SAFE Code-Red Attack Mitigation

Introduction

This document discusses the recently released Code-Red worms and their effects on the network and its hosts. Today numerous technologies are available for use in Cisco products that mitigate the detrimental effects of the worm. These include not only security technologies such as intrusion detection and packet filtering, but also virtual LAN (VLAN) segmentation, packet classification, and content services. In addition to these technologies, this document will describe how the SAFE blueprint combines security best practices and secure network design to mitigate Code-Red and other attacks. Finally, validated Code-Red mitigating configurations for these technologies are included for reference.

Code-Red Worm Origins and Evolution

Code-Red Background

On Thursday, July 19, 2001, Code-Red compromised at least 359,104 hosts in approximately 13 hours. This statistic, from the Cooperative Association for Internet Data Analysis (CAIDA), shows just how damaging and far reaching the Code-Red worm is. For those unfamiliar with worms, they are self-propagating pieces of code that take advantage of flaws in computer software. In this case, the Code-Red worm took advantage of a remotely exploitable vulnerability in Microsoft's Internet Information Server (IIS) Versions 4 and 5. The worm would send a Universal Resource Locator (URL) to a host that would overflow a buffer in Microsoft's Index Server, a part of IIS. This buffer overflow allowed the worm to execute arbitrary code. Code-Red would install a copy of itself into memory on the infected computer, and attempt to infect additional hosts.

There appears to have been two versions of the worm. The first, called CRv1 by eEye Digital Security, the team that initially analyzed the worm, had a flaw in its random generation of target IP addresses. CRv1, first publicly reported in the wild on July 15, used a random number generator that used a static seed to get new IP addresses to attack. This static seed meant that it would be hitting the same machines over and over again, limiting the ability of the worm to spread. CRv1 spread, but not quickly, and not well. Then, sometime in the morning on July 19, CRv2 was discovered. CRv2 was nearly identical to CRv1 except that it used a better random number generator to create target IP addresses. CRv2 also eliminated the temporary Web site defacing. CRv2 was enormously successful at replicating. Infecting 359,104 hosts in such a short time demonstrates the far-reaching capabilities of this type of an attack. According to CAIDA, at its peak, CRv2 was infecting more than 2000 new hosts every minute.

Code-Red II, not to be confused with CRv2, was first widely reported by the Security Focus ARIS analyzer team. They sent a copy of the worm to the team at eEye, who dissected the worm and published an analysis. Code-Red II is similar to CRv1 and CRv2 only in that it is run completely from memory, it uses the same buffer overflow, and it launches multiple threads to spread itself. Similar to how Code-Red did more damage to English systems, Code-Red II does more damage to Chinese systems.

Function of Code-Red CRv1 and CRv2

Code-Red CRv1 and CRv2 function similarly. They both set up an initial 99 threads to infect other systems. CRv1 uses the 100th thread to check to see if the system is an English system. If it is an English system, then the 100th thread sleeps for a while, allowing the system to propagate the worm without bringing too much attention to the system. After a couple hours of this, the Web site is temporarily defaced with the message:

"Welcome to http://www.worm.com!, Hacked By Chinese!"

The worm actually intercepts connections to port 80 to display this message. No files are modified. CRv2 does not deface the Web site, but it does harmless redirects with the 100th thread when it infects an English system. If CRv1 or CRv2 encounter a c:\notworm file, they go dormant. If there is no notworm file, the worm spreads when the days are 1-19, a distributed denial of service (DDoS) attacks the now former White House Web site when the days are 20-28, and goes dormant from the 28th to the end of the month, only to start the cycle all over again the next month. DDoS attacks work by using multiple compromised systems to packet flood a target so that it is unavailable to provide legitimate usage. CRv1 DDoS attacks the hard-coded former IP address of the White House Web site (198.137.240.91), whereas CRv2 resolves www.whitehouse.gov to determine the IP address.

