



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
"Hacker Tools, Techniques, and Incident Handling (Security 504)"
at <http://www.giac.org/registration/gcih>

**The December Storm of WMF: Preparation,
Identification, and Containment of Exploits**

GCIH Gold Certification

Author: James Voorhees, voorhees.james@gmail.com

Adviser: John Bambenek

Accepted: October [??], 2006

Table of Contents

1. Introduction: The Problem 4

2. The Metafile 5

3. The Actors 8

 3.1. Users 10

 3.2. Sysadmins 11

 3.3. Whitehats 12

 3.4. Blackhats 14

4. Preparations 15

 4.1. Whitehat Defenses 17

 4.2. Microsoft and Incident Response 18

5. Discovery 22

 5.1. Websense 23

 5.2. WMF Becomes Public 25

6. Reaction 27

 6.1. The Metasploit Module 28

 6.2. Other Early Whitehat Responses 31

The December Storm of WMF

| | | |
|-------|---|----|
| 7. | After the First Day: Growing Defense, Growing Confusion | 34 |
| 7.1. | Data Execution Prevention | 34 |
| 7.2. | Whitehat Confusion | 36 |
| 8. | Going into the New Year: New Vectors and Defenses | 39 |
| 8.1. | Guilfanov's Patch | 40 |
| 8.2. | The Next Week: Alarums and Excursions | 44 |
| 9. | Microsoft's Patch: Written, Tested, Released | 45 |
| 10. | Conclusions | 47 |
| 10.1. | Communication | 47 |
| 10.2. | A Learning Curve | 49 |
| 10.3. | The Audience | 50 |
| 10.4. | Better Data | 51 |
| 10.5. | Metasploit | 52 |
| 10.6. | Microsoft | 53 |
| 10.7. | Under the Radar | 56 |
| | References | 59 |
| | Author's Note | 74 |

1. Introduction: The Problem

It would be too much to say that all was calm over the Christmas weekend in 2005. All the same, Deborah Hale, a handler at the Internet Storm Center, found it so quiet on 27 December that she speculated that "Perhaps all of the script kiddies got new computers for Christmas and haven't gotten fully up to speed yet." (Hale, 2005). 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 (Symantec, March 2006; Websense, January 5, 2006). 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.

These events raise questions about how prepared the security community is to contain the effects of a zero-day exploit. This paper will look how and when the vulnerability was discovered and made public. It will identify the actors in the security community and examine how they responded. These actors include Microsoft, the vendor of the vulnerable product, of course, but also anti-malware companies like Internet Security Systems (ISS), Symantec, McAfee, and F-secure, and non-commercial analysts like the volunteers of the Internet Storm Center (ISC). The central questions it will try to answer are how the security community—the whitehats—responded to the

vulnerability and its exploits before Microsoft issued its patch, and where did that response leave users?

The distinction between exploit and vulnerability is essential to this paper, but too often confused. A vulnerability is "A flaw or weakness in a system's design, implementation, or operation and management that could be exploited to violate the system's security policy" (SANS 2006). One definition of exploit is "In computer security, an unethical or illegal attack that takes advantage of some vulnerability." (PCMag, 2006). The key part of this distinction is the relationship between the two: a vulnerability exists in an application; an exploit attacks it. In the case of WMF, the vulnerability had existed for years. It only came to prominence when it was exploited.

2. The Metafile

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 (McNamara, 1996). 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 vulnerability discovered in December 2005 was neither the first nor the last vulnerability found in the format. Indeed, by one estimate, at least 22 functions in the format contain vulnerabilities (Ferrie 2006).

WMF files contain both vector data and bitmap data. They store data for vector graphics using commands from the Microsoft Windows Graphic Device Interface (GDI), which lets the

application interact with the hardware.¹ The files are created and played back in memory, though they can be written to disk if they get too large.

Each file contains a header, followed by the records themselves. The header contains a description of the records in the file. These records are binary-encoded GDI function calls that render an image, sending it to a screen, a printer, or another output device. Microsoft says this about the security of GDI:

GDI generally has few security concerns because it deals with display rather than input. However, here are a few issues that you should consider.

Bitmaps, metafiles, and fonts are complex structures that could become corrupted. It is good practice to try to ensure that these items are uncorrupted and from a trustworthy source.

On Windows NT/2000/XP, an application can specify the security descriptor for some of the printing and spooling APIs. You should take care when setting the security descriptor (Microsoft, 2006a).

Indeed, the problem discovered in December has to do with the printing APIs. Specifically, it was in the Escape function,

¹ This description of the WMF format draws heavily on Liston, 2006. The descriptions found on the Microsoft Developers' Network (MSDN) were also essential. In addition, see McNamara, 1996; SANS, July, 2006, pp. 6-40 to 6-42; and Swan, 1993, pp. 111-130.

The December Storm of WMF

which "enables applications to access capabilities of a particular device not directly available through GDI" (Microsoft, 2006b). Rephrased, the Escape function works directly with a device, without going through the GDI. For example, it can start a print job, sets the number of copies, and end the print job. The Windows 3.x API has 64 escape sequence that could be used.

One of these, SetAbortProc, was the source of December's problem. This subfunction to Escape 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 (Ferrie, 2006).

Some applications cannot process Escape records. This includes Internet Explorer (IE). Other applications can, however, notably the Windows Picture and Fax Viewer. The Viewer converts WMF files into EMF records. In doing so, it processes Escape records, which can include SetAbortProc as a parameter (Toulouse, 2006).

Just as IE itself does not process SetAbortProc, so Windows 98 and Windows ME do not process it when it does not print directly to a printer. The code for WMF files is the same, but this quality of the interaction between application and operating system makes these older operating systems less vulnerable than Microsoft's newer ones (Toulouse, 2006).

The vulnerability lends itself to what one might call passive exploits. The victim had to come to the exploit; the exploit could not go to the victim. For example, e-mail was used to send infected files to potential victims, but the victim had to at least look at the message to get caught. In contrast, active exploits—worms—like Slammer and Code Red spanned the world in search of victims. This passive quality of the WMF vulnerability meant that the creator of an exploit also had to create a siren song—a way of attracting victims to the infected files. Many found this easy.

3. The Actors

Intuitively, security professionals divide the world into three parts: the whitehats, the blackhats, and the rest, otherwise known as users. We will look at these groups more formally to make it clearer who was affected by the vulnerability and its exploits.

In setting up the taxonomy that will divide all those who use computers into groups, we cannot use criteria that allow for shades of gray. Each of our criteria, therefore, will be dyadic, or binary, with only two possible answers.

Whitehats, blackhats, and users can be distinguished using two criteria. First, to paraphrase Bill Gates: "How technical are they?" Phrased more dully, this means: Is the group knowledgeable about computer technology? Second, does the group defend or attack computers?

There are many shades of technical knowledge. A 17-year-old

script kiddie may know less than the handlers at the Internet Storm Center, but he will know much more than your Aunt Minnie. Nonetheless, by generalizing broadly, we can say with confidence that whitehats and blackhats, both of whom deal with information technology professionally, know technology; users do not.

The question whether these groups defend or attack computers is more complex than it might seem. After all, penetration testers attack computers, but are usually counted among the whitehats. Blackhats may attack computers, but they are highly likely to defend their own. In a time when blackhats compete with each other over exploitable computers, their need for defense may be more common than we know.

All the same, the generalization that blackhats attack while users and whitehats defend is true enough that we can use it to distinguish between those two groups.

Using these two criteria, then, the groups can be distinguished as in Table 1.

Table 1—Groups and Two Criteria

| <i>Group</i> | <i>Technically Knowledgeable?</i> | <i>Attack or Defend?</i> |
|--------------|-----------------------------------|--------------------------|
| Users | No | Defend |
| Whitehats | Yes | Defend |
| Blackhats | Yes | Attack |

This is fine as far as it goes. But the world of information security surely contains at least one more group. They might be considered to be among the whitehats, because they know the technology and are essential to computer defense, but, like many users, security is not their central concern. I am

referring to systems or network administrators, sysadmins for short. A third criterion is being introduced here: Is information security the group's main concern? In discussing the WMF vulnerability, this may be an important consideration, because, like many things in security, defenses against exploits involve choices, which is to say that some of the defenses can impose a loss of functionality or increased risk if adopted.

We have, then, three criteria for dividing all those who use computers into groups. The following table shows how these criteria can distinguish the four groups that result:

Table 2—Groups and Three Criteria

| <i>Group</i> | <i>Technically Knowledgeable</i> | <i>Attack or Defend?</i> | <i>Primary Concern with Improving Security</i> |
|------------------|----------------------------------|--------------------------|--|
| Users | No | Defend | No |
| Sysadmins | Yes | Defend | No |
| Blackhats | Yes | Attack | Yes |
| Whitehats | Yes | Defend | Yes |

We have our groups distinguished from each other using the criteria described above. Now let us characterize them according to what we find in the real world. Just who are these people?

3.1. Users

Users are the people we meet every day for whom the computer is simply a means to another end. They use it to communicate it with others through e-mail, instant messaging, or internet relay chat. They surf the web. They create documents. They use their computer at work or at home. This is your Aunt

Minnie, Paul in Accounting, or Sanjay at the Internet café in Delhi. These are most of the billion or so people who use the Internet (United States, Central Intelligence Agency, 2006).

On average, they know less about information technology than people in the other three groups. They are not computer professionals, after all. They certainly do not launch attacks on other computers. They will defend their own if given the means. They have better things to do than to spend a lot of time on computer defense.

3.2. Sysadmins

Sysadmins here take care of more than one computer and they do it for someone else. The tasks they perform are varied. As one sysadmin put it:

The problem set in computing and network operations generally includes all those system tasks users might want to offload -- specification, evaluation, installation, configuration, integration, maintenance, data-integrity management, upgrade management, automation, security management, performance analysis, failure analysis, failure mitigation, recovery design, recovery implementation, testing, and more. (Dijker, 1999).²

That list of tasks is several years old. A new one would be similar, with security-oriented work more prominent.

² SAGE, 2006, is also useful. Dijker, 2006, based on an unscientific survey made in 1998, gives some idea about how sysadmins see themselves.

The precise tasks sysadmins carry out and the skills they have depend on several things:

1. the organization they are with, whether it is public or private, small or large,
2. their job responsibilities,
3. the number of computers they are responsible for, and
4. their background in the field.

On average, they know more about the technology than the users they serve, and they know more about their systems than anyone else. Sysadmins usually bear the primary responsibility for the security of their systems, but security is not their primary responsibility. Like users, they often have something better to do.

