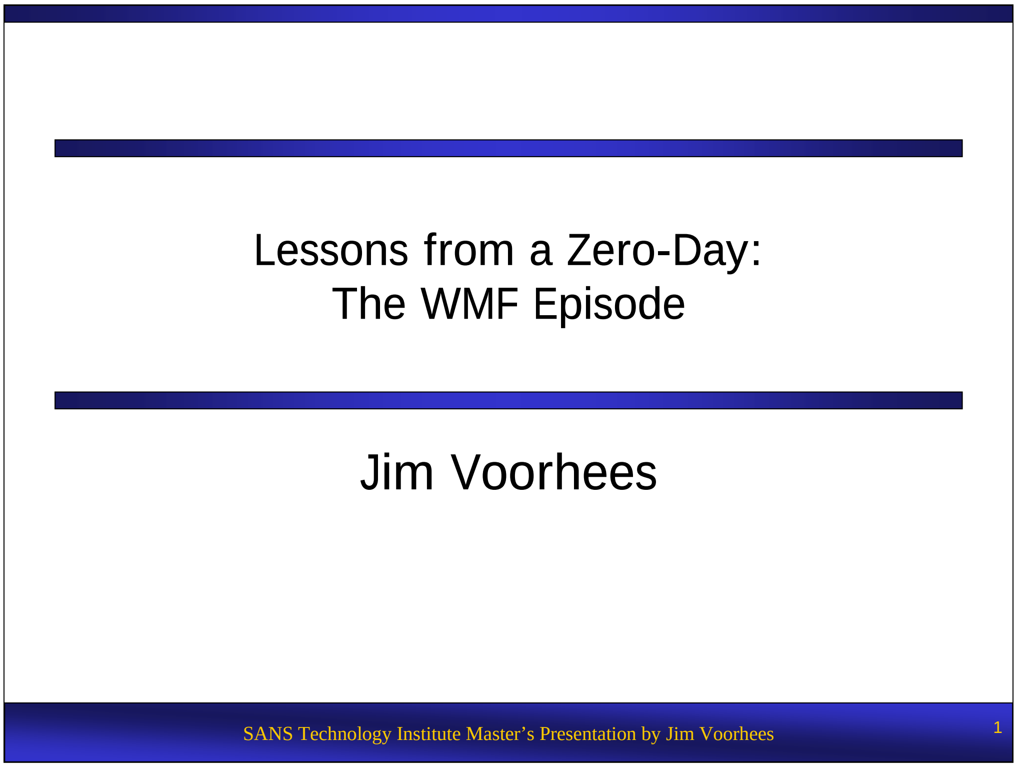

Lessons from a Zero-Day: The WMF Episode

Jim Voorhees

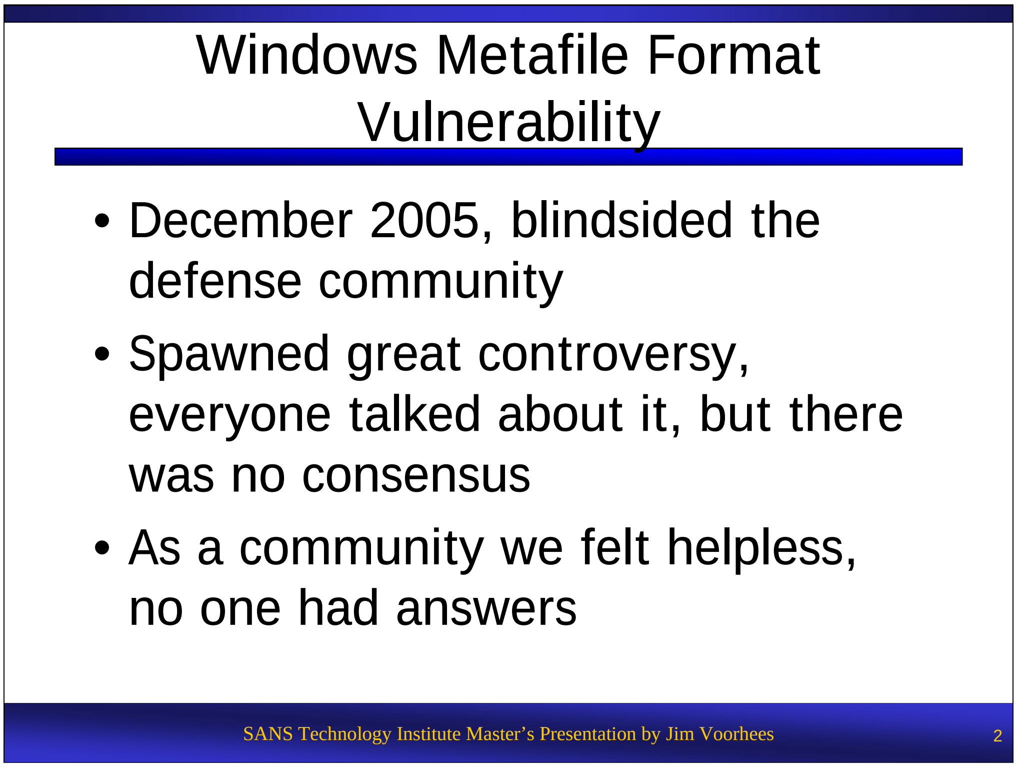

Windows Metafile Format Vulnerability

- December 2005, blindsided the defense community
- Spawned great controversy, everyone talked about it, but there was no consensus
- As a community we felt helpless, no one had answers

On December 27, 2005, a handler at the Internet Storm Center, found it so quiet that she speculated that “Perhaps all of the script kiddies got new computers for Christmas and haven't gotten fully up to speed yet.”

Within hours, however, frenzy would replace that quiet as telephone calls and email messages showed that a vulnerability in Windows Metafile Format (WMF) files, heretofore unknown to most of the world, was being exploited. Exploits multiplied exponentially from that time on, with 200 individual exploits and more than 1100 infectious URLs appearing before Microsoft issued a patch

The vulnerability gained the attention of the entire security community. Extraordinary efforts were made to find a fix for the problem. But no complete fix was available to most users until Microsoft's patch made its patch available more than a week later, on Thursday, 5 January 2006.

