

# Designing, Constructing, and Operating an IPv6 Network at SC23

IPv6 Networking at SC23

A case study in implementing the IPv6 protocol on a heterogenous network that supports the SC23 conference

Kate Robinson

Lawrence Berkeley National Laboratory (LBNL), Energy Sciences Network (ESnet), katerobinson@es.net

Jason Zurawski

Lawrence Berkeley National Laboratory (LBNL), Energy Sciences Network (ESnet), zurawski@es.net

Tom Costello

Argonne National Laboratory (ANL), tcostello@anl.gov

## Abstract

IPv6 is the current version of IP, the protocol that is used to route traffic across internet connections. This standard was originally developed as a new approach to mitigate concerns about address exhaustion and allow for near infinite scalability. While this protocol has gained significant support in mobile and broadband networks, as well as being the default for networks in emerging economies, it has yet to be fully adopted as a standard deployment model. Complications include legacy devices unable to support the proposed changes, as well as potential challenges that exist between devices that may not be able to fully implement current standards or configuration norms.

The SCinet volunteers who deliver advanced networking to support the SC Conference set an ambitious goal of deploying an IPv6-only network at SC23. While the necessary technology is widely available and understood, the implications of deployment to support more than 15,000 users, each with multiple devices of different operating environments and ages, presents a unique technology and policy challenge. This paper will highlight the effort put into designing, implementing, and operating this innovative IPv6-only environment.

**CCS CONCEPTS** • Network protocols • Network layer protocols • Routing protocols

**Additional Keywords and Phrases:** IPv6, Network Monitoring, Network Architecture

## 1 Introduction

Internet Protocol version 6 (IPv6) [1] is the most recent version of the Internet Protocol (IP) [2] that is used to transit traffic across and between networks. The IPv6 standard was developed by the Internet Engineering Task Force (IETF) as a replacement for IPv4 in the late 1990s, was fully ratified in 1998, and was updated to the current standard in 2017 [37, 38]. As of early 2024, the global user base of IPv6 ranges between 40% and 45% [3, 18].

SCinet is a global collaboration of volunteers who design, build, and operate the SCinet network to support the SC Conference. Since its inception at the SC Conference in 1991, SCinet has provided SC attendees with an innovative network platform necessary to connect, transport, and display research at SC from around the world. SCinet is the ideal testbed to demonstrate technologies and allows volunteers an opportunity to test compatibility across platforms and components [4, 36]

SCinet has long been a proponent of IPv6 and has implemented “dual-stack” networks (e.g., availability of both IPv4 and IPv6 addressing) for the conference as far back as 2003 [5]. To this end, SCinet has dedicated resources for over 20 years toward promoting IPv6 use when possible. For SC23, the SCinet team of volunteers set an ambitious goal of deploying an IPv6-only network. This paper will highlight all the successes and challenges that were presented by designing and attempting to implement IPv6-only in a greenfield network. Additionally, we will cover how far translation technologies like NAT64 [9] and DNS64 [7] have come, and how RFC 8925 (DHCPv4 Option 108) [11] is a critical transition technology between IPv4 and IPv6.

## 2 IPV6 PROTOCOL

IPv6 is the current version of the IP standard, the communications protocol that assists in identifying and locating computers that are connected via communications networks, and is used to route traffic across internet connections.

### 2.1 Background

The IPv6 standard was developed by the IETF as a replacement for IPv4. The IETF standard was ratified in 1998 and updated in 2017. As the Internet grew in popularity in the 1990s, exhaustion of address space was a serious concern due to flawed methods of allocation as well as a general lack of planning for the popularity of the technology. The work on the IPv6 standard began in 1995, offering a far wider addressing scheme (e.g., 128-bit addresses) that would allow  $2^{128}$ , or approximately  $3.4 \times 10^{38}$  total addresses [37, 6]. To further illustrate the scale of these numbers, it is estimated that there are  $10^{19}$  grains of sand on Earth.

### 2.2 State of Deployment

Deployment of IPv6 has steadily been increasing for over 20 years. Starting in approximately 2010, all major operating systems for personal computers and other consumer devices were able to utilize the protocol. As of early 2024, the global user base of IPv6 ranges between 40% and 45%. Adoption is strong across mobile telephone networks, but deployment is non-uniform and varies widely by country. Countries including France, Germany, and India can claim greater than 50% deployment for most traffic, with the United States, Brazil, Japan, and other countries approaching 50%. However, Russia and China have less than 10% adoption, with some countries in Africa and Asia having less than 1% IPv6 adoption [3, 17, 18, 19, 20, 21].

### 2.3 United States Federal Guidance

Starting in 2005, the U.S. government specified that the network backbones of all federal agencies had to be upgraded to support IPv6 operation by June 30, 2008. Additional requirements were instituted in 2010, wherein federal agencies must provide dual-stack IPv4/IPv6 access to external/public services by 2012, and internal clients were to utilize IPv6 by 2014. The U.S. Federal Acquisition Rules (FAR) document 11.002 has required that all procurements comply with specific IPv6 technical capabilities to qualify for procurement for over 10 years [43].

Lastly, in November 2020, the U.S. Office of Management and Budget (OMB) issued the latest federal IPv6-only policy in its memorandum (M-21-07) directing all federal government agencies to complete at least 80% of the transition from IPv4 to IPv6-only by 2025 [14, 15]. Beyond the efforts in the United States, other governments have instituted similar policies [39].

