PACKET®

CISCO SYSTEMS USERS MAGAZINE

FIRST QUARTER 2004

Self-Defending Networks 24

Protecting the Network from Human Nature

57 Extending SANs

52 High Availability Campus Networks

39 Securing IP Voice

69 Troubleshooting MPLS VPNs

Cisco Systems

cisco.com/packet

Proactive Protection

New techniques and best practices help service providers counter increase in cyber attacks.

BY DAVID BARRY

When the SQL Slammer worm struck on January 25, 2003, it infected 150,000 to 200,000 servers worldwide within hours of its first appearance, according to Network Associates' AntiVirus Emergency Response Team (AVERT). Customers of the Canadian Imperial Bank of Commerce in Toronto, Canada, and the Bank of America in the US were unable to withdraw funds using automated teller machines during part of the day following the attack and into the next day. And in South Korea, most of Korea Telecom Freetel's and SK Telecom's fixed-line and mobile Internet services failed, stranding millions of the country's Internet users.

Cyber attacks show no signs of abating. Carnegie Mellon University's CERT Coordination Center for reporting Internet security problems states that through the end of September 2003, users and ISPs reported 114,855 security breaches—32,761 more than during all of 2002. These reports include all types of security policy violations, from distributed denial-of-service (DDoS) to hacker attacks.

As the level and sophistication of these attacks increase, they are no longer focused solely on disrupting customer networks and data—many now target the Internet infrastructure itself. As more service providers offer converged services, including voice, the importance of network defense escalates. This migration of the core network is critical to their financial success and competitive differentiation but cannot come at the cost of reliability.

Service providers are quickly realizing that failure to respond to this new threat can be costly in terms of lost revenue caused by network downtime and disruption to customer service. At one time content to provide only "transport," many now realize they must be more proactive in combating these attacks.

Sprint, for example, deploys new tools in its network to better defend itself against attacks. Sprint is specifically focusing on distributed DoS mitigation and intrusion-detection products that it plans to deploy in its backbone within the next year, says John Pardun, senior product manager of network-based IP VPN and security services at the carrier.

Sprint plans to offer customers an "additional level of monitoring and mitigation" as an add-on service that it will charge for, Pardun says. Both MCI and AT&T also say they will charge customers for their planned distributed DoS mitigation and reaction services.

Security as a Process, Not a Point Product

Because service provider infrastructures are so large and the potential for attacks exists at multiple points in the network, Cisco recommends a multifaceted defense approach.

Barry Greene, a consulting engineer at Cisco, is a leading security thinker within the service provider community and a major force in the founding of NSP-SEC (puck.nether.net/ mailman/listinfo/nsp-security), an online forum for security engineers at service providers to exchange information and coordinate response activities during malicious attacks.

"Service provider security is a life-cycle process that must encompass continual preparation, communications with colleagues, and reassessments to keep pace with the changing nature of security attacks," says Greene. "Service providers must continually prepare, test, and deploy new approaches and continually reassess what worked and what didn't after each attack."

Cisco, in collaboration with its service provider customers, has developed a six-stage security framework for service providers that combines industry best practices with specialized routing techniques and tools:

- Preparation
- Identification
- Classification
- Traceback
- Reaction
- Postmortem

Top Priority: Preparation

Service providers have little chance of successful defense against a malicious worm or virus attack unless they have laid the groundwork before an attack. "Trying to respond to an attack without preparation is like trying to line up a militia to combat an invasion when the invaders are already crawling up the beach," says Roland Dobbins, a network engineer in the Cisco IT Internet Services Group.

Aggressive preparedness by Cisco IT enabled the company's entire network to escape unscathed during the Slammer attack that devastated other companies. According to Dobbins, "We did not lose a single packet on the production network due to the Slammer

