Juniper NETWORKS

The IPv6 World Report Series, Volume 4 February 2008

AN IPV6 SECURITY GUIDE FOR U.S. GOVERNMENT AGENCIES

Table of Contents

1 Introduction

1.1 Overview

For well over a decade, the Internet has been proven to be one of the primary cornerstones of the global information technology (IT) infrastructure that has reshaped how civilization communicates and collaborates. The impacts of this affordable and readily available technology have been felt globally and are fundamentally changing the economic foundation of many countries. The underlying technology that drives the Internet, the Internet Protocol (IP), was developed nearly 30 years ago and was never intended to support the robust number of applications and demanding requirements that are continuing to drive its expansion today.

Today, the convergence of voice, video, and data is occurring within the enterprise, and the IP-based infrastructure has become the underlying engine that allows advanced capabilities to be quickly developed and deployed to support a wide range of U. S. Government (USG) capabilities. Many agencies are moving toward the introduction of next-generation systems to support collaborative architectures, geospatial application, net-centric warfare, mobility, continuity of operations (COOP), as well as other numerous applications to better suit their mission.

The issue facing agencies, and the world, is that the current version of IP, IPv4, has nearly reached the end of its useful life cycle and the cost and resources to continue supporting it will become exponentially greater over the next decade. Thus, in 2005, the Office of Management and Budget (OMB) issued a policy requiring all agencies to begin activities for the transition to the next generation of IP, namely IPv6. The transition to IPv6 is not an easy task due to the fact that IPv6 is not backward compatible with IPv4. Thus, careful planning and the use of transition mechanisms will be required throughout the transition cycle to maintain interoperability.

One challenge agencies must face while transitioning to IPv6 is in the area of security. While IPv6 is not directly compatible with its predecessor, it poses many of the same risks associated with IPv4. In addition, IPv6 offers a number of new capabilities that could potentially offer additional vulnerabilities and threats to agencies. However, if properly implemented, IPv6 has the potential to provide a foundation for creating a secure infrastructure for an agency's enterprise as well as the Internet as a whole.

1.2 Scope

This volume of the World Report Series focuses on providing USG agencies with the information that is critical as they plan their transition to IPv6. While the information in this report is useful and valuable to any organization planning to transition to IPv6, USG agencies will find it specifically tailored to their needs.

This volume covers numerous aspects of security related to the transition. In addition to providing a high-level overview of the core concepts of IPv6, it goes into detail on the USG wide policies and the planning that must be accomplished to ensure a successful and secure transition. This volume also goes into greater technical detail not only on the underlying protocols and products that are part of IPv6, but also on related technologies that should be considered when looking at a holistic enterprise security approach.

1.3 Layout

This document is arranged to provide the reader with a progressively deeper understanding of security within IPv6 enterprises. The initial sections provide a higher-level of understanding on the challenges and benefits of IPv6 environments, while the latter sections go into greater detail about the protocols and specific security devices. The sections contained within this document include:

- **Section 1 Introduction:** This section provides the reader with the background and context of the information contained within this volume.
- **Section 2 Overview of IP Security:** This section provides a high-level overview on security for IP-based networks and a core understanding of how security in IPv6 enterprises will be a paradigm shift from how security is currently implemented today.

- **Section 3 Planning for Security During the IPv6 Transition:** This section provides the reader with details on what is required to support the security planning portion of the transition to IPv6. This section not only focuses on meeting near term milestones, but what must be considered as the entire enterprise transitions to IPv6.
- **Section 4 Security within IPv6 Networks:** This section looks at the specific protocols and mechanisms necessary to implement security within an IPv6 environment.
- **Section 5 Conclusion:** This section concludes the volume with a summary.

2 Overview of IP Security

When the original version of Transmission Control Protocol (TCP)/IP was developed, it had very little security inherent to the protocol. In fact, when it was first created, it was developed to operate over protected networks, many of which, if enhanced security was required, would utilize physical layer security encryption mechanisms. After the Advanced Research Project Agency (ARPA) transferred control of the Internet (at the time it was called the ARPANET) to the National Science Foundation (NSF) in 1986 (NSF renamed it the NSFNET), the number of hosts connected to the Internet began to grow exponentially. Security quickly became a major concern as new and un-trusted systems began connecting to the Internet. During the development of IP next generation (IPng), which eventually became IPv6, security requirements were included in the design of the protocol.

2.1 IP-Based Security Today

IP security for enterprises today is primarily boundary focused. The general philosophy is to turn the enterprise into an enclave in which all incoming and outgoing communication channels are tightly controlled and protected. If the enterprise is geographically dispersed, then virtual private networks (VPNs) are employed to create an enclave-to-enclave environment. The common view is that internal users are good, Internet hosts are treated as hostile, and a firewall and other IP based security perimeter devices will provide the security protection necessary, as depicted in Figure 1.

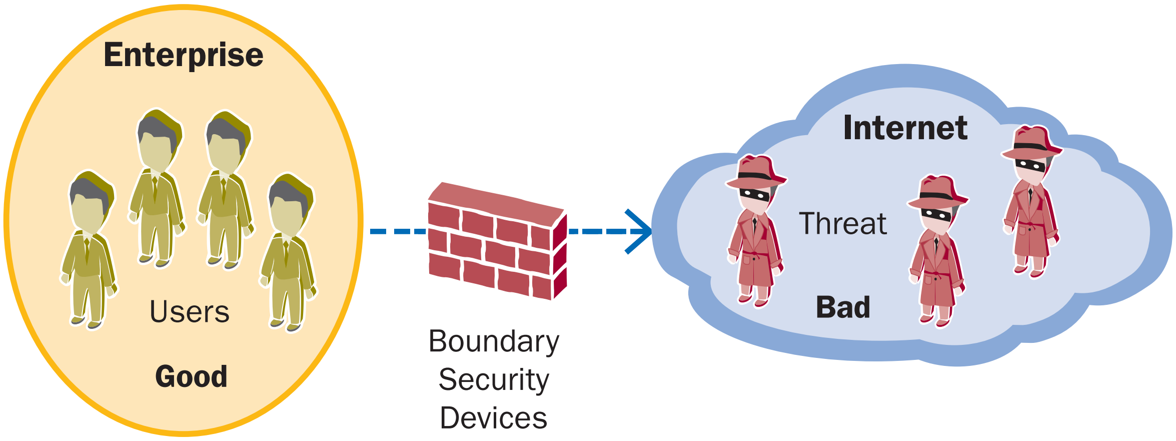

Figure 1: Current View of IP Security

While this approach provides a convenient set of assumptions to work from when developing an enterprise security posture, it is far from reality. Operational experience has proven that attacks on enterprise resources and assets are just as likely (and probably more likely) to occur from internal sources than from the Internet. In fact, internal-based attacks are considered potentially more devastating since the majority of them go unnoticed, and due to limited internal monitoring and logging, it is normally unlikely to fully understand the extent of the compromise. It should be noted that all internal attacks are not malicious actions on an insider's part. The attacks could be generated by mistake, be from corrupt software or viruses accidentally loaded onto user's equipment, or from unaware users being influenced by social hacking. Thus, even significant vetting of the enterprise staff does not always create a trouble-free environment.

Just as it is a fallacy to argue that all internal users and resources should be considered good, it is just as false to argue that all users and resources external to the enterprise should be considered bad. As systems continue to become more integrated, the boundary between what employees, customer, and partners need to access electronically continues to be blurred. It is very common in today's environment for agencies and companies to allow their customers to "self-service" many of their needs, from setting up accounts, filling out and filing forms to making financial transactions. In addition, there is a strong push to decentralize the enterprise to support more collaborative environments, prepare for continuity of operations (COOP), support telework and increase overall organizational efficiency. Productivity within the U.S. has grown tremendously over the past two decades partly due to the ability of skilled resources to effectively complete their work while on travel and at home. Figure 2 presents a more realistic view of the security environment today.

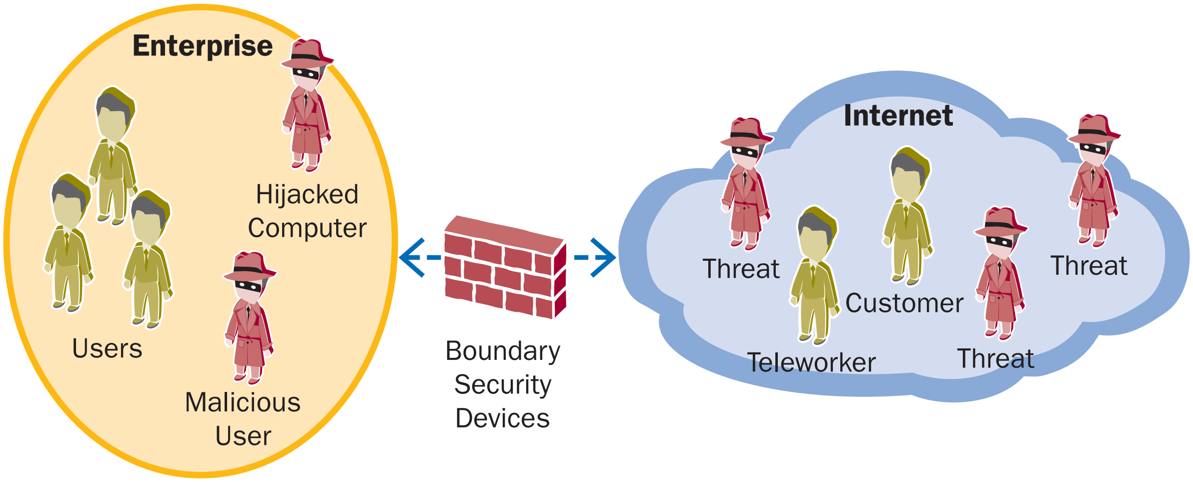

Figure 2: Realistic IP Security Landscape

The use of VPNs have become a common approach to linking geographically disparate offices or users for communication. This entails creating secure tunnels across the Internet, normally using encryption. This provides protection for enterprise traffic as it crosses the Internet. The quality, price point, and manageability of VPNs have become so attractive that many agencies and commercial organizations turn to them to create their own overlay networks across public Internet carriers not only for data traffic, but also to delay sensitive video and voice traffic. Figure 3 shows and example of a VPN network.

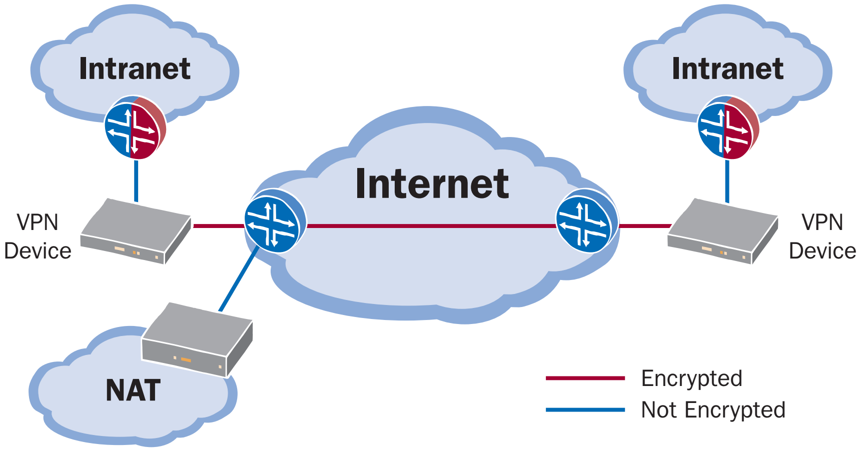

Figure 3: Example of a VPN

Most widely utilized security approaches today focus on prevention as opposed to enablement. The general security practice is to "Deny All" and then allow specific actions to occur. This approach fits well with the concept of turning the enterprise into a protected enclave with barriers all around and only a controlled gate for communications to enter or leave. However, outside of the security and IT offices, most users view security as a "necessary evil" and consider it a "corporate" responsibility – someone else's concern.

One of the core issues with IP security today is that it is an "add-on" and was not built-in from the start. While this is not unique to IP, many IT programs develop their core functionality first and then try to figure out how to meet security requirements. This creates a larger problem as IP is becoming the ubiquitous form of communication that touches the entire enterprise. Many work-a-rounds have been developed to support not only IP security, but also other limitations of IP. The end result? It is almost impossible to achieve the initial intention of the Internet—end-to-end connectivity—and is a challenge to build and rollout new applications.

The good news is that IPv6 can fit into the very same security model that agencies use today, and initial implementations of IPv6 will most likely mirror the agency's IPv4 security architecture and policies. The real challenges will begin when agencies decide to start utilizing advanced features of IPv6. To gain the maximum value, they will have to move away from the enclave-based approaches and begin to implement end-to-end services that will force them to rethink the current enterprise security paradigm.

2.2 Understanding the Holistic Security Model

It is not generally understood that security can be boiled down into a business case issue. Many people view it as a technology problem, or a physical problem, or something that is always solved with some tangible action. It is common sense to protect your agency's valuable assets, resources, and information. But is it always the right course, and how do you determine how much security is enough? This is where a clear understanding of the fundamentals is necessary. Understanding the threats to your enterprise and your vulnerabilities is essential to determining your risk profile, but you must also factor in the value of what you are protecting. Only then can you understand which, if any, of your potential mitigation approaches are acceptable. It would not be wise to spend more money protecting an asset than the asset is worth. Figure 4 shows the business case of security.

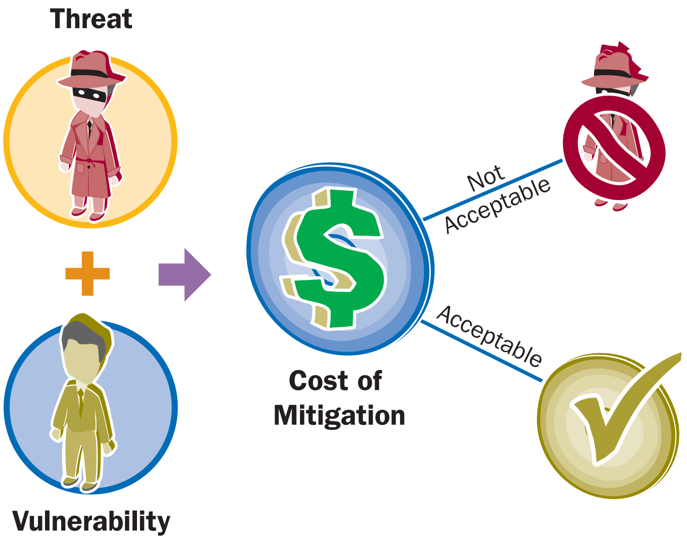

Figure 4: The Business Case of Security

Just as all forms of communications are in the process of converging to a common underlying IP infrastructure, a similar event is occurring with enterprise security. In the past, various aspects of security were segregated across the enterprise with little or no linkage to each other. Thus, while an agency may have excelled at one area of security, other implementations may have left it limited in scope and sorely vulnerable. This is one reason why “social hacking” became effective. In many cases it was easier for a hacker to call someone within an organization to obtain sensitive information, such as user logins and passwords, than it was to fight the information security mechanisms in place.

The term “holistic security” can have a variety of interpretations and typically refers to the combination of physical and information security. In reality, holistic security should concentrate not only on the merging of the physical and information security-oriented silos, but also should take a vertical cut across the entire enterprise as every department, person, and resource has some role in the operational security of the enterprise. The overall process for security should be tied very closely with the mission, goals, and strategy of the organizations and should be clearly spelled out in policy and guidance from agency executives. Figure 5 shows the holistic security model.

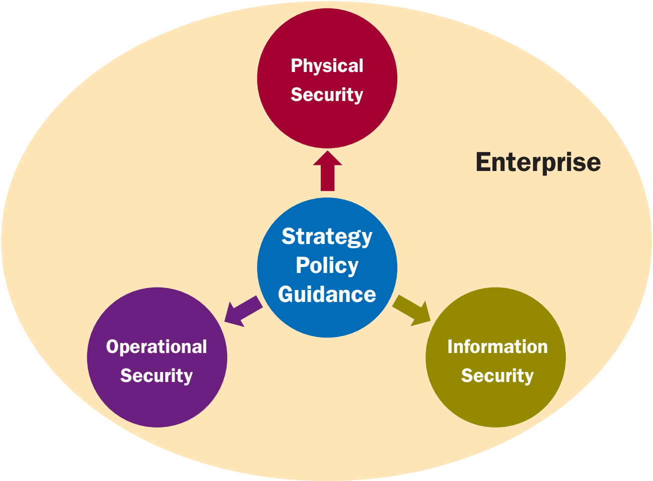

Figure 5: Holistic Enterprise Security Model

While security for the enterprise should be looked at in a holistic manner, so should the components for information security. During the development of technology, it is normal to put blinders on and take an all or nothing approach to security development. While most groups decide security should be handled as a separate issue and do not even consider it, others take the opposite approach and decide every potential security situation must be solved. The Internet standards are no exception to this approach. Using the Open Systems Interconnection (OSI) model as a reference, similar security services can be applied at multiple layers of the communications stack, as shown in figure 6.

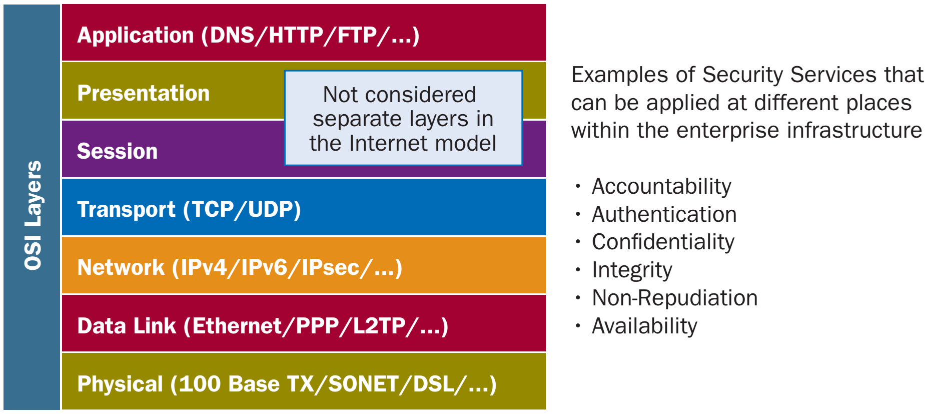
OSI Layers
Application (DNS/HTTP/FTP/...)
Presentation
Session
Not considered separate layers in the Internet model
Transport (TCP/UDP)
Network (IPv4/IPv6/IPsec/...)
Data Link (Ethernet/PPP/L2TP/...)
Physical (100 Base TX/SONET/DSL/...)
Examples of Security Services that can be applied at different places within the enterprise infrastructure
· Accountability
· Authentication
· Confidentiality
· Integrity
· Non-Repudiation
· Availability

Figure 6. Applying Security Services Across Layers

Thus it becomes important when developing technology and deploying security across the enterprise to understand:

- What is being accomplished?
- How is it being accomplished?
- Where is it being accomplished?
- Why is it being accomplished there?