### 3 SCinet

SCinet is a global collaboration of networking experts who provide the fastest and most powerful volunteer-built network in the world for the SC Conference. Designed and created from new technology requirements each year, the SCinet network brings together experts who provide a platform that connects attendees and exhibitors to the world [4].

#### 3.1 Background

Volunteers from academia, government, and industry work together to design and deliver the SCinet infrastructure each year. Industry contributors donate millions of dollars in equipment and services needed to build and support the local and wide area networks. SCinet showcases cutting-edge technologies in network, hardware, protocols, information systems, software, and security — pushing the boundaries of networking technologies.

SCinet has become more than a research network. It provides wired and wireless network connectivity to all conference attendees while in the host city’s convention center. Hundreds of network switches and wireless access points throughout the convention center are deployed in the weeks leading up to the SC Conference. Thousands of attendees and presenters, each bringing numerous devices, expect and depend on SCinet to provide a reliable, high-speed, open network infrastructure.

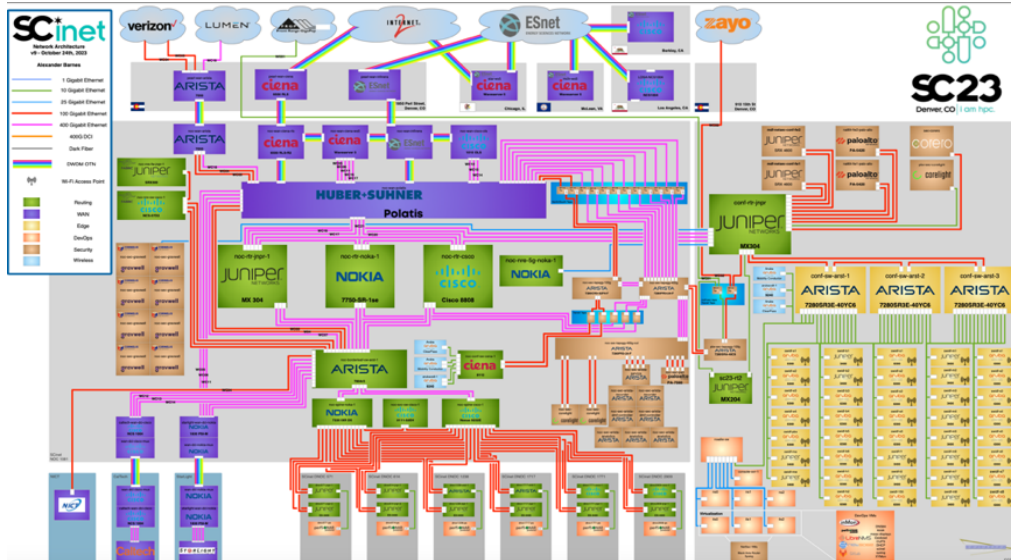


Figure 1: SCinet architecture from SC23, November 2023, Denver CO [45] © The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)

#### 3.2 Network Architecture

The SCinet Network Architecture is designed to address two core use cases:

- Operational network that supports connectivity for approximately 15,000 attendees, volunteers, and staff
- Research-oriented network that supports high-performance demonstrations around the world

Figure 1 shows the SCinet network. This infrastructure relies on optical transport provided by six wide area network (WAN) providers, delivered over four different transportation systems. This heterogeneity of technology is a core strength of SCinet and something the volunteers take pride in yearly: interoperability across platforms helps build understanding of how each will operate in a non-conference scenario.

### 3.3 SC23 Goals & Achievements

At SC23 in Denver, Colorado, SCinet comprised more than 200 volunteers hailing from 9 countries, 31 states, and 113 institutions. The SCinet teams installed nearly 13 miles of fiber, over 400 wireless access points, and delivered a WAN capacity of 6.71 terabits per second (Tbps). All of this was accomplished following the SCinet creed: one year to design, one month to build, one week to operate, and one day to tear down.

To keep with SCinet's theme of innovation, SC23's mission was to promote IPv6 adoption to the fullest extent possible. This meant designing the network to operate primarily as an IPv6-only enterprise network, offering methods of translation to support devices that were unable to natively communicate. There are a number of ways to accomplish this, but due to the varied nature of consumer devices that SCinet can expect to support, a number of mitigations were planned to ensure full coverage [16].

## 4 MOTIVATION & IMPLEMENTATION

Despite the wide availability of IPv6 addressing and numerous sets of instructions and experiences that can be used to assist with deployment, network operators continue leveraging IPv4 and extending its useful life through the use of approaches, such as Network Address Translation (NAT), for many valid reasons. Downtime means lost productivity and revenue, thus a full deployment of a protocol that may not be entirely adopted by all devices that populate the network can be a notable impediment in a critical operations environment. Migrating to IPv6 is not a simple feat, with client devices and operating systems having varying levels of support for IPv6. Commercial equipment manufacturers may not offer support or quality assurance parity with IPv4 in their IPv6 portfolio, which has much to do with the larger market having institutional knowledge and a longstanding installation of IPv4 [22, 23, 24, 25].

Outside factors are spurring the deployment and support for IPv6, however. Cloudflare Radar reported in September 2023 that 37% of connections to cloud-based services are using IPv6 [13]. Amazon Web Services (AWS), a major provider of cloud services, has recently announced that it will begin charging users effective February 1, 2024, \$0.005/hour per IPv4 address (or \$43.80/year per address) [12]. This powerful financial incentive will force many cloud-native applications to reconsider their operating environments.