3.3. Whitehats

Whitehats are usually security professionals. They can be found in a variety of organizations, working as staff members, consultants, or, at places like the Internet Storm Center, as volunteers. The skills they have and the places they hold in the security universe can be varied. They find exploits, configure firewalls, monitor network traffic, conduct penetration tests, audit systems, develop security policies, and manage security departments. One survey estimates that there were 1.4 million security professionals throughout the world (Carey, 2005).³

³ Carey, 2005, does not make it clear how the estimate was derived. It includes "full-time and part-time information security professionals, practitioners, and other employees across a multitude of job titles." Given that the survey is based on a self-selected sample (people who chose to fill

There are divisions among them that have historical roots. For instance, as of 1999, people who work on anti-virus had different attitudes toward disclosure—a topic important to the tale of WMF—than those who work on computer security (Gordon and Ford, 1999). That was long ago in Internet time, but the differences remain. One was the attitude toward the disclosure of code and the technical details of exploits. Indeed, some of the people on one side of the debate placed some of their opponents in the camp of the blackhats.

This suggests a point that should be made. There is no clear divide between any of these groups. We have a continuum of skills and motivation rather than four clear cut blocks of people. The divide between blackhats and whitehats may be the cloudiest. For example, the motivations of “security researchers” are often difficult to divine.

It will be useful later to subdivide the whitehats to examine how groups with different responsibilities responded when the WMF vulnerability became widely known. However, for now, we can assume that all the whitehats are one, whether they work on anti-virus or something else, whether for Microsoft or the Metasploit Project. Our chosen criteria make it possible to distinguish between them and the other three groups. Whitehats are, therefore, as a group, technically skilled professionals dedicated to the protection of computers rather than the

out the survey on a website) and includes people with job titles that include network and systems administrator, it is likely that a significant proportion of that 1.4 million would be placed in one of our other groups, notably the sysadmin group. Nonetheless, the number given suggests in broad terms the number of people in the group.

exploitation of their weaknesses.

3.4. Blackhats

The blackhat, in stereotype, is an asocial American teenage boy who messes with computers in as much of the day as he can manage, working solo, trying to break into computers anywhere for the sport of it, to boost his "cred." Like many stereotypes, this one begs reality. It is becoming less accurate as the nature of the Internet and cybercrime changes.

A survey made more than a decade ago showed that

...virus writers are not a homogeneous group. They have characteristics similar to many populations. They vary in age, income level, location, social/peer interaction, educational level, likes, dislikes and manner of communication" (Gordon, 1994).

An article written to follow up that study two years later, as the visually oriented World Wide Web began to broaden the appeal of the Internet, affirmed the conclusions of the earlier work. It also posited the appearance of a new type of blackhat, older and more skilled, and a legitimization of virus writing, at least within certain communities (Gordon, 1996). Since then, blackhats have become better organized, so that groups ever more frequently act together. Motivations have, if anything, broadened. Political activism—hacktivism—is not uncommon; national capabilities for cyberwar have increased (Wilson, 2004); and the profit motive seems to be driving an increasing amount of malicious activity (The Honeynet Project, 2004, p.

511).⁴

Marcus Sachs wrote recently that there are six threat groups to which the blackhats belong (Parker, et al., 2004, p. 221):

- Nation states
- Terrorists
- Spies
- Organized crime
- Insiders
- Hackers

This list gives some sense of the variety of people at work discovering vulnerabilities and creating exploits. The stereotype teenager would belong to the last group, the most common, but least dangerous. Any of the rest could, conceivably, have discovered the WMF vulnerability and created exploits for it.

4. Preparations

Before examining what happened when the WMF vulnerability was discovered, it will be useful to look at what the three defending groups had available and their usual methods of dealing with exploits in general and malware in particular. These were the technologies and processes that were available for dealing with the exploits of WMF.

⁴ Chapter 16 of this book gives a fascinating profile of blackhats.

The December Storm of WMF

The technologies available to users and sysadmins have been discussed in great detail in many places. These include anti-virus and personal firewalls for users, and firewalls, network and host intrusion protection systems (NIPS and HIPS), and both network and host intrusion detection systems (NIDS and HIDS) for the networks managed by sysadmins. A few notes here will all that will be necessary to set the stage before the exploits are introduced.

First, most anti-virus programs, NIPS, and NIDS and are signature-based, so that they defend against exploits already known. Other techniques are used as well—heuristics and anomaly-based detection primarily—but signatures remain at the center of most defensive technologies. Confronted with exploits of an unknown vulnerability, the people who must update these signatures must play catch-up—there is a race between the blackhats and the whitehats that can only end when a patch for the vulnerability has been universally deployed.

This suggests another point. Like the first one, it is nothing new. It is simply this: no defensive measure is ever universally deployed. According to one set of data, out of 258 million households who access the Internet worldwide, only 51 million subscribe to antivirus (Updata Capital, 2004, p. 3).⁵ In other words, four out of five households that access the Internet do so unprotected.

⁵ The data come from Wachovia.

4.1. Whitehat Defenses

We argued above that all whitehats are one. Of course that is not entirely true when discussing a particular issue, such as the WMF vulnerability. Microsoft, the creator of the WMF format, in a sense owned the vulnerability. In part because of the proprietary nature of the code, no one else could develop a fully adequate patch. Everyone else's actions revolved around Microsoft's.

Other vendors sought defenses against the exploits that were developed. These were the companies that produce the antivirus and the intrusion protection and detection systems that are used by users and sysadmins: antivirus and the intrusion protection and detection systems.

In more general terms, it is useful for our purposes here to divide whitehats into those who provide defenses to users and sysadmins and those who apply them. The vendors mentioned above are in the first group.

This group also includes whitehats that have alerting others as their main task. The Internet Storm Center (ISC) (<http://isc.sans.org>) figures prominently here. So does the U.S. Computer Emergency Readiness Team (US-CERT) (<http://www.us-cert.gov/>), a part of the Department of Homeland Security, and the CERT Coordination Center (<http://www.cert.org/>) at Carnegie-Mellon University. They could not develop solutions, but they could provide information. A commercial counterpart is iDefense, a part of Verisign.

Groups within organizations can fall within this group as

well. I work on a project that monitors the network of a government agency. When an event takes place, we contact sysadmins, users, or both, and suggest how the problem can be resolved or mitigated, but we do not lay hands upon systems ourselves.

Whitehats who apply defensive measures—firewall administrators, for example—can be included among sysadmins for some purposes. Indeed, they often work side-by-side in an organization. When it comes to the trade-off between functionality and security, however, they will usually choose the latter.

Many of these organizations, vendors and alerting groups alike, have programs that monitor the Internet 24/7/365, with established processes for reacting to the discovery of a new vulnerability or exploit.⁶ These processes were tested when the WMF vulnerability was discovered in December. Because the vulnerability was in a Microsoft product, Microsoft's program is especially important.

4.2. Microsoft and Incident Response

Microsoft has a well-developed, formal incident response process, the Software Security Incident Response Process (SSIRP) (Howard and Lipner, 2006, pp. 187-214.⁷ It is centered on the

⁶ For a generic look at such programs and processes, see Hocutt, 2002, pp. 23-25. Hurley, 2003, focuses on TrendMicro. Evergeek focused on Symantec in Sapieha, 2006. Websense addressed its own processes in Websense Security Labs, 2005.

⁷ The SSIRP is described in Howard and Lipner, 2006, pp. 187-214 and in

The December Storm of WMF

Microsoft Security Response Center (MSRC), founded in Redmond in 1996, well before the company began to emphasize security in its products through the Trustworthy Computing initiative. The MSRC deals with all vulnerabilities discovered outside the company.

The MSRC is a monitoring organization. It watches the public email address, `secure@microsoft.com`, 24 hours, every day, all year, and usually provides an answer to any message within 24 hours. The center also monitors news lists, web sites, and other places where word of a new vulnerability or exploit might be found.

Microsoft has, in fact, two incident response processes. One, the Security Response Process (SRP), is for vulnerabilities that do not pose a severe, immediate threat. It is, therefore, more leisurely than its cousin, the Software Security Incident Response Process (SSIRP).⁸ The slower process has two parallel tracks. One is for the response center itself, which receives the report of the vulnerability, makes an assessment of its severity and likely impact, then informs the community of people affected, including, of course, customers and press about its conclusions and any mitigating actions that can be used. The other track is for a development team and a specialized security team. These are the people who create the fix and test it.

When the MSRC or another team identifies an event that

Microsoft, 2005, May.

⁸ Interestingly, Microsoft, 2005, May and other documents online, only refer to the emergency process. The description of the SRP comes from Howard and Lipner, 2006. Also see Andrew Cushman's slides for his talks on Microsoft Security Fundamentals (Cushman, 2006). There are also differences in the description of the SSIRP. We have relied on Howard and Lipner.

poses a "significant and near term threat" to users of its software, Microsoft initiates the SSIRP. It has four phases. In the Watch phase, a small group of "first responders" confirms that an incident is taking place. A team is then assembled as the Alert and Mobilize phase begins. This team has two main groups: one for engineering, the other for communications. The engineering team includes people responsible for the affected product. They work with the Secure Window Initiative Attack Team (SWIAT) to analyze the technical aspects of the incident. The communications team monitors the press and customer support lines.

In the next phase, Assess and Stabilize, the engineering team finds solutions that it can recommend to take care of the problem. The communications team gets the word out to customers, Microsoft's partners, the press, and, of course, Microsoft's sales and support staff.

In the last phase, Resolve, any necessary tools, updates, and information about how to recover from an attack are released and customers are made aware of them. The phase and the SSIRP end with a postmortem to derive any lessons that can be learned from the incident.

While the SSIRP cannot end before an update is released if one is required, the update process is separate. Most Microsoft updates are released by the MSRC, with an accompanying bulletin, at 10:00 Pacific Time on the second Tuesday of each month. This procedure was instituted in October 2003 in response to customer complaints that they wanted to be able to patch less frequently and more predictably (Microsoft, 2003). In other words, users

and sysadmins weighed the costs of frequent, unpredictable patching against the risk that they would be infected by an exploit of an unpatched vulnerability and found a regular patch schedule more compelling. In addition, Microsoft found that blackhats began the race to exploit vulnerabilities after an update was released, not before (Howard and Lipner, 2006, pp. 203-204). Therefore, paradoxically, perhaps, the risk of being infected by an exploit increased after the patch came out, at least if one's machine were left unprotected.

One other aspect of the update process should be mentioned. Microsoft releases the final versions of the updates for all versions of its software, including the version in all 23 languages, simultaneously. The company does not even update its own servers until after 10:00 on the second Tuesday of the month (Howard and Lipner, 2006, p. 204). This means that Microsoft tests all versions before any version is released.

The rationale for this is described by Howard and Lipner:

We've often discussed [this] policy with Chief Information Security Officers of major customers, many of whom believe that their organizations have a critical need to receive security updates or security bulletins before other customers. They make compelling cases, but on examination, it's simply impossible to develop a consistent rationale for giving some customers access to updates before others—you find yourself on a slippery slope at whose bottom everyone receives the updates early. (Howard and Lipner, 2006, p. 204).