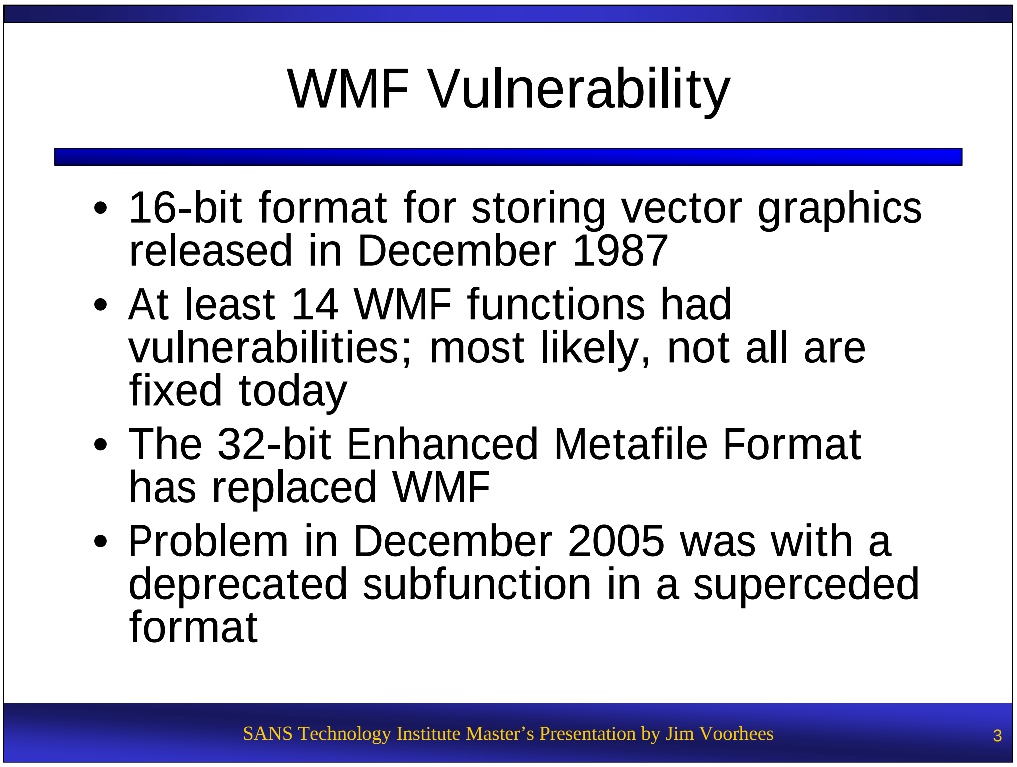

WMF Vulnerability

- 16-bit format for storing vector graphics released in December 1987
- At least 14 WMF functions had vulnerabilities; most likely, not all are fixed today
- The 32-bit Enhanced Metafile Format has replaced WMF
- Problem in December 2005 was with a deprecated subfunction in a superceded format

SANS Technology Institute Master's Presentation by Jim Voorhees 3

WMF is a 16-bit format that first appeared in Windows 2.0. This operating system was designed for the Intel 286 processor and released in December 1987. The format is now out of date. A 32-bit revision of the format, the Enhanced Metafile, incompatible with WMF, was developed later for Microsoft's 32-bit applications, but the earlier format remains popular, common, and fully supported by Microsoft.

The problem discovered in December was with one of the printing APIs. Specifically, it was in the Escape function, which works directly with a device, usually a printer. It can, for example, start a print job, set the number of copies, and end the print job.

One of the subfunctions of Escape, SetAbortProc, was the source of December's problem. This subfunction allows the developer to set the Abort function for a print job. It has been deprecated for more than a decade, as have most of the printer escapes. It was replaced by a function with the same name, but has been retained strictly for backward compatibility with 16-bit versions of Windows, that is, with Windows 3.11 and its predecessors. An important quality of SetAbortProc is that it while its purpose was to refer to a printer, it does not have to. It can refer to another device instead.

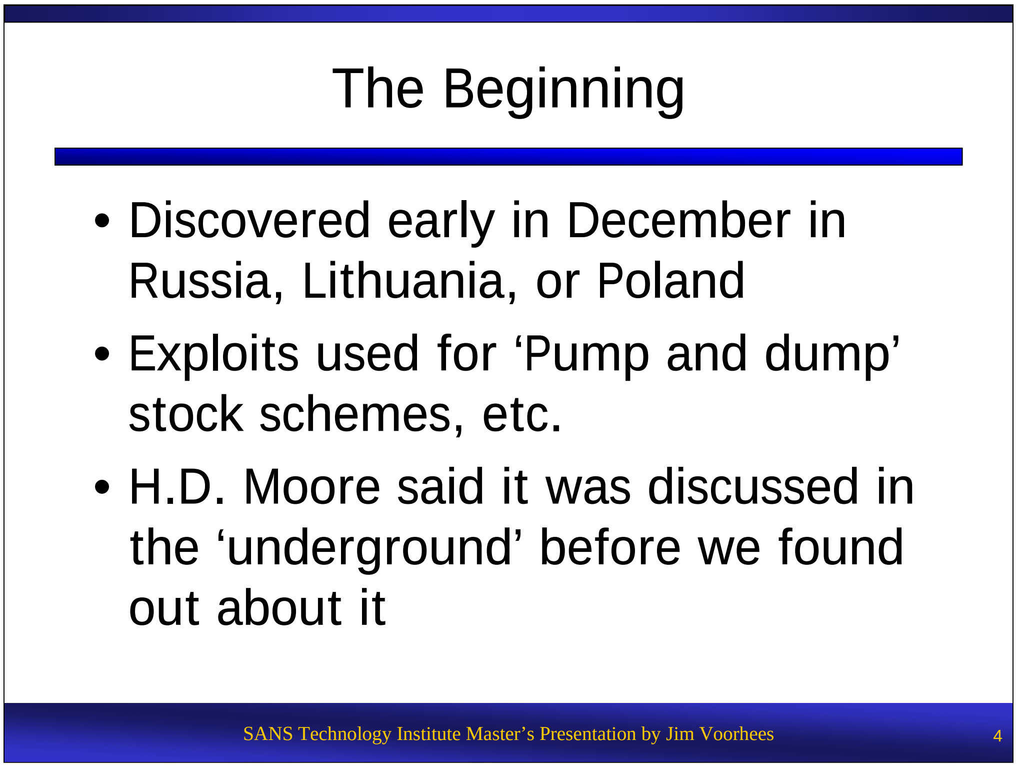

The Beginning

- Discovered early in December in Russia, Lithuania, or Poland
- Exploits used for ‘Pump and dump’ stock schemes, etc.
- H.D. Moore said it was discussed in the ‘underground’ before we found out about it

We do not know who discovered the WMF vulnerability, or when. There is some indication that a Russian found it within two or three days of December 1. Evidence points to a Russian who had put his license plate number in the code of an early exploit. Other evidence points to a Polish educational site. Still other evidence leads to the iframecash group in Russia and Lithuania.

The vulnerability was being exploited by the middle of the month. Unusual activity was seen on Russian hacker sites about this time. It appears that no one outside Russia picked up the portent of that activity. Within Russia, however, it appears that several groups were selling exploits on their web sites for $4,000.00.

It was found out later that the site beehappyy.biz hosted an exploit of WMF that generated spam messages promoting the stock of a Chinese pharmaceutical firm. A spike in the price of the stock of this firm on or about December 15 is one more suggestion that exploits of this vulnerability were extant two weeks before the West discovered them.

Rumors about the vulnerability and its exploits were mentioned on some news sites and among what H.D. Moore terms the “underground”, which he describes as “a loose network of friends that talk about vuln info.” That group, it seems clear, crosses the border between whitehats and black, much like the annual DefCon conference in Las Vegas does. But the whitehats missed the discussion or at least did not understand what it meant.

