Operational Security
Internet-Draft
Updates: 792, 4443 (if approved)
Intended status: Standards Track
Expires: April 10, 2013

F. Baker
G. Van de Velde
Cisco Systems
October 7, 2012

Passive IP Addresses
draft-baker-opsec-passive-ip-address-01

Abstract

This note suggests an approach to minimizing the attack surface of the network elements - routers, switches, and middleware - of a network.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

This note suggests an approach to minimizing the attack surface of the network elements - routers, switches, and middleware - of a network.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

1.2. Problem Statement

The problem, at least in its first instance, is a side effect of diagnostics used in the Internet. Tools such as mtr, traceroute, and pingplotter operate by sending streams of packets to a remote address with varying hop limit values in IPv6 [RFC2460] or Time to Live in IPv4 [RFC0791], and receiving ICMP [RFC0792] or ICMPv6 [RFC4443] messages that indicate which interfaces the packet stream traversed in the forward direction. Path MTU [RFC1191] [RFC1981] discovery depends on ICMP/ICMPv6 Packet Too Big. Various ICMP/ICMPv6 "unreachable" messages respond when routing fails, which are intended to trigger applications to try other peer addresses [I-D.ietf-v6ops-happy-eyeballs], and so on. The IP addresses of these responders can be looked up in Reverse DNS [RFC1033][RFC1912] to build a name that indicates the operator, POP, and equipment in question, which is useful in identifying potential problems in the path.

Unfortunately, those addresses can also be used in another way. A motivated adversary can subject routers to TCP RST attacks, load-based DDOS, and other attacks.

An alternate way to reduce this potential attack vector is to not use addresses that are valid beyond the link it is attached towards. A sollution describing considerations around this is given in [draft-ietf-opsec-lla-only] while passive IPv6 addresses will provide network path visibility withought increasing large extend of vulnerability for the devices using down the traffic path.

1.3. Examples of attacks

To pick one example, attacks are being reported in which residential broadband customer's CPE Router is targeted with large volume SNMP GET Requests. The address of the router is not generally known; in IPv4, that may be a result of NAPT use, with the address being harvested from exchanges. It may be obtained from a traceroute to a

server behind the router, or it may be determined by analysis of SMTP envelopes.

Another example is attacks on BGP peering. BGP neighbors often peer between the loopback addresses of neighboring routers, to make the TCP session stable in the presence of link outages, but may peer using interface addresses. If a router is configured to use interface addresses in ICMP/ICMPv6 messages and to peer using those same addresses, the ICMP response exposes information that can be used in a RST attack on routing. It also facilitates any other kind of attack on the router, such as the previously noted SNMP attack (even if the router knows to refuse the message, it consumes CPU). If global addresses are not used - routers use link-local or private addresses - that makes it harder for an attacker to attack the router, but it means that traceroute and other uses are compromised, which is an attack on network forensics. If link-local addresses are used on the interfaces and ICMP is configured to use the loopback address, the router is again exposed to RST attacks.

2. Proposal

The simplest solution seems to be to enable the router to hide in plain sight - to use an address as the source address in ICMP and other messages that is identifiable using Reverse DNS (and therefore, through the name, useful for network diagnostics and communication between operators), but does not facilitate attacks.

The fundamental theory behind this proposal is the Principle of Least Privilege, which in this application is that an entity in the Internet must be able to access only the information and resources that are necessary for its legitimate purpose. In this case, it is reasonable, for various reasons, to enable a random user to identify the path his or her traffic is using or to identify a system in his path when reporting operational issues to an administration. It is not reasonable, or at least not required, that the user be able to specifically interact with any of those systems in the general case.

We propose that the source IP address in an ICMP/ICMPv6 message, or indeed any message sent to a host that has no inherent need to contact the specific system, be useful for Reverse DNS, but not for touching the system. Ideally, it is not routable to the system in the first place; The passive character of this type of addres address comes to play if a packet with this address as destination address on a targetted device and is delivered to the interface, it is summarily dropped. Such an address is referred to as a "passive address", and if it comes from a specific prefix, the prefix is referred to as a "passive prefix". Addresses that are routable and not dropped on

receipt will, for the purposes of this specification, be called "active" IP addresses.

A passive address is sementically non-disguisable from any other type of address and has no requirement for any new type of address-family. Any IPv4 or IPv6 address can become a passive address by a configuration knob when specifying the interface IP address for the Interface or device.

2.1. Making the address useless

Every interface in the Internet has an address, with the exception of IPv4 unnumbered interfaces; even those have addresses that they use, which are the actual address of some other interface on the same system. Increasingly, this is in fact a list of addresses, some of which are IPv4 and some of which are IPv6.

We propose that any address allocated to an interface on infrastructure equipment be given two binary attributes:

UseInICMP: If the address has this attribute TRUE, the corresponding address may be used as the source address of ICMP or ICMPv6 messages and other messages sent to hosts that have no need to actually touch the system. It is otherwise FALSE.

Respond: If the address has this attribute TRUE, the device will process and respond to packets it receives that have this as a destination address; it is an active address. If the attribute is FALSE, the address is a passive address.