Not only do you want to identify and fix potential security leaks, you also want to avoid layering too many security mechanisms on top of each other. An example of this is with encryption. It is easy to encrypt many places within the OSI model. Encryption is common at the physical, network, and application layer, in addition to any encryption that may occur on the data prior to being communicated across the Internet. While there may be times in which multiple layers of encryption are useful, they can cause significant network overhead and resource utilization issues. As IPv6 is being deployed in the enterprise, agencies need to take a hard look at their overall enterprise security model to understand the right approach for IPv6 security in the near-term and in the future.

2.3 Changing the Security Paradigm

The Internet and other IP-based systems today play a significant role in the way we communicate and do business, but they will play and even bigger role in the future. Thus, it is necessary that as next-generation IP systems are deployed, they are designed to not only provide the necessary security services to support these solutions, but also to provide the levels of assurance that users will trust and rely upon.

As it was earlier discussed, the IP security model typically deployed today focuses on secure enclaves in which the general rules are: users inside the enclaves are good and ones on the outside are bad. As IP-based solutions have become more complex, so have the systems designed to protect the enterprise and the sheer number of security related devices. While each device has a critical purpose, taken as a whole, the added complexity and cost weigh very heavy on the enterprise and slow the ability to respond by developing and deploying new services. Many of the security devices also break the idea of true end-to-end connectivity, either on purpose or as a side-effect.

Many of the new and enhanced capabilities associated with IPv6 focus on moving IP back to an end-to-end connectivity model. This will not only allow greater flexibility and capabilities between communicating nodes, it will also significantly reduce the complexity necessary to support the large number of “boot-strapped” solutions in use today.

The movement back towards end-to-end services will not only impact the way services are supported on the enterprise and across the Internet, it will also radically change the way information security is viewed and implemented within the enterprise. Many agencies utilize firewalls, NAT and other devices to create an “insulation” between their internal assets and the rest of the world. Typical security implementations have hard and crunchy outsides and soft and chewy middles. This is great if you can control everything coming in or leaving the enterprise and have complete trust that nothing malicious will happen inside the enterprise. but this is an impossible scenario to achieve. The greatest weakness of enterprise security has always been associated with the limited security and control that can be effectively implemented within the enterprise. There has been a movement towards “security in depth,” but even so, there is a very high reliance on enclave devices and still not enough emphasis on spreading security capabilities throughout the enterprise.

The transition to IPv6 provides a perfect time for agencies to begin re-architecting their enterprise security solutions to support end-to-end and other enhanced capabilities. While this will not happen overnight, the vision and plans must be developed to achieve maximum value during the initial planning stages of the transition. The two primary aspects that will provide the greatest value will be

- Creating a security plane that will provide an enterprise-wide ubiquitous security infrastructure
- Infusing enterprise nodes with more of the responsibilities typically implemented by enclave security devices

The development of a ubiquitous enterprise security plane leveraging IPv6 will provide agencies with a more robust method of creating security assets and rolling out or incorporating new enterprise-wide services. The lack of a common security plan across the industry has created an environment in which dozens of security solutions are developed, many times based on specific applications. This creates significant overlap in solutions, increases overhead and management requirements, and potentially leaves fatal gaps in enterprise defenses. A common security plan will allow agencies to specify common security solutions across the enterprise and allow applications to leverage and interact with those services through common application programming interfaces (APIs). This approach not only lowers overall cost and management, but also will become critical for effectively utilizing spiral or extreme application development and testing techniques.

As end-to-end services become the norm across the Internet and within the enterprise, the typical methods of protecting the enterprise today will need to evolve in order to maintain and provide greater levels of protection. End nodes will assume a greater degree of security services themselves as opposed to relying on boundary devices. Nodes will have to provide greater firewall, virus, and intrusion detection capabilities. Many of these security capabilities are available today for commonly used operating systems, but nodes will need much greater capabilities, and it will be necessary for them to tie into overall enterprise security management and monitoring systems. While nodes will assume a greater role, boundary devices will not go away, but will continue to play a critical role. All of the devices across the enterprise will need to share information and act as a single prevention system. Thus, while boundary devices will not play the same role and break the end-to-end model, they will still be gatekeepers screening for policy breaches and be the front line to shut down unauthorized streams of communication.

One of the major changes agencies and organizations globally will face is the eventual deprecation of NAT devices. NATs have played a critical role in extending the life of the IPv4 address space and will continue to be critical to support IPv4 portions of the enterprise for a number of years to come. While NATs can be implemented in IPv6 environments, they will most likely go by the wayside as end-to-end services are deployed and they are no longer needed to support the lack of IP addresses. There is also a perception that NATs are a security device, and while they do provide some limited protection, there is more misconception on their overall utility.

2.4 Benefits Associated with Security in IPv6 Networks

IPv6 will provide agencies a foundation to create a number of next-generation services that can deliver major benefits. But agencies will be hard pressed to find short-term operational benefits in their IPv6 deployments. In fact, many issues, especially security will require agencies to spend incremental resources during their initial transition phases.

If agencies view the transition to IPv6 as an opportunity and invest in proper planning, they can realize a number of security benefits associated with IPv6, including:

- **Secure Architectures:** Transitioning to IPv6 provides agencies a chance to significantly modify and enhance their enterprise architecture around the capabilities of IPv6. It affords the opportunity to implement new security architectures that could significantly improve an agency's overall security posture.
- **Ubiquitous Security Layer:** Numerous security protocols have been and are being developed within the IETF to support greater security capabilities within IP, such as IP security (IPsec). While many of these will operate with both IPv4 and IPv6, current entrenched deployments of IPv4 make spending the resources necessary to modify the equipment and architectures to implement them unlikely. In addition, IPSec is considered a mandatory part of IPv6.
- **Node and Topology Hiding:** One of the weaknesses in IPv4 is the ability for malicious entities to quickly scan and identify nodes on the Internet. Once a hacker has access to an organization's subnet, it is a fairly quick and simple process to identify all of the nodes and focus on the ones with the greatest weakness. IPv6 provide a significant advantage due to the sheer number of potential addresses on a single subnet. There are 264 or 18,446,744,073,709,551,616 potential IPv6 nodes on each subnet, making typical network scanning virtually impossible.
- **New Capabilities:** The IPv6 foundation enables the development and deployment of new capabilities and delivers inherent security benefits from utilizing an established and approved framework. Currently, as new services are required for the Internet, inventive companies are identifying issues and design workarounds. While this is great from a service delivery standpoint, many of these workarounds exploit or create new security vulnerability. However, developers and users are left with little recourse to implement their requirement. Thus, IPv6 will provide an environment that can be focused on security and also provide the flexibility for quickly delivering new services.

- **Unique Identification:** Significant security issues on the Internet stem today from the use of NAT and private IPv4 address space. It is virtually impossible to obtain a level of assurance based on IP addresses. Most users sit behind one or more NAT or similar devices that prevent the direct association of an IP address to a specific user or node. With the changes in IPv6 structure and tremendously increased address space, architectures and services can be developed to prevent address spoofing and establish the necessary association to support true network-level access control and authentication.
- **Communities of Interest (COI):** IPv6 makes it easy for nodes to have multiple IPv6 addresses on the same network interface. This can create the opportunity for agencies to establish overlay or COI networks on top of other physical IPv6 networks. Thus, department, groups, or other users and resources can belong to one or more COIs with each having its own established security policies. That way, security can become more granular and easier to implement based on grouping common requirements.
- **High Availability:** One of the major strengths of IPv6 will be the ability to quickly setup and modify networks on the fly. This ad-hoc capability will allow not only nodes on the network, but also entire networks to become much more resistant to denial-of-service scenarios. When deployed in mesh configurations, nodes and potentially networks could quickly identify and establish new routes as existing or preferred routes are disrupted.

3 Planning for Security During the IPv6 Transition

Although understanding the future security architectures and capabilities are critical to the ultimate success for agencies in gaining the maximum benefits from IPv6, the short term security and planning issues are the driving topics for agencies during the transition. Over the next few years agencies will face their greatest security challenges associated with the IPv6 transition. The greatest vulnerabilities agencies face are lack of knowledge and planning. It is critical to understand that agencies have begun the transition process, whether through a formal program or by default. Most operating systems, routers, other networking devices, and some applications are already IPv6 enabled and IPv6 may be applied by default. Thus, by not knowing or planning for IPv6 in an agency's security programs, they are open for innovative attacks utilizing tunneling or other means.

3.1 Phased IPv6 Transition Approach

The security challenges during the transition will change as the method and use of IPv6 changes over time. The initial deployment of IPv6 is expected to operate very similarly to IPv4 in the beginning phases of the transition. However, agencies will most likely move away from pure enclave-based architectures to support the growing requirement for end-to-end services that will be necessary to implement many of the advanced IPv6 capabilities. This will require new thinking around security and a stronger push towards node-based security architectures. The result will be a phased IPv6 security rollout.

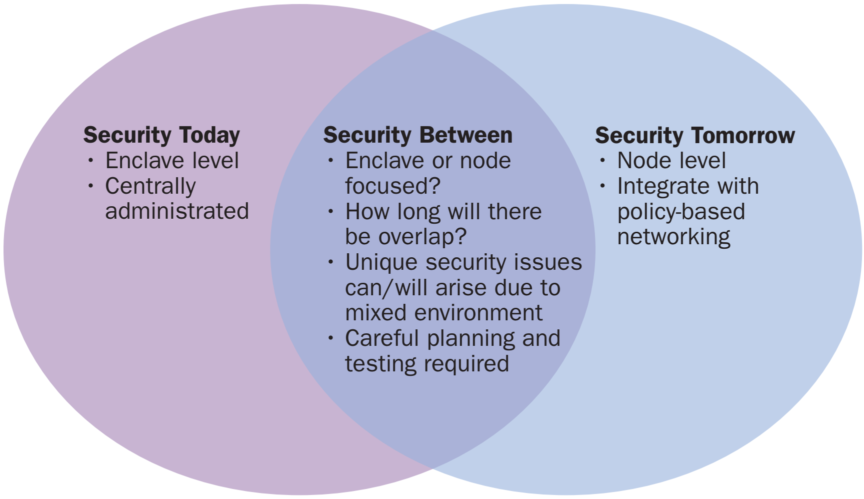
Security Today
• Enclave level
• Centrally administrated
Security Between
• Enclave or node focused?
• How long will there be overlap?
• Unique security issues can/will arise due to mixed environment
• Careful planning and testing required
Security Tomorrow
• Node level
• Integrate with policy-based networking

Figure 7: Phased Security Transition Approach

In addition, the guidance from Office of Management and Budget (OMB) and the CIO council is to move towards a dual-stack (IPv4/IPv6) approach throughout the transition cycle. While this provides for greater interoperability, the dual-stack environment could provide unique risks. Figure 8 shows how IPv6 will most likely be phased in over time, and that throughout the cycle the use of IPv4 will gradually be removed from an agency's environment. Although the phased approach provides greater security, the use of both protocols could provide an extended period where one protocol may be utilized as a point of attack for another.

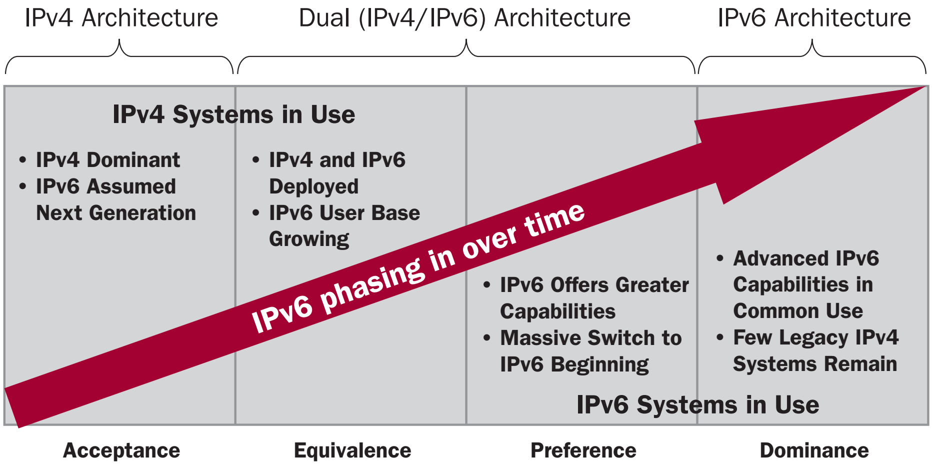
IPv4 Architecture
Dual (IPv4/IPv6) Architecture
IPv6 Architecture
IPv4 Systems in Use
• IPv4 Dominant
• IPv6 Assumed Next Generation
• IPv4 and IPv6 Deployed
• IPv6 User Base Growing
IPv6 phasing in over time
• IPv6 Offers Greater Capabilities
• Massive Switch to IPv6 Beginning
• Advanced IPv6 Capabilities in Common Use
• Few Legacy IPv4 Systems Remain
IPv6 Systems in Use
Acceptance
Equivalence
Preference
Dominance

Figure 8: IPv6 Phased Approach

During the transition process a number of critical implementation issues must be considered from a security perspective, including:

- Governance and Policy Needed
- Training
- Compliance Testing
  - Security Certification
- Institutionalized IPv6, and
  - Make Security Features Available
- New Attack Surface.
  - New technology and processes needed
  - Eliminate IPv4 attack surface ASAP

Some of the first steps any agency should take with regard to security include:

- IPv6 Security Plan
- Policy
- Routers/Switches
  - Disable IPv6/Tunnels
  - ACL to Block IPv6/Tunnels on core/edge/outside enclave
- Network Protection Devices/Tools
  - Contact vendors for IPv6 advice
- Block IPv6 (Type 41) Tunnels
- Enable IPv6 IDS/IPS features
- End Nodes, and
  - Enable IPv6 host firewalls on all end devices
  - Disable IPv6 if not used
- Monitor Core and Enclave Boundaries

3.2 Federal Security Policy and Guidance

While there are not a tremendous number of authoritative sources that provide IPv6 security policy and guidance, some do exist. In addition, there are a number of sources from which agencies can pull general security policy and guidance, and apply it specifically to their IPv6 security planning activities.

The following sources of policy and guidance to agency IPv6 security planning efforts are not exhaustive, but will provide an understanding and basis for the information agencies can refer to for their needs. In addition to the sources below, agencies should look at:

- Internal policies and guidance
- Federal Information Security Management Act (FISMA) policies and directives, Homeland Security Policy Directives (HSPD), and other government-wide and Department of Defense (DoD) specific policy directives
- Other authoritative sources upon which they rely
- Industry best practices

The following federal organizations have released policy and/or guidance of which agencies should be aware as they develop their IPv6 security plans:

- Office of Management and Budget (OMB)
- National Institute of Standard and technology (NIST)
- Chief Information Officer (CIO) Council IPv6 Working Group
- General Services Administration (GSA)
- Off of the Secretary of Defense (OSD)
- National Security Agency (NSA)
- Defense Information Systems Agency (DISA)

3.3 Security Before the IPv6 Transition

It does not matter if an agency is preparing for the transition to IPv6 or not, the transition is occurring through the normal technology acquisition cycles. Almost every router deployed today has the ability to communicate using IPv6 and the majority of operating systems have built IPv6 into them. Many of these implementations have transition mechanisms built into them to operate over IPv4 based networks. Thus, it is possible to pull a new piece of equipment out of the box and set it up and have it operating in IPv6 and not even be aware of it. In fact, it is also possible for that new equipment to establish its own “tunneled” connections to the Internet to begin passing IPv6 traffic, and if an agency’s operational group does not know what to look for, you could have a significant security risk.

The real question agencies must face, then, is not if they plan to transition, but when and what control they can maintain during the process. The most important step agencies need to take regarding IPv6 (and any technology) is becoming aware of it and educating their staff. A better understanding of IPv6 will allow agencies to focus on developing and implementing IPv6 related security policies no matter the stage of the transition.

While most agencies are currently in the IPv6 transition planning stage, most are at various levels of implementation. During this process it would be a mistake for agencies to assume nothing needs to be done regarding IPv6 until they reach a specific point in their IPv6 implementation plan. In fact, the first step agencies should take is protecting their network from unauthorized IPv6 use that could threaten enterprise security. In 2003, when the DoD released its first policies on IPv6, it clearly established a policy that prevented IPv6 from being used on any “operational” network until the DoD’s transition had progressed sufficiently to mitigate the risks.

Implementing IPv6 on a single workstation and establishing a tunnel to the Internet to pass IPv6 traffic is very easy. IPv6 is an integral part of Windows Vista and simple to turn on in Windows XP. With one command, a user can turn on IPv6 in XP and then go to one of several Websites to establish a tunneling account. In all, a less than technical user can have IPv6 operational and an IPv6 tunnel established in less than 10 minutes.

If an agency has not begun the transition process or is in the initial planning stages and has not developed an IPv6 security plan, it should at a minimum consider the following steps:

- **Training:** At a minimum, agencies should provide at least initial amounts of training for key operational personnel and policy makers. The individuals should have enough information to be able to formulate IPv6 policy and guidance for the agency, and to implement enough security safeguards to enforce the policies.
- **Policy:** Each agency should have a specific policy established regarding IPv6. Depending on the need of the agency it may be prudent to limit any IPv6 use in operational networks until certain agency IPv6 transition milestones are met. Typically, focusing IPv6 usage to the lab or pilots that require special approval affords enough latitude for staff to experiment and become better versed in IPv6. This policy should be coupled with the overall agency IPv6 transition efforts.
- **Guidance:** Enough guidance should be provided so that staff can effectively implement policy. This should include identifying in greater detail what is allowed and what is not. Not only should guidance documents identify what not to do regarding IPv6, but how to start the planning for programs and systems.
- **Vendors:** Work with software, hardware and service suppliers to understand the level of IPv6 functionality provided within their applications, equipment, and services. It is critical to not only understand what comes with IPv6 on by default, but also what features are implemented to support IPv6 functionality. Particularly, agencies need to understand when automated tunneling functions are available that provide ease of IPv6 connectivity over existing IPv4 infrastructure—as well as how to disable these functions if policy does not allow for them.
- **Boundary Security:** Utilize existing or even new security equipment to prevent unauthorized or rogue IPv6 traffic. Packet filtering router configurations should be updated to implement IPv6 security policy as well as firewalls and other security devices. Pat particular attention to tunneling traffic, type 41, as well as others.
- **Traffic Monitoring:** Another method agencies should employ is the use of traffic monitors or sniffers to better understand what packets are flowing across the network. Some security equipment does not have a robust IPv6 feature set, thus monitoring traffic at the packet or frame level will allow agencies to better identify security policy transgressors.