The whitehats, then, were prepared for new exploits, new vulnerabilities. Their procedures were in place and practiced. And then, along came WMF.

5. Discovery

We do not know who discovered the WMF vulnerability, or when. There is some indication that some Russian found it within two or three days of December 1. Some evidence points to a Russian in Saint Petersburg who had put his license plate number in the code of an early exploit (Pupkin-Zade, 2006).⁹ Other evidence points to a Polish educational site. Still other evidence leads to the iframecash group in Russia and Lithuania.

It was being exploited by the middle of the month. Unusual activity was seen on Russian hacker sites about this time (see the timeline).¹⁰ 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.

⁹ H.D. Moore received emails from this Russian in nearly January. The Russian bragged that he had discovered the vulnerability and launched the exploit. He told Moore about the license plate. Moore gave that information to iDefense, who confirmed it. iDefense also used fuzzing analysis to find the Polish server. Websense found the connection to iframecash. Iframecash took full advantage of the vulnerability; they exploited about 150 websites. Their own site, iframecash.biz was on most of the early lists of infected sites (Ken Dunham, Director, Rapid Response Team, iDefense, telephone conversation with the author, 23 August 2006; Dan Hubbard, Senior Director of Security and Technology Research, Websense, telephone conversation with the author, 24 August 2005; H.D. Moore, email to the author, 25 August 2006; Dunham, 2006).

¹⁰ Pupkin-Zade seems to draw heavily on Gostev, 2006. Gostev is a senior virus analyst for Kaspersky Labs. Pupkin-Zade also cites Jim Melnick of iDefense as a confirming source.

It was found out later that the site beehappy.biz hosted an exploit of WMF that generated spam messages promoting Habin Pingchuan Pharmaceutical (PGCN), a Chinese pharmaceutical firm. PGCN's stock was over-the-counter stock listed on the Pink Sheets exchange. 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 (iDefense, 2006).¹¹

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" (H.D. Moore, email to the author, 25 August 2006). 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.

5.1. Websense

Before Christmas, Websense Security Labs found iFRAME websites with exploit code that infected fully patched versions of Windows and Internet Explorer, without the need for the user to do anything other than visit the site (Dan Hubbard, telephone conversation with the author, 24 August 2005; Websense Security Labs, Websense Security Labs 2006, January 2). It was a mystery; it seemed that the exploit worked on a vulnerability that had already been patched. Indeed, a number of other people confused

¹¹ For more on "pump and dump" scams like this, using data from Pink Sheets, see Frieder and Zittrain, 2006.

The December Storm of WMF

the new vulnerability with another vulnerability in WMF that had been patched the previous month.¹²

Presented with this mystery, Websense began asking around, trying to solve it. On or about December 26, Dan Hubbard posted a message about it on the vetted Malicious Websites and Phishing (MWP) email list. The next day they posted the following on their blog:

We are in the midst of researching a potential new IE vulnerability which appears to have not being fixed. The exploit uses a WMF file to run without user intervention. We discovered exploit code on a well-known .biz Spyware website. The Trojan Horse code downloads another file and runs it on the machine upon accessing the website. This vulnerability was supposed to be fixed with MS05-053, however our test machine is fully patched and the exploit code still functions." (Websense Security Labs, 2005).

Importantly, they also informed Microsoft. Websense and Microsoft worked closely together then and through the next ten days to understand the vulnerability and to get the patch out. As the quote suggests, Websense and others knew that there was something unusual creeping toward them from the wild, but they knew not what. All would soon become clear.

¹² For example, on December 29, one site published the code for the Metasploit module that came out two days earlier, but said that it corrected a buffer overflow described in Microsoft Security Bulletin MS05-053, which came out in November (<http://www.securiteam.com/exploits/5DP0I2KHHE.html>).

5.2. WMF Becomes Public

While Websense was making its discovery, numerous groups of whitehats were watching, waiting, and analyzing for any malware that might appear in the days after the holidays. Many others had gone on vacation; it was supposed to be a slow period.

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 (E-mail from Alex Eckelberry, President, Sunbelt Software, to the author, 21 August 2006.). They sensed that this was not what they regarded as normal malware:

The major issue was that most of the websites that were using the WMF exploit to install malware were not the normal bad websites like ... porn sites. Most of the websites were normal commercial companies selling real-estate or time shares, etc. ... (Email from Eric Sites, Vice President of Research and Development, Sunbelt Software, to the author, 21 August 2006).

It is worth noting here that when the Internet Storm Center discovered the vulnerability, the user who notified them had become infected through the Knoppix-STD.org website (Ullrich, 2006). This was no ordinary website, but the site for a bootable version of Linux designed for whitehats. The exploits of this vulnerability were not only catching the ignorant and foolish.

After working through the night, Alex Echelberry posted the discovery on the Sunbelt Blog. By that time, they had found more than a handful of websites that were using one particular

The December Storm of WMF

exploit (Echelberry, 2005, December 27). They also 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 Sunbelt Blog was not widely followed by whitehats, users, or sysadmins. So this message on Bugtraq was the first word about the vulnerability that many received:

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).

The time was probably GMT, which would make it about 5:00 pm Eastern time, the same time that Sunbelt was making its discovery.¹³ The author of the email to Bugtraq submitted a sample of the exploit to F-Secure at what appears to be a few minutes before the submission to Bugtraq. F-Secure notified Microsoft, both in Europe and Redmond, other antivirus companies, and Google. Its research had shown that Google Desktop would execute WMF exploits automatically once they were

¹³ Roger Grimes, a subscriber to the list, actually received it somewhat later, at precisely 7:31 EST (Grimes, 2006).

on the hard drive (Email from Mikko Hypponen, Chief Research Officer, F-Secure, to the author 21 August 2006.).

With this message, the existence of the vulnerability and at least one exploit was public. Blackhats and whitehats now began to race each other over the computers in the hands of users and sysadmins.

6. Reaction

By the morning of December 28th, the existence of the vulnerability and at least one exploit were known, but users and sysadmins had nothing to work with. Whitehats swung into action to learn more about what the blackhats were doing and to discover and create defenses against it.

As indicated above, Microsoft learned about the vulnerability on Tuesday night and activated the SSIRP.¹⁴ By Wednesday morning, the technical details of the attack were confirmed and Microsoft "immediately began developing a security update...on an expedited track." (Microsoft Corporation, 2005, December).¹⁵ Teams began to work on the update 24 hours a day (Vijayan, 2006).

The advisory named Microsoft operating systems from Windows

¹⁴ Andrew Cushman, Microsoft's Director of Security Community, says that Microsoft "first noticed" the vulnerability when it saw the Bugtraq post (Cushman, 2006).

¹⁵ The text of this version of the advisory is much fuller than the one currently found on the Microsoft website (Microsoft Corporation, 2006, January 5a). Much of the text was moved to Bulletin MS06-01, which was issued with the patch (Microsoft Corporation, 2006, January 5b).

98 through Windows Server 2003, Service Pack 1 as vulnerable.¹⁶ It was filled with advice for those who used Microsoft products as the authors tried to give them a broad range of ways to avoid getting compromised. The advisory encouraged users "to keep their antivirus software up-to-date," to use Microsoft Windows AntiSpyware (Beta), to visit the Windows Live Safety Center, and to use the "Complete Scan" option there. It said that users of Windows OneCare "...are already protected from known malware..." These suggestions were followed by advice: "follow safe-browsing best practices," "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.

The user who followed the suggestions in the advisory would reduce the risk of infection. But the risk would remain significant, even for experienced users. What was really needed was a patch. Microsoft recognized this; it was working on it. If everything went as expected, it was thought, the patch would appear on the next Patch Tuesday, January 10, which was almost two weeks away (Vijayan, 2006).

6.1. The Metasploit Module

Early on the day after the existence of the vulnerability was announced on Bugtraq, H.D. Moore announced in a reply to the

¹⁶ Notable exceptions were Windows NT and Windows 2000 before Service Pack 4. Security fixes were no longer being made available for them. Support for Windows 98 and its cousins had been extended through June 2006.

The December Storm of WMF

original Bugtraq message that he had issued a module for WMF as a part of the Metasploit Framework, tested on Windows XP (Moore, 2005).¹⁷ This was a matter of some controversy. 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. As Mikko Hypponen at F-Secure put it:

It enables clueless newcomers to easily craft highly variable and hard-to-detect variations of image files. Images that take over computers when viewed. And do this on all common Windows platforms (Hypponen, 2006).¹⁸

H. D. Moore gave his rationale for issuing the module so quickly after it became public in a post to the Full Disclosure list:

The vulnerability was being exploited, in the wild, for at least two weeks (based on email reports) prior to the original BT post. The WMF structure is widely documented. The AV vendors were providing less-than-capable signatures for no reason other than that no public code was available that demonstrated alternate encodings. The IDS vendors were (and some still are) providing signatures that couldn't survive a single legal byte change in the WMF header. The release of a "polymorphic" (not) exploit forced the vendors

¹⁷ The time of the post is 3:34 am. If this is GMT, that would be about 10:30 pm EST. In any case, it was merely hours after the original post.

¹⁸ Note that this posting was responding to an update of the Metasploit module

The December Storm of WMF

to either fix their products or cry "irresponsibility" and give up. IPS vendors realized how SOL they are wrt to client-side HTTP attacks (so many encodings, so many ways to DoS an IPS that tries to decode them). (Moore, 2006).

It is important to note that he was not describing the situation as of January 5, when the post was issued, but as of early December 28, when the module was. The situation then was much as he described it. There were no adequate antivirus signatures. The customers of IPS vendors were similarly helpless. Moreover, exploit code had also been made available on FrSIRT by the next day (Internet Storm Center, 2005).¹⁹ All security vendors began work on that day, so that by the time he wrote his tiny FAQ, the day that Microsoft's patch was issued, matters were much improved. The Metasploit module may have contributed to this, as he intended.

Was the exploit being discussed two weeks before the email from Bugtraq made it public? Here we should take Moore at his word. iDefense and Websense—and I have little doubt others—had a sense that some threat was gathering force. But many seem to have been caught by surprise, including Microsoft. Moreover, this was not the first vulnerability found in WMF. Indeed, Websense's first reaction to the exploits it saw was that they were of the earlier vulnerability.

The module made users and sysadmins more vulnerable to exploits of the vulnerability. Given the multiplication of

¹⁹ Exploit code on FRsIRT was available to everyone at that time. It is now available by subscription only.

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 H.D. Moore.

It is hard not to conclude from H.D. Moore's explanation 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. Was that effect large enough to offset the increased risk posed by the exploits the blackhats created from the module? That depends on how you see Microsoft and the vendors.

