

IPv6: A Primer for Physical Security Professionals

Internet Protocol Version 4 (IPv4) was first published and standardized by the IETF in 1981 as RFC 791. The Internet and most network communications are based on this protocol, which has led to fundamental changes in peoples' personal and business lives. Today, the technology is colliding with limitations that were not envisioned nearly three decades ago, such as limited IP address space.

There are four billion potential IP addresses, but the practical limit is much lower. In 1981, the world's population was approximately 4.5 billion people and there were 200 Internet sites. In 1992, the number of connected hosts had reached 1 million. In 2002, the number of Internet users worldwide exceeded 613 million, and the number of assigned IPv4 addresses approached 2.5 billion.¹ Current address space is nearly exhausted, and the Internet community must deal with not only the increasing number of connected users and devices, but also the new applications and communications technologies that have introduced new ways in which the Internet is used and accessed.

This paper discusses several issues addressed by the successor protocol, IPv6, and the protocol's benefits and implications for physical security applications and products.

Background

The Internet Assigned Numbers Authority (IANA) manages the unallocated IPv4 unicast address pool. IANA allocates blocks of addresses to Regional Internet Registries (RIRs). In 1990, one-eighth of the total available Internet address space was taken. By 2000, one-half had been allocated. In 2007, that figure had grown to roughly 80 percent.

It is currently projected that the IANA unallocated address pool will be exhausted by October 2010.² The RIR's pools of unallocated addresses are expected to be exhausted by July 2011. Contributing to this depletion is the inefficient Class A, B, and C address structure and the uneven allocation of Class A and B addresses around the world.

Several short-term solutions have been devised, including Classless Interdomain Routing (RFC 1518), new allocation policy (RFC 2050), and the setting aside of private IP addresses (intranets) using Network Address Translation (NAT). The industry's more complete response, however, came in the form of IPv6, which was first described in RFC 1752: The Recommendation for IP Next Generation Protocol, RFC 1752, was published in January 1995; IPv6 was finalized in RFC 2460 in December 1998. Among its recommendations were the following:

- **Expanded scale:** The network must allow identification and addressing of at least one billion leaf networks.
- **Robust service:** The network service and its associated routing and control protocols must be at least as robust as IPv4.

¹ For country details, see <http://www.unicttaskforce.org/perl/documents.pl?id=1314>.

² Up to the moment projections may be viewed at <http://www.potaroo.net/tools/ipv4/>.

- **Smooth transition:** The protocol must have a straightforward transition plan from IPv4.
- **Media independence:** The protocol must work across an internetwork of many different LAN, MAN, and WAN media.
- **Multicast support:** The protocol must support both unicast and multicast packet transmission.
- **Mobility support:** The protocol must support mobile hosts, networks, and internetworks.

On May 21, 2007, the American Registry for Internet Numbers (ARIN) Board issued the following advisory to the Internet community on migration to IPv6:

“ARIN and the other Regional Internet Registries have distributed Internet Protocol version 6, IPv6, alongside IPv4 since 1999. To date, ARIN has issued both protocol versions in tandem and has not advocated one over the other. ARIN has closely monitored trends in demand and distribution for both protocol versions with the understanding that the IPv4 available resource pool would continue to diminish.

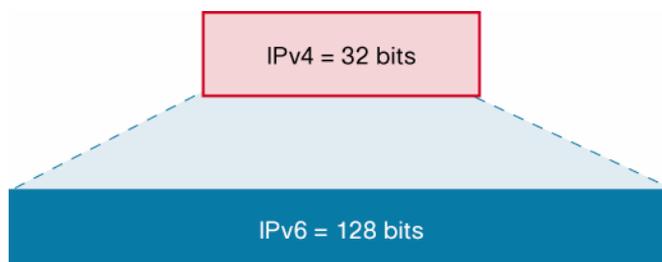
The available IPv4 resource pool has now been reduced to the point that ARIN is compelled to advise the Internet community that migration to IPv6 is necessary for any applications that require ongoing availability from ARIN of contiguous IP number resources.”