Starting in January of 2023, SCinet initiated early planning activities to design an IPv6-only network. This exercise involved identifying the technology that would be deployed, how it would be configured and tested, and ways that SCinet could monitor status and ensure basic connectivity for all connected devices.

There are three major ways to impart IPv6 functionality within a network:

- IPv6-only
- Dual-stack (e.g., accommodating both IPv4 and IPv6)
- Use of DHCPv4 "option 108" to enable an RFC6145 customer-side translator (CLAT), frequently described as "IPv6-mostly" [40, 41, 42]

The following sections describe each of these approaches, including the positive and negative features they offer when deployed on a production network.

### 4.1 Dual-stack Operation

The most common way of deploying IPv6 in network infrastructure is through a process called dual-stack: the simultaneous availability of IPv6 for devices that are capable of using the protocol, along with a fallback mechanism for devices that are unable to utilize function in an IPv6-only environment. In practice, there are still many legacy devices, applications, and services that cannot work properly in an IPv6-only environment. This may be due to the IPv6 protocol not being implemented in the firmware or operating systems of the devices, faulty or incomplete IPv6 implementations, or a lack of motivation to support the protocol. It has been observed that the government mandate is a powerful motivator for some operators, but there is no other mechanism to encourage general purpose consumer devices to adopt the new protocol.

Dual-stack networks provide the best user experience, but at the expense of not responding to the IPv4 address shortage at all. In this mode of operation, a device will be assigned an IPv6 address when applicable, but the network can

fall back to a legacy IPv4 mode of addressing and operation as needed. This does little to force application developers and service providers to switch and has the perverse incentive of allowing older devices to function in perpetuity. Dual-stack networks also have the double-sided effect of masking poor, incomplete, or broken IPv6 implementations by allowing failover to legacy IPv4, thus allowing these implementation flaws and oversights to exist unnoticed for long periods of time. Although common, dual-stack deployments are not a robust way to encourage full IPv6 adoption.

## 4.2 IPv6-only

The state of support for IPv6-only differs mostly by device type or, more precisely, by the underlying operating system. It is routine to provide IPv6-only networks, and in doing so, to rely on backup mechanisms to ensure that the older devices have a way to communicate. As described in [32], the intelligence of a network has historically lived on the edges, within the connecting devices. Recent advancements in technology have slowly pushed more functionality into a smarter network core. The ideal place to address translation between protocols will thus rely on network devices and protocols deployed by operators.

Technologies exist to assist with IP address translation, and many are easily deployed in environments like the SCinet network. NAT64 [10] and DNS64 [9] are designed to support connection between an IPv6-only network and legacy IPv4 networks, and they are commonly used as a way to deliver IPv6 connectivity while transparently passing IPv4 traffic to devices. With these devices, the operating system itself is not the preventative technology, as most major operating systems post-2010 have had support for IPv6, but the applications they use may communicate only with IPv4 addresses. Such applications will not work unless an IPv4 address is provisioned to the device.

The use of Customer-side transLATors (CLAT) can also assist, but this method is far from standard in operating system environments. Translation architecture has been standardized in RFC6145 and RFC6146, but methods to activate CLAT may vary based on vendor implementation. Once activated, CLAT will allow applications initiating connectivity via IPv4 to reach their intended destinations via a NAT64 prefix. This translation is seamless to the user and application as long as NAT64 is properly configured.

## 4.3 Use of DHCPv4 Option 108

One of the biggest challenges of deploying an IPv6-only network is the number and variety of devices that could appear on a general-purpose network such as SCinet. The SC conference could experience attendee counts above 15,000, and it is common for attendees to have at least one connected device — if not more — during conference operations. This can mean connections for anywhere from 5,000 to 20,000 hosts during busy periods of the day. While some hosts are capable of operating in IPv6-only mode, others might still have IPv4 dependencies that require IPv4 addresses to operate properly. To incrementally roll out IPv6-only, we wanted to provide a network where it was possible to have both IPv6-only and dual-stack devices in one network. This is where DHCPv4 option 108 [11] comes in.

Most consumer devices do not care which addressing scheme they are given when negotiating access to a network connection. Aside from devices that require a specific static address or route, connections to a wired or wireless connection will utilize DHCP (Dynamic Host Configuration Protocol) [8] to receive essential communication information. This approach has been widely adopted since 1997. Starting in 2020, with the adoption of RFC 8925, a new way that DHCP operates was designed to allow for the specification of IPv6 operation: an "IPv6-only Preferred" option (e.g., DHCPv4 option type 108). Using this configuration at the DHCP server level allows a connecting device to signal its capability to work properly on an IPv6-only network. This is done by requesting this option during the usual DHCP handshake. If the DHCP server is aware that the particular network supports IPv6-only operation, it will include such an option in the response, which will make the client stop the DHCP handshake before an IPv4 address is assigned. Therefore, despite the fact the network is dual-stack capable, IPv4 addressing is provided only to legacy clients not requesting the IPv6-only Preferred DHCP option.

Using these three major strategies for IPv6 deployment, SCinet set out to implement each of them during design, configuration, and operation of the SC23 conference network. The next section will discuss some of the findings of this exercise and ways that other networks can learn from the positive and negative aspects of the experience.