Function of Code-Red II

Code-Red II was based on an entirely different worm than Code-Red. The functionality, although similar in some ways, is very different. When Code-Red II infects a box, it sets up 300 threads on non-Chinese systems, and 600 threads if the system is Chinese. The worm infects other systems during this time for one day if the system is non-Chinese, and two days if the system is Chinese. After this waiting, the worm reboots the system so that file system protection mechanisms are disabled. Code-Red II, unlike CRv1 and CRv2, affects only Windows 2000. The offset used for the attack does not work on Windows NT systems. Code-Red II also used a very different algorithm to determine its next target. It was more likely to attack systems that shared the first two octets of its own IP address. There was still a chance that it would attack systems with only the first octet being the same, or a completely random IP address. This target selection is the reason that many more corporations were hit internally with Code-Red II. Filters were not in place to disallow their own Web servers from attacking Web servers internal to the company. Code-Red II also left a few backdoors. For one, it left a copy of cmd.exe in two different directories renamed to root.exe. This copy can be used to execute arbitrary commands on the server. It also created virtual mounts to both drives C and D on the Web server. Even if the root.exe files are removed, attackers can still access the cmd.exe file on either C or D.

The Effect of the Worms on the Network and Its Hosts

The DDoS attack disrupts not only the site under attack but also the local network of the compromised host. Depending on the number of infected Web sites in a network, the amount of load generated by these attempts could cause a local network disruption. This disruption varies from a slow network to an unusable network as pipes fill and devices fail from the unexpected load. Services running on any infected system are likely to be slowed, and their legitimate usage possibly blocked.

Cisco's Recommendations for Mitigating Code-Red

Patch every vulnerable system

The most effective manner in which to mitigate the Code-Red and its variants is to patch all systems that are vulnerable. This patching is difficult with uncontrolled user systems in the local network and even more troublesome if they are remotely connected to the network via a virtual private network (VPN) or remote access server (RAS). However, determining which devices are exploitable can be simplified by the use of security auditing tools that look for vulnerabilities. The following links provide information regarding infection mitigation on Microsoft and Cisco Systems products:

- http://www.cisco.com/warp/public/707/cisco-code-red-worm-pub.shtml
- http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS01-033.asp

Many vendors' products install and use IIS to provide Web access for remote management and reporting, and these are also vulnerable unless patched. If it is not possible to patch all systems in a timely manner, consider deploying the technologies discussed in the following section. You should also consider using these technologies proactively to mitigate future attacks by variants of Code-Red or other attacks altogether.

Security Technologies

This section discusses the technologies available in Cisco products to mitigate Code-Red and other attacks. Security technologies, the preferred method, are addressed first. Other technologies that provide attack-mitigation capability are also discussed in the following section. These other technologies may provide an interim solution in case some of the security technologies are not available in your network. To learn more about any of these technologies, or for Code-Red mitigating configurations, refer to the SAFE white papers at http://www.cisco.com/go/safe.

Host-Based Intrusion Detection System (HIDS)

The host-based intrusion detection system (HIDS) operates by detecting attacks that occur on a host on which it is installed. It works by intercepting OS and application calls, securing the OS and application configurations, validating incoming service requests, and analyzing local log files for after-the-fact suspicious activity. It has two modes of operation: monitor (alarm only) and enforce. HIDS performs many security functions:

- Analyzes incoming Hypertext Transfer Protocol (HTTP) traffic and, via the use of generic rules and known attack signatures, determines if it is an attack
- Analyzes the actions of the HTTP server to determine if they reflect its normal mode of operations
- Generally protects the OS, including buffer overflow prevention and binary modification

In the case of Code-Red and its variants, HIDS secures IIS by disabling the indexing service. It also sends an alarm to the centralized monitor console that an exploitation attempt was intercepted.

It may appear that deploying HIDS has the same problem with exploitation mitigation as discussed previously for applying system patches. However, HIDS clients are significantly easier and less obtrusive to install on running systems, and they are less likely to require system interruptions or reboots. To target specific systems for HIDS installation for the current problem, use a network security scanner to identify those systems that are running Web services. To mitigate future attacks beyond Code-Red, consider installing HIDS on critical servers.

Network-Based Intrusion Detection System (NIDS)

The network-based intrusion detection system (NIDS) operates by first detecting an attack occurring at the network level and then either taking a corrective action itself or notifying a management system where an administrator can take action. Attacks are discovered by looking for their signatures in traffic flows in the network. Attack detection triggers NIDS to send an alarm and then take a preconfigured action. The two possible actions are shunning and TCP resets. Because NIDS is not in the data path (meaning it receives a copy of a packet as it traverses through the network versus routing the packet), NIDS cannot filter

the first packet in an attack. Subsequent packets can be filtered via a feature known as shunning, which modifies the upstream access-control device to block any further access from the IP address of the attacking system. TCP resets attempt to tear down the TCP connection by sending a fabricated reset that appears to be from the receiving device to the attacking device.

