



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

Interested in learning more?

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

The Future of Fighting Viruses: A History and Analysis of the Digital Immune System

Toby Bussa

GSEC Practical Assignment version 1.2d

Introduction

IBM's Digital Immune System (DIS) represents a logical step in automated, centralized anti-virus activities. It is an idea that has taken almost a decade to come to fruition, and represents what the future of anti-virus activities may look like. A future where human intervention, in all aspects of the anti-virus process, from virus detection, to signature creation, to virus cleaning, is no longer necessary. Understanding the DIS has become salient now that Symantec has licensed, and incorporated, this technology into their enterprise level anti-virus software Norton AntiVirus (NAV).

This paper will first present a brief overview of how the DIS works. Next, the history of the DIS will be examined. A technical analysis of the DIS in Symantec's commercial product, including strengths and weaknesses, will then be presented. Finally, concluding remarks will be given.

Overview

IBM's Digital Immune System refers to the theoretical model that was proposed by David M. Chess. IBM's DIS was intended to be a fully automated, anti-virus solution from its initial inception. The following is a breakdown of the process of the anti-virus system proposed. We will review the development of the initial design through the history of the DIS in the next section. The initial DIS concept varies from the commercial design, which will be discussed in the technical analysis section later in the paper.

The proposed design was comprised of a system that included: integrity monitors and anti-virus heuristics monitoring at the client level, a virus scanner, an algorithm to detect virus characteristics, a signature extractor, a central database, and a system for communicating new signatures to other computers.

The typical process flow for the system is as follows.

- Integrity Monitors and anti-virus heuristics monitor a system.
- Once an anomaly is detected, the anti-virus scanner scans the system.
- If a virus is found, it is cleaned if a signature exists.
- If no virus is found, decoy files are dropped in an attempt to entice the virus to infect the files.
- The decoy files are periodically checked for changes.
- If changes are found, the files are quarantined.

- An algorithm attempts to determine the infection method and find common byte sequences.
- The signature extractor attempts to create a signature using the common byte sequences the infection method algorithm suggested.
- The virus identity is then added to the database.
- Finally, the system sends signals to other computers informing them of the virus.

This approach, given that it was suggested when viruses still consisted of boot sector infectors and other slow moving viruses, represented a forward-thinking approach. It anticipated the immense increase in the types and quantity of viruses, worms, and other malware that would emerge and dominate anti-virus activities almost ten years later.

In the next section, we trace the development of Chess' idea from the development lab at IBM to its incorporation into an enterprise-level anti-virus software system.

History

David M. Chess, a researcher at IBM in 1991, suggested the idea of the Digital Immune System. It is appropriate to let him summarize the intentions of this idea.

"If transmission bandwidth, CPU cycles, and disk space were free, and programming was easy, every workstation would be protected by a seamless 'immune system.' Objects infected with existing viruses would be detected automatically, the identity of the virus verified and reported to a central location, and the object destroyed or repaired, with minimal user intervention. New viruses would be detected automatically with some high degree of confidence, first-pass signature patterns would be extracted where possible and communicated with a central clearinghouse, along with a sample of the suspicious object. Viruses would very rarely, if at all, spread quickly." [1]

Two years later, in 1993, Kephart et al. analyzed computer viruses in a traditional epidemiological view. There was a relationship made to traditional viral defenses and computer virus defenses. Kephart et al. show in their case study that "[w]ith effective central reporting and response -- where the entire incident is cleaned up as soon as any machine is found to be infected -- the situation is similar to the below threshold." This epidemic threshold is the "relationship between the viral birth and death rates at which a disease will take off and become widespread. Above this threshold, the disease becomes persistent, recurring infection in the population. Below it, the disease dies out." [2]

This represented a further evolution towards a DIS, now that computer viruses have been examined in an epidemiological view and their research supports a centralized system for mitigating computer virus epidemics.

The next year, 1994, Kephart formalizes the ideas that he and Chess put forth in previous papers. Kephart stated the realization that a globally connected network of computers (i.e., the Internet) would exist in the future and allow computer viruses to spread much more rapidly. He also states that IBM is already running pieces of the DIS in their lab. [3] The second paper expands upon the DIS process of automatically extracting virus signatures. [4]

The pieces of the system described in the Overview earlier are already functioning by 1994. Only the following pieces had not been fully implemented at that time: the prototype algorithm that determines the infection method and finds common byte sequences was working on about 90% of viruses, and the ability to send signals to other computers informing them of a virus was not yet developed. This is significant progress to realize in only three years time.

