



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Code Red and the Internet Today

**SANS Security Essentials
GSEC Practical Assignment
Version 1.2e**

By Andres Chiriboga

Introduction

While most Internet users are used by now to spam, banner ads, slow connections, viruses and other annoyances that come with the use of the Internet, few are prepared for the threat of an Internet moving at crawling speeds, or even worse, not moving at all.

Code Red and its latest incarnation, Code Red II, posed that scenario as a possible reality in the last couple of months. But it was not that we could lose data, or have our personal information tampered with or exposed (even though that is a possibility with the vulnerability that Code Red and Code Red II exploit) which created the concern that prompted the FBI, Microsoft and other groups to make a public statement about this issue. It was the fact that the worm could “propagate and clog the Internet to a crawl”¹. CNN carried a full-feature primetime news segment on Code Red, and every major news broadcast did more than simply mention the issue on Monday, August 30th, 2001. Why? Because government officials “feared that when the worm re-awoke it would spread rapidly, scanning the Internet for unprotected servers and in the process flooding the Web with unwanted packets of data, causing it to slow”¹.

But what are Code Red and Code Red II, and how did they become so feared by Internet users?

At the speed of the Internet

Not very fast... that’s what the speed of the Internet is for most users today, until fast connections like DSL, cable and leased-lines become more widely available and affordable. About 92% of the 25.2 million home offices in the U.S. access the Internet using a dial-up connection². The Gartner Group estimates that about 55% of households will still dial-up to connect to the Net by the year 2004³, which probably means connections at 56K or even less. Many new and “modern” neighborhoods in the US today do not have fast Internet access available due to limitations in technology, or the Telco’s lack of funds to implement new technologies in areas that are geographically out of bounds for such access.

However, new web sites come out at neck-breaking speeds to serve the millions of users that “get connected” every day around the world. In order to serve sites to these millions and millions of Web surfers worldwide, new Web servers are installed every day. And this is where the issue starts to become evident. In the race to put out sites as fast as possible and provide hosting for as many customers as possible, many ISPs and hosting companies often sacrifice proper security measures in exchange for a faster return on investment (ROI). That is, until something like Code Red comes around... more on that later.

Code Red (the original) threatened to bring the Internet to a crawl. We have become so dependent on the Internet, that even though the majority of Internet users connect at slow speeds, a potential (major) decrease in the speed of the Internet as a whole, poses a threat to the way we do business on a day-to-day basis. But the concern is not for the average Joe, surfing the Web to find information to write a paper for his certification, or the casual Internet shopper, looking to get a good deal on the latest software late one evening. The concern is the effect of

such a slow down on Corporate America, and the world.

Corporations these days rely on the Internet for B2B, B2C and other acronyms to be profitable, and to survive in this extremely competitive environment, where a one-man shop can compete with the largest of corporations for your business. The backbone of the American economy is shifting and spreading to the Internet at an incredible pace, and it is no longer a matter of inconveniencing your casual Web surfer when there's problems with the Internet. It is a matter of big business, of revenues lost, of shipments not being processed on time, of companies making the headlines when their technical difficulties become a problem for their customers. Time, after all, *is* money.

Enter Internet Information Server

So if time is money, how do we get more of both? Microsoft Corporation, in its competitive spirit, developed and provided, free of charge, the Web server software known as Internet Information Server (IIS) for its Windows NT platform. Windows NT is one of the most popular Operating Systems in the world⁴ according to several research studies, including one from IDC (see the references section) and, arguably, one of the easiest to administer. Free software and easy administration translates into more time and money for corporations eager to setup a Web presence, or hosting companies looking to serve more customers.

Whether IIS is a good Web server software or not depends on who you ask. Is it easy to setup and deploy? I think so. Does it serve its purpose? I think so. Does it come free of bugs or security flaws? Definitely not!

Microsoft has earned a reputation for rushing to put out products that are not necessarily ready for market, and that end up being “buggy” or otherwise flawed. Many people say they would rather wait for the release of the first service pack before they implement a new Microsoft product. It is a reality, whether Microsoft likes it or not, and IIS was no exception. With an estimated 6 million servers worldwide running IIS¹, a major security vulnerability could prove to be a major disaster for much of the Internet. And IIS delivered.

Enter the Vulnerability

On June 18, 2001, eEye Digital Security, in Aliso Viejo, California, discovered a flaw in the .ida ISAPI filter, that makes a default installation of IIS 4.0 or 5.0 susceptible to a buffer overflow attack. But that is not it. An attacker can gain full System access to a vulnerable server, which would allow him/her to pretty much do whatever he/she wanted on that server, including executing arbitrary code with System level permissions, adding and deleting files, accessing databases at will, etc.⁶

The .ida ISAPI filter is a component that allows IIS to interact with Microsoft Index Server (in IIS 4.0) or Indexing Service (in IIS 5.0), to perform content searches and other functionality offered by Index Server. This filter is installed by default with IIS, and is in the file Idq.dll. Originally, it

was thought that only servers with Index Server installed would be vulnerable, but it was soon determined that even if Index Server was not installed, an IIS server could still be vulnerable to this attack.

