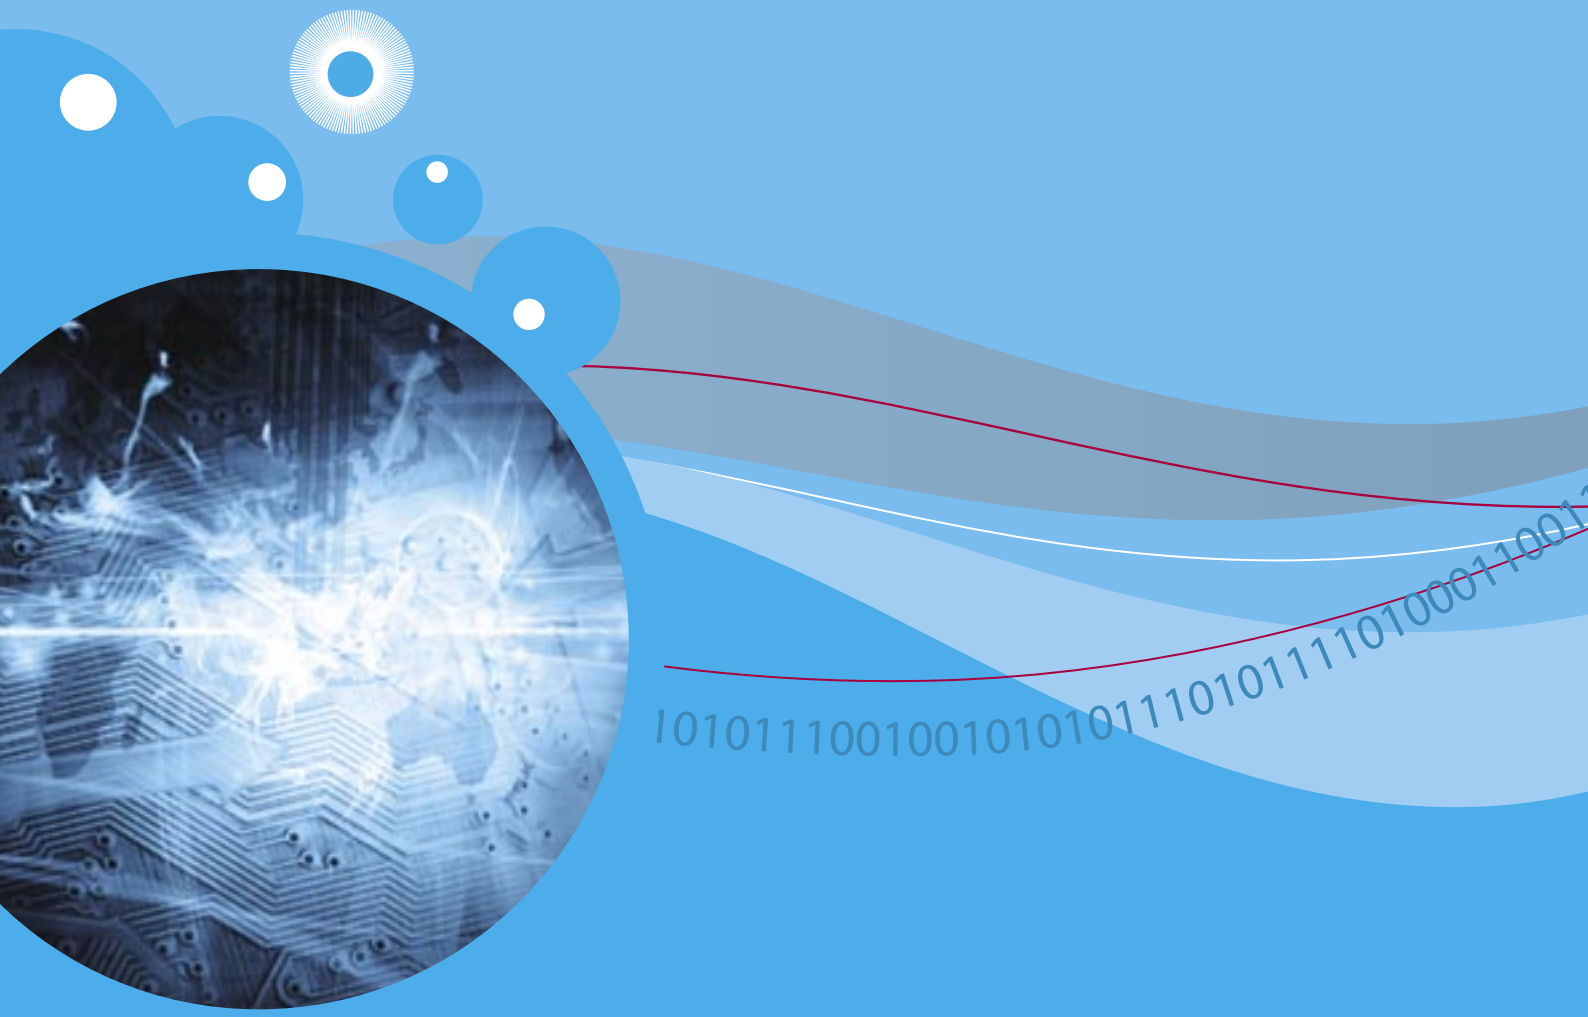

enisa
European Network and Information Security Agency

RESILIENCE FEATURES OF IPv6, DNSSEC AND MPLS

Resilience of communication networks

About ENISA

The European Network and Information Security Agency (ENISA) is an EU agency created to advance the functioning of the internal market. ENISA is a centre of excellence for the European Member States and European institutions in network and information security, giving advice and recommendations and acting as a switchboard of information for good practices. Moreover, the agency facilitates contacts between the European institutions, the Member States and private business and industry actors.

Contact details:

For contacting ENISA or for enquiries on this study, please use the following details:
Technical Department, Security Tools and Architectures Section
Email: sta@enisa.europa.eu
Internet: http://www.enisa.europa.eu/sta/

This study has been prepared by FORTH-ICS. *Authors:* Sotiris Ioannidis,George Apostolopoulos, Kostas Anagnostakis, Nikolaos Nikiforakis, Andreas Makridakis and Charalampos Gkikas with the collaboration of ENISA, STA staff.

Legal notice

Contents

1 Executive Summary

1 Executive Summary

ENISA is evaluating the effectiveness of three key technologies, namely IPv6, DNSSEC, and MPLS, in improving the resilience of public eCommunication networks. This study analyses the characteristics of those technologies and highlights their effect on the resilience of the network. It was found that each of these technologies have properties that can improve both the resilience and security of the Internet. Potential users, however, must also understand how exactly these technologies can be applied and what their limitations are before utilizing them. An overview of the characteristics of IPv6, DNSSEC, and MPLS is given, and the resilience assisting features, as well as other properties that one has to aware of to make an educated decision about their deployment are enumerated.

1.1 IPv6

The Internet Protocol version 6 (IPv6) is the next-generation protocol for the Internet. It is designed to be the successor of IPv4 for general use on the Internet and addresses a number IPv4's shortcomings. IPv6 provides among other functions, a significantly larger address space compared to IPv4, Quality of Service hooks and built-in security features for encryption and authentication of end-to-end communication.

In the 90's, the community started realizing the need for a next-generation protocol. Some of the motivating factors at the time included exhaustion of IPv4 address space, the non-hierarchical nature of address allocation, etc. Furthermore, IPv4 was not designed with security in mind: it was originally implemented in an isolated military network and later adopted by academia and industry. Network-layer security in IPv4 is retrofitted through higher-level protocols such as SSL, HTTPS and IPsec—all optional, meaning that one cannot count on their availability.

IPv6 addresses these shortcomings by offering a vast address space (128), mandatory support for network-layer security in the protocol stack, simplified packet headers, fixed length packet headers, stateless address auto-configuration, new multicast functionality, address scopes, extension headers, flow labels, IP mobility features and jumbograms.

In terms of resilience, IPv6 could address a noteworthy source of vulnerabilities. It is harder to launch opportunistic attacks such as worms against IPv6 hosts and makes reconnaissance probing much more difficult due to the vastness of the address space. The simplified packet headers and the lack of packet fragmentation make packet processing by routers easier and more robustly. The mobile IPv6 features allow for more efficient communication using Route Optimization. Finally, the mandatory implementation of IPSec in the protocol stack of IPv6 gives its users a guarantee that encryption and authentication functionality is always available when necessary, leading to a more resilient network.

1.2 DNSSEC

One of the most critical components of the Internet architecture is the Domain Name System (DNS). DNS is a distributed dynamic database with a hierarchical structure that maps names of machines to protocol-level addresses. DNS is a service that the vast majority of Internet users and services rely upon. The operation of popular Internet services such as e-mail, the web, or instant messaging, depend on it. The basic principle of DNS is the use of human-friendly domain names for Internet service addresses, as names are readable and memorisable, instead of numbers (Internet Protocol addresses), which are practical for computers only. DNS also performs the reverse operation, that is, translate an IP address to a fully qualified domain name. Originally the DNS design was focused on data availability and did not address any resilience or security issues. It is therefore possible to disrupt the operation of DNS by spoofing DNS messages, resulting to loss of integrity in Internet-based applications and services.

DNSSEC was developed to address these critical security shortcomings of DNS. It does so by defining a process

whereby a suitably configured name server can verify the authenticity and integrity of query results from a signed zone. DNSSEC uses public key cryptography and cryptographic hashes to enable a security-aware receiving name server to: (*i*) authenticate that the data received could only have originated from the requested zone, (*ii*) verify the integrity of the data, that is, data that was received at the querying name server was the data that was sent from the queried name server, and (*iii*) verify that if a negative response (NXDOMAIN) was received to a host query, the target record does not exist (denial of existence).

DNSSEC helps to eliminate a certain class of man-in-the-middle attacks, in which the attacker poisons the DNS cache to subvert the hostname to IP address mapping, and subsequently forcing the victim to connect to a malicious host. DNSSEC also protects against domain name hijacking, in which Internet domain names are temporarily "stolen" from the rightful registrants. This counters pharming attacks where hackers redirect web traffic to an attacker controlled website for distribution of malware, dissemination of false information etc. Other target attacks can also be countered, such as attacks that seek to hijack a mailserver or other critical services.

1.3 MPLS

Multi-Protocol Label Switching (MPLS) is a networking technology built around a label based forwarding paradigm. An MPLS header containing one or multiple labels (organized in a label stack) is attached to packets. Label Switch Routers (LSRs) forward these packets based only on the label information. Both Layer 2 (L2) and Layer 3 (L3) packets can be encapsulated in MPLS.

Typically MPLS is deployed in the form of a MPLS backbone or core network. All traffic inside the core network is forwarded using MPLS. Traffic that enters the MPLS core is labelled at the edge router, which also removes labels from the traffic that exits the MPLS core. Typically labels have only local scope. The label mappings in each LSR determine how packets will be forwarded through the MPLS backbone. By properly setting up the label mappings in all the LSRs, one can form a Label Switched Path (LSP) that will carry traffic over a specific path through the network independently of the network's native routing mechanisms.

Providers typically use MPLS to implement: (*i*) L2 point-to-point connections (pseudo-wires) that carry legacy traffic (ATM, Frame Relay) over a common backbone, (*ii*) Various types of Virtual Private Networks (VPNs) (L2 and L3 VPNs and VPLS), (*iii*) Traffic Engineering (TE) of the traffic inside the provider core networks, and (*iv*) improved resilience for provider core networks.

Using MPLS has significant impact on both the security and the resilience of the network and the services it supports. MPLS technology can be used to implement mechanisms that can quickly repair traffic when network failures or sudden variations in traffic patterns occur. More specifically, MPLS offers data protection by providing the ability to survive network link and network node failures, and by guaranteeing traffic repair in less than 50 milliseconds. MPLS enables traffic management and isolation by permitting headers to carry sufficient Quality of Service or Class of Service information. Finally, MPLS offers performance and security monitoring capabilities. It accomplishes this by providing mechanisms for operation, administration and management, for connectivity monitoring and fault isolation.

1.4 Conclusions

Following the analysis, it has been concluded that all three technologies are likely to help improve resilience to some degree, but some of the resilience improving features may be overstated by advocates. In some cases there are even important concerns about increased risks to resilience by using these technologies. With respect to deployment status, it is found that all three technologies have undergone extensive evaluation and trial deployments. However, some important issues may only be exposed with a global-scale deployment.

2 Introduction

2 Introduction

Reliable communications networks and services are critical to public welfare and economic stability. Intentional attacks to the Internet, disruptions due to physical phenomena, software and hardware failures, and human mistakes all affect the proper functioning of public communications networks. Such disruptions reveal the increased dependency of our society on these networks and their services. Experience shows that neither single providers nor a country alone can effectively detect, prevent and effectively respond to such threats.

The European Network and Information Security Agency (ENISA), fully recognizing this problem, devised a Multi-annual Thematic Program (MTP 1) with the ultimate objective to collectively evaluate and improve the resiliency of public eCommunications in Europe[1]. In 2008, one of ENISA's activities in achieving this goal was the assessment of the effectiveness of three key technologies, namely IPv6, DNSSEC, and MPLS, in improving the resilience of those networks.

Resilient are characterised the networks that provide and maintain an acceptable level of service in face of faults (unintentional, intentional, or naturally caused) affecting their normal operation. The main aim of the resilience is for faults to be invisible to users.

Improving the resilience of a network is an issue of risk management which includes risk identification, evaluation and acceptance or mitigation. A wide accepted list of risks to the resilience of networks includes flash crowd events, cyber attacks, outages of other support services, natural disasters and system failings. The mitigation of identified risks involves technical measures such as resilient design, resilient transmission media, resilient equipment and technologies that improve resilience.

The main objective of this document is to highlight the resilience features of IPv6, DNSSEC and MPLS. These are features of the three technologies that when used in specific configurations will improve the resilience of the network. A resilient network can provide and maintain an acceptable level of service in face of events affecting normal operation. Such events include flash crowds, cyber attacks, outages to other services affecting the network, and natural disasters.

In the following three chapters the features of the above technologies that improve network resilience and security are listed and analyzed. Features that may not be as helpful but should be understood by parties interested in these technologies are also discussed. Finally, for each case, examples of existing deployments are presented as well as scenarios where IPv6, DNSSEC, and MPLS have been used in production environments.

[1] http://www.enisa.europa.eu/doc/pdf/management_board/decisions/enisa_wp_2008.pdf

3 IPv6

3 IPV6

The Internet Protocol (IP) provides the core interoperability layer that enables heterogeneous systems to communicate over the Internet. Each network end-point is assigned a unique address and a communication protocol is used to transfer data from one system to another. The number of such unique addresses in the currently deployed protocol (IPv4—Internet Protocol version 4) that are available for new users is slowly being exhausted. As a result, addresses have become a precious resource due to their scarcity and different engineering techniques like Network Address Translation (NAT) are used to share a unique address among several computers. Such ad-hoc mechanisms hamper swift access to data and services. Therefore, expansion of the address space can have a direct impact on the proliferation and accessibility of a wide variety of network-enabled devices.

The Internet Protocol version 6 (IPv6) is the next generation protocol for the Internet. It is the designated successor of IPv4 for general use on the Internet since it was designed to remove the perceived barriers and shortcomings of IPv4 and provide a feature-rich environment for the future of the Internet.

IPv6 provides among other functions, a significantly larger address space compared to IPv4, Quality of Service hooks, and built-in security features for encryption and authentication of end-to-end communication [17].

3.1 Necessity of IPv6

In the 90's, the community started sensing IPv4's inadequacy to handle the growth of the Internet [43], triggering the search for a successor. The candidate proposals at that time were motivated by the following factors:

- Exhaustion of IP Class B address space
- Exhaustion of IP address space in general
- Non-hierarchical nature of address allocation leading to flat routing space

Medium-term remedies for these problems include CIDR (Classless Inter-Domain Routing [66]) which permits the aggregation of Class C networks for routing purposes, as well as assignment policies that allocate Class C network numbers in a fashion that CIDR can take advantage of. Routing protocols supporting CIDR include OSPF and BGP4 [74]. None of these were pre-requisites for the new IP, but were necessary for prolonging the life of the current Internet long enough to work on longer-term solutions. Several other techniques (NAT/PAT—network/port address translation) were adopted as the number of unique IP addresses decreased and the general concept of IP, i.e., direct end-to-end connectivity, compromised. Management and maintenance of IP infrastructure has become complex and inefficient. IPv6 advocates suggest that this has direct consequences to the productivity and growth of today's economy which is so tightly integrated with Internet.

Furthermore, IPv4 was not designed with security in mind: it was originally implemented in an isolated military network and later adapted for public educational and research purposes. Network-layer security in IPv4 is retrofitted through higher-level protocols such as SSL, HTTPS, and IPsec—all optional, meaning that one cannot count on their availability. Modern applications could benefit greatly if the underlying network guarantees features like on-time delivery, availability of bandwidth, and security. If such guarantees are possible, there can be another wave of next generation applications. Retrofitting security and guaranteed service features into IPv4 has high overhead and there is an engineering limit to the amount of retrofitting that can be applied to IPv4. These problems cannot be ignored for long. A natural evolution from IPv4 was required that was designed with extensibility and scalability in mind. In 1998, the IPv6 draft standard [42] came into existence. A brief summary of the important features is provided, noting that some of these features were already present in IPv4 or were retrofitted later.

3.2 Key Features of IPv6

- **Vast Address Space** - unique addresses instead of of IPv4
- **Mandatory Support for Network Layer Security** - Support for IPsec is mandatory in IPv6 and therefore, the IPsec security model is required for all IPv6 implementations in near future. In IPv6, IPsec is implemented using the AH authentication header and the ESP extension header. [56]
- **Simplified Packet header** - The designers of IPv6 decided to simplify the format of the IPv6 packet header in order to minimize processing at the intermediate routers between two hosts in the Internet. They simplified the header by removing certain fields from the IPv4 packet header and replace their functionality through chained extension headers. The exact changes can be seen in Figure 3.1.
- **Fixed size packet header** - The main reason of its introductionwas the will of maximizing processing performance - simple constant size headers can be processed quickly, at or very close to wire-speed.
- **Stateless address auto-configuration** - IPv6 hosts can configure themselves automatically without the need of a DHCP server when connected to a routed IPv6 network only using router discovery messages.
- **Multicast** - New multicast functionality is defined in IPv6 where a host has the ability to send a single packet to multiple destinations. Multicast, in addition to other possible use, will replace the "Broadcast" addresses from the IPv4 specification.
- **Address Types** - IPv6 defines three types of addresses:
  - **Unicast** (individual) - identifies one single network interface (typically a computer or similar device). The packet is delivered to this individual interface.
  - **Multicast** (group) - identifies group of interfaces. Data must be delivered to all group members.
  - **Anycast** (selective) - also identifies a group of network interfaces. But this time the packet is delivered just to one single member of the group (to the nearest one).
- **Address Scopes** - Address scopes is also a new concept in IPv6. It defines the region where an address can be used as a unique identifier of an interface. The scoped addresses can be: link-local, site-local, unique local unicast and global addresses. [50, 49]
- **Routing Header** - Along with the extended authentication and encryption headers, IPv6 introduces a new one, the Routing header. It allows the sender to define a sequence of "checkpoints" (IPv6 addresses), through which the datagram must be routed on its way to the final destination.
- **Flow Labels** - IPv6 can provide QoS for applications that need it (e.g., VoIP) through the use of the "Flow Label" field in the IPv6 packet header.
- **IP mobility** - This feature ensures transport layer connection survivability and allows a host to remain reachable regardless of its location in an IPv6 network. With the help of Mobile IPv6, even when the mobile node changes location and address, its existing connections are maintained [52].
- **Jumbograms** - IPv4 limits packets to 64 KB of payload. IPv6 has optional support for packets over this limit, referred to as jumbograms, which can be as large as 4 GB. The use of jumbograms may improve performance over high-MTU networks. The presence of jumbograms is indicated by the Jumbo Payload Option header.

Figure 3.1 – **Comparison of IPv4 and IPv6 packet headers.**

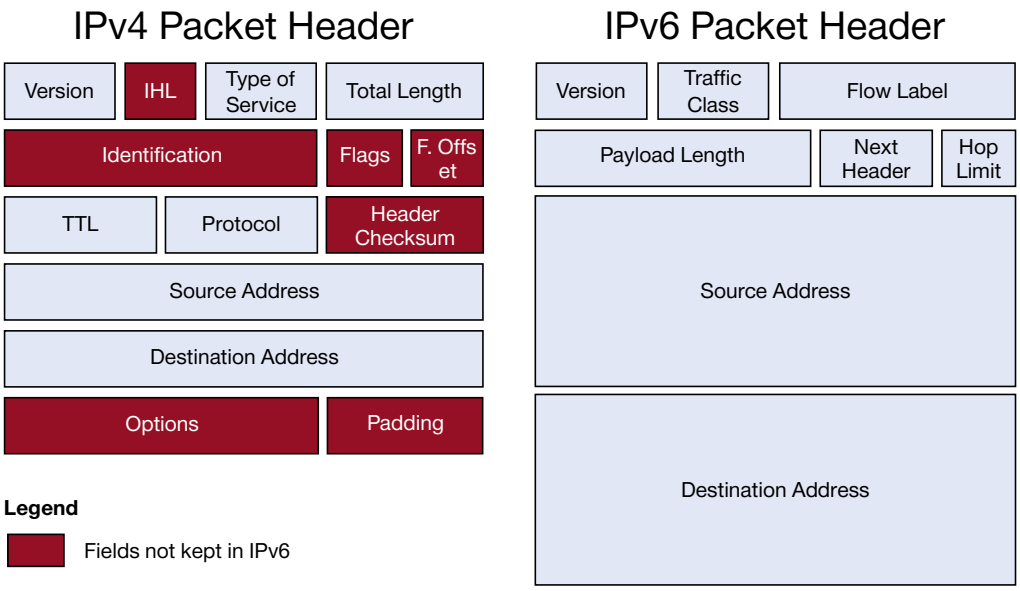

3.3 Analysis of IPv6 Features

Although IPv6 provides many notable features, such as QoS and IP Mobility, only the ones that are related to resilience and security are analysed.

3.3.1 IPv6 Address Space

IPv6 addresses tackle the main problem of IPv4, i.e., the exhaustion of available IP addresses for connecting new computers to the Internet. IPv6 has a significantly larger address space compared to IPv4, allowing the needed freedom and flexibility in the allocation of the available addresses, as well as efficient traffic routing. Furthermore, since IPv6 provides an abundance of IP addresses, it implicitly eliminates the need for NAT (Network Address Translation), which has been the cause of several end-to-end networking problems. IPv6 also simplifies several aspects of address assignment and renumbering.

In more details, the packet header of IPv6 provides 128 bits for addressing, when IPv4 provides only 32 bits. This very large IPv6 address space supports (about) addresses instead of the (about) unique addresses of the IPv4. This means that IPv6 supports more than ten billion billion billion times as many addresses as IPv4. The size of each subnet in IPv6 is bits, which is actually the square of the entire IPv4 address space. This design choice will probably result in sparse address space utilization, but more efficient routing and network management.

While the address space of IPv6 seems to be larger than we would ever need, it was not the intent of the designers of IPv6 to assure geographical saturation with usable addresses. This large address space allows a better and more systematic hierarchical allocation of addresses and more efficient route aggregation, substituting IPv4's complex Classless Inter-Domain Routing (CIDR) techniques that were developed to make the best use of the small address space.

Another advantage of IPv6 regarding addressing is the renumbering process. Renumbering an existing IPv4 network with different routing prefixes is considered a major effort. With IPv6, however, an entire ad-hoc network can be renumbered just by changing the prefix of a few routers thanks to the inner-workings of IPv6's addressing protocol. [42]

3.3.2 Mandatory Support for IPsec

Although IPsec is not an integral part of IPv4, it is a mandatory component for IPv6. The IPsec security model is required to be supported for all IPv6 implementations in the near future. IPsec is a framework of open standards that define policies and ways to implement them, in order to achieve secure communication within a network.

Computers using IPsec can provide data confidentiality, data integrity and data authentication at the network layer of the OSI model.

The main purpose of IPsec is to assure interoperable, high quality, cryptographic security for IPv4 and IPv6. It provides its security services at the IP layer, and therefore, offers protection at both the network as well as higher layers. Its security features include encryption (confidentiality), data origin (authentication), protection against message replay and access control.

IPsec has two different modes: Transport mode (host-to-host) and Tunnel Mode (Gateway-to-Gateway or Gateway-to-host).

- **Transport mode** - In Transport mode, the payload is encapsulated (header is left intact) and the end-host (to which the IP packet is addressed) decapsulates the packet.
- **Tunnel mode** - In Tunnel mode, the IP packet is entirely encapsulated (with a new header). The host (or gateway), specified in the new IP header, is responsible for decapsulating the packet.

3.3.3 Supported IPsec Extension Headers in IPv6

In IPv6, IPsec is implemented using the AH authentication header and the ESP extension header. ESP encrypts the data carried by IP, such as a TCP packet, using a defined algorithm and cryptographic key. The output is the ciphertext that is difficult to decode without the correct key. The receiving IPsec ESP entity uses an associated decryption algorithm and the same key to extract the original data.

The IPsec Authentication Header (AH) provides integrity and authentication but no privacy, since the IP data is not encrypted. The AH contains an authentication value based on a symmetric-key hash function. The combination of both mechanisms, as shown in Figure 3.2, provides integrity, authentication, and privacy.