Other RIR's have issued similar statements in their own areas.

IPv6 Addressing

Perhaps the most well-known feature of IPv6 is its 128-bit address space, which theoretically provides additional addresses by 29 orders of magnitude (Figure 1). The first 64 bits (4 fields) represent the prefix fields and the second 64 bits represent the interface ID.

Figure 1. 128-Bit Address Space in IPv6



Prefixes

The prefix denotes the type of address and where a unit is connected (registry, ISP, site, subnet). It is derived in a hierarchical fashion. There are several address types: unicast, multicast, and anycast.

Unicast (one-to-one) addresses may be global, link local, unique local, special, or compatible.

Unicast global addresses are conventional publicly routable IP addresses, whose prefix designations are administered by the IANA. The assignable global unicast address space is the address block defined by the prefix 2000::/3.

Unicast link local addresses are designed to be used for addressing on a single link for purposes such as automatic address configuration or neighbor discovery, or when no routers are present and the prefix is defined by an initial block FE80::/10 (or binary 1111 1110 10).

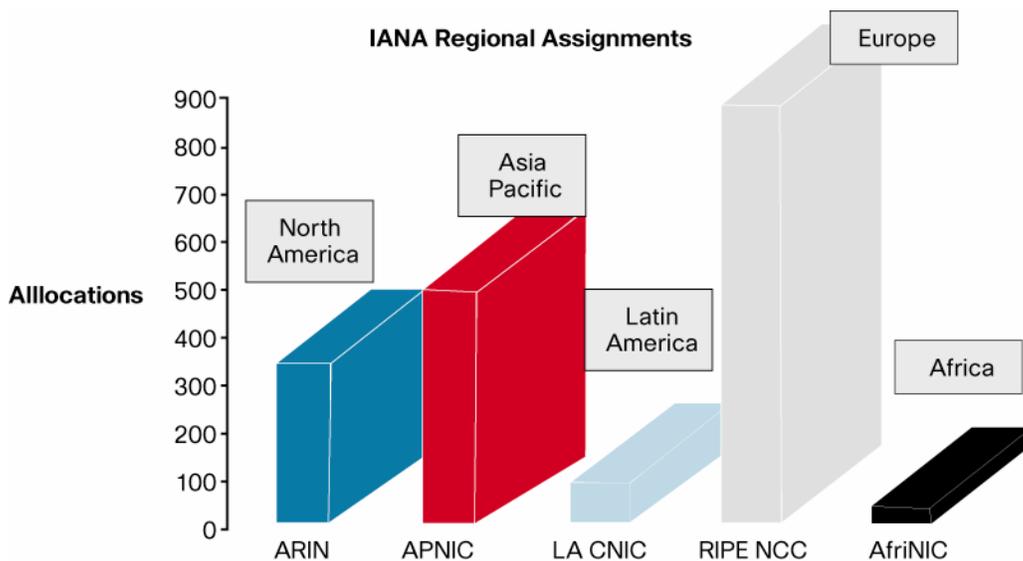
Unicast unique local IPv6 addresses (RFC 4193) are not expected to be routable on the global Internet, but routable inside of a more limited area such as a site or a limited set of sites. These addresses are defined by the prefix FD00::/8 or FC00::/8.

Unicast special addresses cover unspecified addresses (0:0:0:0:0:0:0 or ::) and loopback (0:0:0:0:0:0:0:1 or ::1).

Unicast compatible addresses were intended to aid in the migration from IPv4 to IPv6, but are not normally used.

