

Security Problems and Their Defenses in TCP/IP Protocol Suite

Rajanala. Manasa Valli, Parvathaneni.Tejaswini, A.Rama Krishna

Abstract: The Transmission Control Protocol/Internet Protocol (TCP/IP) is combination of different protocols at various layers. TCP/IP is the basic communication language or protocol of the Internet and private networks either an intranet or an extranet. The TCP/IP suite has many design weaknesses so far as security and privacy are concerned. Some of these are protocol design weaknesses, where as rest are defects in the software that implements the protocols. In this paper, I focused mainly on protocol level issues, rather than implementation flaws. In this paper, we discuss about the security issues related to the some of the protocols in the TCP/IP suite.

Keywords: IP (Internet Protocol), RIP (Routing Information Protocol), TCP (Transmission Control Protocol), UDP (User Datagram Protocol).

I. INTRODUCTION

This paper is an overview of security attacks in the core protocols (IP, UDP and TCP) and other protocols like EGP, BGP, RIP, ICMP and DNS. However, we do not address the exploits in various application protocols. Some of these are protocol design weaknesses per se, whereas the rest are defects in the software that implements the protocols. IP, UDP, TCP and infrastructure protocols were designed at a time when security concerns were almost non-existing and trust was assumed. While this paper summarizes design Weaknesses in the TCP/IP suite from a security point of view, it is important to remember that many implementations have "fixed" these weaknesses, but are not described in RFCs. We assume that the reader is fluent in TCP and IP details. Protocol weaknesses can be divided into those due to i) the design of the protocol itself, and ii) the configuration, deployment and daily operation of the DNS servers. As can be expected, there is a strong interplay between the two. All major OS have made improvements in their implementations of the protocol stack that mitigate or disable many of the attacks described below. Of course, the attack tools also improve. A number of enhancements for TCP/IP have been made that are not yet in common use. Several of them (e.g., DNSSEC and IPV6) involve heavy use of encryption and require more computing power. As computing power in end-user hosts increases, we expect to see these universally deployed.

II. TCP SYN ATTACKS (OR SYN FLOODING)

Manuscript Received on November, 2012

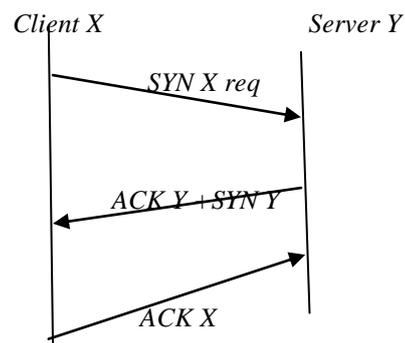
Rajanala. Manasa Valli, Electronics and Computer Engineering, K L University, Guntur, Krishna Dt, Andhra Pradesh, India.

Parvathaneni.Tejaswini, Electronics and Computer Engineering, K L University, Guntur, Krishna Dt, Andhra Pradesh, India.

Prof.A.Rama Krishna, Electronics and Computer Engineering, K L University, Guntur, Krishna Dt, Andhra Pradesh, India.

The TCP uses sequence numbers to ensure data is given to the user in the correct order. The sequence numbers are initially established during the opening phase of a TCP connection in the three way handshake. TCP SYN attacks take advantage of a flaw in how most hosts implement TCP three-way handshake. When host B receives the SYN request from A, it must keep track of the partially opened connection in a "listen queue" for at least 75 seconds and a host can only keep track of a very limited number of connections. A malicious host can exploit the small size of the listen queue by sending multiple SYN requests to a host but never replying to the SYN&ACK the other host sends back. By doing so, the other host's listen queue is quickly filled up, and it will stop accepting new connections, until a partially opened connection in the queue is completed or times out. This availability to affectively remove a host from the network for at least 75 seconds can be used as denial of service attack, or it can be used to implement other attacks like IP spoofing.