Figure 3.2 – **Nested ESP with AH in an IPv6 packet.**

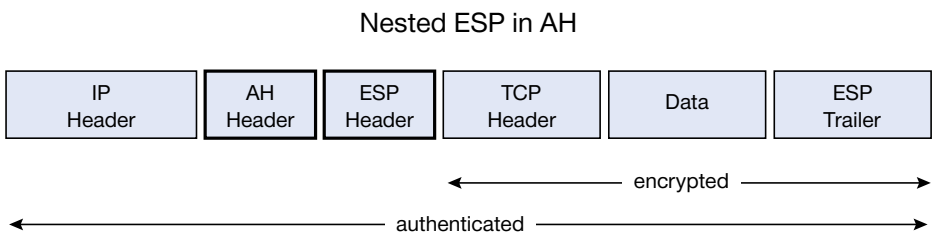

3.4 IPv6 Resilience and Security

3.4.1 Ubiquitous Connectivity

A lot has been written about the future Internet, where many electric appliances will be connected on the global network. Several companies have presented prototypes for refrigerators that will automatically order new groceries when supplies run low; chandeliers which will sense the burning of a lamp and order new ones over the Internet; television sets which will receive content directly from the Internet (IPTV) and automobiles which will use the Internet to download roadmaps and weather reports [32, 30, 16].

Such features will be more easily implementable with the use of IPv6. The abundance of IP addresses gives the freedom to connect many more embedded devices to the Internet. While this "ubiquitous connectivity" will provide us with numerous new capabilities, at the same time it opens the door to attacks that were traditionally targeting only computers.

Until recently, exposed systems that could be attacked over the Internet were primarily servers, personal computers, routers and switches. Until the discovery of the first cellphone worm, no-one expected that his cellphone could be infected and data could be stolen, just by having his bluetooth turned on and accepting connections [79]. Even more recently, an Internet-connected coffee machine had a bug and could be exploited remotely [14]. One can only imagine what can happen when a fridge, a TV set, or a car with an unpatched software vulnerability is connected on the public Internet.

The following scenarios may at first seem hard to believe, but, considering the ubiquity of software bugs and vulnerabilities, should be considered as part of the emerging threats landscape:

- **Fridge hacking** - Attackers gain access on an IPv6-ready fridge and if they can't steal the credit card details, they can empty it by ordering lots of unwanted products.
- **Car hacking** - Attackers take control of an IPv6-ready car and change the roadmaps, leading the passengers to an unknown location instead of the desired one.

While the above examples are not meant to scare, they warn us that with greater connectivity comes greater risk. Manufacturers of all sorts of Internet-connected devices should take every conceivable measure in order to write bug-free code, while customers should think twice before replacing their current electric and electronic equipment with Internet-enabled devices.

3.4.2 To NAT or not to NAT?

Network Address Translation (NAT) is a technique that allows the translation of local and internal network addresses (used within an organization) into global IP addresses which identify an online resource uniquely over the Internet. NAT allows multiple resources within an organization or those connected to a local LAN to use a single IP address to access the Internet.

NAT became popular because it alleviated the IPv4 address shortage. IPv6 eliminates the need for Network Address Translation by offering a much larger address space that allows the network resources to have their own unique public IP address.

While NAT was originally a pragmatic fix to the address space shortage, IPv6 is now technically competing with NAT. IPv6 supporters claim that since the original purpose of NAT was to provide more IP addresses, it is no longer

needed. Also, NAT breaks the end-to-end paradigm and many applications (like VoIP) have gone to considerable lengths in order to provide service over NAT-ed networks. In addition, NAT is one more thing that can be misconfigured or one more device that can fail.

On the other hand, NAT advocates argue that NAT provides security along with address translation. By hiding the presence of internal hosts and by making them unreachable from the rest of the Internet, NAT essentially acts as a firewall protecting them from attacks. With IPv6, every host will be globally accessible and the various networks will have to deploy "proper" firewalls in order to protect their "internal" hosts. This firewall deployment can lead to misconfigurations and security breaches which would not have happened if NAT was in place.

While both NAT naysayers and NAT advocates have valid points, one can argue that security through obscurity (or security as a side-effect) is not a good choice. By hiding a host via NAT, one does not necessarily make it more secure. If an attacker manages to penetrate the perimeter and get access to a local machine then NAT can't protect it anymore. Furthermore, with the use of UPNP (Universal Plug and Play) feature on NAT boxes, hosts already claim ports on the public IP, so the protection is not complete. If a company uses multiple firewalls with carefully crafted and restrictive policies then even a breach like the above could be mitigated. Replacement of implicit firewalling as done by NAT with proper firewalls will be a challenge, and in the transition process may introduce risks, but it is not expected to be a major issue in the long run.

3.4.3 Stateless Auto-Configuration

Stateless Auto-Configuration allows the various devices attached to an IPv6 network to connect to the Internet without requiring any intermediate IP support in the form of a DHCP (Dynamic Host Configuration Protocol) server. A DHCP server holds a pool of IP addresses that are dynamically assigned for a specified amount of time to requesting nodes in a Local Area Network. In contrast, Stateless Auto Configuration is designed to work without the use of a centralized server. Each network device, once connected to the Internet, will generate its own valid IPv6 globally unique address by combining its device identifier with the network prefix broadcasted by the router. The protocol requires testing whether the address is unique by probing the network with possible addresses. Some advantages of Auto Configuration are that it does not require any dedicated protocol server, it allows hot plugging of network devices, and it is cost effective.

While Stateless-Auto Configuration provides many good features it also raises several red flags when it comes to privacy and security.

3.4.3.1 Auto-Configuration Privacy

An 128-bit IPv6 address is made out of the concatenation of the 64-bit network identifier part with the 64-bit host identifier part. The host identifier part of the IPv6 address consists of four segments containing four hexadecimal numbers. This part is the address used to identify the host. The interface ID is a 64-bit (four-segment) Extended Unique Identifier (EUI-64) generated based on the Media Access Control (MAC) address of a network device, as shown in Figure 3.3.

Figure 3.3 – **Creation of an EUI-64 address from a MAC address.**

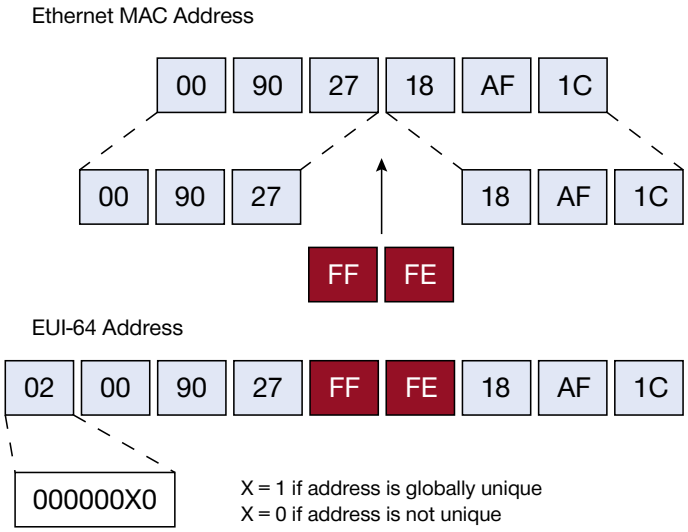

Let's consider a user that connects from three or four different networks daily. Each network will provide the Stateless Auto-Configuration protocol with a unique network interface part, but the host identifier part will remain the same since the user's MAC address does not change. This can result in privacy issues, since an attacker can extract the MAC address from the user's IPv6 addresses and correlate his traffic from all networks. In addition, since most laptops are not custom made, a large company can correlate traffic even further and pin down an IP address to the specific person who bought a specific laptop model.

One solution to this problem is to randomize the 48-bit part of the EUI-64 so that a different host identifier part is generated every time. The problem with this solution is that a network administrator cannot really enforce this rule since the protocol creates the host-identifier part without consulting any network service. Therefore, a randomization solution should be built in the actual IPv6 protocol.

3.4.3.2 Auto-Configuration Security

Stateless Auto-Configuration uses the Neighbor Discovery Protocol in order to perform various functions. An interesting function is the Duplicate Address Detection (DAD). When an interface is initialized or reinitialized, it uses auto-configuration to tentatively associate a link-local address with that interface (the address is not yet assigned to the interface in the traditional sense). At this point, the interface joins the all-nodes and solicited-nodes multicast groups, and sends a neighbour discovery message to these groups through DAD. By using the multicast address, the node can determine whether that particular link-local address has been previously assigned, and choose an alternative address in that case. If a response comes from the network before a predefined timeout, this means that the tested address is already in use.Auto-configuration will choose another address and try again. This process is repeated until the host finds an address that is not used by anyone.

The obvious problem with Stateless Auto-Configuration is that the host relies on the other hosts' good behavior. If a malicious host is already connected to the internal network (link local scope), it can generate a DoS attack by answering "I am already using that address" to all the host's address requests, effectively preventing it from obtaining a valid IPv6 address.

3.4.4 Host Reconnaissance on an IPv6 Internet

If the entire IPv4 address space could be used for addressing (ignoring reserved blocks, private addresses, multicast groups, and so on), we could assign a unique IP address to 4,294,967,296 computers. The fastest Internet worm until now has scanned the entire IPv4 Internet in less than 10 minutes [72]. Even if a web server is not listed on any search engine and no other website points to any of the served pages, this server is scanned several times a day by Internet worms, botnet clients, and always-hoping script kiddies.

In IPv6 though, things are different. The address space is so vast that months or even years could pass by before an attacker randomly finds a specific host. Does this "invisibility" mean that we should stop caring about securing our computers and networks? Randomly finding one vulnerable host among can be very difficult, but it still can be done. Researchers have pointed out that random scanning is not the only way to find vulnerable hosts. A worm could crawl search engines and attempt to attack every host that ends up in the search results. It could also join a P2P file-sharing system and find even more available hosts, or even snoop on the local network for new potential victim addresses once it manages to infect a single host [70]. This means that an attacker could eventually find any reachable host and exploit its vulnerabilities. It is thus important to realize that the vast IPv6 address space provides obscurity but not security.

Network administrators should patch their machines and setup firewall policies as they always did. In addition, they can use the IPv6 address space to their advantage by choosing difficult-to-guess IP addresses. This will result in less scanning and attack attempts per day, reducing this way the background noise in the administered network.

3.4.5 Mobile IP

Mobile IP allows an IP node to change its location and address on an IP network while maintaining existing connections. When a node which doesn't support Mobile IP changes its location and its IP address, the existing connections of the node that are using the address assigned from the previously connected link cannot be maintained and are thus terminated.

In Mobile IP the node can still change location and IP address but the existing connections through which the mobile node is communicating are maintained. To accomplish this, connections to mobile nodes are made with specific addresses that are always assigned only to mobile nodes and through which the mobile nodes are always reachable.

A basic Mobile IP architecture consists of:

- **Mobile Node** - An IPv6 node (laptop/cellphone/pda) which changes networks
- **Home Agent** - A router on the home location of the mobile node which maintains the registrations and bindings of mobile nodes and foreign addresses
- **Correspondent Node** - An IPv6 node which wants to communicate with the mobile node

In Mobile IPv4, all communications between the mobile and correspondent node are tunneled through the Home Agent, creating a "triangle-routing" (a.k.a asymmetric routing) scheme which is inefficient especially when the mobile node is far away (large number of network hops) from his Home Agent and the Correspondent Node is far away from the Home Agent. The routing path created for this communication is far from the optimal. ([51] – red path in Figure 3.4)

In Mobile IPv6 though, a new kind of routing is introduced called Route Optimization. While the protocol of Mobile IPv6 starts the same way as the one in Mobile IPv4, after the authentication is performed (through the use of the ESP and AH extensions), the mobile and correspondent nodes stop using the Home Agent and start

communicating directly. In addition, the protocol stacks at each end re-write the sender and receiver fields in such a way that the applications running on top of both nodes think that the communication still exists between the Correspondent Node and the Mobile Nodes' Home Agent. In this way, IPv6 achieves the same mobility as IPv4 but the communication is direct, eliminating the large latency of Mobile IPv4 (green path in Figure 3.4).

Figure 3.4 – **Route comparison of Mobile IPv4 and Mobile IPv6**

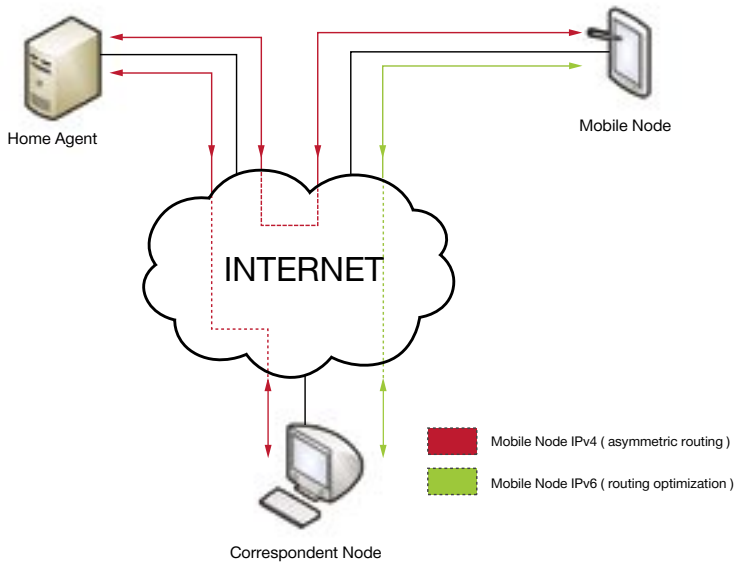

3.4.6 Additional Points

- **IPsec** - Provided that a valid key-infrastructure will be used (like IKE, IPsec Key Exchange), IPsec will provide the necessary protection for end-to-end IPv6 communications. With the use of the Authentication Header and Encryption Header, IPv6 guarantees authentication, integrity, and privacy of data [33].
- **Limited Monitoring** - Given the vastness of the IPv6 address space, chances that a malicious user will attack a honeypot, or be seen through a network telescope are much slimmer than in IPv4. This could degrade security research unless the population of honeypots and telescopes increases and the existing ones are re-configured to use multiple IPv6 addresses.
- **Fragmentation** - In IPv4, routers have the ability to split an incoming packet into smaller units called fragments so that they can "fit" outgoing interfaces' MTU. This is crucial for successfully sending traffic across different types of networks with different capacities and rules. After the split, the routers mark the appropriate fragmentation options in the IPv4 headers of the resulting fragments and then send them to their destinations. The host that receives the IP fragments re-arranges them as needed in order to create the original IP packet. Attackers have used fragmentation to conduct several attacks, mostly DoS attacks and firewall circumvention [55]. In IPv6, the routers no longer split the incoming packets into fragments. The fragmentation can be done only by the sending host which performs a path MTU discovery and fragments its packets accordingly, using the fragmentation extension headers. While this policy of IPv6 won't reduce potential fragmentation attacks (since the fragmentation can still be done by the host), it will certainly eliminate the overhead of fragmentation from routers. In essence this means that the routers will become faster since they will have the ability to process more packets per second without having to worry about fragmentation and MTU discovery.

- **Misuse of IPv6** - IPv6 is not considered "experimental" anymore and thus it has already been implemented in various network devices. It has been observed that attackers use IPv6 in order to hide their presence in IPv4 networks, as well as for evading firewall policies using the Teredo tunneling protocol. Teredo is designed to grant IPv6 connectivity to nodes that are located behind IPv6-unaware NAT devices. Specifically, if a network administrator finds certain IPv4 packets (Teredo packets) or IPv6 packets (auto-configuration packets, router and neighbour discovery, solicitation and advertisement packets) in the IPv4 network, this is an indication of possible mis-configuration, rogue tunnels and routers, or malicious activity [76].
- **Address Spoofing** - The problem of a local host acting as a gateway and performing a Man-in-the-Middle attack still exists. Especially during the Auto-configuration phase, in which a host requests address and gateway information from any router in the local network, IPsec is not yet enabled so anyone can respond with fake information. Furthermore, Duplicate Address Detection protocol is flawed since it relies on end-hosts for answering whether an address is already in use. A non-IPsec solution called SEND (SEcure Neighbour Discovery) has been proposed to address the above security issues but it is currently not implemented by any vendor due to some intellectual property rights concerns [62, 54].
- **Application Layer Attacks** - Even with IPsec, the majority of vulnerabilities on the Internet today are at the application layer (e.g., SQL Injection, Cross-site scripting (XSS), Cross-site request forgery (CSRF)), which cannot be tackled at the network layer, in which IPv6 operates.
- **Rogue Devices** - Rogue devices will be as easy to insert into an IPv6 network as they were in IPv4. Even if DHCP is not used, through IPv6's auto-configuration they will create a 128-bit IP address and proceed as normal.
- **Denial of Service** - With the exception of the aforementioned DAD DoS attack, Denial of Service (DoS) attacks are identical between IPv4 and IPv6. What is interesting to point out is that since more hosts will be globally addressable than they were before, we could see DoS attacks to additional categories of targets like TV sets and refrigerators connected to the Internet.

3.5 IPv6 Deployment

We can measure the percentage of networks running IPv6 by comparing the set of ASes (Autonomous Systems) in the IPv6 routing table to those in the combined set of IPv4 and IPv6 [21] tables. As of November 13, 2008, 3.9% out of a total of 30090 ASes have deployed IPv6.

Table 3.1 – **Autonomous Systems running IPv6. Routing Information Base data as of 13 November 2008.**

IPv4 ASes	**: 30043**
IPv6 ASes	**: 1160**
ASes using only IPv4	**: 28930**
ASes using only IPv6	**: 47**
ASes using IPv4 and IPv6	**: 1113**
ASes using IPv4 or IPv6	**: 30090**

According to a report [46] prepared by RIT for NIST, Internet users incur approximately 90% of the IPv6 transition costs (estimated to USD 25 billion, 2003) and the rest by vendors and ISPs. Although these cost estimates seem large, they are actually small relatively to the overall expected expenditures on IT hardware and software, and even smaller relatively to the expected value of potential market applications. Although the transition from IPv4 to IPv6 is slow, it is expected to gradually increase as the cost-benefit ratio tilts enough towards IPv6 over the time. In the following, we discuss the IPv4 to IPv6 transition mechanisms and few reasons behind current deployment gaps.

Much work has gone into developing standardized IPv6 transition mechanisms to ease the shift from IPv4 to IPv6. SIT ("Simple Internet Transition" or "Six In Tunnel"), 6to4 automatic SIT tunnels, and IPv6 over UDP are common examples of these technologies. These transition mechanisms couple with well-connected and easily available tunnel brokers to make IPv6 readily available to anyone with an IPv4 address, regardless of whether IPv6 is supported on any given network.

3.5.1 6NET

6NET was a three-year European IST project with the aim to demonstrate that continued growth of the Internet can be met using the new IPv6 technology. The project built and operated a pan-European native IPv6 network connecting sixteen countries in order to gain experience with IPv6 deployment and the migration from existing IPv4-based networks. The network was used to extensively test a variety of new IPv6 services and applications, as well as its interoperability with legacy applications. This allowed practical operational experience to be gained, and provided the possibility to test migration strategies, which are important considering that the IPv4 and IPv6 technologies will need to coexist for several years.

6NET involved thirty-five partners from the commercial, research, and academic sectors and represented a total investment of EUR 18 million, 7 million of which came from the project partners themselves, and 11 million from the Information Society Technologies Programme of the European Commission. The project commenced on 1st January 2002 and was due to finish on 31 December 2004. However, its success lead to a six month extension, primarily for dissemination of the findings and recommendations. The network itself was decommissioned in January 2005, handing over the reigns of pan-European native IPv6 connectivity to GEANT.

The principal objectives of the project were:

- Install and operate an international pilot IPv6 network with both static and mobile components in order to gain a better understanding of IPv6 deployment issues.
- Test the migration strategies for integrating IPv6 networks with the existing IPv4 infrastructure.
- Introduce and test new IPv6 services and applications, as well as legacy services and applications on IPv6 infrastructure.
- Evaluate address allocation, routing, and DNS operation for IPv6 networks.
- Collaborate with other IPv6 activities and standardization bodies.
- Promote the IPv6 technology.

6NET used its 3-year experience to publish several papers in conferences as well as whitepapers and guides on a wide range of IPv6-related topics [19].

3.5.2 Euro6IX

Euro6IX [13] is the largest research project up to now funded by the European IST Program (IST-2001-32161). The goal of the Euro6IX project is to support the rapid introduction of IPv6 in Europe. Its objectives are:

- Design an appropriate architecture for the deployment of the first Pan-European non-commercial IPv6 Internet Exchange (IX) Network. It will connect several regional neutral IPv6 Internet Exchange points across Europe, and achieve the same level of robustness and service quality as currently offered by IPv4 Internet Exchange Networks.
- Use the deployed IPv6 IX infrastructure to research, test, and validate IPv6-based applications and services.
- The network built within the Euro6IX project will be open to specific user groups (existing and to be created),

that will be connected to the Euro6IX network through a variety of access technologies such as mobile, xDSL, and cable, and will be connected to legacy IPv4 networks and services. This will allow to test the performance of future IPv6 networks, as well as non-commercial native IPv6 advanced services and applications. The network's Acceptable Use Policy (AUP) excludes the possibility of carrying commercial traffic. The network will be used and tested by the user groups to validate and assess the feasibility, features, and potential of the Next Generation Internet through daily routine use of the services, internal trials, and also in highly visible events and public trials.

- Dissemination, liaison, and coordination with clusters, fora, standards organizations (e.g., the IETF and RIPE) and third parties, with particular consideration for interworking and coordination with peer projects, such as GEANT, 6WINIT, LONG, MIND, 6NET and any other projects related to Euro6IX, that might be available during the project's lifetime.

3.5.3 Summer Olympics 2008

The 2008 Summer Olympic Games were a notable event in terms of IPv6 deployment. For the first time, a major World event had a presence on the IPv6 Internet [31] and all network operations of the Games were conducted using IPv6. Beijing 2008 Olympics provided the largest public showcase of IPv6 technology since the inception of IPv6.

3.5.4 Some IPv6 Projects per Country

Although the following list is a partial list of the various deployments, conferences and research groups related to IPv6, we can easily discern that there is a global movement of IPv6 supporters. Initiatives, task forces and research projects push towards the wide-spread acceptance of IPv6 as the only way for the further expansion of the Internet.

- USA
  - The US Government has issued a mandate to all vendors—both civilian and defense—to make the switch to an IPv6 platform by summer of 2008.
  - Lots of states have assembled their own IPv6 Task Forces for educating and promoting IPv6.
  - Internet2 has created an IPv6 Working Group focused both on understanding how IPv6 will enable Internet2 to achieve its goals and on promoting and coordinating the deployment of IPv6 throughout the Internet2 infrastructure.
  - US IPv6 Summit is the largest IPv6 conference in America that provides the latest status of US Government and Industry progress for IPv6.
- Canada
  - Viaginie of Canada has developed a tunnel server, the freenet6.net to allow any IPv4 node to be connected to the 6Bone.
  - DNS root server IPv6 accepts DNS requests through IPv6.
- Great Britain (UK)
  - Bermuda 2 is a joint project between the UK Universities of Southampton, UCL, and Lancaster. Its aim is to study and report on IPv6 deployment issues in collaboration with Internet 2 partner sites which include ISI and the CRC.
  - The School of Electronics & Computer Science (ECS) at the University of Southampton is working with Virgin Radio to make them the first commercial radio station in the UK to be accessible over IPv6—currently relaying four of the official Virgin Radio UK feeds available over unicast IPv6.

- Japan
  - JGN (Japan Gigabit Network) - IPv6 over ATM and Native IPv6 transport (no tunnels).
  - NTT Communications launched the first commercial IPv6 service in 2001, started DUAL services for ADSL users and a world-wide transport service.
  - NSPIXP-6, IPv6-based Internet Exchange in Tokyo , 1999–2008.
  - KAME project was a joint effort of six companies in Japan to provide a free stack of IPv6, IPsec, and Mobile IPv6 for BSD variants.
  - Major router vendors (Hitachi, Fujitsu, NEC, Furukawa Electric, Yamaha etc) are 'IPv6-ready.'
  - Service providers like Powered Com, Japan Telecom, KDDI have started trials in areas like mobile phones, online gaming, Internet Car/Train and medicine.
  - IPv6style.jp is a Japanese information web-site for people to learn, build and use IPv6.
  - The TAHI Project is the joint effort formed with the objective of developing and providing the verification technology for IPv6.
- China
  - China's Next Generation Internet Project (CNGI) is a five-year plan with the objective of cornering a significant proportion of the Internet space by implementing IPv6 early.
  - Summer Olympics 2008 - All network connections from security cameras to vehicles and the coverage of the event was implemented over IPv6.
- France
  - IPv6 Task Force created in France on 2002. Active involvement of France Telecom.
  - Deployment of "Open Transit," a native IPv6 international commercial network, 2002.
  - France Telecom was assigned a larger IPv6 prefix (/19), 2005.
  - French ISP free.fr gives DSL customers IPv6 addresses customers since 2007.
- Netherlands
  - SURFNet (http://www.ipv6.surfnet.nl/) created an IPv6 Playground.
- Korea
  - KOREAv6, composed with IPv6 Trial Services and Field Test for IPv6 Equipments, is the IPv6 Pilot Project launched in Korea.
  - The government of Korea plans to achieve complete IPv6 Transition in Public Sector and 10M IPv6 users by 2010: total IPv6 Transition in Backbone network by 2010 and in access network by 2013 for ISP.

3.5.5 Mobile IPv6 Deployment scenario

In this section we will present the process of a mobile node changing networks and how Mobile IPv6 adapts to this change. As an example lets consider a user who starts a VoIP phone call from his IPv6 PDA while he is at home. How can the call be transferred seamlessly when he leaves his home and connects to his city's Wireless Metropolitan Area Network?