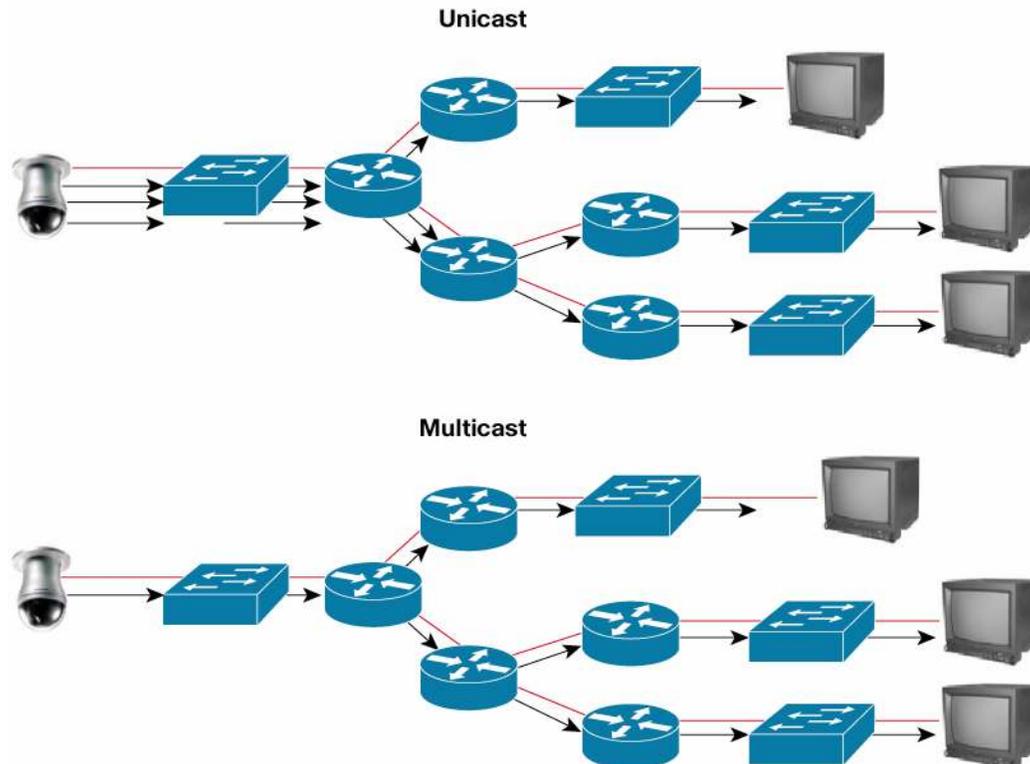
Figure 2 shows the global allocation of IPv6 prefixes as of September 2007. This is a significantly different profile from IPv4, which heavily skewed addresses to North America. The effect of this vastly increased address space is that every device can now have a globally unique IP address and multiple addresses, if necessary. This will simplify configuration in many cases, allow greater flexibility in the deployment of IP devices, and likely expand the population of IP-capable devices (for physical security, this means cameras, door controls, locks, sensors, and other devices). IPv6 also offers entities outside the U.S. much greater flexibility in device addressing.

Figure 2. Global Allocation of IPv6 Prefixes



A **multicast** address (one-to-many communication) features the initial prefix FF00::/16 and is an identifier for a group of interfaces (typically on different nodes). An interface may belong to any number of multicast groups. Multicast is commonly used in CCTV applications to limit the number of network streams, based on the population of users who wish to view a given video stream. These viewers join a multicast group, which is defined by a multicast address and administered through the Internet Group Management Protocol (IGMP), supported in most managed switches and routers. Broadcast is achieved by multicasting to a group consisting of all addresses on the network.

Figure 3 shows unicast and multicast operation where a video stream from a dome camera is viewed at three different locations. The unicast operation requires three separate streams in the network; the multicast operation requires only one stream entering each router and switch.

Figure 3. Comparing Unicast and Multicast Video Streams

An IPv6 anycast address is an address that is assigned to more than one interface (typically belonging to different nodes). A packet sent to an anycast address is routed to the "nearest" interface having that address. Anycast addresses are allocated from the unicast address space, using any of the defined unicast address formats. Thus, anycast addresses are indistinguishable from unicast addresses. Nodes to which the anycast address is assigned must be configured to know that it is an anycast address. An example of anycast would be the transmittal of an alarm from a mobile unit to the nearest responder, where the response mechanism is configured to provide general notification to other interested parties without simultaneous receipt of alarms by all potential responders.