Normal TCP handshake



TCP SYN Flood attack

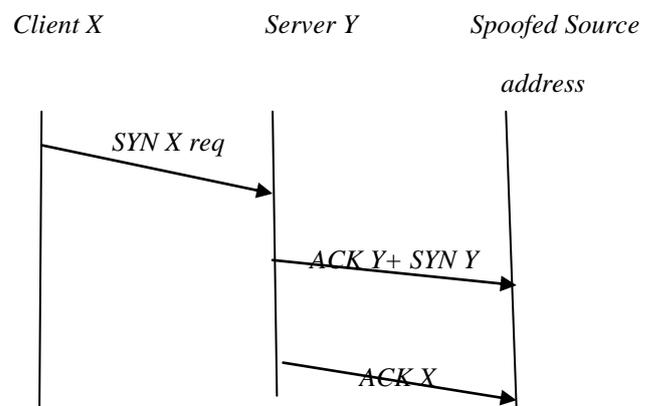


Fig 1. TCP SYN Flood attack

A. Defenses

Both end-host and network-based solutions to the SYN flooding attack have merits. Both types of defense are frequently employed, and they generally do not interfere when used in combination. Because SYN flooding targets end hosts rather than attempting to exhaust the network capacity, it seems logical that all end hosts should implement defenses, and that network-based techniques are an optional second line of defenses that a site can employ.

III. UDP EXPLOITS

UDP is a connectionless protocol belongs to the transport layer. It is a thin protocol on top of IP providing high speed but low functionality. UDP does not guarantee the delivery of datagrams. Messages can be delivered out of order, delayed or even lost. Datagrams may get duplicated without being detected. The UDP protocol is used mostly by application services where squeezing the best performance out of existing IP network is necessary, such as trivial file transfer protocol (TFTP), NFS and DNS. Unfortunately, UDP cannot provide security and privacy of the data flow.

A UDP flood attack sends a large number of UDP packets to random ports. Such ports may be open or closed. If open, an application listening at that port may be open or closed. If closed, the network layer, replies with an ICMP Destination Unreachable Packet. Thus, the victim host will be forced into sending many ICMP packets and wasting computing cycles. If the flooding is large enough, the host will eventually be unreachable by other clients. The attacker will also IP-spoof the UDP packets, both to hide and to ensure that the ICMP return packets do not reach him.

Surprisingly, using legitimate applications or OS services an attacker can generate a storm of packets. On many systems, the standard services known as chargen that listens typically at port 19 and echo that listens typically at port 7 are enabled. Chargen sends an unending stream of characters intended to be used as test data for terminals. The echo service just echoes what it receives. It is intended to be used for testing reachability, identifying routing problems, and so on. An attacker sends a UDP packet to the port 19 with the source address spoofed to a broadcast address and the source port spoofed to 7. The chargen stream is sent to the broadcast address and hence reaching many machines on port 7. Each of these machine will echo back to the victims port 19. This ping-pong action generates a storm of packets.

An attack called fraggle uses packets of UDP echo service in the same fashion as the ICMP echo packets.

A. Defenses

To defend, most host disable many UDP services such as the chargen and echo mentioned above. Because UDP is better suited for streaming applications, there are suggestions to run UDP over SSL or even create a protocol immediately above UDP. Routing protocol is used. Some of these attacks succeed only if the remote host does source address-based authentication; others can be used for more powerful attacks. A number of these attacks described below can also be used to accomplish denial of service by confusing the routing tables on a host or gateway.

IV. ROUTING INFORMATION PROTOCOL (RIP) ATTACKS

The Routing Information Protocol (RIP) is used to propagate routing information on local networks, especially broadcast media. Typically, the information received is unchecked. This allows an intruder to send bogus routing information to a target host, and to each of the gateway along the way, to impersonate a particular host. The most likely attack of this sort would be to claim a route to a particular unused host, rather than to a network; this would cause all packets destined for that host to be sent to the intruder's machine. (Diverting packets for an entire network might be too noticeable; impersonating an idle work-station is comparatively risk-free). Once this is done, protocols that rely on address-based authentication are effectively compromised.

This attack can yield more subtle, and more serious, benefits to the attacker as well. Assume that the attacker claims a route to an active host or workstation instead. All packets for that host will be routed to the intruder's machine for inspection and possible alteration. They are then resent, using IP source address routing, to the intended destination. An outsider may thus capture passwords and other sensitive data. This mode of attack is unique in that it affects outbound calls as well; thus, a user calling out from the targeted host can be tricked into divulging a password. Most of the earlier attacks discussed are used to forge a source address; this one is focused on the destination address. This and are the earliest mentions of routing attacks in the literature. The attacks described here-abusing the routing protocols for eavesdropping and/or packet modification-remain a very serious threat. Indeed, a National Research Council study identified routing attacks as one of the two major threats to the internet. While there are proposals to solve this problem, nothing has been implemented; all of the proposed solutions have their drawbacks. Defense against routing attacks must still be considered a research problem.