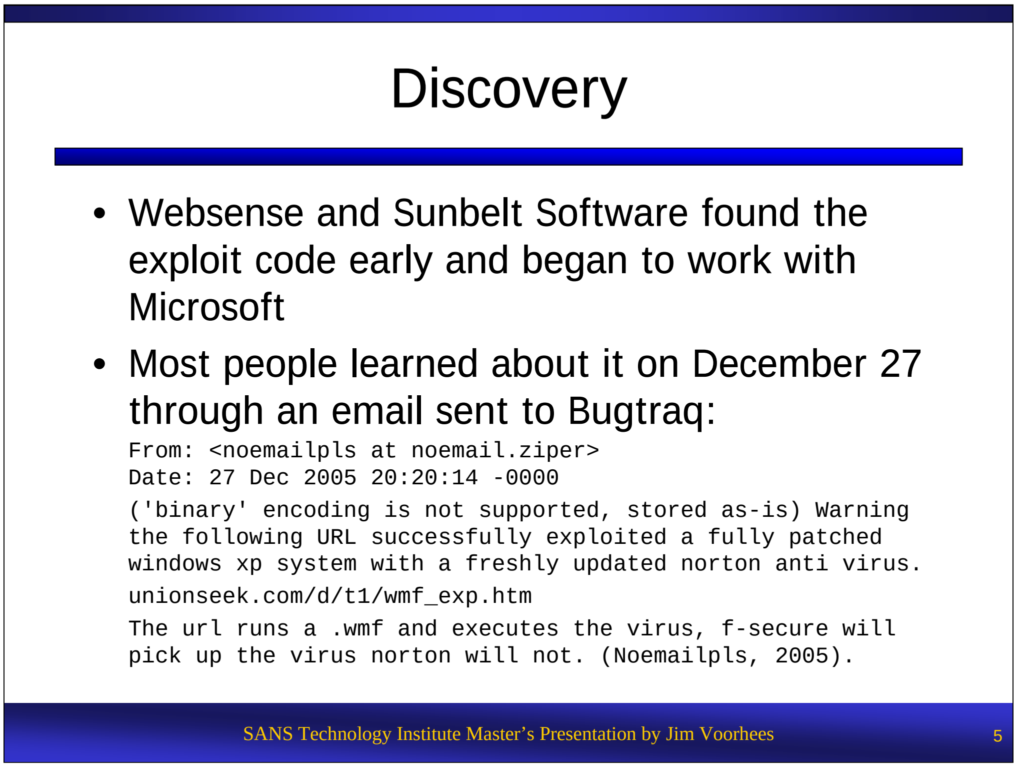

Discovery

- Websense and Sunbelt Software found the exploit code early and began to work with Microsoft
- Most people learned about it on December 27 through an email sent to Bugtraq:

```
From: <noemailpls at noemail.ziper>
Date: 27 Dec 2005 20:20:14 -0000
('binary' encoding is not supported, stored as-is) Warning
the following URL successfully exploited a fully patched
windows xp system with a freshly updated norton anti virus.
unionseek.com/d/t1/wmf_exp.htm
The url runs a .wmf and executes the virus, f-secure will
pick up the virus norton will not. (Noemailpls, 2005).
```

Before Christmas, Websense Security Labs found iFRAME websites with exploits that infected fully patched versions of Windows and Internet Explorer, without the need for the user to do anything other than visit the site. It seemed that the exploit worked on a vulnerability that had already been patched.

Presented with this mystery, Websense began asking around. On or about December 26, Dan Hubbard posted a message about it on a vetted email list. Websense also informed Microsoft.

Sunbelt Software found the exploit on the Mega Man comic book site at about 5:00 pm EST on December 27 and notified Microsoft within a few hours. They sensed that this was not normal malware: Most of the websites that used the exploits to install malware were not the normal bad websites (like porn sites). Most of the websites belonged to companies that sold things like real-estate.

After working through the night, Alex Echelberry posted the discovery on the Sunbelt Blog. They shared the information with other whitehats, the owners of the infected websites, and their own developers. They also wrote Snort signatures for their own Kerio firewall.

The message on Bugtraq was the first word about the problem that many received. With it, the existence of the vulnerability and at least one exploit was public. Blackhats and whitehats now began to race.

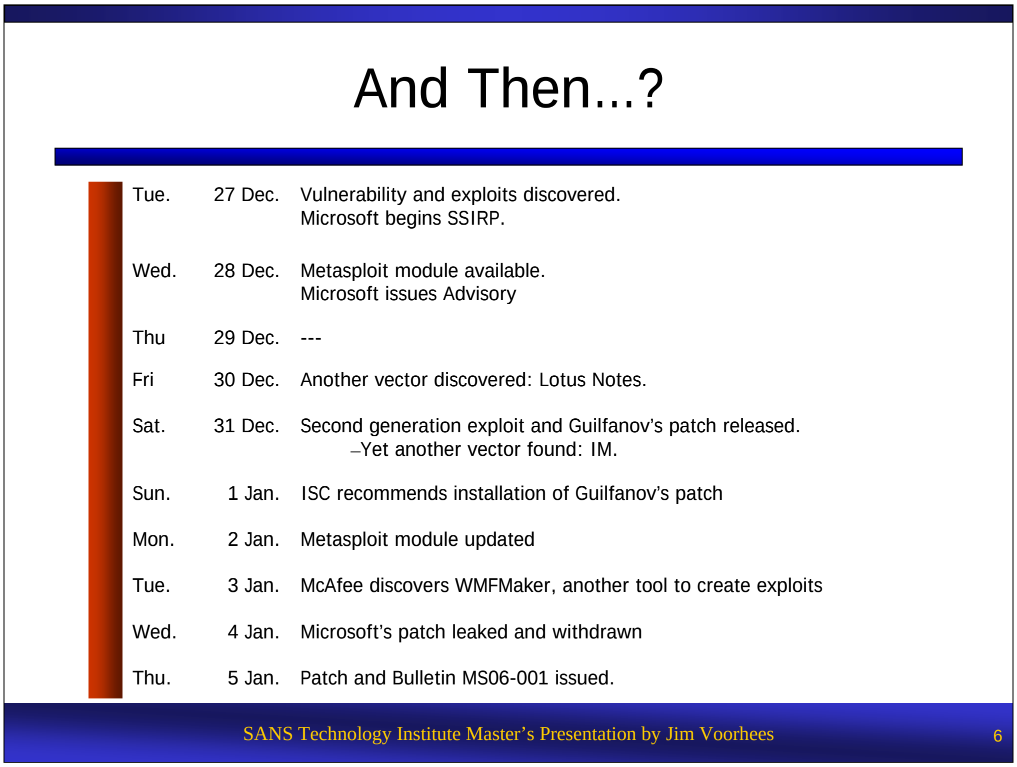

And Then...?

Tue.	27 Dec.	Vulnerability and exploits discovered. Microsoft begins SSIRP.
Wed.	28 Dec.	Metasploit module available. Microsoft issues Advisory
Thu	29 Dec.	---
Fri	30 Dec.	Another vector discovered: Lotus Notes.
Sat.	31 Dec.	Second generation exploit and Guilfanov's patch released. –Yet another vector found: IM.
Sun.	1 Jan.	ISC recommends installation of Guilfanov's patch
Mon.	2 Jan.	Metasploit module updated
Tue.	3 Jan.	McAfee discovers WMFMaker, another tool to create exploits
Wed.	4 Jan.	Microsoft's patch leaked and withdrawn
Thu.	5 Jan.	Patch and Bulletin MS06-001 issued.

