IPv6 Operations T. Chown
Internet-Draft University of Southampton
Expires: April 30, 2006 October 27, 2005

IPv6 Implications for TCP/UDP Port Scanning
draft-chown-v6ops-port-scanning-implications-02

Status of this Memo

Copyright Notice

Abstract

The 128 bits of IPv6 address space is considerably bigger than the 32 bits of address space in IPv4. In particular, the IPv6 subnets to which hosts attach will by default have 64 bits of host address space. As a result, traditional methods of remote TCP or UDP port scanning to discover open or running services on a host will potentially become far less computationally feasible, due to the larger search space in the subnet. This document discusses that property of IPv6 subnets, and describes related issues for site administrators of IPv6 networks to consider, which may be of

importance when planning site address allocation and management strategies.

Table of Contents

1. Introduction

One of the key differences between IPv4 and IPv6 is the much larger address space for IPv6, which also goes hand-in-hand with much larger subnet sizes. This change has a significant impact on the feasibility of TCP and UDP based port scanning probing, which is something that most of today's IPv4 sites are subjected to routinely around the clock.

The 128 bits of IPv6 [1] address space is considerably bigger than the 32 bits of address space in IPv4. In particular, the IPv6 subnets to which hosts attach will by default have 64 bits of host address space. As a result, traditional methods of remote TCP or UDP port scanning to discover open or running services on a host will potentially become far less computationally feasible, due to the larger search space in the subnet. This document discusses that property of IPv6 subnets, and describes related issues for site administrators of IPv6 networks to consider, which may be of importance when planning site address allocation and management strategies.

This document complements the transition-centric discussion of the issues that can be found in Appendix A of the IPv6 Transition/ Co-existence Security Considerations [5] text, which takes a broad view of security issues for transitioning networks.

It must be remembered that the defense of a network must not rely on the obscurity of the hosts on that network. Such a feature or property is only one measure in a set of measures that may be applied. However, with a growth in usage of IPv6 devices in open networks likely, and security becoming more likely an issue for the end devices, such considerations should be given some weight where to implement appropriate measures is of little cost to the administrator.

Port scanning is quite a prevalent tactic from would-be attackers. The author observes that a typical university firewall may generate many tens of megabytes of log files on a daily basis purely from port scanning activity.

It is also worth noting that worms that spread by scanning target networks for hosts to re-attack have become more common in recent times. Thus a much more sparsely address-populated IPv6 network will have a more innate defense to such forms of worm infection, although there may still be significant scanning traffic generated.

2. Target Address Space for Port Scanning

2.1 IPv4

A typical IPv4 subnet may have 8 bits reserved for host addressing. In such a case, a remote attacker need only probe at most 256 addresses to determine if a particular open service is running on a host in that subnet. At one probe per second, such a scan may take under 5 minutes to complete.

2.2 IPv6

A typical IPv6 subnet will have 64 bits reserved for host addressing. In such a case, a remote attacker needs to probe 2^64 addresses to determine if a particular open service is running on a host in that subnet. At a very conservative one probe per second, such a scan may take some 5 billion years to complete. A more rapid probe will still be limited to (effectively) infinite time for the whole address space.

2.3 Reducing the IPv6 Search Space

The IPv6 host address space through which an attacker may search can be reduced in at least two ways. First, the attacker may rely on the administrator conveniently numbering their hosts from [prefix]::1 upwards.

Second, in the case of statelessly autoconfiguring [1] hosts, the host part of the address will take a well-known format that includes Ethernet vendor prefix and the "fffe" stuffing. For such hosts, if the Ethernet vendor is known, the search space may be reduced to 24 bits (with a one probe per second scan then taking 194 days). Even where the exact vendor is not known, using a set of common vendor prefixes can reduce the search space.

Further reductions may be possible if the attacker knows the target is using 6to4, ISATAP, Teredo, or other techniques that derive low-order bits from IPv4 addresses (though in this case, unless they are using IPv4 NAT, the IPv4 addresses may be probed anyway). For example, the current Microsoft 6to4 implementation uses the address 2002:V4ADDR::V4ADDR while older Linux and FreeBSD implementations default to 2002:V4ADDR::1. This leads to specific knowledge of specific hosts in the network. Given one host in the network is observed as using a given transition technique, it is likely that there are more.

2.4 DNS considerations

Any servers that are DNS listed, e.g. MX mail relays, or web servers, will remain open to probing from the very fact that their IPv6 addresses will be DNS registered. Where a site uses sequential host numbering, publishing just one address may lead to a threat upon the other hosts.