## 5 PRACTICAL RESULTS

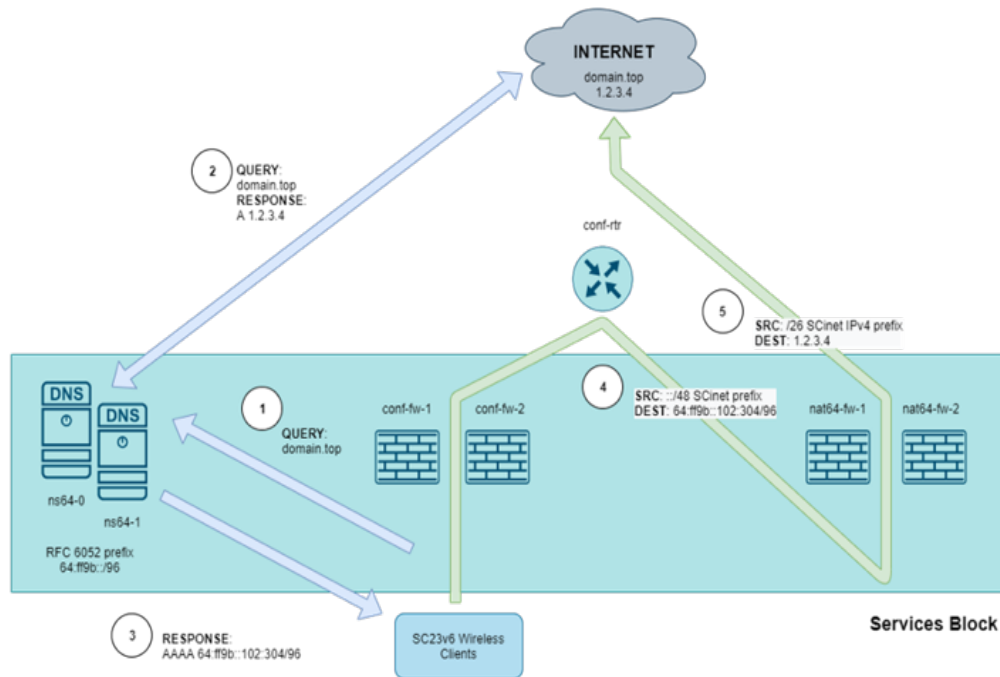


Figure 2: SCinet traffic flow.

The annual SCinet design process takes approximately six to nine months. During this time the volunteers are tasked with ensuring that all of the use cases for connectivity are adequately documented and supported. This process produced the traffic pattern shown in [Figure 2](#). This design was tested using several simulations in a cloud environment, using virtual network hardware and a virtual test client machine, in an effort to understand where bottlenecks for traffic and service may exist as well as points of the network that needed to have additional high-availability components added. It was understood and accepted that even with the designed mitigations to provide IPv6-only services, an IPv4 operating mode (e.g., dual-stack) must still be available as a fallback to support legacy devices and services that could not be supported.

With this block diagram in mind, the following sections will describe the configuration experience and practical results for the major parts of SCinet that were employed to manage IPv6 services and traffic.

### 5.1 Basic Connectivity

It is critical to ensure basic IPv6 connectivity is functional before standing up any other services within a networking environment. In practice, it is not uncommon, either by policy or by mistake, to filter Internet Control Message Protocol (ICMP) [33], a protocol that makes it possible to pass administrative-level messages, along with supporting a number of monitoring and debugging tools. In many cases, operators may configure a network to block ICMP completely, since it has developed a reputation as being a security risk due to the way that devices consume and handle the protocol, and in the past, it has led to certain types of attacks (e.g., the “ping of death”) [47]. Nevertheless, it is critical when initially setting up IPv6 to ensure that components can communicate clearly and effectively.

Preventing the use of ICMPv6 makes finding problems with IPv6 transition mechanisms much more challenging. While IPv4 Address Resolution Protocol (ARP) [34] was found to be working between the main conference router and

NAT64 firewall, as shown in [Figures 1 and 2](#), it was observed that the control plane was not functioning as expected. It was discovered that this was related to the protocol policies, which resulted in faulty operation of the “IPv6 neighbor discovery,” a process that leverages ICMPv6 messages and solicited-node multicast to discover the link layer (MAC) addresses of a neighbor on the same Layer 2 segment, as well as to discover and keep track of local routers [48]. In this scenario, SCinet’s IPv4 prefix was advertised, and provisioned to the hosts as expected, but all IPv6 connectivity was not functional. After debugging, a few BGP filtering changes were required, which resulted in successful reconfiguration.

## 5.2 NAT64 and DNS64

As described in [Section 4.2](#), NAT64 must be used as a transition mechanism for IPv6-only hosts to reach IPv4-only servers, to ensure true end-to-end reachability. Working in concert with NAT64, should an IPv6-only client without a CLAT enabled try to reach a Fully Qualified Domain Name (FQDN) that has only an IPv4 A record, DNS64 will synthesize an AAAA record that corresponds to the NAT64 equivalent IPv6 address. The SCinet volunteers chose to have both IPv4 and IPv6 DNS resolvers on our “SC23v6” wireless network set to our DNS64 servers. This ensured that even dual-stacked clients without CLAT support would prefer IPv6 and NAT64 at all times.

