Practical Interdomain Routing Security

Rick Kuhn, ***US National Institute of Standards and Technology***

Simon Liu, ***US National Library of Medicine***

Hart Rossman, ***SAIC***

Abstract: This article reviews risks and vulnerabilities in interdomain routing, and best practices that can have near-term benefits for routing security. It includes examples of routing failures and common attacks on routers, and coutermeasures to reduce router vulnerabilities.

Keywords: interdomain routing; network attacks; network vulnerabilities; routing security

Routing—the process of determining paths to move packets from source to destination—is fundamental to network operation. Everyone in IT is familiar with routers and firewalls; they're the essential components of every organization's network defense. But what about routing *between* organizations? Interdomain routing vulnerabilities can lead to denial of service or compromise of sensitive information, but many system administrators know little about the risks or what can be done to improve routing security. In this installment of Insecure IT, we review interdomain routing and best practices that can have near-term impacts on security.

Routing within an Organization

The systems that packets pass through need to know where to forward them based on the destination address and information contained in routing tables in each router. The routing table says, for example, that packets with a destination of A can be sent to system H, which can forward the packets to A, possibly through intermediate nodes. Because the Internet changes continuously as systems fail or get replaced, an organization's routing tables must be updated many times a day. The Border Gateway Protocol (BGP) serves this purpose for the global Internet; when BGP fails, portions of the Internet can become unusable for a time period ranging from minutes to hours.

So far, most major BGP incidents have been unintentional. In April 1997, a small Florida Internet service provider (ISP) triggered an Internet-wide instability that caused routers to crash and communications to slow dramatically for more than an hour when it accidentally sent messages indicating that it had the most direct route to large portions of the Internet [1]. The large backbone ISP to which it was connected allowed these faulty messages to be forwarded to other large providers. Consequently, millions of packets were incorrectly sent to the Florida ISP, and errors propagated throughout the Internet when they couldn't be properly routed. More recently, YouTube's address space was "hijacked" in 2008 through a similar process, making it temporarily inaccessible to millions of users [2]. Lessons for security administrators are clear: significant routing

vulnerabilities exist, and outages or other incidents can be triggered either accidentally or with malicious intent.

Many organizations use ISPs that take care of interdomain routing management functions, but many organizations such as universities or companies with large networks run BGP for this task. The collection of routers, computers, and other components within one administrative domain is known as an *autonomous system* (AS); each AS can process packets for thousands of IP addresses within an organization—for example, the IEEE operates AS 13462, which serves a block of 65,792 addresses. The Internet Corporation for Assigned Names and Numbers (ICANN) authorizes Internet registration organizations to assign AS numbers. As of December 2008, the Internet included more than 30,000 of them.

Packets in an Internet transmission, such as an email message, pass from one AS to another until they reach their destination—BGP's task is to maintain lists of efficient paths between them. The paths must be as short as possible and loop-free. BGP routers exchange and store tables of reachability data, which are lists of AS numbers that packets can use to reach a particular destination network—for example, an interdomain router could announce that it can reach addresses in the block 129.6.0.0/16, that is, addresses where the first 16 bits designate decimal address prefix 129.6. Suppose another router announces 129.6.2.0/23. If a packet contains an address of 129.6.3.164, a forwarding router would normally prefer to send the packet to the second router because the /23 address is more specific—we would expect fewer hops for the message to reach its destination. This is one reason why routers are configured to give preference to the most specific addresses.

Normally this practice makes routing more efficient, but when an AS announces overly specific addresses by mistake, routers can become overloaded, as happened with the 1997 incident described earlier. Active BGP entries (that is, the number of reachable address prefixes) are currently approaching 300,000. Each AS uses the reachability information it sends back and forth to other ASs to construct graphs of Internet paths that are loop-free and as short as practical.

Potential Attacks

Although it's not an exhaustive list, the attacks discussed in this section are some of the most common that are likely to be a concern. Because BGP runs on TCP/IP, any TCP/IP attack can be applied to BGP, but here we focus on factors specific to routing security.

Malicious Route Injection

In the absence of security controls, a malicious party can send updates with incorrect routing information. The US National Institute of Standards and Technology's (NIST's) address space is 129.6.0.0/16, for example, so an attacker who announces a more specific route (such as a /24 address in NIST's IP address space) could divert packets that should be sent to NIST. This occurs because other routers would view the /24 as a more direct route to some of the addresses within NIST, so packets would be routed to the attacker's

machine, which could then “blackhole” (drop) them. The attacker could also sniff packets by attacking other routers to manipulate path length and force packets through the attacker’s router. Malicious route injection of this kind is possible because standard BGP has no authentication to guarantee the identity of BGP peers and no authorization mechanism to ensure that a BGP peer has the authority to update routes to particular prefixes.