3.4 Developing an IPv6 Security Plan

As agencies transitions to IPv6, security will be a major concern across not only its core networks, but also across the entire enterprise. While IPv6 provides the foundation for the development and implementation of a more secure enterprise, agencies must be concerned with potential issues the new protocol may create. Examples of these issues may include:

- Poorly implemented IPv6 stacks
- Few network protection devices/tools support IPv6
- New attacks
- Poorly implemented IPv6 routing protocols
- Inconsistent IPv4/IPv6 security features
- Few IPv4 network management tools ported to IPv6
- Organizations not leveraging new security features
- New/existing applications unable to leverage new IPv6 features to solve mission issues

One of the critical steps agencies must take in the IPv6 transition process is the development of an IPv6 Security Plan. The IPv6 security plan should not be developed in isolation, but should be a close, collaborative effort between the agency's IPv6 transition team and the agency's Information Security (IS) team, and should comply with agency-established IS policy and procedures. The objective of the IPv6 Security Plan is to facilitate the insertion of IPv6 while maintaining the security posture of the agency's IT infrastructure as required by Federal Information Security Management Act (FISMA). The development of the IPv6 Security Plan should include a core understanding of all of the components necessary to secure the agency's enterprise during the transition, including those shown in figure 9.

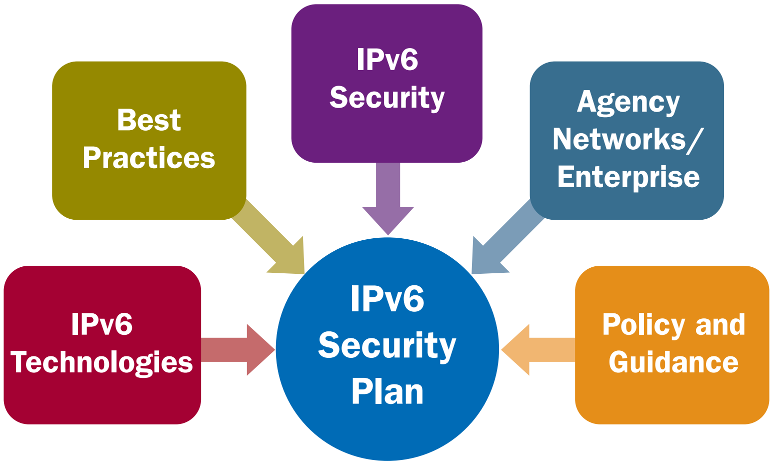

Figure 9: IPv6 Security Plan Inputs

The primary target of the IPv6 Security Plan in most cases will be to support the security necessary to achieve the June 2008 OMB mandated milestone. However, the plan should go further to provide a basis for the ongoing IPv6 transition for the agency's entire enterprise. The IPv6 Security Plan should be considered a living document and continually updated as necessary. At a minimum, the plan should:

- Provide a single guidance document that provides the processes and approvals for deploying IPv6 networks that meet the security requirements and supports the Agency transition strategy
- Identify IS roles and responsibilities for IPv6 transition
- Create a security policy and enforcement architecture which will identify the various IPv6-related security policies along with enforcement mechanisms and procedures and with specific dates required to enactment
- Identify technical and coordination requirements
- Incorporate and reference all Agency applicable IS policy directives and documents
- Identify the IPv6 IS procedures and methods used for IS testing, analysis and documentation

The IPv6 Security plan should provide an approach to all the security-based steps and tasks necessary for an agency to implement IPv6 in its core network, as directed by OMB, in a way that will maximize security to the agency and minimize any potential risks. In addition, the IPv6 Security Plan should provide a basis to support the security aspects of the agency's IPv6 transition beyond the June 2008 milestone date. The IPv6 Security Plan will be a living document and should be continued to be updated throughout the IPv6 transition cycle.

Due to the compressed IPv6 transition cycle and the nearing June 2008 milestone date, many agencies will be pressured for sufficient time to develop a comprehensive IPv6 Security Plan. However, it is critical that agencies carefully consider their approach and allocate appropriate resources to mitigate potential risk associated with the IPv6 transition.

While many agencies will require expertise outside their agency for support, it is critical that the IPv6 Security Plan be developed in a collaborative environment and be heavily integrated into not only the agency's IPv6 Transition Plan, but to the agency's current security planning approach. The goal of the IPv6 Security Plan should not be to develop a completely new approach to the topic of security within the agency, but to closely resemble the existing policies and systems in place today and where possible, not employ new processes or technologies. Limiting the impact to the organization will not only reduce costs and risks to the IPv6 implementation schedule, it will also minimize the potential for security mishaps that come from a conglomeration of disjoint processes and security practices.

The IPv6 Security Plan should be viewed as a project and implemented utilizing a process that takes a methodical approach. A critical aspect in developing the IPv6 Security Plan will be the gathering and analyses of relevant data, including:

- Deduce OMB Policies to be Meaningful
  - Directives
  - Memoranda
- Evaluate Guidance
  - CIO Council
  - NIST
  - DOD and NSA
  - Industry Best Practice
- Review Existing Agency Security Requirements and Solutions
- Review Existing and Emerging RFCs
- Review Existing and Emerging Industry Solutions
- Focus First on Security Policy
- Understand Security Architecture
- Develop Potential Solution Sets For:
  - Transition
  - IPv6 Capability
  - Beyond.

3.5 Initial IPv6 Deployment Architectures

Many agencies are making significant progress towards achieving the OMB June 2008 milestone, while some are already looking beyond being able to just pass IPv6 traffic across their core networks to applications and mission improvements. No matter the case, it is clearly understood within the government and the commercial community that the transition to IPv6 is a necessity and needs to occur sooner rather than later. Thus, deploying initial capabilities, such as those discussed in the OMB Memo M-05-22 creates not only the first step in a long roadmap of transition, but also allows agencies to become familiar with IPv6 in an operational environment.

Agencies will face significant challenges as they begin inserting IPv6 into their enterprise architectures, and it will be critical for them to take approaches that will provide minimal disruption to their operational environment. This approach will vary widely by agency. Some agencies have significant and highly technical operational talent and will be able to integrate IPv6 operational procedures into their daily routing as a matter of course; other agencies may have greater challenges as their resources are more limited.

Another factor agencies must face is the varying degree of support from their products and service providers. While the majority of router and operating systems support IPv6 today, security and network management products have not incorporated robust IPv6 capabilities as quickly. This could lead to many challenges as agencies may look to replace their current systems altogether or supplement them with IPv6 specific security devices.

Agencies should consider a number of IPv6 deployment approaches or philosophies as they design their initial IPv6 implementations, including:

- **Keep it Simple:** Designing complex solutions for initial IPv6 deployments within agencies will only add to operational issues and potential security problems. Initial IPv6 deployment designs should focus on achieving success goals, providing a baseline approach from which agencies can grow future enterprise solutions based on IPv6, and providing the means from which their staff can gain first hand operational experience.

- **Segregate IP Connections:** Enclave approaches dictate a limited number of ingress and egress points for the enterprise architecture that need to be closely guarded and monitored. Thus, in many cases, these have become significant traffic points for interconnecting to the outside world. Agencies should consider installing parallel sets of IPv6 connections to the outside as opposed to running both IPv4 and IPv6 on their primary peering or Internet service provider (ISP) connections. This approach has no operational impact on established connections and allows agencies to be much more vigilant in analyzing the connections and traffic on IPv6 connections.
- **Divide and Conquer:** IP-based attacks have become increasingly sophisticated. This has driven the need for better and better security capabilities housed not only within specific security devices such, as firewalls and intrusion detection systems, but in all devices connected to the network. This not only increases the complexity of the security devices, but also requires that they support much more processor intensive applications. Compounding the issue to an even greater degree is the fact that many security products have been slower to adopt IPv6 capabilities than other networking devices.

Similar to the rationale for segregating IP connection, agencies should consider implementing security devices specifically to meet their IPv6 needs. This approach allows agencies to maintain their current IPv4 security posture and utilize specific security equipment solely focused on IPv6. While some modification will be necessary to the IPv4 equipment, the majority of the work and complexity can be relegated to the IPv6-specific devices. This approach provides minimal processing impact on existing infrastructure and can be used to significantly mitigate risks in deploying IPv6 within an agency's enterprise.

- **Information is Power:** One of the most effective security tools available is information. Not just raw data, but a thorough comprehension of information, preferably in real-time, of what is happening now. Agencies should focus on incorporating IPv6 data gathering and reporting solutions into their routines as early as possible. In fact, agencies may want to focus on providing an additional focus on this area during their IPv6 deployment and for a period of time following deployment. The use of network sniffers and analysis tools designed to collect frame and packet level information can be utilized to supplement other data collection techniques and provide an in-depth benchmark to determine how well the agencies' implementation of their security policy is occurring.
- **Slow (but Steady) Growth:** The aim in deploying initial IPv6 services within agencies in June 2008 primarily focuses on building the confidence and success agencies will need to transform their entire enterprise to support IPv6. Many agencies will not have a hard operational requirement in 2008 other than the ability to transport IPv6 packets. But agencies should develop a plan that promotes steady growth of IPv6 across their enterprise, but limits the deployments based on needs and the agency's ability to provide sufficient security to implement new capabilities. Initial deployments can utilize limited bandwidth and support a limited number of applications that can be closely monitored, while trials of other new capabilities such as VoIPv6 can be cordoned off to specific enclaves that will not create greater risks to the remainder of the enterprise.

3.6 Security in a Dual IPv4/IPv6 Environment

The deployment of IPv6 presents agencies with an interesting challenge, but one they have most likely have faced in the past. That is, operating separate logical networks on top of a single physical infrastructure. While IPv6 has many similarities to IPv4 in the way it operates, and these can be implemented simultaneously within the same physical infrastructure without interfering with each other, they are not directly interoperable. Thus, a host that speaks only IPv4 cannot directly communicate with a host that speaks only IPv6. The IETF foresaw this challenge and understood that the transition to IPv6 would require the operation of both protocols on the Internet and enterprise networks for a number of years. To handle this issue they developed three primary classes of transition mechanisms as shown in figure 10.

- **Dual-Stack:** A single network node which operates with both IPv4 and IPv6 on the same or different interfaces.
- **Tunneling:** The ability to transport IPv6 packets across IPv4 networks as data encapsulated within the IPv4 packet and vice versa.
- **Translation:** The ability to convert IPv6 packets to IPv4 and vice versa.

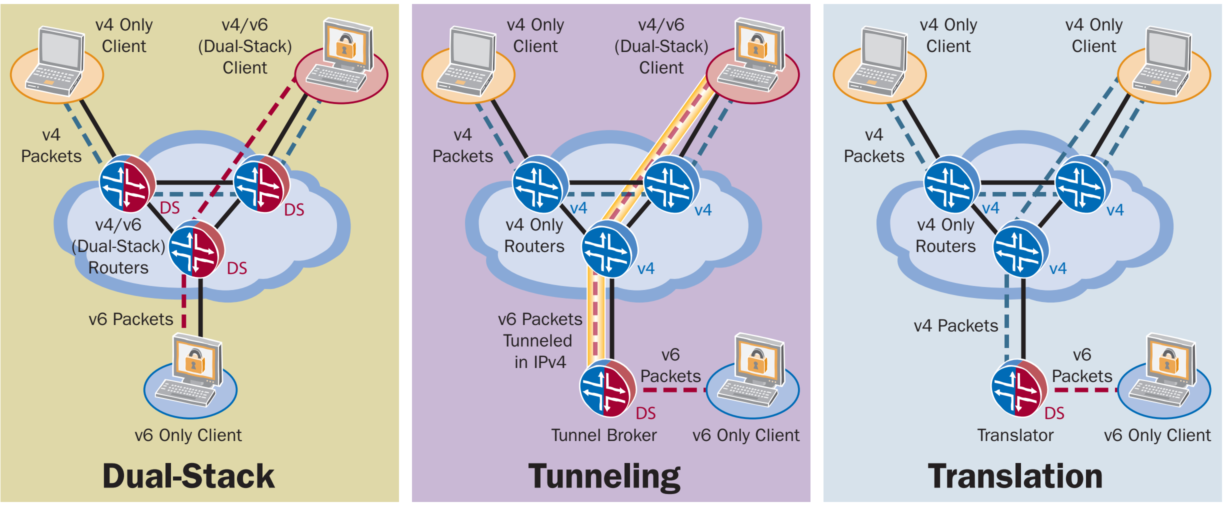

Figure 10: IPv6 Transition Mechanisms

While there are numerous methods agencies can use to transition, the bottom line is that they will be challenged with operating physical networks that must support both IPv4 and IPv6 at the same time What's more, there will be a growing prevalence of equipment attached to the network that will need to support both protocols simultaneously. This can potentially create unique challenges from a security perspective. Where IPv4 threats and vulnerabilities have been studied and understood for decades, IPv6 threats and vulnerabilities are not as well known. There has been significant research into this area and there will be a mapping in some cases of IPv4 threats and vulnerabilities to IPv6. But the more complex issue will be the potential for creating new threat and vulnerability categories for systems operating IPv4 and IPv6 at the same time. It is likely these issues will appear more from vendor implementations as opposed to protocol designs. In either case, agencies need to be aware and be vigilant for these types of risks.

Another area of concern will be for potentially overlapping security policies. During the initial IPv6 implementation and application growth, the agency should be concerned about maintaining a distinct security policy and implementation associated with IPv6. If differing security policy and implementation are employed, they should be closely coordinated and mapped between the IPv4 and IPv6 instances to ensure that one policy is not providing a greater degree of risk than what is acceptable to specific nodes.

3.7 Tunneling and Security

While implementing IPv6 through a pure dual-stack approach has certain advantages, most agencies will find that it will be cost- and resource-prohibitive to perform large-scale, enterprise-wide conversions over a short period of time. Thus, to provide end-user IPv6 access, they will need to employ transition mechanisms, such as tunneling. Tunneling has been used by numerous organizations across the world to support their transition efforts. If carefully implemented, this can provide cost-effective and secure solutions to enabling IPv6 across a wide user population while the enterprise infrastructure is steadily converted to IPv6.

Although tunneling does provide certain advantages, it can also cause significant problems to the security posture of an agency if not closely monitored and controlled. In many cases, agencies will have "rogue" or unauthorized tunnels in use today without even knowing it. Agencies should develop the use of tunneling into their overall transition approach, and craft agency-wide policy and guidance that will specifically identify when and where tunneling technology is allowed or denied. Agencies should then go a step further to utilize tools, such as network sniffers, to ensure their policies are being followed. In some cases, equipment may be IPv6-enabled and can utilize default tunneling mechanisms that will establish links to the outside world.

4 Security within IPv6 Networks

4.1 The IETF and IP Security

The IETF is the primary international body responsible for developing Internet related standards to make the Internet work better. The IETF produces technical and engineering documents to support the better design, use, and management of the Internet. Examples of documents produced by the IETF include protocol standards, best current practices, and informational documents. The IETF operates under the following cardinal principals:

- Open process
- Technical competence
- Volunteer core
- Rough consensus and running code
- Protocol ownership

The IETF is open to anyone, and is comprised of an international community of network designers, operators, vendors, and researchers whose primary concern is with the evolution and smooth operation of the Internet. The IETF is split into several topic areas, with each area containing multiple working groups. The actual work is done in the working groups during meetings or via mailing lists. Topic areas and working groups are created and retired as needed and the current topic areas within the IETF are shown in Figure 11.

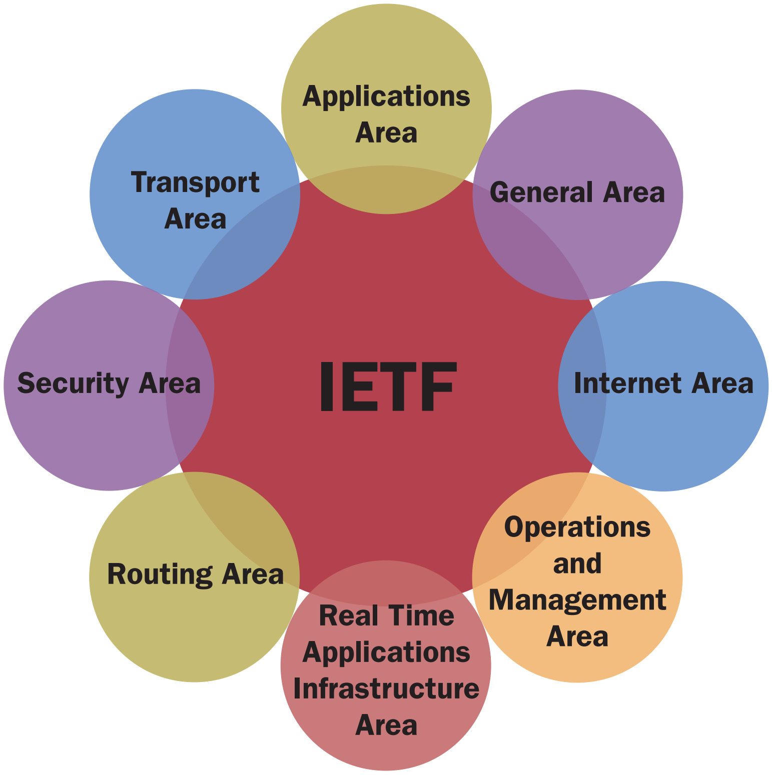

Figure 11: IETF Topic Areas

Security has been a significant focus within the IETF as can be seen in Table 1 and the sheer number of active and past working groups within the security area. The range of security working group topics has been very wide and has made significant contributions toward securing the Internet architecture.

Active	Concluded	
Better-Than-Nothing Security (btns)	Authorization and Access Control (Aac)	One Time Password Authentication (otp)
Domain Keys Identified Mail (dkim)	Authenticated Firewall Traversal (aft)	Privacy-Enhanced Electronic Mail (pem)
EAP Method Update (emu)	Common Authentication Technology (cat)	Profiling Use of PKI in IPSec (pki4ipsec)
Handover Keying (hokey)	Commercial Internet Protocol Security Option (cipso)	Securely Available Credentials (sacred)
Integrated Security Model for SNMP (isms)	Domain Name System Security (dnssec)	Secure Shell (secsh)
Provisioning of Symmetric Keys (keyprov)	Credential and Provisioning (enroll)	SNMP Security (snmpsec)
Kitten GSS-API Next Generation (kitten)	TCP Client Identity Protocol (ident)	Simple Public Key Infrastructure (spki)
Kerberos (krb-wg)	Intrusion Detection Exchange Format (idwg)	Internet Security Policy (spwg)
Long-Term Archive and Notary Services (ltans)	Extended Incident Handling (inch)	Secure Network Time Protocol (stime)
Multicast Security (msec)	IP Authentication (ipauth)	Trusted Network File Systems (tnfs)
Network Endpoint Assessment (nea)	IP Security Protocol (ipsec)	XML Digital Signatures (xmldsig)
An Open Specification for Pretty Good Privacy (openpgp)	IPSEC KEYing information resource record (ipseckey)	Web Transaction Security (wts)
Public-Key Infrastructure X.509 (pkix)	IP Security Policy (ipsp)	
Simple Authentication and Security Layer (sasl)	IP Security Remote Access (ipsra)	
S/MIME Mail Security (smime)	Internet Secure Payments Protocol (ispp)	
Security Issues in Network Event Logging (syslog)	Kerberized Internet Negotiation of Keys (kink)	
Transport Layer Security (tls)		

Table 1: Active and Concluded IETF Security Working Groups

The primary outputs of the working groups are Request for Comments (RFCs) which can be released as Internet standards, best practices or for informational reasons. Appendix A provides a list of RFCs that are relevant to IPv6 security.