A High Availability (HA) pair of firewalls was configured within SCinet for use solely as the NAT64 translation appliance. By having the NAT64 service exist on a dedicated device and using the well-known NAT64 prefix of 64:ff9b::/96, the routing function that SCinet provides was greatly simplified [35]. This approach was extensively tested in multiple testbeds to understand the impacts of the configuration in practice.

Given that the majority of users are connected to the internet through multiple devices, and each device relies heavily on services that are remotely deployed and locally accessed, DNS performance is critical. Due to the large number of clients that rely on the service during the operation of the conference, SCinet maintains direct control over DNS servers located both on-site and remotely at volunteer institutional sites. This design was implemented in the 2010s to combat network congestion and add reliability, and to date has resulted in high availability of the service and zero observed outages. SCinet deployed dedicated DNS and DNS64 servers, via BIND [44], to address these needs. As a last resort, backup systems were configured to use well-known IPv4 and IPv6 DNS servers located at Google and Cloudflare.

## 5.3 DHCPv4, DHCPv6, and SLAAC

As described in [Section 4.3](#), a fully functional DHCPv4 server and scope of operation are required at all times. This is done to ensure that legacy devices still have a way to connect seamlessly, as well as to support newer devices that can understand new options and behaviors, such as DHCPv4 option 108. The SCinet DHCPv4 configuration was automatically generated using configuration information logged via SCinet’s internal database of services, links, and capabilities.

The SCinet network, now with DHCPv4 option 108 enabled, announced DNS64 server addresses within the DHCPv4 scope of operation. In most situations, networks only choose to announce DNS64 servers on the IPv6-only networks that need them; all other dual-stack networks should use the regular DNS servers so that IPv4 traffic stays native to the IPv4 stack without an unnecessary NAT64 translation. The SC23 network comprises a special circumstance, however: anyone connecting is effectively opting-in to have any legacy IPv4 traffic translated in flight. This approach was progressive and did result in the majority of devices functioning without issue for the duration of the show. Several issues were encountered and mitigated with manual changes to device configuration, defined in [Section 6](#).

The use of Stateless Address Auto-Configuration (SLAAC) was adopted as the primary method to assign IPv6 addresses. This was designed to work in tandem with the SCinet routers and firewalls. The network devices were reconfigured to send first-hop Router Advertisements (RAS) supported Recursive DNS Server (RDNSS) [31], and thus there was no immediate need for DHCPv6.

While DHCPv6 would make network security investigations easier in some cases, the ubiquitous use of IPv6 privacy addresses on most devices makes DHCPv6 less desirable than it was in the past: most hosts will still generate privacy addressing and source traffic from these temporary addresses even in the presence of DHCPv6. Some institutions may have a strong desire to use DHCPv6 within environments where group policies and mobile device management have

been configured to disable privacy addresses, as this may simplify firewall rules and rogue-device management. Although DHCPv6 was not configured at SC23, we may choose to further utilize DHCPv6 in the future.

## 6 LESSONS LEARNED

SCinet made a number of important observations during the deployment, testing, and operation of this network. During operation, although IPv6 connection count eventually outpaced the IPv4 connection count as the conference started (Figure 3), IPv4 throughput remained strong, due to a number of external services that had to utilize IPv4.

The following sections serve as cautionary guidance for those who are contemplating deploying an IPv6-only network in practice in the near term. It is expected that as technology matures in the coming years at the consumer device level, the service level, and the application level, many of these problems will cease to be major blockers to IPv6 use. Several resources were consulted during the build process that provided insights into previous approaches [26, 27, 28, 29, 30].

### 6.1 Virtual Private Networks (VPNs)

A number of SCinet users who rely on Virtual Private Network (VPN) connections to reach home institutions experienced connectivity problems. This problem was not present for all VPN manufacturers and configurations, but was experienced by several users, resulting in their choice to forgo the use of IPv6 connectivity to support these applications. SCinet’s own VPN, deployed as a service for those who were working remotely, had underlying configuration settings altered to support operation. It is common for many institutions to configure VPN clients in a manner that results in all DNS traffic directed towards their internal DNS servers. While this results in VPN clients having no issues with internal DNS queries to their home institution, this was problematic for DNS64 operations.

Specifically, the use of “split tunneling” of IPv4 routes on a host computer, with only an IPv6 IP address, resulted in the largest set of problems. One instance involved the use of remote cloud-based productivity resources that had to be accessed through the home institution’s connection to maintain license requirements.



Figure 3: SCinet connection statistics at the start of SC23.



## 6.2 Non-OEM Operating System Environments

A number of users who had customized operating system environments (e.g., systems that were specifically configured to disable portions of the networking subsystem for security or privacy reasons) failed to connect to the IPv6 network over wireless or wired connections. Given that these environments were specifically designed to function in limited circumstances, SCinet did not spend resources trying to work with their specific requirements. SCinet volunteers did find that some institutionally managed devices had corporate policies in place to either disable or reduce the functionality of the IPv6 network stack, which resulted in no possibility of connection. This issue was particularly difficult to troubleshoot, as the computer would receive an IPv6 address and DHCPv4 option 108 would disable IPv4 networking, but then IPv6 would not function at all, leaving the computer in a partially connected state. These issues were to be reported back to the institution by the user.