A single interface may be assigned multiple addresses of any type (unicast, multicast, anycast). For more information, see RFC 3513.

Interface IDs and Configuration Options

The interface ID (the second 64 bits in an IPv6 address) is usually assigned using stateless or stateful autoconfiguration. It may also be manually assigned (static configuration).

RFC 2462 describes stateless address autoconfiguration:

Stateless autoconfiguration requires no manual configuration of hosts, minimal (if any) configuration of routers, and no additional servers. The stateless mechanism allows a host to generate its own addresses using a combination of locally available information and information advertised by routers. Routers advertise prefixes that identify the subnet(s) associated with a link, while hosts generate an "interface identifier" that uniquely identifies an interface on a subnet. An address is formed by combining the two. In the absence of routers, a host can only generate link-

local addresses. However, link-local addresses are sufficient for allowing communication among nodes attached to the same link.

The locally available information described above is an interface identifier generated by a host that uniquely identifies an interface on a subnet, based on that unit's 64-bit Extended Unique Identifier (EUI-64). For Ethernet units with a 48-bit MAC address, the EUI-64 address is derived by adding the hexadecimal characters "FFFE" in the middle of the MAC address. An address is formed by combining the two.

Stateless autoconfiguration allows networks to create their own IP addresses and to check for address duplication, but it takes some of the control away from the network administrator and may enable the misuse of network resources. This is particularly important when considering more bandwidth-intensive video appliances, where CCTV cameras could be placed on the network without consideration of proper network configuration for bandwidth control or user access privileges.

RFC 2462 describes stateful autoconfiguration as follows:

In the stateful autoconfiguration model, hosts obtain interface addresses and/or configuration information and parameters from a server. Servers maintain a database that keeps track of which addresses have been assigned to which hosts. The stateful autoconfiguration protocol allows hosts to obtain addresses, other configuration information, or both from a server.

Dynamic Host Control Protocol (DHCPv6) supports stateful autoconfiguration. While IPv6 for Windows XP does not support stateful address configuration or DHCPv6, Windows Vista and the next version of Windows Server include a DHCPv6 client.

Stateless autoconfiguration should make it easier to deploy surveillance-only networks, but may complicate or delay the ultimate convergence into a broader network fabric (and the benefits that accrue from this convergence) until IT has its own policies and procedures worked out.

Address Form

Each address is in the form $x:x:x:x:x:x$ /prefix length

where **x** is a 16-bit (2-byte) field (4 hexadecimal symbols)

and **prefix-length** is a decimal value specifying how many of the

Further, the addressing protocol allows leading zeros to be dropped and consecutive zero fields can be abbreviated by double colons (::):

For example, the address 2001:0DB8:C003:0001:0000:0000:F00D

can be simplified to 2001:DB8:C003:1:0:0:F00D

and then to 2001:DB8:C003:1::F00D

Tunneling

As IPv6 is being phased in, a significant portion of the existing infrastructure will remain until IPv4 is phased out. IPv6 packets can be tunnelled over IPv4 networks by encapsulating the IPv6 data in an IPv4 packet with the use of an IPv4 header, effectively making the IPv6 header and payload invisible to the IPv4 equipment. However, this potentially allows users to exchange IPv6 packets before network administration has fully allowed for it.

Information Security

One of the main objectives in IPv6 is to provide an enhanced level of network security that is embedded into the standard, as opposed to the “add-on” security features that IPv4 required. Security features developed for IPv6 are known as IP Security (or IPsec).

As described in RFC 2401, IPsec “is designed to provide interoperable, high-quality, cryptographically based security for IPv4 and IPv6. IPsec services include access control, connectionless integrity, data origin authentication, protection against replays (a form of partial sequence integrity), confidentiality (encryption), and limited traffic flow confidentiality. These services are provided at the IP layer, offering protection for IP and/or upper-layer protocols.” These objectives are met through the use of two traffic security protocols, the authentication header and the Encapsulating Security Payload (ESP), and through the use of cryptographic key management procedures and protocols.