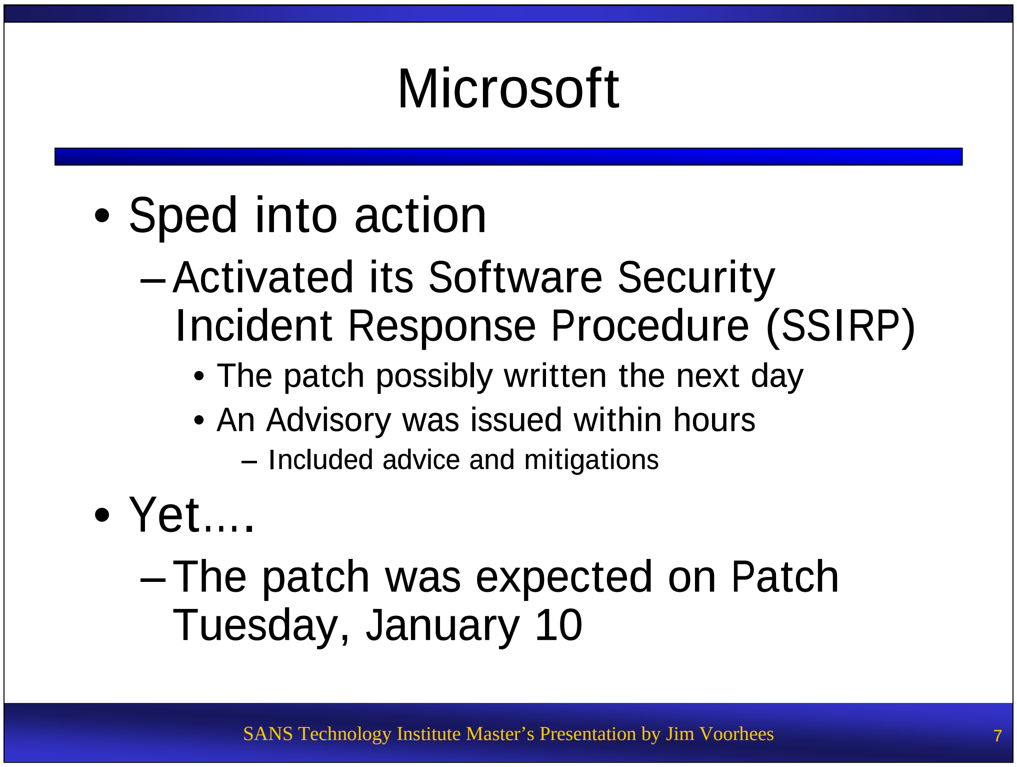

Microsoft

- Sped into action
  - Activated its Software Security Incident Response Procedure (SSIRP)
    - The patch possibly written the next day
    - An Advisory was issued within hours
      - Included advice and mitigations
- Yet….
  - The patch was expected on Patch Tuesday, January 10

SANS Technology Institute Master's Presentation by Jim Voorhees 7

Microsoft learned about the vulnerability on Tuesday night and activated the Software Security Incident Response Procedure (SSIRP). This is the company's rapid incident response procedure, invoked when the threat is immediate and severe.

By Wednesday morning, the technical details of the attack were confirmed and Microsoft "immediately began developing a security update...on an expedited track." Teams began to work on the update 24 hours a day.

An advisory issued on December 28, Wednesday, named Microsoft operating systems from Windows 98 through Windows Server 2003 as vulnerable. It was filled with advice for those who used Microsoft products, such as keep antivirus software up-to-date, visit the Windows Live Safety Center, "follow safe-browsing best practices," and "exercise caution" with email. The advisory also described four mitigating factors. Two noted that the attacker could not succeed without an action by the user, either visiting a "malicious web-site" or opening up an infected email message or attachment.

While suggestions in the advisory reduced the risk of infection, the risk would still remain significant, even for experienced users. What was really needed was a patch. Microsoft recognized this; it was working on it. Furiously. If everything went as expected, the patch would appear on the next Patch Tuesday, January 10.

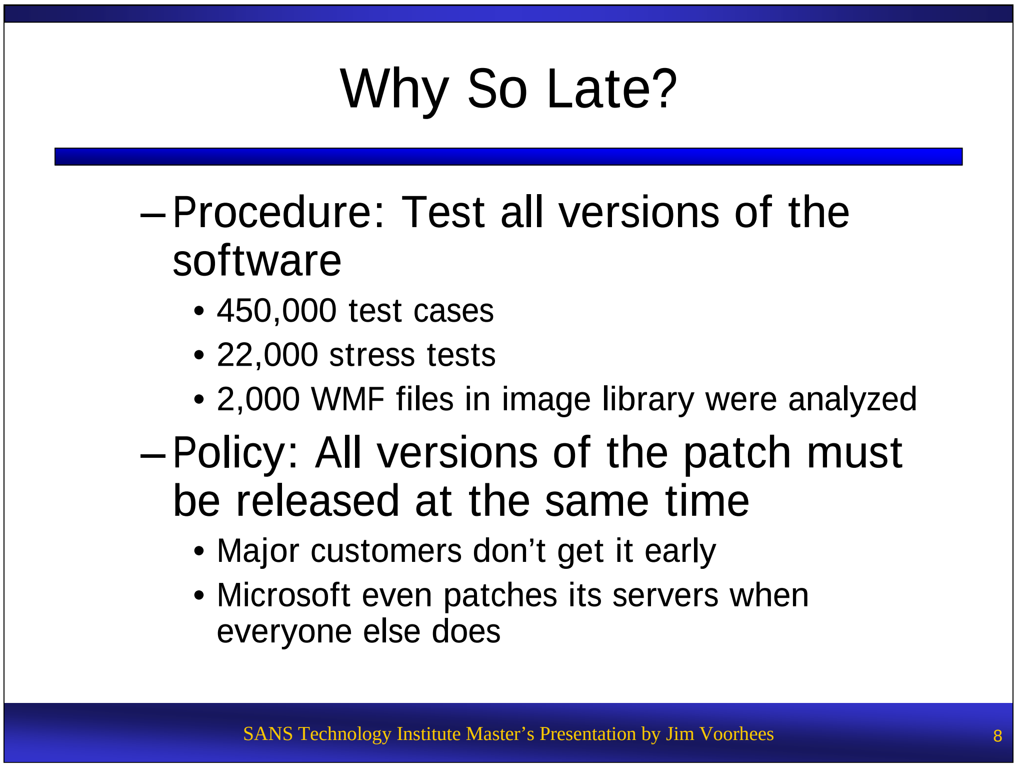

Why So Late?

- Procedure: Test all versions of the software
  - 450,000 test cases
  - 22,000 stress tests
  - 2,000 WMF files in image library were analyzed
- Policy: All versions of the patch must be released at the same time
  - Major customers don't get it early
  - Microsoft even patches its servers when everyone else does

SANS Technology Institute Master's Presentation by Jim Voorhees 8

Microsoft did react quickly, springing into action almost immediately . Indeed, there are indications that the patch itself was completed the next day.

But that was when the time-consuming work began. The test team ran the patch through an exhaustive, perhaps also exhausting, series of tests. There were more than 400 applications tested on the six Windows platforms that Microsoft then supported (Windows Server 2003, Windows XP, Windows 2000, Windows 98, Windows 98 SE, and Windows ME). The team tested versions in all 23 languages that Microsoft's software appears in. They ran through more than 450,000 test cases, subjected the patch to 22,000 stress tests, analyzed 2,000 WMF files from Microsoft's image library, verified that more than 125 malicious WMF files were fixed, and verified 15,000 printing-specific variations and 2,800 pages

Microsoft releases the final versions of the updates for all versions of its software, including the version in all 23 languages, simultaneously. The company has even refused requests from major customers with compelling cases to get patches early. Yet, the company does not even update its own servers until after 10:00 on the second Tuesday of the month. This means that Microsoft completes its tests on all versions of the patch before any version is released.