1. The user initiates the call while he is at home. This network is his Home Network where a host (or a router) acting as his Home Agent is located. The users' PDA is connected to his home network through a Wireless

Access Point. The VoIP connection is between the user's Home Address and the CN's (Correspondent node) address.

2. Once the user starts moving away from his home, the signal of his access point becomes weaker. At this time the signal of the wireless metropolitan area network becomes stronger.
3. The PDA decides it is time to change wireless networks (due to the difference in signal strength). It disconnects from his Home Network and connects to the Wireless MAN thus changing IP address.
4. Once the PDA is connected to the Wireless MAN it obtains a CoA (Care of Address) which will be used for binding with his Home Agent.
5. The PDA contacts his home network and sends a Binding Update to his Home Agent which binds the PDAs' CoA with his Home Address. In order for this step to be executed securely, the whole message sequence should be encrypted and authenticated using the AH and ESP headers of IPsec.
6. At this point the communication between the PDA and the CN is tunneled through the users' Home Agent. The CN doesn't know that the PDA changed network and IP address.
7. Now, the PDA will attempt to connect directly to the CN in-order to avoid the assymetric routing caused by tunneling all its traffic through his Home Agent. Using his foreign (CoA) address it contacts the MN with a "Binding Update" message which contains the PDAs' home address in a Home Address Option in a Destinations Options extension header. The whole process is once again protected using the IPsec extension headers to prevent from address spoofing and session hijacking.
8. The CN logically replaces the source address of the message (the CoA address) with the Home Address contained in the message.
9. Likewise, when the CN contacts our users' PDA, the PDA logically replaces the destination address of the message (the CoA address) with the Home Address. In this way, the connectivity is maintained and the upper layers of the protocol stack do not get involved with the transition.

Using the process described above, a mobile node (such as a PDA or a laptop) can change multiple networks and IP Addresses without shutting down its open connections. At the same time, using the Routing Optimization provided by Mobile IPv6, the connections do not suffer from high latencies as was the case with Mobile IPv4.

3.6 Summary

IPv6 is the next generation IP protocol, designed to overcome IPv4 barriers and shortcomings and at the same time help the Internet expand even more. IPv6 provides a wealth of benefits such as a vast address space, more efficient routing, Quality of Service, and improved security. What is important to understand though, is that when it comes to resilience and security, there is no silver bullet. IPv6 does help in certain security issues, but it is not and thus it should not be considered as an all-inclusive security solution. Companies and organizations should train their network administrators on the specifics of IPv6 and create an IPv6 task force so that the transition from IPv4 to IPv6 (including the intermediate stages) to be as secure and painless as possible.

4 DNSSEC

4 DNSSEC

4.1 Overview of the DNSSEC protocol

4.1.1 About DNS

One of the most critical components of the Internet architecture is the Domain Name System (DNS). DNS is a distributed dynamic database with a hierarchical structure, perhaps one of the largest and most active distributed databases on the world, that maps names of machines to protocol-level addresses. DNS is a service that the vast majority of Internet users and services relies upon, since the operation of popular Internet services such as e-mail, the web, or instant messaging, depend on it. The basic principle of DNS is the use of human-friendly domain names for Internet service addresses, as names are human readable and memorable, instead of numbers (Internet Protocol addresses), which are practical for computers only. In reality, whenever a user uses the domain name of a service (web page, email address, or other) the operating system must translate it to a numeric address in order to be able to connect to the service the user wants to use. DNS also performs the reverse operation, that is, translate an IP address to a fully qualified domain name.

Originally the DNS design was focused on data availability and did not address any resilience or security issues. The main category of DNS threats is data corruption, namely unauthorized modifications of DNS data. It is therefore possible to disrupt the operation of DNS by spoofing DNS messages, resulting to loss of integrity in Internet-based applications and services. Simply, if someone is able to spoof an IP address, a user may connect to a different server than the one initially intended, without any way of noticing it.

4.1.2 From DNS to DNSSEC

In 1990, Steve Bellovin discovered serious flaws in DNS. In 1995, a study was published by Steve Bellovin, and the Internet Engineering Task Force (IETF) started the discussions about DNSSEC. In 1997, RFC-2065 [44] published by the IETF was a first attempt to develop DNSSEC. In March 2005, RFCs 4033 to 4035 [36, 38, 37], led to the current form of DNSSEC. Simply put, DNSSEC defines a process whereby a suitably configured name server can verify the authenticity and integrity of query results from a signed zone. DNSSEC uses public key cryptography and cryptographic hashes to enable a security-aware receiving name server to:

- Authenticate the data received – verify that they could only have originated from the requested zone.
- Verify the integrity of the data. The data that was received at the querying name server was the data that was sent from the queried name server.
- Verify that if a negative response (NXDOMAIN) was received to a host query, that the target record does not exist (denial of existence).

To implement the three operations listed above, DNSSEC introduces special sets of Resource Records (RRs), specifically DNSKEY RRs, Resource Record Signatures (RRSIGs), Next Secure (NSEC) RRs, Delegation Signer (DS) RRs, and DNSSEC Lookaside Validation (DLV) RRs. Also, DNSSEC adds the following new message header bits: AD (for authenticated data) [78] and CD (checking disabled). To provide space for all the overhead added by the additional RRs, DNSSEC capable name servers and resolvers must support the EDNS0 extension [75], that implies that every DNSSEC packet contains an extra OPT pseudo record. Finally, DNSSEC requires support for the DNSSEC OK (DO) EDNS header bit [41] so that a security-aware resolver can indicate in its queries that it wishes to receive DNSSEC RRs in response messages.

4.1.3 DNSSEC related Resource Records

- **DNSKEY**: Every DNSSEC-secured zone has an associated private and public key pair, as generated by the zone's administrator. The private key remains secret, while the associated public key is published in the zone file, in the form of a DNSKEY resource record. The zone is digitally signed using the private key of the key pair mentioned above. DNSSEC allows for the use of RSA-SHA-1, DSA-SHA-1 and RSA-MD5 digital signatures. Users of DNSSEC are encouraged to use SHA-2, because it is widely believed to be more resilient to attacks than SHA-1 [48]. Two types of keys are identified for use in zone signing operations. The first type is called a Zone Signing Key (ZSK), and the second type is called a Key Signing Key (KSK). The ZSK is used to sign the RRsets within the zone, and this includes signing the ZSK itself. The KSK is used to sign the keys of the zone, which includes the ZSK and the KSK. The KSK is also defined in a DNSKEY RR.
- **RRSIG**: A Resource Record set (RRset) is a collection of RRs in a DNS zone that share a common name, class and type. In DNSSEC, RRsets are digitally signed by the zone administrator. This signature is computed by generating a hash of the RRset, then encrypting the hash using the zone administrator's private key. For a zone that contains SOA, NS, A, MX, DNSKEY resource records, there are minimally 5 distinct RRsets, and each RRset has its own RRSIG Resource Record. This implies that the granularity of DNSSEC signing is not at the level of an entire zone, but is aligned to a unit of a DNS query response.
- **NSEC & NSEC3**: When a zone is signed, an NSEC RR is added after each RR to chain together the valid host names appearing in the zone file. The last NSEC RR will point back to the zone apex or root. The responsibility of this new type of RRs is to authenticate "denial of existence" of a DNS query. For example, if the zone contains the names "alpha" and "beta", then there would be a NSEC RR for "alpha", and the RR value would be"beta", indicating that there are no defined names that lie between "alpha" and "beta". In addition, the NSEC record defines the set of RR types for this domain name, namely the NSEC record for "alpha", would have as a value field the enumeration of the RR types that are defined for "alpha". A side effect of this new RR is the implicit revealing of the enumeration of the entire zone file. With judicious use of the NSEC response it is possible to reconstruct the contents of a domain zone file, analogous to the outcome of a DNS list operation. Driven from this vulnerability, the NSEC3 [60] approach uses a hash algorithm on the names within a zone, and then uses a hashed ordering of these names. The next name references in the NSEC3 RR is the next name corresponding to this hashed name order. The objective of this approach is to increase the cost of zone enumeration using NSEC3 responses.
- **DS**: Once a zone is secured, it can then be added to an existing chain of trust or can be used to secure delegation to a sub-domain. In both cases, this is accomplished using a Delegation Signer RR. This resource record helps a client to validate a zone's public key (DNSKEY RR). The approach adopted by DNSSEC is to use a chain of trust within the hierarchical delegation structure of the DNS itself. The Delegation Signer (DS) RR contains the hash of the public key of the child zone. This record is signed by the parent zone's private key with a matching RRSIG RR. To validate a zone's DNSKEY, the associated DS, RRSIG (DS) and DNSKEY of the parent zone is retrieved. The DS record is validated by using the DNSKEY to decrypt the RRSIG (DS) record, and then checking that the result matches the DS record. This is the zone public key, according to the zone's parent. It can be compared to the DNSKEY record of the zone in question. This relies on the parent zone key, and the question is "how can this key be validated?". The same process can be applied here. The process stops when the DNSSEC client encounters a "trusted" key. The ideal "trust key" would be the DNSKEY of the root zone.
- **DLV**: The DNSSEC Lookaside Validation (DLV) service is an alternative method by which a chain of trust may be created and verified without the need to sign the parent zone file. The service is described in details, in RFC-5074 [77]. The DLV RR, described in RFC-4431 [35], is functionally identical to the DS RR and is placed in a special signed zone called a lookaside zone instead of the DS RR that would normally be added to the parent zone, thus removing the need to sign the parent zone. Assume that a name server is trying to verify the chain

of trust for the secured/signed zone example.com. Normally, it will first check for a trusted anchor provided for this zone, and in a absence of one, it will issue a query to find a DS RR at the parent zone .com. If a DS RR is not found, the DLV service will allow the name server, if the lookaside feature is enabled, to query a lookaside zone for the DLV RR of the zone being verified (example.com). The lookaside zone must be signed and the name server must have a trusted anchor for this zone. If the queried DLV RR is found, the zone example.com will be characterized as secure.

4.1.4 Dan Kaminsky's vulnerability - Overview

On July 8th 2008, the US-CERT issued a notice about DNS implementations being vulnerable to cache poisoning attacks. This vulnerability is a twist on classic DNS spoofing/cache poisoning, where false DNS records are injected into a DNS server's cache, causing domain names to resolve to incorrect IP addresses and potentially re-routing Internet traffic to malicious servers. The new vulnerability, discovered by security researcher Dan Kaminsky [12], is a more effective variant of the classic approach, namely flooding the DNS server with spoofed responses. According to the security vendor Websense, one of China's largest ISPs has fallen victim to the above dangerous vulnerability. Driven from this, one can say that DNSSEC may be a powerful measure one can take to prevent such types of attacks and increase the resilience of the network.

In summary, with DNSSEC, a zone administrator digitally signs Resource Record Sets (RRSets), and publishes the digital signatures produced, along with the zone administrator's public key, in the DNS. When checking a DNS response, a security aware DNS resolver can retrieve the related RRset digital signature and then check this signature using the public key against the locally calculated hash value of the RRset. The final step is to validate the zone administrator's public key against a hierarchical signature path that leads to a point of trust (trust anchor). If all these checks succeed, then the client has some confidence that the DNS response was complete and authentic. It is important to note that a security-oblivious name server, or a security-aware name server can continue obtaining results for all the domains, both secure-signed and insecure.

The main operations of DNSSEC are shown in Figure 4.1

Figure 4.1 – **The main operations of the DNSSEC protocol.**

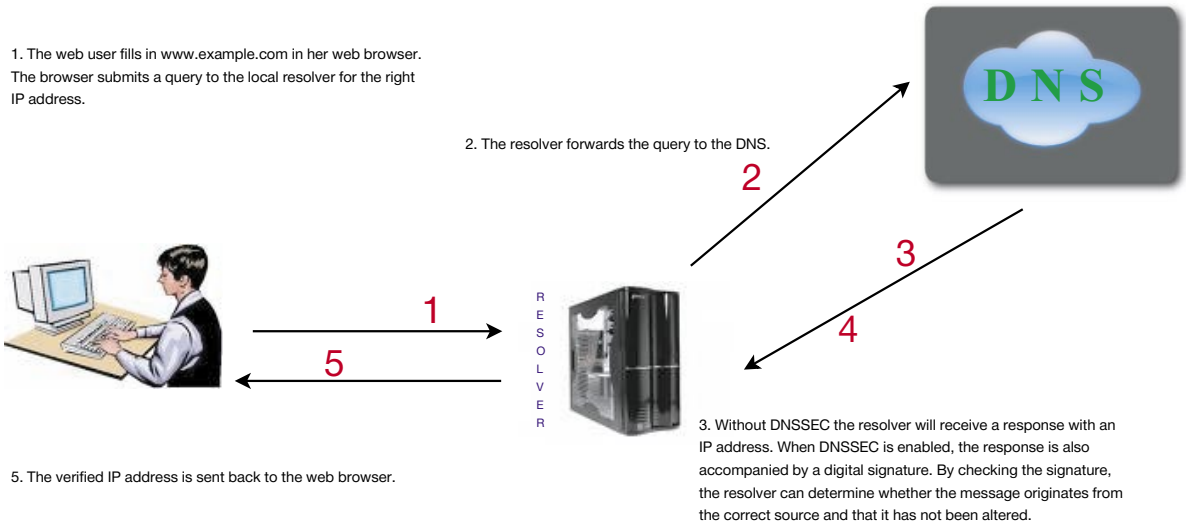

4.2 Resilience provided by DNSSEC

4.2.1 What is provided now

In this section we are going to discuss the points shown on the left side of Figure 4.2 in more detail. Security issues on the right side of the figure, cannot be addressed by DNSSEC.

Figure 4.2 – **Main threats against DNS.**

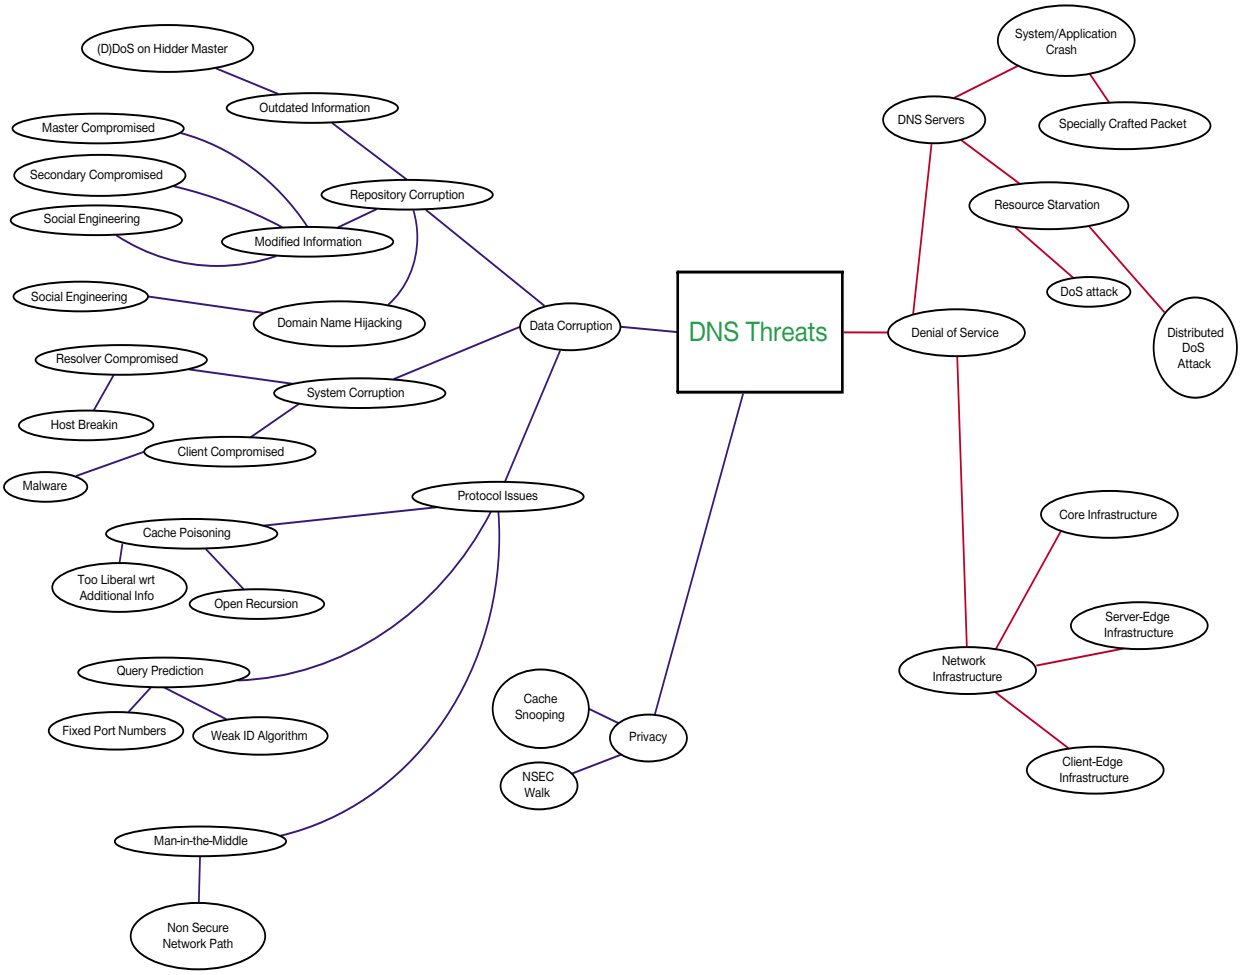

DNSSEC can be used by recursive name servers to validate the authenticity of a DNS response, the data integrity of the response, and to check whether the response indicates that no such domain or resource type exists (this negative information can also be authenticated). If an attacker forges a DNS response based on an original authentic response, and then attempts to propagate the response as an authentic response, then a DNSSEC-aware resolver should be able to detect that the response has been altered and does not correspond to the authoritative DNS

information for that zone. In other words, DNSSEC is intended to protect DNS clients from forged DNS responses. This protection does not eliminate the potential to inject false data into a DNS resolution transaction, but it augments DNS responses with additional information that allows clients to check that the response is authentic and complete.

The validation process performed by DNSSEC-aware resolvers is a powerful solution to DNS cache poisoning attacks, a concept known for over a decade. In short, to perform a cache poisoning attack, the attacker manages to inject bogus data into a recursive name server's cache, causing it to give out forged information to unsuspecting local clients. To do so, the attacker must send forged replies to the recursive name server by exploiting query prediction. Specifically:

- The Question section of the forged reply packet should match one of the sent question packets that exist in a pending queue, waiting for replies.
- The query ID field of the forged reply packet should match that of the question packet.
- The forged reply packet is sent to the same network address and port number from which the query was originally sent.
- The Authority and Additional sections represent names that are within the same domain as the question: this is known as "bailiwick checking."

Furthermore, the attacker must win the race against the authoritative name server, that is, the name server that is intended to answer the query. If the authoritative name server manages to answer the client's query before the attacker, then the attack will fail. If the victim name server does not correctly validate DNS responses to ensure that they have come from an authoritative source, the name server will end up caching the incorrect entries locally and serve them to users that will make the same request. With DNS cache poisoning, the nature of DNS has been subverted such that hostname-to-IP address lookups can no longer be trusted. The hostnames being visited are genuine, making DNS cache poisoning a problem different than phishing. DNSSEC was designed to protect from such types of attacks. Even if an attacker manages to win the race against the authoritative name server, he will not have the Digital Signatures of the resource records he would like to insert into the cache, so the attacked name server will discard the attacker's responses. Concerning open resolvers, which exist in a high manner out there, the problem of cache poisoning becomes more and more serious and frequent, as they are more susceptible to these attacks. Thus, DNSSEC can protect the users issuing queries to this kind of name servers. In Sections 4.4.3.1 and 4.4.3.2 we present two different scenarios, where cache poisoning attacks are performed against recursive name servers.

Concerning the man-in-the-middle attacks, where an attacker is between two hosts, DNSSEC can provide a prevention mechanism. The attacker has knowledge about the connection between the DNS server and the DNS client, and can use it to eavesdrop the connection or even inject data into it. The rapid growth and usage of wireless networks, where the network paths are most of the times insecure, increase the potential for this type of attacks. With DNSSEC, the injected data will not pass the data integrity and authenticity verification on the client side, so the attack will be prevented.

DNSSEC can also be useful in protecting against domain name hijacking, that stands for the situation in which Internet domain names are "stolen" from their rightful registrants. Hackers that have hijacked a domain can do anything with that name, including putting up their own website or redirecting its visitors to another malicious site. Thus, by hijacking a domain name an unauthorized DNS configuration change can lead to a pharming attack, explained below.

DNSSEC is an effective way to prevent pharming attacks, in which hackers aim to redirect a website's traffic to another bogus website, which is under the control of the attacker. Hackers employ pharming attacks for four

primary reasons: identity theft, distribution of malware, dissemination of false information and man-in-the-middle attacks. Unlike phishing attacks, this technique does not require the user to follow a hyperlink in a fake email message. In Section 4.4.3.3 we present a scenario based on a pharming attack.

Nowadays, Internationalized Domain Names (IDN) provide a backwards compatible way for domain names to use the full Unicode character set and this standard is already widely supported. An attacker could register a domain name that looks just like the one of a legitimate website, but in which some of the letters have been replaced by homographs in another character set. This creates many opportunities for phishing and other forms of fraud. For example, the attacker could send e-mail messages pretending to come from the original site, but directing to the fake website. However, if the domain name that corresponds to the original site is secure-signed, the redirection to the fake website will generate an alert if the resolver supports DNSSEC-aware queries. On the other hand, if DNSSEC is not in use, the fake site could record information such as usernames and passwords, while passing traffic through to the real website. The user may never notice the difference, until abnormal activity happens with their accounts.

The trust introduced by DNSSEC in the current DNS infrastructure can be also leveraged for other purposes than just trusting DNS responses. For example, DNS could be used to pass other keys, SSL certificates, and other data objects in general. All of the above can be secured by DNSSEC. DNSSEC promises a system in which a user can validate a key from an unknown host with only one (public) key. If the validation is successful, the user has some confidence that the key comes from the host from which it claims to come. DNSSEC provides a basis to build trust on the Internet to support higher level protocols facilitating IP telephony and web services.

DNSSEC can be used in conjunction with IPsec for establishing secure tunnels with remote entities. If a host wishes to communicate with some remote entity in a secure manner, the host will need to obtain a public key to authenticate the remote entity. Also, the host must have guidelines about how to contact the entity, directly or use a gateway along the path to the entity. All this needed information is stored in the IPSECKEY resource record [67]. Often, the host that launches the communication only has access to the IP address of the target node. Thus, the best way of looking up IPSECKEY RRs is to perform reverse lookups, based on the known IP address. Apparently, this is the time where DNSSEC is applicable. An IPSECKEY resource record should be used in combination with DNSSEC unless some other means of authenticating the IPSECKEY resource record is available. The information contained in the above resource record should be delivered to the end client in an integral manner. The trust relationship between the client and the server may be an end-to-end DNSSEC validation. When DNSSEC is not available, an attacker can:

- replace the public key stored in the IPSECKEY RR, which may compromise the resulting IPsec channel.
- replace the gateway address to point to a node under his control, which may result in a man-in-the-middle attack.

Note that RFC-4025 [67] suggests arguably that a deployment scenario supporting IPsec tunnels, where DNSSEC is obsolete, is not recommended.

4.2.2 Resilience Features

- **End-to-end data integrity check**: DNSSEC provides the ability to validate the authenticity and integrity of DNS messages in such a way that tampering with the DNS information anywhere in the DNS system can be detected. Namely, DNSSEC only secures traffic to the local recursive name server and typically cannot protect traffic all the way down to the end host (e.g., desktop, laptop). Thus, a malicious user can still attack DNS traffic that is in transit from the local recursive name server to the end host. In order to achieve end-to-end security, the current stub-resolvers, installed on most user computers, would need to be replaced with secured versions. But that would probably be very difficult.

- **Use TSIG to ensure the integrity with a recursive name server**: Within the scope of DNSSEC, DNS protocol interactions remain in the clear. Encryption of DNS exchanges using mechanisms such as TSIG for data protection between DNS servers are orthogonal to the relatively focused DNSSEC objective of allowing clients to authenticate the DNS response. Currently, TSIG (Transaction SIGnatures) is used to secure zone transfers (e.g. from a primary server to a secondary server). TSIG is based on a shared secret, which is used to sign the content of each DNS packet. The generated signature can be used for both authentication and integrity checking of the data.
- **Deploying DNSSEC without a Signed Root**: Many can argue that we don't need a signed root . A single ccTLD trust anchor is all that is required to secure the DNS transactions that occur within a country. Similarly, a large enterprise can deploy and include its own trust anchor to its caching name servers in order to fully secure its internal DNS. But, one way or another, the desired and ideal solution is to sign the root zone, and distribute only one trust anchor. In Section 4.4.2.1, we will provide the details concerning the progress of signing the root zone.

4.3 Potential Weak Points of DNSSEC

4.3.1 Denial of Service attacks

DNSSEC authoritative name servers, which are serving signed zones, send more zone records (including RRSIGs and NSEC RRs) per request, and thus, at worst, are marginally more vulnerable to DoS attacks. Also, an attacker could consume more resources in a security aware name server that supports DNS dynamic updates, by sending a stream of update messages forcing the security aware name server to resign some RRsets in the zone more frequently. Furthermore, an attacker has the opportunity to consume more resources in a security-aware resolver's signature validation process by interfering with RRSIGs in response messages sent to the resolver. Many researchers have the opinion that not deploying DNSSEC will not be a prohibitive factor for (D)DoS attacks, because the possible benefits of DNSSEC outweigh the additional (D)DoS risks. However, DoS attacks are not addressed by DNSSEC standards, as they are difficult to protect against, in a protocol like DNS. As mentioned in RFC-4033, there are no protocol features in DNSSEC to protect against DoS attacks.