6.2. Other Early Whitehat Responses

By the time Microsoft issued its advisory and as blackhats were discovering the Metasploit module, most antivirus vendors were already finding defenses against the exploits.

Symantec received its first samples of exploits of the vulnerability on December 27. The samples they received then and later were mostly downloader trojans, with different files downloaded. Peter Ferrie had become quite familiar with the WMF parsing code from his work on exploits of vulnerabilities discovered earlier (Peter Ferrie, email to the author, 29 August 2006.). Owing to his knowledge, a heuristic was developed based

The December Storm of WMF

on the presence of the SetAbortProc subfunction. It was given the name Bloodhound.Exploit.56 (Symantec, 2006, February). The very presence of SetAbortProc was considered malicious. The risk from the vulnerability was labeled high. Updates were offered through LiveUpdate Weekly and Intelligent Updater on the 28th.

McAfee discovered it on December 27. A note giving that date said that they had found two exploits that downloaded a Trojan. These exploits were hosted on two web sites. Their first DAT file to detect the Trojan (4661) was issued on December 28; DAT 4662, issued the next day, provided "Enhanced Detection" of the exploit.

At 8:38 GMT on December 28, Mika Tolvanen posted an entry on the F-Secure blog that alerted its readers to exploits that used Trojan downloaders and were able to infect fully patched Windows XP machines with Service Pack 2 (Tolvanen, 2005). This, no doubt, reflected the information passed on in the email sent by the Bugtraq poster. The early advice given was to block access to the site cited in the Bugtraq message and "to filter all WMF files at HTTP proxy and SMTP level." They also noted that versions of Firefox and Opera that used "Windows Picture and Fax Viewer" were vulnerable as well as Internet Explorer.

At 12:07 GMT on the 28th, Kaspersky raised its alert level and gave the vulnerability its highest rating, extremely critical (Kaspersky, 2005). It had seen only Trojan downloaders used in exploits from two sites. Users were urged to update their antivirus, not to open files with the WMF extension, and to set their security settings in Internet Explorer to "high." It listed only Windows XP and Windows Server 2003 as vulnerable.

The December Storm of WMF

Three hours later, at 15:22 GMT, Jerome Athias, a French contributor to the Full Disclosure list, offered two workarounds (Athias, 2005). Both deregistered shingvw.dll in order to disable Windows Picture and Fax Viewer, which was often the immediate source of infection. These were picked up by Microsoft and others and included in the advice given to customers. This was the first time, but not the last, when an individual offered a defense for the vulnerability.

A result of the workarounds was a reduction of functionality; the user needed to have another way to see media files. Moreover, deregistering the shimvw.dll required the user to have technical skills, or a least confidence, that most lacked. It would also be learned in the next few days that there were ways of exploiting the vulnerability that did not use the dll.

Both vendors and individuals began making Snort signatures available. We noted SunBelt's contributions above. Verisign iDefense did so at 20:48 GMT (Dunham, 2006). The first cut at a signature on BleedingSnort appeared a few hours earlier (Mmlange, 2005). Updated signatures were posted often as the week went on, partly to improve the effectiveness and performance of the signatures against existing exploits, and partly to respond to changes in the blackhats' exploits that made them invisible to the earlier signatures. Snort signatures, of course, identified exploits, a vital step in handling intrusions. But no IDS stops them.

On the 28th, other vendors made their announcements that the vulnerability and its exploits were extant. The exploits

identified were usually Trojan downloaders. That is, the blackhats were trying to exploit the vulnerability with programs that seemed harmless, but would download malware that itself seemed innocuous.

By the end of the day on the 28th, the vulnerability had been identified, Microsoft was mobilizing—the whitehat community more generally was, too—and some basic but only partially effective defenses had been identified. On the other hand, a new source of exploits had been created. There was no evidence that a patch would appear quickly.

7. After the First Day: Growing Defense, Growing Confusion

The whitehat response after the 28th largely built on what had been started on that first day. 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.

7.1. Data Execution Prevention

On Thursday, December 29, both Sunbelt Software and the ISC found, independently, that Data Execution Prevention (DEP) helped to prevent infection. H.D. Moore had mentioned it the previous day (Moore, 2006, December 28a).

DEP prevents applications from executing code in regions of

the computer's memory not normally given to running code (Microsoft Corporation, 2004). It is enforced in hardware, if the processor supports it, and in software. This is a new technology, first included as one of the many security improvements in Service Pack 2 of Windows XP. Intel supports it in their more recent processors, beginning with the Pentium 4.20 AMD supports it in the Athlon 64, Socket 754 and 939 processors.

Microsoft activates DEP for Windows system files by default. That also includes programs that "opt-in" to DEP. Not all processes are covered, however, unless the user changes the default settings, found on the **Advanced** tab of the **System** applet in **Control Panel**, or makes changes to the boot.ini file. Neither procedure is likely to be undertaken by the ordinary user. Indeed, a late version of Microsoft's advisory on WMF tiptoes around how effective DEP is, suggesting that the interested user consult the hardware manufacturer for "more information about how to enable this feature and whether it can provide mitigation" (2005, December). DEP was not included among the mitigating factors, nor was changing DEP settings one of the actions suggested. It is not mentioned at all in MS06-001, the bulletin issued with the WMF patch.

That is unfortunate, because hardware DEP did work and could have been one more mitigating defense available to sysadmins. As noted above, H.D. Moore knew that as early as December 28. More unfortunate still was the confusion about DEP.

²⁰ For a list of all processors that supported DEP in January 2006, see Ou, 2006.

Did it work at all? Did it work with the default settings? Did software DEP work, or just hardware DEP? Different answers to all these questions were thrown around, even by Microsoft (Ou, 2005). This probably explains why the advisory was so circumspect about it. By Friday, 30 December, it was concluded that only hardware DEP was probably effective, but only if a broader setting than the default were chosen (Echelberry, 2006, December 30).

7.2. Whitehat Confusion

DEP provides but one example of the erroneous or confusing information that could be found sprinkled liberally among the multitude of messages about WMF. In part, this is to be expected in the communications environment that we have. It is the flip side of the good that comes when anyone can post to widely read lists like Bugtraq and Full Disclosure or create a blog accessible to the world. 'Caveat lector' (Let the reader beware) should be a mantra for those who browse through the world of information security.

Another example was the registry fixes suggested by Jerome Athias on December 28. Within two days, it was shown that they did not work, just as it was being shown that deregistering shimgvw.dll was only partly effective. Still another example was the confusion that could be found as late as Saturday about where the vulnerability lay—was it in shimgvw.dll or gdi32.dll? (Schouwenberg, 2006).

The estimate of the risk that users and sysadmins faced

differed widely among vendors and other whitehats. SANS, Secunia, and ultimately Microsoft, among others, told users that the risk was high, the vulnerability was critical (Secunia, 2005). McAfee and Symantec, to give contrary examples, found the risk from the vulnerability low, even in advisories published before the patch was issued (McAfee, 2006, and Symantec, 2006, February).²¹ The user could believe whomever he or she wanted. Widely differing definitions of what risk means in this context may account for part of the discrepancy. If so, the differences are not at all clear.

Perhaps the most significant confusion was over which operating systems were affected. Microsoft's advisory, in the version issued before patch was released, "discusses" Windows 98 and its brethren, Windows 98SE and Windows ME, Windows 2000, Windows XP, and Windows Server 2003. One has to assume that it skipped Windows NT only because Microsoft no longer supported it. No distinction was made about how the vulnerability affected the different systems. Because this was Microsoft, this extensive list was considered by many to be authoritative. F-Secure did Microsoft one better, noting that it even had a DOS box that became infected. iDefense, on the other hand, was cited as claiming that only Windows XP and Windows Server 2003 were vulnerable (Seltzer, 2006).²² The Metasploit module, of course, targeted only those two operating systems. Some said, emphatically, that Linux and other non-Windows systems were

²¹ Symantec does not address risk per se, but it rated the geographic distribution as low and both threat containment and removal as easy.

²² Later documentation made it clear that iDefense had found that more operating systems were vulnerable, but only the two named had been exploited.

unaffected, which later turned out not to be true (The Debian Project, 2006).²³

How much of this was wrong? Some, but not much. The truth was that programs on almost all systems were vulnerable, some more than others. But the user or sysadmin reading the list of operating systems that needed his or her full attention was poorly served by the lack of distinction between the nature of the threat and its severity. As became clear later, Windows 98 and its siblings were affected differently than Windows XP and Windows Server 2003. In the bulletin that accompanied the patch issued on January 5, Microsoft made this distinction clear by noting that the threat to the earlier operating systems was not critical. Even at the time, most of the attacks that had been seen were on the latter two operating systems.

The nature of the threat seems not to have been clear to all whitehats, but even knowing the focus of the blackhats' efforts would have made it possible for the sysadmins to prioritize. A sysadmin who has to add defenses to 10,000 machines would find it helpful to know where to start, particularly when the proper priority is counterintuitive. Who would believe that a flaw in a file format dating back to 1993 would be less likely to affect Windows 98 than the newest products from Microsoft, fully patched?

²³ Interestingly, this flaw, like so much else about WMF, was discovered by H.D. Moore.

8. Going into the New Year: New Vectors and Defenses

The blackhats' efforts seemed to pick up going into the New Year's weekend. Not only were new exploits discovered—that was old hat by now—but so were new vectors for exploits and new tools for creating exploits.

On Friday, December 30, John Herron of the National Institute of Standards and Technology, found that Lotus Notes was also vulnerable, as it used `shimgvw.dll` to view image files attached to messages. The message conclude despairingly that "...all Lotus Notes users are vulnerable to the WMF zero-day exploit. At this point there is little that can be done except block all incoming images at the perimeter" (United States, National Institute of Standards and Technology [NIST], 2006).²⁴

The next day, New Year's Eve, yet another vector for exploits was found—instant messenger (IM). Kaspersky Labs received reports from the Netherlands that an IM-worm in MSN used a link to an HTML page named `xmas-2006 FUNNY.jpg` to infect the victim in a manner typical for a WMF exploit (Schouwenberg, 2005). The page contained a WMF file. When touched, the file downloaded and executed a VBScript file. This file, in turn, downloaded an Sdbot. That bot, in its turn, was instructed to download another IM-worm in order to spread across MSN. The bot armies were ready to form, using WMF.

As if that were not enough, what the ISC termed a second

²⁴ This includes an update on 3 January. The original advisory came out on 30 December. Also see Laurio, 2005.

generation exploit was released that same day. F-Secure discovered that a blackhat, or a group of blackhats, decided to celebrate the holiday by releasing a new version of the exploit, unrelated to earlier exploits derived from Metasploit. It was sent through email as spam with a WMF file attached named HappyNewYear.jpg. When accessed, it downloaded a common backdoor from a particular website. The ISC found that only three AV programs detected it (Frantzen, 2006).