Routing attacks have happened frequently by accident. In the most famous case, known as the "AS 7007" incident, an ISP started advertising that it had the best routes to most of the internet. Even after they powered down their router, it took more than four hours for the global routing tables to stabilize.

As suggested here, most subtle routing problems are harder to diagnose. AT&T's dial up internet service knocked off the air for many hours when another ISP started advertising a route to small, internal network. There are many other such incidents as well.

Are malicious routing attacks happening? Yes, they are, and the culprits are a very low life form: the spammers. In some cases, they are hijacking a route, injecting spam and then withdrawing a route. The attack is hard to trace, because by the time someone notices it the source addresses of the email are either non-existent or innocent.

A. Defenses

A RIP attack is somewhat easier to defend against than the source routing attacks, though some defenses are similar. A paranoid gateway-one of that filters packets based on source or destination address-will block any form of host spoofing (including TCP sequence number attacks), since the offending packets can never make it through. But there are other ways to deal with RIP problems. Filtering out packets with bogus source

address would help against many forms of attack. Too few ISPs do it, even though it is a recommended practice. One defense is for RIP to be more skeptical about the routes it accepts. In most environments, there is no good reason to accept new routes to your own local networks. A router that makes this check can easily detect intrusion attempts. Unfortunately, some implementations rely on hearing their knowledge of directly-attached networks. The idea, presumably, is that they can use other networks to route around local outages. While fault-tolerance is in general good idea, the actual utility of this techniques is low in many environments compared with the risks.

It would be useful to be able to authenticate RIP packets; in the absence of inexpensive public key signature schemes, this is difficult for a broadcast protocol. Even if it were done, its utility is limited; a receiver can only authenticate the immediate sender, which in turn may have been deceived by gateways further upstream. This paragraph summarizes the essential difficulty in defending against routing attacks: the problem can originate with non local machines. That is, even if your neighbors are authenticated, they may be deceived rather than dishonest.

More and more sites are starting to protect their routing protocols against direct attacks. The most commonly used mechanism is described in 50, caveats on key selection are given in 59. Another mechanism is the so called TTL security hack: if a packet is supposed to originate on link, send it with a TTL of 255, and verify that on receipt. Any off-link packets will have passed through at least one router which would have decremented the TTL.

Even if the local routers don't implement defense mechanisms, RIP attacks carry another risk: the bogus routing entries are visible over a wide area. Any router (as opposed to host) that receives such data will rebroadcast it; a suspicious administrator almost everywhere on local collection of networks could notice the anomaly. Good log generation would help, but it is hard to distinguish a genuine intrusion from the routing instability that can accompany a gateway crash.

V. THE INTERNET CONTROL MESSAGE PROTOCOL (ICMP)

The Internet Control Message Protocol (ICMP) is used by the IP layer to send one-way informational messages to a host. There is no authentication in ICMP, which leads to attacks using ICMP that can result in a denial of service, or allowing the attacker intercept packets.

Denial of service attacks primarily use either the ICMP "Time exceeded" or "Destination unreachable" messages, which can cause a host to immediately drop a connection. An attacker can forge one of these ICMP messages, and send it to one or both of the communicating hosts to disconnect their connection. ICMP "Redirect" message which is commonly used by gateways when a host has mistakenly assumed the destination is not on the local network. If an attacker forges an ICMP "Redirect" message, it can cause another host to send packets for certain connections through the attacker's host. This attack is similar to a RIP attack, except that ICMP messages only apply to existing connections, and the attacker (the host receiving redirected packets must be on local network.

VI. IP ADDRESS SPOOFING

The IP layer of the typical OS simply trusts that the source address, as it appears in an IP packet is valid. It assumes that the packet it received indeed was sent by the host officially assigned that source address. The IP protocol specifies no method for validating the authenticity of this address. Replacing the true IP address of the sender (or, in rare cases, the destination) with a different address known as IP spoofing. Because the IP layer of the OS normally adds these IP addresses to a data packet, a spoofed must circumvent the IP layer and talk directly to the raw network device.

