not want to recommend any specific implementation.

1.3 Requirement Clauses

Each clause in a requirement is described using a table.

Brief			
ID		Priority	
Description			
Motivation			
Test			

Table 1 Requirement Clause

Meaning of each field:

Brief

This is a short one-liner that briefly describes the requirement

ID

This is a unique ID in the document for the requirement that never changes. It is used to track requirements.

Priority

This describes the importance of the requirement. This is not an RFC document but the RFC2119, BCP 14 definition is used.

MUST	This word, or the terms "REQUIRED" or "SHALL", mean that the definition is an absolute requirement.
MUST NOT	This phrase, or the phrase "SHALL NOT", means that the definition is an absolute prohibition of the specification.
SHOULD	This word, or the adjective "RECOMMENDED", mean that there may exist valid reasons in particular circumstances to ignore a particular item, but the full implications MUST be understood and carefully weighed before choosing a different course.
SHOULD NOT	This phrase, or the phrase "NOT RECOMMENDED" mean that there may exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.
MAY	This word, or the adjective "OPTIONAL", means that an item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because the vendor feels that it enhances the product while another vendor may omit the same item. An implementation which does not include a particular option MUST be prepared to interoperate with another implementation which does include the option, though perhaps with reduced functionality. In the same vein an implementation which does include a particular option MUST be prepared to interoperate with another implementation which does not include the option (except, of course, for the feature the option provides.)

Description

The description should provide enough information to understand, and implement the requirement, answering the WHAT question.

Motivation

This describes the motivation of the requirement, answering the WHY question.

Test

Details on how testing of the requirement can be performed.

1.4 Reference model

There are many ways to build a broadband network, too many to describe all of them (and we will probably anyway miss many variations). To make the description of issues etc. simpler we have defined a reference model and some definitions of terms.

1.4.1 Terms

Client

Some type of equipment in a home that uses one or multiple services. This can be a Personal Computer, Set-Top-Box, Voice ATA etc.

End-user

Someone, typically a person with the help of a client consuming a service. It can for example be a person surfing Internet using a PC, or watching TV using a Set-Top-Box and a TV set.

Subscriber

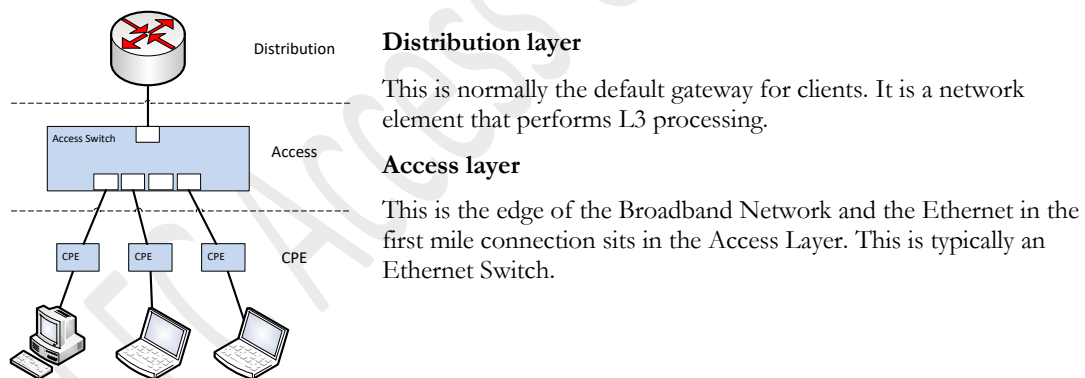
A Subscriber is responsible for the agreement and legal contracts towards a service provider. This can be the same as the end-user. Example: The Mother in a home can be the subscriber and all family members including the Mother are end-users.

Demarcation Point

The point in which a service is delivered to the client. In a broadband network this is normally a RJ45 Ethernet Connection in the wall or at the back of a CPE.

1.4.2 Model

This document describes a network which has three main layers.



CPE layer

If multiple demarcation ports are needed in the home the CPE could be part of the Broadband Network. The CPE can be “dumb” working as a L2 bridge or more intelligent, adding L3 features such as IPv4/IPv6 forwarding, stateful firewall, Network Address Translation, VoIP services etc.

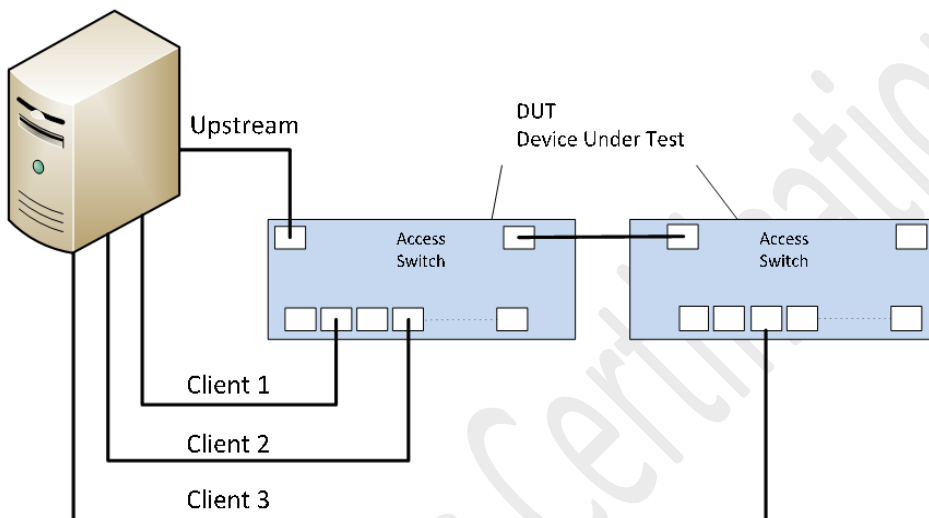
Not all networks have a CPE layer, or let the CPE be part of the home network. In most of the pictures in this document the CPE is not important so it is not included.

1.5 Test

To test the requirements in this document a test setup is needed. Most of the requirements can be tested with one Access devices, in certain cases two Access devices is needed. Due to this a static setup with two devices will be used in all tests.

1.5.1 Test setup

A test server with four gigabit Ethernet interfaces is connected to the DUTs. The test server will simulate both the upstream router/switch and three clients (end users).



All traffic on the client ports will be untagged (no VLAN). Depending on the test being performed, the traffic on the upstream link will be VLAN tagged or untagged.

For testing of a few of the requirements a DHCP server and a Multicast sender is needed. This can be accomplished by setting up a DHCP server on the Upstream interface and letting it run for the duration of the test. It is also possible to use an external DHCP server connected to the Access Switch. A multicast sender can either be the real distribution of multicast traffic in the network or it can be a sender from Upstream interface.

2 SEC - Common


This section contains a description of the various problems that can occur in a shared L2 segment. This section is protocol agnostic and is valid for both IPv4 and IPv6.

2.1 MAC Flooding

Impact: Eavesdropping, DoS

The 802.1 standard for L2 switches specifies that when the L2 forwarding table is overloaded the switch should go into a “hub” mode. In this mode all received packet on one port is flooded on all the other ports on the switch (flooded to all ports in the VLAN). It needs to do this to ensure that all traffic will reach its correct destination.


2.1.1 Normal Operation



Normal unicast traffic (green arrow) between a customer port on an access switch and the default gateway cannot be intercepted by other ports in the same VLAN.

The malicious user only sees broadcast and possibly multicast packets, which themselves does not contain too much useful information.

2.1.2 Attack – MAC flooding



The malicious user runs a special program that sends packets to the Access Switch. The destination is not that important but to avoid policing of flooding the default gateway is a good candidate.

Each packet that is sent has a different source MAC address so each packet adds one entry in the Access Switch L2 forwarding table. By sending >100K packets the malicious user fills the forwarding table to its capacity.

The malicious user then starts sniffing. Since the switch now is in “hub” mode, the sniffer will see all traffic on all ports in the Access Switch.

2.1.3 Requirement

Brief	Limit number of MACs per customer port		
ID	SEC-CM-MACLIMIT-1	Priority	Must
Description	When a customer port on the access switch has learned more than the specified number of MAC addresses, drop received packets with other MAC addresses.		
Motivation	Avoid that the switch goes into “hub” mode, so no sniffing of network traffic can occur.		
Test	Send 100,000 Ethernet frames with different random source MAC addresses from the malicious customer while the friendly customer does not send any traffic. Once that is done, send traffic from the uplink to the friendly customer and verify that the traffic is not received on other ports.		
	The requirement is verified if no traffic meant for other users are received on the malicious customer port.		

Brief	Report when the number of MACs per customer port has been reached		
ID	SEC-CM-MACLIMIT-2	Priority	Should
Description	When the number of learned MACs on a port has been reached this should be noted, for example a counter per customer port and/or log the event using syslog/snmp trap.		
Motivation	Troubleshooting aid: If this counter is increasing there are more MACs on the port than allowed and those will not get any services. Abuse: If the counter increases rapidly a DoS attack is in process that tries to fill the forwarding table in the network. This could potentially lead to the customer port being disabled.		
Test	When performing test in SEC-CM-MACLIMIT-1, verify that the limit has been reached (SNMP trap received, display counter on port, check log etc).		
	The requirement is verified if there is some notification that the limit has been reached.		

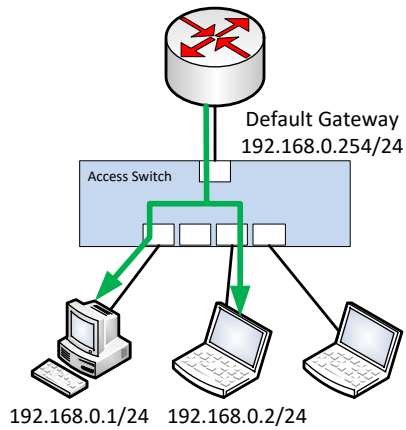
2.1.4 Implementation example

- One simple solution is to track the number of MAC addresses seen on the interface and refuse to learn more than the specified number of addresses.
- An alternative solution is to use source address spoofing filters, and the DHCP/DHCPv6 server tracks per customer port and only hands out a limited number of addresses at a time.

2.2 Multicast injection


Impact: DoS

2.2.1 Normal Operation



The multicast stream received from the upstream router is replicated in the Access Switch and sent to two customer ports.

2.2.2 Attack – Multicast injection



A malicious user sends multicast data (red line) into the access port on the same multicast group as the green line.

This can seriously affect the video and/or audio quality for other customers.

2.2.3 Requirement

Brief	Filter unwanted sending of IPv4 and IPv6 multicast from customer ports.		
ID	SEC-CM-MCAST-1	Priority	Must
Description	Filter out unwanted IPv4 and IPv6 multicast packets when they arrive at the customer access port.		
Motivation	If clients should be able to send multicast into the network, the groups being used need to be filtered so client initiated multicast cannot be sent on the same groups as other multicast services uses.		
Test	Send a multicast stream into one customer access port. At another customer port and the uplink port, join the same multicast group, and measure whether this multicast stream is received at any of those ports. The requirement is verified if no multicast packets are received at the other customer access port or uplink port.		

Brief	Select what multicast groups can be received on customer ports		
ID	SEC-CM-MCAST-2	Priority	Must (only relevant if multicast is used in the network)
Description	Customers should only be able to receive specified multicast channels.		
Motivation	Prevent users from joining and receive unused multicast channels.		
Test	Join two multicast groups, one allowed and one not allowed, and measure whether they can be received or not. The requirement is verified if the valid multicast group can be joined but not the multicast group that is not allowed.		

2.2.4 Implementation example

Filtering can be done on L2 or L3.

If filtering is done on L2 each customer port is configured with a filter, which drops all packets with the multicast bit set in the Destination MAC Address.

If filtering is done on L3 an Access Control List is used that blocks all multicast destination addresses (224.0.0.0-239.255.255.255).


Note that if IPv6 is used certain multicast traffic cannot blindly be dropped. Many IPv6 control protocols are dependent on multicast, for example Router Advertisement, Router Solicitation, Neighbor Solicitation etc.

2.3 Spanning-Tree

Impact: DoS


By injecting spanning-tree packet into a customer port, the topology of the spanning-tree can be modified. By injecting large amounts of BPDUs the CPU of the switches will be overwhelmed and the switch will not perform reliably.

2.3.1 Normal Operation



No BPDUs is received at the customer ports. Traffic flows between client and default gateway.

2.3.2 Attack – Spanning Tree



The malicious user sends fake BDPUs packets to Access Switch. The effect can be somewhat different depending on the amount and content of the packets. One possibility is sending so many random packets that the CPU of the access switch will be overloaded and might cause links to go down.

Another possibility is for the malicious user to get traffic sent by him from some other customers. This can be done by adding himself as a valid bridge and in some cases this can cause MITM attacks.

There are several types of attacks that can be done with BPDUs, some of these are

- Sending RAW Configuration BPDU
- Sending RAW TCN BPDU
- Denial of Service (DoS) sending RAW Configuration BPDU
- DoS Sending RAW TCN BPDU
- Claiming Root Role
- Claiming Other Role
- Claiming Root Role Dual-Home (MITM)

One tool for creating the above attacks is Yersinia.

2.3.3 Requirement

Brief	Disable spanning-tree on customer ports		
ID	SEC-CM-SPT-1	Priority	Must
Description	Spanning-tree packets should not be sent on customer ports and any received spanning-tree packets should be silently discarded. The Access Switch should not participate in spanning tree protocol on any customer port. It is not an error if Access Switch receives Spanning Tree packets on the customer port; customer may have a spanning-tree capable L2 switch in the home/as CPE.		
Motivation	Avoid that customers adds themselves in the path creating DoS or MITM attacks.		
Test	Listen with a sniffer if there are any Spanning Tree packets (BPDUs) transmitted on a customer port. Send 10 spanning-tree BPDUs into the switch and measure if the switch responds. The requirement is verified if the switch does not send any STP BPDUs to the customer access ports		

2.3.4 Implementation example


Each customer port is configured to not participate in the Spanning-Tree protocol, no BPDUs are ever sent out on these ports.