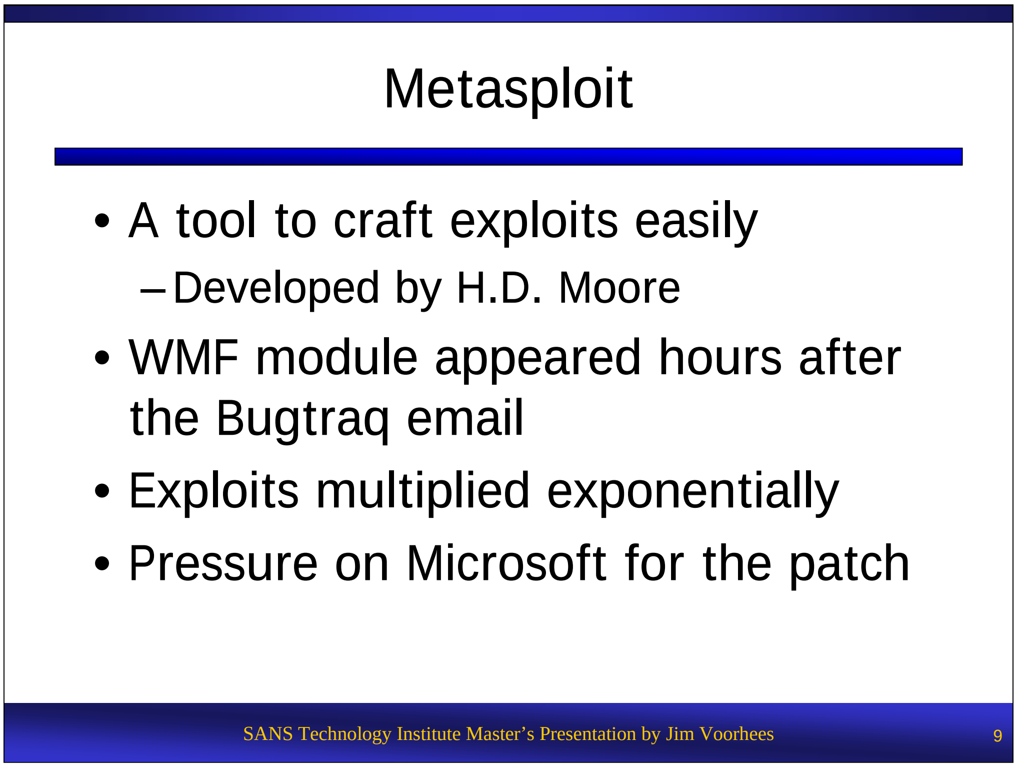

Metasploit

- A tool to craft exploits easily
  - Developed by H.D. Moore
- WMF module appeared hours after the Bugtraq email
- Exploits multiplied exponentially
- Pressure on Microsoft for the patch

Early on December 28, hours after the Bugtraq email came out, H.D. Moore announced in a reply to the original Bugtraq message that he had issued a module for WMF as a part of the Metasploit Framework, tested on Windows XP. There can be little doubt that the module multiplied the number of exploits of the vulnerability exponentially and quickly. Metasploit is designed to make it easy to create an exploit, after all, so that anyone can create and test it.

The module made users and sysadmins more vulnerable to exploits of the vulnerability. Given the multiplication of exploits, that cannot be argued. Sysadmins who had the time and expertise to take advantage of the module and in so doing improve their own defense would have been helped. But it is doubtful that many sysadmins, even those with the skill to do so, would have had the time. They had other priorities. Moreover, many of them would have had more faith in their security vendors than in H.D. Moore.

It is hard not to conclude that a primary target of the module, in fact, was other whitehats: those at Microsoft and the vendors of antivirus software, IPSs, and other security products. The module pushed them to create defenses rapidly. Did the push have a significant effect? Almost certainly.

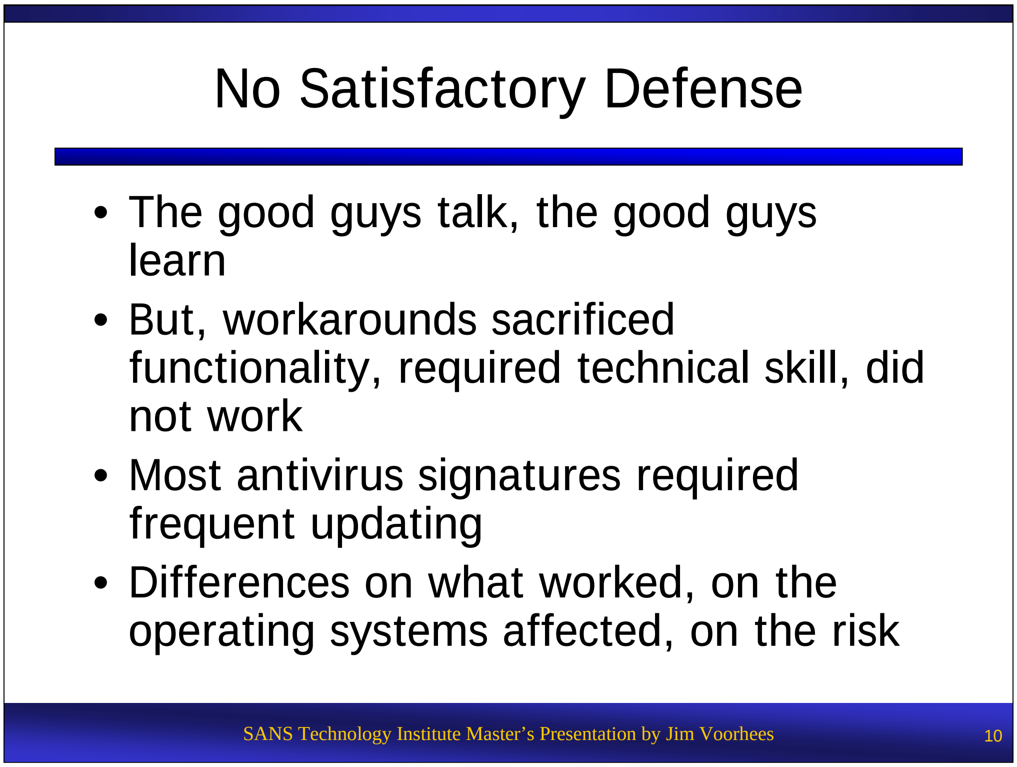

No Satisfactory Defense

- The good guys talk, the good guys learn
- But, workarounds sacrificed functionality, required technical skill, did not work
- Most antivirus signatures required frequent updating
- Differences on what worked, on the operating systems affected, on the risk

SANS Technology Institute Master's Presentation by Jim Voorhees 10

As the week went on, concern, and the public expression of it, was widespread. It grew as the number of exploits developed by the blackhats did. Sometimes new defenses were found or developed. Sometimes old ones were found to be ineffective. New exploits and new vectors for exploits were exposed.

Some workarounds reduced functionality; the user needed to have another way to see media files. One, deregistering the shimgvw.dll, required the user to have technical skills, that most lacked. Moreover, it was learned that it was ineffective: there were ways of exploiting the vulnerability that did not use the dll.

Antivirus vendors came out with signatures as early as December 27th. Some were effective; some became outdated when the bad guys made simple changes in their exploits.

Disagreement was common. There was disagreement over whether one defense, Data Execution Prevention, was effective in part or at all. The estimates of vendors and other whitehats about the risk that users and sysadmins faced differed widely. SANS, Secunia, and ultimately Microsoft, among others, told users that the risk was high, the vulnerability was critical. McAfee and Symantec, on the contrary, found the risk from the vulnerability low, even in advisories published before Microsoft's patch was issued. Finally, there was disagreement over which Microsoft operating systems were vulnerable. Was Windows 98? Was DOS?