The remainder of this section discusses in greater details various protocol and standards development activities completed or in progress by the IETF that impact the security of IPv6 networks. While each section may not be dedicated to IPv6, there is a significant impact in creating an enterprise-wide security implementation that will support the agency's IPv6 security requirements. It should also be noted that there are many more efforts within the IETF and through other organizations to create security solutions that agencies should consider as they transition to IPv6.

4.2 IP Security (IPsec)

This section provides a high-level overview of IPsec. More specific information can be found in the specific IPsec RFCs (listed in Appendix A) and from RFC 4301 Security Architecture for the Internet Protocol.

IPsec is one of the core security technologies for IPv4 and IPv6 and, when combined with other security mechanisms, can create a security infrastructure from which agencies can provide a common set of ubiquitous security services across the enterprise. IPsec is considered a mandatory part of IPv6, but optional for IPv4. IPsec utilizes cryptographically-based mechanisms to implement the following security services at the IP layer and for all protocols carried over IP:

- Access control
- Authentication
- Confidentiality (data traffic flow)
- Integrity

IPsec provides flexibility in the types of security services that are provided through the use of multiple protocols, including the Authentication Header (AH), the Encapsulating Security Payload (ESP) and cryptographic key management procedures and protocols. The IPsec architecture was developed to allow for the deployment of compliant implementations that provide not only the security services, but also the management interfaces needed to meet the security and operational requirements of the user community.

In addition to having multiple protocol options, IPsec can operate in either a transport mode or a tunnel mode.

- **Transport Mode**: When operating in transport mode, only the data portion of the packet (the payload) has the IPsec security services applied to it, except for the hashing function, which may be applied to the entire IP packet. Thus, the header and routing information is left unchanged or encrypted through the IPsec process. But if the authentication header is utilized IP addresses cannot be translated as this will cause an invalidation of the hash value and indicate a loss of packet integrity. The data and higher level protocols are always secured by the hash function and cannot be modified, including translating the port numbers. Transport mode is used for host-to-host communications to provide an end-to-end security service. In practice, IPv4 implementation makes limited use of this mode due to the miniscule deployment of IPsec within IPv4-only hosts.
- **Tunnel Mode**: When operating in tunnel mode, the entire IP packet has the IPsec security services applied to it. The IPsec tunneling mode concept operates very similarly (with the addition of certain security services) to other tunneling concepts. The entire IP pack is encapsulated and carried within the data portion of a new IP packet. Once the new IP packet reaches its destination, it is stripped away and the packet that was being carried as the payload can then be used by the end system or may continue to be routed through the network. Tunneling mode is primarily used in the development of VPN or VPN-like links through a network where end-to-end security is not implemented but a enclave-to-enclave security is desired. It is also used to implement host-to-enclave security for remote users. Tunneling is the IPsec mode predominantly used within IPv4 to create VPNs.

Figure 12 shows how IPsec Transport Mode and Tunnel Mode looks when applied using either IPsec AH or ESP.

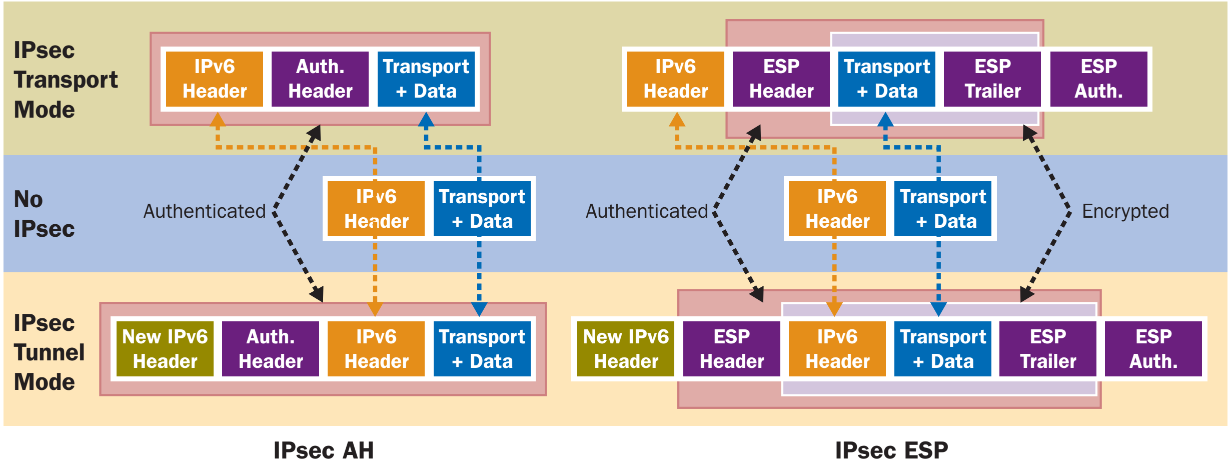

Figure 12: Transport and Tunnel Modes for IPsec AH/ESP

Implementations of IPsec on a host must support both transport and tunnel mode, while IPsec implementation on a security gateway must support only tunnel mode and may optionally support transport mode.

IPsec is designed to provide a flexible and scalable system for IP-based security and should be interoperable across multi-vendor IPsec implementations. In addition, IPsec should not impact interoperability with non-IPsec devices. IPsec is designed to be cryptographically independent, and while specific cryptographic techniques are required to create a certain level of interoperability, additional cryptographic techniques can be applied based on need.

While many of the processing requirements within IPsec are considered local and not subject to the needs for standardization, there are a number of external aspects that must be standardized for interoperability and management. Thus, there are three databases IPsec implementations must support:

- **The Security Policy Database (SPD):** The SPD specifies which security services will be offered to IP packets and how they will be implemented. The SA, in relation to the SPD, is a management construct used to enforce security policy for IP packets crossing the IPsec boundary. While the specific method of implementing the SPD can vary by implementation, a minimum management functionality must be provided to allow control over whether and how IPsec is applied to IP packets transmitted or received.
- **Security Association Database (SAD):** The SAD provides a location to maintain the parameters associated with each SA. Thus, each SA will have one line within the SAD to identify the relevant parameters associated with that specific SA.
- **Peer Authorization Database (PAD):** The PAD provides the necessary link between the SPD and a SA management protocol such as IKE. It includes the following critical functions:
  - Maintains the identity of peers (or groups of peers) authorized to communicate with the IPsec entity
  - Includes location information for peer gateways or for peers behind a security gateway
  - Supplies peer authentication data
  - Identifies the peer authentication protocol and method for each peer
  - Provides limitations on peers with regards to the establishment of child SAs and limits the types/values of IDs for child SA creation

IPsec utilizes the SPD, which can be created and maintained by an agency, to implement one or more security policies. It can also be updated and maintained by applications as agencies roll out new functionality that require direct interaction to implement more security services. When a packet is created or received it can have one of three actions applied to it based on the SPD:

- **Protected:** IPsec security services are applied based on policy
- **Discarded:** The packet is thrown away
- **Bypass:** The packet is passed without applying IPsec security services

In order to support the security of IP traffic crossing the network, IPsec utilizes the AH and ESP protocols, which are described in greater detail in their respective RFCs (note: see appendix A for specific RFC numbers). While it is mandatory that IPsec implementations support ESP, it is optional for them to support AH. Both protocols may be used individually or in combination. However, most security requirements can be met through the use of ESP alone. While ESP provides all of the security services for IPsec, the AH header provides security only for access control, authentication, and integrity.

Figure 13 shows the IPsec ESP header format.

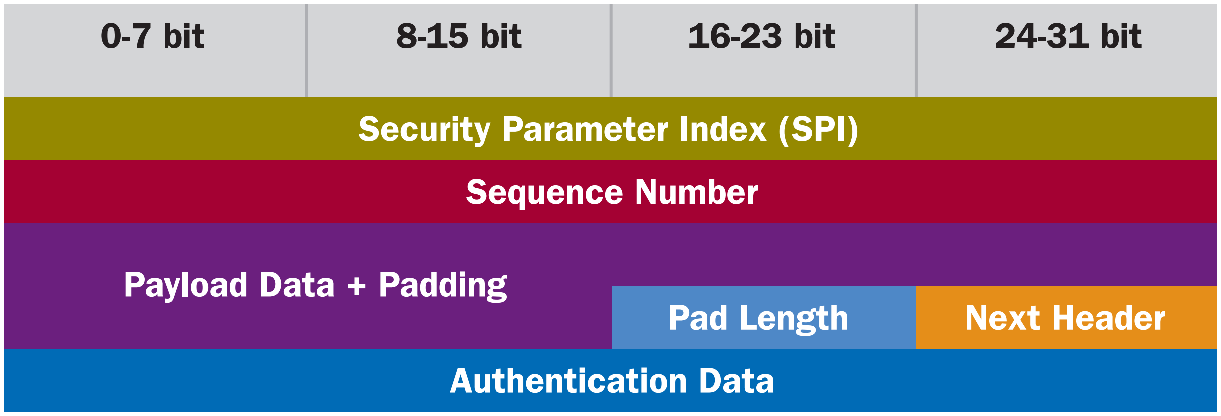

Figure 13: ESP Header

Figure 14 shows the IPsec AH format.

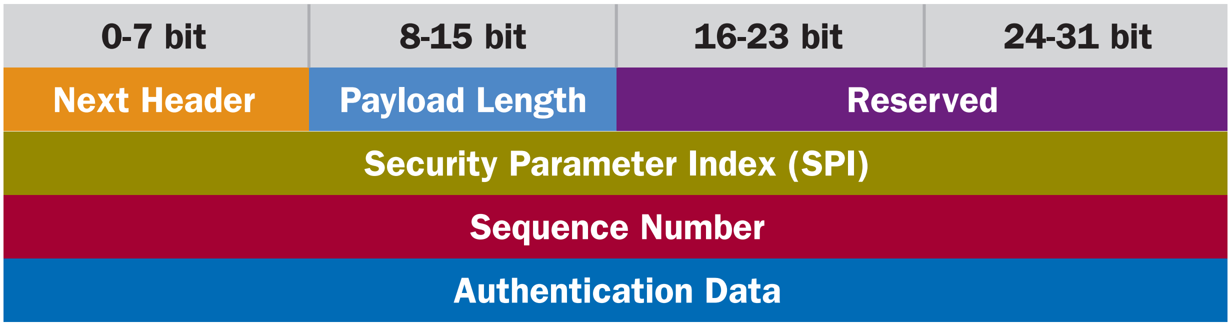

0-7 bit	8-15 bit	16-23 bit	24-31 bit
Next Header	Payload Length	Reserved	
Security Parameter Index (SPI)			
Sequence Number			
Authentication Data			

Figure 14: Authentication Header

The SPD provides the ability to adjust the granularity at which IPsec implementations provide security to IP traffic. Thus, a single tunnel can be used to carry all traffic between end-points or a separate tunnel can be established for each TCP connection.

IPsec relies heavily on encryption to perform its security services, thus, cryptographic keys and their distribution are critical. At a minimum, compliant IPsec implementations require the ability to perform both manual key management and the use of the Internet Key Exchange (IKE) protocol. Other key management methods may be used in addition to these.

IPsec can be implemented in hosts, routers, and other devices connected to the network. There are multiple methods in which IPsec can be implemented to support end-to-end connectivity or to create security gateways, including:

- **Native:** In a native implementation IPsec is built into the IP stack.
- **Bump-in-the-Stack (BITS):** In a BITS approach IPsec is not implemented within the IP stack, just as an add-on below it, between the native IP stack and lower layer protocols. This implementation is common with legacy IP stacks.
- **Inline/Bump-in-the-Wire (BITW):** In an inline approach, a dedicated, security processor is utilized and may have its own IP address.

Security Associations (SA) are fundamental concepts within IPsec and are utilized in the AH and the ESP protocols. A major function of the IKE protocol is the establishment and maintenance of the SAs. An SA is a unidirectional connection that establishes the security services on the packets carried by it. The security services for an SA are implemented by the use of the AH or ESP. A single SA can only be associated with one security connection. Thus, if AH and ESP are used simultaneously, then two separate SAs are required. Since the SA is unidirectional, in order to achieve bidirectional security, a separate set of SAs will need to be established for the return packets. However, IKE will create SA pairs as a common part of its operation.

The Security Parameters Index (SPI) is an arbitrary 32-bit value that is used by the receiving node to bind the SA with the incoming packet. IPsec may be applied to multi-cast traffic in addition to unicast traffic. When applying IPsec to unicast, the SPI is enough to specify the SA but implementers can choose to use the SPI in conjunction with the either the AH or ESP protocols for SA identification. If the use of IPsec is supporting multicast, then a specific algorithm must be utilized for mapping inbound IPsec packet to SAs.

IPsec provides mandatory support for both manual and automated SA and Key Management. Although many of the IPsec security services are supported using both manual and automated techniques, some do require an automated approach. An example of this is with anti-reply, which can be supported under both AH and ESP—but only if automated SAs are utilized. It should also be noted that the level of granularity that can be achieved with regards to authentication depends on the granularities of keys utilized. If manual key management techniques are used, and a group utilizes the same key material, then there will be much less granularity than if specific keys are tied to individuals, interfaces, traffic flows, or other characteristics.

Manual techniques are the simplest method for distributing SA and key management material, and will work well in smaller group or static environments. But they quickly become burdensome under conditions that require scaling.

Automated techniques require more overhead and forethought, and are generally a more costly method for distributing SA and key management material. But they can provide quicker scalability and the ability to build in more security features, such as greater authentication granularity and protection against traffic replay. As IPv6 becomes more widely utilized and widespread, the increased use of IPsec will require a standard solution that is automated and scalable. This will not only make the deployment and management of a large IPsec user community more feasible, but will accommodate on-demand creation of SAs for user/session oriented keying. It will create an infrastructure that can become very granular with regards to the level of security provided by IPsec. The automated key management protocol for IPsec is IKE. While other protocols may be used, IKE is a minimum implementation for IPsec interoperability.

4.3 Internet Key Exchange (IKE)

This section provides a high-level overview of IKE. For more detailed information please refer to RFC 4306. IKE provides an automated SA and key management capability. It accomplishes this by mutually authenticating two parties and then establishing an IKE SA. The IKE SA includes shared secret information which can be utilized to establish SAs for the IPsec ESP and/or AH protocols as well as cryptographic algorithms to be used by the SAs.

During an IKE session, the initiator proposes one or more cryptographic algorithms that can be combined into suites for use in a mix-and-match fashion during communication sessions. In addition to negotiating key material and SA, IKE can also negotiate use of IP Compression in conjunction with an ESP and/or AH SA. Where an IKE SA is referred to "IKE_SA", the SAs established for ESP and/or AH through the IKE_SA is called a CHILD_SA.

IKE communications utilize a pairs of messages: a request and a response. The flow always consists of a request followed by a response. The requestor has the responsibility for reliability and if a response is not received prior to the timeout, the requester either retransmits the request or abandons the connection. During the IKE session setup, the first request/response negotiates security parameters, sends nonces, and sends Diffie-Hellman values. During the second request/response, IKE transmits identities, proves knowledge of the secrets, and sets up an SA for the IPsec AH and/or ESP protocols.

IKE is used to negotiate IPsec ESP and/or AH SAs for a number of different scenarios, each with its own special requirements, including:

- **Security Gateway to Security Gateway Tunnel:** This scenario is a typical method of implementing IPv4-based VPNs as it creates an enclave-to-enclave IPsec tunnel. The endpoints do not need to implement IPsec as their traffic travels to the security gateway to become encapsulated within the tunnel. Once on the other end, the second security gateway removes the packet from the tunnel and continues to the destination in the clear.
- **Endpoint to Endpoint Transport:** This scenario focuses on providing end-to-end security between the two communicating nodes and provides a future basis for security architectures within IPv6. IPsec transport mode will commonly be used with this approach and a single pair of addresses will be negotiated for packets to be protected by this SA. If one or both of the nodes are behind a NAT, the tunneled packets will have to be UDP encapsulated so that port numbers in the UDP headers can be used to identify individual endpoints "behind" the NAT.
- **Endpoint to Security Gateway Tunnel:** This scenario is also very common in IPv4 and will likely continue in some form with IPv6. This is typically thought of as the remote user VPN scenario where and computer reverts back to its organizational enterprise to access a security gateway device. In this case the IP address may be static or dynamically allocated by the security gateway. IKE includes a mechanism for the initiator to request an IP address owned by the security gateway for use for the duration of its SA. If the protected endpoint is behind a NAT, the IP address as seen by the security gateway will not be the same as the IP address sent by the protected endpoint, and packets will have to be UDP encapsulated in order to be routed properly.

Figure 15 shows the IKEv2 header format.