A filter is installed on each customer port that silently drops any received BPDUs

2.4 Other protocols


Impact: MITM, DoS

2.4.1 Normal Operation



Broadband Network is secured according to SEC for IPv4. End-users do not use any other protocols than IPv4 and ARP.

2.4.2 Attack – Other Protocols



If the network does not block other protocols such as IPv6 end-users could communicate directly with each other and create MITM/DoS attacks, as is described in section 4 “

There are more protocols that could be used for unauthorized access to other customer's assets than IPv4 and IPv6. Some examples are IPX, AppletalkAppleTalk and proprietary L2 protocols.

IPv6 is enabled by default in most modern operating system (Windows, Linux, OSX) so it is very important to stop IPv6 if that protocol is not used in the Broadband Network.

The old Ethernet 802.3 uses the EtherType field as length of the packets as opposed to Ethernet II. The packets are usually sent to the correct destination but some packet filters might not work correctly (MAC access lists on Cisco devices for example will not work correctly).

2.4.3 Requirement

Brief	Disable unsupported protocols on customer ports.		
ID	SEC-CM-OTHER-1	Priority	Must
Description	Any unsupported protocols in the Broadband Network must be disabled. All packets must be silently dropped when received.		
Motivation	Avoid non-supported protocols MITM, DoS		
Test	<p>Send the following protocols from customer port to uplink:</p> <ul style="list-style-type: none"> • Frame Relay ARP • Raw Frame Relay • AppleTalk and AppleTalk ARP • DEC LANBridge • YyyMLPS • Ethertype < 1500 • IPX (two different EtherTypes) <p>Measure if the protocols are received at the uplink.</p> <p>The requirement is verified if all protocols are filtered by the switch/network.</p> <p>Depending on if the network is IPv6 enabled, IPv6 packets should be either dropped or allowed.</p>		

2.4.4 Implementation example

On each customer port, install a filter that looks at the EtherType.

If IPv4 is in use, accept packets with EtherType 0x0800 (IPv4) and 0x0806 (ARP).

If IPv6 is in use, accept packets with EtherType 0x86dd.


All packets with other EtherType should be silently dropped, unless needed by the control plane.

2.5 Control Plane Protection

RFC6192 “Protecting the Router Control Plane” has a strong router perspective. A subset of this RFC can be applied on an Access Switch running in L2 bridged mode.

An Access Switch (and most other modern networking equipment) can be broken down into two major parts. One is doing packet forwarding and one is handling the control plane. These two parts are interconnected. If the packet forwarding realizes that the packets should go to the control plane the traffic is either copied or redirected there.

A typical model of this is illustrated here. Packet forwarding is done by the ASIC and the CPU handles the control plane. Again, actual implementation can differ but the model/concept is normally the same.




The bandwidth from the ASIC to the CPU and the CPU itself has limited capacity compared to the ASIC. The ASIC can normally handle any type of load at line rate, independent of the packet size.

2.5.1 Flooding the control plane

Impact: DoS

The switch is normally configured with different VLANs for different purposes. One VLAN (typically VLAN 1, which is the default VLAN) handles the management traffic. This VLAN is used by the Network Operator to configure, troubleshoot and retrieve statistics from the Access Switch. The management VLAN is normally not accessible from customer ports for security reasons. Other VLANs are used to deliver services to customer ports.

There are some types of packets/traffic that the CPU needs to handle. The ASIC is configured to match these packets and sends them to the CPU. When doing this a large window is opened for DoS.



An example, in a customer home a patch cable is incorrectly connected back to back in the CPE.

The client laptop sends a DHCPv4 Discover, which is a broadcast. The DHCPv4 packet is flooded in the L2 domain and sent to the CPU since it is doing DHCPv4 snooping (green line).

The flooding goes out the port with the patch cable and is directly received and sent out again creating a loop. The available bandwidth is quickly saturated, if Fast Ethernet is used this sends 148800 packets per second to the CPU (red line) and all other nodes in the L2 broadcast domain.

Few CPUs in an Access Switch can handle that kind of load.

The CPU due to being overloaded will drop other legit traffic from the other ports, and can severely affect other customer port services. For example if IGMP snooping is active for IPTV and those packets are dropped the other customer ports cannot watch IPTV and zapping between channels becomes impossible.

There can be many reasons traffic is sent to the CPU, a malicious user can run a program that sends huge amount of traffic that the control plane is interested in.

Some of the traffic that the CPU handles are:

- BPDUs (Spanning-Tree, 802.1x, LLDP etc.)
- IGMP/MLD, for snooping
- DHCPv4/DHCPv6, for snooping
- ARP

2.5.1.1 Requirement

Brief	Protect the control plane from excessive traffic		
ID	SEC-CM-CP-1	Priority	Must
Description	When packets are sent from customer ports to the control plane, they must be limited (shaped, policed etc) so the CPU is not overloaded. The limit should be per customer port.		
Motivation	Avoid that a single user uses all resources at the control plane. Effectively creating DoS.		
Test	<p>Enable DHCP snooping with option82 insertion on Customer Access Port 1 and 2</p> <p>A malicious customer injects more than 5000 DHCPv4 DISCOVER per second.</p> <p>A friendly customer perform a DHCPv4 request.</p> <p>The requirement is verified if the friendly customer successfully retrieves an IPv4 address from the DHCPv4 server.</p>		

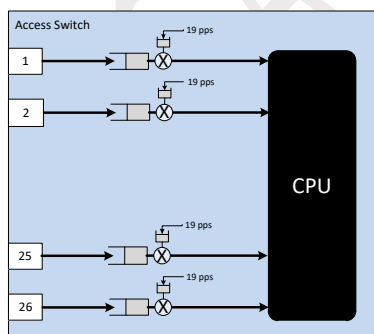
2.5.1.2 Implementation example

Here two implementations are described. To simplify the pictures the Switch ASIC is not drawn.

The Control Plane is usually limited by the number of packets per second (pps), bandwidth has less implication. Of course both could be used when deciding on dropping packets.

Policing

Here a policer, one per port in the Access Switch is used. If the pps are above the limiter, packets are dropped. This usually works well but configuring the limiter can be hard. One question directly arises, how many pps can the CPU handle and to what value should the policers be configured?



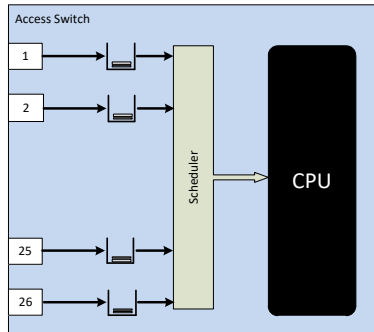
An example: Assume that the CPU can handle 500 pps. In a 26-port Access Switch each ports policer for traffic to the control plane would then be set to $500/26 = 19$ pps.

A scenario in which all ports will have problems at the same time is very unlikely so the control plane could be oversubscribed, for example 50 pps per port.

The capacity of the control plane will vary over time depending on what it does (spanning-tree, 802.1x etc), and can change between software versions, so the pps policer method is a rough one. It needs to be monitored and possibly adjusted over time.

Queuing

Another approach is to implement a queue for each port that is used when sending packets to the control plane. That means that instead of protection the Control Plane with a fixed value per port the Control Plane asks for more work when it is ready.



A scheduler gives fair priority between the queues. When the Control Plane is ready for the next packet it fetches one. This makes the processing of control plane packets much more efficient and self-adjustable.

If the Control Plane is busy with other stuff it retrieves less pps but still gives the different queues (ports) a fair share of the CPU capability.


If a port sends too much traffic to the control plane the queue for that port will fill up and start tail-dropping packets. No other queue/port will be affected by this.

The work conserving nature makes the capacity of the Control Plane evenly spread over the queues that have packets in them. If just two ports send packets to the control plane, they will share the available processing power, each will get $500 \text{ pps} / 2 = 250 \text{ pps}$, compared to the policing method in which they never get more than the policer max value (19 or 50 pps as in previous example).

2.5.2 Packet Buffering

Impact: DoS

To avoid buffer starvation, each port in the Access Switch must have a minimum amount of buffers available, independent of the amount and type of traffic other ports handles.



If there is more traffic received than the egress port can handle, a queue is built up as illustrated by the red line. Traffic received on a gigabit Ethernet port exiting a Fast Ethernet port can create such a scenario.

If all buffers will be assigned to the red egress port, other traffic being received (green line) will be dropped due to the lack of buffers.

2.5.2.1 Requirement

Brief	Avoid starvation of packet buffers due to excessive traffic		
ID	SEC-CM-PB-1	Priority	Must
Description	Make sure that customers use separate buffers.		
Motivation	Avoid that filled buffers cause dropped packets on multiple ports.		
Test	From uplink port send twice the link speed to one of the customer ports for a duration of 60 seconds. At the same time, send 80% of the link speed from the uplink to another customer port. Measure packet loss on both customer ports. The test is verified if no packet loss occurs on the customer port where 80% of the link speed is sent and about 50% packet loss occurs on the overloaded customer port.		

2.5.3 Fragmentation of packets

Impact: DoS

If IPv4 and IPv6 packets that are sent to the control plane are fragmented, the CPU/control plane needs to reassemble those before processing them. This can consume packet buffers and CPU cycles, potentially creating a DoS, when CPU runs out of steam or the packet buffers are full.

The control plane is normally in a controlled environment, so the MTU is known. It is therefore no reason that packets should be fragmented and no reassembly should be needed.

Brief	Drop fragmented IPv4 and IPv6 packets before reaching the control plane.		
ID	SEC-V4-CP-FRAG-1	Priority	Must
Description	When packets are being redirected to the control plane, check if they are fragmented and if so just drop them.		
Motivation	Avoid DoS from fragmented packets to the control plane.		
Test	On a customer port, activate DHCPv4 snooping with option-82 insertion. Send a standard DHCPv4 request into the customer port, verify that the packet is correctly tagged with option-82 on uplink port. Send a DHCPv4 request fragmented to 40 bytes and one fragmented to 104 bytes from the start of the UDP packet header into the customer port. Measure on the uplink if the fragmented packets are received at the uplink port.		

Brief	Drop fragmented IPv4 and IPv6 packets before reaching the control plane.		
ID	SEC-V4-CP-FRAG-1	Priority	Must
	The requirement is verified if the fragmented packet is not received at the uplink port.		

SEC Access Certification

2.6 Network Protection

2.6.1 Broadcast, Multicast, Unicast Flooding limiter

Impact: DoS

Protect against a loop on L2. For example a customer can incorrectly connected an RJ45 patch cable back to back in two CPE ports. This generates linerate traffic into the customer port at the access switch.

Note that valid multicast traffic identified with IGMP/MLD snooping should not be included in the Multicast flooding limiter.

2.6.1.1 Requirement

Brief	Limit broadcast, multicast and unicast flooding traffic, per port		
ID	SEC-CM-NP-1	Priority	Must
Description			
Motivation	Avoid DoS		
Test	<p>Activate the limit for broadcast, multicast and unicast flooding traffic on customer port 1, for example set the limit to 100 pps or 500 Kbit/s.</p> <p>Inject ARP packets (broadcast) on customer port 1, twice the pps limit (for example 200 pps). Verify that no more than the pps limit is forwarded out on uplink port.</p> <p>Inject multicast packets corresponding to destination group 224.1.1.1 on customer port 1, for example set the limit to 10000 per second. Verify that no more than 100pps or 500 Kbit/s is forwarded out on uplink port.</p> <p>Inject IPv4 packets with a destination MAC address of C8:0A:A9:A8:D9:04. It is important that the destination MAC is unknown so no host will reply with this address, forcing the Access Switch to flood the packets. Set the limit to for example 100pps or 500 Kbit/s. Send twice the limit from the customer port to the uplink. Measure the pps or bandwidth received at the uplink.</p> <p>The requirement is verified when no more than the limit is forwarded on the uplink port, for broadcast, multicast and unicast flooded packets.</p>		

2.6.1.2 Implementation example

Typically three policers per port are used to measure the amount of broadcast, multicast and unknown flooding traffic. If the limit is reached the packets are dropped.

2.6.2 Loop-Detect

Impact: DoS

Protect against a loop on L2. Loops generates linerate traffic and can severely affect the control plane in one or multiple systems in the L2 broadcast domain.

There are several loops that need to be detected

- Loop in the customer home/CPE
 - Customer has bridged two ports together in the home, creating a forwarding loop.
- Loop between two ports in the same Access Switch
 - Two customers has connected their networks to each other with a cable
- Loop between two ports in different Access Switches
 - Two customers has connected their networks to each other with a cable

2.6.2.1 Requirement

Brief	Detect loops on customer ports, if loop detected disable the port.		
ID	SEC-CM-LD-1	Priority	Must
Description	Disable customer ports that contains loops.		
Motivation	Protect the network from loops in customer equipment.		
Test	<p>Start sending traffic from two customer ports towards the uplink. Make sure the traffic goes through from both customers.</p> <p>Create a loop on one of the customer ports and measure received traffic at the uplink.</p> <p>The requirements are verified if the access switch detects the loop, closes down the customer port that has a loop, and thereby that the uplink port stops receiving traffic from that customer port while still receiving traffic from the other customer port without a loop.</p> <p>Another example of how to do this is: On Client port 1, connect a CPE with at least three ports. CPE port 1 goes to Client port 1. Make sure all ports are in the same VLAN, and no broadcast, multicast and/or unicast flooding limits is configured. Connect a RJ45 cable between CPE port 2 and CPE port 3, effectively creating a loop. A good tip is to have broadcast, multicast and unicast flooding limits in place in the Access Switch, so the test server is protected from all the traffic the loop generates.</p> <p>The requirement is verified if the Access Switch detects the loop.</p>		

Brief	If a loop is detected, log the event		
ID	SEC-CM-LD-2	Priority	Should
Description	To indicate to the Network Owner that a customer port is disabled, this should be reported using for example syslog and/or SNMP trap.		
Motivation	Proactive troubleshooting		
Test	The requirement is verified when SEC-CM-LD1 is detected and this is logged (syslog, SNMP trap etc.)		