IP spoofing is used as technique aiding an exploit on the target machine. For example, an attacker can silence a host A from sending further packets to B by sending a spoofed packet announcing a window size of zero to A as through it originated from B.

The attacker's machine cannot simply be assigned the IP address of another host T, using ifconfig or a similar configuration tool. Other hosts, as well as T, will discover (through ARP, for example) that there are two machines with the same IP address.

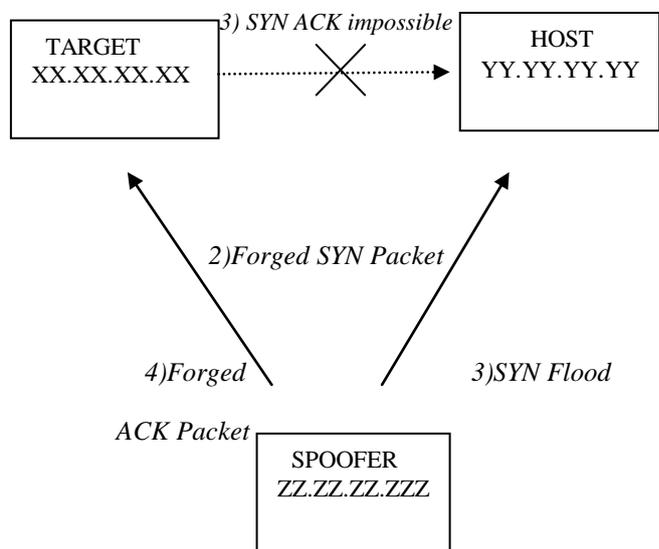


Fig2. IP Address Spoofing

A. Detection of IP Spoofing

We can monitor packets using network-monitoring software. A packet on an external interface that has both its source and destination IP addresses in the local domain is an indication of IP spoofing. Another way to detect IP spoofing is to compare the process accounting logs between systems on your internal network. If the IP spoofing attack has succeeded on one of your systems, you may get a log entry on the victim machine showing a remote access; on the apparent source machine, there will be no corresponding entry for initiating that remote access.

B. Prevention of IP Spoofing

All routers must employ proper IP filtering rules. They should only route packets from source that could legitimately come from the interface the packet arrives on. Most routers now have options to turn off the ability to spoof IP source address by checking the source address of a packet against the routing table to ensure the return path of the packet is

through the interface it was received on.

VII. DOMAIN NAME SYSTEM

The Domain Name System(DNS) provides for a distributed database mapping host names to IP addresses. An intruder who interferes with the proper operation of the DNS can mount a variety of attacks, including denial of service and password collection. There are a number of vulnerabilities. The original DNS specifications did not include security based on the fact that the information that it contains, namely host names and IP addresses, is used as a means of communicating data. As more and more IP based applications developed, the trend for using IP addresses and host names as a basis for allowing or disallowing access (i.e., system based authentication) grew. Unix saw the advent of Berkeley "r" commands (e.g., rlogin,rsh etc.) and their dependencies on host names for authentication. Then many other protocols evolved with similar dependencies, such as NFS, HTTP etc. The existence of widespread use of such protocols as the r-commands put demand on the accuracy of information contained in the DNS. False information within the DNS can lead to unexpected and potentially dangerous exposures. The majority of weaknesses within the DNS fall into the following categories: Cache Poisoning, client flooding, dynamic update vulnerability, information leakage and comprise of the DNS server's authoritative database.

A. Security Threats of the DNS

DNS zone transfers Questioning the legitimacy of a zone transfer request is left out of the protocol. It is also possible to include a zone transfer gratuitously as part of response to a legitimate query.

DNS Cache Poisoning Cache poisoning happens whenever a DNS server does not have the answer to a query within its cache, the DNS server can pass the query onto another DNS server on behalf of the client. If the server passes the query onto another DNS server that has the incorrect information, whether placed there intentionally or unintentionally, then cache poisoning can occur. Malicious cache poisoning is commonly referred to as DNS spoofing.