A couple of years later, in 1996, TIME visits IBM and reviews the progress of the DIS. The article first mentions the ability of the DIS to automatically send a snippet of suspect code to IBM across the Internet, and then for IBM to automatically create a vaccine and send it back to the sending computer. [5]

We now began to see more pieces of the DIS. The small bits of insight here indicate that IBM is looking away from a peer-to-peer update scheme and opting for a centralized anti-virus analysis center.

The next year, 1997, Kephart et al. propose that the DIS meet the following criteria: innate immunity, adaptive immunity, delivery and dissemination, speed, scalability, safety and reliability, security and customer control. [6] These requirements provide a framework for the technical details of the DIS. The technical aspects of these requirements are discussed, especially those related to virus disinfection.

There is a realization that the Internet represents an even more fertile ground for viruses than was imagined in 1993. Kephart et al. state that "the nature of computer viruses and their ability to propagate is on the cusp of a fundamental, qualitative change -- one that demands an equally fundamental change in the way we must defend against them." They were right.

In May 1999, Symantec announces that they have licensed IBM's DIS technology. [8] Around the same time, IBM researchers describe the pilot of their technology in Symantec's Norton AntiVirus software.

White et al. pose the requirements of a commercial-grade solution as: detecting new and previously unknown viruses at the client-level, the ability to handle epidemics and floods, an automated system for responding to new viruses, the ability to scale at the architecture and implementation levels to deal with future virus threats, maintaining a high degree of stability, and allowing enough flexibility for the customer to integrate the product in existing infrastructures. [9]

This implementation of IBM's technology is called the Symantec Digital Immune System. More pieces of the DIS are discussed in detail: loads on the central immune system (average, peak, and overloads), the "active network" design to deal with the loads on the central immune system, security and reliability, and the automated virus analysis center. [9] The types of issues being dealt with are those that are necessary to implement the DIS on a much larger scale than a development lab. These are discussed later in the technical analysis.

White indicates that there is still work left to do on the technology to handle various types of viruses, such as bimodal and polymorphic, as well as Microsoft Access and PowerPoint macro viruses.

In October of 2000, Symantec announces the commercial availability of Norton AntiVirus, the first commercial product to use IBM's DIS. [10]

Technical Analysis

In this section we provide a technical analysis of the components of Symantec's DIS. The DIS is comprised of the following components: Norton AntiVirus at the client level, the Active Network (comprised of the Administration Consoles, Gateway Servers, and Central Virus Analysis Center), and Cure Distribution.

It is easiest to understand each of these components in the DIS and how they interact with the other components to provide an automated, end-to-end anti-virus solution. The following sequence describes the process shown in Figure 1 when a virus is detected at the client level.

© SANS Institute 2000 - 2002
Author retains full rights.

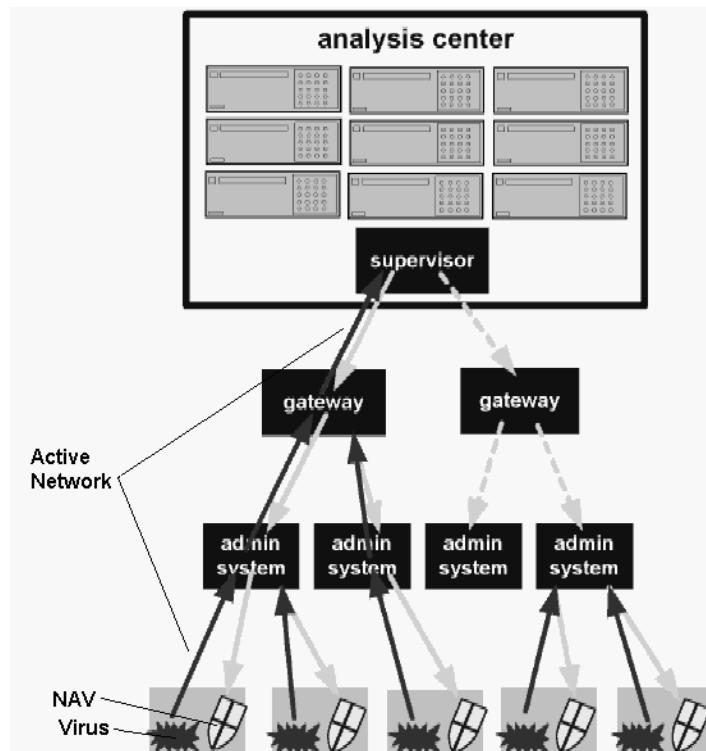


Figure 1. Digital Immune System described by White et al. [9]

Norton AntiVirus (NAV) detects a potentially new virus on a client workstation or server via a heuristic or a signature file that cannot disinfect or verify the virus. If enough suspicion is raised, NAV isolates a sample and sends it to the Administrative Console.