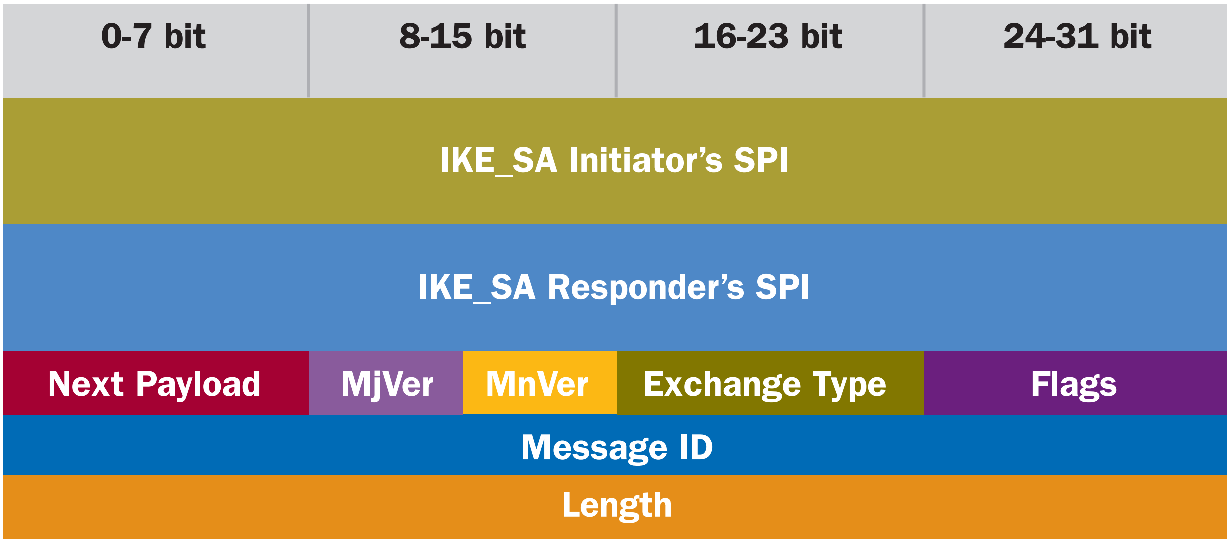

0-7 bit	8-15 bit	16-23 bit	24-31 bit
IKE_SA Initiator's SPI			
IKE_SA Responder's SPI			
Next Payload	MjVer MnVer	Exchange Type	Flags
Message ID			
Length			

Figure 15: IKEv2 Header

4.4 Secure Neighbor Discovery (SEND)

This section provides a high-level overview of SEND. A more detailed description can be found in RFC 3971. The Neighbor Discovery Protocol (NDP) is defined in RFCs 2461 and 2462 and is utilized by IPv6 for nodes to find other nodes on the same link. This allows nodes to understand who their neighbors are, their link-layer addresses, and to find routers and maintain reachability information. NDP provides IPv6 with the functionality that was achieved in IPv4 with ARP, ICMP, and ICMP Redirect. Since IPv6 does not have a broadcast capability like IPv4, it uses multicast for data link layer address resolution. NDP is utilized by hosts and routers and, in addition to Neighbor Discovery (ND), provides the following functionality:

- **Router Discovery (RD):** Provides a method for IPv6 hosts to discover local routers on local-link. This is primarily used for the purpose of establishing connections with routers that will forward packet on the host behalf. It also provides subnet prefix discovery to determine which destinations are directly on a link.
- **Address Autoconfiguration:** Automatically assigns addresses to a host, which allows hosts to operate without explicit IP configurations related to connectivity. The default autoconfiguration mechanism is stateless, and hosts use prefix information provided during RD to establish their IPv6 address and then test for uniqueness.

Dynamic Host Control Protocol version 6 (DHCPv6) can provide additional autoconfiguration features. DHCPv6 is defined in RFC 3315 and assigns IPv6 addresses and provides other configuration information to nodes when they attach to the network. A reserved, link-scoped multicast address is primarily used for DHCP clients to communicate with DHCP servers. This helps alleviate the need for clients to have the IPv6 address of the DHCP server. However, the client can communicate with the DHCP server utilizing its IPv6 address, once it is known. DCHP relay agents can be used to allow clients to communicate with DHCP servers not on their link. The relay communication process is transparent to the DHCP client. DHCP may be used to provide additional configuration information to nodes that already have globally unique IPv6 addresses (from stateless autoconfiguration), or to assign IPv6 addresses and additional configuration information to clients that do not already have a globally unique IPv6 address. Nodes that do not need the DHCP server to assign it an address can obtain configuration information for available DNS or NTP servers through a single message and reply exchanged with a DHCP server.

- **Address Resolution:** Provides the ability for nodes on a link to resolve IPv6 addresses to the corresponding link-layer address on the same link. The source link layer address on the frames is not checked via the Address Resolution function allowing for easier addition of network elements.
- **Neighbor Unreachability Detection (NUD):** Provides for tracking/reachability of neighboring hosts and routers.
- **Duplicate Address Detection (DAD):** Prevents address collisions by verifying that no other node is using the same address. IPv6 requires addresses be found unique before higher layer traffic can be proceed.
- **Redirection:** Provides the ability to redirect the host to a better route or let it know that the destination is on the local-link and does not require routing.

NDP is not only vulnerable to a number of potential threats, but the risks associated with utilizing NDP and not security can be substantial. Originally the intention was to utilize IPsec to protect NDP traffic, but during the development of the standard, little details were provided to support actual implementations. In addition, practical experience has shown there are a number of issues with utilizing IPsec for NDP. Specifically, there is a bootstrapping problem with IKE, and NDP can only be used with manual configuration of IPsec. This becomes an issue as the number of security associations is large enough to make it impractical to use in a manual configuration, and it removed many of the advantages of NDP.

Given the critical nature of NDP and the issues faced by the implementation community, the IETF developed SEND to counter the threats to NDP. SEND is primarily developed for environments where the security of physical links cannot be assured and create concerns for NDP attacks. SEND provides a set of new ND options to secure the various functions in NDP and are used to protect NDP messages. SEND specifies these options and an authorization delegation discovery process, an address ownership proof mechanism, and requirements for the use of these components in NDP. Some of the solution components include:

- Router authority is established by certification paths, anchored on trusted parties. The host and router must be configured with a trust anchor before the host can set the router as its default.
- Cryptographically Generated Addresses (CGAs) are used to verify the ND sender is the owner of the address. The CGA option carries the public key and associated parameters. Non-CGAs with certificates that authorize their use can be utilized.

Figure 16 shows the SEND CGA option.

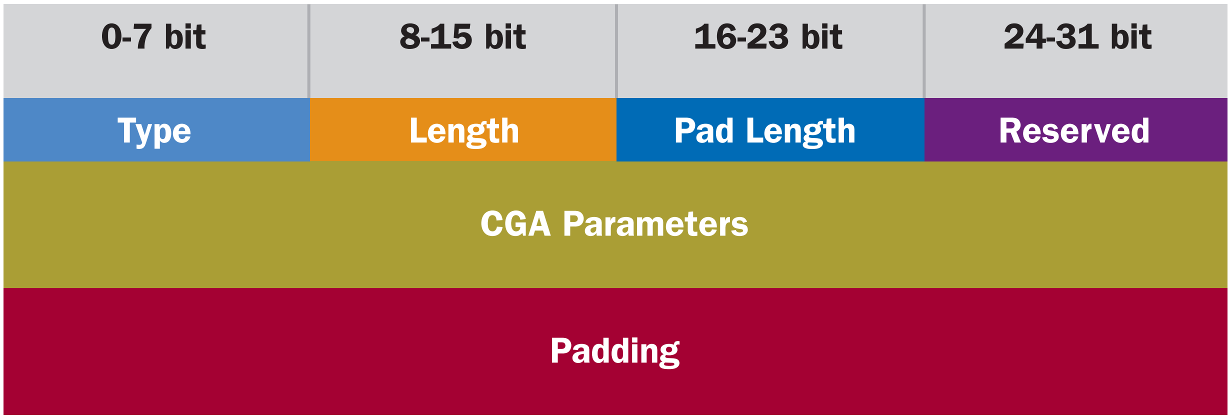

0-7 bit	8-15 bit	16-23 bit	24-31 bit
Type	Length	Pad Length	Reserved
CGA Parameters			
Padding			

Figure 16: SEND CGA Option

- An RSA Signature option is used to protect all messages relating to Neighbor and Router discovery to protect the message integrity and provide sender identity authentication. Public key authority is established through an authorization delegation process using certificates, CGAs, or both.

Figure 17 shows the SEND RSA option.

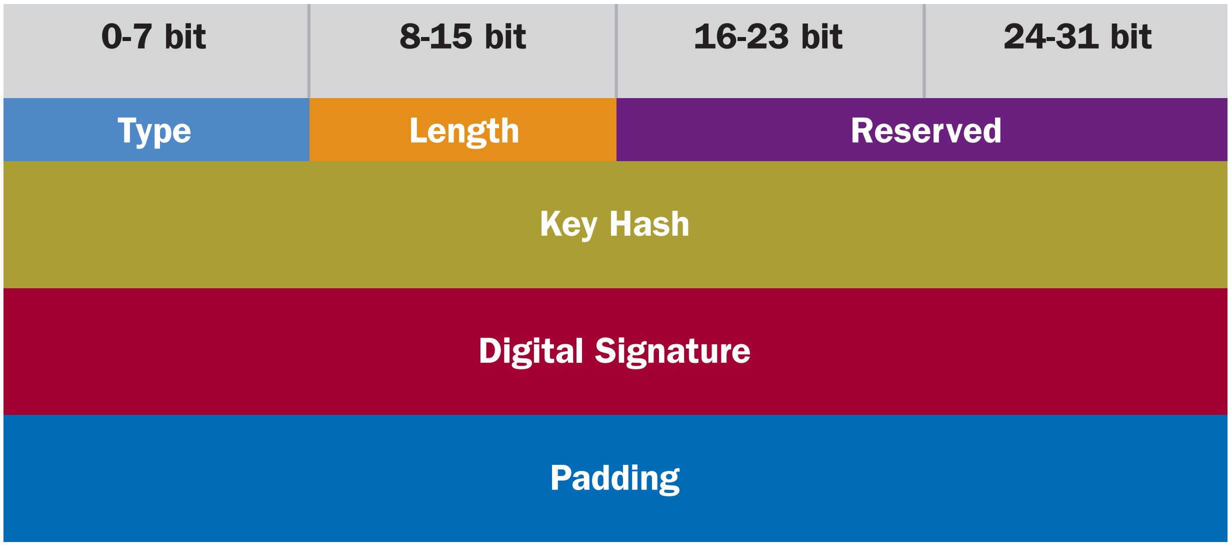

0-7 bit	8-15 bit	16-23 bit	24-31 bit
Type	Length	Reserved	
Key Hash			
Digital Signature			
Padding			

Figure 17: SEND RSA Signature Option

- Introduces Timestamp and Nonce to prevent replay attacks. To support multicast addresses, Timestamp options do not require previously established state or sequence numbers. Solicitation-advertisement pairs are protected with the Nonce option.

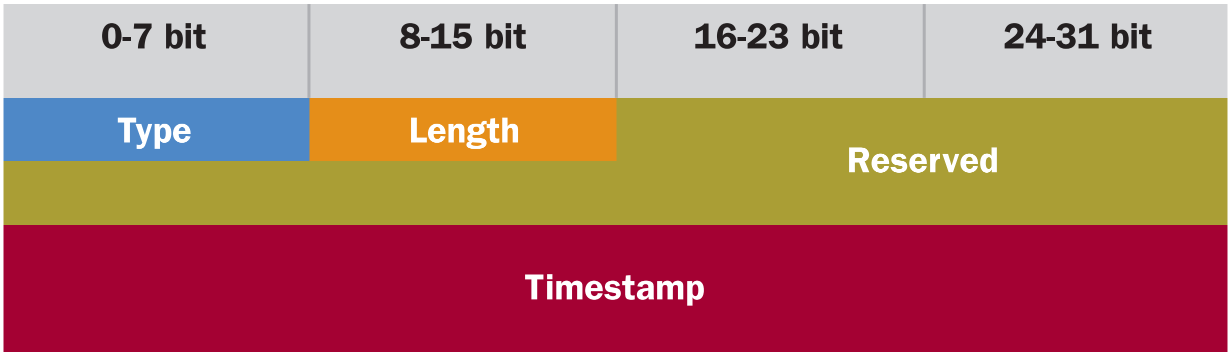

0-7 bit	8-15 bit	16-23 bit	24-31 bit
Type	Length	Reserved	
Timestamp			

Figure 18: SEND Timestamp Option

Figure 19 shows the SEND Nonce option.

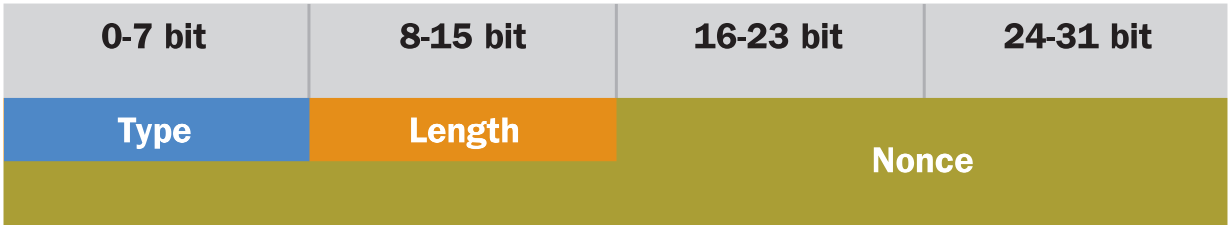

0-7 bit	8-15 bit	16-23 bit	24-31 bit
Type	Length	Nonce	

Figure 19: SEND Nonce Option

The SEND protocol is designed to counter the threats to NDP, including:

- **Neighbor Solicitation/Advertisement Spoofing:** The ability to cause a false entry in a node's Neighbor Cache. SEND utilizes RSA Signature and CGA options in the solicitations to counter these threats.
- **Neighbor Unreachability Detection Failure:** SEND requires a node responding to Neighbor Solicitations sent as NUD probes include an RSA Signature option and proof of authorization to use the interface identifier in the address being probed. If these are not included, the node performing NUD discards the responses.
- **Duplicate Address Detection DoS Attack:** SEND requires Neighbor Advertisements sent as responses to DAD to include an RSA Signature option and proof of authorization to use the interface identifier in the address being tested. If these are not included, the node performing DAD discards the responses. In addition, when a node using SEND performs DAD, it can listen for address collisions from nodes that do not use SEND for the first address generated, but not for new attempts. This provides DoS protection for nodes using SEND.
- **Router Solicitation and Advertisement Attacks:** SEND requires that Router Advertisements contain an RSA Signature using the public key of the node that can prove its authorization to route the subnet prefixes contained in any Prefix Information Options.
- **Replay Attacks:** SEND includes a Nonce option in the solicitation and requires that the advertisement include a matching option. This will create a challenge-response when combined with the signatures. SEND also includes a Timestamp option to protect against unsolicited messages.

While SEND provides significant support against the attacks above, it does not cover every potential scenario. SEND does not prevent ND Denial-of-Service (DoS) attacks where routers are bombarded and kept busy performing Neighbor Solicitations for addresses that do not exist. However, this can be handled using rate limiting techniques when implementing ND on a router. SEND does not provide confidentiality for NDP communications or compensate for unsecured link layers. Unsecured link layers could allow nodes to spoof the link layer address of other nodes. Also, SEND does not require that addresses and Neighbor Advertisements correspond. Multicast Listener Discovery (MLD) contains no provision for security, and an attacker could send a Done message to unsubscribe a node from a Multicast address. However, the node should be able to detect this and respond to avoid being dropped.

In addition to some threats SEND does not cover, there are some DoS attacks against NDP and SEND itself. If an attacker sends a high number of packets that a node has to verify through asymmetric methods, it could leave SEND non-operational. Another vulnerability is in the Authorization Delegation Discovery process where routers are targeted with requests for a large number of certification paths to be discovered for different trust anchors. However, routers can defend against this attack through caching discovered information and limiting the number of different discoveries. Similarly, hosts can also be targeted by sending a large number of unnecessary certification paths. This can force the host to spend useless memory and verification resources. The host can defend itself by limiting the amount of resources devoted to the certification paths and their verification prioritizing advertisements.

4.5 Domain Name Services Security (DNSSEC)

This section provides a high-level overview of DNSEC. A more detailed understanding can be found in RFC 4033. DNSSEC incorporates greater security to the DNS through the use of security extensions to provide DNS data with origin authentication and integrity services. It also includes mechanisms for authenticated denial of the existence of DNS data.

In order to achieve the additional security services, DNSSEC adds four new resource record types:

- Resource Record Signature (RRSIG)
- DNS Public Key (DNSKEY)
- Delegation Signer (DS)
- Next Secure (NSEC)

It also adds two new message header bits:

- Checking Disabled (CD)
- Authenticated Data (AD)

DNSSEC requires Extension Mechanism for DNS (EDNS0) to support larger DNS message sizes that result from adding the DNSSEC RRs. It also requires support for the DNSSEC OK (DO) EDNS header bit. This will allow the security-aware resolver to indicate which queries to receive DNSSEC RRs in response messages.

DNSSEC provides for Data Origin Authentication and Data Integrity. Authentication is achieved through associating cryptographically generated digital signatures with DNS RR sets, which are stored in the RRSIG record. It is possible to use multiple keys, but most implementations will rely on a single private key. Thus, once a resolver knows a zone's public key, that zone's signed data can then be authenticated. The key signing the zone's data is associated with the zone and not with the authoritative name servers. While keys for DNS transaction authentication mechanisms may also appear in zones, DNSSEC itself provides security for the object data and not for DNS transactions.

Zone keys can be distributed to resolvers through a trust anchor or by normal DNS resolution. If DNS resolution is used to distribute keys, the public keys are stored in the DNSKEY RR. Because discovering keys reliably via DNS requires the target key to be signed, this may occur through an authentication chain. But the resolver must be configured with at least one trust anchor.

DNSSEC also provides the ability to authenticate name and type non-existence. It accomplishes this through the NSEC record which allows a security-aware resolver to authenticate a negative reply for either name or type non-existence with the same mechanisms used to authenticate other DNS replies.

DNSSEC is primarily designed to allow delivery of DNS data with enough assurance that can be trusted, but does not provide security services for:

- Confidentiality
- Access control lists
- Differentiating between inquirers
- DoS
- Protecting operations such as zone transfers and dynamic update

4.6 Transport Layer Security (TLS)/Secure Sockets Layer (SSL)

This section provides a high-level overview of TLS (SSL). A more detailed description can be found in RFC 4346. The TLS Protocol (also known as SSL) provides privacy and data integrity between two communicating applications. While some of the security services provided by TLS are similar to that of IPsec, they are applied in a much different fashion and at different layers. Where IPsec has an advantage of working at the IP layer and can thus provide a potentially more ubiquitous service, TLS is more widely deployed and available on just about every Web browser. In practice, the use of TLS is much more widespread than IPsec for three primary reasons.

- **Mass availability:** TLS is deployed on just about every Web browser and Internet facing server. It is utilized to secure transactions from simple logon routines to credit card processing.
- **Key Infrastructure:** TLS has an operational infrastructure for exchanging key material. IPsec is lacking a widely deployed key distribution technique.
- **NAT Traversal:** TLS integrates much better into the existing Internet infrastructure and is more readily able to traverse multiple layers of NAT devices versus IPsec.

In more recent years, TLS has moved beyond supporting individual transaction processing into creating a new capability in private networking called SSL VPNs. SSL VPNs have made significant market penetration for agencies and other organizations that want to provide a secure, remote capability to access IP-based applications.

The TLS protocol is composed of the TLS Record Protocol (shown in figure 20) and the TLS Handshake Protocol (shown in figure 21). The TLS Record Protocol is normally layered on top of TCP and provides the following security services.