As December ended, then, it appeared that users and sysadmins were becoming more vulnerable rather than less. To many it seemed clear that the blackhats were winning the race while the whitehats waited for Microsoft's patch.

8.1. Guilfanov's Patch

That same day, New Year's Eve, Ilfak Guilfanov issued his own patch (Guilfanov, 2005). This was no ordinary code jockey, which made his patch all the more important. He was the architect and main developer of IDA Pro, a disassembler and debugger used widely by whitehats to analyze malware. Other patches were created in the next few days, notably one by ESET, a Slovakian antivirus company, but none received the notoriety of this one (Internet Storm Center, 2006, and ESET, 2006).²⁵

His 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

²⁵ Note that the ESET patch appeared several days after Guilfanov's patch.

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 (Murphy, 2006). Tom Liston of the ISC reviewed the code in detail, published his analysis, and helped to modestly extend its capabilities (Liston, 2006).

What did it do? Simply put, it prevented any application from calling SetAbortProc. So if the escape function were called, it would be rendered "invalid." No other escape was affected. The patch was tested on Windows 2000, Windows XP, and Windows Server 2003. It was updated later, on January 1 and January 3, to give it some important but minor improvements.

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. Liston concluded that:

This is a bad situation that will only get worse. The very best response that our collective wisdom can create is contained in this advice - unregister shimgvw.dll and use the unofficial patch. You need to trust us. (Liston, 2006, January).

With Guilfanov's permission, the ISC began to make his patch available. It was available on his own site as well. It

The December Storm of WMF

soon became more popular than Guilfanov or, it is likely, anyone else anticipated. So popular, in fact, that Guilfanov had to move the patch to a different site (set up by Castle Cops) owing to the number of people trying to download it.

Not everyone was thrilled to see the patch issued. On Tuesday, 3 February, Microsoft added a question on third party patches to its advisory. The answer included this:

Microsoft recommends that customers download and deploy the security update for the WMF vulnerability that we are targeting for release on January 10, 2006.

As a general rule, it is a best practice to utilize security updates for software vulnerabilities from the original vendor of the software. With Microsoft software, Microsoft carefully reviews and tests security updates to ensure that they are of high quality and have been evaluated thoroughly for application compatibility. In addition, Microsoft's security updates are offered in 23 languages for all affected versions of the software simultaneously.

Microsoft cannot provide similar assurance for independent third party security updates (Microsoft Corporation, 2005, December).²⁶

In most situations, there would be nothing objectionable to Microsoft's position. The first paragraph was an echo of what everyone else had said for a week. The third paragraph was

²⁶ The same language appears in Bulletin MS06-001.

simply fact. Microsoft tested the software thoroughly and could stand by it. Guilfanov's patch had been tested; its code had been reviewed thoroughly. But as solid as it was, Guilfanov could not be certain what it might break. In the blog entry he used to introduce the patch, he wrote: "I'd like to know what programs are crippled by the fix, please tell me" (Guilfanov, 2005). And, indeed, there was at least one report that the patch caused problems (Echelberry, 2006). Microsoft could not afford to be as uncertain as Guilfanov was, or the ISC. Its audience, after all, was predominantly users. These were people who looked to Microsoft for the security of the products that Microsoft built and they bought. They were much less technically sophisticated than the whitehats and sysadmins who read the *Handler's Diary*.

All the same, a less strictly oppositional approach could have been taken, and might be considered in the future. The principle of *caveat emptor*, after all, is alive and well in the computer industry. A passage could have been added along the lines of the following: "A patch has been produced and tested by respected experts in the industry. It may offer an alternative to some customers until our own security update has been issued. Microsoft can take no responsibility for the consequences."

Why would that be valuable? For the same reason that some felt it necessary to install the patch. Microsoft's customers had pushed the company to go to a regular schedule of patching because they believed that the risk they faced from a delay in most patches was small when placed against the benefit that having a predictable, monthly "Patch Day." Similarly, the risks that Microsoft sees in third party patches are also seen by most

whitehats. Aside from problems of quality, they can also be sources of malware, dangled before the desperate. But there are times when the situation makes the risks worth running.²⁷ For many sysadmins, such was the case in regard to the Guilfanov patch. Such situations may well be seen again.

8.2. The Next Week: Alarums and Excursions

After the storm of the weekend, the first days of the week brought little new. Not that the blackhats or 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. iDefense decided to sound the alarm. They sent off the code samples and results of other research they had conducted to antivirus vendors and others (Ken Dunham, telephone conversation with the author, 23 August 2006).²⁸

Also on Monday, H.D. Moore informed the ISC that The Metasploit Project had issued a new version of its WMF module. As he told SANS, the new module

...uses some header padding tricks and gzip encoding to bypass all known IDS signatures. Consider this "irresponsible" if you like, but it clearly demonstrates that a run-of-the-mill signature-based IDS (or A/V) is not

²⁷ Jesper M. Johanson, Senior Security Strategist in the Security Technology Unit at Microsoft, made this point about the patch on his blog three days before Microsoft's patch came out. See Johanson, 2006.

²⁸ They saw their efforts produce results. Between the alarm they sounded and the return of people from vacation, antivirus detections were up significantly within 24 hours.

going to work for this flaw. (Sachs, 2006).

That same day, Panda Software discovered a new application to develop WMF-based malware, WMFMaker (Panda Software, 2006, February). 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” (Ready Rangers Information Front, 2006). The blackhats had no intention of letting the problem fade away.

That same day, as cries for an official patch continued, a patch was leaked from Microsoft, becoming available through several sources (Evers, 2006; Eckelberry, January 4a; and Eckelberry, January 4b). It was tested by some who got it. It was good. The time had almost come.

9. Microsoft’s Patch: Written, Tested, Released

We left Microsoft gearing up to work on a patch for the vulnerability on December 27, shortly after the Bugtraq message appeared. There are indications that the patch itself was completed the next day (Seltzer, 2006, January 5).²⁹ It was “smoke tested” for obvious defects, then turned over to the test team.

This 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

²⁹ Seltzer found that the file date on a leaked version of the patch was December 28 at 21:54 EST.

The December Storm of WMF

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 (Cushman, 2006).

It was a massive effort, originally designed to last a full two weeks, until January 10. Yet this work was completed and the patch issued five days early, faster than Microsoft had written and tested an update (Vijayan, 2006). The process, following procedures that had been planned carefully and tested, appears to have worked as intended. The release of the patch and the accompanying bulletin did not go as smoothly.

Releasing it ahead of schedule was a last-minute decision, made the afternoon after testing was completed, according to the director of the MSRC (Vijayan, 2006). As late as Wednesday, January 4, Microsoft affirmed that the patch would be released on "Patch Tuesday" (Reavey, 2006). No official announcement was made until after 3:00 EDT on Thursday, January 5. But the news was leaked to Ron Trent, who announced on his blog on myITforum that Microsoft would release the patch at 5:00 EST. Sunbelt, like others, picked up Trent's post and made their own report at 12:03 EST (Trent, 2006). 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 6 January at 2:00 PM PST, but 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). Then at 3:20 EST, the blog of the Microsoft Security Response Center issued an announcement from Mike Nash, the Corporate Vice President responsible for security at Microsoft. The official announcement was made at the same time (Nash, 2006; and Microsoft, 2006, January 5b).

The official patch was finally out. Users and sysadmins now had the tools they needed to make their system safe from exploits of the WMF vulnerability.

10. Conclusions

With the release of Microsoft's patch, the concern that many whitehats felt about the threat from exploits of the WMF vulnerability could be relaxed. Not that the threats were gone. As noted above, most people do not patch. Exploits remain common even now, nine months after the patch was released.

But how did the whitehats perform before the patch was released? How well did they provide defenses for the sysadmins and users? The answer has to be mixed. The security community as a whole did some things well, but there is room for improvement on the whole. Some parts of the community did better than others.

10.1. Communication.

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. Defense in depth and following best practices by, for example, not running a computer as administrator, provided incomplete, but significant protection.

Yet communication was also a central problem in the way the WMF vulnerability was handled. Not all whitehats communicated well. Some vendors were deafeningly silent. Google Desktop was named as a vector for the vulnerability on the first day. Yet 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 produces security equipment that is used widely, yet issued no advisory. The vulnerability was not in a CISCO product, but CISCO products helped to mitigate it, and the sysadmins who rely on CISCO were likely to have found CISCO's advice useful.³⁰ In contrast, companies that produce competing products, Internet Security Systems (ISS), for

³⁰ The only documentation that CISCO has produced about WMF was a piece showing how effective the CISCO Security Agent was against a single exploit. It reads more like a sales piece than the result of determined testing. See CISCO Systems, 2006.

example, were not as shy. They made certain that their customers knew about the vulnerability and how to mitigate it (Internet Security Systems, 2005).

10.2. A Learning Curve.

It may seem like faint praise to say that the whitehats made incomplete protection available to users and sysadmins. But its not. Whitehats worked long hours to find ways to mitigate the threat. On the whole they communicated well. Yet there are two caveats that go along with that word of praise.

First, few understood the nature of the vulnerability or how to defend against its exploits from day 1. Early exploits were at first misidentified by some as targeted at an earlier vulnerability. It was only after several days and some fairly optimistic postings that the weaknesses of DEP and of deregistering shimvw.dll became clear to all taking part in the discussion. In addition, the blackhats clearly glommed onto something in the nature of the vulnerability that escaped the whitehats when the "second generation" exploits began to appear on New Year's Eve. Even Microsoft had to change its advisory as it learned from others about defenses its own technology provided. In short, there was a great deal of ignorance about what was going on and how to stop it.

There was, in fact, learning curve for this vulnerability. It took time to learn about features of files that used this format, even though it had already been exploited and was more than a decade old. In the future must expect that it may be difficult to respond quickly and adequately to the rapidly

spreading exploits of a vulnerability that is more recent and more complex.

10.3. The Audience.

The other caveat gets back to communications. Much of what was written about WMF was by and for the cognoscenti. Users and, to a lesser extent, sysadmins, have neither the time, the knowledge, nor the inclination, to fruitfully spend time reading Bugtraq, *The Handler's Diary*, or even Microsoft Advisories about how to take care of their computer. As we have seen elsewhere, as in updating antivirus signatures and patching, many of those the whitehats sought to protect from WMF exploits undoubtedly skipped reading about or acting on the good advice that was available, if they knew about it at all.

