Securing the Edge

Abstract

At every edge of the global Internet are the hosts who generate and consume the packet flows which, together, form the overall Internet traffic load. By number, most of these hosts are not secure, leading to dangerous, untraceable traffic flows which can be used to attack other hosts. This memo describes some of the security problems "at the edge" and makes some recommendations for improvement.

1 - Connection Taxonomy

1.1. The Internet is a "network of networks", where the component networks are called Autonomous Systems (AS), each having a unique AS Number (ASN).

1.2. Connections inside an AS are called "Interior" (or sometimes "backbone"), and their security policies are set according to local needs, usually based on business or technical requirements.

1.3. Connections between ASs are called "Border" (or sometimes "peering"), and their security policies are set bilaterally according to the joint needs of the interconnecting parties.

1.4. Connections between an AS and its traffic sources (generators) and traffic sinks (consumers) are called "Edge" (or sometimes "customer"), and their security policies are generally, by long standing tradition, inconsistent.

2 - DDoS Vulnerability

2.1. The most common attack on Internet hosts or infrastructure at the time of this writing is to cause the receipt of too much traffic, consuming all available resources on a victim's host or Internet connection. This is often called a "Denial of Service" (DoS) attack.

2.2. For a DoS attack to succeed, the source or "launch point" must not be trivially detectable. Therefore, successful attacks employ large numbers of weak attackers. An attack launched from ten thousand hosts who each sent ten packets per second would be called a Distributed Denial of Service (DDoS) attack.

2.3. For a DDoS attack to succeed more than once, the launch points must remain anonymous. Therefore, forged IP source addresses are used. From the victim's point of view, a DDoS attack seems to come from everywhere at once, even from many IP addresses that are unallocated or otherwise invalid.

2.4. A successful DDoS can last for minutes or weeks. Because there is no way to determine who launched it, because the process of identifying and correcting each compromised host cannot be practically undertaken as a means of mitigating the attack, and because filtering out "attack flows" invariably has the side effect of damaging valid traffic, every "cure" is nearly as expensive as just "waiting it out."

2.5. While most DDoS attacks are by bad actors against other bad actors, it is quite common to select a high profile victim for no better reason than bragging rights. At the time of this writing there is virtually always an attack in progress somewhere, and in the foreseeable future these attacks will represent a large permanent share of the global Internet's traffic.

3 - DDoS Vector

3.1. The typical vector for DDoS launches is a personal computer (PC) running operating system and application software that purposely trades off security for convenience. These computers are usually poorly managed, such that there are weak passwords or no passwords, known security "holes" that are never patched or closed, and services offered to the global Internet that the owner has no knowledge and no use for.

3.2. From the point of view of almost any single purveyor -- or consumer -- of operating system and application software, convenience will almost always have more perceived value than security. It is only when viewed in the aggregate that the value of security becomes obviously higher than the value of convenience.

3.3. With the advent of high speed "always on" connections, these PCs add up to either an enormous global threat, or a bonanza of freely retargetable resources, depending upon one's point of view.

3.4. Bad actors, in teams or acting alone, exert constant background effort to locate these hosts, probe them for known weaknesses, and subvert them in any way possible. There are software "kits" available that make all of this trivially easy, so no actual technical skill is needed to locate, subvert, and direct an army of thousands of high performance drones.

4 - Remediation

4.1. The foundation of DDoS is anonymity. Even if thousands of hosts are involved, it is both desirable and possible to filter them out, report them to their owners, and repair them, one by one -- if and only if it is possible to learn their identities.

4.2. Source addresses that appear at Border or Interior connections are nonrepudiable by nature, since flows from an alleged source could validly occur in either direction at any Border or Interior connection.

4.3. Source addresses that appear on ingress flows from the edge are generally repudiable, since a typical edge host has no valid reason to use any source address other than one from the pool assigned by the "upstream" or "transit" provider.

4.4. Edge source address repudiation -- the dropping of packets with invalid source addresses upon their ingress across a network edge -- has more immediate beneficial impact than improving PC security. In addition to the difference in complexity and variety, PCs outnumber network edges by at least three orders of magnitude.

5 - Corner Cases

5.1. Multihomed networks who use address space from multiple upstream providers will occasionally emit packets into upstream "A" using source addresses that were assigned by upstream "B". In this case, upstream "A" must be prepared to accept source addresses in address space "B", and vice versa. This is only a slight complication and does not invalidate the approach.

5.2. Networks who have their own address space and ASN, and who speak a dynamic routing protocol such as BGP4, should have their offered routes filtered by their upstream provider(s) where practical in order to prevent bad actors from injecting temporary routes to unassigned or contested address space, from which to launch untraceable attacks.

6 - References

[BCP38]
D. Ferguson, D. Senie, ``Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing,'' BCP 38 and RFC 2827, IETF Best Current Practice, May 2000. 6.1.

7 - Author's Address

Paul Vixie
Internet Software Consortium
950 Charter Street
Redwood City, CA 94063
+1 650 779 7000
vixie@isc.org