4.3.2 Resolver's work load

When a resolver receives a reply from an authoritative name server regarding a previous issued query, it should validate the received data based on the RRSIGs and the public key published on a DNSKEY RR, corresponding to the queried domain. This increased workload will also increase the time it takes to get an answer back to the original DNS client. Furthermore, if the resolver is open, which means that it provides recursive domain name resolution for clients outside of its own organization, the problem becomes even more serious because the resolver has to serve a larger clients' community. Thus, its workload is increasing in a higher manner.

4.3.3 Hierarchical trust model

Each zone parent has a role of signing every delegated child's zone key, so that a zone's signing key can be verified by confirming the parent zone's signature of that key. In turn, the parent's zone signature can be verified by its parent's signature of this key and so on, until either a trust anchor or the root of the DNS is reached, or preferably both at once.

4.3.4 Size of a DNS response

The average size of a DNS response message increases due to the additional signature records that are attached to the responses. These larger UDP packets may exceed the path's (from the resolver to the name server) maximum transmission unit (PMTU), resulting to fragmented or dropped packets. There may be a specific case where a local

firewall on the resolver side could be configured to disallow fragmented DNS packets. As a result, this resolver will never receive possibly fragmented responses. In case the message size exceeds the maximum UDP message size, the name server will need to set the truncated response flag, causing the query to fall back to the use of TCP or try a larger message size. However, TCP has higher overheads in terms of client and server state and the number of network messages required to manage the TCP connection. Various opinions of researchers in the community discourage the use of TCP for DNS transactions [71]. In addition, authors in [40], stated that only one of 22 residential routers could proxy DNS queries/responses over TCP.

4.3.5 Zone file size increase

The zone file size usually increases due to the addition of the DNSSEC related resource records. The major contributors for the increase are the NSEC and RRSIG resource records. The increase size depends on the content of the zone file, but it has been noted in the DNSSEC literature that it increases of a factor up to seven.

4.3.6 Workload of zone administrators

Private keys need to be protected and rolled over. All changes to the zone file need to be signed with the Zone Signing Key (ZSK), so the busier the zone is in terms of additions, removals, and changes, the more work this entails.

4.3.7 Key rollover at the root

Key rollover at the root is difficult to achieve as it affects the whole hierarchical database.

4.3.8 Firewalls can be an obstacle

As stated in [65], during the early deployment of DNSSEC, some resolvers could not obtain signed DNS data from secure zones. This was caused by intervening firewalls that blocked any responses containing digital signatures. To some firewalls, these unknown signatures were clearly some sort of attack and the responses were dropped.

4.3.9 Routers may fail to process DNSSEC

There may be a situation where a number of broadband routers on the consumer market do not support DNSSEC, especially the AD bit, leading zone administrators to discontinue the use of DNSSEC. The report [1] mentions that a local broadband router could not handle properly the AD bit and therefore clients could not reach the gavle.se or ockelbo.se domains. This happened on September 21th, 2007. In another study [40], authors tested 24 residential router and firewall devices in two separate modes:

- Route mode: DNS queries were addressed to an upstream DNSSEC aware recursive resolver to verify that DNS packets could be routed transparently.
- Proxy mode: DNS queries were addressed directly to the unit under test to exercise router/firewall DNS proxies.

The clients performed the DNS queries, the routers/firewalls, an upstream recursive name server and an authoritative name server containing both signed and unsigned resource records, were all located in a closed test bed. On the one hand, all 24 routers/firewalls could successfully **route** DNSSEC queries addressed directly to an upstream recursive resolver. On the other hand, 16 of 22 DNS **proxies**, roughly 73%, could successfully pass DNSSEC queries and return validated responses of some size. For example, two proxies simply dropped any DNS queries that had the CD bit set. [2] One proxy dropped any DNS response that had the AD bit set. Furthermore, 6 proxies did not support the EDNS0 OPT RR and they failed to handle any DNSSEC queries.

[2] The CD (Checking Disabled) bit is used to inform an upstream validating resolver that full DNSSEC validation is not required and that any DNSSEC related resource records should be returned to the client.

4.3.10 Stale RRsets

An RRset is stale if an administrator has changed data values in new sets, but there exists a digital signature that covers the previous values and it has not expired yet. In this case, the stale RRset can be replayed by an attacker, who exploits the lack of a revocation mechanism in DNSSEC. Let's consider the example in Figure 4.3, presented by Eric Osterweil et al. in their study [65].

Figure 4.3 – **If data changes, such as in Modified RRset 1, then the old RRset 1 will still be verified by the zone's keys (even though the data is no longer valid).**

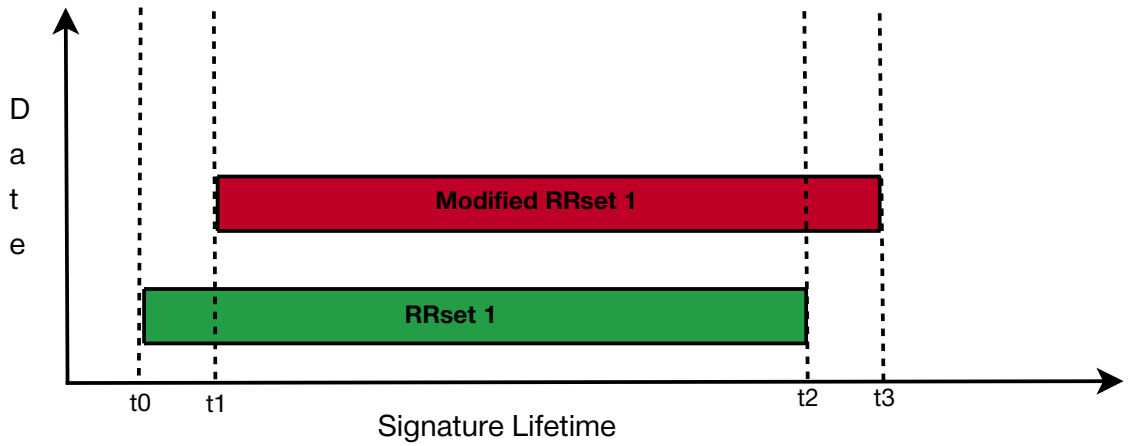

At time 0, RRset1 is created and signed. The signature includes an expiration date of time 2. At time 1, RRset1 value is modified. For example, the IP address of a host may have changed. The modified RRset is distributed to all authoritative servers and the previous value is flushed from caches after the TTL expires. However, an attacker can continue to replay the old record until the signature expires at time 2. Resolvers that receive the stale RRset will verify the signatures and declare that the set is valid.

4.4 Deployment

4.4.1 Current Deployment Status

At this time, DNSSEC is officially deployed and supported in Sweden (**.se**), Puerto Rico (**.pr**), Brazil (**.br**), Bulgaria (**.bg**), and Czech Republic (**.cz**).

4.4.1.1 Sweden

Sweden was the first country in the world that in late 2005 announced the signing of .se ccTLD domain [27].
In Sweden, DNSSEC was part of a pilot program by the Swedish registry of ICANN to implement DNSSEC as a commercial service. Participants were the Ministry of Enterprise, Energy and Communications, the registry of .SE, Swedish ISP TeliaSonera, Swedish bank Swedbank group and the Swedish National Post and Telecom Agency. .SE-DNSSEC is a supplemental service to .SE's domain name system. The objective of the service is for the domain name registrant to be able to secure his/her domains with DNSSEC. The service consists of .SE providing the customers with the possibility of administrating and publishing their DNSSEC keys, a hash of the customer's public keys - DS resource records, in the .SE zone, while .SE verifies them in accordance with .SE's policy for its handling of DNSSEC. The service provides an interface, the Domain Manager, where the domain name holder can publish the DS records to be stored in the .se zone. Using the above interface, the domain name holder has the opportunity to cancel the .SE-DNSSEC service for his/her domain name.

Just like .SE's regular domain name system, the domain name registrant is required to administer his/her own domain, namely administering his/her DNSSEC keys and signing their own DNS data with them. With respect to fee of the service, the first year's fee is charged by the .SE registrar when the service is ordered and varies among the different .SE Registrars. After the first year, the domain name holder is charged an annual fee of 80 SEK for every domain name with DNSSEC support. However, the next year, .SE will not invoice this fee and DNSSEC support will be a natural part of the domain name. The domain name fee currently is 120 SEK per domain name and year. At present, the following .SE registrars offer the service: Frobbit AB, Interlan Gefle AB, Gotlandica Internet (BRS - Intron AB), Leissner Data AB, Loopia AB, NEware AB, Melbourne IT CBS, Yask, City Network Hosting AB, Larsen Data v/Peter Larsen, and TDC Song AB. Technical issues for signing .se include:

- Generation of KSK takes place once a year. This key has 2048 bits length and it is generated using the RSA algorithm. It is stored on smart cards (FIPS certified). Its validity period is two years, which means that there exist keys with a one year overlap in validity. Lastly, the public key of KSK key-pair is published and distributed to the Internet community.
- Generation of ZSK takes place once a month. This key has 1024 bits length and it is generated using the RSA/SHA-1 algorithm. It is stored on portable storage media (USB). Its validity period is one month.

4.4.1.2 Puerto Rico

On August 1th, 2006, Puerto Rico (.pr) as the second ccTLD, began the transition to DNSSEC, a process leaded by nic.pr, a research laboratory in Puerto Rico. Nic.pr started signing the zones on July 2006, and on August 2006 started transmitting DNSSEC zones to the public server for the first time. Currently 19 zones are digitally signed, the .pr zone and 18 second level zones. Nic.pr provides a web portal [22], where the public Zone Signing Keys are listed, for .pr zone and the 18 child zones, mentioned above. Nic.pr is currently developing tools for: a web based DNS authentication approach, automated key rotations, Key Signing Key (KSK) support, and a dynamic tutorial of deploying DNSSEC.

4.4.1.3 Brazil

On July 2007, Brazil's .br top-level domain became a DNSSEC-signed zone. Chief technical officer Frederico Neves reports that .br, as well as four of its child zones, were signed on June 4th, 2007 and additional zones are expected to be signed in the near future. On April 2008, Brazil joined the top 10 top-level domains with DNSSEC zones when its Ministry of Justice signed all its subzones. Jus.br is the ministry's new top-level domain, and it requires DNSSEC as a mandatory feature. Registro.br [26] uses 3 key pairs for DNSSEC signatures:

- KSK BR (.br zone's Key Signing Key): Its private key is used to sign the set of public keys of the .br zone. The key pair is generated in a server totally separated from the network and stored on a removable media with encrypted file system. This removable media is stored in a safe place. The algorithm used to produce the key pair is the RSA/SHA-1, the key size is 1280 bits. The Digital Signatures generated with KSK keys are valid for 4 months. As regards the key rollover mechanism, programmed KSK rollovers are performed once a year. The keys are in use for 14 months and the double-signing technique is used [57]. During a period of 2 months there will be 2 active KSK key pairs. Lastly, the public key of the KSK key pair is published at: https://registro.br/ksk.
- ZSK BR (.br zone's Zone Signing Key): Its private key is used to sign .br zone's all resource records, except from the DNSKEY records. The key pair is generated in the on-line signer, a server connected only to the DNS publication server. The connection from the DNS publication server to the on-line signer is done with an exclusive cable and is used only for sending the records to be signed and to receive the signatures. The algorithm used to produce the key pair is the RSA/SHA-1, the key size is 1152 bits. The Digital Signatures

generated with ZSK keys are valid for 7 days. Programmed ZSK rollovers are performed every 3 months. The keys are in use for a little over than 3 months and the pre-publishing technique is used [57].

- ZSK *.BR (.br childs' Zone Signing Key): Its private key is used to sign authoritative records of .br child zones. The key pair is generated in the on-line signer. The algorithm used to produce the key pair is the RSA/SHA-1, and the key size is 1024 bits. The Digital Signatures generated with ZSK keys are valid for 7 days. Programmed ZSK rollovers are performed every 1 month. The keys are in use for a little over than 1 month and the pre-publishing technique is used [57].

4.4.1.4 Bulgaria

On October 30th, 2007, Bulgaria (.bg) successfully completed the 60-day period of the integration of the automated DNSSEC administration interface in the .bg zone and the subzones. On August 29th, 2007, register.bg [25] provided the users access to DNSSEC interface which allows each domain registrant to manage and setup DNSSEC using digital signature certificate. Register.bg accepts digital signature certificates issued by the following Certificate Authorities: Service B-Trust by Bankservice PLC, Service InfoNotary by InfoNotary PLC, Service StampIT by Information Services PLC, and Service Spektar by Spektar PLC. Register.bg is the first ccTLD registry in the world that has introduced a technology for automated DNSSEC management and setup. The DNSSEC service is provided by register.bg without an extra fee. As of August 2008, over 80 DNSSEC enabled zones exist in .bg.

4.4.1.5 Czech Republic

On September 30, 2008, the cz.nic association [11], allowed to use DNSSEC for securing DNS records. The public key used by cz.nic in order to validate the correctness of DNS records, for .cz domain, can be found at https://www.nic.cz/dnssec page. Furthermore, Czech Republic is the first country to implement DNSSEC for ENUM domains, which consists of telephone numbers, published by their owners in order to allow to call them over the Internet. The cz.nic association, as 0.2.4.e164.arpa domain registry, publishes DS records of this domain to an parent authority, that is e164.arpa domain, which is administered by the RIPE organization. To setup DNSSEC validation for ENUM domains, a client will need e164.arpa keys, which can find them on the RIPE website, on DISI project page [24].

4.4.1.6 in-addr.arpa domain

The RIPE NCC has the task to sign the reverse **in-addr.arpa** domain, which is delegated to it from the Internet Assigned Numbers Authority (IANA) and VeriSign (a trusted provider of Internet infrastructure services).

4.4.1.7 Testbeds

DNSSEC test beds exist in Russia (**.ru**), Mexico (**.mx**) and the United Kingdom (**.uk**). The test bed in Netherlands (**.nl**) is no longer active [47]. The world-wide DNSSEC deployment is presented in [34].

4.4.2 Future Deployment Status

According to a survey conducted in October 2007, 85% of ccTLD registries planned to deploy DNSSEC and 45% of these planned to deploy within two years.

4.4.2.1 root zone

With respect to the signing of the **root zone**, the Internet Governance Project (IGP) points at the US government's insistence on maintaining control of the root zone file as one of the reasons why it is not signed yet. More broadly, there are important concerns in the Internet operations and policy world about the assignment who should be given the supervision over the root DNS signing keys, and whether it would be possible to put in place the measures needed to avoid giving too much control to a single party. This debate has had negative impact on progress in

deploying DNSSEC compared to how it was originally envisioned during the standard's development.

The initial idea was to sign the DNS root zone and then distribute a single DNS trust anchor. Due to the recently increased awareness on security in general and broader exposure of DNS flaws in particular, the NTIA (National Telecommunications and Information Administration), a bureau of the U.S. Department of Commerce, announced on October 7 to international network leaders, that it was requesting comments on DNSSEC and the subject of root zone signing. Comments were due on November 24, 2008. With a properly signed root file, your browser can repeatedly ask, "How do I know this is the real answer? ", until the question reaches the root file, which says, "Because I confirm it."

4.4.2.2 .org domain

The Public Interest Registry (PIR), a non-profit corporation created by the Internet Society in 2002 for managing the .org top-level domain, wants to implement DNSSEC in the .org zone. The .org domain is the third largest generic top-level domain, with more than 6 million registered domain names worldwide. Specifically, ICANN has announced a Request for Comments on implementing DNSSEC on the Public Interest Registry's. As the first gTLD authorized to implement DNSSEC, .org is preparing an education and adoption plan within the Internet infrastructure community. Steve Crocker, chair of ICANN's Security and Stability Advisory Committee and CEO of Shinkuro, Inc, said, "Deployment of DNSSEC in .org is a big step forward for the security of the Internet. This will make secure DNS service available in the global gTLD community and thus available to everyone. I applaud .org, The Public Interest Registry for their leadership." "The argument we're trying to make is that there is a very real problem that DNSSEC solves and once we implement it within .org, it will be secure," said Alexa Raad, .ORG's CEO. The problem he mentioned is the critical flaw that researcher Dan Kaminsky exposed in the DNS system.

4.4.2.3 .gov domain

The U.S. government recently decided to sign .gov zone and issued a mandate to all federal agencies to deploy DNSSEC. The White House released a memo, signed by Karen Evans, the Administrator for the Office of E-Government and Information Technology, which instructs all government agencies to prepare for securing the federal government DNS infrastructure over the next year [45]. The deployment will start in January 2009 and is a response to the DNS cache poisoning attack that Dan Kaminsky made public a few months ago. In this regard, all top level .gov domains will be secured with DNSSEC by January 2009 and all the .gov sub-domains need to be secured by December 2009.

4.4.2.4 Windows 7 & DNSSEC

On October 30th, 2008, Shyam Seshadri, the Program Manager for Windows DNS at Microsoft, wrote a short entry in his blog [23] about DNSSEC in Windows 7. Indeed, it seems that Microsoft recognized the important role that DNSSEC will play in securing the DNS infrastructure, in the coming years. Windows Server 2008 R2 DNS server, will offer support for DNSSEC as per these RFCs 4033 through 4035. The DNS server is capable of generating keys and signing DNS zones using a sign-tool that Microsoft is providing with the product. The server will also be able to host these signed zones either as a primary or secondary zone. On the DNS client, a non-validating security-aware stub resolver relies on its local DNS server to perform DNSSEC validation and will check to make sure that the server has indeed done so. One positive side effect of this is that trust anchors do not need to be configured on the clients, thus saving a big chunk of the deployment efforts. However, the DNS client is security-aware, so it will expect the configured DNS server to indicate results of the validation when returning the response back. This is done by setting the AD bit in the response. If the DNS server failed to validate successfully, the DNS client will fail and discard the query. The security-aware behaviour of the client is a policy based mechanism whereby the Name Resolution Policy

Table (NRPT) will tell the client on which domains it is to expect DNSSEC functionality. Only for those domains will the DNS client set the DO bit in the query and expect the AD bit in the response. Additionally, the DNS client can use IPsec when issuing a query to a local DNS server, so the last-hop communication is also secured. An example NRPT table, in a simplified version, is shown in Table 4.1.

Table 4.1 – **An example NRPT table in a simplified version.**

Namespace	DNSSEC validation	Last Hop Security	IPsec encryption level
***.example.com**	Client sets DO bit. Server performs validation and sets AD bit. Client checks AD bit in response	Secure last hop using IPsec	High encryption
***.sub.example.com**	Client does not set DO bit. Server does not perform validation	Do not secure last hop using IPsec	N/A

Rule #1 is applied to the example.com domain and all its subdomains. For example, if a web browser queries the DNS client for www.example.com, the query will match the rule. Rule's details indicate that the DNS client must set the DO bit when issuing the query and check for the AD bit in the response. The rule also says it must use IPsec when issuing this query to the DNS server. Rule #2 is a more specific rule, thus any query under *.sub.example.com will match Rule #2 and not Rule #1. This rule indicates no DNSSEC validation, hence the DNS client would not set the DO bit, would not look for the AD bit in the response and would not use IPsec either.

4.4.3 Possible deployment scenarios

4.4.3.1 Cache poisoning attack against www.mybank.com

In this scenario, an attacker attempts to poison the recursive name server of a public ISP with a fake IP address for the legitimate banking website *www.mybank.com*. The attacker aims to force all customers of the ISP to visit a malicious website operated by him, instead of the real one operated by www.mybank.com.

The following steps are required for carrying out this attack:

1. The attacker sends a DNS query to the victim name server, asking for the IP address of www.mybank.com.
2. The victim name server will ask one of the root name servers and then will follow the referral (NS records) to ask one of the name servers which is responsible for the .com domain. The latter, will drive the victim name server to ask the name server of mybank.com, e.g., *ns.mybank.com*.
3. The attacker knows that at some point, the victim will ask ns.mybank.com for an IP address. Thus, he starts flooding the victim with forged DNS replies. All these packets are supposed to originate from ns.mybank.com, but they include a fake A record for www.mybank.com, containing the IP address of the attacker's web server.
4. The real name server (ns.mybank.com) provides a legitimate response to the victim's query. However, if the attacker has successfully predicted the query of the victim name server to the ns.mybank.com name server (according to the four requirements listed in Section 4.2.1), he will probably win the race against ns.mybank.com. If the attacker is the winner, the legal reply will be dropped from the victim name server.
5. The attacker asks the victim name server for the IP address of www.mybank.com and when he receives the bogus IP address, he knows that the attack has succeeded and then stops flooding the victim with forged replies.

The outcome of the above steps is a poisoned cache in the victim name server, with the bogus IP address of the attacker's web server. All future DNS clients asking for www.mybank.com will receive the same fake answer, until the

bogus entry expires from the cache, as denoted by the TTL value. Then, all users of this ISP that access www.mybank.com, will provide their credentials to the attacker, as they do not know that the name server of the ISP is under attack. The attacker can steal these credentials and use them to steal money from users' accounts. To some extent, the above problem can be addressed with the correct deployment of digital certificates. However, this will only address a part of the problem. Due to inefficiencies with this process, many users perceive certificate warnings as an annoyance and accept them automatically, without reading the details accompanied to them.

As discussed in Section 4.2.1, DNSSEC is a powerful solution for defending against cache poisoning attacks. If mybank.com had signed its zone and stored the digitally signed resource A records in ns.mybank.com, in conjunction with a security-aware recursive resolver at the ISP side, the above attack could be prevented. The administrator of mybank.com zone has to generate the keys and then sign the zone, producing the RRSIGs. Also, a DS resource record can be added to the parent .com zone, to create secure delegation. This step requires that the parent zone is also secured. Alternatively, a DLV RR can be added in a lookaside signed zone, thus removing the need to sign the parent zone. In the DNS query performed by the victim name server, the DO bit [41], will be turned on to indicate that the resolver wishes to receive DNSSEC data in return. Also, the victim name server will try to verify the chain of trust for www.mybank.com, by first checking for the existence of a trust anchor in its trusted-keys clause. If this is not found, it will issue a query to the parent .com zone for a DS resource record. If it is still not found, then the DNSSEC Lookaside Validation service will give the answer. After the verification of the chain of trust to www.mybank.com, the resolver can use the DNSKEY RRs of mybank.com zone and the RRSIGs in order to verify the data integrity and authenticity of the replies. The attacker has to modify the A records **and** the RRSIGs in order to pass the verification step in the resolver. But now, the chances to perform the attack are nearly to zero, because it is really hard to modify a RRSIG to match the digest generated at the resolver side. Remember that the digital signatures are checked by decrypting the signature, using the public key - DNSSKEY RRs, regenerating a new digest from the received data, e.g., A records, and comparing the decrypted digest and the newly generated digest. These two digests must match in order to accept the reply.

Note that DNSSEC can prevent the Dan Kaminsky's attack, in which instead of poisoning a single A record, the authority records are hijacked, giving the opportunity of controlling an entire zone.

4.4.3.2 Mail attack against www.mybank.com

In this scenario, an attacker attempts to mask copying of an entire email, again poisoning by a recursive name server of a particular ISP. This attack was presented by Dr. Paul V. Mockapetris [63]. Let us suppose that customers of www.mybank.com can send email messages to bank's employees, informing them for various issues, such as buy or sell stocks on behalf of the customer. Thus, a customer can contact a bank's employee and list, in the body of the message, at what price he would like to sell or buy particular stocks. The attacker's aim is to mask copying of this message and then send another email to the bank, which will cancel the use of the previous one. For example, the attacker can change the price of buying or selling a stock.

The following steps are required for carrying out this attack:

1. One bank's customer wants to send an email to the employee whose email address is account@mybank.com. The mail server of the customer will try to SMTP connect to the mail sever of www.mybank.com (e.g., mail.mybank.com) However, the customer's mail server has to know the IP address of the bank's mail server. Thus, it will issue a query to the victim recursive name server in order to find this IP address.
2. Similarly, as the attack explained in Section 4.4.3.1, the attacker sends a DNS query to the victim name server for the mail server of www.mybank.com. The steps followed by the victim name server and the attacker are nearly the same. As a result, the victim name server will have two MX records, one containing the valid mail server of the bank (mail.mybank.com) and one containing a mail server under the attacker's control (mail.hackerbank.com). The illegitimate mail server will have higher priority than the valid name server.

3. The customer's mail server will try to SMTP connect to mail.hackerbank.com, as it is the one that has the higher priority. mail.hackerbank.com will respond back that it can not store email. But, the profit gained is the body and the header info of the email.
4. As the connection to the fake mail server failed, the customer's mail server will try to SMTP connect to the valid mail server, mail.mybank.com. The connection now is established and the email will successfully delivered to account@mybank.com.

The above steps have many implications, which can lead to a mask copying of the entire message sent to account@mybank.com. As stated in Step 3 the attacker knows that the customer wants to contact the bank and perform a stock exchange. Thus, he can send another email to the bank, which will have the same sender (spoofing of the original sender), a later date and its body will contain the previous message sent to the bank. For example, the spoofed email's body can contain the phrase: *"Go ahead and ignore the previous message. I updated the information"*, and then the content of the previous valid email sent to account@mybank.com. As a result, the bank's employee will ignore the first message and will take into account the invalid second message sent. This could have serious implications for the customer's stocks.

Like the previous attack scenario, this new one can also be prevented if mybank.com has signed its zone and stored the digitally signed resource MX records in ns.mybank.com, in conjunction with a security-aware recursive resolver at the ISP side. The steps required by the administrator of ns.mybank.com and the resolver are exactly the same, as these presented in the previous cache poisoning attack. The attack is presented in Figure 4.4.