- **Confidentiality:** Symmetric cryptography is used for data encryption (e.g., DES, RC4, etc.). Unique keys are generated for each connection and are based on a secret negotiated by the TLS Handshake Protocol.
- **Integrity:** An integrity check employing a secure hash functions is utilized.

<table>
<tr><th>0-7 bit</th><th>8-15 bit</th><th>16-23 bit</th><th>24-31 bit</th></tr>
<tr><td>Content Type</td><td>Version (MSB)</td><td>Version (LSB)</td><td>Length (MSB)</td></tr>
<tr><td>Length (LSB)</td><td colspan="3" rowspan="2">Protocol Messages</td></tr>
<tr><td></td></tr>
<tr><td colspan="4">MAC</td></tr>
<tr><td colspan="4">Padding</td></tr>
</table>

Figure 20: TLS Record Protocol

The TLS Handshake Protocol provides the ability for the server and client to authenticate each other, negotiate an encryption algorithm, and generate cryptographic keys. The TLS Handshake Protocol provides the following security services.

- **Authentication:** The peer's identity can be authenticated using asymmetric cryptographic techniques (e.g., RSA, DSS, etc.).
- **Confidentiality:** Provides for the secure negotiation of a shared secret that is unavailable to eavesdroppers. For any authenticated connection, the secret cannot be obtained, even by a man-in-the-middle attack.
- **Integrity:** The negotiation communication cannot be modified without being detected.

<table>
<tr><th>0-7 bit</th><th>8-15 bit</th><th>16-23 bit</th><th>24-31 bit</th></tr>
<tr><td>22</td><td>Version (MSB)</td><td>Version (LSB)</td><td>Length (MSB)</td></tr>
<tr><td>Length (LSB)</td><td>Message Type</td><td colspan="2">Message Length</td></tr>
<tr><td>Msg. Lngth. Cont.</td><td colspan="3">Handshake Message</td></tr>
<tr><td>Handshake Msg.</td><td>Message Type</td><td colspan="2">Message Length</td></tr>
<tr><td>Message Length</td><td colspan="3">Handshake Message</td></tr>
</table>

Figure 21: TLS Handshake Protocol

While the TLS application protocol (shown in figure 22) is independent for higher level protocols transparency, the TLS standard does not specify how protocols add security with TLS. Thus, vendors will make their own decisions on how to initiate TLS handshaking and how to interpret the authentication certificates exchanged.

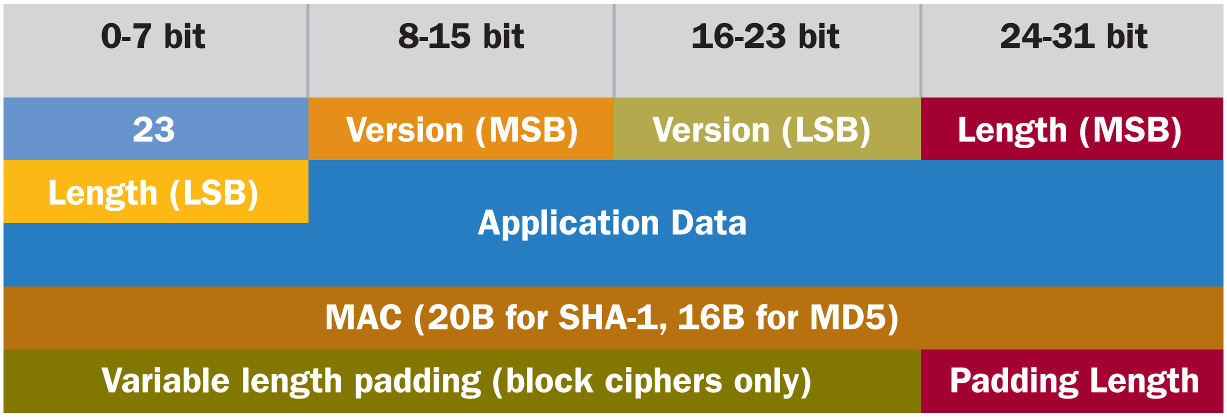

0-7 bit	8-15 bit	16-23 bit	24-31 bit
23	Version (MSB)	Version (LSB)	Length (MSB)
Length (LSB)	Application Data		
MAC (20B for SHA-1, 16B for MD5)			
Variable length padding (block ciphers only)			Padding Length

Figure 22: TLS Application Protocol

4.7 Routing and Addressing Security Considerations

Routing and addressing are critical to the successful deployment of IPv6 in any organization's enterprise, but in addition to having architectural and operational impacts, they can also play a relevant factor in the overall security of an agency's IPv6 implementation.

Traditionally on the Internet, security of the routing infrastructure has been secured through various means and vendor-specific implementations. More recently though, security of the routing infrastructure has been identified as vulnerable at a high level and several efforts have begun to incorporate greater security. The Department of Homeland Security's CyberSecurity and Science and Technology office has been funding research to improve routing security. In addition, NIST also has teams looking at security of routing infrastructure, not to mention many other agencies that are looking for solutions to their immediate and future security needs.

The IETF has two working groups focused on routing security. The Routing Protocol Security Requirements (rpsec) and the Secure Inter-Domain Routing (Sidr) working groups are in the process of defining the threats to the routing infrastructure and are working on methods and protocols to mitigate those threats. RFC 4593 provides an overview of generic threats to routing protocols and several Internet-Drafts have been created to support routing security.

At a minimum, agencies should be aware of threats to their routing infrastructure and factor these into their specific IPv6 implementations. In the short term, many of the same techniques utilized to protect the IPv4 routing infrastructure will be utilized to protect he IPv6 routing infrastructure. Many of these techniques focus on protection of the physical router and the operating systems as opposed to the exchange of routing information.

IPv6 addressing can play a role in overall security of an agency's IPv6 implementation in multiple ways, including:

- **Address Obscurity**: Given the enormous size potential of a single IPv6 subnet, the concept of brute-force scanning methods for hackers to identify targets of opportunity could quickly become history. The typical IPv6 subnet will consist of 264 potential IPv6 addresses, which would be impossible for a hack to scan because of the sizeable amount of address space. Thus, as long as agencies take care in determining IPv6 addresses (don't use obvious selections or a set numbering scheme hackers can identify) and keep assets, such as DNS servers that maintain IPv6 address information protected, it will be very difficult for hackers to identify hosts to attack.
- **Network Based Authentication**: The deployment of globally unique addresses coupled with widespread IPsec availability could provide agencies the ability to establish trust relationships based on IPv6 addresses. This could be established as a network centric service that could then be applied to higher level protocols and applications.

- **Communities of Interest (COI):** The ability to implement multiple IPv6 address on a single interface provides for the potential to efficiently build scalable COI. Thus, each node could belong to one or more groups each of which can have its own security policy. Groups could then have unique settings and characteristics at a very granular level that could be enforced in part through the management of the IPv6 address and packet flows within the enterprise and entering/exiting the enterprise.
- **Subnet Topology Masking:** This can be accomplished in smaller deployment by utilizing centrally assigned pools of addresses in combination with host routes in the interior gateway protocols. For larger deployments, the Home Agent tunneling approach can be employed with firewalls configured to block outbound binding update messages. In addition, a border Home Agent can utilize internal tunneling to the node and completely mask all internal topology.

4.8 Network Address Translation (NAT)

This section provides an overview of NAT and its relationship to IPv6. A more detailed analysis of NAT functionality regarding IPv6 can be found in RFC 4864.

NAT has played a critical role within the Internet community and throughout enterprises across the globe. NAT is commonly utilized as a standard method of operation in nearly every federal, state, and local agency. While NAT has been extremely valuable in many regards, it has also created numerous challenges in developing scalable solutions and end-to-end services. NAT has been a major point of discussion in the IPv6 community about how to remove NAT and replace its functionality and perceived benefits with other mechanisms within the IPv6 protocols.

As agencies build their IPv6 Transition Plans and are considering security, they will need to consider which aspects of their IPv4 security architecture they will want to maintain, which they will change, and over what timeframe. It is important to note that many of the IPv6 advanced capabilities and features will require agencies to transition their approach to security. Thus, while implementing a security architecture for IPv6 similar to their IPv4 approach is appealing and may happen in the short term, agencies need to consider how that architecture will evolve to support the future needs of the agency and next-generation applications.

NAT implementations provide a number of potential benefits for IPv4 networks and, while some will argue that some of the benefits are more perception than reality, they are important to agencies and need to be considered when transitioning to IPv6. Some of the perceived benefits of NAT include:

- **Simple Gateway Between Internet and Private Network:** One of the most common benefits from NATs is the ability to establish an interface or choke point between a private network and the Internet (or another network). In some cases, organizations use NATs internally to create separate enclaves of control or influence. On the internal or private network, the administrators can use either non-routable (private) addresses or globally unique addresses. If NATs are used with private address space, they are typically easy to setup and a simple interface or default configuration is sufficient for configuring both device and application access rights.
- **Simple Security Due to Stateful Filter Implementation:** There is a common belief that NATs provide protection to internal networks by providing an “air gap” between the outside world and the internal networks. While a NAT does provide a measure of protection against some rudimentary attacks, it is more likely to provide a false sense of complacency regarding security. Hackers understand NATs very well and have been effective at designing attacks that easily circumvent NAT-only devices. Many NAT devices have begun building in basic firewall capabilities to offer better protection. Thus the real protection is generated by the firewalling and not the NAT functionality.
- **User/Application Tracking:** Some NAT devices can be utilized to track usage and map Internet session back to specific users. Depending on the actual device and configuration, this can be a straight forward task or potentially be very complicated.
- **Privacy and Topology Hiding:** NATs are relatively effective at hiding internal network topology and providing anonymity to users at an IP layer. Since NATs normally are configured to support a many-to-one (or few) approach of many internal private addresses to one (or few) external addresses, outside entities are not able to see the structure of the internal network from the IP packets once they leave the NAT device. Thus, the outside world only sees the publicly facing IP addresses associated with the outside interface of the NAT.

- **Independent Control of Addressing in a Private Network:** Many networks utilizing NATs employ private (or non-routable) address space. This allows them to utilize extremely large blocks of IPv4 addresses without having to provide justification or to report on the management of their address spaces. The downside is that these addresses cannot be routed on the public Internet. This NAT functionality also provides organizations with the ability to more easily switch providers or renumber IP addresses of public facing devices, because the internal network infrastructure is immune to those modifications. The address space available for private networks include:
  - 10.0.0.0 - 10.255.255.255 (10/8 prefix),
  - 172.16.0.0 - 172.31.255.255 (172.16/12 prefix), and
  - 192.168.0.0 - 192.168.255.255 (192.168/16 prefix).
- **Global Address Pool Conservation:** NAT has been very successful in supporting the conservation of the IPv4 address space. While the number of IPv4 addresses available for distribution has dropped to a critical level, and the industry has been warned about IPv4 address exhaustion within the next three to five years, NAT has been a key tool in the slowing of significant IPv4 address space allocations, thus extending the useful life of IPv4 over an extra decade.

Some tools within IPv6 that are effective in replacing the functionality found in NAT devices include:

- **Privacy Addresses (RFC 4941):** Nodes use IPv6 stateless address autoconfiguration to generate addresses using a combination of locally available information and information advertised by routers. Addresses are formed by combining network prefixes with an interface identifier. On an interface that contains an embedded IEEE identifier, the interface identifier is typically derived from it. On other interface types, the interface identifier is generated through other means, for example, via random number generation. This RFC describes an extension to IPv6 stateless address autoconfiguration for interfaces whose interface identifier is derived from an IEEE identifier. Use of the extension causes nodes to generate global scope addresses from interface identifiers that change over time, even in cases where the interface contains an embedded IEEE identifier. Changing the interface identifier (and the global scope addresses generated from it) over time makes it more difficult for eavesdroppers and other information collectors to identify when different addresses used in different transactions actually correspond to the same node.
- **Unique Local Addresses (RFC 4193):** This is an IPv6 unicast address format that is globally unique and is intended for local communications. These addresses are called Unique Local IPv6 Unicast Addresses and are not expected to be routable on the global Internet. They are routable inside of a more limited area such as a site. They may also be routed between a limited set of sites. Local IPv6 unicast addresses have the following characteristics:
  - Globally unique prefix (with high probability of uniqueness),
  - Well-known prefix to allow for easy filtering at site boundaries,
  - Allow sites to be combined or privately interconnected without creating any address conflicts or requiring renumbering of interfaces that use these prefixes,
  - ISP independent and can be used for communications inside of a site without having any permanent or intermittent Internet connectivity,
  - If accidentally leaked outside of a site via routing or DNS, there is no conflict with any other addresses, and
  - In practice, applications may treat these addresses like global scoped addresses.
- **DHCPv6 Prefix Delegation:** Provide for a mechanism to automatically delegate IPv6 prefixes using DHCPv6. This is intended for longer terms delegations from delegating routers or possibly DHCPv6 services to routers.
- **Untraceable IPv6 Addresses:** Used to create an infrastructure that looks anomalous to the outside world and can resist attempts to map network activities. Given the large availability of IPv6 addresses, particularly within a subnet structure that utilizes the 264 lower order bits for unique host identification, random address generators can be used effectively. Thus, random IPv6 addresses can be assigned under the organization's routing prefix, and when combined with the ability to apply multiple IPv6 addresses to a single interface, they can allow unique IPv6 addresses for individual TCP sessions while maintaining a more permanent IPv6 address structure for other activities.

Table 2 below shows a comparison of some of the IPv4 functions that can be achieved in conjunction with NAT and how they can be achieved within IPv6 without NAT.

NAT Functionality	IPv6 Implementation
Simple Gateway Between Internet and Private Network	Basic IPv6 routers should have default configurations to advertise random ULA prefixes within a local site that are independent from external connection states. Thus, local nodes in multi-link topologies would be able to communicate without dependence on global connectivity. When global or outside connectivity is established, DHCP-PD can be utilized to obtain the routing prefix from the upstream delegating source.
Simple Security Due to Stateful Filter Implementation	While some of the specific threats may vary, IPv6 hosts connected to the Internet are vulnerable as are IPv4 hosts connected to the Internet. The following protections are available without the use of NAT while maintaining end-to-end reachability. · IPv6 nodes will not respond to packets destined for an address once the IPv6 address lifetime has expired. Thus, shorter privacy extension lifetimes minimize attack profiles. · Greater availability of IPsec in IPv6. While IPsec can be used with IPv4, there are more implementation issues, specifically regarding the use of IPsec across NAT devices. As mentioned previously, IPsec provides some basic firewalling features at the node level and provides an infrastructure for network-level authentication. · The sheer size of the IPv6 address space makes it very unlikely subnet scans will be successful in detecting potential IPv6 targets. Thus, malicious intruders would need to identify internal topology through other methods, for which other IPv6 protections can be applied.
User/Application Tracking	The unique nature of IPv6 addresses provides agencies the ability to track user and application information at a granular level. In addition, when combined with IPsec security features, the data correlation can be provided with a significantly greater level of integrity. Thus, it is easier to match specific data flow with nodes or even users with a higher degree of confidence.
User/Application Tracking **Privacy and Topology Hiding**	Privacy addresses (RFC 4941) can be utilized to achieve partial host privacy. By limiting a session address to a specific lifetime, a session can only be traced to the originating subnet. When combining subnet and interface identifier randomization, hackers able to see IPv6 packets will not be able to garner much information about the networks topology. The ability for an attacker to gain a single address, or even multiple addresses with limited lifetimes, will reveal little information. While this method should suffice for most cases, an agency could completely hide its internal IPv6 topology and internal use addresses using ULAs. In the extreme situation where an agency wants to fully conceal all or a portion of its internal IPv6 topology, it can utilize Mobile IPv6. Thus, the home agent will act as the interface to the outside world and specific nodes and internal topology will remain hidden.
Independent Control of Addressing in a Private Network	ULAs provide the ability for organizations to have local use of autonomy for IPv6 addresses. This can be combined with the ability to layer multiple IPv6 addresses onto a single interface. Thus, ULAs can be utilized for the internal addressing structure and globally routable IPv6 can be deployed later or in an as-needed capacity.
Global Address Pool Conservation	IPv4 addresses are quickly becoming a scarce commodity. Over the past decade, the IPv4 address space management and the requirements for obtaining new or incremental space have become more challenging. While care should be taken to not needlessly waste IPv6 address space, the focus on international and regional addressing policies is on hierarchical routing and address aggregation to maintain a reasonable size in local and global routing tables. Therefore, this is not considered an issue with IPv6.

Table 2: Comparison of NAT (IPv4) and IPv6 Functionality

5 Conclusion

Convergence is forcing not only the USG, but the entire industry to have a greater reliance on IP-based infrastructures. IP-based infrastructures are quickly becoming the underlying engine that allows advanced capabilities to be rapidly developed and deployed. As agencies move toward the introduction of next generation systems to support collaborative architectures, geospatial application, netcentric warfare, mobility, and COOP, they will face significant issues with a legacy IPv4 enterprise. Understanding the future ramifications of this situation, OMB has issued a policy requiring all agencies to begin activities for the transition to the next generation of IP, IPv6.

During the transitioning cycle to IPv6, agencies will face a number of challenges and one of the greatest will be in the area of security. The movement back towards end-to-end services will not only impact the way services are supported in the enterprise and across the Internet, it will also radically change the way information security is viewed and implemented within the enterprise. The transition to IPv6 provides a perfect time for agencies to begin re-architecting their enterprise security solutions to support end-to-end and other enhanced capabilities. During the transition process a number of critical implementation issues must be considered from a security perspective, including:

- Governance and Policy Needed,
- Training,
- Compliance Testing,
- Institutionalized IPv6, and
- New Attack Surface.

Some of the first steps any agency should accomplish with regard to security include:

- IPv6 Security Plan,
- Policy,
- Routers/Switches,
  - Disable IPv6/Tunnels
  - ACL to Block IPv6/Tunnels on core/edge/outside enclave
- Network Protection Devices/Tools,
  - Contact vendors for IPv6 advice
- Block IPv6 (Type 41) Tunnels,
- Enable IPv6 IDS/IPS features,
- End Nodes, and
  - Enable IPv6 host firewalls on all end devices
  - Disable IPv6 if not used
- Monitor Core and Enclave Boundaries.

Appendix A: IPv6 Security RFCs

The following table provides a list of RFCs with relevance to IP based security and IPv6. While all of the RFCs listed below are not solely focused on security or IPv6, the list may provide agencies with a better understanding of how they need to approach their IPv6 transition effort. Note that this list is not intended to be a listing of all the standards an agency should look to include in their solution sets, but more of a listing of the more relevant RFCs they may wish to review and consider during their transition.