SECURING POP AND CORE INFRASTRUCTURE

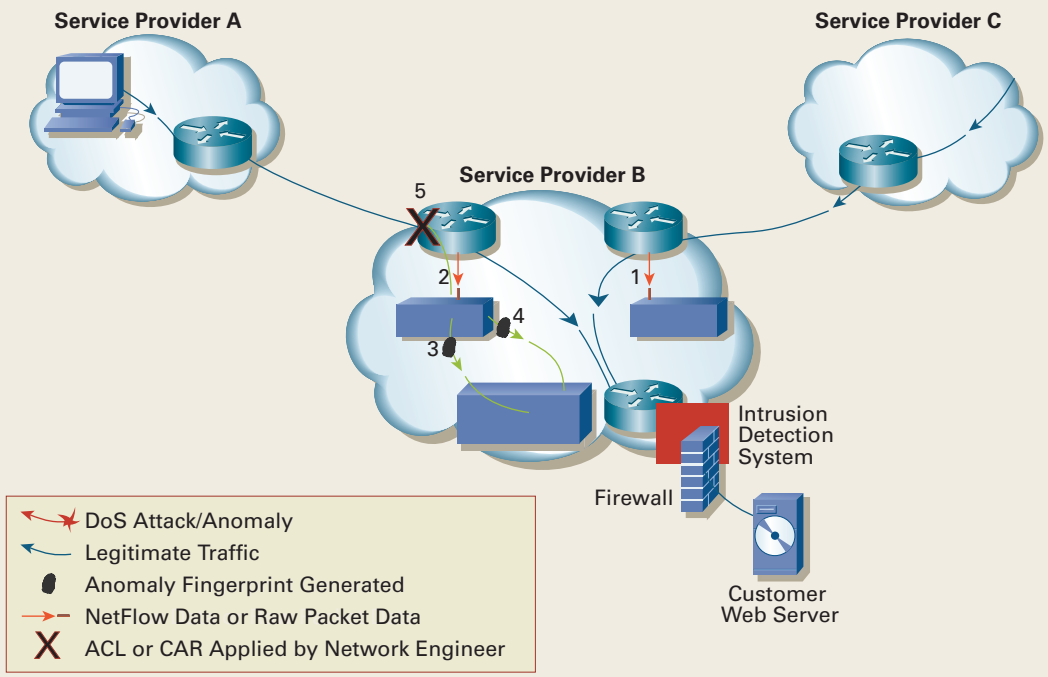

FIGURE 1: Anomaly-detection systems, such as those from Arbor Networks, use raw packet data from NetFlow to detect and characterize undesirable traffic, such as DoS attacks. In addition to detecting suspicious traffic, they can also recommend where filters should be applied to limit attacks.

worm." As a major content provider, e-commerce site, and ISP for 32,000 people worldwide (and with 93 percent of its revenue booked online), Cisco takes security of its worldwide network seriously (see sidebar, "Cisco Security Best Practices," page 67).

"Our motto is that 'proactive work buys us time to be reactive,' " says Dobbins. "We owe our success during Slammer to the fact that we had the systems, communications path, escalation path, and processes already in place. Our rules state that if the attack is rated a P1—the highest level—management must be updated every 15 minutes, with additional network engineers brought on continually until the attack is under control."

A critical part of preparation is the creation of a security or incident response team, which must be given the necessary resources and authority to combat attacks, and a team member must be available 24 x 7. As part of a response plan, team members should also establish relations with security team members at other peering partners and IXCs so that security engineers can be reached quickly at any hour. This enables teams to rapidly share information in the early stages of an attack and can help to determine the real source of the attack and possible remediation.

Relationships with key infrastructure vendors should also be established. Cisco maintains the Product Security Incident Response Team (see "Security Advocates," page 36) as a key resource for both reporting security incidents and for obtaining information about newly discovered vulnerabilities (security-alert@cisco.com for emergencies or psirt@cisco.com for nonemergencies).

One of the problems with communication among top engineers at various service providers is that often these engineers are not easily reachable—and they might not answer phone calls during an emergency. To help ensure that top security engineers at service providers can reach one another during an attack, a dedicated Session Initiation Protocol (SIP)-based network operations center (NOC) hotline system was created. This voice over IP (VoIP) system, managed by the nonprofit research institute Packet Clearing House (pch.net) and sponsored by Cisco, provides a direct connection to the engineers on the incident response teams at each service provider. When the NOC hotline rings, the NOC engineer knows it is one of the NOC engineer's peers.

The online security forum NSP-SEC maintains a closed security operations alias and provides a vehicle for security engineers within NSPs and service providers to communicate and share information before or during attacks. During the Slammer attack, many of the participants in what is known as the "Skitter Group,"—the top Tier 1 and Tier 2 service providers—credited the NSP-SEC with helping to mitigate the impact of the worm. NSP-SEC was the first to report the worm to the general public and, in fact, CERT/FIRST teams received their alerts from NSP-SEC.