Figure 4.4 – **Mail attack against www.mybank.com.**

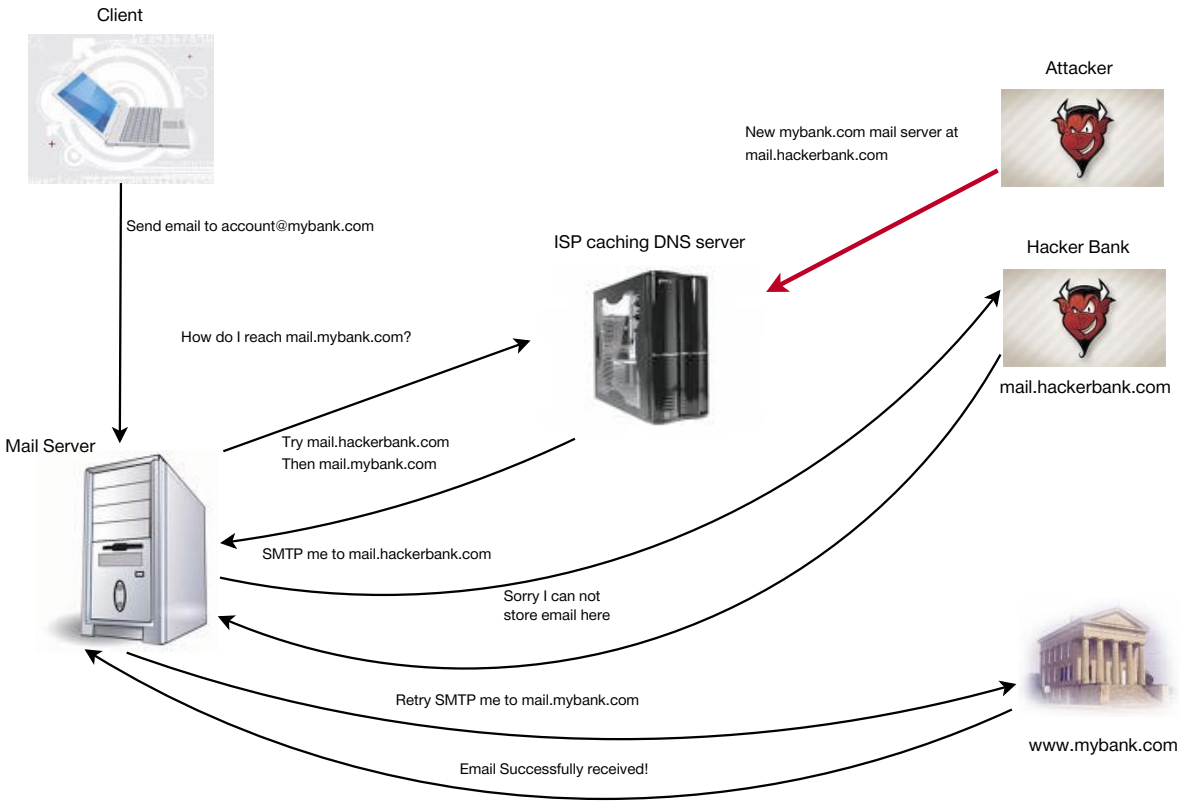

4.4.3.3 Pharming attack against www.mybank.com

In this scenario, an attacker attempts to poison the recursive name server of a public ISP with a fake IP address for the legitimate banking website www.mybank.com. As in Section 4.4.3.1, the attacker wants to direct traffic to his fake web site for www.mybank.com, by poisoning the cache in a different way.

The following steps are required for carrying out this attack:

1. The attacker registers a dummy domain name (e.g, mydomain.com) and sets up a DNS server for this domain.
2. The attacker sends out spam emails, trying to advertise his domain name out there.
3. A customer of the target ISP receives the spam email and clicks a hyperlink to visit www.mydomain.com.
4. Attacker's DNS server, configured in Step 1, returns the IP address of www.mydomain.com accompanied with an A resource record containing a fake mapping of www.mybank.com to the IP address of attacker's fake web site.
5. The recursive name server, at the ISP side, accepts the bogus information for www.mybank.com and store it in its cache.

All future DNS clients, being serviced by the target ISP, asking for www.mybank.com will receive the same fake answer, until the bogus entry expires from the name server's cache, as denoted by the TTL value. All these clients will be redirected to the attacker's fake web site.

Like the previous attack scenarios, this new one can also be prevented if the DNS records in the authoritative name server for mybank.com (ns.mybank.com) are digitally signed by the zone administrator. Again, the resolver at the ISP side must be a security-aware resolver. If the attacker changes the A resource record for mybank.com, the responses delivered by a DNS client will be marked as bogus because the checking process of the attached signatures, using the public key of ns.mybank.com, will fail. Again, if the attacker tries to replace the valid signature with a new one, it will not be verified using the public key of ns.mybank.com.

4.4.4 Operational Status of the DNSSEC Deployment

Eric Osterweil et al., in their study [65], used DNSSEC deployment data collected using the SecSpider monitoring project [28]. Their dataset, discussed in the paper, covers October 2005 through January 2008 and includes 871 secure "production" zones. The monitoring infrastructure consists of monitoring points, called pollers that send DNS queries to the authoritative name servers of zones and use the DNS responses to form the raw data used in the study. The pollers are located in the United States, Europe and Asia, and on networks comprised of universities, home access and enterprises. Authors analyzed the collected monitoring data, using three measurement metrics: availability, verifiability and validity.

On the subject of **availability**, they are trying to measure whether the system (name servers) can provide all the data to the end systems (resolvers) requesting it. To perform this, at regular intervals all pollers query all secure zones. To overcome transient networks problems, the pollers will each issue up to three queries with timeout thresholds set to a conservative 10 seconds. Their results found that at least one poller could receive a response from a zone at a particular time in 99.925% of the experiments. Namely, out of the 871 secure zones, only 44 zones could not send a response to none of the pollers. The next point to look is the availability dispersion, namely how many resolvers - pollers can reach the zone. As shown in Figure 4 from [65], roughly 20% of the monitored zones suffer availability dispersion. This means that some resolvers may not be able to receive critical data (DNSKEY RRsets) from a zone based solely on where they query from.

Regarding **verifiability**, they are trying to measure whether the end systems can cryptographically verify the data they receive. This metric captures the amount of configuration needed to verify DNSKEY RRsets, in terms of trust anchors. Recall that in order to verify the integrity of the responses delivered to a resolver, the latter must be configured with some initial set of keys from trusted zones, referred to as trust anchors. A DNSKEY RRset for zone z is covered by trust anchor T if there is an authentication chain leading from T to z. In their study, authors realized that out of 871 secure zones, 662 have no authentication chain leading to them, which means that a resolver would need to manually configure 662 trust anchors in order to verify all existing signed DNSSEC data from all the 871 zones. This process clearly becomes not feasible as the number of secured zones moves from hundreds to thousands and the majority of them have no authentication chain leading to them.

Lastly regarding **validity**, they are trying to measure whether the verified data is actually valid. There are four possible combinations, which are shown in Table 4.2, dealing with verification and validity of data.

Table 4.2 – **Verification vs validity matrix.**

	Verified	**Unverified**
Valid	Ideal Behavior	False Negative
Invalid	False Positive	Intended Defense

Verification refers to the cryptographic process, performed by a resolver, where a data unit (DNSSEC data) is determined to be either verified or not. Validity refers to whether the data delivered (DNSSEC data) actually corresponds to what the zone administrator intended. In particular, false negatives can occur when a zone breaks its own secure delegation from its parent. A DS resource record stored at a parent zone must match a DNSKEY stored at the child zone. In a different situation, the authentication chain is broken. As of January 17th, 2008, the monitoring pollers had observed 1.730 DS resource records and 1.573 of these matched DNSKEY resource records in the child zones, namely roughly 9% of the authentication chains were broken and data verification would have failed for all DNSSEC data in the affected child zone and all its descendants. Taking into account false positives, an attacker can replay stale RRsets long after these RRsets have been removed from the zone's authoritative name servers and have been flushed from all the caches. The above happens when a zone administrator selects a long signature lifetime. For example, if a RR is signed using a one year lifetime and changes a few days later, the stale RR can be replayed until the year long signature expires and verified by unsuspecting resolvers. The attackers perform these types of attacks exploit the lack of a revocation mechanism in DNSSEC protocol. Figure 9 from [65] shows the number of zones that have a number of stale RRsets associated with them, for all the 871 secure zones.

4.4.5 Resource Requirements

On October 5th, 2005, RIPE NCC in collaboration with NLnet Labs published a report, where they presented their measurements of the effects of deploying DNSSEC on CPU, memory and bandwidth consumption of authoritative name servers [58]. They did their experiments by replaying query traces captured from ns-pri.ripe.net and k.root-servers.net, one of the 13 root name servers, in a controlled lab environment. Namely, they assume that DNSSEC is deployed in ns-pri.ripe.net and k.root-servers.net. For zones signing, they created a 2048 bits RSA/SHA-1 key signing key (KSK) and two zone signing keys (ZSK) varying from 512 to 2048 bits.

Regarding memory load, they measured the virtual memory size (VSZ) as reported by ps command, taking as parameters the size of zone signing keys, the number of RRSIGs generated and the operating system of the name server. The memory increase after signing the root zone was measured to be about 156 KB for BIND v9.3.1 on FreeBSD 6.0. Regarding CPU load, it does not seem to be a function of the zone signing key size and would grow

from 4 to 5%, so this is of no concern. To collect bandwidth statistics, they ran the iostat command on the name server machine, while the query trace was replayed, using tcpreplay tool. As shown in Figure 4 from [58], the egress bandwidth for measurements about the ns-pri.ripe.net traces, against BIND v9.3.1, doubles for small and triples for larger zone signing keys.

Figure 5 from [58] displays the amount of packets sent out from the name server for zone signing keys ranging from 512 to 2048 bits. For a 512-bit key the amount of packets is the same as for unsigned zones. Lastly, for 768-bit to 2048-bit signed zones the amount of packets is 10% more, which probably indicates IP fragmentation.

4.5 Summary

DNS Security Extensions (DNSSEC) provide a more secure way of doing lookups of Internet addresses for services such as the web and e-mail, through the definition of additional DNS Resource Records. In contrast to the current domain name system (DNS), lookups with DNSSEC are signed cryptographically, which makes it possible for DNS clients to validate the authenticity of a DNS response, the data integrity of the response and authenticated denial of existence. The DNSSEC service protects against many of the threats to the Domain Name System, but it does not provide confidentiality of data and does not protect against DDoS attacks. It is the best way to address the recent large-scale exposure of DNS vulnerabilities. The research community tries to deploy DNSSEC world-wide, but the absence of a signed root zone has a negative impact on these deployment efforts. In general, full deployment of DNSSEC requires global cooperation across many entities in both the public and private sectors. These entities include those known as registries and registrars that provide DNS services, as well as Internet service providers, non-profit and professional organizations, equipment and software manufacturers, standards and coordinating bodies, research labs, universities, and large enterprises.

5 MPLS

5 MPLS

5.1 Overview

Multi-Protocol Label Switching (MPLS) is a networking technology built around a label based forwarding paradigm. An MPLS header containing one or multiple labels (organized in a label stack) is attached to packets. Label Switch Routers (LSRs) forward these packets based only on the label information. Both Layer 2 (L2) and Layer 3 (L3) packets can be encapsulated in MPLS. Typically MPLS is deployed in the form of a MPLS backbone or core network. All traffic inside the core network is forwarded using MPLS. Traffic that enters the MPLS core is labeled at an edge router and edge routers also remove the labels from the traffic that exits the MPLS core. Typically labels have only link local scope and label mappings in each LSR determine how packets will be forwarded through the MPLS backbone. By properly setting up the label mappings in all the LSRs, one can form a Label Switched Path (LSP) that will carry traffic over a specific path through the network independently of the network's native routing mechanisms. Labels are distributed through label distribution protocols that are either specifically developed for MPLS (LDP) or extensions of pre-existing protocols (RSVP-TE or MP-BGP). The MPLS header can carry Quality of Service (QoS) precedence information that allows differentiated treatment of labeled packets inside the MPLS core.

Figure 5.1 – **An overview of the MPLS architecture**

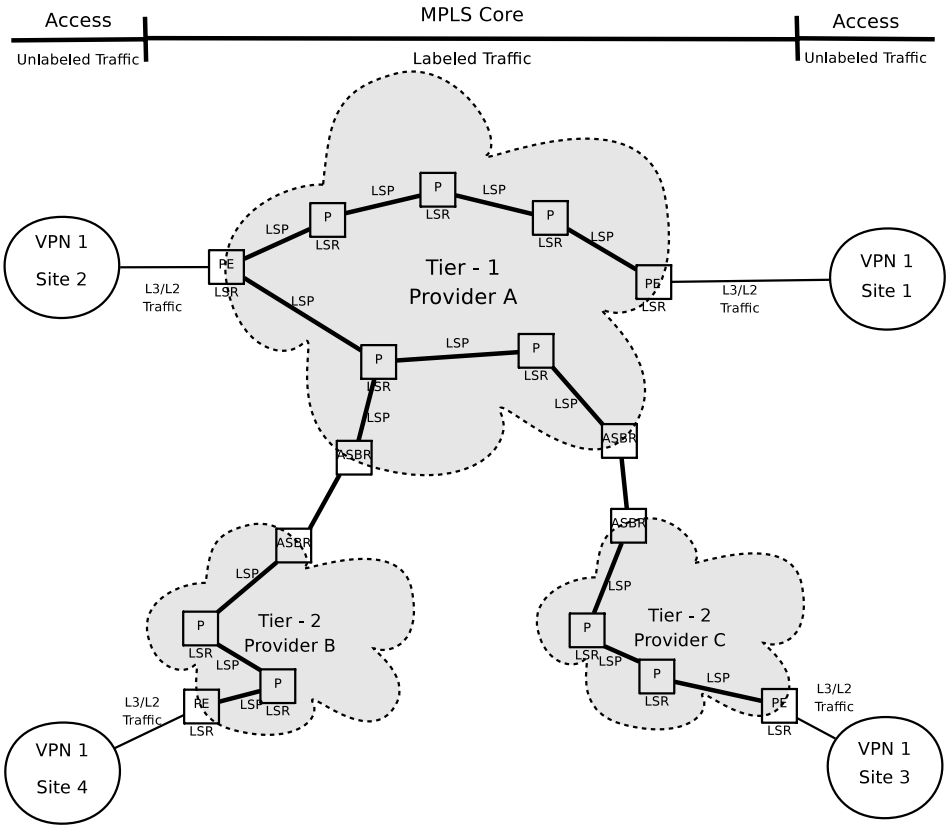

Figure 5.1 shows a simplified overview of the MPLS architecture. A provider's core network consists of internal provider (P) LSRs and provider edge (PE) LSRs. Customer sites connect to the PE routers through L3 and L2

interfaces. The LSPs are setup between the PE routers inside the MPLS core and typically all traffic inside the MPLS is labelled. The networks of multiple providers can be interconnected. Smaller tier-2 providers (like regional ISPs) are connected through the networks of larger tier-1 providers (national or international ISPs). When multiple provider networks are interconnected PE routers act are Autonomous System Border Routers (ASBRs). A customer VPN can now have sites attached to different tier-2 and tier-1 providers.

MPLS has gradually become very popular and widely deployed by multiple service providers and carriers. Indeed, MPLS offers multiple advantages to a service provider. Its multi-protocol nature allows a provider to carry multiple types of traffic such as Asynchronous Transfer Mode (ATM), Frame Relay (FR), and Time Division Multiplexed (TDM) over the same MPLS/IP infrastructure. This allows providers to consolidate their networks while protecting the investment of their customer that can continue using their ATM/FR equipment. MPLS allows providers to take advantage of cheaper new generation IP and Ethernet equipment when they upgrade their core networks. The circuit-like forwarding model of MPLS is very convenient since providers are very familiar with provisioning and running circuit-based ATM and Frame Relay networks and most of the Operation Support Systems (OSS) currently used by providers are circuit-based. MPLS packets carry QoS information so it is possible to differentiate traffic inside the provider network and create flexible QoS mappings between the unlabeled customer traffic and the MPLS traffic inside the core. Furthermore, unlike ATM, MPLS is a packet based technology that allows statistical multiplexing of traffic and can increase network utilization and further reduce capital expenditure for the providers. Finally, MPLS is very flexible: various services can be constructed by combining building blocks and protocols allowing fast deployment of new service offerings with relatively little cost to the provider.

In addition to the label based forwarding paradigm, MPLS has introduced flexible and open signalling protocols that have become the basis for offering services across provider boundaries that use equipment from different vendors. This is a clear improvement from the closed networks of the past with their proprietary signalling systems. In fact, the usefulness of the MPLS signalling is such that it has been used over non-IP or even non packet switched networks that do not support label based forwarding. Generalized MPLS (GMPLS) is a set of signalling protocol extensions for setting up LSPs over a variety of non packet switched technologies such as Time and Wavelength Division Multiplexing, and fibre switching. Transport MPLS (T-MPLS) is a subset of MPLS signalling and management mechanisms that can be used in non-IP packet transport networks. While the standards for MPLS and GMPLS where developed by the Internet Engineering Task Force (IETF), initial T-MPLS development was undertaken by the International Telecommunication Union (ITU) and for a while this caused concerns about potential interoperability issues that could endanger further adoption of MPLS. Recently, IETF and ITU agreed to work together in developing the standards for T-MPLS.

Providers typically use MPLS to implement:

- L2 point-to-point connections (pseudo-wires) that carry legacy traffic (ATM, Frame Relay) over a common backbone. The multi-protocol capabilities of MPLS are used to consolidate the provider's network while maintaining customer investment in ATM/FR technology.
- Various types of Virtual Private Networks (VPNs) (L2 and L3 VPNs and VPLS). This service replaces previous VPN offerings based on ATM and FR technology. MPLS L3 VPNs make it possible to offer L3 VPNs to a customer and Internet access through a single L3 interface. This was not possible with ATM or FR VPNs where customer had to have a separate L2 connection for each remote site and a separate connection for Internet access. Furthermore, now when providers run IP networks and common protocols, it is much easier to offer services across multiple provider networks. The label based forwarding model of MPLS can elegantly support traffic isolation between VPNs through a combination of labels with an outer label determining how traffic will be forwarded through the core and an inner-label providing per-VPN traffic multiplexing.

- Traffic Engineering (TE) of the traffic inside the provider core networks. The circuit-like nature of MPLS forwarding along with its QoS support allows providers to have fine control over how the traffic flows over their networks. Essentially, providers can get traffic management and QoS functionality similar to what is offered by ATM but over a statistically multiplexed and cheaper IP network.
- Improved resilience of provider core networks. The circuit like nature of MPLS allows providers to quickly shift traffic around node and link failures and guarantee minimal service impact when such failures occur. Before MPLS this was very difficult to do over an IP network.

5.1.1 Resilience and Security of MPLS

Using MPLS has significant impact on both the security and the resilience of the network and the services it supports. As we will see MPLS technology can be used to implement mechanisms that can quickly repair traffic when there are network failures or sudden variations in traffic patterns. In this light, it is not hard to argue that MPLS improves the resilience of the network. In terms of "traditional" security, the picture is not so clear. MPLS was not designed with security in mind. In most MPLS deployment cases (for example L2 and L3 VPNs) it just replaces the incumbent technology (ATM/FR) without any claims of improving the security of the service. Most of the studies of MPLS security strive to show that MPLS is as secure as the technology it replaces. Some of the major advantages of a MPLS based VPN service like Internet access and service that spans multiple provider networks come with security risks. Also, the IP nature of the backbone enables the easier connectivity between customers and providers but also makes the network more vulnerable to generic IP attacks. MPLS is a good example of a technology that enables new services but also creates new security risks.

The security/resilience of any network technology depends as much on how the technology is deployed and operated as on its technical details. Although the fundamentals of the MPLS technology (i.e. the label based packet forwarding) are appealingly simple, the reality is that MPLS deployments rely on a variety of routing and signalling protocols and specific provider-centric network architecture. Each protocol used in an MPLS deployment can be attacked so the security and resilience of the underlying protocols has to be taken into account. MPLS's extreme flexibility allows providers to create very complex deployments that are hard to configure and manage/operate properly. High configuration complexity can mean increased risk of misconfiguration that can result in service downtime or security breaches. Furthermore, flexible VPN services like extra-nets and inter-provider services create very complex trust models that are hard to fully understand and enforce.

In the following we will attempt to provide a brief but comprehensive overview of the issues around MPLS resilience and security. This overview will by necessity be brief since the MPLS architecture comprises of multiple protocols and network architectures. There are multiple IETF work groups devoted to various aspects of the MPLS architecture that have produced a large number of Requests for Comments (RFC) and Internet drafts that cannot be summarized in the space available in this study. Also, some of the protocols and mechanisms used in MPLS are adapted from pre-existing technologies. In this study we will focus more on what is unique and new to the MPLS architecture and ignore the more mature and familiar elements.

5.2 Resilience

In the context of a service provided over a network, resilience is the ability of the network to continue providing the service even in the face of various types of unplanned adverse conditions. Typically, these are equipment failures but they also include sudden shifts of traffic load due to unplanned events or attacks. Typically, providers offer Service Level Agreements (SLAs) to their customers that provide them with some guarantees on the level of service provided. As a result, the provider must not only ensure that the service is available when unplanned events occur but that it also satisfies the agreed SLAs. Even when service interruptions occur it is important to repair the service fast. This typically requires tools that help localize the failure fast and repair it.

The MPLS technology provides mechanisms that can be used to improve the resilience of provider networks and consequently the resilience of the service offered to customers. MPLS supports mechanisms that can protect from link and node failures inside the backbone network as well as at the interconnection points between providers or customers and providers. MPLS also provides extensive Operation, Administration and Management (OAM) functionality that can be used to verify the levels of service offered and localize and repair failures fast.

One of the major drivers for MPLS deployment is its ability to carry multiple types of traffic over the same backbone. As more traffic types share the same network resources the potential for unwanted interactions between them increases and problems in one of the traffic types can negatively affect other traffic types. Strong traffic isolation and management mechanisms are needed in order to prevent these interactions. Traffic management mechanisms also allow providers to control how traffic flows over their networks. This ensures that changes in network traffic and failures will not compromise the quality of the offered services.

In the following sections we present a summary of the MPLS resilience mechanisms.

5.2.1 Data protection

A fundamental requirement on any type of network is its ability to survive link and node failures. The previous generation of transport networks achieved this at Layer 1 with Synchronous Optical Networking (SONET) that supports elaborate ring topologies and can guarantee traffic repair in less than 50 milliseconds. IP networks recover from failures by re-converging to a new set of routes that bypasses the failure. This convergence happens at the protocol level and can take a significant amount of time. Clearly, IP mechanisms are not sufficient for the emerging multimedia types of IP traffic. MPLS uses the label based forwarding paradigm to achieve effective traffic repair at the sub-IP level in packet networks of arbitrary topology. Furthermore, these traffic repair mechanisms are signalled and based on standard protocols and can be used across provider networks, something that is difficult to do at the SONET level.

MPLS supports both protection and restoration. Restoration reacts to a failure by determining an alternate path for the traffic that is affected by the failure and redirecting the traffic over this path. Protection pre-computes the alternate paths for a variety of failures and simply redirects the traffic when the failure occurs. The backup path must be diverse from the path that is being protected, since otherwise a single failure may affect both. Protection is typically much faster than restoration. The ability of MPLS to force traffic to follow an alternate LSP is fundamental for fast traffic recovery. In IP networks traffic cannot be forwarded along any path other than the one that all routers have agreed upon when the routing protocol has converged. It is not possible to send traffic over another path immediately after a failure occurs since this traffic may form routing loops. MPLS uses explicitly routed paths that do not depend on the underlying IP routes and does not have this problem. Nevertheless, recently there have been efforts for addressing fast re-routing in IP networks and a variety of IP tunnelling techniques have emerged that can achieve very fast data protection without MPLS [15].

MPLS can offer end-to-end or local protection to an LSP through extensions to the RSVP-TE label distribution protocol. In end-to-end protection the LSP is rerouted at its source. Since the failure notification has to propagate to the source first, the traffic repair may take longer. In local protection, also known as Fast Re-Route (FRR), traffic is switched at the node that is next to the failure. In this way, recovery can be very fast, and in these cases MPLS can deliver guaranteed sub-50 millisecond recovery. MPLS allows great flexibility when configuring how traffic will be protected. Dedicated backup LSPs can be created for each protected LSP, or multiple protected LSPs can share the same backup. On the other hand, local protection from all potential node and link failures will require setting up a potentially large number of backup LSPs. MPLS provides the Constrained Shorted Path (CSPF) mechanism that allows the automatic computation of the proper paths for the backup LSPs and can automatically set them up through RSVP-TE signalling. Even if CSPF makes the setup of a large number of backup LSPs easy, the cost of these

LSPs is not negligible. The backup LSPs will have to be monitored for correct operation and will consume resources in the control plane.

The complexity of FRR increases considerably when we consider point-to-multi-point (P2MP) LSPs that are used for services that require the multi-casting of high volumes of information, e.g. television broadcasting. P2MP LSPs require multiple backup tunnels for protection from failures of a node that replicates traffic to multiple outgoing links. In addition, since P2MP trees typically carry high volumes of traffic, the placement of the backup LSPs must be done very carefully otherwise the network can become congested during a failure. Furthermore, protecting the source of a P2MP requires special procedures and adds more complexity. P2MP FRR is a relative new area and solutions are still in their standardization phase in IETF.