It is important to note that implementation of the IPsec features embedded in IPv6 does not automatically achieve robust security. While certain threats are mitigated, the opportunity for new threats has increased. Types of threats include:

- Reconnaissance (unauthorized retrieval of information)
- Unauthorized access
- Header manipulation and fragmentation
- Layer 3 and Layer 4 spoofing
- Address Resolution Protocol (ARP) and DHCP attacks
- Broadcast amplification attacks (smurf attacks)
- Routing attacks that disrupt or redirect traffic flows
- Viruses and worms
- Sniffing
- Application-layer attacks
- Rogue devices connected to the network
- Man in the middle attacks
- Flooding

An organization’s IT group must adopt best practices to deal with the issues applicable to their networks, including appropriate use of IPsec, cryptographic procedures, and filtering policies.

For a more in-depth look at IPv6 security issues, please refer to:

http://www.cisco.com/application/pdf/en/us/guest/products/ps6553/c1161/cdccont_0900aecd8057a244.pdf

Mobility

Mobility is one of the most important enhancements of IPv6. All IPv6 networks and nodes are Mobile IP-ready. The following terminology relates to Mobile IPv6, as defined in RFC 3775.

- **Home link:** The link on which a mobile node's home subnet prefix is defined.
- **Home address:** The permanent address of the mobile node, within the mobile node's home link. The home address is a unicast routable address. Standard IP routing

mechanisms will deliver packets destined for a mobile node's home address to its home link.

- **Foreign link:** Any link other than the mobile node's home link.
- **Foreign subnet prefix:** Any IP subnet prefix other than the mobile node's home subnet prefix.
- **Care-of address:** A unicast routable address associated with a mobile node used while the node is on a foreign link; the subnet prefix of this IP address is a foreign subnet prefix. While a mobile node may have multiple care-of addresses at any given time (e.g., with different subnet prefixes), the one registered with the mobile node's home agent for a given home address is called its primary care-of address.
- **Correspondent node:** A remote node with which a mobile node is communicating.
- **Binding:** The association between a mobile node's home address and care-of address
- **Home agent:** A router on a mobile node's home link with which the mobile node has registered its current care-of address. While the mobile node is away from home, the home agent intercepts packets on the home link destined to the mobile node's home address and sends them to the mobile node's registered care-of address.

RFC 3775 summarizes Mobile IPv6 as follows:

"Each mobile node is always identified by its home address, regardless of its current point of attachment to the Internet. While situated away from its home, a mobile node is also associated with a care-of address, which provides information about the mobile node's current location. IPv6 packets addressed to a mobile node's home address are transparently routed to its care-of address. The protocol enables IPv6 nodes to cache the binding of a mobile node's home address with its care-of address, and to then send any packets destined for the mobile node directly to it at this care-of address. To support this operation, Mobile IPv6 defines a new IPv6 protocol and a new destination option. All IPv6 nodes, whether mobile or stationary, can communicate with mobile nodes."

In bidirectional tunneling mode, the home agent remains the packet forwarding agent if the correspondent node is not IPv6-capable or traffic needs to be inspected by the home network. In route optimization mode, the mobile node is required to register its current binding at the correspondent node. After an initial series of packet exchanges between the mobile node and the correspondent node using the home agent to update the link, the home agent does not need to be involved in the communication path any further. IPv6 must be enabled on the correspondent node. Route optimization mode allows the shortest communication path to be achieved, while eliminating congestion at and possible failure of the home agent.

Important differentiators between Mobile IPv6 and IPv4 include:

- No need to deploy special routers as "foreign agents" or to require special support from the local router
- Route optimization mode with secure operation
- Reduced overhead in IPv6 packets
- Robust roaming connections
- Built-in security features using Ipsec

Mobile IPv6 enables direct, “always-on” connectivity, without location constraints or transport dependence. This improves overall system performance, continuity, and response, and is expected to expand the use of mobile technology in safety and security scenarios, such as law enforcement, emergency vehicles, rapid deployment situations, and personal monitoring devices.

