M^{3}AAWG Network Address Translation Best Practices

The Implications of Large Scale NAT for Security Logging

August 2012

Note: Any revisions or updates to this document will be published
at https://www.maawg.org/published-documents

Table of Contents

Executive Summary

This M^{3}AAWG best practices document provides guidance for system operators, network designers, security professionals, and Internet Service Providers about potential issues associated with Large Scale Network Address Translation systems. The document explains how the prevalence of Large Scale NAT (LSN) systems changes the best practices for security logging throughout the Internet and outlines the information that abuse reports need to contain for LSN system operators to be able to act upon them. It also clarifies how constraints on the information that reporters are prepared to include within abuse reports should influence decisions about LSN designs.

Note: This document does not prescribe how LSN systems should be built. Instead, it offers commentary on some possible designs and potential issues that may arise.

Introduction

Limitations on the availability of IPv4 address space and delays in adopting IPv6 have led to increased interest in Network Address Translation (NAT). Multiple devices can share the same public IP address space because NAT equipment is able to relay IP packets between private and public networks. NAT deployment has been growing since it first became available to consumers in the early 1990s.

Most broadband Internet users today rely on NAT for sharing their Internet connection among the various computers and mobile devices in their household or business. The NAT functionality is typically embedded within their DSL router or cable modem, with a single IP address allocated per household.

In recent years, Internet Service Providers (ISPs) have deployed Large Scale NAT (LSN) technology to allow multiple customers to share the same set of IPv4 addresses. LSN systems are already commonplace on the Internet and their widespread use will continue for many years. Originally deployed for mobile Internet access, LSN (also called Carrier Grade NAT or CGN) is now used for all types of Internet connections.

Traceability

In traditional NAT implementations, IP addresses were shared by computers in a single location and administrative environment, such as an office building or home. Accountability for the use of these shared IP addresses remained stable and they were used by just a single group of people.

In contrast, while some Large Scale NAT designs may provide users with stable IP addresses, LSN no longer upholds a one-to-one mapping between a single owner and an IP address. LSN systems generally share IP addresses between devices that are administratively and geographically independent. This has significant implications for the *traceability* process, important in determining who is responsible for specific events on the Internet. Traceability is a fundamental requirement for investigating all types of Internet abuse because mitigation can seldom occur until the person or business responsible for a particular Internet connection has been traced and identified.

Traceability is inversely related to privacy; however, the scope of this document does not include privacy. Instead, this document explains how to ensure that the deployment of LSN does not inadvertently jettison pre-existing levels of traceability.

In RFC 6302[1] (Internet Best Current Practice 162), *Logging Recommendations for Internet-Facing Servers*, the IETF describes how traceability is expected to work in the presence of LSN and sets out requirements for the logging of IP addresses. In this M^{3}AAWG best practices document, we explain the implications of the IETF work and extend the IETF requirements very slightly. The primary motivation is to ensure that we can continue to rely on traceability in dealing with abuse and other operational problems.

As IP packets pass across the boundary between private and public networks, NAT equipment translates IP addresses. It is common for the number of computers on the private network to exceed the number of assigned public IP addresses. To resolve this, NAT equipment maps the source IP address and protocol port number from the private network to a public IP address and port number. Typically, the port numbers are different. The NAT equipment keeps a record of this mapping and is able to direct packets returning from the public Internet to the correct computer on the private network.

Traceability is affected by the presence of NAT because logging of events elsewhere on the Internet can only record the public IP addresses and ports used by the NAT equipment, rather than the private IP addresses

[1] IETF RFC 6302 (BCP 162), *Logging Recommendations for Internet-Facing Servers*, June 2011, http://tools.ietf.org/html/rfc6302.

and ports of individual computers. However, isolating an issue to a single customer account owner and a single location is usually sufficiently accurate for the original problem to be addressed.

Large Scale NAT systems operate in much the same manner as the NAT systems just described. However, instead of the NAT equipment being on the customer premises at the *edge* of the network, in LSN systems, the NAT equipment is *within* the network and is operated by the ISP. Although any particular IP address is still shared between several different computers, in LSN systems those computers are no longer associated with a single account holder and are no longer contained within a single geographic location.

The practical impact is that when LSN is in use, traceability breaks down unless both of the following conditions are satisfied:

- The ISP keeps persistent records of the translation between public and private address space.
- The tracing requests made to the ISP contain the information necessary to distinguish between different users of the IP address.

IETF Security Logging Recommendations

In RFC 6302 (BCP 162), the IETF recommends that Internet-facing servers that log incoming IP addresses from inbound IP traffic also log the following items:

- The source port number
- The transport protocol and destination port number (when the server application is defined to use multiple transports or multiple ports)
- A timestamp, accurate to the second, from a traceable time source

This document explains the implications of these recommendations in more detail. It also adds an LSN design recommendation to allow traceability to occur when the source IP address and port number are not accompanied by details of the destination IP address and port number.