2.6.2.2 Implementation example

Since spanning-tree cannot be used on customer ports another method needs to be used. The author of this document does not know any standardized way to implement the loop-detect, which creates interoperability problems if multiple proprietary solutions are in use in the L2 broadcast domain.

One solution is to periodically send unicast frames (probes) with a non-existing destination MAC address. These frames should be flooded by all L2 switches, and if such probe packets are received on a customer

port a loop has been detected. One such frame that could be used is the Ethernet Loopback (Ethertype 0x9000) as originally specified in EthernetII.

2.7 Network Equipment Management

Impact: MITM, DoS, Abuse

2.7.1 Normal operation

An administrator can access network equipment over SSH or HTTPS from a special management VLAN. Other relevant protocols are ICMP Echo, Telnet, HTTP, SNMP and NTP. They are all used for configuration of some kind (except for ICMP) and all might be used to overload a switch.

2.7.2 Attack Equipment Management

A malicious customer connects to the network equipment and manages to gain access due to weak password or similar. He can also overload the switches control plane by creating multiple requests for connection.

2.7.3 Requirement

Brief	Prevent users from communicating with equipment management.		
ID	SEC-CM-NEM-1	Priority	Must
Description	Network equipment must ignore incoming management traffic from customer ports.		
Motivation	Prevent DoS by control plane overload and configuration altering.		
Test	Connect to a management IP address from customer ports using TCP connection for Telnet, SSH, HTTP and HTTPS. Send ICMP Echo request to management IP (Ping). Send a SNMP get request to the management IP. The test is verified if no answer is received for any of the tests above.		

2.7.4 Implementation example

Filter management traffic from customer ports to the switch using an access list.

Configure management not to be available from other than specified ports/IP addresses/VLANS.

2.8 Further Study

2.8.1 GARP/GVRP / GMRP

Impact: MITM, DoS

Filter out from customers

Never transmit to customer port.

If an Access Switch is listening on these packets a malicious user can force the customer port to join other VLANs – for example the management VLAN.

2.8.2 MRP, MMRP, MVRP

Impact: MITM, DoS

The GARP/GVRP/GMRP replacement, which is compatible with MSTP, as defined in 802.1ak.

There are the same issues and solution as GARP / GVRP / GMRP.

2.8.3 LLDP

Impact: DoS

Filter out from customers

Never transmit to customer port.

2.8.4 CDP

Impact: DoS

Filter out from customers

Never transmit to customer port.

2.8.5 ISL

Impact: DoS

ISL is a Cisco proprietary protocol that is replaced by 802.1Q VLANs.

Make sure it is disabled on customer ports.

2.8.6 802.3ad

Impact: DoS

Link Aggregation (LAG) and Link Aggregation Control Protocol (LACP).

Never use this towards customer, don't accept packets and don't send them on a customer port.

3 SEC – IPv4

This section contains a description of the various problems that can occur in a shared L2 segment using the IPv4 protocol.

3.1 IPv4 overview

To better understand the issues, we begin with some generic information.

3.1.1 Host Address Configuration

A host that wants to use IPv4 on the global internet needs an IPv4 unicast address. There are several methods the host can use to retrieve the address

- Static configuration
- BOOTP (RFC1531), not that common anymore, DHCPv4 is built on BOOTP
- Dynamic Host Configuration Protocol – DHCPv4 (RFC2131)
- Dynamic Configuration of IPv4 Link-Local Addresses (RFC3927)

Static configuration

- Static configuration is not very practical. IPv4 addresses are hard to remember. There are also additional addresses, such as DNS servers, NTP servers etc. that normally needs to be configured. When a host moves it is easy to forget to reconfigure it.

BOOTP

- BOOTP is not that common today, it has some drawbacks, and for example it cannot reclaim an assigned address in a well-behaved way.

DHCPv4

- DHCPv4 is supported and enabled by default by most mainstream operating systems and DHCPv4 is the most common way to assign addresses to hosts today.

DHCPv4 Link Local Addresses

- IPv4 Link-Local Addresses (169.254/16) is normally used when no connection exist to another network so it is not really applicable in a Broadband Network.

Brief	Log all assigned dynamic addresses		
ID	SEC-V4-HAS-1	Priority	Must
Description	All assignment of IPv4 addresses must be logged.		
Motivation	Used when troubleshooting, abuse cases, IPRED etc.		
Test			

3.1.2 IPv4 Address configuration

Since the IPv4 addresses are a scarce resource it is important to avoid wasting them. This is one of the major reasons to put several customers into the same VLAN/L2 Broadcast Domain, fewer addresses will be consumed.

There are techniques to conserve address space in networks with separate VLAN/L2 broadcast domains per customer port. Some of these techniques are called proxy-ARP, unnumbered interfaces and host routing. Describing those is out of the scope for this document.

Example 1, separate broadcast/L2 domain per customer:

24 port Access Switch, six switches, five usable IPv4 addresses on each customer port. Separate VLAN/L2 Broadcast Domain on each customer port.

Each port needs eight IPv4 addresses. Network, default gateway, five hosts and broadcast.

24 customer ports * 8 IPv4 addresses per port * 6 switches

=> A total of 1152 IPv4 addresses. A /21 prefix is needed, which has 2048 IPv4 addresses.

Utilization: $1152 / 2048 = 56\%$

Example 2, shared broadcast domain for all customers:

24 customer ports * 5 usable IPv4 addresses per customer port * 6 switches.

1 IPv4 address for default gateway

2 IPv4 addresses for network and broadcast

=>A total of 720 IPv4 addresses. A /22 prefix could be used for a total of 1024 addresses.

Utilization: $720 / 1024 = 70\%$

3.1.3 Customer identification

There are several reasons to always identify subscribers and end-users. It describes what services should be delivered, and avoid abuse issues. Without some type of identification it is not possible to distinguish customers. This identification can be done by using the DHCP option 82 that is available in most access equipment. The feature usually adds option 82 to DHCP packets sent to the DHCP server and provides information such as the identification of the switch, subscriber id and port. The exact information differs between vendors.

Requirements SEC-V4-DHCP-1 and SEC-V4-DHCP-2 is only valid if DHCPv4 is used. There exist other methods to identify customers such as 802.1x and web login. For such solutions it should be verified that the identification cannot be spoofed.

Brief	DHCPv4 requests tagged with customer port identification		
ID	SEC-V4-DHCP-1	Priority	Must
Description	Traceability must be added to the DHCP-requests in "L2-mode", typically option 82 circuit-id, remote-id, subscriber-id.		
Motivation	Assist in identifying users in abuse cases.		
Test	Send a DHCP discover from Customer and verify that Uplink receives it correctly marked with option-82.		

Brief	Drop packet or replace customer specified port identification if present in DHCPv4 requests		
ID	SEC-V4-DHCP-2	Priority	Must
Description	When a DHCPv4 packet is received on a customer port that already contains customer identification, either drop the IP packet, or overwrite with the correct values.		
Motivation	Avoid customers faking wrong customer identification, which would prevent traceability.		
Test	Send a DHCP discover from Customer without option-82 specified. Verify that the discover is received at the Uplink. Then send a DHCP discover with option-82 set by the customer. At uplink verify that this packet is either dropped, or that this marking is replaced with the correct one.		


SEC Access Certification

3.2 ARP

Impact: MITM, DoS

3.2.1 Normal operation

Here is a typical configuration. The PCs and Laptops are in the same broadcast L2 domain (same VLAN). The Router at the top is the default Gateway – it is the path to the rest of the Internet.



When a client PC wants to send packets to Internet it first needs to find out the Ethernet MAC address of the default gateway.

The client PC broadcasts an ARP request that contains the question “who has IPv4 address 192.168.0.254”. Since it is a broadcast all devices in the L2 domain receives the query.


The router that has the IPv4 address 192.168.0.254 answers with an ARP Reply (green line). The ARP reply contains “My MAC address is [whatever it is]”.

From now on, the client knows how to send packets directly to the default gateway using the routers MAC address.

The ARP protocol has no built-in security. Most computers, routers and switches update its ARP table even if they have not sent any ARP Requests. This lack of security is a problem.

3.2.2 Attack - ARP Poisoning

This is also sometimes called “ARP Cache Poisoning” and “ARP spoofing”.



A malicious user (red laptop) starts with sending to ARP replies. The first ARP reply is sent to the router with the client PCs IPv4 address and its own MAC address. This makes the router believe that the malicious user is the client PC


The second ARP reply is sent to the client PC with the default gateway IPv4 address and its own MAC address. This makes the client PC believe that the laptop is the default router.

Depending on what the malicious user does a MITM or DoS attack can be created.

If the malicious user sends its own MAC address as the default gateway, the PC will send all traffic to the malicious users computer, which can eavesdrop and send the traffic to the real default gateway. The PC is unaware of this since the Internet connection works as usual.

The malicious user computer could send an ARP reply with an unknown MAC address. The client PC would then try to use this to communicate with Internet but that would fail since nothing will receive and handle those packets.

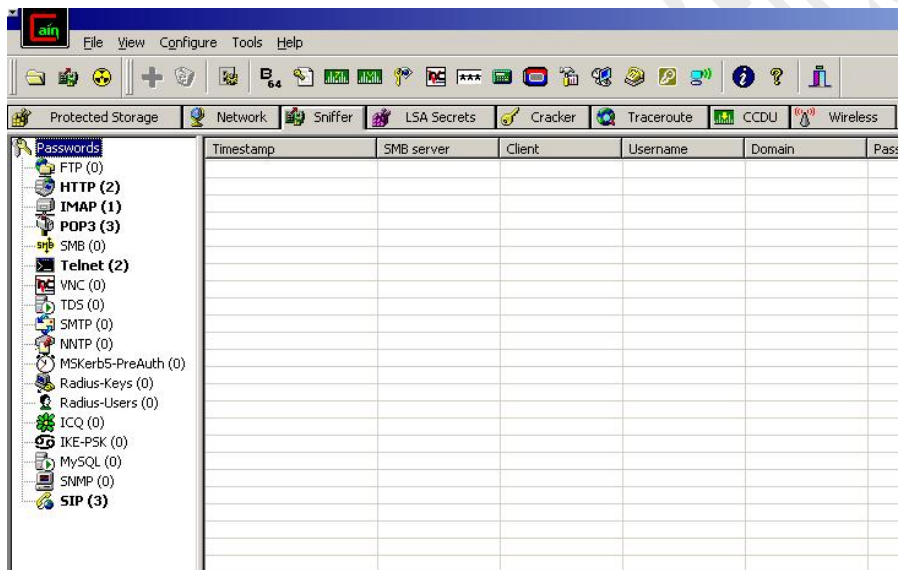
Hacker PC can now eavesdrop all logins (Facebook, email etc.), voice over IP phone calls, redirect traffic so when you go to one web site you end up on another fake web site and more. There is software available for this, which are very simple to use. See for example www.oxid.it.



There are other L2 link layers that may have the same ARP poisoning problem, e.g. some configurations of wireless hotspots.

Here the malicious user can do the exact same ARP Poisoning attack as described above.

Here is an example of tests done with the tool “Cain” in a broadband network. After just a couple of seconds traffic capturing you can see sniffed logins for mail, ftp and IP telephony.



3.2.3 Requirement

Brief	Filter out ARP Poisoning packets		
ID	SEC-V4-ARP-1	Priority	Must
Description	Access Switch must inspect all received ARP packets and only allow those that has a correct source MAC in the Ethernet header and Source MAC/Source IP in the ARP payload		
Motivation	Avoid ARP poisoning, which can lead to Man In The Middle and Denial of Service attacks.		
Test	<p>From the uplink, send a continuous UDP stream to a friendly customer.</p> <p>From a malicious customer send spoofed ARP packets containing the friendly customer source IP address to the uplink using broadcast and unicast, ARP request and ARP response.</p> <p>Verify:</p> <ol style="list-style-type: none"> 1. That the friendly customer does not receive any spoofed ARP packets. 2. That the malicious customer does not receive the continuous UDP stream from the uplink. <p>If the network is built with a ring structure the test need to consider this.</p>		

3.2.4 Implementation example


By intercepting each ARP packet being received on a customer port and checking the source MAC address and the source IPv4 address both in the packet and the ARP payload, incorrect ARP packets can be screened and dropped.

3.3 ICMPv4

Impact: MITM, DoS

ICMPv4 is the IPv4 Control Protocol. There are several issues with the protocol in a L2 shared broadcast domain.

3.3.1 Normal operation – ICMPv4 Redirect




Redirect informs a client of a better way to reach the destination. It is used to optimize traffic.

In a broadband network there is normally only one way to the rest of the network so ICMPv4 redirect messages is not that useful.

Client communicates with the rest of the network through the default gateway (green line). When sending packets to 172..16.1.0/24 the default gateway will send an ICMPv4 redirect back to the client (yellow line), asking it to use 192.168.0.253 instead for that destination.

This removes the packet forwarding task from the default router, the client sends packets in a more direct way to that destination (dashed green line).

3.3.2 Attack –ICMPv4 Redirect



A malicious user can send an incorrect redirect message to a client (red line), informing the client of a better path to the default gateway, using itself as the redirect target.

This effectively creates an MITM attack, in which all traffic from the client to the default gateway passes through the malicious user (green line).

If an unknown nexthop is used instead of the malicious user a DoS attack is created.

3.3.3 Requirement

Brief	Filter out ICMP redirect packets.		
ID	SEC-V4-REDIR-1	Priority	Must
Description	Access Switch must inspect all received ICMP redirect packets and only allow those that are valid for that customer port. It is very seldom an ICMP redirect is needed in a Broadband Network, so they could potentially be dropped.		
Motivation	Avoid that malicious user redirects traffic using ICMP.		
Test	Send an ICMP packet with redirect message from malicious customer to a friendly customer. Verify that the packet does not arrive at the friendly customer.		