Moreover, defenses like the Guilfanov patch and deregistering shimvw.dll were in essence useless for most of the potential victims of the exploits. Users by and large do not have the technical competence or confidence to apply such solutions. That should be so obvious that it may be unnecessary to say it. The advice given was useful to other whitehats and to many sysadmins. It could not take the place of an easily installed patch from Microsoft. And the continued importance of WMF exploits shows that a patch is a necessary, but by no means sufficient means for stopping the blackhats, as the continuing success that blackhats have with WMF shows. This lesson is old, of course, with hundreds of examples.

So, what can be done about vulnerabilities like WMF and their exploits? Ultimately, it will take a combination of good

coding practices and automatic defenses to bring what seems to be the explosive spread of malware under control. Like other ills—floods, tornadoes, scam artists, and economic recessions—malware and its purveyors may always be with us. Our job is to minimize the harm.

10.4. Better Data.

How bad was the WMF episode? The truth is that we do not know. Exploits were, and are, remarkably widespread. Someone is finding it worthwhile to continue producing them. Someone is making money. By implication, thousands, hundreds of thousands, perhaps millions of computers have been infected. But while there are numerous and conflicting estimates of the number of infected sites, there are few estimates of the number of users infected, and none that are clearly complete.

Some useful data is available. McAfee was quoted as saying on Saturday, December 31 that 6 percent of their customers were infected. By Monday, that figure was up to 7.4 percent. Their own bulletin said that as of January 3, "McAfee is aware of over 120,000 McAfee VirusScan Online customers who have reported detecting Exploit-WMF files attempting to execute on their systems" (McAfee, 2006). Trend Micro provides data as a matter of course on the number of infected computers (Trend Micro, 2006). Panda offers its Global Malware Observatory to look at malware in real time (Panda Software, 2006, October). All these sources are valuable, but they all depend on a narrow range of essentially self-chosen sources. For the vendors, the data depends on who uses their products and agrees to report. The ISC faces similar constraints on the global data it provides.

In short, there is no source that can tell us, with confidence, how many computers of any kind were infected, much less distinguish infections by operating system, by location (geography or size of network) or other characteristics of the machine or its owner.³¹ Even now, there is some dispute over how serious the threat from the vulnerability was. Vendors have tended to call it a minor threat now that the patch is out.

If such data exists, it is not publicly available, which gets us back to the communication problem. Like other statistics about the victims of crime, reliable data will be difficult to get. But it cannot be impossible. Having such data early on would have made it easier for users, sysadmins, and the whitehats trying to help them to focus their efforts to find defenses. In trying to fight off a more active attack it could make the difference between success and failure.

It may be necessary to have a third party—the government, perhaps, or a private, non-profit organization—to produce such data.

10.5. Metasploit.

The question whether the Metasploit module should have been released on December 28 is a key one that the episode raises.

³¹ Cert/CC tried to keep statistics on the number of incidents per year, but gave up, saying that "Given the widespread use of automated attack tools, attacks against Internet-connected systems have become so commonplace that counts of the number of incidents reported provide little information with regard to assessing the scope and impact of attacks" (CERT Coordination Center, 2006). I argue that more precise data can be quite useful, particularly as the blackhats focus on targets that can prove lucrative.

Some have found the answer easy, depending on how they view the vendors of security tools and, especially, Microsoft. It seems likely that the module and its update, released after New Year's Day, helped speed the issuance of the patch. The cost to users, in particular, however, was probably high.

H. D. Moore's reasons for issuing the module were given above. In many respects, they come down to a question of trust and communication. Moore, like many others, did not trust Microsoft—or other vendors—to act responsibly in the face of the brewing trouble over the new vulnerability. There may be more trust in the future; the rest of us would benefit if there is. Both Moore and Microsoft proffered something of an olive branch this spring, with the latter extending an invitation to its Bluehat briefings and the latter accepting it.

10.6. Microsoft.

Microsoft took a lot of criticism, both before and after the patch came out. It should be recognized, however, that Microsoft did a lot of things right. The SSIRP was invoked early, and a tremendous effort was put into getting a patch written and tested. The advisory on the vulnerability was issued early and updated several times in the week that followed. The information in it was vital. The decision to issue the patch early rather than to stick to the schedule of even the day before can only be applauded.

There are, however, three areas where improvements might be made.

The December Storm of WMF

The patch came out earlier than expected, but later than it might have. With passive exploits like WMF, this might be acceptable or at least tolerable. Given the lack of data about the number of machines infected before the patch came out, it is hard to argue persuasively against arguments like this:

[The value of an automated updating process] was proven when the famed Windows Meta File (WMF) flaw was uncovered in late 2005...; there was very little damage. In fact, it was a bit of a non-event because updates were applied rapidly to the vast majority of computers, thereby making the systems secure from attack quickly. (Howard and Lipner, 2006, p. 134).

Yet, as the same authors point out, there are situations when a delay of more than week can be disastrous, such as when there is an attack on the scale of Code Red or Nimda (Howard and Lipner, 2006, pp. 221-222). Most of the time taken to prepare the patch was in testing, and Microsoft's procedures, as noted above, call for all tests to be complete before a patch is issued. There are commendable reasons for this. But Microsoft needs an addendum to its procedures that describes when and how to introduce shortcuts to the process to allow a patch or, at least, a beta patch, to be issued quickly.³² At the same time, Microsoft needs to retain flexibility in its approach. After

³² The creation of the Zeroday Emergency Response Team (ZERT) in September 2006, picking up where Guilfanov left off, provides an potential alternative to a change in Microsoft's procedures. Its track record may make it a reputable, credible, temporary alternative to an out-of band Microsoft patch. See <http://isotf.org/zert>.

all, we don't know when or how the next massive attack will come, or which blackhats will launch it.

Second, the information provided in the initial security advisory was incomplete. Perhaps that was inevitable, given the speed with which things happened. But, as argued above, sysadmins and whitehats may have found it helpful to know where to focus their efforts. Bulletin MS06-001 gave them some of that information. The advisory did not tell them that Windows 98 and its ilk were less likely to be infected.

Nor did it describe the vulnerability in enough detail for the source of the vulnerability to be known. Consequently there was some confusion about whether which DLL the vulnerability could be found in. Close reading of the advisory would have made that clear, but close reading should not be required.

Third, arguably, Microsoft should have known more about the vulnerability before it was discovered by the blackhats or at least made public on Bugtraq. Vulnerabilities had been found in WMF files before. Microsoft has had an effective program of code review in place since 2002. Given the amount of code that needs to come under review, no one should expect all of it to have been done in four years. All the same, should the WMF have been reviewed after the first vulnerability was discovered? If not that, then some effort might have been made to finally retire what was a legacy format anyway. The ubiquity of WMF would argue against that, but it has become clear that the security risk of keeping it is high.

In addition, Microsoft was helped significantly by Dan

Hubbard and his Websense team. But why doesn't Microsoft have its own early warning research group? That is not simply a rhetorical question; the business case for it may not be strong enough to warrant the expense. But there was a cost to Microsoft in not knowing what was about to come down the pike. The cost to Microsoft and its customers may be much higher next time.

10.7. Under the Radar.

Most of the discussion about WMF took place in English. Some of it was translated, but translation, understandably, took a back seat to ascending the learning curve—in English—and getting the information out in the easiest form possible, which meant English. This may have hindered the ability of users and sysadmins whose speak English poorly or not at all to get the information they needed. Given the growing number of users who do not speak English, that is a topic worthy of research on its own.

But another implication of this is that the blackhats can operate under the radar. Exploits of the WMF vulnerability were extant for about two weeks before they were made public—leaked—in an email message to an English language list. What if the leak had not been made? Websense may have been on the verge on discovering the vulnerability that the exploits it saw was exploiting, but few others began to move until they saw that email and the Metasploit module was released. It could easily have been several days before the vulnerability was recognized. Almost assuredly, Microsoft would have held off a patch until the January 10th, assuming that it knew enough to make one even then.

More broadly, how much is going on that we do not know about because the whitehat community as a whole has neither the language skills nor the predisposition to linger where blackhats talk in languages other than English? There are regions of the world where blackhats are developing exploits—writing them, testing them, using them—that may easily become global threats.³³ Businesses and governments in Europe and North America will be targeted by such exploits. We have some capability to create an early warning capability that can try to learn about what is happening in hidden hotbeds of blackhat activity like Brazil, China, and the former Soviet empire. WMF showed a little about what can be done, with iDefense and Kaspersky in particular producing essential information about the origins of the exploits. But it also showed that a stronger capability is needed, either from companies like iDefense and Symantec, or from government. Unfortunately, the required combination of skills—linguistic and technical—is scarce.

³³

A useful perspective on this can be found in Hayashi, 2006.

Timeline

- 12 December (app)** Evidence shows exploits of WMF from approximately this date.
- 23 December Friday** Websense finds an unexplained exploit.
- 26 December (app) Monday** Dan Hubbard of Websense posts a message about it on the vetted Malicious Websites and Phishing (MWP) email list.
- 27 December Tuesday** Sunbelt Software learns about an exploit about 5:00 PM EST. They notified Microsoft later that day. Sample of exploit of WMF vulnerability sent to F-Secure. An email to Bugtraq sent at about 5:20 pm EST makes the vulnerability and its exploits public. Sender had sent sample to F-Secure moments before. F-Secure notified Microsoft, both in Europe and Redmond, other antivirus companies, and Google.
- 28 December Wednesday** H.D. Moore notifies the full disclosure list that an exploit of the vulnerability has been ported to the Metasploit Framework. Microsoft begins the SSIRP (Software Security Incident Response Process). Microsoft issues Microsoft Security Advisory (912840).
- 29 December Thursday** ...
- 30 December Friday** John Herron at NIST.org discovers that Lotus Notes is a vector.
- 31 December Saturday** Second generation exploit released. Patch released by Ilfak Guilfanov. New vector: IM Worm found using WMF.
- 1 January 2006 Sunday** SANS recommends installation of Guilfanov's patch.
- 2 January Monday** Update of Metasploit module released.
- Tuesday 3 January** McAfee discovers WMFMaker, another tool to create exploits.
- 4 January Wednesday** Microsoft's patch is leaked, but quickly withdrawn.
- 5 January Thursday** Microsoft issues its patch and Bulletin MS06-001.

References

Athias, Jerome. (2005, December). "Someone wasted a nice bug on spyware...." Retrieved October 19, 2006, from <http://lists.grok.org.uk/pipermail/full-disclosure/2005-December/040699.html>.

Carey, Allen. (2005, December). "2005 Global Information Security Workforce Study." IDC. Retrieved on October 18, 2006, from <https://www.isc2.org/download/workforcestudy05.pdf>.

CERT Coordination Center. (2006, October). "CERT/CC Statistics 1988-2006." Retrieved October 20, 2006, from http://www.cert.org/stats/cert_stats.html#incidents).