At the Administrative Console, control exists to determine what is sent in and out of the internal network by the DIS. There are several reasons why the administrator might not automate having the sample sent to Symantec for processing. The administrator can remove sensitive information from the file, such as a Microsoft Word document. The administrator might have newer signature files that have not been deployed which can be tested against the virus.

Next, the virus can be submitted to Symantec's Central Virus Analysis Center for processing. This is done via secure protocols across the Internet using a proprietary transaction protocol called AVIS, which is used on the Active Network. Data transmissions actually occur over HTTP using SSL via TCP/IP. DES, RSA, and DSA are all used as the underlying cryptographic primitives. [9]

The sample is actually sent to a gateway, which acts as an intermediary between the clients' Administrative Console and the Supervisor system in the Central Analysis Center. The gateways are the primary nodes in the Active Network with the Administrative Consoles acting as "leaves" and the Analysis Center as the "root" in a "tree hierarchy."

Under peak loads, the gateways act as filters for the Central Analysis Center. They have a couple of functions. They check to see if they can handle the request by themselves, comparing the virus sample to a database of previously examined files, which means the file is either clean or can be handled by a previous virus signature.

Under normal loads, the gateways will check the sample against currently known signatures and then forward the virus samples to the Analysis Center. This process works since they rely on checksums, which are sent along with the samples from clients. The gateway only needs to check the CRC's which makes the process very efficient. [9]

At the Central Virus Analysis Center, the sample is examined and a new virus signature is automatically constructed. The analysis center consists of a series of IBM RISC/6000 running AIX and Windows NT systems behind a firewall. A Supervisor system coordinates the activities of the other worker systems in a modular fashion. These systems attempt to get the virus to spread by using "goat" files. If an infection does occur, these files are examined and used to produce a signature for that virus. The signature is then tested against the original sample and if successful, sent back out to the client. If there is any problem in the process, the sample is sent to a human for manual examination.

The virus signature is then sent back to the Administrative Console via the gateway where it can be tested by the administrator or automatically deployed to all clients on the internal network running NAV.

Strengths

This approach has a number of strengths over conventional anti-virus processes. Signature files, especially those written in high-level languages, are as good or better than human-produced signature files. These signature files are created in a significantly faster period of time with fewer instances of false-positives. [4] Additionally, as the number of viruses in the wild grows, an automated system with scaling capabilities will be able to grow faster and handle many more viruses than a system which relies on humans that possess a special knowledge for decoding viruses and creating signatures.

An emphasis has been placed on building a system that is scalable, has high availability, and can deal with different load situations. These are important issues that present bottlenecks to the DIS and represent an obvious acknowledgement by IBM that as the DIS has developed, so has the understanding that viruses are going to continue to be produced at an increasing rate in the foreseeable future.

Since the Internet is used as the transport medium, strong security is mandatory to transport samples and signatures between the Central Virus Analysis Center and the Administrative Console. This consideration has been dealt with by using commonly used network protocols. The DIS uses encryption and secured channels on HTTP port 80 to minimize the potential of virus samples and signatures being intercepted in route.

Finally, IBM and Symantec have provided flexibility in the application in order to tailor the system to existing infrastructures, especially with regards to Administrator-level controls. The DIS technology in a fully automated solution, becomes a necessity for smaller companies that lack dedicated IT support to deal with viruses. A properly configured system would require little maintenance and provide a large return on investment as the cost of cleaning up viruses (usually done by outside vendors on contracts) would be greatly reduced. As a company's size, computing environment, and network grows in complexity, then so does the need for more control over these systems. By providing that control at the Administrators Console, large companies can tailor the system to provide the most effective level of control and allow dedicated IT staff to oversee anti-virus activities.

Weaknesses

The Digital Immune System has been well thought out. All appearances indicate that the number of weaknesses in the system is small. However, there are still some issues which should be mentioned.