Basically, the idq.dll file contains an unchecked buffer in its code that allows for a malicious user to request a specially formed URL that would create a buffer overflow. Specific technical details of this vulnerability can be found in the eEye site's June 18, 2001 advisory, at <http://www.eeye.com/html/Research/Advisories/AD20010618.html>.

Enter the Exploit

Code Red is the name of the worm that first exploited this IIS vulnerability. It came soon after it was discovered by eEye, and it was also discovered and named by eEye. According to eEye, the name "Code Red" was given to this worm in honor of the "Code Red" Mountain Dew that kept them up while they were dissecting the worm code.

This is, in a nutshell, how the worm infects a system and what it does afterwards:

1. A vulnerable IIS server receives an HTTP GET request with the worm code. This request calls any name, usually "default" (or NULL) with the ".ida" extension appended to it, and passes it a query string that is approximately 240 bytes in length. Code Red uses "N" as the character to create the overflow. The worm code is appended to this.
2. The initial worm code contained in the HTTP request executes, creating a new memory stack for its own use. At this point, the base address of w3svc.dll is stored for later use in defacing the website.
3. The worm now creates 99 threads, which are an exact replica of the worm code.
4. Each thread checks for the existence of c:\notworm or d:\notworm. If the file exists, then the worm will "go to sleep".
5. If "notworm" does not exist, and the date is less than the 20th of the month, UTC time, it will start spreading itself to other vulnerable servers by sending the same request with the worm code to a sequence of random IP addresses. Since there is a static seed for the generation of these random IP addresses, every infected computer will potentially send requests to the same list of IP addresses, which would create excessive packet traffic, and could result in a denial-of-service (DOS) effect.
6. The 100th thread will perform a slightly different job at first. It will check if the language of the infected system is English. If so, it takes over w3svc.dll to intercept all requests to the server, and return code of its own instead, which would read: "Welcome to <http://www.worm.com/>, Hacked by Chinese!" This message will not be returned until 2 hours after the initial infection time, and will be returned for all pages requested for 10 hours. After that, the message will disappear, until/if the system is re-infected again. Meanwhile, w3svc.dll is returned to its original state. If the codepage is not in English, the 100th thread will simply act as another infecting thread. As an interesting note, eEye believes that the thread waits 2 hours before "defacing" the site, to allow for some time for the worm to spread, before calling attention to itself.
7. Once the date reaches the 20th of the month, UTC time, all threads within the worm shift their

attack to www.whitehouse.gov, by sending it 100K bytes of data, one byte at a time. This is clearly a distributed-denial-of-service attack attempt.

Interestingly, this worm was hard-coded to attack 198.137.240.91, which is only one of the IP addresses for www.whitehouse.gov. Soon after the attack was discovered, that IP address was simply disassociated from the URL, which rendered the attack useless. However, the major issue here is the amount of traffic generated by the estimated 760,000 computers infected worldwide, and the cost associated with it.

Good Coverage, Less Damage

While the potential for major damage of this worm was clearly there, good press coverage, quick reaction by security industry groups and Microsoft itself, helped to contain what could have amounted to a major “cyber-catastrophe”. In a rare show of concern, the Federal Bureau of Investigation’s National Infrastructure Protection Center (NIPC), along with security industry groups and Microsoft, held a press conference on July 30th, to urge systems administrators and businesses worldwide to install the free patch that Microsoft has made available since June. This created headlines around the world, and the coverage given to this threat was unlike any other for this type of incident.

With millions of servers around the world running IIS, it is almost impossible to reach, and sometimes persuade, all systems administrators to install the patch. The coverage given to the story, though, helped reach thousands more than would have been possible by word-of-mouth and specialized media alone, resulting in thousands or even millions of servers being “patched” before or soon after the worm began to spread again on August 1st UTC time. Microsoft states that more than 1 million downloads of the patch had been recorded between the time it was released and July 31st. Reportedly, another million or so downloads took place in the few days that followed, after the press conference took place. The number of downloads does not necessarily correspond to the number of servers patched, since most server administrators handle more than one server, and would likely download it once and use it on several servers.

This wide spread coverage also had a side effect. Regular computer users, running Windows 95, 98 and ME systems panicked, thinking that they were also at risk. E-mails circulated the Internet with warnings that users should stay off the Internet on Tuesday (July 31st) night, when the worm would start spreading again (August 1st UTC time), and that everyone with Windows machines was at risk. Microsoft reported several users calling because they were having problems installing the patch on systems other than Windows NT or Windows 2000.