5.2.2 Traffic management and isolation

MPLS allows multiple types of traffic to share the same packet based IP infrastructure. It is very important to prevent problems in one type of traffic from resulting in service degradation for some other traffic. Certain types of traffic, like video and voice have strict requirements in terms of service received from the network, i.e. delay, loss and jitter. The network must be able to fulfill these requirements at all times. Today it is common for providers to offer SLAs where there is a financial cost to them if the SLA is not met. At the same time, the provider has to ensure the sufficient utilization of its network in order to reduce its capital costs.

The MPLS architecture provides the necessary building blocks for achieving all the above goals. These are:

- The MPLS header's capability to carry sufficient Quality of Service (QoS) or Class of Service (CoS) information in the so called Experimental (EXP) bits. LSRs can provide various types of queuing and scheduling according to the QoS/CoS information carried on the MPLS headers of packets that can provide end-to-end guarantees on delay, jitter and loss through the MPLS core.
- The processing at the edge of the MPLS core provides a natural point where non MPLS QoS/CoS per-packet information (like IP ToS, DSCP markings, dot1.q P bits, ATM QoS information) can be mapped to the MPLS domain's EXP bits. This allows the provider to implement a number of different policies and mappings that control how the various types of traffic will be carried over its network. Incoming traffic is mapped into various internal CoS classes and groomed into a single or multiple LSPs for transport over the MPLS core and is re-mapped to its initial or a different QoS class as it exits the MPLS network.
- MPLS allows providers to control how traffic will flow over their MPLS network. Traffic with certain QoS requirements can be sent over an appropriately routed LSP, for example traffic will very low latency requirements can be sent over an LSP that has a small hop count.
- The MPLS edge provides a natural point for performing rate limiting (on ingress) or shaping (on egress). This is fundamental for the proper operation of the network since it implements the admission control part of the QoS architecture.
- MPLS allows operators to react fast to changing network conditions. Since LSPs are typically dynamically signalled it is possible to re-route them fast when network conditions change due to failures, changing traffic patterns, or attacks.

Some of the above capabilities are not unique to MPLS. ATM backbones offered extensive traffic management capabilities and traffic scheduling but in a very discrete granularity (since ATM virtual circuits had pre-determined "sizes"). MPLS provides equally sophisticated traffic management capabilities over a packet based network. Furthermore, ATM backbones required complex and mostly non signalled (i.e. through an OSS system) configuration which typically resulted in higher configuration turnaround times.

The ability of MPLS to force traffic to follow certain paths in the network is being used to replace the ATM circuit-oriented mechanism that the providers have used in the past to control how traffic flows over the backbones. Using the much simpler and controlled MPLS forwarding model it is possible to formulate the routing problem in a provider's network as a flow optimization problem and derive "optimal" solutions based on certain definitions of optimality. Initially, it was thought that this was a fundamental advantage of the MPLS forwarding model that could never be emulated by IP. More recently, research has shown that it is possible to achieve similarly optimal routings by tweaking the IGP weights of an IP network although the situation becomes more complex if one considers network failures. Still, most providers are familiar with the circuit based provisioning systems and are not willing to replace their provisioning and traffic management systems.

There are still some issue with MPLS traffic management though. One important limitation is that the MPLS traffic management system can provide only differentiated service, i.e. packets are classified into a small number of discrete service levels. This model seems to be appropriate for current services that only handle a small set of well defined traffic types (low latency video, low jitter and loss video, customer data and low priority non customer data). It is not clear how well this model will carry forward as new applications and traffic types emerge.

Another potential issue is the complexity of network provisioning. Providers must continuously ensure that their network resources are sized properly so that there is no congestion in their networks. Network provisioning is a complex problem since usually there is limited information about the incoming traffic and how it flows over the network. To a certain degree MPLS simplifies network provisioning since traffic flows over a set of known LSPs. On the other hand, the provisioning problem becomes much more complex when one considers the various backup tunnels used for FRR and has to ensure that the network will have sufficient resources to handle a variety of potential failures. MPLS provides mechanisms where paths without enough resources can be found automatically for backup LSPs through constrained path computation algorithms (CSPF). This simplifies the provisioning problem but typically the paths found are only locally optimal. Global optimality requires off-line computation of all the LSP paths in the network by sophisticated optimization software. It appears that provisioning an MPLS network is much harder than provisioning an ATM network, since the later relies on the SONET layer for traffic protection and traffic flows over explicitly routed circuits with known capacity and traffic characteristics. To make matters worse, the ability to operate across provider boundaries increases even further the complexity of provisioning QoS service levels and TE paths in the network.

5.2.3 OAM

A necessary dimension of providing a resilient service is the validation that the service operates properly, the quick detection of failures and the equally quick service restoration that involves fast fault isolation. This function is commonly referred to as OAM (Operation, Administration, and Management). The MPLS architecture provides mechanisms for OAM that are based on existing OAM mechanisms for IP networks and it operates in two main modes: (a) always on, where the OAM function is always running without explicit operator intervention and (b) on demand, where the operator initiates the OAM function.

The OAM mechanisms used in MPLS are typically:

- BFD: this is a low level HELLO protocol that can be run over a single link, or over LSPs. It is meant to be an always on OAM mechanism and can quickly detect link level failures or failures of the logical connection between protocol or forwarding plane peers.
- MPLS ping: is an extension of the ICMP ping used in IP networks and is used in on-demand mode to check the reachability of various destinations in the MPLS network. Its operation relies on a Router Alert (RA) Label that instructs the MPLS LSRs to intercept the packets carrying it and process them at the control plane. MPLS ping relies on responses from the destination to determine if it is reachable or not.

- traceroute: is based on MPLS ping and can be used to isolate faults after reachability problems have been detected by MPLS ping.
- VCCV (Virtual circuit connectivity verification) is a generic mechanism used to monitor the health of a LSP. It establishes a control channel between the endpoints of the LSP and allows for the exchange of connectivity verification information. It is designed to support multiple profiles that allow it to operate in both on-demand and always on mode and use a variety of control channel models, which include in-band where the verification traffic in over the data path and out-of-band where the verification traffic follows the control path using the MPLS router alert label. The actual OAM function is performed through MPLS ping or ICMP ping.

MPLS OAM presents some interesting problems. Unlike the IP world, MPLS LSPs are uni-directional, i.e. the end point of the LSP cannot reach the head node of the LSP over the LSP itself. This presents a problem for MPLS ping, since the destination needs to send a reply to the ping initiator. In IP/MPLS networks this happens over IP, so there must be connectivity at the IP layer between source and destination. Failures in the IP layer connectivity will affect the correct operation of MPLS Ping and may lead one to assume that the data layer is not working properly. Also, unlike IP, in a MPLS network there are multiple logical "destinations" that can be pinged, i.e. LSP endpoints, IP VPN prefixes, pseudo-wire endpoints or segmented pseudo-wire stitching points. Pinging all these different types of destinations requires pushing the right labels in the Ping packet and ensuring that a proper return path exists in all cases. Cases where the forwarding plane performs load balancing across multiple links are also problematic since it may be difficult to force the ping probe to follow a particular link from the link group. All these constrains make MPLS ping a relatively complex mechanism.

Typically MPLS pseudo-wires provide symmetric service and consist of two LSPs each in opposite direction emulating bi-directional service. VCCV can handle this configuration and automatically provides bi-directional OAM functions.

OAM may be contained within a provider's network or can span the link between customer and provider or more commonly the link between providers. In these cases common policies for what is considered a failure and how to react to it must also be in place across the provider/customer boundaries. Some of the OAM mechanisms used in MPLS require coordinated configuration across provider/customer boundaries.

Finally, in applications like pseudo-wires, it is necessary to inter-work MPLS OAM mechanisms with the L2 OAM mechanisms used at the attachment circuits in order to provide end-to-end OAM capabilities. For example when two ATM attachment circuits are inter-connected over a MPLS pseudo-wire, a failure on the remote attachments circuit or in the MPLS section of the network must be communicated to the local ATM circuit and be converted to the appropriate ATM OAM messages.

Overall, MPLS OAM tools have reached a good level of standardization and support from vendor equipment and are now widely deployed. They provide a quite effective way to monitor the basic operation of the network. The complexity of the various MPLS services and building blocks makes using these OAM mechanisms, interpreting their results, and isolating failures difficult. In most cases only specialized personnel can perform the OAM functions and even then only with the support of OSS systems. Certain areas like OAM between different providers are both very well understood and difficult to configure due to both technical and business reasons.

5.2.4 Performance and security monitoring

In order to offer a comprehensive security and resilience solution, providers must implement high level mechanisms that monitor the operation of their network. In addition to basic connectivity monitoring and fault isolation that is provided by the OAM mechanisms discussed above, the providers need to implement a security/performance monitoring system that ensures that the provided levels of services are normal and attack are detected and mitigated quickly. Such systems are relatively well understood and currently widely deployed in most providers.

MPLS based services provide some significant new challenges though. One is the multi-provider nature of some MPLS services. This forces the security/performance monitoring systems of different providers to work together even in a limited way, something that most of the existing systems were not designed to do. In some other MPLS services like L2 VPNs, performance monitoring must span both the MPLS backbone and the L2 attachment points at the two ends of the pseudo-wire. For example an ATM node that is connected to a remote ATM node over a MPLS pseudo-wire should be able to receive performance information using the ATM performance model. At the edges of the MPLS network appropriate inter-working operations must be performed so that MPLS specific performance information is mapped to ATM specific performance information and vice versa.

5.3 Security

In this section we provide a brief analysis of the security properties of the MPLS architecture. We present an overview of the MPLS trust and threat models, the basic mechanisms used for providing security and then we study the specific security characteristics of the most popular MPLS services.

5.3.1 Trust model

MPLS is always deployed in a customer/provider architecture with MPLS used by the provider to efficiently implement revenue driving services. In most of the common MPLS applications (L3, L2 VPNs, and VPLS) the customers do not necessarily know that the service they get is implemented using MPLS and their interface to the provider is non-MPLS (IP, or L2). Multiple customers get can get service from one or more providers. Some MPLS services require building a hierarchy of providers, where a tier-N provider is a customer of tier-(N-1) provider while others require providers to form a peer-to-peer relationship. In a L3 VPNs service the customers may also access the Internet through the same service interface.

In this complex environment, there can be many types of "trust":

- customers always trust their provider for providing the service: there are contractual agreements that the provider will implement the agreed service, with certain quality (i.e. Service Level Agreements)
- customers may or may not trust their provider with the contents of their traffic: certain security conscious customer like banks may encrypt their data before sending them to the provider. Note that MPLS does not provide any form of data encryption
- customers do not trust other customers. Like in legacy VPN services, MPLS VPNS deployments strive to completely isolate customers from each other so that there is no possibility of attack or data leakage between customer networks
- in extra-net deployments customers may trust other customers in a limited way. This means they want limited connectivity between some selected systems or sites of the two networks. In non MPLS deployments this would be achieved through a private link between the two networks. MPLS L3 VPNs can support some amount of connectivity between the two customer networks but it is difficult to provide fine grained access control and the customers will need to take their own measures to ensure that only the intended parts of their networks are indeed exposed
- providers do not trust their customers and try to protect their infrastructure from customer attacks
- providers do not trust their peer providers but they have to allow them some minimum amount of access in order to implement a service that spans the provider boundaries. This involves setting up of a link layer connection and a protocol peerings, and accepting and installing control plane information like routes from other providers. On the other hand providers do not want to expose details of their internal topologies and policies that are considered trade secrets

- providers and customers do not trust the Internet and try to protect their networks from external attacks. If a customer network is compromised the impact is relatively low since this does not mean that the MPLS service has been compromised. When the provider's network is compromised this can have catastrophic consequences since all the MPLS services and customers can be compromised. If a customer is under DoS attack the whole provider network can be affected and other customer's service can be disrupted

5.3.2 Threat model

Like IP routing, MPLS operates at both the control and the data plane so we must consider both aspects to form a complete picture of the security issues. In brief, the common threats of any MPLS deployment are:

- compromise of a router
- compromise of the protocols involved in providing the MPLS service
- link layer attacks
- DoS attacks at the protocol, link, and control plane level
- forwarding plane attacks
- misconfiguration
- information leakage, i.e. exposing the topology or real IP addresses of the provider's or customer's network to unauthorized parties
- attacks from the Internet

Attacks can be initiated by

- insiders to the provider's network
- customers
- other providers
- external entities (i.e. "from the Internet")

Most MPLS security studies simplify the problem by assuming that "'the core is secure"' so there can be no insider attacks. This is generally not true but it is useful for reducing the scope of the issues that have to be considered and we will mostly follow this assumption in the rest of this study.

5.3.2.1 Router compromise

Routers that belong to the service provider's domain are typically secured following sets of rules commonly referred to as Best Current Practices (BCP), i.e. a set of configuration guidelines that if followed can improve the security of a system. For an overview of the various BCP documents see [29]. The most effective mechanism for protecting the provider's routers is to apply packet filtering at the edges of the provider's network, so that outside traffic is not allowed to reach the routers. Although this mechanism is effective, it may be hard to configure consistently on all the external interfaces of the network provider. Packet filtering may also have performance implications since it requires more resources from the forwarding plane of the routers.

PE routers, i.e. routers that are directly connected to customers are potentially more exposed, since their address may be known to customers in certain configurations. Depending on the type of service provided the customer may have L2 or L3 access to the router and in some cases may setup routing protocol sessions with the PE router. Packet filtering is extensively used to drop all un-authorized traffic from the customer. Filtering can be static or dynamic (i.e. with more filtering rules added dynamically based on the traffic to be filtered). In certain deployment scenarios the provider's router may be co-located in the customer's premises. This means that the physical security of the PE

router cannot be ensured and a malicious or compromised customer can physically tamper with the PE router.

Routers can always be compromised easily by insiders. Extensive logging of all configuration actions on all routers and a real time system that monitors these actions may be a way to quickly detect such incidents. Proper password management is also crucial, especially changing passwords immediately after authorized personnel leaves the company. It is also important to ensure that all the available protection mechanism are configured properly, e.g. all passwords are set to reasonably strong values. According to various studies it is not uncommon to find, even in the networks of large service providers, devices that still have their manufacturers default passwords or other trivial passwords.

Common protocols used for securing administrative access to systems are SSH, SNMPv3 that provides encryption and authentication (while SNMPv1 does not), TLS, IPSec, and secure SNMP over SSH. Some of these mechanisms are still in the standardization stage so their deployment may still be limited. A recent carrier survey [59] suggests that there is still a large percentage of carriers and network operators that use SNMPv1 for both read and write access to their equipment. Another survey [53] suggests that most operators find IPSec too complex to configure and do not use it for securing access to their equipment. [59] also indicates that a significant percentage of operators (at least among the survey participants) still uses telnet for configuration access, which is alarming. Note that both [59] and [53] do not focus exclusively on MPLS providers, but it is reasonable to assume that a large percentage of the tier-1 and tier-2 providers they participated offers MPLS based services.

5.3.2.2 Attacks at the protocol level

Attackers can target the protocols that are used to support the MPLS service. MPLS deployments rely on MP-BGP for exchanging routing information and labels, LDP and RSVP-TE for exchanging labels and SNMP and other protocols for configuration management. If the attacker manages to compromise one of the provider's routers then it is likely that he can interfere with the protocols used by the provider and disrupt the services provided. Even without compromising one of the provider's routers an attacker can try to attack the protocols operation if he can observe and tamper with the traffic on the provider's links. The PE router again is a weak point in the provider's network since it may participate in protocol sessions with the customer. As a result the customer can attack the protocol using malformed protocol packets or just brute force DoS attacks like sending storms of HELLO packets or attempting to establish a high number of TCP sessions. Protocol peerings between providers are equally exposed. Given the sometimes high volume of protocol traffic and state exchanged between providers, the stress on protocols there is much higher.

Routing and signalling protocol security has received a lot of attention and there are many BCP documents [73, 39] on how to secure protocol operation both in general and for specific protocols. Here we provide a very brief overview of the general security mechanisms typically employed by protocols:

- do not enable the protocol on interfaces that is not needed. Install packet filtering so that protocol packets are dropped as early as possible on the interfaces the protocol is not activated
- rate limit the rate of incoming messages to avoid meltdown under DoS attacks
- rate limit the rate of generated messages to avoid destabilizing the network
- rate limit the rate that new protocol information is accepted, i.e. number of OSPF Link State Advertisements (LSAs) or BGP routes
- rate limit the rate the new protocol information is generated
- if the identity of the peers is known, make sure that protocol traffic from other peers is dropped early. This can be susceptible to spoofing so it is mostly used in TCP based protocols

- Use some form of MD5 signatures in an attempt to validate the originator and the contents of received protocol packets. The security of the MD5 hashes has come under scrutiny recently and may not be sufficient to deter sophisticated attackers and MD5 inherently suffers from replay attacks [18]. In order to increase security, MD5 keys must be configured on a per protocol peering and be changed frequently but this presents significant management challenges. Although according to [59] MD5 is the most popular option for security protocol exchanges some of the MPLS protocols are still not very mature in terms of their key management or re-keying operations that are used in conjunction with MD5 [61]. Furthermore, securing a protocol session between a customer and the provider requires them to share a password and agree to the key rollover details. The same applies when securing a protocol session between different providers making inter-provider provisioning complex. Note that the MD5 approach attempts to protect protocol packets and not the logical protocol information contained inside these packets. The latter is a much harder problem due to the potential volume of this information and the cost of providing cryptographic protection to it. The issue of routing information authenticity has received a lot of attention in the inter-domain context where there are multiple proposals on how to secure the BGP routing information. These efforts are still in the experimentation stage and have not been considered in the context of MPLS services.
- If protocol exchanges are over TCP, there is some degree of protection from various attacks but possible attacks have been shown to exist. TCP implementations must be modified so they are more resilient to these attacks [18].
- use IPSec to secure protocol sessions. According to [59] and [53] this option does not seem to be very popular due to its perceived configuration complexity
- use Time to Live (TTL) to ensure that protocol peers are one hop away. Using this mechanism protocol packets that are directly sent to a peer are sent with a TTL of 1, so the receiver can be fairly confident that packets were sent from a directly connect peer. According to [59] this does not seem to be very popular. [53] found that vendor support for this functionality is inconsistent
- use audits and test suites to ensure that the protocol implementation is reasonably robust against malformed packets
- use graceful restart mechanisms that allow quick recovery from protocol crashes
- use a distributed control plane implementation where each protocol is isolated from the others and can fail without affecting the rest. Note that in some router architectures this has been carried to a more extreme form with the virtual router concept where the router is logically partitioned in multiple "different" virtual routers, one for each customer. This partition can become very sophisticated with operating system support for fair sharing of system resources like memory and CPU. These virtual router implementations though fail to scale to the number of customers required by modern MPLS deployments and most implementations use a single instance of a protocol for all customers

A large number of protocols may be used in a MPLS deployment but the importance of each protocol for the service may be very different and the way the protocol is deployed will determine the potential risks. Naturally, attacks again protocols like LDP, RSVP-TE, MP-BGP that are directly used in providing the MPLS service are the most dangerous. On the other hand, these protocols are typically (but not always) deployed within the provider's core network and are more difficult to attack from "outside" of the core. ICMP and ICMPv6 are commonly used in OAM mechanisms and their compromise could cause disruption to all the offered services. In many cases customers interact with the provider's OAM system and this exposes the provider to customer initiated attacks. In some MPLS services a customer may peer with the provider over OSPF, RSVP-TE or LDP. In this case, the danger for the provider is much higher. SNMP is used for configuration management of the provider's system so it is very sensitive but also hard to attack from outside the core. Other common protocols like RADIUS, DNS, NTP are commonly used inside the

customer's networks or inside the provider's network with limited interaction between the two so they are harder to attack for outside the provider's core. Any time a provider establishes a protocol peering with its customers it opens itself to potential attacks.

If a service is offered across multiple providers then even protocols like MP-BGP and RSVP-TE that are typically are used internally must establish peerings across the inter-provider links enabling a whole new set of attacks from malicious or improperly secured peer providers.

5.3.2.3 Link Layer attacks

Typically the connection between customer and provider or between providers is a dedicated point to point link. If in certain deployments this link is over a shared medium (eg an Ethernet switch) then a number of well known link layer attacks are possible. Ethernet has gained popularity as an interconnection technology between customer and provider but in most cases the customer provider connection is strictly point to point between a customer router and a provider router. Interconnection over Ethernet has become more popular in Internet exchange points and PoPs where multiple providers peer with each other. Given the cost of a full mesh of point-to-point connections, in many cases these connections are over a switched network.

Ethernet in particular is very sensitive to attacks since the Address Resolution Protocol (ARP) that is used to maintain the mapping between L3 and L2 addresses is not secure. L2 attacks especially the ones related to ARP are well known and there exist tools that allow unsophisticated malicious parties to mount L2 attacks very easily. These toolkits not only enable L2 attacks but also facilitate the escalation of the attacks in the application layer, and there are examples where traffic can be injected in TCP sessions or SSL sessions can be attacked in a man in the middle fashion. A comprehensive discussion of Ethernet attacks and techniques for their mitigation can be found in [69]. Certain attacks, like attacks on the spanning tree algorithm used in bridged networks, are not applicable to the PoP environment since there is no spanning tree used there. Also, some solutions, like manual configuration of MAC address, although infeasible in larger networks can be very effective in the much simpler PoP interconnect networks.

5.3.2.4 Forwarding plane attacks

External entities and customers can try to attack the network by spoofing the network information, i.e. spoofing L2 and L3 address information or MPLS labels. Customer data traffic with spoofed L2 and L3 addressing information is typically not a big concern since the the PE routers decide how to handle the incoming traffic based on the interface it receives it and not the address information of the packet. If for example two customers are connected to the same PE router, their connections will be over different interfaces on the PE router, so the PE router will be always able to determine the source customer for all incoming packets. Then it will forward the traffic based on the per-VPN specific forwarding information, so even if the destination address information is spoofed, this will only affect that specific VPN. It is not possible to inject traffic to a different VPN by manipulating the L2 or L3 addressing information on the data packet. In some cases the customer-provider connection is over a logical rather than a physical point to point connection, e.g. over a VLAN. In these cases if the customer-provider connection is not secure (i.e. it is over a unsecured switched Ethernet network for example) it may be possible that the customer or an attacker can spoof the VLAN information in the packets it sends and in this way impersonate traffic from other customers.

Spoofing labels is not very effective since in most MPLS services the customer does not exchange labelled packets with the provider and it is simple for the provider to drop all the labelled packets that arrive from customers. Labelled packets are typically exchanged on the inter-provider links though. A common method for protecting against spoofed labels is to check that each incoming packet contains a label that was distributed to the originating network. This functionality is rather complex to implement since it requires coordination between the signalling

protocols and the forwarding plane and increase processing at the forwarding plane. Such tests will be even more expensive in the interfaces between carriers where traffic rates are high, but this is the place where this checks are needed the most. If the L2 connection between the providers is not secured, an attacker can replace labels with other labels that are valid but correspond to e.g. a different VPN. In this way, the packets will pass the label test but will be delivered to the wrong VPN.

5.3.2.5 Denial of Service attacks

Since the MPLS infrastructure is shared among all customers, any kind of DoS attack can potentially affect all customers. Control plane DoS attacks are possible against the PE routers. For this reason, PE routers use all available protocol mechanisms to reduce this exposure. This typically involves limiting the number or rate of routing messages and the number of routes that the PE will accept from the customer. BGP provides several such mechanisms but other protocols like OSPF may not be as robust. ASBR routers that are used to interconnect different providers in multi-provider topologies also need to deploy similar mechanisms to protect themselves from resource exhaustion.

More subtle resource exhaustion attacks are possible when customers or peer providers are allowed to create higher level MPLS constructs like LSPs or VPN routes. MPLS resources can be exhausted if the ability of the customers to create these constructs is not effectively controlled.

In certain cases, even normal service operation can cause a significant control plane load. For example, in VPLS services, signalling in the MPLS core is triggered to resolve ARP requests originating from customer networks. Without proper limiting mechanisms in the core it is trivial for a customer to overwhelm the core network with ARP resolution requests. OAM present a similar problem since customer initiated OAM requests can result in OAM activity over the MPLS core.

Typically the data rate that a customer can send to the provider is controlled through QoS policies and ingress rate limiting, so the provider will have adequate protection against the customer sending too much traffic. Still, data plane attacks can become a more serious issue in the links between providers. If a provider is to protect itself from attacks originating from a peering provider, its border equipment (i.e. the ASBR) should be able to rate limit traffic at the LSP and potentially QoS class granularity. This may require significant amount of computing power especially given the high data rates of the inter-provider links, and in most cases will require advanced (and expensive) equipment.

Like in all DoS cases, it should be possible for the operator to intervene even when the DoS attack is going on, in order to mitigate this. This means that the operator should be able to access the equipment that is being attacked and the equipment must have enough resources reserved so that it can allow the operator to log in and perform certain configuration actions (i.e. install new traffic filters or shut down interfaces). This capability typically requires careful engineering of the software of the equipment in question and it may not be available in some low capacity or older equipment.

5.3.2.6 Information leakage

The typical MPLS architecture with the clear boundaries between the customer and the provider network is suitable for hiding the topology of the provider's network. Very little information about this network needs to be known to the customer so the addresses of the provider's routers and other provider internal information can be kept hidden. An external observer can try to derive information about the provider's network through probing or by monitoring the changes in the TTL of the packets that cross the provider's network. MPLS protocols solve this by allowing the operator to control how to update the TTL of forwarded packets and by packet filtering at the edge of the MPLS core.