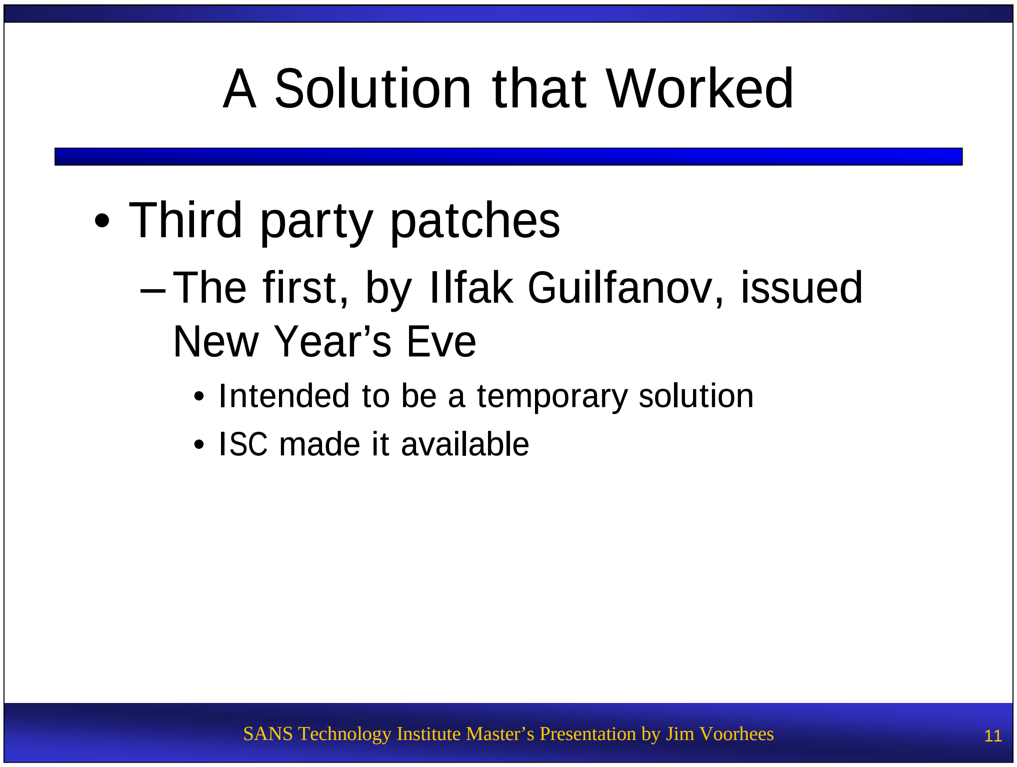

A Solution that Worked

- Third party patches
  - The first, by Ilfak Guilfanov, issued New Year’s Eve
    - Intended to be a temporary solution
    - ISC made it available

SANS Technology Institute Master’s Presentation by Jim Voorhees 11

New Year’s Eve, Ilfak Guilfanov issued his own patch, the first but not the last to be issued. This was no ordinary code jockey, which made his patch all the more important. He was the architect and main developer if IDA Pro, a disassembler and debugger used widely by whitehats to analyze malcode.

This patch was modest: it took a modest effort on his part and was offered modestly—complete with source code and an uninstaller, so the patch could be reviewed before installation and removed if it was found harmful or after Microsoft's patch appeared. Indeed, Guilfanov consistently advised those who installed his patch to replace with the official patch when they could. Tom Liston of the ISC reviewed the code in detail, published his analysis, and helped to modestly extend its capabilities.

On a weekend when blackhat exploits seemed to be running rampant and Microsoft's patch seemed distant, the patch seemed to offer an important means of defense. The ISC saw the situation as serious enough to warrant an extraordinary measure. Over Friday and Saturday, going into Sunday, the “rag-tag group of volunteers” as Liston called them, analyzed the new exploit, the risk it created, and the possible ways of defending against it.

With Guilfanov's permission, the ISC began to make his patch available. It was available on his own site as well. It soon became more popular than Guilfanov or anyone else anticipated. So popular that Guilfanov had to move the patch to a different site (set up by Castle Cops).

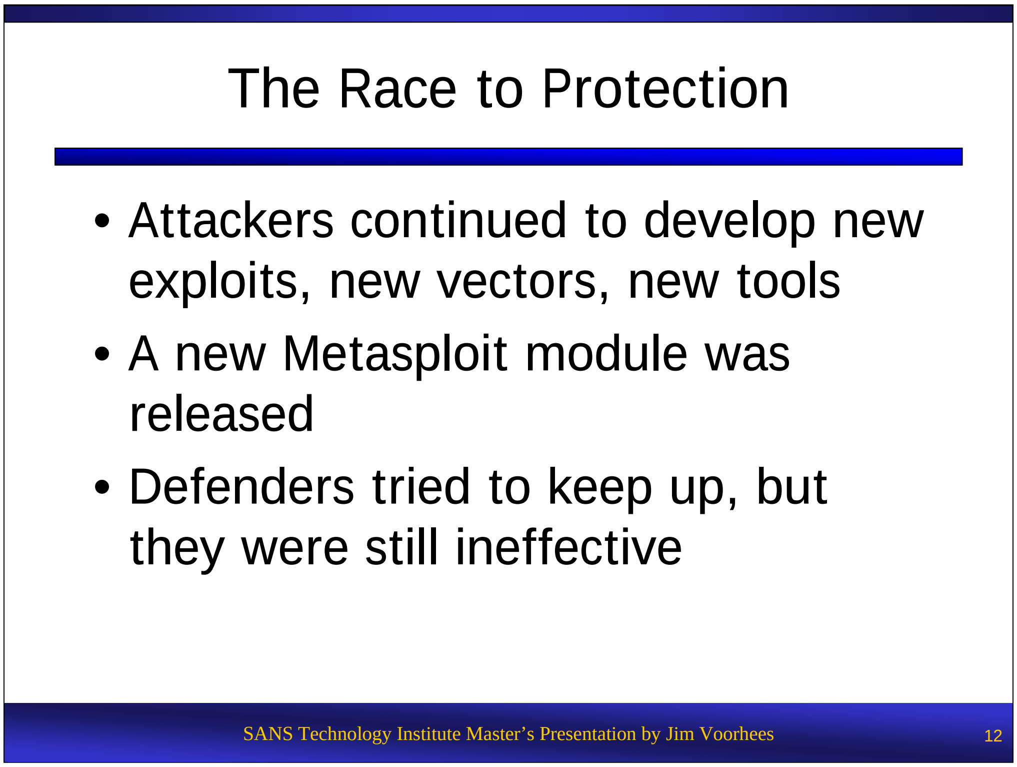

The Race to Protection

- Attackers continued to develop new exploits, new vectors, new tools
- A new Metasploit module was released
- Defenders tried to keep up, but they were still ineffective

SANS Technology Institute Master's Presentation by Jim Voorhees 12

The first days of the week brought little that was new, with the problems of the previous week extending into the new.. Neither the blackhats nor the defenders opposing their efforts rested. iDefense, for one, by Monday saw the actions taken by antivirus companies as insufficient. Few were detecting the new exploits of the vulnerability. For example, what the ISC termed a second generation exploit, released on New Years Eve, was detected by only three antivirus programs.

On Monday, H.D. Moore informed the ISC that The Metasploit Project had issued a new version of its WMF module that, he said, was designed to bypass all to bypass all known IDS signatures.

That same day, Panda Software discovered a new application to develop WMF-based malware, WMFMaker. On Wednesday, another group—the Ready Rangers Liberation Front—began a competition to create a WMF worm, that is, a WMF-based exploit consisting of "shellcode that replicates itself." The blackhats did not intend to let the problem fade away.