Implications for Product Design

IPv4 and IPv6 will need to coexist for many years. Transitional strategies and mechanisms have been developed; there is no danger of immediate mass obsolescence of installed IP cameras, encoders, network video recorders (NVR’s), and other physical security devices.

Some device manufacturers have incorporated IPv6 capabilities into their products and have published configuration guidelines. However, products that are currently in design should either support the IPv6 protocol stack or have the capability for field upgrade, preferably over the network. Manufacturers should make sure that their products’ hardware (memory, processor speed, and performance) can support both IPv6 and IPv4. Host operating systems must be IPv6-enabled, either through a dual IPv4/IPv6 protocol stack or an integrated IPv4/IPv6 stack.

IPv6 is commonly supported on current operating systems. The IPv6 implementation in Windows XP and Windows Server 2003 is a dual-stack architecture. Both Microsoft Windows Vista and the newest Windows Server, termed “Longhorn”, include a redesigned TCP/IP protocol stack with integrated IPv4/IPv6. Sun Microsystems released Solaris 8 with IPv6 support. The USAGI Project is addressing a production-quality IPv6 and IPsec (for both IPv4 and IPv6) protocol stack for the Linux operating system. Many of the long-term IPv6-related patches are integrated into Linux kernel series 2.6.x.

Cisco and IPv6

Cisco has been a leader in defining and implementing the IPv6 architecture within IETF and various working groups. Cisco is also a founding member of the IPv6 forum. Cisco’s core switching, routing, and physical security products support IPv6. For a full listing of these products, visit http://www.cisco.com/en/US/tech/tk872/technologies_white_paper09186a00802219bc.shtml.

Conclusion

Currently, IPv6 has a low number of users relative to IPv4. Hardware support, middleware, management tools, trained technical staff, and applications all must be further developed to enable widespread deployment. Users that are likely to mandate IPv6 support include major corporate entities, the U.S. federal government and other countries. This is in large measure due to the current address shortage with IPv4, but also moves client entities towards the ultimate standard and the benefits it provides..

It is uncertain when physical security hardware will use IPv6 to communicate on a broad basis. Dedicated physical security networks may move more quickly due to the ease of autoconfiguration or the move may be slower, depending on the role of IT in their deployments. Converged physical security networks will be subject to local IT policies.

The benefits that IPv6 will bring to the physical security industry include:

- Expanded number of IP-capable devices, many directly addressable from almost anywhere (cameras, doors, sensors)
- Faster expansion of IP networks outside North America

- Ease of network configuration
- More robust multicast support
- New uses for anycast communication, allowing one of a selected group of addresses to receive messages
- Richer set of integral information security features
- More reliable, more responsive mobile IP networks

For More Information

- **Cisco.com:** [IPv6](#) (White papers, data sheets, further technical information)
- **IPv6 Forum:** <http://www.ipv6forum.com>
- **Cisco Press:** <http://www.ciscopress.com> (Search for “IPv6”)
- **RFCs:** <http://www.ietf.org/rfc.html>



Americas Headquarters
 Cisco Systems, Inc.
 170 West Tasman Drive
 San Jose, CA 95134-1706
 USA
www.cisco.com
 Tel: 408 526-4000
 800 553-NETS (6387)
 Fax: 408 527-0883

Asia Pacific Headquarters
 Cisco Systems (USA) Pte. Ltd.
 168 Robinson Road
 #28-01 Capital Tower
 Singapore 068912
www.cisco.com
 Tel: +65 6317 7777
 Fax: +65 6317 7799

Europe Headquarters
 Cisco Systems International BV
 Haarlerbergpark
 Haarlerbergweg 13-19
 1101 CH Amsterdam
 The Netherlands
www-europe.cisco.com
 Tel: +31 0 800 020 0791
 Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCVP, the Cisco logo, and Welcome to the Human Network are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, IQ Expertise, the IQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0711R)