RFC#	Title	Abstract
4945	The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2, and PKIX	The Internet Key Exchange (IKE) and Public Key Infrastructure for X.509 (PKIX) certificate profile both provide frameworks that must be profiled for use in a given application. This document provides a profile of IKE and PKIX that defines the requirements for using PKI technology in the context of IKE/IPsec. The document complements protocol specifications such as IKEv1 and IKEv2, which assume the existence of public key certificates and related keying materials, but which do not address PKI issues explicitly. This document addresses those issues. The intended audience is implementers of PKI for IPsec.
4941	Privacy Extensions for Stateless Address Autoconfiguration in IPv6	Nodes use IPv6 stateless address autoconfiguration to generate addresses using a combination of locally available information and information advertised by routers. Addresses are formed by combining network prefixes with an interface identifier. On an interface that contains an embedded IEEE Identifier, the interface identifier is typically derived from it. On other interface types, the interface identifier is generated through other means, for example, via random number generation. This document describes an extension to IPv6 stateless address autoconfiguration for interfaces whose interface identifier is derived from an IEEE identifier. Use of the extension causes nodes to generate global scope addresses from interface identifiers that change over time, even in cases where the interface contains an embedded IEEE identifier. Changing the interface identifier (and the global scope addresses generated from it) over time makes it more difficult for eavesdroppers and other information collectors to identify when different addresses used in different transactions actually correspond to the same node.
4894	Use of Hash Algorithms in Internet Key Exchange (IKE) and IPsec Use of Hash Algorithms in Internet Key Exchange (IKE) and IPsec	This document describes how the IKEv1 (Internet Key Exchange version 1), IKEv2, and IPsec protocols use hash functions, and explains the level of vulnerability of these protocols to the reduced collision resistance of the MD5 and SHA-1 hash algorithms.
4891	Using IPsec to Secure IPv6-in-IPv4 Tunnels	This document gives guidance on securing manually configured IPv6-in-IPv4 tunnels using IPsec in transport mode. No additional protocol extensions are described beyond those available with the IPsec framework.
4890	Recommendations for Filtering ICMPv6 Messages in Firewalls	In networks supporting IPv6, the Internet Control Message Protocol version 6 (ICMPv6) plays a fundamental role with a large number of functions, and a correspondingly large number of message types and options. ICMPv6 is essential to the functioning of IPv6, but there are a number of security risks associated with uncontrolled forwarding of ICMPv6 messages. Filtering strategies designed for the corresponding protocol, ICMP, in IPv4 networks are not directly applicable, because these strategies are intended to accommodate a useful auxiliary protocol that may not be required for correct functioning. This document provides some recommendations for ICMPv6 firewall filter configuration that will allow propagation of ICMPv6 messages that are needed to maintain the functioning of the network while dropping messages that are potential security risks.
4882	IP Address Location Privacy and Mobile IPv6: Problem Statement	This document discusses location privacy as applicable to Mobile IPv6. We document the concerns arising from revealing a Home Address to an onlooker and from disclosing a Care-of Address to a correspondent.
4877	Mobile IPv6 Operation with IKEv2 and the Revised IPsec Architecture	This document describes Mobile IPv6 operation with the revised IPsec architecture and IKEv2.
4869	Suite B Cryptographic Suites for IPsec	This document proposes four optional cryptographic user interface suites for IPsec, similar to the two suites specified in RFC 4308. The four new suites provide compatibility with the U.S. National Security Agency's Suite B specifications.

4868	Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec	This specification describes the use of Hashed Message Authentication Mode (HMAC) in conjunction with the SHA-256, SHA-384, and SHA-512 algorithms in IPsec. These algorithms may be used as the basis for data origin authentication and integrity verification mechanisms for the Authentication Header (AH), Encapsulating Security Payload (ESP), Internet Key Exchange Protocol (IKE), and IKEv2 protocols, and also as Pseudo-Random Functions (PRFs) for IKE and IKEv2. Truncated output lengths are specified for the authentication-related variants, with the corresponding algorithms designated as HMAC-SHA-256-128, HMAC-SHA-384-192, and HMAC-SHA-512-256. The PRF variants are not truncated, and are called PRF-HMAC-SHA-256, PRF-HMAC-SHA-384, and PRF-HMAC-SHA-512.
4864	Local Network Protection for IPv6	Although there are many perceived benefits to Network Address Translation (NAT), its primary benefit of amplifying available address space is not needed in IPv6. In addition to NAT's many serious disadvantages, there is a perception that other benefits exist, such as a variety of management and security attributes that could be useful for an Internet Protocol site. IPv6 was designed with the intention of making NAT unnecessary, and this document shows how Local Network Protection (LNP) using IPv6 can provide the same or more benefits without the need for address translation.
4852	IPv6 Enterprise Network Analysis - IP Layer 3 Focus	This document analyzes the transition to IPv6 in enterprise networks focusing on IP Layer 3. These networks are characterized as having multiple internal links and one or more router connections to one or more Providers, and as being managed by a network operations entity. The analysis focuses on a base set of transition notational networks and requirements expanded from a previous document on enterprise scenarios. Discussion is provided on a focused set of transition analysis required for the enterprise to transition to IPv6, assuming a Dual-IP layer (IPv4 and IPv6) network and node environment within the enterprise. Then, a set of transition mechanisms are recommended for each notational network.
4843	An IPv6 Prefix for Overlay Routable Cryptographic Hash Identifiers (ORCHID)	This document introduces Overlay Routable Cryptographic Hash Identifiers (ORCHID) as a new, experimental class of IPv6-address- like identifiers. These identifiers are intended to be used as endpoint identifiers at applications and Application Programming Interfaces (API) and not as identifiers for network location at the IP layer, i.e., locators. They are designed to appear as application layer entities and at the existing IPv6 APIs, but they should not appear in actual IPv6 headers. To make them more like generic IPv6 addresses, they are expected to be routable at an overlay level. Consequently, while they are considered non-routable addresses from the IPv6 layer point-of-view, all existing IPv6 applications are expected to be able to use them in a manner compatible with current IPv6 addresses. This document requests IANA to allocate a temporary prefix out of the IPv6 addressing space for Overlay Routable Cryptographic Hash Identifiers. By default, the prefix will be returned to IANA in 2014, with continued use requiring IETF consensus.
4835	Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)	The IPsec series of protocols makes use of various cryptographic algorithms in order to provide security services. The Encapsulating Security Payload (ESP) and the Authentication Header (AH) provide two mechanisms for protecting data being sent over an IPsec Security Association (SA). To ensure interoperability between disparate implementations, it is necessary to specify a set of mandatory-to-implement algorithms to ensure that there is at least one algorithm that all implementations will have available. This document defines the current set of mandatory-to-implement algorithms for ESP and AH as well as specifying algorithms that should be implemented because they may be promoted to mandatory at some future time.
4818	RADIUS Delegated-IPv6-Prefix Attribute	This document defines a RADIUS (Remote Authentication Dial In User Service) attribute that carries an IPv6 prefix that is to be delegated to the user. This attribute is usable within either RADIUS or Diameter.
4809	Requirements for an IPsec Certificate Management Profile	This informational document describes and identifies the requirements for transactions to handle Public Key Certificate (PKC) lifecycle transactions between Internet Protocol Security (IPsec) Virtual Private Network (VPN) Systems using Internet Key Exchange (IKE) (versions 1 and 2) and Public Key Infrastructure (PKI) Systems. These requirements are designed to meet the needs of enterprise-scale IPsec VPN deployments. It is intended that a standards track profile of a management protocol will be created to address many of these requirements.
4807	IPsec Security Policy Database Configuration MIB	This document defines a Structure of Management Information Version 2 (SMIv2) Management Information Base (MIB) module for configuring the security policy database of a device implementing the IPsec protocol. The policy-based packet filtering and the corresponding execution of actions described in this document are of a more general nature than for IPsec configuration alone, such as for configuration of a firewall. This MIB module is designed to be extensible with other enterprise or standards-based defined packet filters and actions.

4798	Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)	This document explains how to interconnect IPv6 islands over a Multiprotocol Label Switching (MPLS)-enabled IPv4 cloud. This approach relies on IPv6 Provider Edge routers (6PE), which are Dual Stack in order to connect to IPv6 islands and to the MPLS core, which is only required to run IPv4 MPLS. The 6PE routers exchange the IPv6 reachability information transparently over the core using the Multiprotocol Border Gateway Protocol (MP-BGP) over IPv4. In doing so, the BGP Next Hop field is used to convey the IPv4 address of the 6PE router so that dynamically established IPv4-signaled MPLS Label Switched Paths (LSPs) can be used without explicit tunnel configuration.
4754	IKE and IKEv2 Authentication Using the Elliptic Curve Digital Signature Algorithm (ECDSA)	This document describes how the Elliptic Curve Digital Signature Algorithm (ECDSA) may be used as the authentication method within the Internet Key Exchange (IKE) and Internet Key Exchange version 2 (IKEv2) protocols. ECDSA may provide benefits including computational efficiency, small signature sizes, and minimal bandwidth compared to other available digital signature methods. This document adds ECDSA capability to IKE and IKEv2 without introducing any changes to existing IKE operation.
4739	Multiple Authentication Exchanges in the Internet Key Exchange (IKEv2) Protocol	The Internet Key Exchange (IKEv2) protocol supports several mechanisms for authenticating the parties, including signatures with public-key certificates, shared secrets, and Extensible Authentication Protocol (EAP) methods. Currently, each endpoint uses only one of these mechanisms to authenticate itself. This document specifies an extension to IKEv2 that allows the use of multiple authentication exchanges, using either different mechanisms or the same mechanism. This extension allows, for instance, performing certificate-based authentication of the client host followed by an EAP authentication of the user. When backend authentication servers are used, they can belong to different administrative domains, such as the network access provider and the service provider.
4718	IKEv2 Clarifications and Implementation Guidelines	This document clarifies many areas of the IKEv2 specification. It does not introduce any changes to the protocol, but rather provides descriptions that are less prone to ambiguous interpretations. The purpose of this document is to encourage the development of interoperable implementations.
4671	RADIUS Accounting Server MIB for IPv6	This memo defines a set of extensions that instrument RADIUS accounting server functions. These extensions represent a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. Using these extensions, IP-based management stations can manage RADIUS accounting servers. This memo obsoletes RFC 2621 by deprecating the MIB table containing IPv4-only address formats and defining a new table to add support for version-neutral IP address formats. The remaining MIB objects from RFC 2621 are carried forward into this document. This memo also adds UNITS and REFERENCE clauses to selected objects.
4670	RADIUS Accounting Client MIB for IPv6	This memo defines a set of extensions that instrument RADIUS accounting client functions. These extensions represent a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. Using these extensions, IP-based management stations can manage RADIUS accounting clients. This memo obsoletes RFC 2620 by deprecating the MIB table containing IPv4-only address formats and defining a new table to add support for version-neutral IP address formats. The remaining MIB objects from RFC 2620 are carried forward into this document. This memo also adds UNITS and REFERENCE clauses to selected objects.
4669	RADIUS Authentication Server MIB for IPv6	This memo defines a set of extensions that instrument RADIUS authentication server functions. These extensions represent a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. Using these extensions, IP-based management stations can manage RADIUS authentication servers. This memo obsoletes RFC 2619 by deprecating the MIB table containing IPv4-only address formats and defining a new table to add support for version-neutral IP address formats. The remaining MIB objects from RFC 2619 are carried forward into this document. This memo also adds UNITS and REFERENCE clauses to selected objects.
4668	RADIUS Authentication Client MIB for IPv6	This memo defines a set of extensions that instrument RADIUS authentication client functions. These extensions represent a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. Using these extensions, IP-based management stations can manage RADIUS authentication clients. This memo obsoletes RFC 2618 by deprecating the MIB table containing IPv4-only address formats and defining a new table to add support for version-neutral IP address formats. The remaining MIB objects from RFC 2618 are carried forward into this document. The memo also adds UNITS and REFERENCE clauses to selected objects.

4659	BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN	This document describes a method by which a Service Provider may use its packet-switched backbone to provide Virtual Private Network (VPN) services for its IPv6 customers. This method reuses, and extends where necessary, the BGP/MPLS IP VPN method for support of IPv6. In BGP/MPLS IP VPN, Multiprotocol BGP is used for distributing IPv4 VPN routes over the service provider backbone, and MPLS is used to forward IPv4 VPN packets over the backbone. This document defines an IPv6 VPN address family and describes the corresponding IPv6 VPN route distribution in Multiprotocol BGP. This document defines support of the IPv6 VPN service over both an IPv4 and an IPv6 backbone, and for using various tunneling techniques over the core, including MPLS, IP-in-IP, Generic Routing Encapsulation (GRE) and IPsec protected tunnels. The inter-working between an IPv4 site and an IPv6 site is outside the scope of this document.
4651	A Taxonomy and Analysis of Enhancements to Mobile IPv6 Route Optimization	This document describes and evaluates strategies to enhance Mobile IPv6 Route Optimization, on the basis of existing proposals, in order to motivate and guide further research in this context. This document is a product of the IP Mobility Optimizations (MobOpts) Research Group.
4621	Design of the IKEv2 Mobility and Multihoming (MOBIKE) Protocol	The IKEv2 Mobility and Multihoming (MOBIKE) protocol is an extension of the Internet Key Exchange Protocol version 2 (IKEv2). These extensions should enable an efficient management of IKE and IPsec Security Associations when a host possesses multiple IP addresses and/or where IP addresses of an IPsec host change over time (for example, due to mobility). This document discusses the involved network entities and the relationship between IKEv2 signaling and information provided by other protocols. Design decisions for the MOBIKE protocol, background information, and discussions within the working group are recorded.
4615	The Advanced Encryption Standard-Cipher-based Message Authentication Code-Pseudo-Random Function-128 (AES-CMAC-PRF-128) Algorithm for the Internet Key Exchange Protocol (IKE)	Some implementations of IP Security (IPsec) may want to use a pseudo-random function (PRF) based on the Advanced Encryption Standard (AES). This memo describes such an algorithm, called AES-CMAC-PRF-128. It supports fixed and variable key sizes.
4593	Generic Threats to Routing Protocols	Routing protocols are subject to attacks that can harm individual users or network operations as a whole. This document provides a description and a summary of generic threats that affect routing protocols in general. This work describes threats, including threat sources and capabilities, threat actions, and threat consequences, as well as a breakdown of routing functions that might be attacked separately.
4555	IKEv2 Mobility and Multihoming Protocol (MOBIKE)	This document describes the MOBIKE protocol, a mobility and multihoming extension to Internet Key Exchange (IKEv2). MOBIKE allows the IP addresses associated with IKEv2 and tunnel mode IPsec Security Associations to change. A mobile Virtual Private Network (VPN) client could use MOBIKE to keep the connection with the VPN gateway active while moving from one address to another. Similarly, a multihomed host could use MOBIKE to move the traffic to a different interface if, for instance, the one currently being used stops working.
4554	Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks	Ethernet VLANs are quite commonly used in enterprise networks for the purposes of traffic segregation. This document describes how such VLANs can be readily used to deploy IPv6 networking in an enterprise, which focuses on the scenario of early deployment prior to availability of IPv6-capable switch-router equipment. In this method, IPv6 may be routed in parallel with the existing IPv4 in the enterprise and delivered at Layer 2 via VLAN technology. The IPv6 connectivity to the enterprise may or may not enter the site via the same physical link.
4543	The Use of Galois Message Authentication Code (GMAC) in IPsec ESP and AH	This memo describes the use of the Advanced Encryption Standard (AES) Galois Message Authentication Code (GMAC) as a mechanism to provide data origin authentication, but not confidentiality, within the IPsec Encapsulating Security Payload (ESP) and Authentication Header (AH). GMAC is based on the Galois/Counter Mode (GCM) of operation, and can be efficiently implemented in hardware for speeds of 10 gigabits per second and above, and is also well-suited to software implementations.
4494	The AES-CMAC-96 Algorithm and Its Use with IPsec	The National Institute of Standards and Technology (NIST) has recently specified the Cipher-based Message Authentication Code (CMAC), which is equivalent to the One-Key CBC-MAC1 (OMAC1) algorithm submitted by Iwata and Kurosawa. OMAC1 efficiently reduces the key size of Extended Cipher Block Chaining mode (XCBC). This memo specifies the use of CMAC mode on the authentication mechanism of the IPsec Encapsulating Security Payload (ESP) and the Authentication Header (AH) protocols. This new algorithm is named AES-CMAC-96.

4487	Mobile IPv6 and Firewalls: Problem Statement	This document captures the issues that may arise in the deployment of IPv6 networks when they support Mobile IPv6 and firewalls. The issues are not only applicable to firewalls protecting enterprise networks, but are also applicable in 3G mobile networks such as General Packet Radio Service / Universal Mobile Telecommunications System (GPRS/UMTS) and CDMA2000 networks. The goal of this document is to highlight the issues with firewalls and Mobile IPv6 and act as an enabler for further discussion. Issues identified here can be solved by developing appropriate solutions.
4472	Operational Considerations and Issues with IPv6 DNS	This memo presents operational considerations and issues with IPv6 Domain Name System (DNS), including a summary of special IPv6 addresses, documentation of known DNS implementation misbehavior, recommendations and considerations on how to perform DNS naming for service provisioning and for DNS resolver IPv6 support, considerations for DNS updates for both the forward and reverse trees, and miscellaneous issues. This memo is aimed to include a summary of information about IPv6 DNS considerations for those who have experience with IPv4 DNS.
4449	Securing Mobile IPv6 Route Optimization Using a Static Shared Key	A mobile node and a correspondent node may preconfigure data useful for precomputing a Binding Management Key that can subsequently be used for authorizing Binding Updates.
4434	The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)	Some implementations of IP Security (IPsec) may want to use a pseudo-random function derived from the Advanced Encryption Standard (AES). This document describes such an algorithm, called AES-XCBC-PRF-128.
4380	Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)	We propose here a service that enables nodes located behind one or more IPv4 Network Address Translations (NATs) to obtain IPv6 connectivity by tunneling packets over UDP; we call this the Teredo service. Running the service requires the help of Teredo servers and Teredo relays. The Teredo servers are stateless, and only have to manage a small fraction of the traffic between Teredo clients; the Teredo relays act as IPv6 routers between the Teredo service and the native IPv6 Internet. The relays can also provide interoperability with hosts using other transition mechanisms such as 6to4.
4346	The Transport Layer Security (TLS) Protocol Version 1.1	This document specifies Version 1.1 of the Transport Layer Security (TLS) protocol. The TLS protocol provides communications security over the Internet. The protocol allows client/server applications to communicate in a way that is designed to prevent eavesdropping, tampering, or message forgery.
4339	IPv6 Host Configuration of DNS Server Information Approaches	This document describes three approaches for IPv6 recursive DNS server address configuration. It details the operational attributes of three solutions: RA option, DHCPv6 option, and well-known anycast addresses for recursive DNS servers. Additionally, it suggests the deployment scenarios in four kinds of networks (ISP, enterprise, 3GPP, and unmanaged networks) considering multi-solution resolution.
4322	Opportunistic Encryption using the Internet Key Exchange (IKE)	This document describes opportunistic encryption (OE) as designed and implemented by the Linux FreeS/WAN project. OE uses the Internet Key Exchange (IKE) and IPsec protocols. The objective is to allow encryption for secure communication without any pre-arrangement specific to the pair of systems involved. DNS is used to distribute the public keys of each system involved. This is resistant to passive attacks. The use of DNS Security (DNSSEC) secures this system against active attackers as well. As a result, the administrative overhead is reduced from the square of the number of systems to a linear dependence, and it becomes possible to make secure communication the default even when the partner is not known in advance.
4312	The Camellia Cipher Algorithm and Its Use With IPsec	This document describes the use of the Camellia block cipher algorithm in Cipher Block Chaining Mode, with an explicit Initialization Vector, as a confidentiality mechanism within the context of the IPsec Encapsulating Security Payload (ESP).
4309	Using Advanced Encryption Standard (AES) CCM Mode with IPsec Encapsulating Security Payload (ESP)	This document describes the use of Advanced Encryption Standard (AES) in Counter with CBC-MAC (CCM) Mode, with an explicit initialization vector (IV), as an IPsec Encapsulating Security Payload (ESP) mechanism to provide confidentiality, data origin authentication, and connectionless integrity.
4308	Cryptographic Suites for IPsec	The IPsec, Internet Key Exchange (IKE), and IKEv2 protocols rely on security algorithms to provide privacy and authentication between the initiator and responder. There are many such algorithms available, and two IPsec systems cannot interoperate unless they are using the same algorithms. This document specifies optional suites of algorithms and attributes that can be used to simplify the administration of IPsec when used in manual keying mode, with IKEv1 or with IKEv2.