Brief	Pass on valid ICMP traffic		
ID	SEC-V4-REDIR-2	Priority	Must
Description	ICMP is the IPv4 control protocol and is used for several tasks. It is important that the Access Switch do not blindly drop ICMP packets, since this would break other protocols, such as Path MTU discovery, traceroute, ping etc.		
Motivation	Don't limit possibilities to troubleshoot the network.		
Test	Send an ICMP Echo (ping) from a customer to the uplink. Verify that a reply is received.		

3.3.4 Implementation example


- Use an ACL to filter out incorrect ICMP packets
- Use the section 3.8 “Forced Forwarding” feature so clients cannot communicate with each other over L2.

3.4 IPv4 Source Address Spoofing

Impact: DoS, Abuse


Source Address Spoofing is when a program or a person masquerades as another by falsifying the source IPv4 address. A typical attack is when someone does some sort of crime, and later when the source of the offender is searched for (Abuse, logs) points at the incorrect subscriber.

3.4.1 Normal Operation



Each host uses a unique and approved IPv4 source address when communicating over the broadband network.

3.4.2 Attack – IPv4 source address spoofing



The malicious user uses the same source IPv4 address as the desktop when sending packets to the network, effectively impersonating the desktop.

The malicious user could also just grab a “free” IPv4 address in the subnet being used and start to communicate with the rest of the network. If a crime is being done and the offender is searched for, no record will exist that describes where this “free” IPv4 address exist in the network.

3.4.3 Requirement

Brief	Filter out spoofed IPv4 packets in the access node.		
ID	SEC-V4-SPOOF-1	Priority	Must
Description	When a packet is received on the access port, it needs to be verified that the source IPv4 address is legit and if it not the packet must be dropped. How the address is verified is an implementation detail but it must work with all forms of IPv4 address configuration seen in section “3.1.2 IPv4 Address ”		
Motivation	Mitigate any chance of IPv4 Source Address Spoofing. Avoid Abuse cases.		
Test	Send valid IPv4 packets from a customer to the uplink. Verify that the packets are received. Send IPv4 packets using spoofed source IP and verify that they do not arrive at the uplink.		

Brief	Track when spoofed IPv4 packets are received		
ID	SEC-V4-SPOOF-2	Priority	Should
Description	When a spoofed packet is received it should be notified in a counter (per customer port) or a log.		
Motivation	Used for troubleshooting. If the spoofed packet counter is incremented this indicates that packets with incorrect IPv4 source addresses are received.		
Test			

3.4.4 Implementation example

By using the IPv4 whitelist table an access filter can be applied on customer ports. If a packet with an IPv4 source address is received that has a source address that doesn't exist on the specific customer port in the whitelist table, the packet is dropped.

3.5 IGMP

Impact: Avoid sending multicast traffic to uninterested ports, DoS

3.5.1 Normal Operation

According to the IEEE 802.3 standards, packets that have the multicast bit set in the destination address must be flooded in the L2 broadcast domain. This means that if one client in the L2 domain watches a multicast IPTV channel, all other customer ports in the L2 domain also gets a copy of the IPTV channel.

This has several issues:

- Wastes bandwidth
- DoS when the amount of multicast traffic is more than available bandwidth on the Customer Ports.

By looking at the IGMP membership reports/queries, the Access Switch can learn what IPTV channels each customer port needs, and only copy packets to those.

3.5.2 Requirement

Brief	Do not use IGMPv1		
ID	SEC-V4-IGMP-1	Priority	Should
Description	IGMPv1 is an old protocol and lack the proper function to signal that a receiver is no longer interested of a multicast source. It will time out approx. 120 seconds later when no one is asking for it. This creates huge bandwidth requirements if a user zaps quickly through a bunch of IPTV channels.		
Motivation	Avoid inefficient usage of available bandwidth		
Test			

Brief	Use the IGMPv2/IGMPv3 messages to learn about interesting multicast receivers.		
ID	SEC-V4-IGMP-2	Priority	Should
Description	By monitoring IGMPv2/IGMPv3 packets, the switch can learn what ports are interested in a certain (*,G) or (S,G) multicast group. This can then control the packet forwarding/replication.		
Motivation	Avoid inefficient usage of available bandwidth		
Test			

3.5.3 Implementation example

There has been several implementations of IGMP, with various problems, some of them were:

- IPv6 did not work (IPv6 ND and RA multicast messages did not get through)
- RIP, OSPF and other IPv4 protocols that use IPv4 multicast did not work

Due to the lack of a standard, IETF released RFC4541 [Link](#) – “Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches”, which describes how a correct implementation should be done.

3.6 Multicast group overload

Impact: DoS

A client / STB can ask for any number of multicast groups and have all of them at the same time. This can exceed the available bandwidth in the network and on the customer port. It can also consume all available entries in the multicast forwarding table, effectively creating a DoS for multicast on other ports.

3.6.1 Normal Operation

Each multicast receiver is joining one or a few concurrent multicast channels.

3.6.2 Attack – Multicast Group Overload

A malfunctioning device or malicious user sends joins using IGMP protocol for all possible multicast groups there are. This is a huge number and no existing router/switch can track that many multicast groups.

When another customer port wants to change or join a multicast channel (for example watching a TV channel) there is no way the Access Switch can distribute this new channel to the port, effectively creating a DoS attack.

3.6.3 Requirement

Brief	Limit number of concurrent multicast groups a customer port can join using the IGMP protocol		
ID	SEC-V4-MCAST-1	Priority	Should
Description	Assure that users can receive the wanted multicast channels without problems.		
Motivation	Avoid that users join too many multicast channels effectively creating DoS.		
Test	Join existing multicast groups in the network one by one and verify that the number of concurrently received multicast groups never exceeds the specified limit. Note: Requires a multicast sender in the network.		

3.6.4 Implementation example

Count the number of active multicast groups on each customer port. If the limit is reached, refuse to add the requested group to the customer port.

3.7 Protocol protection

Impact: MITM, DoS, Abuse


A number of protocols are used by the Network Owner for running and maintaining the broadband network and its customers. Dynamic Host Configuration Protocol (DHCPv4) is one of these protocols. If a malicious user could inject DHCPv4 responses it would affect other customer's services.

3.7.1 DHCPv4 Server Spoofing

Impact: MITM, DoS

This can be triggered both intentionally by a malicious user and unintentionally (for example misconfiguration, enabling Windows Connection Sharing).

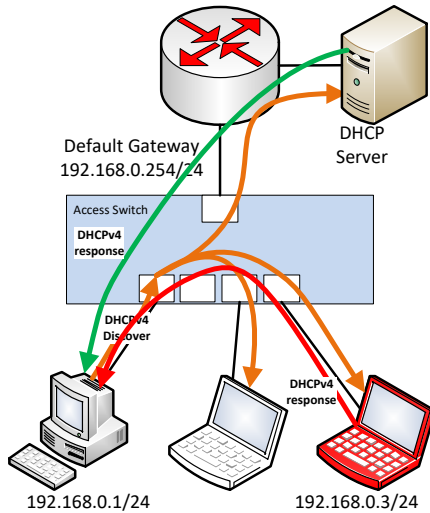
3.7.1.1 Normal Operation



Client asks networks for an IPv4 address by broadcast a DHCP discover. This is responded to by the network, eventually ending up with the client having an IPv4 address.

SEC Access Certification

3.7.1.2 Attack – DHCP Server Spoofing



When a DHCP server spoofing attack is in place the malicious user sees the DHCPv4 discover packet (it is a broadcast one and will be transmitted to all ports) and responds to the discovery in the same way as the real DHCPv4 server does.

The malicious user responds with default gateway 192.168.0.3 (itself) so all traffic from the desktop to the rest of the network is passed through the malicious user, creating a MITM attack

If an invalid default gateway is returned, the desktop will not be able to communicate with the rest of the network, creating a DoS attack.

3.7.1.3 Requirement

Brief	Filter out all DHCP server traffic coming from customer ports
ID	SEC-V4-DHCP-1 Priority Must
Description	So-called DHCP-snooping to filter out DHCP messages between subscribers must be located on the access ports. I.e. a subscriber must not be able to act like a DHCP-server for another subscriber.
Motivation	Avoid DHCPv4 server spoofing attacks. Avoid abuse cases.
Test	From a friendly customer send a DHCP discover and verify that no answer is received at the malicious customer.

Brief	Track invalid DHCPv4 server packets
ID	SEC-V4-DHCP-2 Priority Should
Description	When an invalid DHCPv4 response packet is received on a customer port, this should be counted and/or logged.
Motivation	Simplify troubleshooting. If the invalid counter is increasing this indicates a local DHCPv4 server on that customer port. This will probably interfere with the service providers DHCPv4 responses, affecting services.
Test	

3.7.1.4 Implementation example

- Use an Access Control List on customer ports to filter and drop “DHCPv4 Offer” and “DHCP Acknowledge” received packets.
- Use “Forced Forwarding” in section 3.8.

3.7.2 DHCPv4 Starvation

Impact: DoS

DHCP Starvation is an attack that works by broadcasting vast numbers of DHCP requests with spoofed MAC addresses simultaneously, exhausting the DHCP server IP pool. If the IP pool is empty the DHCP server will not answer DHCP discover traffic and either cause a DoS or help the attacker perform DHCP Server Spoofing.

3.7.2.1 Requirement

Brief	Limit the number of IPv4 addresses that can be assigned to a customer port.		
ID	SEC-V4-DHCPSTARV-1	Priority	Must
Description	Don't let a single customer take all available IP addresses.		
Motivation	Avoid that a single user takes all DHCP addresses. Effectively creating DoS.		
Test	At one customer, request several IP addresses using different source MAC in the DHCP discover. Verify that the customer is not assigned more than the specified limit of addresses. Note if limitation is done at DHCP server the real DHCP server must be used to perform the test.		

3.7.2.2 Implementation example

- Limit the number of allowed MAC addresses on a customer port (section 2.1 “MAC Flooding”)
- If the Access Switch snoops DHCP traffic (section 3.1.3: “Customer identification”) it can track the number of assigned addresses per customer port and drop new requests if there are too many clients.
- The DHCP server can limit the number of clients on a port, using some customer identification as described in section 3.1.3 “Customer identification”.

3.7.3 IPv4 Multicast Discovery


Impact: Unauthorized access to resources

UPnP, Universal Plug And Play and its AV extensions is a protocol that permits networked devices, such as PCs, printers, Wi-Fi Access Points, ADSL modems, Broadband modems, TV sets, DVD and Blu-ray players etc. to seamlessly discover and share data. Apple's Airplay uses a similar protocol called Bonjour.

UPnP can for example stream a Video from a media server to a TV set or open access to certain protocols in your Broadband modems firewall (SIP, BitTorrent). UPnP is a common protocol and is very often enabled by default so users in a home are not aware of it, it just works. This is also one of the problems with the protocol, there are very little or no safety.

Note: DLNA – Digital Living Network Access – popular in some mobile phones and other home equipment is another name for UPnP AV (Audio and Video). If UPnP is blocked, DLNA is also blocked.

3.7.3.1 Normal Operation




The UPnP server and the TV set finds each other using the Simple Service Discovery Protocol (SSDP), which is multicast based. This makes it easy for the TV set to find the UPnP server without any configuration.

In this picture, the TV set has asked for and receives a Video stream using UPnP from the server. Server could for example be a PC running Windows 7.

No authentication or other access control is in use.

3.7.3.2 Attack – IPv4 UPnP, LLMNR mDNS and Bonjour



Here the TV set on one customer port asks for servers and incorrectly receives answers from another customer port.

This makes it possible for one customer to browse, download and watch content from the other customer's server.

This can be either "by accident" or a malicious user using the fact that there is no authorization in the UPnP, LLMNR, mDNS and Bonjour .

3.7.3.3 Requirement

Brief	IPv4 UPnP (SSDP), LLMNR, mDNS and Bonjour must be blocked between customer ports.		
ID	SEC-V4-UPNP-1	Priority	Must
Description	By blocking the Multicast Discovery traffic no incorrect access to other customers content can be done.		
Motivation	Avoid Abuse, illegal content access.		
Test	<p>All of the protocols uses Multicast traffic for discovery and have specific channels and ports. UPnP: Address: 239.255.255.250 Port: 1900 mDNS: Address: 224.0.0.251 Port: 5353 LLMNR: Address: 224.0.0.252 Port: 5355</p> <p>For these destination addresses and ports. Send a UDP packet to the address and port and verify that it does not arrive at other customers.</p> <p>Note: The multicast group should be joined by the "other customer".</p>		

3.7.3.4 Implementation example

- Use an Access Control List to identify and drop UPnP, LLMNR, mDNS and Bonjour packets
- Use "Forced Forwarding" in section 3.8.

3.7.4 IPv4 Fragmentation attacks

Impact: DoS, Abuse, Illegal access to content.

By using fragments when sending packets, for example an ACL could be traversed that normally denies such traffic.

3.7.4.1 Requirement

Brief	Drop packets with too small fragment offset		
ID	SEC-V4-FRAG-1	Priority	Must
Description	When a packet is received the fragmentation offset should be checked, if it is too small – pointing into the IP/UDP/TCP header it should be dropped. There is NO reason to forward such a packet.		
Motivation	Avoid Abuse, illegal content access.		
Test	Send packets that are not fragmented from a customer to the uplink. Verify that they are received. Send packets fragmented in headers from the customer to the uplink and verify that the fragments are not received. TCP packets can be fragmented in 8 and 16 byte fragments to fragment in the TCP header.		

3.7.4.2 Implementation example

Use an ACL that matches a fragment offset less than the maximum IPv4 header (20 bytes + options)

See also RFC1858 - Security Considerations for IP Fragment Filtering

3.7.5 Filter out IPv6

If there are no intention to use IPv6 in the network IPv6 should be filtered out.