If you are considering enabling shunning in your network, refer to the SAFE white papers for more information, because you need to consider special considerations when using this feature. Because CRv1 and CRv2 contain the attack in a single packet, NIDS cannot stop the attack. NIDS does, however, provide visibility by sending an alarm when CRv1 and CRv2 attacks traverse the network. NIDS is able to stop Code-Red II attacks with high probably through the use of TCP resets as Code-Red II uses multiple packets. For more information in NIDS, refer to: http://www.cisco.com/go/ids.

Access Control

Stateful firewalling provides numerous security features to proactively mitigate Code-Red. First, the stateful inspection engine can control connection attempts at a level more granular than normal by validating proper protocol adherence. This filtering could be used to allow only inbound connections to a Web server and at the same time disallow that Web server from initiating outbound connections, thus limiting the ability of the worm to self-propagate. This filtering is particularly applicable for DMZ Web-server deployments. As discussed in SAFE, your Web servers do not normally need the ability to establish outbound connections to say, surf the Web. In most cases they need to respond only to incoming Web requests. Second, stateful firewalling has the capability of limiting the number of permitted inbound connections to a server so that the server does not become overwhelmed. In the case of Code-Red, this limiting blocks inbound exploitation connection attempts.

Ingress filtering is typically carried out by access control on the perimeter of the network. It is used to block access to hosts and services that should not be publicly available. For instance, it is a security best practice to disallow incoming connection requests to hosts or networking devices unless those hosts or devices are actively participating in providing a publicly accessible service. As it pertains to Code-Red, incoming HTTP connections would be blocked from accessing any possibly exploitable user systems or nonpublicly available Web servers. These same filters, however, would need to allow access to a publicly available Web presence or e-commerce server. Ideally the public servers are under tight administrative control and have the latest patches. Ingress filtering would in effect block Code-Red exploitation attempts targeted at user systems.

Egress filtering is also typically carried out by access control on the perimeter of the network. This filtering blocks a local host's access outbound out of your network. Devices that don't need outbound Internet access, such as most of the networking devices in your network or Web servers that serve only the internal environment, should not be allowed to initiate outbound connections. As this pertains to Code-Red, if a device is compromised it will not be able to launch a DDoS attack against an external network because the traffic will be intercepted and dropped at the perimeter of your network. This setup will also guard against the DDoS attack flooding the Internet link and interfering with legitimate inbound or outbound traffic. Additional layers of egress filtering in the network in addition to those at the WAN edge could also be used to disallow an infected public Web server (or its entire segment for the case of a Web farm) from infecting private internal servers that were protected by the edge ingress filtering. For more information on access control and filtering, refer to the SAFE white papers.

Private VLANs

Private VLANs work by limiting which ports within a VLAN can communicate with other ports in the same VLAN. Typically private VLANs are deployed so that the hosts on a given segment can communicate only with their default gateway and not the other hosts on the network. For instance, if a Web server is compromised by Code-Red, it will not be able to initiate infection attempts to other Web servers in the same VLAN even though they exist in the same network segment. This access control, carried out by assigning hosts to either an isolated port or a community port, is an effective way to mitigate the effects of a single compromised host. Isolated ports can communicate only with promiscuous ports (typically the router). Community ports can communicate with the promiscuous port and other ports in the same community.

For more information on private VLANs, refer to: http://www.cisco.com/warp/public/473/90.shtml

Additional Cisco Networking Technologies to Assist in Mitigating Code Red

Network-Based Application Recognition

Network-based application recognition (NBAR) is a classification engine in Cisco IOS® Software that can recognize a wide variety of application level protocols, including HTTP via URL/Multipurpose Internet Mail Extensions (MIME) type and protocols that utilize dynamic port assignments. After the traffic has been classified by NBAR, appropriate quality-of-service (QoS) policies can be applied to the traffic classes. NBAR recognizes the CRv1 and CRv2 URL request but not the Code-Red II URL request because Code-Red II spreads the GET request over multiple packets and NBAR today inspects only the first packet. Unlike NIDS, NBAR can immediately classify the CRv1 and CRv2 traffic and drop the packet before reaching the server. NBAR can be used inbound and outbound to mitigate the effects of Code-Red.