Unfortunately, implementation of inter-provider services requires providers to reveal some information about their network. Critical MPLS functions like establishing diverse LSPs in FRR applications, or establishing a TE LSP that spans multiple provider networks and OAM across provider networks inherently require relatively detailed information about the remote provider's network. MPLS protocols go to great lengths in order to minimize the amount of information revealed and this in many cases results in suboptimal service implementation and increased configuration complexity.

5.3.2.7 Misconfiguration

MPLS services require the configuration of a large number of protocols and consistent configuration of a large number of network elements. Moreover, large providers have large networks with 1000s of customers, some of them with 100s of VPN sites. Certain MPLS services such as multicast VPNs or carrier-of-carriers L3 VPNs have even higher configuration complexity. Typically, providers depend on Operation Support Systems to automate the configuration tasks and reduce the probability of configuration errors. MPLS protocols provide some help in the form of various auto-discovery mechanisms that can automatically perform the necessary configuration steps when new entities are added in a MPLS service, e.g. a new VPN site is added in a L3 MPLS VPN. These auto-discovery mechanisms have been controversial in the past in IETF and may have not been universally implemented by all vendors yet.

MPLS functions like FRR require the establishment of potentially very large number backup LSPs. Again MPLS protocols provide some help through the Constrained SPF (CSPF) that allows the automated routing and establishment of the backup LSPs, but it is not clear if this mechanism can provide a globally optimal placement of backup LSPs and providers still depend on their OSS for traffic engineering.

5.3.2.8 Attacks from the Internet

It is common that a customer that has subscribed for L3 VPN service also receives Internet access as part of the same service. The customer (and for that matter the provider that offers the service) are exposed to all the threats that exist on the broader Internet. The standard methods of defence are used in this case (Firewalls, IDS). A difficult issue with Internet access is that of DoS. Depending on the implementation of the MPLS service a DoS attack from the Internet may also disrupt the L3 VPN service for the customer, resulting in a more severe impact. To avoid these problems, some providers use different infrastructure (i.e. PE router) for the VPN and the Internet service. This typically results in higher costs but it improves the resilience of the customer's VPN service.

5.3.3 Security analysis of the most common MPLS services

Since the most popular MPLS services are rather different in terms of operation and configuration here we provide a brief security analysis of each service.

5.3.3.1 Pseudo-wires and L2 VPNs

Pseudo-wires are used to implement a one-to-one replacement of a previous ATM/FR service where the customer must maintain a separate L2 circuit for each remote VPN site. The customer accesses the service through the same interface (ATM, FR, TDM or recently Ethernet). The L2 traffic is appropriately encapsulated and carried across the provider's MPLS backbone to a remote L2 circuit.

Customers access the service through a L2 interface and have limited opportunities to attack the provider's core. Customer traffic is unlabeled so there is no potential for label spoofing. Spoofing the L2 addressing information is going to affect only the attacker's connection since the PE determines how to handle the incoming traffic based on the interface it arrives and not its source or destination addresses. In a pseudo-wire deployment the incoming

packet can be only transported to the remote L2 circuit, so the destination information on the packet is completely ignored. In L2 VPNs, the destination information on the packet will determine the destination VPN site that will receive the packet, but this site will always be on the same VPN so there is no potential for traffic injection to other VPNs. Only if the PE receives traffic over logical interfaces (i.e. VLANs) and the L2 connection between customer and provider is not secured it may be possible to achieve traffic injection attacks.

Typically Pseudo-wire services require relatively simple network structures and fewer protocols so the dangers from misconfiguration and attacks at the protocol level are smaller. Nevertheless, the trend is towards more complexity as services emerge where pseudo-wires can be setup across multiple providers. The requirement of having pseudo-wires that span the boundaries of provider networks resulted in building the end-to-end pseudo-wire as concatenation of segments that are "stitched" together at various points, mostly the provider boundaries. Recently there have been IETF standardization efforts for the automatic discovery and placement of multi-segment pseudo-wires. Standardization work for protection of pseudo-wires has not even started yet in IETF.

Another important difference between the pseudo-wire service and the L2 and L3 VPN services described below is that some pseudo-wire deployments allow the customers to dynamically create pseudo-wires. This opens the door for resource exhaustion attacks where the customer maliciously or due to faulty operation will attempt to create too many pseudo-wires. Appropriate protocol level and service level limiting mechanisms must be deployed to protect from such attacks. Similar issues exist also on the interfaces between providers, where a provider can dynamically request the creation of pseudo-wires over the network of another provider.

Finally, the pseudo-wire and L2 VPNs services (similarly to the legacy ATM and FR services) do not offer data confidentiality. The customers must use other mechanisms to ensure that their traffic can not be viewed or altered by the provider.

5.3.3.2 Virtual Private LAN Service (VPLS)

VPLS is a L2 VPN service where the provider's network operates as a big L2 switch allowing multiple customer sites to be connected over a single familiar L2 interface (in most cases Ethernet). If the L2 protocol supports broadcast, the MPLS core may use point-to-multipoint LSPs to efficiently carry this broadcast traffic to all destinations.

As in the pseudo-wire case the customers connect to the provider over a L2 interface so they have limited opportunities to attack the provider's network and protocols. The PE determines the VPN of the incoming traffic through the interface the traffic was received on so spoofing L2 addressing information can not lead to traffic injection.

Unlike the L2 VPN case, in VPLS customers may peer with the provider for some protocols (most likely the spanning tree protocol STP). In this case, protection from protocol attacks should be implemented by the provider.

L2 VPN network topologies can get pretty complex since there is support for hierarchical networks so the risk of misconfiguration is relatively high. Furthermore, the issues with the security of the L2 connection between the customer and the provider apply in L2 VPNs as well.

L2 VPNs can also suffer from high processing load on the provider's equipment. Due to the nature of the service provided, L2 VPNs may result in frequent control plane updates, for example ARP resolution for new destination IP addresses. The volume and frequency of these control operations may be considerably higher than those in a L3 VPN, and this increased activity will result in a larger load on the provider's equipment. Some studies have shown that it is relatively easy in a medium size deployment to consume up to 30% of the CPU resources of a high end router with control plane activity. Proper network design and dimensioning of router CPU resources is necessary to ensure that the risk of accidental or malicious DoS is contained.

Note that in some cases, even non-malicious operation can result in DoS like situations. For example in L2 VPNs each time a system moves in the customer's network signalling over the provider's network is required to update the information about the location of its MAC address. Thus, misconfigurations in the customer's network can directly cause control overload in the provider's network, this is much less likely to do accidentally in a L3 VPN service. Also, in L2 VPNs, it is much easier for even an unsophisticated malicious party in the customer's network to mount an attack on the provider's network. For example initiating a high number of pings or telnet sessions to a number of bogus (customer) addresses can result in high signalling overheads in the provider's network as it tries to determine the MAC addresses of these non-existent L3 addresses.

Since in some cases L2 VPNs offer a service identical with that of a L2 switch, certain L2 attacks like ARP spoofing are now possible. Since the MPLS VPN service ensures isolation between customers these attacks can only exist within a single customer network and they do not affect the provider or other customers.

The L2 VPN service does not provide data confidentiality and the customers must use other mechanisms to ensure that their traffic cannot be viewed or altered by the provider.

5.3.3.3 L3 VPNs

L3 VPNs offer a L3 service where customer sites connect with each other over a routed infrastructure. The customer is offered a single L3 interface towards the provider and in many cases the customer maintains a protocol peering with the PE over the customer-provider link. This offers increased opportunities for attacks to the provider's network and its protocols. Providers typically deploy all the measures we discussed in Section 5.3.2.2 in order to protect the PE router.

Spoofed IP address will be routed based on the VPN specific routing table that is determined based on the interface traffic arrives and cannot do any harm to other VPNs. Label spoofing the CE-PE link is not labelled so the PE can protect itself trivially by dropping any packet with labels on it. This is not the case for the CoC and inter-AS applications. In these cases inter-carrier interfaces exchange labelled packets that contain a label stack. The design is such that the downstream label is allocated by the PE and the PE can verify if an incoming label is one that it allocated and if not it drops the incoming packet.

L3 VPNs also have reached very high levels of network complexity due to applications like Carrier-of-Carriers (CoC) and inter-AS where there is a hierarchy of providers and tier-N providers can be themselves customers of a tier-(N-1) provider. This complexity increases the cost and the operational risk of the L3 VPN service. Moreover, it requires coordination between multiple providers which makes it even more complex and error prone. Since different providers treat a lot of information about their networks as trade secrets, cooperation has to be achieved without leaking too much information about the topology or the policies of the providers.

The implementation of an inter-AS L3 VPN service is a good example of the complexity of MPLS based networks. Thanks to the flexibility of the MPLS architecture there are three different options [68] on how to implement the exchange of L3 VPN information between carriers. Option (a) operates like a simple L3 VPN customer-provider connection and it is safe but does not scale. Options (b) and (c) allow exchange of labelled packets, and all the VPN routes go into the same ABRs, so a malicious provider can inject packets in any VPN of the other provider. Options (b) and (c) require exchange of information at the routing protocol level and this also can cause problems with increased complexity and increased probability of misconfiguration as well as revealing too much information about a provider's network.

In L3 MPLS VPNs, a number of BGP mechanisms (route distinguishers, extended communities and route targets) are used to control how VPN BGP routes are distributed across the provider's networks. In addition QoS handling of traffic depends on a complex set of per packet information that includes IP TOS and DHCP bits, MPLS EXP

bits, as well as layer 2 information such as dot1q.p bits. Both providers must agree on common values for all this information if they are to jointly provide a VPN service across their networks. Accomplishing this requires careful configuration of all these values across the provider network boundaries. A lot of this information can be considered proprietary (since for example it can reveal the number of QoS service classes that a given provider implements in its network). To avoid this exposure all the internally used values are converted to values that are considered appropriate for exporting to other providers. This conversion requires more protocol machinery, and even more configuration complexity. In some cases (as for example the settings of QoS bits on transiting packets) this re-mapping has to happen at the forwarding plane of the router, requiring more complex operations and more sophisticated silicon. In certain cases, it is not possible to completely prevent the leaking of some provider internal information. For example in Option (c) the addresses of PE routers are exchanged between providers vs only the ASBR addresses in (b), thus (c) reveals more topology information.

Like in the other MPLS services, L3 VPNs also suffer from L2 interconnection issues (CE-PE and inter-provider links). If one interface that allows labelled packets is unsecured then it may be possible for third parties to inject traffic into VPNs.

A unique service that can be offered only over L3 VPNs is that of Internet access. The customer not only can contact other sites of its own network but can also get Internet access over the same provider interface. In these cases customers must use firewalls and Intrusion Detection Systems to protect their networks since the MPLS VPNs does not offer must more than IP connectivity. In some cases providers sell centralized firewall or IDS services to their L3 customers.

Another issue is the separation of the VPN and the Internet access services. If both are offered over the same infrastructure (PE router and provider-customer links) then it is possible that during an Internet attack the VPN service will also suffer. Providers offer a range of connection strategies (from complete separation of the two services over diverse infrastructure to shared infrastructure) that can provide the service level that the customer expects.

The L3 VPN service does not provide data confidentiality and the customers must use other mechanisms to ensure that their traffic cannot be viewed or altered by the provider.

5.3.3.4 L3 Multicast VPNs

This is a variation of the L3 VPN service where multicast traffic is also supported. This service has become very important due to the popularity of video delivery. Providing this service efficiently requires extensions to the MPLS forwarding plane in order to support Multicast Labels. Furthermore, the label distribution protocols need to be extended to support the creation of point-to-multi-point tunnels (for example extensions to RSVP-TE) or existing multicast protocols need to be extended to offer label distribution. The multicast extensions to the control and forwarding plane do not create any new security problems, but make the potential impact of a successful attack much more serious since more traffic and customers may be affected (for example compromising a live video stream that is delivered to 10000s of viewers can damage a provider's reputation significantly).

Multicast L3 VPNs externally appear like a regular L3 service but now the complexity is dramatically increased since customers will need to run multicast routing protocols, and in some deployments the core will also need to multicast routing protocols. The highly complex cases of CoC and Inter-AS must also be supported. As a result, the design and provisioning of resources in the provider's network becomes significantly more complex that the unicast L3 VPN service. Running a network of such complexity requires very effective tools and high levels of expertise but even then it is not clear which impact this extreme complexity has on the resilience and security of the service.

Fast Re-Route is also significantly more complex in P2MP LSPs. Standardization work in still ongoing in IETF and

interoperability between vendors is unclear at this stage. This has not prevented some providers from deploying FRR protected Video distribution services (based on a equipment from a single vendor).

A unique problem with L3 multicast VPNs is that there are multiple possible ways to build certain elements of the network and there are many "options" when designing the network. This can cause potential problems with configuration and interoperability. Moreover, currently some of the technical details of the L3 multicast VPN implementations are hotly debated in IETF where a vendor war is being waged. This causes lack of clear standards at this moment, making the future of this service unclear.

5.3.3.5 Traffic Engineering/QoS

To the degree that this service is used mostly inside the provider's networks, it is less exposed to external attacks. Since the QoS that traffic will receive inside the MPLS core depends on the QoS information carried in the incoming customer packets, it is possible that a malicious customer would spoof the QoS information of its packets in an attempt to manipulate the service it receives from the core. Such attacks will not be very effective. Providers typically rate limit incoming rates per QoS class, so they will not be exposed to unexpected traffic rates and other customers will not be affected. By spoofing QoS information on its packets the customer can force its traffic to receive lower service or higher service. Both are not very meaningful for the customer given that he is probably paying more when he sends traffic that needs higher levels of QoS.

Things can get more complex when multiple providers must coordinate to provide traffic engineering across their networks. Since traffic engineering requires a good view of the available resources, it is very difficult to accomplish across provider boundaries since providers are not willing to disclose detailed information about their networks. There has been some work that tries to balance the amount of information exposed about the internal network structure with providing enough information to make some useful traffic engineering decisions. Furthermore, RSVP-TE when used for traffic engineering carries the Explicit Router Object (ERO) and the Record Route Object (RRO) which are used to choose the path the the RSVP-TE signalled LSP will take in the network. These objects contain a list of interface of node addresses, and may be hard or even impossible to use across provider boundaries since the source provider may not have enough visibility in the destination provider's network so that it can create an end-to-end ERO and the destination provider may not want to return to the source provider a detailed RRO object. In order to compensate for these restrictions partial signalling is used and the end-to-end LSP is created by "stitching" together LSP segments that are entirely within the boundaries of a single provider. This creates management and configuration complexity and may affect the optimality of the end-to-end TE paths used since the originator of the TE path does not have an accurate global picture of the network and the resource availability in it.

Another twist is that it is possible to provide a TE service to a customer, for example let the customer use RSVP-TE to perform traffic engineering over its own network across all its VPN sites. If this will require the provider to accept labelled packets from the customer, this could provide opportunities for label spoofing attacks from the customer. Such types of services are not very common currently.

5.3.3.6 Fast Re-Route

Fast Re-Route is most commonly used inside the provider's network and in this respect it is similar to the TE service discussed above. The FRR mechanisms must be appropriately extended to support resilient operation over the links between different providers or between customers and providers. This is complicate since FRR typically requires detailed information about the network topology in order to find diverse paths for establishing the various backup LSPs. Since this type of visibility is generally not available, it is very difficult to setup backup LSPs that extend into the network of a peer provider. This makes it particularly difficult to protect against failures of the ASBR routes that are at the endpoints of the inter-provider links. Solutions to these problems exist but they tend to be

specially formulated for the constraints of the inter-provider connection environment and they require complex configuration.

It is also possible to attempt to provide FRR as a service to a customer for use inside its own network. This can give more opportunities for attack to the customer. This kind of service is not very common currently. Instead LDP sessions between customer sites are used over RSVP-TE provider tunnels to allow L3 VPNs to take advantage of the FRR protection and traffic engineering support provided by RSVP-TE. This means that the provider will have to accept labelled packets from the customer and as discussed earlier can create opportunities for label spoofing attacks from the client.

5.3.3.7 OAM

Attacks to the OAM mechanisms of the network can have very serious consequences since they can make entire sections of the network appear out of service. Like in IP networks, the OAM mechanism always carry the risk of DoS attacks (remember all the ICMP based attacks in the IP world) and must be carefully managed by the operators. To avoid DoS attacks the reception and processing of OAM messages must be carefully rate limited and enabled only on the interfaces/LSPs/QoS class where this function is required. This may be difficult for MPLS ping since it relies on the router alert mechanism to reach the router's control plane. Honouring the router-alerts means that the filtering will have to be performed at the control plane but this may be too late since even sending a packet to the control plane and dropping it there may have already consumed too many resources. Not honouring the router-alert may interfere with other services that are based on it, and may affect the correct operation of the MPLS OAM itself. Other ways to protect against DoS attacks include using the TTL in order to ensure that the incoming OAM packets are initiated from nodes that are only one hop way and accepting OAM packets only from neighbours that are authorized/explicitly configured.

Rate limiting can also interfere with the correct operation of always on OAM mechanism like BFD. Depending on the configuration it may be hard to determine the right levels of rate limiting. A more serious concern is that a DoS attack may fail to overload the control plane but may cause legitimate MPLS OAM packets to be dropped so that the OAM layer will conclude that there is a network failure. Luckily, BFD, which is most sensitive to these attacks since it is always on, has mechanisms for authenticating the BFD packets. This will allow the receiver to distinguish valid BFD packets from the ones that are part of a DoS attack. Still, there may be problems with this approach since typically the security processing of received packets does not happen at the forwarding plane but in the control plane, so this may require all BFD packets (including the ones that are part of the DoS attack) to be sent to the control plane. In general effective usage of packet authentication requires router equipment that has appropriate crypto-acceleration hardware either in the linecards or in a co-processor in the controllers.

MPLS ping and VCCV do not have any support for message authentication. Since MPLS ping is an on-demand mechanism it is relatively difficult for an attacker to predict when there will be MPLS ping activity so that it can attack it. If the attacker has the capability to monitor the line though in principle it is possible to observe when MPLS ping activity is occurring and then mount a mini DoS attack by sending enough spoofed OAM traffic so that the OAM rate limiting will cause the legitimate MPLS ping packets to be dropped and result in service outage. Also, note that since the return path of MPLS ping is often over IP, similar attacks are possible at the IP layer. This appears to be a significant hole in the security architecture of MPLS and at this point it is not clear that there have been attempts to address it. Deployment guidelines often recommend to use IPSec tunnels or GRE encapsulation to secure MPLS ping traffic [64].

In order to perform OAM attacks, the attacker must be able to inject packets into the MPLS layer. This may be relatively difficult to do for links that are part of the provider's network, but as discussed earlier, if the L2 link between providers or between provider and customer is not secured, third parties can mount attacks over it. If the

customer does not exchange labelled packets with the provider then there is no issue, since the provider typically makes sure that it drops all the labelled packets arriving by the customer. The situation is more complex at the boundary between providers where OAM traffic must be exchanged and a third party that has compromised the L2 link between providers or a compromised provider can mount OAM attacks. In certain cases where the OAM peers are a single hop away, the TTL can be used to ensure that the OAM packets do not proceed further in the network than their intended destination.

In applications like pseudo-wires, the nodes at the boundary between the L2 network and the MPLS backbone may perform OAM inter-working functions, converting OAM messages from the MPLS to the L2 (eg ATM or FR) format and vice versa. This means that the inter-working node will be exposed to OAM based attacks from the customer, i.e. a OAM message based DoS attack. These attacks do not only affect the boundary node that receives the OAM messages but can also affect the MPLS backbone since these L2 specific OAM messages are typically translated to MPLS OAM messages that are sent across the MPLS backbone to the remote attachment circuit. As in the other cases of DoS attacks, the boundary node must protect itself with careful rate limiting of incoming messages, enabling reception of OAM messages only on selected physical interfaces and circuits and where applicable only from authorized peers. Typically in legacy technologies like ATM and FR, the OAM messages are not secured or authenticated thus it may be impossible to validate them at the receiving node.

Another security complication with MPLS OAM is the potential of information leakage. Pinging certain elements will result in response packets that may contain information that a provider may consider confidential (i.e. certain labels or label stacks). Also the traceroute ability can be used to explore the network topology of a provider. Typically providers limit the ability of other providers/customer to perform these operations in their network. This may not be easy though since it may require deep packet inspection of the OAM packets at the high volume peering links between providers. Also, any interference and selective dropping of OAM traffic has to be considered very carefully since it will affect the effectiveness of the OAM functionality. If a provider disables all MPLS traceroutes through its network, it may not be possible to debug a cross-provider problem. Overall, there exist relatively little understood issues around network information leakage through OAM and OAM across provider boundaries seems to be more of a fine art at this point. The IP/MPLS forum has recently tried to provide a more solid framework with the publication of the MPLS-ICI technical specification [20].

5.4 Deployment scenarios

MPLS is a widely deployed technology. Many of the tier-1 and tier-2 providers in the world deploy MPLS in some form, and MPLS based L3 and L2 VPN services are becoming a common offering in many countries. Many examples of existing deployments can be found in the provider and vendor literature [3, 9, 2, 4, 5, 6, 7, 8, 10]. Very large global providers like NTT, Verizon, and France Telecom have deployed MPLS protection features.

Below we present some typical MPLS deployment scenarios.

5.4.1 Point-to-point legacy services

A customer relies on a point-to-point ATM connection between two sites in different parts of the country. The provider's previous network uses a SONET infrastructure to establish SONET circuits between the two sites and the ATM circuit is established over these circuits. Configuration of the SONET circuits is through the OSS system. The provider now migrates to a nation-wide IP/MPLS core. The two customer sites are now connected to PE routers but without changing the details of their ATM circuits. Two LSPs are established between the PE, one in each direction. Incoming ATM traffic is encapsulated into the LSP that in this case acts as a pseudo-wire. The provider and customer negotiate a SLA with a delay bound of 40 milliseconds, traffic loss less than .01% and availability of 99.99%. The provider's OSS system determines the paths for the two LSPs based on the network load and the SLA. In order

to meet the latency requirements the provider picks a path with few hops. RSVP-TE is used for signalling the two LSPs across the predetermined path. In order to achieve the high availability requirements the provider chooses to implement end-to-end protection of the two LSPs, and the PEs will switch to a backup LSP if one of the primary LSPs fails. The customer does not require protection from PE failures or failures of the customer to PE link so there is no need for redundant connections between the customer and the provider. The provider enables OAM inter-working for the pseudo-wire so that failures on the ATM circuit on the remote site will be communicated to the local site. The customer is charged a flat fee per month and the size of the pseudo-wire if fixed.

In this scenario, MPLS enables the provider to offer a legacy service over a less expensive core network that can be used for other services. MPLS mechanisms ensure that the service has the same resilience as when implemented over the legacy network. The OAM support of MPLS hides the fact that the ATM circuit is actually composed of two access ATM circuits connected with a MPLS LSP.

5.4.2 L3 VPN service with QoS

The provider offers a L3 VPN service. The core MPLS network implements a full mesh of explicitly routed LSPs between all the PEs. These paths are protected by backup LSPs that are also explicitly routed so they are diverse from the LSPs they are protecting. End-to-end protection is implemented with the PE switching to the backup LSP if the primary LSP fails. Customers connect to the PE routers over an IP interface. The provider implements three classes of service on its core: Standard (for e-mail, file transfer, and non-critical Internet access), Priority (for critical Internet access, point-of-sale, and streaming video), and Near Real-time (for voice over IP and video-conferencing). The customer and provider negotiate how customer IP traffic will be mapped to these three classes of service based on the DSCP markings on the IP packets and an SLA for each class of service. The provider configures the PE with the QoS mapping rules that will map the DSCP markings of the customer traffic into the EXP bit settings inside the core. The customer is charged based on the amount of traffic sent for each traffic class. All customer traffic is put in the same LSP to the destination PE irrespective of its traffic class. The provider uses the OSS system and off-line optimization software to determine the routes for the full mesh of LSPs to ensure there is no congestion in the network even when there are failures so that it can satisfy all the customer SLAs. P routers inside the core queue and schedule traffic based on the setting of the EXP bits in the MPLS headers. If a customer requires a very strict SLA the provider can direct this customer's traffic in specially established LSPs that are placed over links with light loads. The provider charges more for this service.

In this scenario MPLS again allows the provider to offer multiple services over a shared backbone network but also gives the provider flexibility on the QoS services it offers. The provider can support both, aggregated SLAs that use a shared set of LSPs, or specialized SLAs that use per customer LSPs for more optimized traffic handling. In this way the provider can maximize its revenue.

5.4.3 Video distribution

A television network buys nationwide data services from a MPLS provider. 100s of live TV streams must be distributed from the national headquarters to multiple regional stations. The provider sets up a Point-to-multipoint LSP for efficient distribution of this data stream that can be multiple gigabits/second. The customer requires a very strict SLA with minimum jitter and loss and less than 50 millisecond of disruption in case of network failures. The provider implements FRR to ensure fast traffic protection. The provider must ensure that in cases of failure the traffic flowing over the backup LSPs will not cause network congestion. They use off-line optimization software to careful determine the placement of the primary P2MP LSPs and its backups.

In this scenario, MPLS allows the provider to offer a very high bandwidth and strict SLA service over a mesh network using multicast. Multicast is a major advantage over the existing transport technologies that are strictly point-to-

point. Multicast allows for very efficient usage of the network resources of the provider and is dramatically cheaper for the customer that would otherwise have to pay for a large number of P2P circuits between the headquarters and the regional stations.