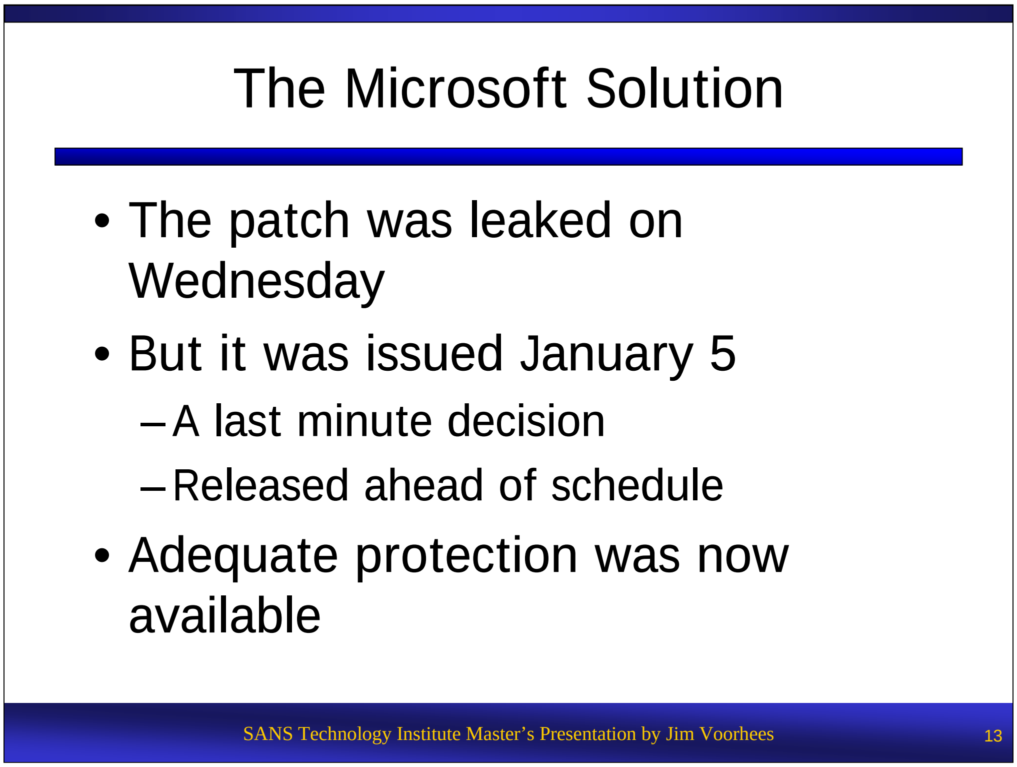

The Microsoft Solution

- The patch was leaked on Wednesday
- But it was issued January 5
  – A last minute decision
  – Released ahead of schedule
- Adequate protection was now available

Wednesday, as cries for an official patch continued, a patch was leaked from Microsoft, becoming available through several sources. It was tested by some who got it. It was good. The time had almost come. But that same day, Microsoft reaffirmed that it would not appear until Patch Tuesday.

It was finally released the next day, ahead of schedule. This was a last-minute decision, made the afternoon after testing was completed.

No official announcement was made until after 3:00 EDT on Thursday, January 5. But, interestingly enough, news was leaked to Ron Trent’s myITforum blog that Microsoft would release the patch later that day. At that time there was nothing about it on Microsoft's home page. The home page of the Security Center had the next release set for the next afternoon, but there was no other indication that something was about to come out. A Security Headline on the WMF vulnerability still had the old release date (10 January). Finally, however, the official announcement was made.

With the release of the patch, a satisfactory defense was now available.

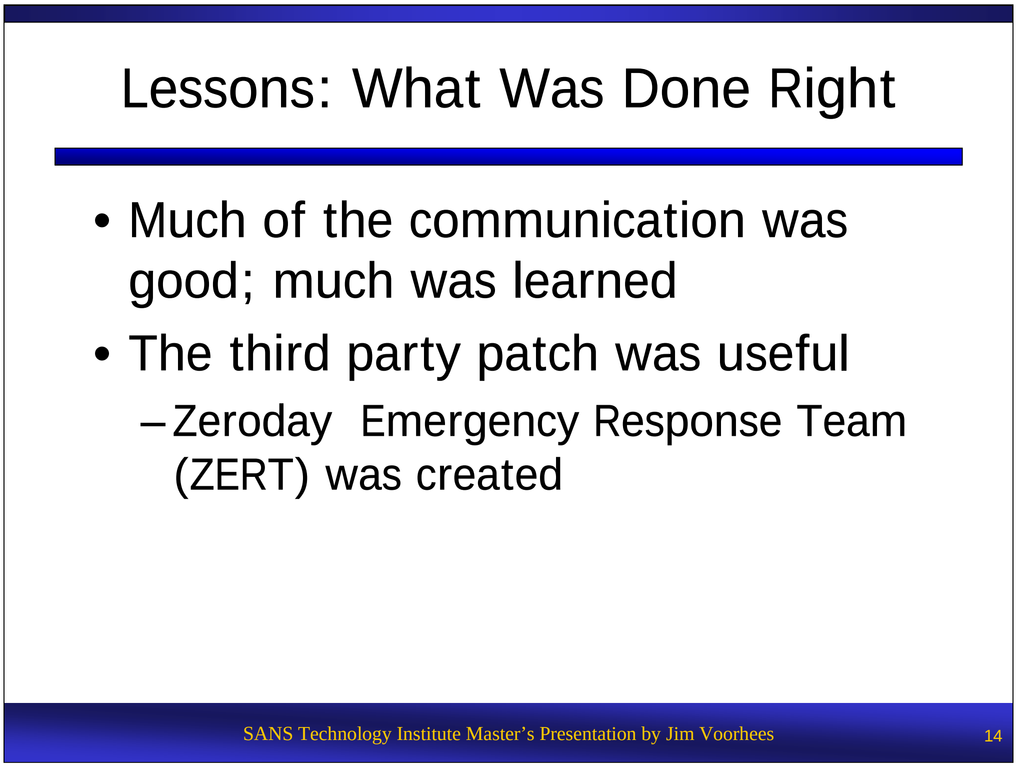

Lessons: What Was Done Right

- Much of the communication was good; much was learned
- The third party patch was useful
  - Zeroday Emergency Response Team (ZERT) was created

SANS Technology Institute Master’s Presentation by Jim Voorhees 14

One of the remarkable things about the episode was the amount of communication among the whitehats. Using blogs and e-mail lists, the telephone and, who knows, perhaps old-fashioned face-to-face meetings, once word of something came out somewhere, it came out everywhere. Moreover, this was an international effort, with FrSIRT in France, Secunia in Denmark, and F-Secure, in small, distant Finland cited prominently and often. The ISC itself spans the globe, with handlers in Belgium and Brazil, among other places.

The user or sysadmin who read the alerts and advisories available would have found as much protection available as could be had in the absence of a patch from Microsoft.

One valuable result of the WMF episode was that the usefulness of the Guilfanov patch led to the creation of the ZeroDay Emergency Response Team (ZERT; http://zert.isotf.org), which can make a third party patch available before the vendor can.