4307	Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)	The IPsec series of protocols makes use of various cryptographic algorithms in order to provide security services. The Internet Key Exchange (IKE (RFC 2409) and IKEv2) provide a mechanism to negotiate which algorithms should be used in any given association. However, to ensure interoperability between disparate implementations, it is necessary to specify a set of mandatory-to-implement algorithms to ensure that there is at least one algorithm that all implementations will have available. This document defines the current set of algorithms that are mandatory to implement as part of IKEv2, as well as algorithms that should be implemented because they may be promoted to mandatory at some future time.
4306	Internet Key Exchange (IKEv2) Protocol	This document describes version 2 of the Internet Key Exchange (IKE) protocol. IKE is a component of IPsec used for performing mutual authentication and establishing and maintaining security associations (SAs). This version of the IKE specification combines the contents of what were previously separate documents, including Internet Security Association and Key Management Protocol (ISAKMP, RFC 2408), IKE (RFC 2409), the Internet Domain of Interpretation (DOI, RFC 2407), Network Address Translation (NAT) Traversal, Legacy authentication, and remote address acquisition. Version 2 of IKE does not interoperate with version 1, but it has enough of the header format in common that both versions can unambiguously run over the same UDP port.
4305	Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)	The IPsec series of protocols makes use of various cryptographic algorithms in order to provide security services. The Encapsulating Security Payload (ESP) and the Authentication Header (AH) provide two mechanisms for protecting data being sent over an IPsec Security Association (SA). To ensure interoperability between disparate implementations, it is necessary to specify a set of mandatory-to-implement algorithms to ensure that there is at least one algorithm that all implementations will have available. This document defines the current set of mandatory-to-implement algorithms for ESP and AH as well as specifying algorithms that should be implemented because they may be promoted to mandatory at some future time.
4304	Extended Sequence Number (ESN) Addendum to IPsec Domain of Interpretation (DOI) for Internet Security Association and Key Management Protocol (ISAKMP)	The IP Security Authentication Header (AH) and Encapsulating Security Payload (ESP) protocols use a sequence number to detect replay. This document describes extensions to the Internet IP Security Domain of Interpretation (DOI) for the Internet Security Association and Key Management Protocol (ISAKMP). These extensions support negotiation of the use of traditional 32-bit sequence numbers or extended (64-bit) sequence numbers (ESNs) for a particular AH or ESP security association.
4303	IP Encapsulating Security Payload (ESP)	This document describes an updated version of the Encapsulating Security Payload (ESP) protocol, which is designed to provide a mix of security services in IPv4 and IPv6. ESP is used to provide confidentiality, data origin authentication, connectionless integrity, an anti-replay service (a form of partial sequence integrity), and limited traffic flow confidentiality. This document obsoletes RFC 2406 (November 1998).
4302	IP Authentication Header	This document describes an updated version of the IP Authentication Header (AH), which is designed to provide authentication services in IPv4 and IPv6. This document obsoletes RFC 2402 (November 1998).
4301	Security Architecture for the Internet Protocol	This document describes an updated version of the Security Architecture for IP, which is designed to provide security services for traffic at the IP layer. This document obsoletes RFC 2401 (November 1998).
4285	Authentication Protocol for Mobile IPv6	IPsec is specified as the means of securing signaling messages between the Mobile Node and Home Agent for Mobile IPv6 (MIPv6). MIPv6 signaling messages that are secured include the Binding Updates and Acknowledgement messages used for managing the bindings between a Mobile Node and its Home Agent. This document proposes an alternate method for securing MIPv6 signaling messages between Mobile Nodes and Home Agents. The alternate method defined here consists of a MIPv6-specific mobility message authentication option that can be added to MIPv6 signaling messages.
4225	Mobile IP Version 6 Route Optimization Security Design Background	This document is an account of the rationale behind the Mobile IPv6 (MIPv6) Route Optimization security design. The purpose of this document is to present the thinking and to preserve the reasoning behind the Mobile IPv6 security design in 2001 - 2002. The document has two target audiences: (1) helping MIPv6 implementors to better understand the design choices in MIPv6 security procedures, and (2) allowing people dealing with mobility or multi-homing to avoid a number of potential security pitfalls in their designs.
4219	Things Multihoming in IPv6 (MULTI6) Developers Should Think About	This document specifies a set of questions that authors should be prepared to answer as part of a solution to multihoming with IPv6. The questions do not assume that multihoming is the only problem of interest, nor do they demand a more general solution.

4218	Threats Relating to IPv6 Multihoming Solutions	This document lists security threats related to IPv6 multihoming. Multihoming can introduce new opportunities to redirect packets to different, unintended IP addresses. The intent is to look at how IPv6 multihoming solutions might make the Internet less secure; we examine threats that are inherent to all IPv6 multihoming solutions rather than study any specific proposed solution. The threats in this document build upon the threats discovered and discussed as part of the Mobile IPv6 work.
4213	Basic Transition Mechanisms for IPv6 Hosts and Routers	This document specifies IPv4 compatibility mechanisms that can be implemented by IPv6 hosts and routers. Two mechanisms are specified, dual stack and configured tunneling. Dual stack implies providing complete implementations of both versions of the Internet Protocol (IPv4 and IPv6), and configured tunneling provides a means to carry IPv6 packets over unmodified IPv4 routing infrastructures.
4205	A Method for Storing IPsec Keying Material in DNS	This document describes a new resource record for the Domain Name System (DNS). This record may be used to store public keys for use in IP security (IPsec) systems. The record also includes provisions for indicating what system should be contacted when an IPsec tunnel is established with the entity in question. This record replaces the functionality of the sub-type #4 of the KEY Resource Record, which has been obsoleted by RFC 3445.
4196	The SEED Cipher Algorithm and Its Use with IPsec	This document describes the use of the SEED block cipher algorithm in the Cipher Block Chaining Mode, with an explicit IV, as a confidentiality mechanism within the context of the IPsec Encapsulating Security Payload (ESP).
4193	Unique Local IPv6 Unicast Addresses	This document defines an IPv6 unicast address format that is globally unique and is intended for local communications, usually inside of a site. These addresses are not expected to be routable on the global Internet.
4109	Algorithms for Internet Key Exchange version 1 (IKEv1)	The required and suggested algorithms in the original Internet Key Exchange version 1 (IKEv1) specification do not reflect the current reality of the IPsec market requirements. The original specification allows weak security and suggests algorithms that are thinly implemented. This document updates RFC 2409, the original specification, and is intended for all IKEv1 implementations deployed today.
4106	The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP)	This memo describes the use of the Advanced Encryption Standard (AES) in Galois/Counter Mode (GCM) as an IPsec Encapsulating Security Payload (ESP) mechanism to provide confidentiality and data origin authentication. This method can be efficiently implemented in hardware for speeds of 10 gigabits per second and above, and is also well-suited to software implementations.
3972	Cryptographically Generated Addresses (CGA)	This document describes a method for binding a public signature key to an IPv6 address in the Secure Neighbor Discovery (SEND) protocol. Cryptographically Generated Addresses (CGA) are IPv6 addresses for which the interface identifier is generated by computing a cryptographic one-way hash function from a public key and auxiliary parameters. The binding between the public key and the address can be verified by re-computing the hash value and by comparing the hash with the interface identifier. Messages sent from an IPv6 address can be protected by attaching the public key and auxiliary parameters and by signing the message with the corresponding private key. The protection works without a certification authority or any security infrastructure.
3971	SEcure Neighbor Discovery (SEND)	IPv6 nodes use the Neighbor Discovery Protocol (NDP) to discover other nodes on the link, to determine their link-layer addresses to find routers, and to maintain reachability information about the paths to active neighbors. If not secured, NDP is vulnerable to various attacks. This document specifies security mechanisms for NDP. Unlike those in the original NDP specifications, these mechanisms do not use IPsec.
3948	UDP Encapsulation of IPsec ESP Packets	This protocol specification defines methods to encapsulate and decapsulate IP Encapsulating Security Payload (ESP) packets inside UDP packets for traversing Network Address Translators. ESP encapsulation, as defined in this document, can be used in both IPv4 and IPv6 scenarios. Whenever negotiated, encapsulation is used with Internet Key Exchange (IKE).

3884	Use of IPsec Transport Mode for Dynamic Routing	IPsec can secure the links of a multihop network to protect communication between trusted components, e.g., for a secure virtual network (VN), overlay, or virtual private network (VPN). Virtual links established by IPsec tunnel mode can conflict with routing and forwarding inside VNs because IP routing depends on references to interfaces and next-hop IP addresses. The IPsec tunnel mode specification is ambiguous on this issue, so even compliant implementations cannot be trusted to avoid conflicts. An alternative to tunnel mode uses non-IPsec IPIP encapsulation together with IPsec transport mode, which we call IIPtran. IPIP encapsulation occurs as a separate initial step, as the result of a forwarding lookup of the VN packet. IPsec transport mode processes the resulting (tunneled) IP packet with an SA determined through a security association database (SAD) match on the tunnel header. IIPtran supports dynamic routing inside the VN without changes to the current IPsec architecture. IIPtran demonstrates how to configure any compliant IPsec implementation to avoid the aforementioned conflicts. IIPtran is also compared to several alternative mechanisms for VN routing and their respective impact on IPsec, routing, policy enforcement, and interactions with the Internet Key Exchange (IKE).
3776	Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents	Mobile IPv6 uses IPsec to protect signaling between the home agent and the mobile node. Mobile IPv6 base document defines the main requirements these nodes must follow. This document discusses these requirements in more depth, illustrates the used packet formats, describes suitable configuration procedures, and shows how implementations can process the packets in the right order.
3756	IPv6 Neighbor Discovery (ND) Trust Models and Threats	The existing IETF standards specify that IPv6 Neighbor Discovery (ND) and Address Autoconfiguration mechanisms may be protected with IPsec Authentication Header (AH). However, the current specifications limit the security solutions to manual keying due to practical problems faced with automatic key management. This document specifies three different trust models and discusses the threats pertinent to IPv6 Neighbor Discovery. The purpose of this discussion is to define the requirements for Securing IPv6 Neighbor Discovery.
3715	IPsec-Network Address Translation (NAT) Compatibility Requirements	This document describes known incompatibilities between Network Address Translation (NAT) and IPsec, and describes the requirements for addressing them. Perhaps the most common use of IPsec is in providing virtual private networking capabilities. One very popular use of Virtual Private Networks (VPNs) is to provide telecommuter access to the corporate Intranet. Today, NATs are widely deployed in home gateways, as well as in other locations likely to be used by telecommuters, such as hotels. The result is that IPsec-NAT incompatibilities have become a major barrier in the deployment of IPsec in one of its principal uses.
3686	Using Advanced Encryption Standard (AES) Counter Mode With IPsec Encapsulating Security Payload (ESP)	This document describes the use of Advanced Encryption Standard (AES) Counter Mode, with an explicit initialization vector, as an IPsec Encapsulating Security Payload (ESP) confidentiality mechanism.
3602	The AES-CBC Cipher Algorithm and Its Use with IPsec	This document describes the use of the Advanced Encryption Standard (AES) Cipher Algorithm in Cipher Block Chaining (CBC) Mode, with an explicit Initialization Vector (IV), as a confidentiality mechanism within the context of the IPsec Encapsulating Security Payload (ESP).
3585	IPsec Configuration Policy Information Model	This document presents an object-oriented information model of IP Security (IPsec) policy designed to facilitate agreement about the content and semantics of IPsec policy, and enable derivations of task-specific representations of IPsec policy such as storage schema, distribution representations, and policy specification languages used to configure IPsec-enabled endpoints. The information model described in this document models the configuration parameters defined by IPSec. The information model also covers the parameters found by the Internet Key Exchange protocol (IKE). Other key exchange protocols could easily be added to the information model by a simple extension. Further extensions can further be added easily due to the object-oriented nature of the model. This information model is based upon the core policy classes as defined in the Policy Core Information Model (PCIM) and in the Policy Core Information Model Extensions (PCIMe).
3566	The AES-XCBC-MAC-96 Algorithm and Its Use with IPsec	A Message Authentication Code (MAC) is a key-dependent one way hash function. One popular way to construct a MAC algorithm is to use a block cipher in conjunction with the Cipher-Block-Chaining (CBC) mode of operation. The classic CBC-MAC algorithm, while secure for messages of a pre-selected fixed length, has been shown to be insecure across messages of varying lengths such as the type found in typical IP datagrams. This memo specifies the use of AES in CBC mode with a set of extensions to overcome this limitation. This new algorithm is named AES-XCBC-MAC-96.
3554	On the Use of Stream Control Transmission Protocol (SCTP) with IPsec	This document describes functional requirements for IPsec (RFC 2401) and Internet Key Exchange (IKE) (RFC 2409) to facilitate their use in securing SCTP (RFC 2960) traffic.

3457	Requirements for IPsec Remote Access Scenarios	IPsec offers much promise as a secure remote access mechanism. However, there are a number of differing remote access scenarios, each having some shared and some unique requirements. A thorough understanding of these requirements is necessary in order to effectively evaluate the suitability of a specific set of mechanisms for any particular remote access scenario. This document enumerates the requirements for a number of common remote access scenarios.
3281	An Internet Attribute Certificate Profile for Authorization	This specification defines a profile for the use of X.509 Attribute Certificates in Internet Protocols. Attribute certificates may be used in a wide range of applications and environments covering a broad spectrum of interoperability goals and a broader spectrum of operational and assurance requirements. The document establishes a common baseline for generic applications requiring broad interoperability as well as limited special purpose requirements. The profile places emphasis on attribute certificate support for Internet electronic mail, IPSec, and WWW security applications.
3226	DNSSEC and IPv6 A6 aware server/ resolver message size requirements	This document mandates support for EDNS0 (Extension Mechanisms for DNS) in DNS entities claiming to support either DNS Security Extensions or A6 records. This requirement is necessary because these new features increase the size of DNS messages. If EDNS0 is not supported fall back to TCP will happen, having a detrimental impact on query latency and DNS server load. This document updates RFC 2535 and RFC 2874, by adding new requirements.
3193	Securing L2TP using IPsec	This document discusses how L2TP (Layer Two Tunneling Protocol) may utilize IPsec to provide for tunnel authentication, privacy protection, integrity checking and replay protection. Both the voluntary and compulsory tunneling cases are discussed.
3053	IPv6 Tunnel Broker	The IPv6 global Internet as of today uses a lot of tunnels over the existing IPv4 infrastructure. Those tunnels are difficult to configure and maintain in a large scale environment. The 6bone has proven that large sites and Internet Service Providers (ISPs) can do it, but this process is too complex for the isolated end user who already has an IPv4 connection and would like to enter the IPv6 world. The motivation for the development of the tunnel broker model is to help early IPv6 adopters to hook up to an existing IPv6 network (e.g., the 6bone) and to get stable, permanent IPv6 addresses and DNS names. The concept of the tunnel broker was first presented at Orlando's IETF in December 1998. Two implementations were demonstrated during the Grenoble IPng and NGtrans interim meeting in February 1999.
3041	Privacy Extensions for Stateless Address Autoconfiguration in IPv6	Nodes use IPv6 stateless address autoconfiguration to generate addresses without the necessity of a Dynamic Host Configuration Protocol (DHCP) server. Addresses are formed by combining network prefixes with an interface identifier. On interfaces that contain embedded IEEE Identifiers, the interface identifier is typically derived from it. On other interface types, the interface identifier is generated through other means, for example, via random number generation. This document describes an extension to IPv6 stateless address autoconfiguration for interfaces whose interface identifier is derived from an IEEE identifier. Use of the extension causes nodes to generate global-scope addresses from interface identifiers that change over time, even in cases where the interface contains an embedded IEEE identifier. Changing the interface identifier (and the global-scope addresses generated from it) over time makes it more difficult for eavesdroppers and other information collectors to identify when different addresses used in different transactions actually correspond to the same node.
2410	The NULL Encryption Algorithm and Its Use With IPsec	This memo defines the NULL encryption algorithm and its use with the IPsec Encapsulating Security Payload (ESP). NULL does nothing to alter plaintext data. In fact, NULL, by itself, does nothing. NULL provides the means for ESP to provide authentication and integrity without confidentiality. Further information on the other components necessary for ESP implementations is provided by [ESP] and [ROAD].

CORPORATE HEADQUARTERS
AND SALES HEADQUARTERS
FOR NORTH AND
SOUTH AMERICA

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER
(888.586.4737)
or 408.745.2000
Fax: 408.745.2100

www.juniper.net

EAST COAST OFFICE

Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886-3146
USA
Phone: 978.589.5800
Fax: 978.589.0800

ASIA PACIFIC REGIONAL
SALES HEADQUARTERS

JJuniper Networks
(Hong Kong) Ltd.
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EUROPE, MIDDLE EAST,
AFRICA
REGIONAL SALES
HEADQUARTERS

Juniper Networks (UK)
Limited
Building 1
Aviator Park
Station Road
Addlestone
Surrey, KT15 2PG, U.K.
Phone: 44.(0).1372.385500
Fax: 44.(0).1372.385501

905009-001 Jan 2008