BACKSCATTER TRACEBACK TECHNIQUE

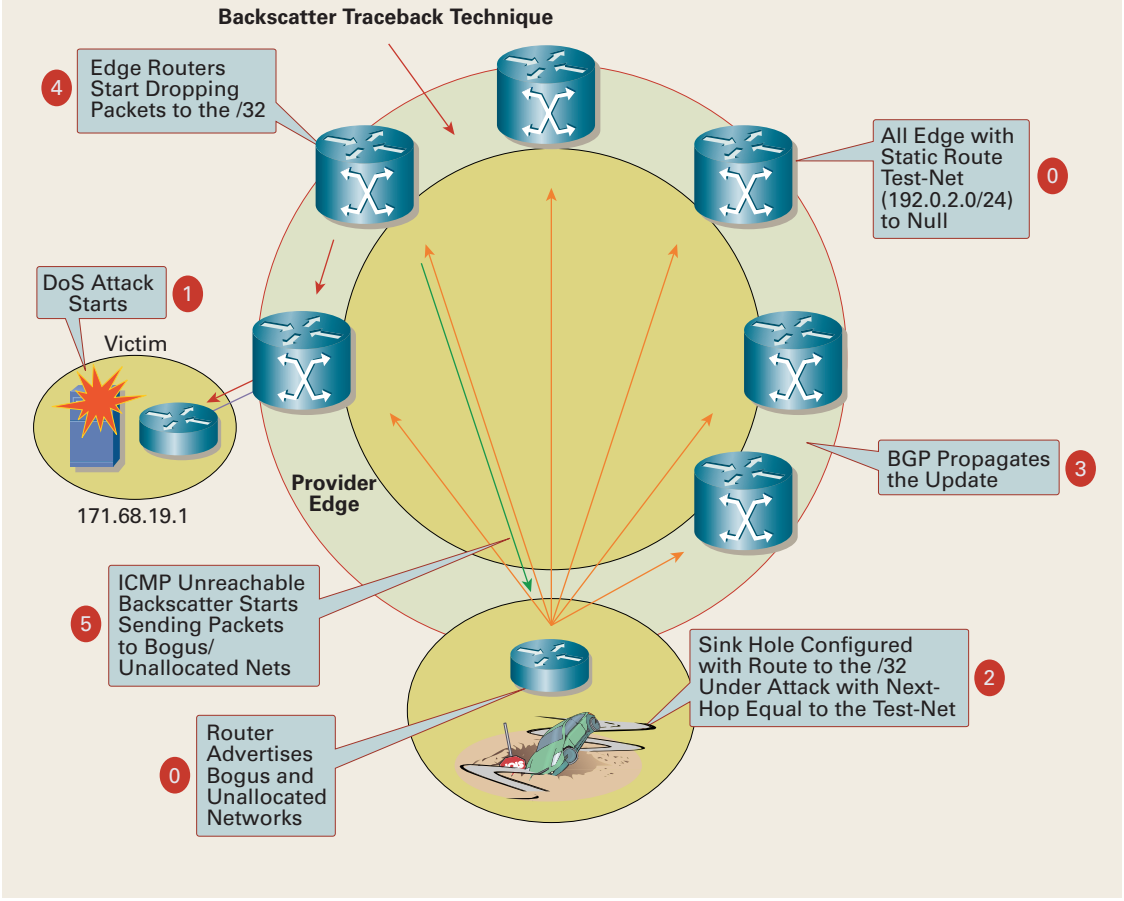

FIGURE 2: Backscatter traceback uses a network's own routing protocols to identify where an attack enters the network when the attack is using "spoofed" source addresses. When the source of the attack is identified, service providers can apply mitigation techniques. This technique was pioneered by UUNET Operational Security (OPSEC) engineers Chris Morrow and Brian Gemberling, and is offered to the greater Internet community as a tool to trace back an attack to the edge of an organization's network.

Detect the Threat

The ability to quickly identify an attack is critical to minimizing the damage that the virus or worm can ultimately cause. But aside from waiting for customers to deluge your NOC with complaints or your network management alarms to begin ringing in the NOC, how can you detect an attack?