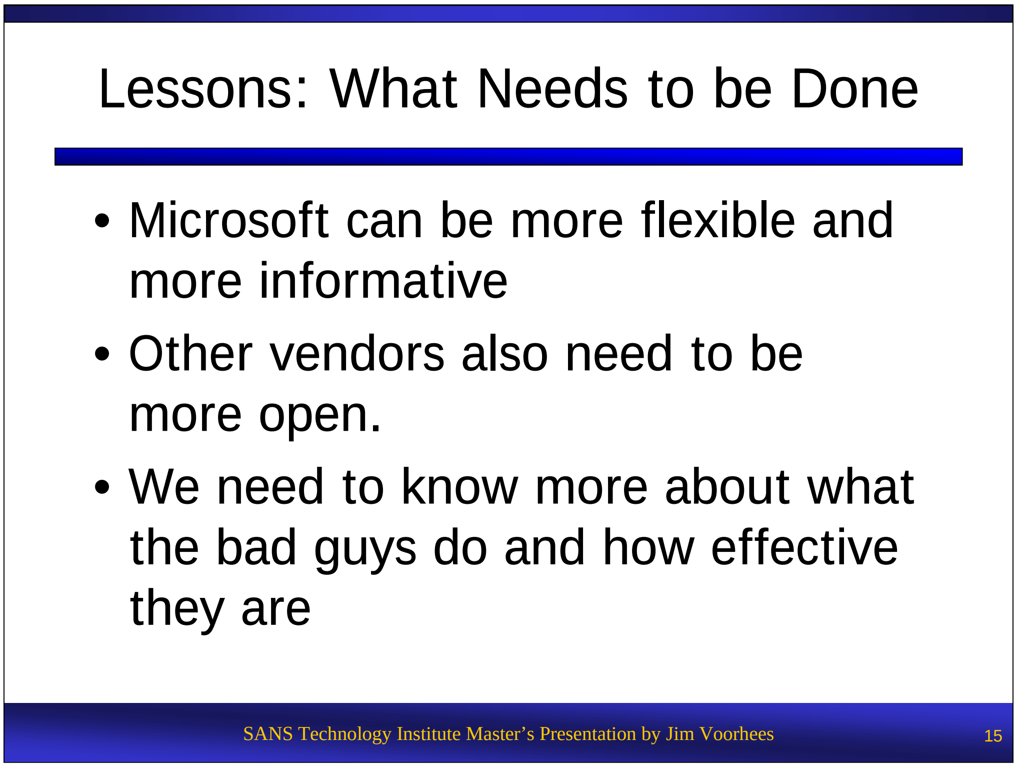

Lessons: What Needs to be Done

- Microsoft can be more flexible and more informative
- Other vendors also need to be more open.
- We need to know more about what the bad guys do and how effective they are

Microsoft needs to make its procedures more flexible. There are cases where a 10-day delay can be disastrous. It could introduce shortcuts, a beta patch, for example.

Second, Microsoft needs to be more forthcoming with information. The initial security advisory was incomplete. For example, the advisory did not describe the vulnerability in enough detail for the source of the vulnerability to be known. This hampered the defense.

Third, Microsoft should have known more about the vulnerability before the bad guys found it. A review of the code could have revealed the problem.

Microsoft was not the only firm that could have told more. Google Desktop was named as a vector for exploits on the first day. Its users have yet to hear from the company. In certain conditions Firefox became vulnerable, yet there was no word from the Mozilla project. CISCO issued no advisory, though CISCO products could have helped to mitigate the problem.

Finally, we need more data about what the bad guys can do and more information about what they do do. The security community has no source that can tell us, with confidence, how many computers were infected. In addition, the bad guys were able to operate under the radar. There are too few good guys are willing or able to communicate with them on their own turf in their own language. The information is available, we need some way to get it and make it available.

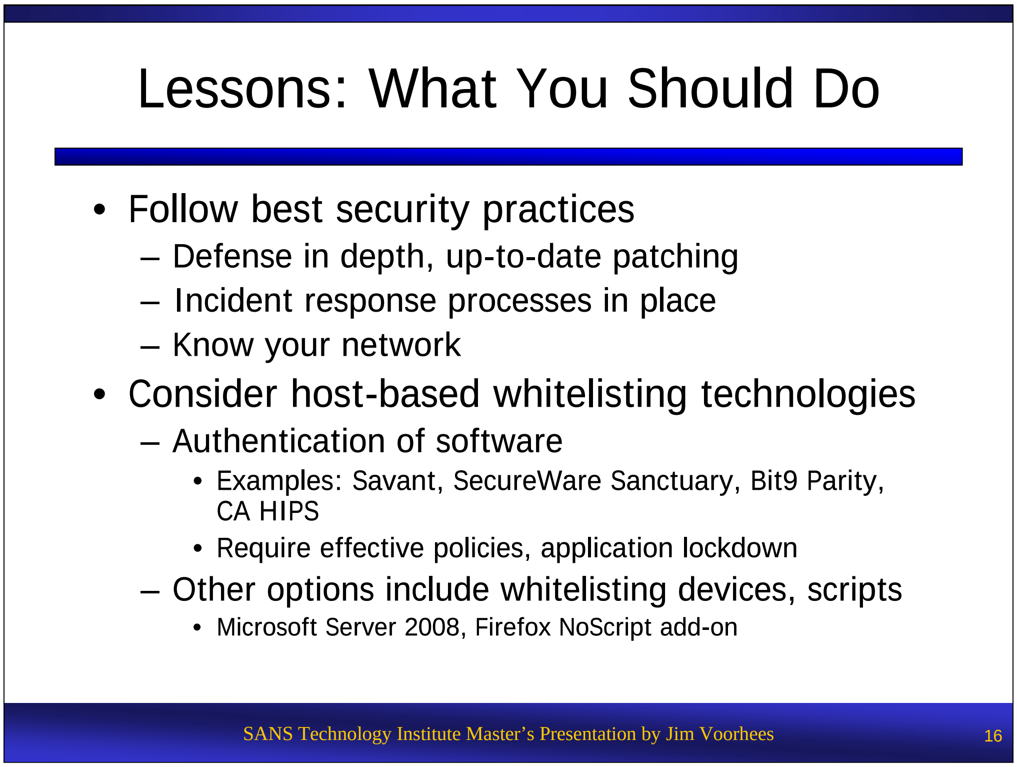

Lessons: What You Should Do

- Follow best security practices
  - Defense in depth, up-to-date patching
  - Incident response processes in place
  - Know your network
- Consider host-based whitelisting technologies
  - Authentication of software
    - Examples: Savant, SecureWare Sanctuary, Bit9 Parity, CA HIPS
    - Require effective policies, application lockdown
  - Other options include whitelisting devices, scripts
    - Microsoft Server 2008, Firefox NoScript add-on

The best thing you can do to prepare for a zero-day attack is to follow best security practices. See NIST publications or, of course, the lessons taught at SANS to find out what they are. Adopt defense in depth, including up-to-date firewall rules are router ACLs. Install and monitor IDSs or IPSs on your internal network. Put antivirus on workstations and servers. HIDS can be quite effective. Harden all machines. Of course, keep patches and signatures up to date

It it important to have incident response processes in place. Something will break through sometime—you need to know how to handle it

It is vital to become aware of what goes on both inside and outside your network. Know your network, document it, watch it. Develop a baseline of network activity so that you know when something is wrong.

The WMF episode showed, again, that signature-based protection is ineffective. Whitelisting—that is, allowing only the known—seems to be the best alternative. There are several things to focus on: devices, scripts, and applications. The last is fairly new, but promising. Doing it effectively requires strong, effective policies and application lockdown, forbidding users access to unapproved applications.