Timestamp Accuracy

A timestamp is "accurate to the second" when its value deviates from the actual time by no more than 1.0 seconds. When a value recorded in a log is created by rounding a more accurate value to the nearest second, it may differ from the actual time by an additional 0.5 seconds. If the recorded value is formed by truncation, the total error may reach 2.0 seconds.

When an *accurate to the second* value is compared with a reference value, it might differ by as many as 2.0 seconds. However, if it is compared with another accurate to the second value, then the two values could differ by up to 4.0 seconds. This must be kept in mind to ensure that over-rapid reuse of source port numbers does not prevent matches from being made.

Security Logging Best Practices

Where logs are generated for security reasons and there is a potential need to trace the source of a connection, it is not possible to predict whether the traffic will traverse an LSN system. As a best practice, always log the following items:

- The source IP address
- The source port number
- The exact time (accurate to the second)

In addition, where there might be doubt, also log the following:

- The protocol used
- The destination port number

Note: The source port number is a significant new addition to many existing logging schemes. The protocol and destination port are also new; however, they are usually fixed values.

Although the allocation of large blocks of IP address space to ISPs is a matter of public record, ISPs do not generally publish the details of their sub-allocations to customers. In order to trace who is responsible for a given event, the relevant ISP will necessarily have to be contacted. As a best practice, the following details should always be provided to identify each specific connection:

- Source IP address
- Source port number
- Exact time (accurate to the second) along with details of the time zone
- The protocol used

Whenever possible, also report the following:

- The destination port number
- The destination IP address

Note: Since the last two items (the destination port and destination IP address) can disclose the identity of the detecting computer, in many circumstances—particularly when spam is detected—the reporter may not wish to reveal them.

Large Scale NAT Deployment Best Practices

M^{3}AAWG best practices for LSN deployments are as follows:

- Maintain logs of the mapping from IP address and port number to customer identity. Store these logs for a sufficient duration—long enough to serve their purpose. The exact period you choose should address not only technical but also regulatory requirements.
- Maintain logs with accurate timestamps that are synchronized with standardized sources. This practice supports the dynamic allocation of IP addresses.
- Permit customers to be disambiguated when incoming reports only attempt to be accurate to the second.
- Permit customers to be identified when the destination IP address and port number have not been revealed.

Large Scale NAT Design

A number of different types of LSN have been proposed. However, this document does not prescribe how LSN systems should be built. Instead, it offers commentary on some possible designs and potential issues.

Many of the earliest LSN systems, deployed by mobile phone companies in Europe, dynamically allocate a small pool of IP addresses to users upon demand with no preset allocations. This is how traditional NAT systems operated. These systems now have hundreds and even thousands of users sharing each individual IP address. These systems can be extremely efficient in their use of IP address space, but they have substantial logging requirements. Performance issues often lead to the disabling of logging, and consequently, traceability suffers or disappears altogether.

Although perhaps better suited for DSL and cable connections, an alternative static partitioning scheme establishes a many-to-one mapping between a set of customers and a single public IP address. If the port number ranges that each customer uses are also statically allocated, then the logging requirements are greatly reduced. This scenario is very similar to the current logging of customer IP address allocation. The main disadvantage of this approach is the relative inefficiency of IP address space usage.

The most efficient use of IP address space is to allow simultaneous use of the *same* IP addresses and port numbers, provided that their destination details differ. Recording both the source and destination IP addresses and port numbers enables NAT equipment to perform the appropriate translation. However, this will only permit disambiguation of customers if the details of the destination accompany any traceability requests, and, as noted above, disclosure of this information can be problematic in some circumstances.

It is specifically ***not*** a best practice to deploy any design where traceability depends on having knowledge of the remote IP address or port number. Additionally, avoiding any necessity to log the identity of the remote site will mean that the system cannot inadvertently record sensitive personal data.

Conclusion

M^{3}AAWG recommends the following best practices when deploying Large Scale NAT:

- Maintain accurately timestamped logs of the mapping from IP address and port number to customer identity.
- Permit customers to be disambiguated when incoming reports only attempt to be accurate to the second.
- Permit customers to be identified when the destination IP address and port number have not been revealed.
- Consider the volume of logs needed to ensure that LSN traceability is not impaired when compared to traceability in a traditional system.

M^{3}AAWG recommends the following best practices when logging and reporting security information to ensure that traceability is not adversely affected in cases where traffic has traversed Large Scale NAT systems:

- When creating log entries for security purposes, systems must record both an IP address and a source port for each entry.
- Timestamps must be accurate to the nearest second.
- Reports to other ISPs must always include the source port number of a connection along with the IP address and timestamp.