For more information on NBAR, refer to:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121limit/121e/121e2/nbar2e.htm

Content Engines and Accelerators

Content-aware devices handle content distribution as they offload some of the functionality normally performed by a content device, such as a Web server, onto a high-speed network appliance. Because these devices broker the connections between the content requester and content server, they can alter the connection establishment. These devices have the capability to identify the Code-Red exploitation based on its unique HTTP request and drop it before it reaches the Web server. This scenario can be accomplished with two different methods, internal block/allow lists as well as Websense integration. With the internal block/allow list, the lists of URLs may be explicitly blocked or allowed using the internal block/allow feature. An administrator may upload text files containing lists of URLs to be allowed or blocked. The URLs contained within a list are matched using a substring match from the beginning of a URL. Cisco does not recommended use of this technology at moderate to high rates of speed. If the content engine (CE) becomes overtaxed, the Web Cache Control Protocol (WCCP) switches into bypass mode and ceases filtering until the performance condition is resolved.

For more information on content engines, refer to:
http://cco/warp/public/779/largeent/learn/technologies/content_networking.

Sink-Hole Routers

Setting up a sink-hole router will assist in determining which systems in your environment are infected when NIDS is not available. This scenario works by using addresses not yet allocated by the Internet Assigned Numbers Authority (IANA) that Code-Red will inadvertently attempt to exploit. The sinkhole router advertises these networks locally (only), and any attempts at reaching them are then routed to the router. When received, they can be logged and discarded. The results of the logs will provide a list of infected hosts.

For more information on how to configure this functionality, refer to: http://www.cisco.com/public/cons/isp/security/

The SAFE Blueprint

The SAFE blueprint utilizes all the security technologies listed above to mitigate Code-Red. For this reason, the SAFE blueprint is “Code-Red safe.” Ingress and egress filtering is applied not only at the network edge but also between virtually all SAFE modules. This filtering restricts outbound access from infected servers and inbound infection attempts against user systems. Stateful firewalling protects both the user and server segments in addition to the filtering and provides DDoS connection rate limiting for the public servers. NIDS is deployed not only in all public segments to identify Code-Red infection attempts but also behind the network edge filtering and stateful inspection to determine if any exploitation attempts made it through the edge. HIDS is installed on all publicly available servers and even critical internal servers that do not have Internet access to guard against possible infection from uncontrolled user systems. Private VLANs are deployed in public-service segments where multiple public servers are available to guard against trust exploitation.

Conclusion

The technologies discussed in this document mitigate not only the potential damage done by Code-Red and its variants but also virtually any attack. It is important to remember that security has its place throughout the infrastructure, and the discussed technologies prove this. Protecting your network and its resources against Code-Red is only the first step. It is necessary to be proactive when it comes to security so that you can protect the network not only against Code-Red but future attacks as well. Establishing a security policy, implementing some of the discussed features, and regular in-house or outsourced posture assessments will secure your network and keep it secure. This document has addressed a small sampling of the documented security and network design best practices available from Cisco Systems. For additional information on securing your network, refer to the SAFE blueprint at www.cisco.com/go/safe.

As with any feature, if you are considering enabling some of the discussed features, ensure that your devices have sufficient CPU resources available. Also realize, however, that the increased load brought on by enabling these features is significantly less than that of the load brought on by an internal Code-Red infection.

As a special note, the SAFE Blueprint was released in October 2000. No design or implementation modifications were required to deal with Code-Red. Only NIDS signature updates at regular intervals were necessary to detect the IIS exploit and Code-Red. As Code-Red and other high-profile network exploits constantly remind us, designing network security reactively is not recommended. Only by taking a comprehensive approach to network security founded on good security policy decisions can your organization be assured that the risks you are taking are known, and that virtually any potential threat can be effectively contained.

Configuration Information

This section provides sample configurations for some of the technologies discussed in this document that were not tested for attack mitigation capabilities as part of SAFE or that required later configuration changes. HIDS is not discussed because the mitigation capability it provides is available out-of-the-box and requires no additional configuration beyond placing the system in active mode.

NIDS Attack Signatures

The signatures provided below were added to NIDS systems (Cisco Secure IDS 4210 Sensor, Cisco Secure IDS 4230 Sensor, Intrusion Detection System Module) in many modules of the SAFE Blueprint.

Index Server Access with Attempted Exploitation

String:
“[Gg][Ee][Tt].*[.][Ii][Dd][Aa][\x00-\x7f]+[\x80-\xff]”

Occurrences: 1

Port: 80

If you have Web servers listening on other TCP ports (for example, port 8080), you will need to create a separate custom string match for each port number.