3.7.5.1 Requirement

Brief	Filter out IPv6 when IPv6 isn't intended to be used		
ID	SEC-V4-FILTER-1	Priority	Must
Description	IPv6 – ethertype 0x86dd should be dropped		
Motivation	Avoid rouge IPv6 routers, MIM attacks, DOS attacks		
Test			

3.8 Forced Forwarding

Some vendors can't filter out specific traffic as described in section 3.7 "Protocol protection". They instead rely on a feature in which customer ports in the same L2 domain cannot communicate with each other.


Pros:

- Protocol agnostic, all IPv4 protocols will be sent to the upstream router
- No risk of new protocols being forgotten, by default zero traffic is allowed

Cons:

- All traffic between clients in the same L2 domain needs to be relayed through the default gateway router.
 - This can give bandwidth constraints.
- Proxy ARP needs to be enabled.
 - If this is enabled in the upstream router it will answer to all customer ARP requests, even if the host sending ARP request and the host sending the ARP reply is on the same port. Traffic between these two hosts will then flow through the default gateway router, probably being limited to the speed of the internet service.
 - If the proxy ARP is done in the Access Switch this problem does not occur.

3.8.1 Forced Forwarding, switch based



Ports are tagged as customer or network ports. Traffic received on network ports can be freely transmitted to other network or customer ports.

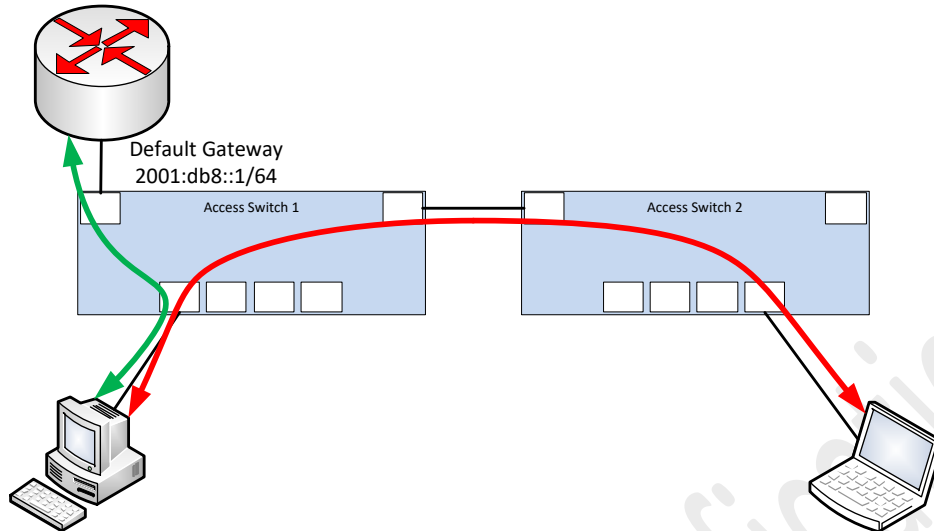
Traffic received on customer ports can only be sent to network ports, independent of protocol or type of traffic (unicast, multicast, broadcast).

As can be seen here, traffic between the desktop and laptop to the default gateway is passed.

Traffic between customer ports are dropped.

3.8.2 Forced Forwarding, network based

The forced forwarding feature in one switch is not enough if multiple switches are in the same L2 broadcast domain. This can easily be seen in the picture below.



If the rules between network and customer ports are followed, traffic from the desktop on switch 1 can go through network port over to switch 2. Here the traffic is then forwarded to the laptop and this is ok since the traffic came from a network port.

There is additional functionality / features needed to get the proper isolation between customer ports when multiple switches in same L2 broadcast domain is connected.

3.8.3 Requirement

Brief	Drop packets between customer ports in the same L2 broadcast domain		
ID	SEC-V4-FF-1	Priority	Must
Description	No layer 2 broadcast should be sent between customer ports.		
Motivation	Prevent unfiltered traffic between customers that can be used to create attacks.		
Test	From a customer send layer 2 broadcast and verify that it does not arrive at other customers. Both customers on the same access switch, and also at two different access switches.		

3.8.4 Implementation example

Different vendors have different names/techniques to accomplish this, some are

- Private VLAN (RFC5517, Independent Submission)
- MAC Forced Forwarding – MACFF (RFC4562)
- Source port filter
- Traffic Segmentation


These features are not 100% standardized, there may be interoperability problems between vendors.

3.9 Secure Boot, Access Switch

Impact: MITM, DoS, Abuse


If the Access Switch uses BOOTP or DHCPv4 for address configuration, the switch MUST NOT accept BOOTP/DHCPv4 packets from customer ports.

3.9.1 Normal Operation



The Access Switch boots and negotiates an IPv4 address with the DHCPv4 server. The IPv4 address is used for management purposes.

3.9.2 Attack – Switch boot, BOOTP/DHCPv4 server spoofing



If the BOOTP/DHCPv4 negotiation is done on all ports in the Access Switch, a malicious user can respond to those and give the switch another IPv4 address and potentially configuration.

After that the malicious user can do most of the MITM and DoS attacks, and complicate the Abuse handling.

3.9.3 Requirement

Brief	When Access Switch boots, only accept DHCPv4 / BOOTP packets from uplinks		
ID	SEC-V4-BOOT-1	Priority	Should
Description	To avoid incorrect IPv4 address configuration and possible switch configuration from a malicious user.		
Motivation			
Test			

3.9.4 Implementation example

- Do not send BOOTP/DHCPv4 packets on customer ports at boot

- Verify that any BOOTP/DHCPv4 responses are not from customer ports, if they are drop them.

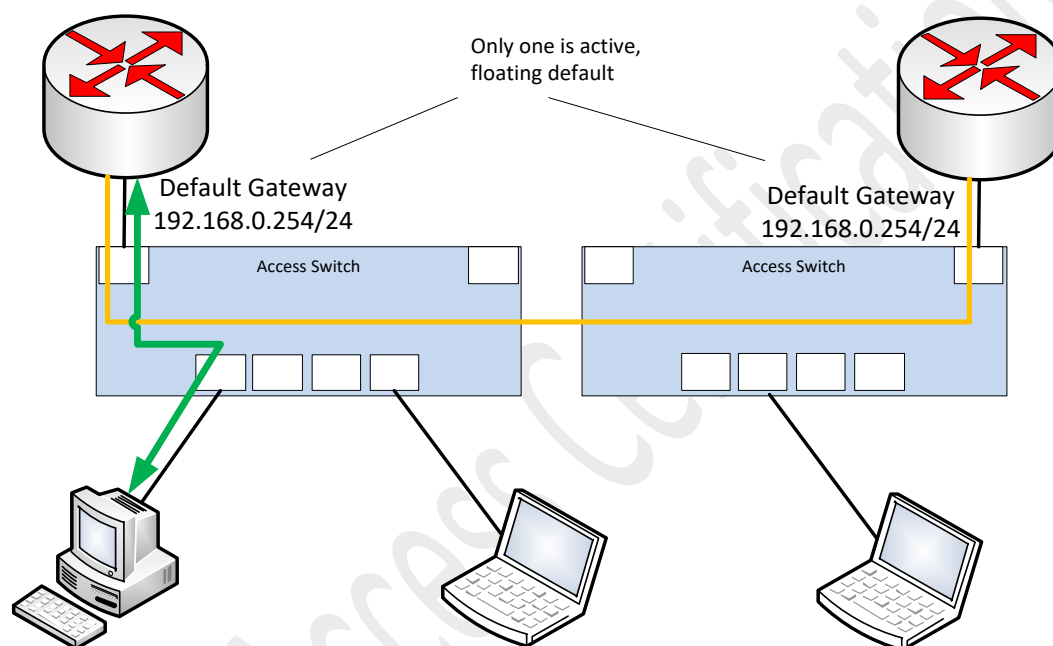
3.10 VRRP / HSRP spoofing

Impact: MITM, DoS

There exist other protocols with basically the same purpose such as CARP, NSRP, GLBP, SMLT and ESRP. Some of them work similarly, and should possibly be added.

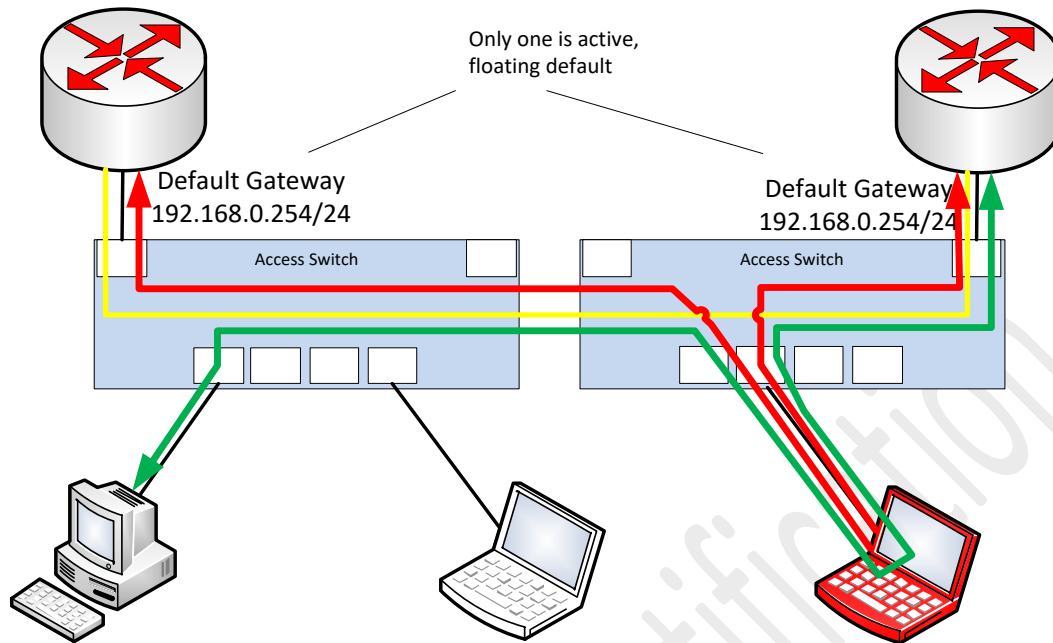
3.10.1 Normal Operation

Many networks today use VRRP, HSRP or similar to create a virtual default gateway for redundancy purposes. If one default gateway goes down the other takes over.



It is important that the protocol used to setup and control this feature is protected by encrypted shared secrets or similar so a malicious user cannot be a member of the VRRP/HSRP group. If that would be possible the malicious user can force customer traffic to point their default route to the malicious user and create an MITM attack.

3.10.2 Attack – VRRP / HSRP spoofing



A malicious user first sniffs the VRRP/HSRP packets, then creates and sends similar packets into the VLAN/L2 broadcast domain (red line), becoming a member of the VRRP/HSRP group. By changing the priority it can force the other peers to give up, effectively forcing all clients to send all packets through the malicious user (green line).

3.10.3 Requirement

Brief	VRRP/HSRP from customer ports should be blocked		
ID	SEC-V4-xxRP-1	Priority	Must
Description			
Motivation	A man-in-the middle attack can otherwise be created.		
Test	Sniff traffic on a customer port for 60 seconds.		
	The requirement is verified if no router redundancy protocol traffic is sent to customer ports.		
	Note: The test cannot be performed when only a switch is tested in a lab, only in a real network.		

3.10.4 Implementation example

Protect the communication between the VRRP/HSRP neighbors, using hashes, cryptographic signatures or similar.

3.11 Routing Protocols

Impact: MITM, DoS

3.11.1 Normal operation

Routers in a network send route updates between each other using one of many protocols. Some of the protocols uses multicast to communicate these updates and messages sent from one router will be received by many others, some protocols using unicast to send specific updates from one router to another.

3.11.2 Attack Routing Protocols

A malicious user can by sending updates using the current routing protocol alter the path taken by packets to internet. This can be accomplished by listening for this type of packets and send similar ones adding himself as a path.

3.11.3 Requirement

Brief	Protect routing protocols from malicious users.		
ID	SEC-V4-ROUTE -1	Priority	Must
Description	Routing protocol traffic should not be received nor sent to customer ports.		
Motivation	Avoid DoS by exploiting routing protocols.		
Test	Issue multicast join messages for the groups used by common routing protocols. Sniff traffic on a customer port for 60 seconds. The requirement is verified if no routing protocol traffic is sent to customer ports. Some examples of routing protocols; IS-IS, OSPF, BGP, IGRP and RIP. Note: The test cannot be performed when only a switch is tested in a lab, only in a real network.		

3.11.4 Implementation example

Disable routing on access ports in the routers.

Filter routing traffic on customer ports in the access switch.

4 SEC - IPv6

This section contains a description of the various problems that can occur in a shared L2 segment using the IPv6 protocol.

4.1 IPv6 overview

To better understand the issues, we begin with some generic information.

4.1.1 Host Address Configuration

A host that wants to use IPv6 on the global internet needs a global IPv6 unicast address. There are several methods the host can use for this.

- Static configuration
- IPv6 Stateless Address Auto Configuration - SLAAC (RFC2462)
- Dynamic Host Configuration Protocol IPv6 - DHCPv6 (RFC3315)

Static Configuration

- Static configuration is not very practical. IPv6 addresses are long and hard to remember. There are several addresses, such as DNS servers, NTP servers etc. that also needs to be manually configured.

SLAAC

- SLAAC is supported by all operating systems that support IPv6. It is a simple and stateless protocol in which the client uses the Ethernet MAC address as a base for selecting the host part. The host listens on Router Advertisement messages to learn what prefix and what default gateway that should be used.
- SLAAC has several drawbacks
 - The Ethernet MAC address is identical and globally unique, which means the host part of the IPv6 address always will be the same. This opens up for privacy concerns, wherever the computer is connected and uses IPv6 it can be tracked where it is. This was considered so bad that there is an extension to SLAAC that allows the host to randomly change the host part of the IPv6 address - RFC3041 Privacy Extensions for Stateless Address Auto Configuration in IPv6.
 - RFC3041 itself has some issues. The host never asks for a new address it just takes one, at a certain time the host may have several IPv6 addresses that all need to have full network access.
- Using SLAAC as the address configuration method is not recommended, there are no clients today with default settings that can retrieve the DNS servers over SLAAC (RFC 5006).