Even with all the information available, there is one factor that is harder to manage than the installation of a patch on a server: the regular everyday users. In Windows 2000 systems, IIS 5.0 is installed by default, making it an instant vulnerable system. A server administrator might have done all his/her work properly and in a timely manner. But what happens when an Executive takes his Windows 2000 laptop home, connects to the Internet from home and gets infected with the worm? The next day he comes back to work plugs-in to the network, and the worm goes to work. Hopefully, all internal servers were patched, but there might be a few computers lying

around that few people knew about, and that had IIS 4.0 installed, or maybe other users' Windows 2000 machines. All of a sudden, the internal network is flooded with requests from these infected computers. The network comes to "a crawl", maybe e-mail goes down, and frustrated users flood the help desk with calls. Education and control are the keys here. Education for the regular non-technical user, so that he/she knows what to do and not to do. Control of what machines exist, what software is installed, and who owns them, and a good inventory to make sure patches are applied to all vulnerable computers before they become a problem. Sharing of specific information by the IT departments with the general users is very important, just as it was important for information to be disseminated to the general public.

Overall, the media coverage proved to have a positive effect on the whole incident. That is, unless you count disseminating excessive information about the worm.

When is it too much information?

eEye is credited for discovering both the IIS vulnerability and the Code Red worm that exploits it. However, several industry experts have criticized eEye for providing too much information about it to the general public. Even though eEye alerted Microsoft of this vulnerability immediately, almost at the same time, it released its own [advisory](#), with detailed information on the IIS flaw.

Is it really necessary for security experts and security firms to release so much information about a flaw? It could be interpreted as a trait of the typical hacker: the need for recognition and to demonstrate vast knowledge. Security experts are usually, and should always be, cautious about what information is provided to whom. Are some security-experts falling prey to a mutation of the so-common social engineering? There are plenty of very smart kids out there with lots of time on their hands who I am sure, welcome the opportunity to take a publicized flaw like this one, and make themselves known for exploiting it (in the underground world, of course). Providing too much detail on how a flaw can be exploited is like giving them a recipe to do it, with the added bonus that they have to figure out the exact way to do it. Sounds like a challenge. Isn't that what hackers thrive on?

On the other hand, one could argue that disseminating this information allows for quicker development of effective counter-measures for an exploit. But shouldn't that dissemination be less public? It is a debatable issue, but an important one that the security industry should look into.

Would the Code Red worm exist if eEye had not made the information public? What if only Microsoft had been informed, a patch had been released, and no one else found out about it? How different would things be now? Then again, how would a systems administrator that is not necessarily up on the latest patches have found out about this important issue?

Conclusion

This IIS vulnerability is only one of many flaws that exist in the software that runs the Internet and most businesses today. It is not only Microsoft, it is not only IIS. I get about half a dozen e-mails every day with warnings on security flaws found in all sorts of software and hardware, from Linux, to Apple, to Apache, to Windows NT, to Windows 2000, to firewall software, hardware, etc. After all, is there a perfect piece of software? I don't think so. Is there a computer system that is 100% secure? Yes, it is the one that is unplugged from the wall and locked in a closet.

Otherwise, there would be no need for Information Security departments, and quite frankly, most computer/software improvements that we've seen since the beginning of the "computer era" would probably not have happened without those "curious, inquisitive, poking minds" that invented them. Those are the same type of minds that now create worms like Code Red. Why do they do it? Who knows! It is basic instinct, part of the balancing of the species, just like there is crime in the real world, and we don't like it, it is a reality that there is and will always be crime in the "virtual world".

As security experts, systems administrators and software developers, we need to maintain the awareness, and walk the fine line between maximum security and minimum functionality. Users want convenience, speed and functionality. We want control, security and audit trails. Users want all these things too, but are usually not willing to sacrifice their convenience, speed and functionality for them. Educating users is a key part in the whole security field. The better educated the user, the better protected the resources under his/her control. This also goes for systems administrators.

Code Red II is now out in the wild. It functions very similar to the original Code Red, but it is not a variant of it. Code Red II works only on Windows 2000 machines, and propagates more efficiently and faster. It also creates a backdoor to the infected system using Trojan functionality. Code Red II exploits the same vulnerability as the original Code Red, so a system that has been patched already, is safe from it. But the question remains how many systems have not been patched yet and will be compromised? We will find out soon enough.

References

1. "Code Red wakes up with a whimper" – From NetworkWorld Fusion
(<http://www.nfusion.com/news/2001/0801codedred.html>)
 2. "Are dial-up networks endangered?" – From CNN.com and NetworkWorld Fusion
(<http://www.cnn.com/2001/TECH/internet/07/31/dial.up.ignored.idg/index.html>)
 3. "Is dial-up Web access dead?" – From CNN.com
(<http://www.cnn.com/2000/TECH/computing/03/23/dial.up.idg/>)
 4. IDC's IT Forecaster (Aug. 2000)
(<http://www.idc.com/itforecaster/itf20000808.stm>)
 5. Microsoft Security Bulletin MS01-033
(<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/ms01-033.asp>)
 6. eEye Advisory (<http://www.eeye.com/html/Research/Advisories/AD20010618.html>)
 7. "Security Firm Blamed for Code Red Costs" - WashingtonPost.com
(<http://www.newsbytes.com/news/01/168934.html>)
-

© SANS Institute 2000 - 2005, Author retains full rights.