Recommended alarm severity level:
– High (Cisco Secure Policy Manager [CSPM])
– 5 (UNIX Director)

Index Server Access Buffer Overflow Code-Red Worm

String:
"[/]default[.]ida[?][a-zA-Z0-9]+%u"
Note that there are no blank spaces in the above string.

Occurrences: 1

Port: 80

If you have Web servers listening on other TCP ports (for example, port 8080), you will need to create a separate custom string match for each port number.

Recommended alarm severity level:
– High (CSPM)
– 5 (UNIX Director)

NBAR Marking

A Cisco 7206 VXR was used in these configurations. Three methods to discard the packet are addressed below; it should be noted that testing showed that NBAR policing had the least effect on CPU utilization. The following commands classify Code-Red traffic and mark it with the first differential-services-control-point (DSCP) value.

```
class-map match-any http-hacks
  match protocol http url "*default.ida*"

policy-map mark-inbound-http-hacks
  class http-hacks
  set ip dscp 1

interface FastEthernet 2/0
  service-policy input mark-inbound-http-hacks

interface ATM 4/0
  service-policy input mark-inbound-http-hacks
```

NBAR Marking with ACL Block and Log (optional)

The following commands use the DSCP marking to deny the packet from traversing outbound from the device and to log to that effect. Take special consideration when enabling access-control-list (ACL) logging, ensuring that significant packet loads do not overwhelm the router.

```
access-list 105 deny ip any any dscp 1 log
access-list 105 permit ip any any

interface ATM 4/0
  ip access-group 105 out

interface FastEthernet 2/0
  ip access-group 105 out
```

NBAR Marking with Policy Route to Null0

The following commands use the DSCP marking in combination with policy routing to discard the packet.

```
access-list 106 permit ip any any dscp 1

route-map null_policy_route 10
 match ip address 106
 set interface Null 0

interface ATM 4/0
 ip policy route-map null_policy_route

interface FastEthernet 2/0
 ip policy route-map null_policy_route
```

NBAR Policing Drop

The following commands use NBAR policing to discard the packet.

```
policy-map drop-inbound-http-hacks
 class http-hacks
 police 100000000 50000 50000 conform-action drop exceed-action drop
 exit

interface ATM 4/0
 service-policy input drop-inbound-http-hacks

interface FastEthernet 2/0
 service-policy input drop-inbound-http-hacks
```

Content Engine (transparent mode)

The following commands configure a block rule that uses a regular expression to match the Code-Red signature and then discard the packet. This functionality was tested using Cisco CE-550 and CE-590 content engines.

```
!CE blocking filter rule
rule block url-regex ^http://.*/default\.ida$
```

Links to Additional Information

Cisco Systems response to Code-Red and required patches:
http://www.cisco.com/warp/public/707/cisco-code-red-worm-pub.shtml

Microsoft response to Code-Red and required patches:
http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS01-033.asp

Cisco Technical Assistance Center (TAC) technical tips for dealing with Code-Red:
http://www.cisco.com/warp/customer/63/codered_index.shtml

eEye documentation on Code-Red:
http://www.eeye.com/html/Research/Advisories/AL20010717.html
http://www.eeye.com/html/Research/Papers/DS20010802.html

Description of the IIS vulnerability and NIDS signature ID:
http://www.cisco.com/go/csec. Search for ID 3394.

Computer Emergency Response Team (CERT) information on Code-Red:
http://www.cert.org/advisories/CA-2001-19.html
http://www.cert.org/advisories/CA-2001-23.html

Information on the SAFE blueprint: www.cisco.com/go/safe

Links to Cisco Products/Services

NIDS: http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz

Network scanners: http://www.cisco.com/warp/public/cc/pd/sqsw/nesn/index.shtml

Information on Cisco security products and security consulting:
http://www.cisco.com/go/security
http://www.cisco.cm/go/securityconsulting

The Cisco PIX® Firewall: http://www.cisco.com/warp/public/cc/pd/fw/sqfw500/

Content engine/content services switch (CE/CSS) content networking devices:
http://www.cisco.com/warp/public/779/largeent/learn/technologies/content_networking/switch.html

Websense content filtering server:
http://www.cisco.com/warp/public/779/largeent/partner/esap/profiles/Websense_entv3.html

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy Les Moulineaux
Cedex 9
France
www.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 17, 99 Walker Street
North Sydney
NSW 2059 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

Cisco Systems has more than 190 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the

Cisco.com Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

0801/DA