DHCPv6

- Using DHCPv6 is the most flexible way to assign addresses to clients and gives the most control to the Broadband Network/Service Providers.
- DHCPv6 is an extension to SLAAC, there is no standardized way to only use DHCPv6 without SLAAC. DHCPv6 for example does not return default gateway, which is part of RA.

4.1.2 IPv6 Address allocation

In the IPv4 world, IPv4 addresses are a limited resource and must be effectively utilized. This is one of the major benefits of placing multiple customers into the same VLAN. The IPv4 addresses can be freely spread over the customers.

When implementing IPv6 this sharing of one or a few prefixes over many customers can of course be copied to the IPv6 world. One question directly arises; is this really necessary? In IPv6 there is huge amount of available addresses.

Each Customer port could have its own IPv6 /64 prefix. That would simplify several things, for example

- IPv6 source address validation
- Duplicate Address Detection

Instead of tracking each host on a port and building a table dynamically from static, DAD snooping and/or DHCPv6 snooping a single filter entry the /64 prefix could be used. That filter would be very static and typically installed when the port/service is activated on the port.

4.1.3 Customer Identification

Brief	Identification info must be added to the DHCPv6 snooped packets		
ID	SEC-V6-DHCP-1	Priority	Must
Description	Typically RFC3315 DHCPv6 Interface-ID Option 18 and RFC4649 Remote-ID option 37		
Motivation			
Test			

4.2 IPv6 Whitelist Database

To verify received packets on customer ports that they have the correct source MAC and source IPv6 address, a Whitelist Database can be built. The whitelist database typically contains customer port, VLAN, MAC Address and IPv6 address.

There are several methods to generate the whitelist database; some of them may need to be used together.

4.2.1 Static configuration

Each port is manually configured with the correct information. The manually configured entries are added to the Whitelist Database.

4.2.2 Strict ports

Allow one or a couple of source MAC/source IPv6 sources to connect, when limit is reached any other sources are dropped. The sources that are allowed are added to the Whitelist Database.

4.2.3 DHCPv6 snooping

By listening on the DHCPv6 packets, the Access Switch can build a table based on the assigned addresses by the DHCPv6 server. This is sometimes called DHCPv6 snooping.

[draft](#) SAVI Solution for DHCP

4.2.4 SLAAC snooping

Since SLAAC IPv6 address configuration is not a negotiation SLAAC cannot directly be snooped. Indirect methods could be used. See implementation below.

4.2.5 Requirement

Brief	IPv6 whitelist database		
ID	SEC-V6-DHCP-2	Priority	Must
Description	DHCPv6 and SLAAC/NDP packets are snooped and builds a whitelist database with mac and IPv6-address If DHCPv6 only should be used the SLAAC/static addresses should be filter out. If SLAAC/static addresses is use the equipment must log the assigned/generated address(es).		
Motivation	Make sure a IPv6 address can not be spoofed/hijacked		
Test			

Brief	DHCPv6 Prefix Delegation must be snooped and added to the whitelist database.		
ID	SEC-V6-DHCP-3	Priority	Must
Description	In addition to the DHCPv6 client sources, the RFC3633 addresses (DHCPv6 options) must also be snooped		
Motivation			
Test			

4.2.6 Implementation example

SLAAC Snooping

The Access Switch can listen for Duplicate Address Detection packets and other NDP messages. Using these it creates entries in the Whitelist Database.

Entries created need to be purged by a timeout since there is no protocol that informs the switch if an IPv6 address is not used anymore.

RFC3041 – “Privacy Extensions for Stateless Address Auto configuration in IPv6” can complicate this a little, since a client can have a number of IPv6 addresses concurrently and this is valid.

DHCPv6 snooping

The original RFC3315 that specifies the DHCPv6 protocol has no notion of a DHCPv6 snooper that can add options when operating in L2 mode. This is later fixed and specified in RFC 6221 – “Lightweight DHCPv6 Relay Agent” (LDRA).

The LDRA MUST intercept and process all IP traffic received on any client-facing interface that has:

- destination IP address set to All_DHCP_Relay_Agents_and_Servers(ff02::1:2);
- protocol type UDP; and
- Destination port 547.

The LDRA MUST also prevent the original message from being forwarded on the network-facing interface.

SEC Access Certification

4.3 IPv6 source address spoofing

Impact: DoS, Abuse

Corresponding IPv4 issue: Section 3.4 “IPv4 Source Address Spoofing”. The only difference is that IPv6 addresses are used instead of IPv4.

4.3.1 Requirement

Brief	Filter out spoofed IPv6 packets in the access switch.		
ID	SEC-V6-SPOOF-1	Priority	Must
Description			
Motivation			
Test			

Brief	Track when spoofed IPv6 packets are received		
ID	SEC-V6-SPOOF-2	Priority	Should
Description	When a spoofed packet is received it should be notified in a counter (per customer port) or a log.		
Motivation	Used for troubleshooting. If the spoofed packet counter is incremented there are some issues with IPv4 address configuration and the verification of source addresses.		
Test			

4.3.2 Implementation example

By using the IPv6 whitelist table an access filter can be applied on customer ports. If a packet with an IPv6 source address is received that has a source address that doesn't exist in the whitelist table, the packet is dropped.

IETF has a working group, Source Address Validation Improvements – SAVI that works on how to control and only allow valid IPv6 source addresses into the network. SAVI charter can be found at [Link](#).

The SAVI working group is working on several ways to control and only permit valid source addresses. Some of them are:

- FCFS – First-Come First-Serve Source address validation for locally assigned IPv6 addresses [Link](#)
 - o The first IPv6 source address seen on the port is the only that will be allowed.
- SEND-based Source-Address Validation Implementation [Link](#)
 - o Depends on cryptography and pre-shared information.
 - o This is not feasible to implement in a Broadband network, it will generate too much administrative work, and give the end-users a negative experience when connecting devices to the network.
- SAVI Solution for DHCP [Link](#)
 - o Snooping DHCPv6 messages
-


The working group has identified a couple of issues with their current drafts, please see the IETF draft [Link](#) section

4.4 Neighbor Solicitation / Advertisement

Impact: MITM, DoS

Corresponding IPv4 issue: ARP Poisoning

4.4.1 Normal Operation




Neighbor Solicitation / Advertisement create a binding between an L3 IPv6 address and the L2 address (in broadband network normally the Ethernet MAC Address).

Client sends a Neighbor Solicitation using multicast, asking for 2001:db8::1 (default gateway). All IPv6 capable nodes are listening on that multicast group.

The addressed IPv6 node replies with the Ethernet MAC Address, in this case the L3 router/default gateway.

4.4.2 Attack - Neighbor Solicitation / Advertisement



A malicious user can incorrectly reply to the Neighbor Solicitation message, specifying its own MAC address. This has the effect that the client will use the malicious user as the default gateway, creating a MITM attack.

If the malicious user replies with an unknown MAC address the client cannot communicate with the rest of the network, nothing will receive the packets.

4.4.3 Requirement

Brief	Filter out ND/NS messages with unknown source IPv6 addresses		
ID	SEC-V6-NS-1	Priority	Must
Description			
Motivation			
Test			

4.4.4 Implementation example

IETF

[draft](#) IPv6 Router Solicitation Driven Access Considered Harmful

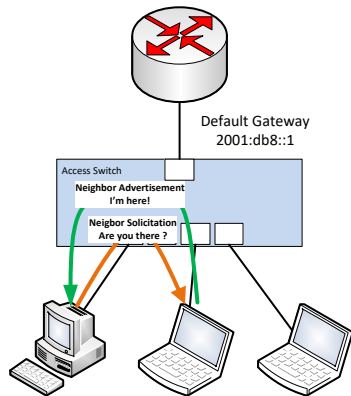
SEC Access Certification

4.5 Neighbor Unreachability Detection

Impact: DoS

Corresponding IPv4 issue: None

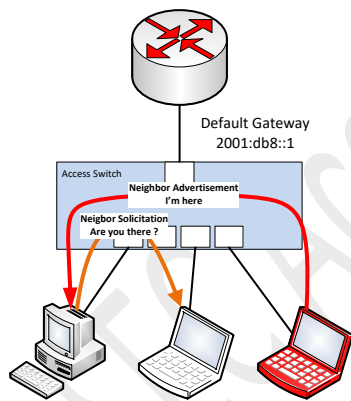
4.5.1 Normal Operation



Each entry in the Neighbor cache contains a timer. When a packet is received from a neighbor the timer for that neighbor is reset. When the timer times out, the cache entry is considered stale.

The NUD procedure is invoked and an NS is sent, probing for the stale host. If the stale host responds with an NA, the ND cache entry is put in Reachable state again, and the timer is reset.

4.5.2 Attack – Neighbor Unreachability Detection



A malicious user can respond to the NS sent by the host that have initiated the NUD procedure. If the response contains incorrect data the communication between the two end-users will be stopped, effectively creating a DoS attack.

4.5.3 Requirement

Brief	Neighbor Unreachability Detection [NUD, RFC4861] filtering		
ID	SEC-V6-NUD-1	Priority	Must
Description	There must be a NUD filtering function to ensure that false NUD messages cannot be sent into the network.		
Motivation			
Test			

4.5.4 Implementation example

- Any NUD packets received on a customer port must be screened against the whitelist database and if they are unknown drop the packet.
- Any NUD packets received on a customer port is dropped.

IETF

[draft](#) Neighbor Unreachability Detection is too impatient


SEC Access Certification

4.6 ICMPv6 Redirect Messages

Impact: MITM, DoS

Corresponding IPv4 issue: Section 3.3 "ICMPv4"

4.6.1 Normal Operation



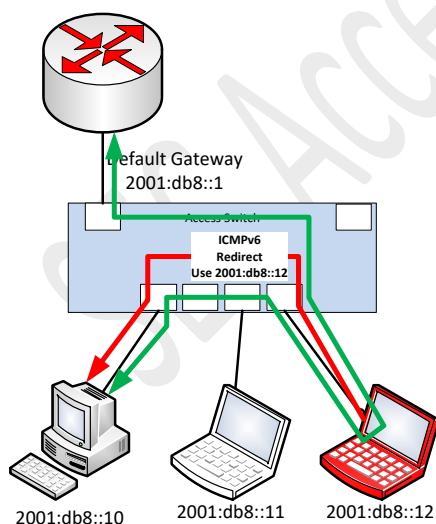
Redirect informs a client of a better way to reach the destination. It is used to optimize traffic.

In a broadband network there is normally only one way to the rest of the network so ICMPv6 redirect messages is not that useful.

Client communicates with the rest of the network through the default gateway (green line). When sending packets to 2001:db8:1::/64 the default gateway will send an ICMPv6 redirect back to the client (yellow line), asking it to use 2001:db8::2 instead for that destination.

This removes the packet forwarding task from the default router, the client sends packets in a more direct way to that destination (dashed green line).

4.6.2 Attack - Redirect Messages



A malicious user can send an incorrect redirect message to a client, informing the client of a better path to the default gateway, using itself as the redirect target

This effectively creates an MITM attack, in which all traffic from the client to the default gateway passes through the malicious user.

If an unknown nexthop is used instead of the malicious user a DoS attack is created.

4.6.3 Requirement

Brief	Filter out incorrect ICMPv6 redirect messages on customer ports.		
ID	SEC-V6-REDIRECT-1	Priority	Must
Description	When an ICMPv6 redirect message is received on a customer port it needs to be verified that it actually is correct.		
Motivation	Avoid MITM, DoS		
Test			

4.6.4 Implementation example

- Any ICMP Redirect packets received on a customer port must be screened against the whitelist database and if they are unknown drop the packet.
- Any ICMP Redirect packets received on a customer port is dropped.
- Use “Forced Forwarding” in section 4.14.


SEC Access Certification

4.7 Duplicate Address Detection - DAD

Impact: DoS

Corresponding IPv4 issue: None

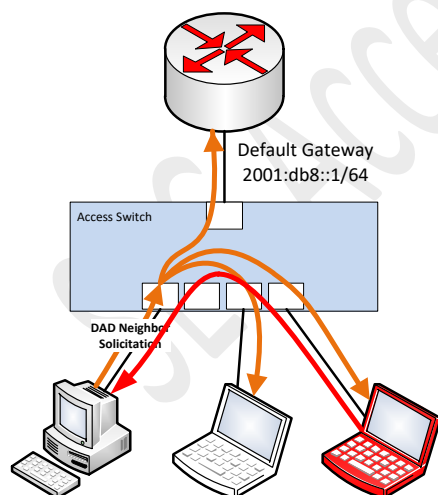
4.7.1 Normal operation



When a host has completed section 4.1.1 “Host Address ” the host asks everybody else in the same L2 segment if the IPv6 address is already in use, verifying the address uniqueness.

If everything is ok, there are no responses and the host completes the host address configuration and starts to use IPv6.

4.7.2 Attack – Duplicate Address Detection



The malicious user answers the DAD packet informing the client that the IPv6 address is already in use. This results in the client disabling IPv6 processing on its interface (according to the standard)

4.7.3 Requirement

Brief	Drop invalid replies to IPv6 Duplicate Address Detection packets		
ID	SEC-V6-DAD-1	Priority	Must
Description	When a DAD response is received, ensure that the packet is sent from a known IPv6 address on the Switch Access port. The detection of a valid address can be identified by several mechanisms.		
Motivation	Avoid Denial of Service by malicious users sending fake DAD responses.		
Test			

4.7.4 Implementation example

- Implement a duplicate address detection in the Access Switch
- Filter out incorrect DUD packets, which contain an unknown IPv6 source address, both in the IPv6 header and the DUD payload.
- Use “Forced Forwarding” in section 4.14.