CISCO Systems. (2006). *CISCO Security Agent and the Microsoft WMF Exploit*. Product Bulletin 324. Retrieved on October 20, 2006, from http://www.cisco.com/en/US/products/sw/secursw/ps5057/prod_bulletin0900aecd80420fde.html.

Cushman, Andrew. (2006, February 20). "Microsoft Security Fundamentals." Presented at EUsecWest. London. Retrieved on October 18, 2006, from eusecwest.com/esw06/esw06-cushman.ppt.

The Debian Project. (2006, February). "Debian Security Advisory DSA-954-1 WINE -- Design Flaw." Reported 25 January 2006. Retrieved on October 20, 2006, from <http://www.debian.org/security/2006/dsa-954>.

Dijker, Barbara. (Fall, 1999). "Careers in System Administration: The future looks bright for problem solvers."

The December Storm of WMF

Dr. Dobbs Portal, July 2001. Originally published in *Dr. Dobbs Journal*, Fall 1999. Retrieved on October 18, 2006, from <http://www.ddj.com/showArticle.jhtml;jsessionid=10EXOXM3XUGVQSNDBECKHSCJUMEKJVN?articleID=184411171>.

_____. (Dijker, 2006). "A Day in the Life of System Administrators (aka DITL)." Retrieved October 18, 2006, from <http://www.sage.org/field/ditl.pdf>.

Dunham, Ken (2006). *Anatomy of WMF Zero-Day Attacks*. May 3, 2006.

Echelberry, Alex. (2005, December 27). "New Exploit Blows by Fully Patched Windows XP Systems." *Sunbelt Blog*. Retrieved October 19, 2006, from <http://sunbeltblog.blogspot.com/2005/12/new-exploit-blows-by-fully-patched.html>.

_____. (2005, December 30). "Microsoft clarifies "DEP" issue." *Sunbelt Blog*. Retrieved October 19, 2006, from <http://sunbeltblog.blogspot.com/2005/12/microsoft-clarifies-dep-issue.html>.

_____. (2006, January 3). "One report of network printing problems." *Sunbelt Blog*. Retrieved October 19, 2006, from http://sunbeltblog.blogspot.com/2006_01_01_sunbeltblog_archive.html.

_____. (2006, January 4a). "Microsoft security patch has leaked." *Sunbelt Blog*. Retrieved October 20, 2006, from http://sunbeltblog.blogspot.com/2006_01_01_sunbeltblog_archive.html.

The December Storm of WMF

_____. (2006, January 4b). "We take a quick look at the Microsoft hotfix." *Sunbelt Blog*. Retrieved October 20, 2006, from http://sunbeltblog.blogspot.com/2006_01_01_sunbeltblog_archive.html.

ESET. (2006, January 4). "No Microsoft Patch Currently Available." Retrieved January 20, 2006, from <http://www.eset.com/company/article.php?contentID=947>.

Evers, Joris. (2006, 4 January) "Microsoft inadvertently leaks WMF patch." C|Net News.com. Retrieved on October 20, 2006, from http://news.com.com/Microsoft+inadverdently+leaks+WMF+patch/2100-1002_3-6018263.html.

Ferrie, Peter. (2006, February). "Inside the Windows Meta File Format." *Virus Bulletin*. Retrieved October 18, 2006, from <http://www.virusbtn.com/virusbulletin/archive/2006/02/vb200602-wmf>.

Frantzen, Swa. (2006, January 1). "2nd generation WMF exploit: status of the anti-virus products after one day." *Handler's Diary*. Retrieved October 20, 2006, from <http://isc.sans.org/diary.php?storyid=998>.

Frieder, Laura L., and Zittrain, Jonathan L. (2006, August). "Spam Works: Evidence from Stock Touts and Corresponding Market Activity." Social Science Research Network. August 2006. Retrieved October 18, 2006, from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=920553#PaperDownload.

The December Storm of WMF

Gordon, Sarah. (1994, September). "The Generic Virus Writer." Paper presented at The 4th International Virus Bulletin Conference, Jersey, UK. Retrieved on October 18, 2006, from <http://www.research.ibm.com/antivirus/SciPapers/Gordon/GenericVirusWriter.html>.

_____. "The Generic Virus Writer II." (1996, September). Paper presented at the 6th International Virus Bulletin Conference, Brighton, UK. Retrieved on October 18, 2006, from <http://www.research.ibm.com/antivirus/SciPapers/Gordon/GVWII.html>.

Gordon, Sarah, and Ford, Richard. (1999, September). "When Worlds Collide: Information Sharing for the Security and Anti-virus Communities." Virus Bulletin Conference, Vancouver, British Columbia. Retrieved October 18, 2006, from <http://www.windowsecurity.com/uplarticle/10/vb99final.pdf>.

Gostev, Alexander. (2006). "Malware Evolution: October-December 2005." Retrieved October 18, 2006, from <http://www.viruslist.com/en/analysis?pubid=178619907>.

Grimes, Roger. (2006, January 6). "WMF Warnings: I Wasn't Calling Wolf." Retrieved October 19, 2006, http://www.infoworld.com/article/06/01/06/73582_020Psecadvise_1.html.

Guilfanov, Ilfak. (2005, December 31). "Windows WMF Metafile Vulnerability HotFix." *Hex Blog*. Retrieved on October 20, 2006, from <http://www.hexblog.com/2005/12>.

Hale, Deborah. (2005, 27 December). "Quiet Weekend - Not

much news." Retrieved October 17, 2006, from <http://isc.sans.org/diary.php?storyid=968>.

Hayashi, Kaoru. (2006, April). *Regional Threats*. White Paper: Symantec Security Response. Originally published by *Virus Bulletin*. Retrieved October 20, 2006, from <http://www.symantec.com/avcenter/reference/regional.threats.pdf>.

Hocutt, Ronald. (2002, April). "The Virus Hunters: A Day in the Life." *SmartComputing: Learning Series*, pp. 8:23-8:25. Retrieved on October 18, 2006, from <http://www.smartcomputing.com/editorial/article.asp?article=articles/archive/10804/08104/08104.asp&guid>.

The Honeynet Project. (2004). *Know Your Enemy: Learning about Security Threats*. Second Edition. Boston, MA: Addison-Wesley.

Howard, Michael, and Lipner, Steve. (2006). *The Security Development Lifecycle: SDL: A Process for Developing Demonstrably More Secure Software*. Redmond, WA: Microsoft Press.

Hurley, Edward. (2003, June 18). "A Day in the Life of a Virus Researcher: Speed, Accuracy Paramount When Wrestling With Worms." Retrieved on October 18, 2006, from http://searchsecurity.techtarget.com/originalContent/0,289142,sid14_gci906881,00.html.

Hyponnen, Mikko. (2006, January). "Bad Behaviour." Retrieved on October, 19, 2006. <http://www.f-secure.com/weblog/archives/archive-012006.html#0000075>.

The December Storm of WMF

iDefense. (2006, January 3). "State of the Hack: WMF Attacks: Code for Cash." Document ID 434705.

Internet Security Systems (ISS). (2005). "Microsoft Shared DLL WMF graphics Rendering Code Execution." Internet Security Systems Protection Alert. Original: 28 December 2005, updated 4 January 2006. <http://xforce.iss.net/xforce/alerts/id/211>.

Internet Storm Center. (2005, December). "Update on Windows WMF 0-day." *Handler's Diary*. Retrieved October 19, 2006, from <http://isc.sans.org/diary.php?storyid=975>.

_____. (2006, January) "WMF FAQ (NEW)." *Handler's Diary*. Published: January 3, 2006; Last Updated: January 5, 2006. Retrieved on October 20, 2006, from <http://handlers.dshield.org/jullrich/wmffaq.html>.

Johanson, Jesper M. (2006, January 2). "Conscientious Risk Management and WMF." *Jesper's Blog*. Retrieved October 20, 2006, from http://blogs.technet.com/jesper_johansson/archive/2006/01/02/416762.aspx.

Kaspersky Labs. (2005, December 28). "Windows Meta File Vulnerability," 28 December 2005 Retrieved October 19, 2006, from <http://www.viruslist.com/en/alerts?alertid=176701669>.

Laurio, Juha-Matti (2005, December 30). "Lotus Notes WMF File Handling Code Execution Vulnerability. Retrieved October 20, 2006, from <http://www.security.nnov.ru/Kdocument842.html>.

Liston, Tom. (2006). "The WMF SETABORTPROC Vulnerability

and Ilfak Guilfanov's Patch: A Technical Explanation of the Issues and Available Workarounds for the WMF Flaw." Retrieved in January 2006 from handlers.sans.org/tliston/WMFTech.pdf.

_____. (2006, January 1). "Trustworthy Computing," *Handler's Diary*, Published: January 1, 2006; Last Updated: January 1, 2006. Retrieved October 20, 2006, from <http://isc.sans.org/diary.php?storyid=996>.

McAfee. (2006, January 5). "Exploit-WMF." Modified January 5, 2006. http://vil.nai.com/vil/Content/v_137760.htm.

McNamara, Caolan. (1996). "Microsoft Windows Metafile." Retrieved October 17, 2006, from <http://wvware.sourceforge.net/caolan/ora-wmf.html>.

Microsoft Corporation. (2002). "Windows History." Published June 30, 2002, updated March 7, 2006. Retrieved October 18, 2006, from <http://www.microsoft.com/windows/WinHistoryDesktop.aspx>.

_____. (2003, October). "Revamping the Security Bulletin Release Process." Published October 1, 2003; updated February 14, 2005. Retrieved on October 18, 2006, from <http://www.microsoft.com/technet/security/bulletin/revsbwp.aspx>.

_____. (2004, September 15). "Changes to Functionality in Microsoft Windows XP Service Pack 2; Part 3: Memory Protection Technologies." Updated September 15, 2004. Retrieved October 19, 2006, from <http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/sp2mempr.aspx>.

The December Storm of WMF

_____. (2005, May). "Responding to Security Incidents." Published May 6, 2005; updated: February 9, 2006. Retrieved on October 18, 2006, from http://www.microsoft.com/security/msrc/incident_response.aspx.

_____. (2005, December). "Microsoft Security Advisory (912840) Vulnerability in Graphics Rendering Engine Could Allow Remote Code Execution," published: December 28, 2005; updated: January 3, 2006. Retrieved on January 4, 2006, from <http://www.microsoft.com/technet/security/advisory/912840.aspx>.

_____. (2006a). "Security Considerations: GDI." Retrieved October 18, 2006, from [http://windowssdk.msdn.microsoft.com/en-us/library/ms536711\(VS.80\).aspx](http://windowssdk.msdn.microsoft.com/en-us/library/ms536711(VS.80).aspx).

_____. (2006b). "Windows GDI, Escape." Retrieved October 18, 2006, from http://msdn.microsoft.com/library/default.asp?url=/library/en-us/gdi/prntspol_0d6b.asp.