Cisco and several service providers extensively deploy a NetFlow-based anomaly detection system that uses NetFlow data exported from Cisco routers to Arbor Networks PeakFlow Traffic and PeakFlow DoS systems. The system allows organizations to detect and characterize undesirable traffic such as DoS attacks (see Figure 1).

NetFlow Data Export must be enabled and the Arbor Peakflow system must be deployed before an attack. The first step is to use Arbor Peakflow Traffic to characterize network traffic; that information is used to set initial thresholds for the Arbor Peakflow DoS system, which dynamically updates traffic profiles and alerts designated personnel when anomalous traffic patterns such as SYN-floods, fragmented traffic streams, and the like, are detected. Based on information from the Arbor system (anomaly type, sources, destinations, protocols, ports, packet sizes, packets per second, specific routers and interfaces involved in the attack, for example) operational personnel instantly receive the information they need to respond effectively.

"Our Arbor-based anomaly detection system is a tremendous force multiplier," says Dobbins. "During the Slammer attack, it took us straight to the router and interfaces where we could see an abnormally high amount of UDP traffic flows coming in through port 1434. We could characterize the type of attack and where it was coming from within minutes. This gave us a tremendous time advantage to apply ACLs [access control lists] to our border routers at our POPs [points of presence] worldwide—while many organizations were still working to characterize the attack and understand its ramifications."

Traceback

Once a service provider has detected an attack, often the next step is traceback—trying to determine the source of the attack so that the service provider can apply mitigation techniques or, if the source of the

Cisco Security Best Practices

When the SQL Slammer virus hit last year, Cisco used the following six-phase implementation to prevent damage to its own network.

Preparation. People, processes, procedures, lines of communication, architecture, and automation tools all need to be in place *before* an event occurs. Security and networking teams must work together smoothly, with no contention over authority. A duty manager is available around the clock to make business decisions as needed, and engineers are empowered to take decisive actions. Professional relationships were established with Cisco's Product Security Incident Response Team; the Cisco Technical Assistance Center (TAC); Cisco Advanced Services; and ISPs, peers, and customers that run other large networks. Cisco participates in FIRST. A detailed communications and escalation plan, facilitated by the operations group and used daily, is in place and well understood.

Identification. Use of Cisco and Cisco partner products and technologies (NetFlow on routers and switches exporting to Arbor Peakflow Traffic and Peakflow DoS anomaly-detection system) allows Cisco to know what is normal for its network, and what is abnormal (for example, unusually high numbers of UDP/1434 traffic flows) and potentially hostile.

Classification. Knowledge of Cisco's own network architecture, network traffic patterns, systems and input from Arbor/NetFlow instrumentation allows Cisco to quickly classify and scope threats.

Traceback. Instrumentation plus knowledge of the Cisco network allows Cisco to identify all presently visible and potential sources and vectors of the attack. This proves to be critical; in the case of Slammer, many organizations failed to account for indirect vectors, such as virtual private networks (VPNs) and laptops, that carried the virus into companies on Monday morning.

Reaction. Cisco immediately "dropped the shutters" through the use of ACLs at all Internet POPs worldwide. A well-designed, "bulkheaded" network allowed a pause so that Cisco could determine its follow-on actions. Knowledge of the virulence and threat level caused Cisco to push ACLs down to the desktop level in every Cisco facility worldwide, along with strategically placed ACLs in the Cisco WAN backbone. This ensured that Cisco wasn't affected on the following Monday. Operations teams provided focal point and bridges for all intergroup communications—network engineers didn't have to search for telephone numbers while dealing with mitigation. Thorough, complete, draconian, and pervasive ACLs were essential.

Postmortem. Cisco conducted daily followup sessions for two weeks to ensure that the threat was eradicated and to discuss lessons learned. Management issued the directive to prioritize and implement lessons learned in concrete, measurable ways, with meaningful followup to ensure future success.

attack is from another network, inform the respective peer.

Attacking sources fall into two groups: valid source addresses and spoofed source addresses.