## 6.3 Multicast Router Advertisements

SCinet follows the recommendations of RFC 4286 for Multicast Router Discovery. This approach involves routers periodically sending a router advertisement packet that announces availability to the multicast group, and allows a given host to receive router advertisements from all routers, building a list of default routers. Generally, this operation is done frequently, and full lists can propagate in minutes. [Section 5.1](#) outlined the need to support ICMPv6, which is typically the reason why this operation may fail. After still experiencing problems, SCinet made the choice to fall back to unicast advertisements: the list of SCinet routers (and upstream peers) is a manageable quantity to maintain directly, and this approach led to more operational stability.

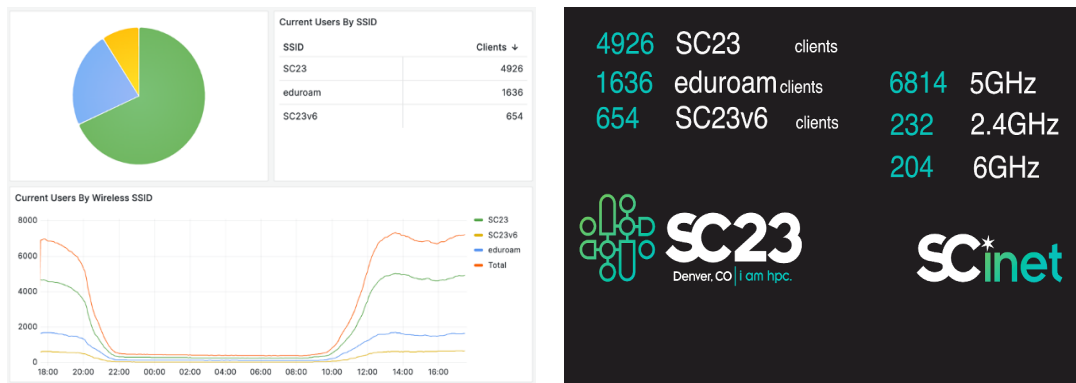
Router advertisements are by default sent via multicast to all clients in a network. SCinet witnessed a large quantity of multicast messages either delayed or missing. It was the case that the router was configured to send out router advertisements at a default minimum of 3 seconds and a maximum of 4 seconds. When clients connected to the wireless network, it was common to see the router solicitation going out, but then the router advertisement wouldn't arrive for almost a minute. To combat this, SCinet configured the router to send router advertisements out via unicast. This made the IPv6 connectivity significantly faster, less than 1 second vs a minute or more.

Additionally, a regularly deployed feature termed "client isolation" was reconsidered. Typically, SCinet would allow this, as it allows clients to interact with each other directly. As the number of connected clients increased, multicast traffic began to inundate the network, which led to performance issues on the wireless side not just related to IPv6. Disabling this facilitated a reduction in traffic and a better overall user experience.

In addition to the use of multicast as a core function of IPv6, deploying IPv6 over wireless entails specific considerations, which can be found in the IETF neighbor discovery considerations document [\[46\]](#).

## 6.4 Wireless SSIDs

SCinet has traditionally advertised two SSIDs for wireless: a general-purpose SSID (e.g., "SC23") that has no default password or restrictions, and "eduroam," which uses the IEEE 802.1X protocol (e.g., WPA2-enterprise) and a system of interconnected RADIUS (Remote Authentication Dial-In User Service) servers to manage access for authentication and authorization back to institutional accounts. On average, 40% of conference attendees will use the eduroam SSIDs by default, since many have access to this resource at their home institution. This can be seen in [Figure 4](#).



**Figure 4: SCinet SSID connection statistics at the end of SC23.**

One of the goals of the year-long IPv6 project was to ensure transparent operation where possible. Both of these SSIDs were designed to work in a way that would function as cleanly as possible for all users (e.g., operating in dual-stack mode and failing back to IPv4 when possible). For future years, the addition of a new network that indicates it is IPv6-only will be added, with the goal of offering IPv6 connectivity only via the DHCPv4 option 108 approach. Unlike the SC23v6 SSID, this new network might not feature any IPv4 internet access and offer a captive portal policy, resulting in a user being notified gracefully to join a different SSID with dual-stack capabilities. This will lead to more accurate IPv6-only client statistics — but at the cost of a more complicated user experience.

## 7 SCINET NEXT STEPS

SCinet begins planning for the next year immediately after the conclusion of the previous year, and the volunteers are actively engaged in ways to learn from the list of findings presented in [Section 6](#). SCinet teams will be tasked with moving toward full adoption of IPv6 within the SCinet network as well as for supporting all conference attendees with reliable and scalable network solutions.

## ACKNOWLEDGMENTS

The authors would like to thank the contributions of the SCinet volunteers who helped to define the vision and assistance in executing this project: Hans Addleman, Adam Bertsch, Nick Buraglio, James Dickerson, Jeremy Duncan, Kalina Dunn, Corey Eichelberger, David Ediger, Jeff Hagley, Britt Huff, Lance Hutchison, Scott Kohlert, Tom Kroeger, Ross Lindsay, Neil McKee, Brenna Meade, Nathaniel Mendoza, Nathan Miller, Hallie Mull, and Greg Veldman.

ESnet is stewarded by Lawrence Berkeley National Laboratory (Berkeley Lab), which is operated by the University of California for the U.S. Department of Energy, Office of Science, under contract DE-AC02-05CH11231.