If UseInICMP is set TRUE on a Global Unicast Address or Unique Local Address, the address will be available for use in ICMP messages. If "Respond" is set TRUE, traffic sent to the address will be served in the usual way. This describes the present Internet usage. If Respond is set FALSE, traffic sent to the address will be summarily discarded, in effect presenting a "local firewall" blockage related to the address.

An address that has UseInICMP set FALSE will not be used as the source address of an ICMP message. That address will be indiscoverable via ICMP messages. If Respond is TRUE and the address becomes known by other means, such as DNS, traffic sent to the address will be served in the usual way. If Respond is set FALSE, traffic sent to the address will be summarily discarded, in effect presenting a "local firewall" blockage related to the address.

The scenario in view here is that

- an address that is used to access the system would have UseInICMP FALSE (the address is not leaked in such messages) and Respond TRUE (messages sent to the address MAY be operated on by the system).

- an address that is used in ICMP and similar messages would have UseInICMP TRUE (the address MAY be leaked in such messages) and Respond FALSE (messages sent to the address will be dropped on receipt).

2.2. ICMP/ICMPv6 handling

Per [RFC4443], an ICMP Response such as Time Exceeded or Parameter Problem is sent from "the" source address of the interface that detected the issue. This specification narrows that: it SHOULD use one of the source addresses that have the attribute UseInICMP set to TRUE. If no address has that attribute TRUE, it SHOULD NOT send the message.

2.3. Removing the address from routing

If the passive address is taken from any prefix that is not advertised in routing, it will be difficult for an adversary to route to the address, which simplifies the treatment of certain forms of attacks. It is not impossible; a system on the same LAN could send a crafted packet that would arrive anyway. However, especially in inter-domain routing, it is often quite reasonable to believe that addresses exist that need not be advertised to a neighboring network.

One example of such an address, in IPv6, might be a Unique Local IPv6 Unicast Address [RFC4193], or a global unicast address or prefix. There are obvious operational issues in the use of a global prefix; it is easy to accidentally advertise it. In an IPv4 network, the counterpart might be to use an [RFC1918] address, or to use another prefix that one chooses to not advertise.

Link Local addresses SHOULD NOT be used in this context; while they are obviously unroutable except on the local LAN, they are not useful in Reverse DNS.

One problem with this relates to Ingress Filtering [RFC2827]. If the prefix used for passive addresses is not advertised to the neighboring network and the neighboring network is using unicast reverse path filtering, it will filter these responses. For this reason, a network doing this SHOULD advise neighboring networks of passive prefixes for the purpose of inclusion in ingress filters.

2.4. DNS and Reverse DNS

[RFC1912] recommends that "For every IP address, there should be a matching PTR record in the in-addr.arpa domain." In IPv6, there is an important special case, in that link-local addresses are not reflected there, and are used in routing protocols for local communication among IPv6 routers. Like other addresses, passive IP addresses SHOULD have a corresponding Reverse DNS entry; these names help with traceroute and in fault diagnosis. While active addresses may be expected to have A or AAAA records in the administration's own DNS, there is little point for doing so for passive addresses, as they are unresponsive and very likely unreachable.

However, the names given to passive addresses SHOULD NOT be directly similar to the names given active IP addresses. For example, it may be useful to name the interfaces on a certain router so as to identify the router - "ethernet7.card3.router5.lax.example.com". If the correlation to the name of the loopback interface ("router5.lax.example.com") is obviously derivative, the security value is largely forfeit, although it might require human interaction. Such names should differ enough that they are not readily intuited, such as "rack12.lax.example.com".

3. IANA Considerations

This memo asks the IANA for no new parameters.

Note to RFC Editor: This section will have served its purpose if it correctly tells IANA that no new assignments or registries are required, or if those assignments or registries are created during the RFC publication process. From the author's perspective, it may therefore be removed upon publication as an RFC at the RFC Editor's discretion.

4. Security Considerations

This entire note could be described as addressing a set of security considerations. It is not a complete solution to attacks on infrastructure - if loopback addresses, which are used for network management and other purposes are generally known, the infrastructure can still be attacked. However, it is an important reduction of the attack surface. It creates no attack surface that did not already exist.

4.1. Privacy Considerations

This proposal also introduces no new privacy issues.

5. Acknowledgements

This document grew from a conversation among the authors, John Brzozowski, and Thienpondt Hans. Merike Keao's review was very helpful.

6. Change Log

Initial Version: 1 March 2012

2th version: 7 October 2012

7. References

7.1. Normative References

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.

7.2. Informative References

[I-D.ietf-v6ops-happy-eyeballs] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", draft-ietf-v6ops-happy-eyeballs-07 (work in progress), December 2011.

[RFC1033] Lottor, M., "Domain administrators operations guide", RFC 1033, November 1987.

[RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, November 1990.

[RFC1912] Barr, D., "Common DNS Operational and Configuration Errors", RFC 1912, February 1996.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC1981] McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6", RFC 1981, August 1996.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

[draft-ietf-opsec-lla-only]
, M., "Using Only Link-Local Addressing Inside an IPv6 Network", 20012.

Authors' Addresses

Fred Baker
Cisco Systems
Santa Barbara, California 93117
USA

Email: fred@cisco.com

Gunter Van de Velde
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2704 5473
Email: gvandeve@cisco.com