There is a relation between port scanning and DNS zone transfers. In the IPv4 world, this relationship is very weak because the IPv4 space is densely populated and a DNS zone transfer (usually) doesn't help an attacker target a port scan significantly. In the IPv6 world, a zone transfer is much more likely to narrow the number of targeted hosts. This implies restricting zone transfers is (more) important for IPv6, even if it is already good practice to restrict them in the IPv4 world.

2.5 Dual-stack networks

Full advantage of the increased IPv6 address space in terms of reslience to port scanning may not be gained until IPv6-only networks and devices become more commonplace, given that most IPv6 hosts are currently dual stack, also with (more readily scannable) IPv4 connectivity. However, many applications or services (e.g. new peer-to-peer applications) on the (dual stack) hosts may emerge that are only accessible over IPv6, and that thus can only be discovered by IPv6 port scanning.

2.6 Defensive Scanning

The problem faced by the attacker for an IPv6 network is also faced by a site administrator looking for vulnerabilities in their own network's systems. The administrator may have the advantage of being on-link for scanning purposes though, or be able to deduce information about on-link hosts through queries to managed Ethernet switching equipment.

3. Alternatives for Attackers

If IPv6 port-scanning becomes infeasible, attackers will need to find new methods to identify IPv6 addresses for subsequent port scanning. One such method would be the harvesting of IPv6 addresses, either in transit or from recorded logs such as web site logs. Another may be to inspect the Received from: or other header lines in archived email or Usenet news messages.

IPv6-enabled hosts on local subnets may still be discovered through probing the "all hosts" link local multicast address. This implies

that if an attacker can compromise one remote host, they may then learn addresses of the hosts in the same subnet on the remote network.

In IPv6 networks, attackers may also switch to using more aggressive yet subtle methods of attack, e.g. by using worms or viruses that may attach to or attack the new IPv6 applications (e.g. peer-to-peer messaging).

4. Recommendations for Site Administrators

There are some methods that site administrators can apply to make the task for IPv6 port scanning attackers harder. We describe such methods in this section.

The author notes that at his current (university) site, there is no evidence of general port scanning running across subnets. However, there is port-scanning over IPv6 connections to systems whose IPv6 addresses are advertised (DNS servers, MX relays, web servers, etc), which a presumably looking for other open ports on these hosts to probe.

4.1 Use of IPv6 Privacy Addresses

By using the IPv6 Privacy Extensions [3] the hosts in the network may be able to only ever connect to external sites using their (temporary) privacy address. While an attacker may be able to port scan that address if they do so quickly upon observing the address, the threat or risk is reduced. An example implementation of RFC3041 already deployed has privacy addresses active for one day, but such addresses reachable for seven days.

Note that an RFC3041 host may well also have a separate static global IPv6 address by which it can also be reached, and that may be DNS-advertised if an externally reachable service is running from it. However, for client-only systems, RFC3041 offers some level of defence.

4.2 DHCPv6 Configuration

The administrator could configure DHCPv6 so that the first addresses allocated from the pool begin much higher in the address space than [prefix]::1.

DHCPv6 also includes an option to use Privacy Extension [3] addresses, i.e. temporary addresses, as described in Section 12 of the DHCPv6 [4] specification.

4.3 Rolling Server Addresses

Given the huge address space in an IPv6 subnet/link, and the support for IPv6 multiaddressing, whereby a node or interface may have multiple IPv6 valid addresses of which one is preferred for sending, it may be possible to periodically change the advertised addresses that certain long standing services use (where 'short' exchanges to those services are used).

For example, an MX server could be assigned a new primary address on a weekly basis, and old addresses expired monthly. Where MX server IP addresses are detected and cached by spammers, such a defense may prove useful, especially as such IP lists may also be passed between potential attackers for subsequent probing.

5. Security Considerations

There are no specific security considerations in this document outside of the topic of discussion itself.

6. Acknowledgements

Thanks are due to people in the 6NET project for discussion of this topic, including Pekka Savola (CSC/FUNET), Christian Strauf (JOIN Project, University of Muenster) and Martin Dunmore (Lancaster), as well as Tony Finch (Cambridge) and David Malone (TCD, Dublin).

7. Informative References

[1] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[2] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

[3] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.

[4] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[5] Davies, E., "IPv6 Transition/Co-existence Security Considerations", draft-ietf-v6ops-security-overview-03 (work in progress), October 2005.

Author's Address

Tim Chown
University of Southampton
Southampton, Hampshire SO17 1BJ
United Kingdom

Email: tjc@ecs.soton.ac.uk

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.