Argonne National Laboratory's work was supported by the U.S. Department of Energy, Office of Science, under contract DE-AC02-06CH11357.

## REFERENCES

- [1] S. Deering, R. Hinden. 2017. RFC 8200: Internet Protocol, Version 6 (IPv6) Specification. RFC Editor, USA.
- [2] J. Postel. 1981. RFC 791: Internet Protocol. RFC Editor, USA.
- [3] Google. 2024. IPv6 Statistics. Retrieved March 8, 2024 from <https://www.google.com/intl/en/ipv6/statistics.html>
- [4] The International Conference for High Performance Computing, Networking, Storage, and Analysis. 2024. SCinet. Retrieved March 8, 2024 from <https://sc24.supercomputing.org/scinet/>
- [5] SDSC - UC San Diego. 2003. SDSC Networking Experts Contribute to Success of SCinet at SC2003. Retrieved March 8, 2024 from <https://www.sdsc.edu/News%20Items/PR121503b.html>
- [6] S. Deering, R. Hinden. 2006. RFC 4291: IP Version 6 Addressing Architecture. RFC Editor, USA.
- [7] S. Thomson, C. Huitema, V. Ksinant, M. Souissi. 2003. RFC 3596: DNS Extensions to Support IP Version 6. RFC Editor, USA.
- [8] T. Mrugalski, M. Siodelski, B. Volz, A. Yourtchenko, M. Richardson, S. Jiang, T. Lemon, T. Winters. 2018. RFC 8415: Dynamic Host Configuration Protocol for IPv6 (DHCPv6). RFC Editor, USA.