IETF

RFC4429 – Optimistic Duplicate Address Detection

[draft](#) Duplicate Address Detection Proxy

[draft](#) IPv6 DAD Enhancements for handling Layer1 Loopbacks


[draft](#) Enhanced Duplicate Address Detection

4.8 Router Advertisement attack

Impact: MITM, DoS

Corresponding IPv4 issue: None

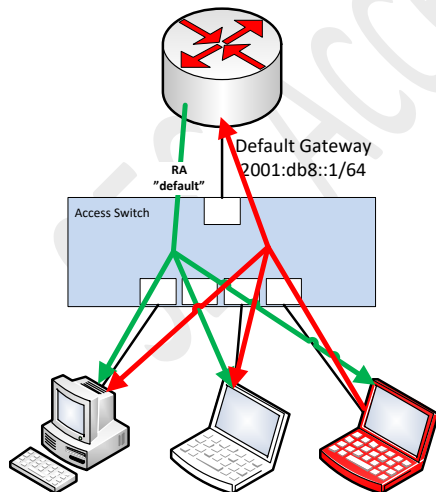
4.8.1 Normal Operation



Periodically and on demand by clients, each router announces itself on the L2 segment as a candidate for being a default gateway, using Route Advertisement (RA) messages.

Clients listen after Route Advertisement and use the information in those to add a default route.

4.8.2 Attack - Router Advertisement



A malicious user can send false Route Advertisement packets, forcing other hosts on the L2 segment to either use the malicious user as the default route – creating a MITM attack, or using an unknown host/MAC address as default route, creating a DoS attack (which makes it impossible for clients to communicate).

4.8.3 Requirement

Brief	Block Received Router Advertisement packets on customer ports.		
ID	SEC-V6-RA-1	Priority	Must
Description			
Motivation			
Test			

There is a standard protocol called SEND (RFC3971) that can be used to limit this problem, but it has a fairly high installation cost, in which shared secrets/certificates need to be exchanged between each host and the access switches. This is not practical in a Broadband Network, where the network operator has no control over the client environment.

4.8.4 Implementation example

RFC6104 Rouge IPv6 Router Advertisement Problem Statement

RFC6105 IPv6 Router Advertisement Guard

IETF Draf

[draft](#) IPv6 Router Advertisement Guard (RA-Guard) Evasion

[draft](#) Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)

4.9 IPv6 Routing Header

Impact: DoS

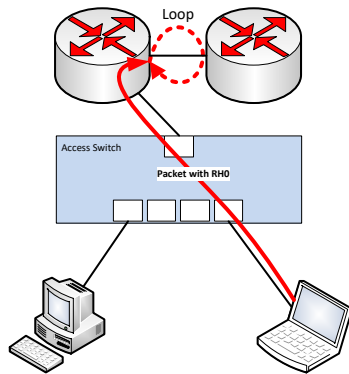
Corresponding IPv4 issue: None

RFC2460 specifies the IPv6 Routing Header. Note [RFC5095](#) – Deprecation of Type 0 Routing Headers in IPv6. This was a badly constructed part of the IPv6 protocol and should not be possible to use anymore.

4.9.1 Normal Operation

Since Routing Header type 0 (RH0) is deprecated there is no Normal Operation.

4.9.2 Attack - IPv6 routing header



A single RH0 may contain multiple intermediate node addresses, and the same address may be included more than once in the same RH0. This allows a packet to be constructed such that it will oscillate between two RH0-processing hosts or routers many times. This allows a stream of packets from an attacker to be amplified along the path between two remote routers

This attack is particularly serious in that it affects the entire path between the two exploited nodes, not only the nodes themselves or their local networks.

A malicious user can send a specially crafted packet with a RH0 header that points to two routers, the routers will bounce the packet between them until Hop Count becomes zero, effectively consuming bandwidth and possibly create a DoS attack.

4.9.3 Requirement

Brief	Drop all received packets on customer ports that contains a type 0 Routing Header		
ID	SEC-V6-RH-1	Priority	Must
Description	IPv6 Routing Header messages must not be allowed between subscriber ports and subscriber and uplink. The routing header cannot be disabled because it is Required for mobile IPv6.		
Motivation	Routing loops are avoided		
Test			

4.9.4 Implementation example

Use a filter on customer ports that identifies IPv6 packets with a type 0 routing header and drop those packets.

4.10 MLD/MLDv2

Impact: Avoid sending multicast traffic to uninterested ports

Corresponding IPv4 issue: Section 3.5 “IGMP”

MLD is the IPv6 version of IGMPv2

MLDv2 is the IPv6 version of IGMPv3

4.10.1 Normal Operation

See Section 3.5.1 “Normal Operation”

4.10.2 Requirement

Brief	Use MLD/MLDv2 messages to learn about interesting multicast receivers		
ID	SEC-V6-MLD-1	Priority	Should
Description	By monitoring MLD/MLDv2 packets, the switch can learn what ports are interested in a certain (*,G) or (S,G) multicast group. This can then control the packet forwarding/replication.		
Motivation	Efficient multicast distribution without wasting bandwidth.		
Test			

4.10.3 Implementation example

MLD is specified in RFC 3810.

RFC4541 [Link](#) – “Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches” describes how to implement MLD snooping.

4.11 IPv6 Multicast Group Overload

Impact: DoS

Corresponding IPv4 issue: Section 3.6 “Multicast group overload”

A client / STB can ask for any number of multicast groups and have all of them at the same time. This can exceed the available bandwidth in the network and on the customer port. It can also consume all available entries in the multicast forwarding table, effectively creating a DoS for multicast on other ports.

4.11.1 Normal Operation

Each multicast receiver is joining one or a few concurrent multicast channels.

4.11.2 Attack – Multicast Group Overload

A malfunctioning device or malicious user sends joins using MLD/MLDv2 protocol for all possible multicast groups there are. This is a huge number and no existing router/switch can track that many multicast groups.

When another customer port wants to change or join a multicast channel (for example watching a TV channel) there is no way the Access Switch can distribute this new channel to the port, effectively creating a DoS attack.

4.11.3 Requirement

Brief	Limit number of concurrent multicast groups a customer port can join using the MLD/MLDv2 protocol		
ID	SEC-V6-MCAST-1	Priority	Should
Description			
Motivation			
Test			

4.11.4 Implementation example

Count the number of active multicast groups on each customer port. If the limit is reached, refuse to add the requested group to the customer port.

4.12.3 Requirement

Brief	Protect against filling up ND cache		
ID	SEC-V6-NDCACHE-1	Priority	Should
Description			
Motivation			
Test			

4.12.4 Implementation example

- Don't create L3->L2 mappings for packets going downstream to a customer port. If DHCPv6 is used for address configuration, the client has already communicated with the Default router so the L3->L2 binding already exist.

[draft](#) Mitigating Neighbor Discovery Based Denial of Service Attacks

[draft](#) Operational Neighbor Discovery Problems

SEC Access Certification

4.13 Protocol Protection

4.13.1 DHCPv6 server spoofing

Impact: MITM, DoS

Corresponding IPv4 issue: Section 3.7.1 “DHCPv4 Server Spoofing”

4.13.1.1 Requirement:

Brief	DHCPv6 must be blocked between subscriber ports		
ID	SEC-V6-DHCP-1	Priority	Must
Description			
Motivation			
Test			

4.13.1.2 Implementation example

- Use an Access Control List on customer ports that identifies and drops DHCPv6 response packets
- Use “Forced Forwarding” in section 4.14.

4.13.2 DHCPv6 Starvation

Impact: DoS

Corresponding IPv4 issue: Section 3.7.2 “DHCPv4 Starvation”

4.13.2.1 Requirement

Brief	Limit the number of IPv6 addresses/prefixes that can be assigned to a customer port.		
ID	SEC-V6-DHCPSTARV-1	Priority	Should
Description			
Motivation			
Test			

4.13.2.2 Implementation example

- Limit the number of allowed MAC addresses on a customer port (section 2.1 “MAC Flooding”)
- If the Access Switch snoops DHCPv6 traffic (section 4.2.3 “DHCPv6 snooping”) it can track the number of assigned addresses/prefixes per customer port and drop new requests if there are too many clients.
- The DHCPv6 server can limit the number of clients on a port, using some customer identification as described in section 4.1.3.

4.13.3 IPv6 UPnP, LLMNR, mDNS and Bonjour filtering

Impact: Unauthorized access to resources

Corresponding IPv4 issue: Section 3.7.3 “IPv4 ”

4.13.3.1 Requirement

Brief	IPv6 UPnP, LLMNR, mDNS and Bonjour must be blocked between subscriber ports		
ID	SEC-V6-UPNP-1	Priority	Must
Description			
Motivation			
Test			

4.13.3.2 Implementation example

- Use an Access Control List on customer ports to identify and drop any received UPnP packets.
- Use “Forced Forwarding” in section 4.14.

SEC Access Certification

4.13.4 IPv6 Fragmentation attacks

Impact: DoS, Abuse, Illegal access to content.

Corresponding IPv4 feature: Section 3.7.4 "IPv4 Fragmentation attacks"

By using fragments when sending packets, an ACL could for example be traversed that normally denies the traffic.

4.13.4.1 Requirement

Brief	Drop IPv6 packets with too small fragment offset.		
ID	SEC-V6-FRAG-1	Priority	Should
Description	When a packet is received the fragmentation offset should be checked, if it is too small – pointing into the IP/UDP/TCP header it should be dropped. There is NO reason to forward such a packet.		
Motivation			
Test			

4.13.4.2 Implementation example

[RFC5722](#) - Handling of Overlapping IPv6 Fragments

[draft](#) Processing of IPv6 "atomic" fragments

[draft](#) Tiny Fragments in IPv6

4.14 Forced Forwarding

The forced forwarding feature for IPv6 is very similar to IPv4 section 3.8 “Forced Forwarding”.

The blocking of packets between customer ports are the same. IPv6 is different so the IPv4 features cannot be used.

4.14.1 Requirement

Brief	Drop packets between customer ports in the same L2 broadcast domain		
ID	SEC-V6-FF-1	Priority	
Description			
Motivation			
Test			

4.14.2 Implementation example

Since packets between customer ports cannot be exchanged, additional features are needed so clients can communicate.

- Neighbour Solicitation / Advertisement
 - [RFC4389](#) Neighbor Discovery Proxies (ND Proxy)
 - This can be done in the access switch or in the default router (as Proxy ARP)
- Neighbor Unreachability Detection
 - Probably not needed in a broadband network, ok to filter them out.
- Redirect Messages
 - Probably not needed in a broadband network, ok to filter them out.
- Duplicate Address Detection – DAD
 - Makes sure no clients using SLAAC gets the same IPv6 source address.
 - If DHCPv6 is used for address configuration, the DHCPv6 server ensures that two clients don't get the same IPv6 source address. The client still need to verify that the received DHCPv6 assigned address is unique using DAD.
 - Duplicate Address Detection Proxy [draft](#)
- Router Advertisement Attack
 - No problems, client cannot send packets directly to each other.
- IPv6 Routing Header
 - No problems, client cannot send packets directly to each other.

4.15 Secure Boot, Access Switch

Impact: MITM, DoS, Abuse

Corresponding IPv4 Issue: Section 3.9 “Secure Boot, Access Switch”

Requirement

Brief	When Access Switch boots, only accept DHCPv6 packets from uplinks		
ID	SEC-V6-BOOT-1	Priority	Must
Description	To avoid incorrect IPv6 address configuration and possible wrong configuration from a malicious user.		
Motivation			
Test			

4.15.1 Implementation example

Same as for IPv4.

SEC Access Certification

4.16 VRRPv6 / HSRPv6

Impact: MITM, DoS

Corresponding IPv4 issue: Section 3.10 “

VRRP / HSRP spoofing”

4.16.1 Requirement

Brief	VRRPv6/HSRPv6 or similar protocol must be protected against malicious users		
ID	SEC-V6-xxRP-1	Priority	Must
Description			
Motivation			
Test			


4.16.2 Implementation example

Protect the communication between the VRRPv6/HSRPv6 neighbors, using hashes, cryptographic signatures or similar.

5 Sample configurations

This section describes several network designs, what is required to implement them and pros and cons for each.

5.1 Sample 1, L2 in access, client traffic exchanged on L2. shared VLAN/L2 for customers



Traffic between clients is exchanged on L2 in each Access Switch. Certain type of traffic is filtered out or snooped to ensure no MITM, DoS and/or Abuse issues can occur.

Client exchange of traffic in Access: Yes

IPv4 and IPv6 protocol support.

IPv4 and IPv6 address configuration is done using DHCPv4 and DHCPv6.

Common

- Spanning-tree between Access Switches, to block any potential loop
- Limit of number of MAC addresses on each customer port.
- Multicast injection filter
- Spanning-tree filter
- Filter on customer ports that drops all received non IPv4/IPv6 packets.
- Protection of control plane against flooding
- Protection of packet buffer starvation
- Drop fragmented IPv4/IPv6 packets to the control plane
- Broadcast, Multicast and unknown destination flooding rate limiting
- Loop-detect

IPv4

- DHCPv4 option-82 tagging. Customer ports are configured as untrusted
- DHCPv4 snooping for building IPv4 whitelist database
- Filter out ARP poisoning packets
- Filter out ICMP redirect packets
- Filter out spoofed IPv4 packets
- IGMP snooping for efficient multicast control/distribution of traffic
- Limit number of concurrent multicast groups a customer port can join using the IGMP protocol
- Filter out all DHCP server traffic on customer ports.
- Limit the number of IPv4 addresses that can be assigned to a customer port.