First, how are repairs handled if a virus is destructive? What if data is destroyed? Although the virus could be cleaned, what if a Love Letter-type virus is introduced that overwrites all jpeg and mpeg files with a copy of virus? This step could still require extensive human intervention in order to recover from a virus outbreak. Thus, there are instances where a fully automated system would break down.

Second, there is the issue of new viruses and infection methods. How does the DIS scale to accommodate new virus technologies, such as Win2k stream viruses? Thankfully, many new viruses are concept viruses that are first given to anti-virus software companies. This allows anti-virus vendors time to evaluate these viruses and update their products accordingly. But what happens when a new type of virus appears in the wild that hasn't been seen by anti-virus labs?

Finally, it is important that anti-virus companies spread this technology out among multiple research locations. This would require an additional layer of technology in the Central Virus Analysis Centers that would allow all the locations to remain connected in order to maintain consistency. This would add an additional layer of security as multiple locations are less likely to be a target of denial of service attacks and other types of network outages and attacks. White et al. [9] recommend this approach once the use of multiple Central Analysis Centers is necessary, but there is no technical indication of how multiple Centers would remain synchronized with one another.

Conclusion

The Digital Immune System invented by IBM represents an impressive attempt to provide a future system for dealing with viruses. There are issues that should be

addressed by Symantec and IBM in order to improve the DIS, but the information on the DIS indicates that many of the numerous issues and problems that could occur with the system have been evaluated and appropriate measures have been taken. The strengths of this system outweigh any current or potential issues. IBM has brought an interesting concept and delivered it as a commercially viable product which represents one of the few new technological innovations in the area of anti-virus activities and technologies.

References

1. David M. Chess. "Virus Verification and Removal Tools and Techniques." Virus Bulletin, November 1991 (1991).
URL: www.research.ibm.com/antivirus/SciPapers/Chess/CHES3/chess3.html
2. Kephart, Jeffrey O., David M. Chess, and Steve R. White. "Computers and Epidemiology". IEEE Spectrum, May 1993 (1993): pp. 20-16.
URL: www.research.ibm.com/antivirus/SciPapers/Kephart/Spectrum/Spectrum.html
3. Kephart, Jeffrey O. "A Biological Inspired Immune System for Computers." Artificial Life IV, Proceedings of the Fourth International Workshop on Synthesis of Living Systems, Rodney A. Brooks and Pattie Maes, eds. Cambridge: MIT Press, 1994. pp. 130-139.
URL: www.research.ibm.com/antivirus/SciPapers/Kephart/ALIFE/alife4.distrib.html
4. Kephart, Jeffrey O. and William Arnold. "Automatic Extraction of Computer Virus Signatures." Proceedings of the 4th International Virus Bulletin Conference, Jersey, UK, 1994. pp. 179-194.
URL: www.research.ibm.com/antivirus/SciPapers/Kephart/VB94/vb94.html
5. Ramo, Josua Cooper. "The Next Big Thing? Digital think tanks helped invent the information age. Now they're working on the future." Time, 148:22 (1996).
URL: www.time.com/time/magazine/archive/1996/dom/961111/cover_story.the_next_big38.html
6. Kephart, Jeffrey O., Gregory B. Sorkin, Morton Swimmer, and Steve R. White. "Blueprint for a Computer Immune System." Proceedings of the 1997 International Virus Bulletin Conference, San Francisco, California, 1997.
URL: www.research.ibm.com/antivirus/SciPapers/Kephart/VB97/index.html
7. Kephart, Jeffrey O., Gregory B. Sorkin, David M. Chess, and Steve R. White. "Fighting Computer Viruses." Scientific American, November 1997 (1997).
URL: www.sciam.com/1197issue/1197kephart.html

8. "Symantec Unveils Digital Immune System Strategy for Unprecedented Level of Managed, Intelligent Protection and Control." Symantec Press Release, May 11, 1999.
URL: www.symantec.com/press/1999/n990511.html
9. White, R. Steve, Morton Swimmer, Edward J. Pring, William C. Arnold, David M. Chess, and John F. Morar. "Anatomy of a Commercial-Grade Immune System."
URL: www.research.ibm.com/antivirus/SciPapers/White/Anatomy/anatomy.html
10. "Symantec's Norton AntiVirus Corporate Edition 7.5 Integrates Digital Immune System Technology to Provide Automatic Virus Protection Enterprise-wide."
Symantec Press Release, October 31, 2000.
URL: www.symantec.com/press/2000/n001031.html

© SANS Institute 2000 - 2002, Author retains full rights.