- [9] G. Chen, Z. Cao, C. Xie D. Binet. 2014. RFC: 7269: NAT64 Deployment Options and Experience. RFC Editor, USA.
- [10] M. Bagnulo, P. Matthews, I. van Beijnum. 2011. RFC 6146: Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers. RFC Editor, USA.
- [11] L. Colitti, J. Linkova, M. Richardson, T. Mrugalski. 2020. RFC 8925: IPv6-Only Preferred Option for DHCPv4. RFC Editor, USA.
- [12] Amazon Web Services. 2023. New – AWS Public IPv4 Address Charge + Public IP Insights. Retrieved March 8, 2024 from <https://aws.amazon.com/blogs/aws/new-aws-public-ipv4-address-charge-public-ip-insights/>
- [13] Cloudflare Blog. 2023. Welcome to Birthday Week 2023. Retrieved March 8, 2024 from <https://blog.cloudflare.com/welcome-to-birthday-week-2023/>
- [14] Federal Energy Regulatory Commission. 2024. Internet Protocol Version 6 (IPv6) Policy. Retrieved March 8, 2024 from <https://www.ferc.gov/internet-protocol-version-6-ipv6-policy>
- [15] The White House. 2020. M-21-07 Completing the Transition to Internet Protocol Version 6 (IPv6). Retrieved March 8, 2024 from <https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-07.pdf>
- [16] The International Conference for High Performance Computing, Networking, Storage, and Analysis. 2023. SCinet Brings IPv6 to the Blue Bear and SC23. Retrieved March 8, 2024 from <https://sc23.supercomputing.org/2023/08/scinet-brings-ipv6-to-the-blue-bear-and-sc23/>
- [17] Amogh Dhamdhere, Matthew Luckie, Bradley Huffaker, kc claffy, Ahmed Elmokashfi, and Emile Aben. 2012. Measuring the deployment of IPv6: topology, routing and performance. In Proceedings of the 2012 Internet Measurement Conference (IMC '12). Association for Computing Machinery, New York, NY, USA, 537–550. <https://doi.org/10.1145/2398776.2398832>
- [18] G. Fioccola, P. Volpato, J. Palet Martinez, G. Mishra, C. Xie. 2023. RFC 9386: IPv6 Deployment Status. RFC Editor, USA
- [19] Jakub Czyz, Mark Allman, Jing Zhang, Scott Iekel-Johnson, Eric Osterweil, and Michael Bailey. 2014. Measuring IPv6 adoption. In Proceedings of the 2014 ACM conference on SIGCOMM (SIGCOMM '14). Association for Computing Machinery, New York, NY, USA, 87–98. <https://doi.org/10.1145/2619239.2626295>
- [20] kc claffy. 2011. Tracking IPv6 evolution: data we have and data we need. SIGCOMM Comput. Commun. Rev. 41, 3 (July 2011), 43–48. <https://doi.org/10.1145/2002250.2002258>
- [21] Elliott Karpilovsky, Alexandre Gerber, Dan Pei, Jennifer Rexford, and Aman Shaikh. 2009. Quantifying the Extent of IPv6 Deployment. In Proceedings of the 10th International Conference on Passive and Active Network Measurement (PAM '09). Springer-Verlag, Berlin, Heidelberg, 13–22.
- [22] Sho Fujimura and Masaru Okumura. 2023. IPv6 and Network Security Deployment Use Cases. In Proceedings of the 2023 ACM SIGUCCS Annual Conference (SIGUCCS '23). Association for Computing Machinery, New York, NY, USA, 53–57. <https://doi.org/10.1145/3539811.3579580>
- [23] Ciprian Popoviciu, Eric Levy-Abegnoli, and Patrick Grossetete. 2006. Deploying IPv6 Networks (1st. ed.). Cisco Press.
- [24] Jianping Wu, Jessie Hui Wang, and Jiahai Yang. 2011. CNGI-CERNET2: an IPv6 deployment in China. SIGCOMM Comput. Commun. Rev. 41, 2 (April 2011), 48–52. <https://doi.org/10.1145/1971162.1971170>
- [25] APNIC. 2024. Deploy IPv6. Retrieved March 4, 2024 from <https://www.apnic.net/community/ipv6/deploy-ipv6/>
- [26] Thomas A. Limoncelli and Vinton G. Cerf. 2011. Successful Strategies for IPv6 Rollouts. Really.: Knowing where to begin is half the battle. Queue 9, 3 (March 2011), 20. <https://doi.org/10.1145/1952746.1959015>
- [27] Adeel Ahmed and Salman Asadullah. 2009. Deploying IPv6 in Broadband Access Networks. Wiley Publishing.
- [28] Y. Tian et al., "Traffic Engineering in Partially Deployed Segment Routing Over IPv6 Network With Deep Reinforcement Learning," in IEEE/ACM Transactions on Networking, vol. 28, no. 4, pp. 1573-1586, Aug. 2020, doi: 10.1109/TNET.2020.2987866.
- [29] F. Guo, C. Liu, S. Hao, C. Bao and X. Li, "ADIW: A Solution on General Deployment in IPv6-only WLANs," 2021 IEEE 4th Advanced Information Management, Communicates, Electronic and Automation Control Conference (IMCEC), Chongqing, China, 2021, pp. 915-922, doi: 10.1109/IMCEC51613.2021.9482007.
- [30] Pujol, E., Richter, P., Feldmann, A. (2017). Understanding the Share of IPv6 Traffic in a Dual-Stack ISP. In: Kaafar, M., Uhlig, S., Amann, J. (eds) Passive and Active Measurement. PAM 2017. Lecture Notes in Computer Science(), vol 10176. Springer, Cham. [https://doi.org/10.1007/978-3-319-54328-4\\_1](https://doi.org/10.1007/978-3-319-54328-4_1)
- [31] J. Jeong, S. Park, L. Beloeil, S. Madanapalli. 2017. RFC 8106: IPv6 Router Advertisement Options for DNS Configuration. RFC Editor, USA.
- [32] J. H. Saltzer, D. P. Reed, and D. D. Clark. 1984. End-to-end arguments in system design. ACM Trans. Comput. Syst. 2, 4 (Nov. 1984), 277–288. <https://doi.org/10.1145/357401.357402>
- [33] J. Postel. 1981. RFC 792: Internet Control Message Protocol. RFC Editor, USA.
- [34] D. Plummer. 1982. RFC 826: An Ethernet Address Resolution Protocol or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware. RFC Editor, USA.
- [35] Palo Alto Networks. 2024. NAT64. Retrieved March 8, 2024 from <https://docs.paloaltonetworks.com/pan-os/9-1/pan-os-admin/networking/nat64>
- [36] The International Conference for High Performance Computing, Networking, Storage, and Analysis. 2022. SCinet History. Retrieved March 8, 2024 from <https://sc22.supercomputing.org/scinet/scinet-history/>
- [37] S. Deering, R. Hinden. 1995. RFC 1883: Internet Protocol, Version 6 (IPv6) Specification. RFC Editor, USA.
- [38] S. Deering, R. Hinden. 1998. RFC 2460: Internet Protocol, Version 6 (IPv6) Specification. RFC Editor, USA.
- [39] SIDN Fund. 2022. Governments everywhere make IPv6 mandatory. Retrieved March 8, 2024 from <https://www.sidn.nl/en/news-and-blogs/governments-everywhere-make-ipv6-mandatory>
- [40] E. Nordmark. 2000. RFC 2765: Stateless IP/ICMP Translation Algorithm (SIIT). RFC Editor, USA.
- [41] X. Li, C. Bao, F. Baker. 2011. RFC 6145: IP/ICMP Translation Algorithm. RFC Editor, USA.
- [42] C. Bao, X. Li, F. Baker, T. Anderson, R. Linpro, F. Gont. 2016. RFC 7915: IP/ICMP Translation Algorithm. RFC Editor, USA.
- [43] ACQUISITION.GOV. An official website of the General Services Administration. FAR 11.002 Policy. Retrieved May 15, 2024 from <https://www.acquisition.gov/far/11.002?searchTerms=ipv6>
- [44] Internet Systems Consortium, Inc. BIND9. Retrieved May 15, 2024 from <https://www.isc.org/bind/>
- [45] The International Conference for High Performance Computing, Networking, Storage, and Analysis. 2024. SCinet Architecture. Retrieved May 15, 2024 from <https://sc23.supercomputing.org/scinet/scinet-technology/>
- [46] IETF Datatracker, Selectively Isolating Hosts to Prevent Potential Neighbor Discovery Issues and Simplify First-hops. Retrieved May 2, 2024 <https://datatracker.ietf.org/doc/draft-ietf-v6ops-nd-considerations/>
- [47] Fortinet. 2024. Ping of Death. Retrieved June 10, 2024 from <https://www.fortinet.com/resources/cyberglossary/ping-of-death>
- [48] T. Narten, E. Nordmark, W. Simpson, H. Soliman. 2007. RFC 4861: Neighbor Discovery for IP version 6 (IPv6). RFC Editor, USA.