5.4.4 MPLS in the aggregation network

A national cellular provider uses metropolitan TDM networks to connect its base stations in each region to its core network. The core network has already been converted to MPLS. The provider migrates to a much less expensive Metro-Ethernet network that supports MPLS. Some base stations also migrate to Ethernet but the remaining base stations continue to use TDM. The old base stations continue to use TDM connections to the first metro-Ethernet node and this node encapsulates the TDM traffic into a MPLS pseudo-wire that terminates at the PE of the core network. Pseudo-wires are setup using RSVP-TE and end-to-end protection is used to handle network failures. The metro-Ethernet aggregation network implements a L2 VPN service for connecting the newer base stations to the core. Ethernet traffic from the new base stations is received by the first metro-Ethernet node that encapsulates it into a base station specific VLAN and sends it to the core. The L2 VPN service also uses RSVP-TE signalled LSPs with end-to-end protection. The core's PE routers terminate both the TDM pseudo-write traffic from the old base stations as well as the Ethernet traffic from the new base stations and forward it over the core.

In this scenario MPLS allows the provider to replace existing costly infrastructure with a much cheaper metro-Ethernet network. In addition, MPLS allows the coexistence of older generation TDM base stations and newer Ethernet based ones allowing the provider to gradually phase out its older equipment.

5.5 Summary

The above brief study of the security and resilience of MPLS showed that:

- The MPLS technology provides mechanisms that can be used to improve the resilience of provider networks and consequently the resilience of the service offered to customers. MPLS supports mechanisms that can protect from link and node failures inside the backbone network as well as at the interconnection points between providers or customers and providers. MPLS also provides extensive Operation, Administration and Management (OAM) functionality that can be used to verify the levels of service offered and localize and repair failures fast.
- One of the major drivers for MPLS deployment is its ability to carry multiple types of traffic over the same backbone. As more traffic types share the same network resources the potential for unwanted interactions between them increases and problems in one of the traffic types can negatively affect other traffic types. Strong traffic isolation and management mechanisms are needed in order to prevent these interactions. Traffic management mechanisms also allow providers to control how traffic flows over their networks. This ensures that changes in network traffic and failures will not compromise the quality of the offered services.
- A necessary dimension of providing a resilient service is the validation that the service operates properly, the quick detection of failures and the equally quick service restoration that involves fast fault isolation. The MPLS architecture provides mechanisms for OAM.
- MPLS uses an IP based infrastructure, this by itself makes it more vulnerable to attacks when compared with earlier implementations that used isolated legacy technologies. These security risks can be mitigated with careful administration of the systems involved.
- MPLS deployments are built by combining existing building blocks, i.e. access technologies, routing and label distribution protocols, IPsec and more. This makes MPLS a very flexible technology but also increases the costs and risks of configuration and operation.

- Although the basics of the MPLS technology are a decade old, to a large extent MPLS is still a young technology, many MPLS services like VPLS or Multicast MPLS have emerged in the last 2-3 years, and many of the related protocols are still in the standardization phase in IETF. For some of these services (most notably multicast MPLS L3 VPNs) there have been conflicting implementation proposals in the IETF (aligned across vendor lines) and this has caused confusion and slower adoption. Security was not one of the considerations built into the development of the MPLS standards and although it has received a lot of attention recently to some degree it is an after-thought. Some security work on MPLS protocols (keying support in RSVP-TE, security in OAM) is still in its early standardization phases and it will need time to mature.
- MPLS based services are complex. Even when providers can merge multiple disparate networks into a single MPLS/IP core, they still need to support the legacy protocols on the edges of their network and do the necessary adaptation functions between the legacy data plane and the MPLS/IP data plane. In addition, the ability to offer services over multiple different provider networks generates a lot of configuration and operational complexity in OAM, QoS mapping, TE and FRR. Even today, offering TE and FRR across provider boundaries is not well understood. Complexity can endanger both resilience and security of the network and the services it supports and it will remain one of the main challenges as MPLS becomes even more popular and supports an ever expanding range of services.
- Not surprisingly, security is a financial issue for service providers. From the results of various provider surveys, most notably [53], providers operate based on a "management vs. risk" trade-off. If certain technologies and protocols that can provide increased security are complex (and thus expensive to operate) providers may be willing to live with reduced security. Furthermore, if certain security functions are not uniformly supported across all vendors, providers may take the "least-common-denominator" approach since this allows them to have uniform configuration and reduced management complexity, even if in certain equipment they could achieve more secure operation. Standardization and adoption of new standards by vendors seems to be a very important factor for improving the security and resilience of MPLS networks.

Overall, MPLS is a successful technology that has been embraced by a large number of providers due to its ability to provide multiple services over a shared infrastructure and its open and standards based signalling. The deployment of MPLS is strongly driven by commercial motivations and it appears to be poised for further expansion. Nevertheless, MPLS has some important shortcomings. Certain protocol issues like security operations in MPLS signalling protocols and OAM, FRR for M2MP LSPs and pseudo-wires are not fully standardized yet and operational issues like OAM/TE/FRR/QoS across provider boundaries are still very complex to configure and operate and not well understood. The overall resilience of an MPLS/IP network still has some way to go until it catches up with the resilience provided by the incumbent technologies. While the standardization issues will be resolved over time, the configuration and operational complexity will remain a challenge. Recently there has been standardization and research work that shows how non-MPLS IP networks can achieve some of the main advantages of MPLS, i.e. fast data protection and traffic engineering. This shows that IP will be a potential competitor to MPLS in the future and if the current shortcomings of MPLS are not resolved, MPLS's dominance in the near term future is not assured.

5.6 Terminology

This is a brief explanation of the common MPLS abbreviations and acronyms

MPLS	Multi Protocol Label Switching
QoS	Quality of service
L2	Layer 2 of the OSI model
L3	Layer 3 of the OSI model
LSR	Label Switch Router
P	Provider router, a router inside the MPLS core
PE	Provider Edge router, a router at the edge of the MPLS core
RSVP-TE	Resource reservation protocol - Traffic Engineering, a signalling and label distribution protocol used to setup LSPs
LDP	Label Distribution Protocol, a signaling and label distribution protocol used to setup LSPs
MP-BGP	Multi Protocol - Border Gateway Protocol
VPN	Virtual Private Network
VPLS	Virtual Private LAN Service
L3 VPN	Layer 3 VPN
L2 VPN	Layer 2 VPN
TE	Traffic Engineering
ICMP	Internet Control Message Protocol
ICMPv6	Internet Control Message Protocol for IPv6
RADIUS	Remote Authentication Dial In User Service
DNS	Domain Name System
NTP	Network Time Protocol
SNMP	Simple Network Management Protocol
GRE	Generic Routing Encapsulation, an IP tunnelling protocol
Tier-2 provider	a smaller regional provider
Tier-1 provider	a large national or international provider
CoC	Carrier-of-Carriers, a network c onfiguration where a tier-2 provider is a customer of a tier-1 provider
Inter-AS	a network configuration where two providers are peers
Extranet	a network configuration where two separate customer networks have limited connectivity
SONET	Synchronous Optical networking, a Layer-1 transport technology that offers very effective data protection
VPLS	Virtual Private LAN Service
ASBR	Autonomous System Border Router, the router at the border of two provider networks
OAM	Operations Administration and Management
STP	Spanning Tree protocol
IETF	Internet Engineering Task Force the main standardization body for MPLS
FRR	Fast Re-Route, an MPLS mechanism that allows for very fast traffic protection
DSCP	Differentiated Services Code Point, a set of bits in the IP header used for packet classification

6 Conclusions

6 Conclusions

In this study, three key technologies, namely IPv6, DNSSEC (that are at the pre-deployment stage) and MPLS have been outlined. An analysis of their resilience properties and the current deployment status has been performed. The key findings are the following:

- All three technologies are likely to help improve resilience to some degree. IPv6 will address a noteworthy source of vulnerabilities, by making it somewhat harder to launch opportunistic attacks such as worms, and will deter reconnaissance probing; DNSSEC will effectively prevent spoofing and malicious domain takeovers; MPLS provides better traffic isolation and control.
- In all three technologies some of the resilience features may be overestimated by advocates. For instance, the mandatory nature of IPsec in IPv6 does not help to address the key problems in IPsec deployment compared to how it is used today in the IPv4 world; IPv6 doesn't help at all against the rise of targeted attacks; reconnaissance and opportunistic attacks are not completely eradicated; critical services already use SSL/TLS for server authentication, so DNSSEC will only help to expand the same flavour of security to less critical services that are not "secure by default".
- In some cases, there are valid concerns about increased risks that those technologies present to resilience. For instance, the wide address space in IPv6 may allow attackers to improve their techniques of escaping address blacklisting used for SPAM and DDoS prevention; the introduction of IPv6 may create privacy problems with tracking the location of users in the network and the known remedy to this issue (address hopping and randomization) may further amplify the problems of blacklisting-based security solutions.
- It has been found that all three technologies have undergone extensive evaluation and trial deployments. However, some important issues may surface during the deployment on the global scale and some non-technical issues, such as the DNSSEC root signing, are still to be resolved.

7 Bibliography

7 Bibliography

[1] A DNSSEC incident in Sweden (.se secure zone). http://www.dnssec-deployment.org/wg/materials/20071107/dnssec_incident_en.pdf.

[2] Case Study: BAT Italia Boosts Security at Factory and New Headquarters. http://www.cisco.com/en/US/prod/collateral/routers/ps5854/prod_case_study0900aecd8039fc1b_ps6557_Products_Case_Study.html.

[3] Case Study: Bell Canada. http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6557/prod_case_study0900aecd80312171.pdf.

[4] Case Study: BT Group MPLS-Based IP VPNs Built on Cisco IOS Software. http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6557/prod_case_study0900aecd80312171.pdf.

[5] Case Study: Cisco IP/MPLS Helps BT Infonet Reap the Benefits of an Interprovider. http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6557/prod_case_study0900aecd803701eb.html.

[6] Case Study: Global One - First to Offer Global MPLS-Based IP VPN Service. http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6557/prod_case_study0900aecd80312173.pdf.

[7] Case Study: Midwestern Health Insurance Company Builds High-Speed Network. http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6557/case_study_c36_373156.pdf.

[8] Case Study: Sprint-Layer 2/Layer 3 Services Converged Over Common IP Backbone. http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6557/prod_case_study0900aecd80312074.pdf.

[9] Case Study: State Street Corporation Banks on a High-Availability MPLS Network. http://www.cisco.com/en/US/prod/collateral/routers/ps5854/prod_case_study0900aecd80424090_ps6557_Products_Case_Study.html.

[10] Case Study: Webpartner Chooses Cisco VPLS to Interconnect Enterprise Customers. http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6557/prod_case_study0900aecd803f861e.pdf.

[11] CZ.NIC. http://www.nic.cz.

[12] Dan Kaminsky - Short Bio. http://www.doxpara.com/?page_id=1159.

[13] Euro6ix: European ipv6 internet exchanges backbone. http://www.euro6ix.org.

[14] Hacking coffee makers. http://www.securityfocus.com/archive/1/493387.

[15] IP Fast Reroute Framework. draft-ietf-rtgwg-ipfrr-framework-09. http://www.ietf.org/internet-drafts/draft-ietf-rtgwg-ipfrr-framework-09.txt.

[16] Ipv6 car project. http://newsroom.cisco.com/dlls/ts_102003.html.

[17] Ipv6.com - the source for ipv6 information, training, consulting and hardware. http://www.ipv6.com.

[18] Issues with existing Cryptographic Protection Methods for Routing Protocols. http://www.ietf.org/internet-drafts/draft-ietf-opsec-routing-protocols-crypto-issues-00.txt .

[19] Large-Scale International IPv6 Pilot Network. http://www.6net.org/publications/.

[20] MPLS-ICI technical specification. http://www.ipmplsforum.org/tech/IPMPLSForum19.0.0.pdf.

[21] Networks running IPv6. http://bgp.he.net/ipv6-progress-report.cgi.

[22] Nic.pr DNSSEC. http://www.dnssec.nic.pr.

[23] Port 53 - Shyam Seshadri's blog on Windows DNS and more. http://blogs.technet.com/sseshad.

[24] Projects-DISI. https://www.ripe.net/projects/disi/keys.

[25] Register.bg. https://www.register.bg.

[26] Registro.br. http://www.registro.br.

[27] .SE. http://www.iis.se.

[28] SecSpider - A globally distributed polling system. http://secspider.cs.ucla.edu.

[29] Security Best Practices Efforts and Documents. http://tools.ietf.org/id/draft-ietf-opsec-efforts-08.txt.

[30] Slashdot — world's smallest ipv6 stack by cisco, atmel, sics. http://tech.slashdot.org/tech/08/10/15/1839209.shtml.

[31] Summer Olympics 2008 - IPv6 Website. http://en.beijing2008.cn/ipv6/.

[32] The wait is over for IPv6. http://archives.cnn.com/2001/TECH/industry/06/11/IPv6.wait.idg/index.html.

[33] Top 10 features that make ipv6 greater than ipv4. http://archives.cnn.com/2001/TECH/industry/06/11/IPv6.wait.idg/index.html.

[34] Wold Wide DNSSEC deployment. http://www.xelerance.com/dnssec.

[35] M. Andrews and S. Weiler. The DNSSEC Lookaside Validation (DLV) DNS Resource Record. RFC 4431 (Informational), February 2006. http://www.ietf.org/rfc/rfc4431.txt.

[36] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. DNS Security Introduction and Requirements. RFC 4033 (Standards Track), March 2005. http://www.ietf.org/rfc/rfc4033.txt.

[37] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. Protocol Modifications for the DNS Security Extensions. RFC 4035 (Standards Track), March 2005. http://www.ietf.org/rfc/rfc4035.txt.

[38] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. Resource Records for the DNS Security Extensions. RFC 4034 (Standards Track), March 2005. http://www.ietf.org/rfc/rfc4034.txt.

[39] A. Barbir, S. Murphy, and Y. Yang. Generic Threats to Routing Protocols. RFC 4593 (Informational), October 2006. http://www.ietf.org/rfc/rfc4593.txt.

[40] Ray Bellis and Lisa Phifer. Test Report: DNSSEC Impact on Broadband Routers and Firewalls. Technical report, September 2008. http://download.nominet.org.uk/dnssec-cpe/DNSSEC-CPE-Report.pdf.

[41] D. Conrad. Indicating Resolver Support of DNSSEC. RFC 3225 (Standards Track), December 2001. http://www.ietf.org/rfc/rfc3225.txt.

[42] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Specification. RFC 2460 (Draft Standard), December 1998. Updated by RFC 5095.

[43] T. Dixon. Comparison of Proposals for Next Version of IP. RFC 1454 (Informational), May 1993. http://tools.ietf.org/html/rfc1454.

[44] D. Eastlake and C. Kaufman. Domain Name System Security Extensions. RFC 2065 (Standards Track), January 1997. http://www.ietf.org/rfc/rfc2065.txt.

[45] Karen Evans. Securing the Federal Government's Domain Name System Infrastructure, August 2008. http://www.whitehouse.gov/omb/memoranda/fy2008/m08-23.pdf.

[46] Michael P. Gallaher and Brent Rowe. IPv6 Economic Impact Assessment, Oct 2005. http://www.6journal.org/archive/00000282/01/report05-2.pdf.

[47] R. Gieben. DNSSEC in NL, January 2004. http://www.nlnetlabs.nl/downloads/publications/dnssec/dnssecnl/secreg-report.pdf.

[48] W. Hardaker. Use of SHA-256 in DNSSEC Delegation Signer (DS) Resource Records (RRs). RFC 4509 (Standards Track), May 2006. http://www.ietf.org/rfc/rfc4509.txt.

[49] R. Hinden and S. Deering. Internet Protocol Version 6 (IPv6) Addressing Architecture. RFC 3513 (PROPOSED STANDARD), 2003. http://tools.ietf.org/html/rfc3513.

[50] R. Hinden and B. Haberman. Unique Local IPv6 Unicast Addresses. RFC 4193 (PROPOSED STANDARD), 2005. http://tools.ietf.org/html/rfc4193.

[51] Hyun-Ho Choi and Dong-Ho Cho. Mobility management based on mobile IP in mixed IPv4/IPv6 networks. In *Proceedings of Vehicular Technology Conference*, volume 3, 2003.

[52] D. Johnson, C. Perkins, and J. Arkko. Mobility Support in IPv6. RFC 3775 (Proposed Standard), June 2004. http://www.ietf.org/rfc/rfc3775.txt.

[53] M. Kaeo. Operational Security Current Practices in Internet Service Provider Environments. RFC 4778 (Informational), January 2007. http://www.ietf.org/rfc/rfc4778.txt.

[54] Merike Kaeo, David Green, Jim Bound, and Yanick Pouffary. IPv6 Security Report. http://www.ipv6forum.com/dl/white/NAv6TF_Security_Report.pdf, July 2006.

[55] David Karig and Ruby Lee. Remote denial of service attacks and countermeasures, 2001. http://palms.ee.princeton.edu/PALMSopen/karig01remote.pdf.

[56] S. Kent and K. Seo. Security Architecture for the Internet Protocol. RFC 4301 (Informational), December 2005. http://tools.ietf.org/html/rfc4301.

[57] O. Kolkman and R. Gieben. DNSSEC Operational Practices. RFC 4641 (Informational), September 2006. http://www.ietf.org/rfc/rfc4641.txt.

[58] Olaf M. Kolkman. Measuring the resource requirements of DNSSEC. ftp://ftp.ripe.net/ripe/docs/ripe-352.pdf, October 2005.

[59] Craig Labovitz. 2008 Worldwide Infrastructure Security Report. http://asert.arbornetworks.com/2008/11/2008-worldwide-infrastructure-security-report.

[60] B. Laurie, G. Sisson, R. Arends, and D. Blacka. DNS Security (DNSSEC) Hashed Authenticated Denial of Existence. RFC 5155 (Standards Track), March 2008. http://www.ietf.org/rfc/rfc5155.txt.

[61] F. Le Faucheur M. Behringer. Applicability of Keying Methods for RSVP Security, July 2008. http://tools.ietf.org/html/draft-ietf-tsvwg-rsvp-security-groupkeying-01.

[62] Eric Marin. Ipv6 security. In *In Proceedings of the 1st 6NET Workshop*, 2003. http://www.6net.org/events/workshop-2003/marin.pdf.

[63] Paul V. Mockapetris. Keynote "DNS attacks and user protection", November 2008. http://www.enisa.europa.eu/sta/files/ws2008/pvm_enisa.pdf.

[64] T. Nadeau and C. Pignataro. Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires. RFC 5085 (Proposed Standard), December 2007. http://www.ietf.org/rfc/rfc5085.txt.

[65] Eric Osterweil, Michael Ryan, Dan Massey, and Lixia Zhang. Quantifying the operational status of the dnssec deployment. In *IMC '08: Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, pages 231–242, New York, NY, USA, 2008. ACM.

[66] Y. Rekhter. An Architecture for IP Address Allocation with CIDR. RFC 1518 (Historic), September 1993. http://tools.ietf.org/html/rfc1518.

[67] M. Richardson. A Method for Storing IPsec Keying Material in DNS. RFC 4025 (Standards Track), February 2005. http://www.ietf.org/rfc/rfc4025.txt.

[68] E. Rosen and Y. Rekhter. BGP/MPLS IP Virtual Private Networks (VPNs). RFC 4364, December 2003. http://www.ietf.org/rfc/rfc4364.txt.

[69] Louis Senecal. Layer 2 Attacks and Their Mitigation. http://www.cisco.com/web/CA/events/pdfs/L2-security-Bootcamp-final.pdf.

[70] Bill Cheswik Steven M. Bellovin and Angelos D. Keromytis. Worm propagation strategies in an ipv6 internet. *;login: The Usenix Magazine*, 31(1), 2006.

[71] Katsuyasu Toyama, Keisuke Ishibashi, Tsuyoshi Toyono, Masahiro Ishino, Chika Yoshimura, and Kazunori Fujiwara. DNS Anomalies and Their Impacts on DNS Cache Servers. http://www.nanog.org/mtg-0410/pdf/toyama.pdf.

[72] Gregory Travis, Ed Balas, David Ripley, and Steven Wallace. Analysis of the sql slammer worm and its effects on indiana university and related institutions. http://www.anml.iu.edu/downloads/SLAMMER.pdf.

[73] H. Tschofenig and R. Graveman. RSVP Security Properties. RFC 4230 (Informational), December 2005. http://www.ietf.org/rfc/rfc4230.txt.

[74] K. Varadhan, S. Hares, and Y. Rekhter. BGP4/IDRP for IP—OSPF Interaction. RFC 1745 (Historic), December 1994. http://www.ietf.org/rfc/rfc1745.txt.

[75] P. Vixie. Extension Mechanisms for DNS (EDNS0). RFC 2671 (Standards Track), August 1999. http://www.ietf.org/rfc/rfc2671.txt.

[76] Michael H. Warfield. Security Implications of IPv6. http://documents.iss.net/whitepapers/IPv6.pdf, 2003.

[77] S. Weiler. DNSSEC Lookaside Validation (DLV). RFC 5074 (Informational), November 2007. http://www.ietf.org/rfc/rfc5074.txt.

[78] B. Wellington and O. Gudmundsson. Redefinition of DNS Authenticated Data (AD) bit. RFC 3655 (Standards Track), November 2003. http://www.ietf.org/rfc/rfc3655.txt.

[79] Ollie Whitehouse. War nibbling: Bluetooth insecurity, 2003. http://www.wardriving.ch/hpneu/blue/doku/atstake_war_nibbling.pdf.

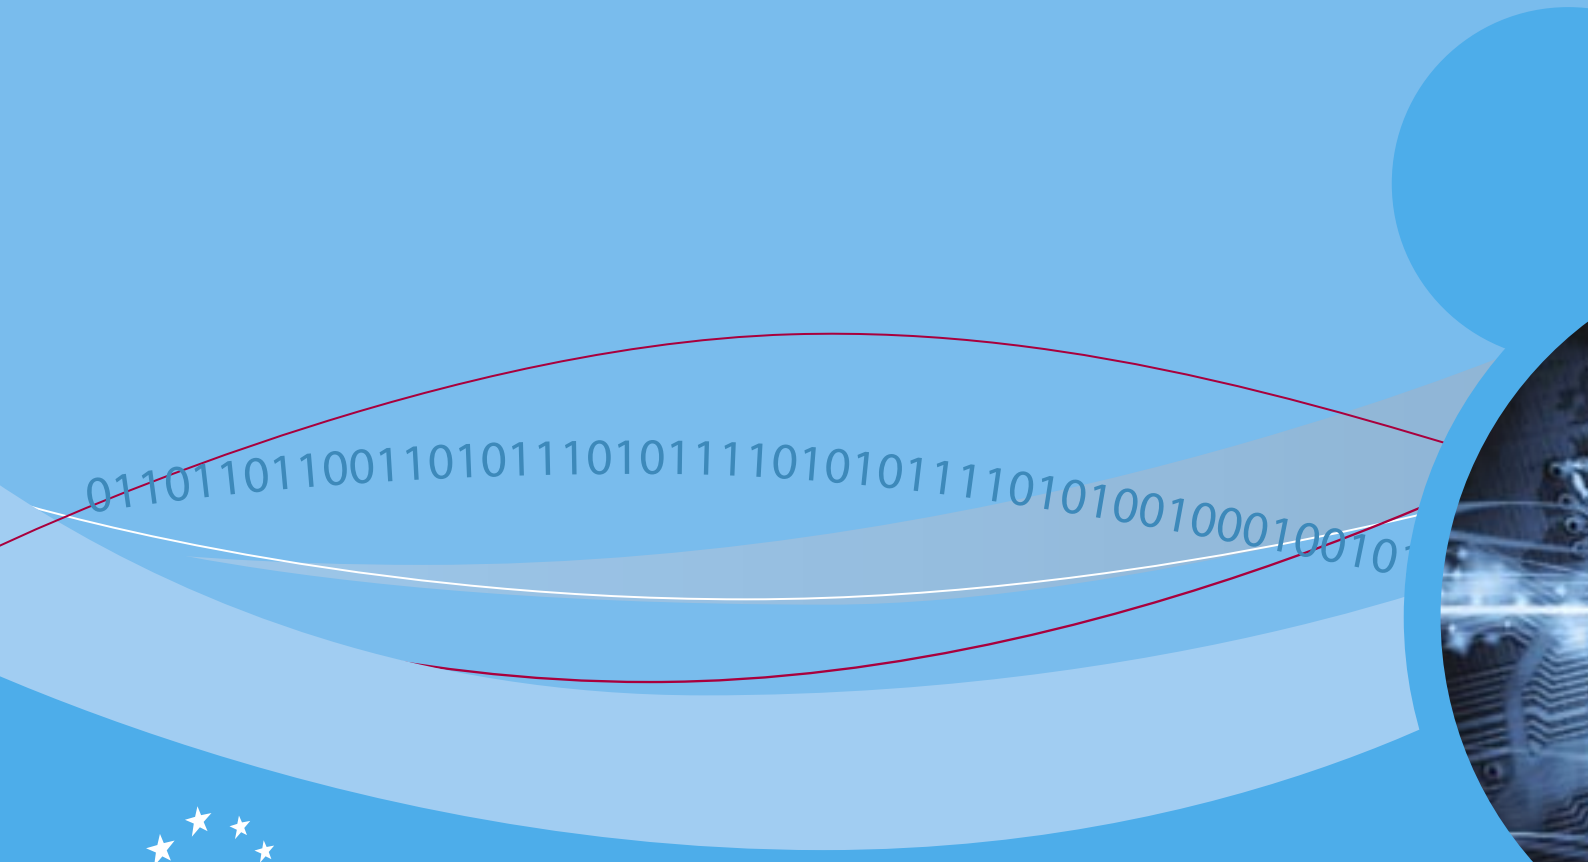

enisa
European Network and Information Security Agency

P.O. Box 1309 71001 Heraklion - Crete - Greece
Tel: +30 28 10 39 1280, Fax: +30 28 10 39 1410
Email: resilience@enisa.europa.eu