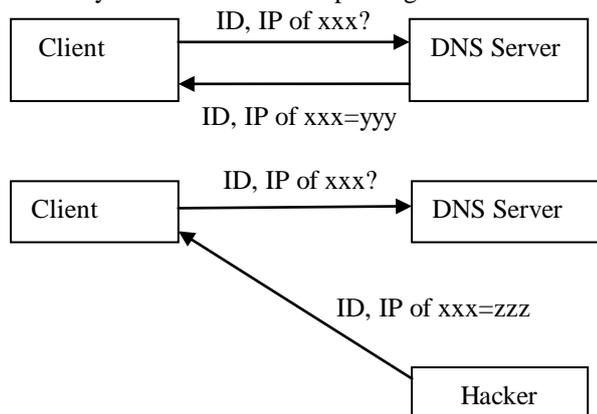


Fig 3. Cache Poisoning

DNS Forgery The DNS answers that a host receives may have come from an attacker who sniffs a query between the victim resolver and the legitimate name servers and responds to it with misleading data faster than the legitimate name server does. The attacked host may in fact be a DNS server. DNS forgery is also called spoofing.

Domain Hijacking A domain is hijacked when an attacker is able to redirect queries to servers that are under the control

of the attacker. This can happen because of cache poisoning, forgery or a domain server has been compromised. DNS hijacking is also known as redirection.

B.Defenses In 1994, the IETF formed a working group to provide security extensions to the DNS protocol in response to the security issues surrounding to the DNS i.e, DNSSEC.

DNSSEC provides authentication and integrity to the DNS. With the exception of information leakage; these extensions address the majority of problems that make such attacks possible. Cache poisoning and client flooding attacks are mitigated with the addition of data origin authentication for RRsets as signatures are computed on the RRsets to provide proof of authenticity. Dynamic update vulnerabilities are mitigated with the addition of transaction and request authentication, providing the necessary assurance to the DNS servers that the update is authentic. Even the threat from compromise of the DNS server's authoritative files is almost eliminated as SIGRR are created using a zone's private key that is kept off-line as to assure key's integrity which in turn protects the zone file from tampering. Keeping a copy of the zone's master file off-line when SIGs are generated takes the assurance one step further. DNSSEC cannot provide protection against threats from information leakage. This is more of an issue of controlling access, which is beyond scope of coverage for DNSSEC. Adequate protection against information leakage is already provided through such things as split DNS configuration.

VIII. CONCLUSION

The TCP/IP suite has many design weaknesses so far as security and privacy are concerned, all perhaps because in the era (1970s) when the development took place network attacks were unknown. The flaws present in many implementations exacerbate the problem. A number of these are due to the infamous buffer overflow which is preventable by better programming practices. However, considerable blame belongs to the many ambiguous RFCs. In this paper, we highlighted the protocol attacks and their defenses.

REFERENCES

- [1]. S.M.Bellovin. A look back at "Security problems in the TCP/IP Protocol suite".
- [2]. S.M.Bellovin, security problems in the TCP/IP protocol suite, Computer Communications Review. <http://www.research.att.com/~smb/papers/ipext.pdf>
- [3]. Larson, M.and Liu, C.," Using BIND: Don't get spoofed again", SunWorld, November 1997 <http://www.sunworld.com/swol-11-1997/swol-11-bind.html>
- [4]. Stevens, Glenn."The Domain Name Service". June 21, 1995.<http://eeunix.ee.usm.maine.edu/guides/dns/dns.html>
- [5]. I.Arce, Attack trends: More bang for the bug: AN account of 2003's attack trends, IEEE Security & Privacy
- [6]. R.Barden, editor.Requirements for Internet hosts-communication layers.RFC 1122, Internet Engineering Task Force, oct.1989.
- [7]. Defense Communication Agency. Defense data subscriber security guide, 1983.

AUTHOR PROFILE



Rajanala. Manasa valli was born in 1990 at Krishna District, Andhra Pradesh. She is pursuing her B.Tech in Koneru Lakshmaiah University. Her interested areas are Networking and Wireless Communications.



Parvathaneni. Tejaswini was born in 1992 at Krishna District, Andhra Pradesh. She is pursuing her B.Tech in Koneru Lakshmaiah University. Her interested areas are Networking, Data Communications and Wireless Communications.



Akella. Rama Krishna was born in 1987, Andhra Pradesh. He is working as a Assistant Professor, Dept. of ECM, K.L. University, Vaddeswaram, A.P, India.