TCP Resets

Attackers can use the Internet Control Message Protocol (ICMP) to produce session resets; current IETF specifications don’t require routers to check received ICMP messages’ sequence numbers. Such attacks require knowledge of the victim’s IP address and port number, but the nature of BGP requires that they be public. Consequently, attackers can easily send spoofed ICMP error messages, which cause TCP session reset (hard errors) or signal performance/throughput degradation (soft errors). TCP resets drop BGP peering sessions, forcing routers to rebuild routing tables. Router vendors are addressing this issue, but fixes aren’t universally implemented yet.

Unallocated Route Injection

One variety of malicious route injection involves the transmission of routes to unallocated prefixes (that is, they aren’t yet assigned to any organization): no one should be using these addresses, so no traffic should be routed to them. A related attack involves using routes on subnets that are allocated but not used by a target organization.

Resource Exhaustion

Because BGP is implemented on TCP/IP, SYN flooding and other attacks on TCP can affect BGP processing. Moreover, in addition to the storage that the underlying TCP/IP processing requires, routers use a large amount of storage for path prefixes as well. These resources can be exhausted if a router receives updates too rapidly, or if the router has too many path prefixes to store due to malicious prefix announcements. Excessive route updates can also occur, due to compromise or a trusted neighbor’s technical issues.

Countermeasures

Although researchers have proposed various protocols for comprehensive security in interdomain routing, none has gained acceptance. However, several immediately practical options are available, including the following.

Generalized TTL Security Mechanism (GTSM)

Often referred to as the “time-to-live (TTL) hack,” this procedure sets the TTL (hop count) to 255 on outgoing packets and forces neighboring routers to ignore packets with a TTL of less than 254 (to allow for some variations in router implementations), thus ensuring that incoming packets are one hop away. GTSM isn’t universally implemented, but cooperating organizations can gain some security by adopting it.

Filtering

System administrators can specify filtering of both incoming prefixes (ingress filtering)

and outgoing prefixes (egress filtering) by using a syntax similar to that for firewalls. Specifically, they set filters to accept only certain blocks of address prefixes, reject unallocated prefixes (using continuously updated lists), and reject obviously invalid prefixes, such as those used in private networks (for example, 192.168.0.0/16). Normally, neighboring routers should have matching prefix filters—that is, an AS's egress filters should match the ingress filters of the peers with which it communicates. Checking that the TCP sequence number is within the range of packets sent but not yet acknowledged can also help resist malicious route injection.

Digital Signatures

Commercial routers offer MD5 digital signatures, which can help ensure that received packets only come from authorized routers. A disadvantage is that every pair of peers must share must a secret key that must be updated periodically to prevent brute-force cracking by an attacker who has accumulated a large volume of messages.

Access Control Lists (ACLs)

Although relatively basic, system administrators can use ACLs to limit connections to the router to only authorized neighbor routers.

Many research projects are working on ideas to improve interdomain routing security, and the IETF has both mature and developing specifications for routing security [3]. But as with many aspects of IT, compatibility with the installed base limits adoption. In addition, some specifications call for significant cryptographic processing, which can impact performance. Consequently, many newer activities have focused on things that individual organizations can do to improve their own security while still contributing to incremental adoption of better Internet-wide mechanisms. These include the US Department of Homeland Security's Secure Protocols for the Routing Infrastructure program (www.cyber.st.dhs.gov/spri.html) and many other active working groups within the IETF.

The network vulnerability landscape changes rapidly, and some of the most common attacks today were relatively unknown just a few years ago. As administrators shore up defenses in one area, attackers look for other means of entry. Interdomain routing has little security today, and economic pressures make it difficult to adopt enhanced versions of BGP, so defenders must concentrate on practical tools at hand. Interdomain routing vulnerabilities are a target of opportunity still not exploited widely today, so adopting practical defenses now might keep administrators one step ahead.

Disclaimer

We identify certain software products in this document, but such identification doesn't imply recommendation by the US National Institute for Standards and Technology or other agencies of the US government, nor does it imply that the products identified are necessarily the best available for the purpose.

Reference

1. W. Aiello, J. Ionnidis, P. McDaniel, "Origin Authentication in Interdomain Routing", Proc., 10th ACM Conf. on Computer and Comm. Security, Washington, D.C., pp. 165-178.

2. RIPE NCC. YouTube hijacking: A RIPE NCC RIS case study. http://www.ripe.net/news/ study-youtube-hijacking.html, 2008.
3. D. Montgomery and S. Murphy, “Toward Secure Routing Infrastructures,” *IEEE Security & Privacy*, vol. 4, no. 5, 2006, pp. 84–87.