_____. (2006, January 5a). "Microsoft Security Advisory (912840) Vulnerability in Graphics Rendering Engine Could Allow Remote Code Execution," published: December 28, 2005; updated: January 5, 2006. Retrieved on October 19, 2006, from <http://www.microsoft.com/technet/security/advisory/912840.aspx>.

_____. (2006, January 5b). "Microsoft Security Bulletin MS06-001: Vulnerability in Graphics Rendering Engine Could Allow Remote Code Execution (912919)." Retrieved on October 19, 2006, from <http://www.microsoft.com/technet/security/Bulletin/MS06-001.aspx>.

The December Storm of WMF

Mmlange. (2005, December 28). "WMF Exploit Signature." Retrieved October 19, 2006, from <http://www.bleedingthreats.net/forum/viewtopic.php?forum=3&showtopic=1544>.

Moore, H.D. (2006, January 5). "Exploiting WMF (tiny) FAQ." Retrieved October 19, 2006, from <http://archives.neohapsis.com/archives/fulldisclosure/2006-01/0151.html>.

_____. (2005, December 28a). "Re: [Full-disclosure] Someone wasted a nice bug on spyware..." Retrieved October 19, 2006, from <http://archives.neohapsis.com/archives/fulldisclosure/2005-12/1363.html>

_____. (2005, December 28b). "Re: Is this a new exploit?" Retrieved October 19, 2006, from <http://www.securityfocus.com/archive/1/420352/30/0/threaded>.

Matthew Murphy. (2006, January 4). "Interview: Ilfak Guilfanov." *SecuriTeam Blogs*. Retrieved October 20, 2006, from <http://blogs.securiteam.com/index.php/archives/176>.

Nash, Mike. (2006, January 4). "Mike Nash on the Security Update for the WMF Vulnerability." *Microsoft Security Response Center Blog*. Retrieved October 20, 2006, from <http://blogs.technet.com/msrc/archive/2006/01/05/416980.aspx>.

Noemailpls. (2005, December) "Is this a New Exploit?" Retrieved October 19, 2006, from <http://www.securityfocus.com/archive/1/420288>.

The December Storm of WMF

Ou, George. (2006, January 23). "Guide to Hardware-Based DEP Protection." Retrieved October 19, 2006, from <http://blogs.zdnet.com/Ou/?p=150>.

_____. (2005, December 30). "Lots of bad advice for critical WMF vulnerability!" Retrieved October 19, 2006, from <http://blogs.zdnet.com/Ou/?p=143>.

Panda Software. (2006, 11 February). "WMFMaker." Virus Encyclopedia. Updated on 11 February 2006. Retrieved October 20, 2006, from http://www.pandasoftware.com/virus_info/encyclopedia/overview.aspx?IdVirus=103242.

_____. (2006, October). "Global Malware Observatory." Retrieved October 20, 2006, from http://www.pandasoftware.com/virus_info/map/observatory.htm?site_panda=empresas.

Parker, Tom; Sachs, Marcus; Shaw, Eric; and Stroz, Ed. (2004). *Cyber Adversary Characterization: Auditing the Hacker Mind*. Syngress.

PCMag. (2006). "Exploit." Encyclopedia. Retrieved October 17, 2006, from http://www.pcmag.com/encyclopedia_term/0,2542,t=exploit&i=42871,00.asp.

Pupkin-Zade. (2006, 2 February). "Eksplloit k WMF uyazvimosti prodali na chernim rinke." Khaker.ru. Retrieved on October 18, 2006, from <http://www.xakep.ru/post/29944/default.asp>.

The December Storm of WMF

Ready Rangers Information Front. (2006, January). "WMF Virus Competition?" Last updated January 23, 2006. Retrieved October 20, 2006, from <http://home.arcor.de/vxdia/misc/wmfrulez.htm>.

Reavey, Mike. (2006, January 4). "WMF Vulnerability Security Update." Microsoft Security Response Center Blog. Retrieved October 20, 2006, from <http://blogs.technet.com/msrc/archive/2006/01/04/416847.aspx>.

Sachs, Marcus. (2006, January 2). "More .wmf woes." Handler's Diary. January 2, 2006 Published: January 2, 2006; Last Updated: January 2, 2006. Retrieved October 20, 2006, from <http://isc.sans.org/diary.php?storyid=1002>.

SAGE. (2006). "Core Job Descriptions." Last changed: 30 Sept. 2006 Retrieved October 18, 2006, from <http://www.sage.org/field/jobs-descriptions.html>.

SANS. (2006). "Glossary of Terms Used in Security and Intrusion Detection." Retrieved October 17, 2006, from <http://www.sans.org/resources/glossary.php>.

_____. (2006, July). *Network Traffic Analysis Using TCPDump*, parts 1 and 2. Bethesda, MD: The SANS Institute.

Sapieha, Chad. (2006, May 14). "Symantec Uses NASA-Like Security Room to Protect Your PC." Toronto Star. Retrieved October 18, 2006, from <http://evergeek.thestar.com/Features/2060.aspx>.

Schouwenberg, Roel. (2006, December). "More on WMF

The December Storm of WMF

Exploitation" *Analyst's Diary*. Retrieved October 19, 2006, from <http://www.viruslist.com/en/weblog?discuss=176892530&return=1>.

Securiteam. (2005, December 29). "Microsoft Windows WMF Buffer Overflow (Exploit Metasploit)." <http://www.securiteam.com/exploits/5DP0I2KHHE.html>.

Secunia. (2005). "Microsoft Windows WMF "SETABORTPROC" Arbitrary Code Execution." Secunia Advisory SA18255. Released December 28, 2005; last updated 28 February 2006. Retrieved October 19, 2006, from <http://secunia.com/advisories/18255>.

Seltzer, Larry. (2006, January 2). "Researchers Dispute Which Windows Versions Are Vulnerable." *Larry Seltzer's Security Weblog*. Retrieved October 19, 2006, from http://blog.eweek.com/blogs/larry_seltzer/archive/2006/01/02/6969.aspx.

_____. (2006, January 5). "How Well Does The Leaked MS Patch Work?" *Larry Seltzer's Security Weblog*. Retrieved October 20, 2006, from http://blog.eweek.com/blogs/larry_seltzer/archive/2006/01/05/6973.aspx.

Swan, Tom. (1993). *Inside Windows File Formats*. Indianapolis, IN: Sams Publishing.

Symantec. (2006, February). "Bloodhound.Exploit.56." Updated February 10, 2006. Retrieved October 19, 2006, from http://www.symantec.com/security_response/writeup.jsp?docid=2005-122814-2600-99.

The December Storm of WMF

_____. (2006, March). *Symantec Internet Security Threat Report Trends for July 05-December 05*. Retrieved October 17, 2006, from http://enterprisesecurity.symantec.com/pdf/ISTR_IX_FullReport.pdf.

Tolvanen, Mika. (2005, December). "New WMF 0-day Exploit." *News from the Lab*. Retrieved October 19, 2006, from <http://www.f-secure.com/weblog/archives/archive-122005.html#00000752>.

Toulouse, Stephen. (2006). "Looking at the WMF issue, how did it get there?" *Microsoft Security Response Center Blog*. Retrieved October 18, 2006, from <http://blogs.technet.com/msrc/archive/2006/01/13/417431.aspx>.

Trend Micro. (2006, October). "EXPL_WM.F.GEN: Statistics." Retrieved October 20, 2006, from <http://www.trendmicro.com/vinfo/virusencyclo/default5.asp?VName=EXPL%5FWMF%2EGEN&VSect=S&Period=All>.

Trent, Rod. (2006, January 5). "Product Support Services - JANUARY 2006 MICROSOFT SECURITY RESPONSE CENTER BULLETIN RELEASE." *myITforum.com*. Retrieved October 20, 2006, from <http://myitforum.com/cs2/blogs/rtrent/archive/2006/01/05/18131.aspx>.

United States, Central Intelligence Agency. (2006). *The World Factbook 2006*. Retrieved October 18, 2006, from <https://www.cia.gov/redirects/factbookredirect.html>.

Update Capital. (2004, Fall). *IT Security Sector Report*.

The December Storm of WMF

Retrieved on October 18, 2006, from
http://www.updata.com/publications/security_111804.pdf.

Ullrich, Johannes. (2006 July). "Networks Under Fire! The Internet Storm Center." Presentation made at SansFire, Washington, D.C. Retrieved October 18, 2006, from
isc.sans.org/presentations/sansfire2006keynote.pdf.

United States, National Institute of Standards and Technology (NIST). (2006, January 3). "Lotus Notes vulnerable to MS Windows graphics rendering engine bug." Retrieved October 20, 2006, from
http://www.nist.org/nist_plugins/content/content.php?content.25.

Vijayan, Jaikumar. (2006, January 6). "Q&A: Microsoft exec explains the early WMF patch release." *Computerworld*. Retrieved on October 19, 2006, from
<http://computerworld.com/securitytopics/security/holes/story/0,10801,107522,00.html>.

Websense Security Labs. (2005, July). *Overview of Our Investigative Process*. Retrieved October 18, 2006, from
<http://www.websense.com/docs/WhitePapers/WSLabsOverview.pdf>.

_____. (2005, December 27) "Potential new unpatched IE exploit?" [12:07 pm]. Retrieved October 18, 2006, from
<http://www.websense.com/securitylabs/blog/blog.php?BlogID=15>.

_____. (2006, January 2). "Malicious Website/Malicious Code: WMF Attack Update/Timeline." Retrieved October 18, 2006, from
<http://www.websense.com/securitylabs/alerts/alert.php?AlertID=39>

0.

_____. (2006, January 5). "Informational Alert: WMF Patch Available from Microsoft." Retrieved October 17, 2006 from <http://www.websense.com/securitylabs/alerts/alert.php?AlertID=39>

2.

Wilson, Clay. (2004, July). *Information Warfare and Cyberwar: Capabilities and Related Policy Issues*. CRS Report for Congress RL31787, Updated July 19, 2004. Retrieved on October 18, 2006, from <http://www.fas.org/irp/crs/RL31787.pdf>.

Author's Note

The author would like to thank the adviser for the paper, John Bambenek, and other advisers who read the paper, for their help. In addition, the following people were generous in answering questions through email or the telephone:

Ken Dunham, Director, Rapid Response Team, iDefense

Alex Echelberry, President, Sunbelt Software

Peter Ferrie, Senior Anti-Virus Researcher, Symantec
Security Response

Mikko Hypponen, Manager of Anti-Virus Research, F-Secure

Dan Hubbard, Senior Director of Security and Technology
Research, Websense

H.D. Moore, founder, The Metasploit Project

Eric Sites, Vice President of Research and Development,
Sunbelt Software