- IPv4 UPnP must be blocked between customer ports.
- Drop packets with too small fragment offset.
- When Access Switch boots, only accept DHCPv4 / BOOTP packets from uplinks

IPv6

- Identification info must be added to the DHCPv6 snooped packets, with option 18 and 37 to identify customer ports.
- DHCPv6 snooping to build the IPv6 whitelist database
- DHCPv6 Prefix Delegation must be snooped and added to the whitelist database.
- Filter out spoofed IPv6 packets in the access switch.
- Filter out ND/NS messages with unknown source IPv6 addresses
- Neighbor Unreachability Detection [NUD, RFC4861] filtering
- Filter out incorrect ICMPv6 redirect messages on customer ports.
- Drop invalid replies to IPv6 Duplicate Address Detection packets
- Block Received Router Advertisement packets on customer ports.
- Drop all received packets on customer ports that contains a type 0 Routing Header
- Use MLD/MLDv2 messages to learn about interesting multicast receivers
- Limit number of concurrent multicast groups a customer port can join using the MLD/MLDv2 protocol
- DHCPv6 must be blocked between subscriber ports
- Limit the number of IPv6 addresses that can be assigned to a customer port.
- IPv6 UPnP must be blocked between subscriber ports
- Drop IPv6 packets with too small fragment offset.
- When Access Switch boots, only accept DHCPv6 packets from uplinks

Distribution router

- VRRP for redundant default router.
- VRRP must be protected against malicious users
- VRRPv6 in distribution routers for redundant default router.
- VRRPv6 protocol must be protected against malicious users
- Protect against filling up ND cache


Pros

- Efficient bandwidth usage in the network. Traffic between clients does not need to go through the default gateway (depending on spanning tree topology).

Cons

- Many tricky and non-standardized features to make the network operational.
- Sensitive to new protocols, they will not be filtered between customer ports unless the network operator adds them/upgrades software in access switch.

5.2 Sample 2, L2 in access, client traffic exchanged on L3, shared VLAN/L2 for customers



Traffic between clients is exchanged on L3 in the default router. All direct communication between the clients on L2 is prohibited. To allow the communication on L3, proxy ARP / proxy ND is used.

Client exchange of traffic in Access: No

IPv4 and IPv6 protocol support.

IPv4 and IPv6 address configuration is done using DHCPv4 and DHCPv6.

Common

- Spanning-tree between Access Switches, to block any potential loop
- Limit of number of MAC addresses on each customer port.
- Multicast injection filter
- Spanning-tree filter
- Filter on customer ports that drops all received non IPv4/IPv6 packets.
- Protection of control plane against flooding
- Protection of packet buffer starvation
- Drop fragmented IPv4/IPv6 packets to the control plane
- Broadcast, Multicast and unknown destination flooding rate limiting
- Loop-detect

IPv4

- MACFF to isolate customer ports from each other on L2.
- Proxy ARP in each access switch (part of MACFF)
- DHCPv4 option-82 tagging. Customer ports are configured as untrusted
- DHCPv4 snooping for building IPv4 whitelist database
- Filter out ARP poisoning packets
- Filter out spoofed IPv4 packets
- IGMP snooping for efficient multicast control/distribution of traffic
- Limit number of concurrent multicast groups a customer port can join using the IGMP protocol
- Limit the number of IPv4 addresses that can be assigned to a customer port.
- Drop packets with too small fragment offset.
- When Access Switch boots, only accept DHCPv4 / BOOTP packets from uplinks

IPv6

- Identification info must be added to the DHCPv6 snooped packets, with option 18 and 37 to identify customer ports.
- DHCPv6 snooping to build the IPv6 whitelist database
- DHCPv6 Prefix Delegation must be snooped and added to the whitelist database.
- Filter out spoofed IPv6 packets in the access switch.
- Filter out ND/NS messages with unknown source IPv6 addresses
- Neighbor Unreachability Detection [NUD, RFC4861] filtering
- Filter out incorrect ICMPv6 redirect messages on customer ports.
- Drop invalid replies to IPv6 Duplicate Address Detection packets
- Block Received Router Advertisement packets on customer ports.
- Drop all received packets on customer ports that contains a type 0 Routing Header
- Use MLD/MLDv2 messages to learn about interesting multicast receivers
- Limit number of concurrent multicast groups a customer port can join using the MLD/MLDv2 protocol
- DHCPv6 must be blocked between subscriber ports
- Limit the number of IPv6 addresses that can be assigned to a customer port.
- IPv6 UPnP must be blocked between subscriber ports
- Drop IPv6 packets with too small fragment offset.
- When Access Switch boots, only accept DHCPv6 packets from uplinks
- Duplicate Address Detection Proxy, so DAD works in the L2 broadcast domain

Distribution router

- Neighbor Solicitation Proxy, so NS works in the L2 broadcast domain
- VRRP for redundant default router.
- VRRP must be protected against malicious users
- VRRPv6 in distribution routers for redundant default router.
- VRRPv6 protocol must be protected against malicious users
- Protect against filling up ND cache


Pros:

- Not sensitive on new protocols, all are denied by default.

Cons:

- Many tricky and non-standardized features to make the network operational.
- All traffic between clients is exchanged at L3 in the default gateway, limiting the available bandwidth

5.3 Sample 3, L2 in access, client traffic exchanged on L3, separate VLAN/L2 per customer port



Each customer port has a separate L2 broadcast domain, implemented with a unique VLAN.

Access Switches are interconnected on L2 trunking the VLANs between them.

Client exchange of traffic in Access: No

IPv4 and IPv6 protocol support.

IPv4 and IPv6 address configuration is done using DHCPv4 and DHCPv6.

Common

- Spanning-tree between Access Switches, to block any potential loop
- One unique VLAN per customer port
- Limit of number of MAC addresses on each customer port.
- Multicast injection filter
- Spanning-tree filter
- Protection of control plane against flooding
- Protection of packet buffer starvation
- Drop fragmented IPv4/IPv6 packets to the control plane
- Broadcast, Multicast and unknown destination flooding rate limiting
- Loop-detect

IPv4

- IGMP snooping for efficient multicast control/distribution of traffic
- Limit number of concurrent multicast groups a customer port can join using the IGMP protocol
- When Access Switch boots, only accept DHCPv4 / BOOTP packets from uplinks

IPv6

- Identification info must be added to the DHCPv6 snooped packets, with option 18 and 37 to identify customer ports
- Use MLD/MLDv2 messages to learn about interesting multicast receivers
- Limit number of concurrent multicast groups a customer port can join using the MLD/MLDv2 protocol
- When Access Switch boots, only accept DHCPv6 packets from uplinks

Pros:

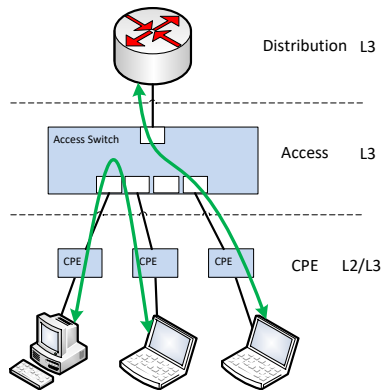
- No proprietary protocols.
- Very simple, no complicated L2/L3 protocol magic such as Proxy ARP, DHCP snooping, MACFF etc.

Cons

- All traffic between clients is exchanged at L3 in the default gateway
- Distribution need many logical L3 interfaces, one for each VLAN
- Wasteful on IPv4 address space, unless proprietary extensions are used. See section 3.1.2

SEC Access Certification

5.4 Sample 4, L3 in Access, client traffic exchanged on L3, separate L2 per customer port, no VLANs



Separate L2 broadcast domain for each customer port.

Access Switches are interconnected on L3.

Client exchange of traffic in Access: Yes

IPv4 and IPv6 protocol support.

IPv4 and IPv6 address configuration is done using DHCPv4 and DHCPv6.

Common

- Multicast injection filter
- Protection of control plane against flooding
- Protection of packet buffer starvation
- Drop fragmented IPv4/IPv6 packets to the control plane
- Loop-detect

IPv4

- OSPFv2 or ISIS between the Access Switches and Distribution router
- PIM-SM or hierarchical proxy IGMP for efficient multicast distribution
- IGMPv2/IGMPv3 to learn about interested receivers
- Limit number of concurrent multicast groups a customer port can join using the IGMP protocol
- DHCPv4 relay with option-82 for customer port identification
- Filter out spoofed IPv4 packets
- Drop IPv4 packets with too small fragment offset.
- When Access Switch boots, only accept DHCPv4 / BOOTP packets from uplinks

IPv6

- OSPFv3 or ISIS between the Access Switches and Distribution router
- PIM-SMv6 for efficient multicast distribution
- MLD/MLDv2 to learn about interested receivers
- Limit number of concurrent multicast groups a customer port can join using the MLD/MLDv2 protocol
- DHCPv6 relay with option 18 and 37 for customer port identification
- Filter out spoofed IPv6 packets
- Drop IPv6 packets with too small fragment offset.
- When Access Switch boots, only accept DHCPv6 packets from uplinks

Distribution router

- OSPFv2 + OSPFv3 or ISIS
- Protect against filling up ND cache

Pros:

- No proprietary protocols. High interoperability between switch vendors
- Very simple, no complicated L2/L3 protocol magic such as Proxy ARP, DHCP snooping, MACFF etc.
- Traffic always takes the shortest path since no paths are disabled (as Spanning Tree will do)

Cons

- Wasteful on IPv4 address space, unless proprietary extensions are used.

6 IPv6 Standards

6.1 General

What	Description
RFC1887	An Architecture for IPv6 Unicast Address Allocation
RFC1981	Path MTU Discovery for IP version 6
RFC2450	Proposed TLA and NLA Assignment Rule
RFC2460	Internet Protocol, Version 6 (IPv6) Specification
RFC 2462	IPv6 Stateless Address Auto configuration
RFC2464	Transmission of IPv6 Packets over Ethernet Networks
RFC2473	Generic Packet Tunneling in IPv6 Specification
RFC2526	Reserved IPv6 Subnet Anycast Addresses
RFC2675	IPv6 Jumbograms
RFC2711	IPv6 Router Alert Option
RFC3041	Privacy Extensions for Stateless Address Autoconfiguration in IPv6
RFC3122	Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification
RFC 3315	Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
RFC3531	A Flexible Method for Managing the Assignment of Bits of an IPv6 Address Block
RFC3587	IPv6 Global Unicast Address Format
RFC 3633	DHCPv6 Prefix Delegation
RFC3879	Deprecating Site Local Addresses
RFC4007	IPv6 Scoped Address Architecture
RFC4191	Default Router Preferences and More-Specific Routes
RFC4193	Unique Local IPv6 Unicast Addresses
RFC4291	IP Version 6 Addressing Architecture
RFC4292	IP Forwarding Table MIB
RFC4293	Management Information Base for the Internet Protocol (IP)
RFC4294	IPv6 Node Requirements
RFC4443	Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification
RFC4489	A Method for Generating Link-Scoped IPv6 Multicast Addresses
RFC4620	IPv6 Node Information Queries
RFC4779	ISP IPv6 Deployment Scenarios in Broadband Access Networks
RFC4861	Neighbor Discovery for IP version 6 (IPv6)
RFC4862	IPv6 Stateless Address Autoconfiguration
RFC4870	Recommendations for Filtering ICMPv6 Messages in Firewalls
RFC4941	Privacy Extensions for Stateless Address Autoconfiguration in IPv6
RFC4943	IPv6 Neighbor Discovery On-Link Assumption Considered Harmful
RFC5075	IPv6 Router Advertisement Flags Option
RFC5095	Deprecation of Type 0 Routing Headers in IPv6
RFC5156	Special-Use IPv6 Addresses
RFC5157	IPv6 Implications for Network Scanning
RFC5375	IPv6 Unicast Address Assignment Considerations
RFC5453	Reserved IPv6 Interface Identifiers

RFC5722	Handling of Overlapping IPv6 Fragments
RFC5871	IANA Allocation Guidelines for the IPv6 Routing Header
RFC5942	IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes
RFC6018	IPv4 and IPv6 Greynets
RFC6036	Emerging Service Provider Scenarios for IPv6 Deployment
RFC6104	Rogue IPv6 Router Advertisement Problem Statement
RFC6105	IPv6 Router Advertisement Guard
RFC6127	IPv4 Run-Out and IPv4-IPv6 Co-Existence Scenarios
RFC6164	Using 127-Bit IPv6 Prefixes on Inter-Router Links
RFC6177	IPv6 Address Assignment to End Sites
RFC6214	Adaptation of RFC 1149 for IPv6

6.2 Address Selection

What	Description
RFC3484	Default Address Selection for Internet Protocol version 6 (IPv6)
RFC5220	Problem Statement for Default Address Selection in Multi-Prefix Environments: Operational Issues of RFC 3484 Default Rules
RFC5221	Requirements for Address Selection Mechanisms

6.3 DHCP

What	Description
RFC3315	Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
RFC3769	Requirements for IPv6 Prefix Delegation
RFC6153	DHCPv4 and DHCPv6 Options for Access Network Discovery and Selection Function (ANDSF) Discovery
RFC6221	Lightweight DHCPv6 Relay Agent

6.4 DNS

What	Description
RFC2874	DNS Extensions to Support IPv6 Address Aggregation and Renumbering
RFC6106	IPv6 Router Advertisement Options for DNS Configuration

6.5 Multicast

What	Description
RFC2375	IPv6 Multicast Address Assignments
RFC2710	Multicast Listener Discovery (MLD) for IPv6
RFC3019	IP Version 6 Management Information Base for The Multicast Listener Discovery Protocol
RFC3306	Unicast-Prefix-based IPv6 Multicast Addresses
RFC3956	Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address

6.6 Transition

What	Description
RFC6144	Framework for IPv4/IPv6 Translation
RFC6145	IP/ICMP Translation Algorithm
RFC6146	Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers
RFC6147	DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers
RFC6156	Traversal Using Relays around NAT (TURN) Extension for IPv6
RFC6157	IPv6 Transition in the Session Initiation Protocol (SIP)
RFC6180	Guidelines for Using IPv6 Transition Mechanisms during IPv6 Deployment

6.7 CPE

What	Description
RFC6092	Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service
RFC6204	Basic Requirements for IPv6 Customer Edge Routers