With valid sources, traceback is not required. Instead, service providers can query online databases such as ARIN's WHOIS database (www.whois.net) to determine address ownership. Additionally, they can use network utilities such as Domain Name Service and traceroute. But these techniques take time and might not be able to determine the source address if the attack is using spoofed addresses. *Backscatter traceback* is the commonly used technique for rapid traceback (see Figure 2). Backscatter takes advantage of Border Gateway Protocol (BGP), the routing protocol pervasively deployed in service provider networks, to drop traffic originally destined for the victim and enables the creation of unreachable Internet Control Message Protocol (ICMP) messages to identify routers that are transmitting data intended for the victim. Once ingress routers have been identified, upstream peers can be contacted to continue traceback on their networks.

The key for backscatter traceback is the existence of a *sink hole*, which is a dedicated portion of the network that attracts traffic. The sink hole advertises large quantities of unused address space, typically referred to as _BOGON or _DarkIP (that is, unused or "unlit" portions of the overall IP address space that have not been allocated by the regional registries and therefore are not valid sources or destinations).

When traffic is dropped at the edge of the network, and if the source of this dropped traffic corresponds to advertised space from the sink hole, the ICMP unreachables generated from the router performing the drops are routed to the sink hole. The sink hole can be monitored to record these ICMP messages and identify the ingress routers, tracing the attack to the edge of an autonomous system.

Reaction or Containment

When an organization knows where an attack is coming from it can apply containment mechanisms such as ACLs. When attack traffic has been detected and classified, appropriate ACLs can be created and deployed to the necessary routers. Because this manual process can be time consuming and complex, many service providers use BGP to propagate drop information to all routers quickly and efficiently. This technique, *remote triggered drop*, sets the next hop of the victim's IP address to the null interface. Traffic destined to the victim is dropped on ingress into the network.

Another option is to drop traffic from a particular source. This is similar to the drop described above but relies on the pre-existing deployment of unicast RPF (uRPF), which drops a packet if its source is "invalid"; invalid includes routes to null0. Using the same mechanism of the destination-based drop, a BGP update is sent, and this update sets the next hop for a SOURCE to null0. Now all traffic entering an interface with uRPF enabled drops traffic from that source.

Although scalable, the BGP-triggered drops limit the level of granularity available when reacting to attack: they drop all traffic to the black-holed destination or source, as described above. In many cases this is an effective reaction to a large attack.

Depending on the type of attack, service providers can use additional reaction methods, other than simply dropping traffic. They can rate-limit using committed access rate (CAR), which limits the rate of administratively identified traffic. For instance, a provider can, upon detecting a high rate of ICMP traffic, rate-limit that traffic to alleviate the effects of the attack. As with traffic filtering, a provider can use QoS Policy Propagation via BGP (QPPB) to remotely trigger CAR configurations.

Postmortem

The final phase of security best practices is postmortem—reviewing what was most effective during an attack and what could be improved. Postmortems should be conducted not only internally, but with other providers as well.

The Skitter Group—a group of the top 10 service providers within the larger NSP-SEC group—came out of an industry-wide postmortem after Cisco issued a major Product Security Incident Response Team advisory. When Cisco issued the advisory, it wasn't communicated consistently to all of the top-tier service providers.

"As a result, some service providers implemented the changes while others didn't, and this led to problems," says Greene. "My contacts at these top-tier service providers later told me that they needed more direct, encrypted communications among themselves. The resulting Skitter Group has further improved coordination against future attacks."

As organizations implement these best practices, the industry is gaining a new edge in the security battle. While attacks are destined to continue unabated, service providers stand a much greater chance of mitigating their most damaging effects through close collaboration and ongoing internal due diligence. ▲▲

◆ ◆ ◆

Joe Dallatore (jdallato@cisco.com) and Paul Quinn (paquinn@cisco.com) of Cisco's Advanced Services engineering group provide security support to service providers and contributed to this article.

FURTHER READING

- **NSP-SEC SP security organization:** puck.nether.net/mailman/listinfo/nsp-security
- **Cisco NetFlow product information:** cisco.com/packet/161_8b1
- **Arbor Networks anomaly detection:** arbornetworks.com/products_platform.php
- **Cisco Catalyst® Network Analysis Module:** cisco.com/packet/161_8b2
- **SQL Slammer Mitigation white paper:** cisco.com/packet/161_8b3
- **Internet security site:** www.cymru.com
- **ISP Best Practices Tutorial:** www.getitmm.com/bootcampflash/launch.html
- **NetWorm.org information center:** www.networm.org
- **Renesys Internet monitoring:** renesys.com/
- **ISP Resource Center:** ispbook.com/