BLUE COAT®

Security Empowers Business

DEFEATING NEXT-GENERATION MALWARE WITH NEXT-GENERATION ANALYSIS

WHITEPAPER

The success of some recent malware attacks has made headlines, crippled corporations, robbed shareholders, and damaged the credit of thousands of consumers. It is also abundantly clear that cybercriminals have evolved adaptive next-generation malware that is capable of bypassing the security defenses that many organizations rely on today. This advanced malware is capable of sensing sandboxing environments and of mutating in situ like a biological virus.

At the same time, hackers show endless ingenuity in penetrating corporate networks. As reported in the media, they have gained entry through a variety of devices and third-party suppliers, including health insurance providers, printers, thermostats, vending machines, and even the online menu of a popular Chinese restaurant.

The purpose of this paper is to describe how next-generation malware has evolved, how it functions, and how it can be identified, neutralized and blocked by next-generation malware analysis. First, let's look at the way a blizzard of advanced malware alarms is affecting IT security teams.

The Car-Alarm Syndrome

To a busy security team, dealing with dozens or hundreds of alarms every day can be numbing. Compare this with car alarms. As far back as 1997, a study of insurance claims covering 73 million cars indicated no overall reduction in theft loss due to car alarms. A law enforcement study of car alarm noise in New York City found that up to 99 percent of alarms were false. We can reasonably say that almost nobody pays much attention to them anymore. Unfortunately, we can find parallels in network security today.

Take the well-known case of a large U.S. retailer. Using an opening unwittingly provided by an HVAC vendor, hackers introduced malware into the store's system just prior to the 2013 holiday shopping season. It soon became active, siphoning off customers' credit card data from point-of-sale checkout operations. The company's offshore security analysts detected the activity and sent an alarm to corporate headquarters. The company did not act. Days later, the security team sent another alarm, and again IT failed to take action. Weeks later, the U.S. Department of Homeland Security informed the retailer that approximately 40 million credit card numbers had been plundered from the company's system and were being offered for sale on the black market. Customers began staying away; sales plummeted; and their stock declined.

So the need is clear: a next-generation defense that deals with advanced threats and generates fewer and more meaningful alarms. This paper describes both the tools and the architecture that let companies achieve this. First, let's look at the way threats are evolving.

How Next-Generation Malware is Evolving to Avoid Detection

As network security tightens, malware is becoming more aware and adaptive, mutating like a biological virus to evade behavior detection.

Virtual machine awareness

One powerful way to discover new malware attacks is sandboxing – isolating suspected or unknown files in a virtual environment that mimics a company's desktop systems. While they're in the sandbox the suspect files are examined. If the file exhibits malicious behavior in the sandbox, it is recognized as malicious, and information about the file is used to prevent further attacks from the newly discovered malware. The malware's activity exposes its identity.

An increasing number of attackers, however, are creating malware that can detect when they're operating in a virtual environment. If the VM-aware code senses a sandbox it will disguise itself by going dormant or performing non-malicious acts that reduce the utility of the sandbox.

BLUE COAT®

Security Empowers Business

How does malware detect virtual environments? One way is by looking to see if a human is interacting with it. If the malware contains a dialog box it will expect a human to respond to it. But if nothing happens, it will presume it's in a sandbox and go dormant. Another way is by checking for virtual device drivers, telltale registry entries, and other giveaways. One example is unexpected system timing elements. In a real Windows environment, events happen quickly and with known predictability. Hackers do performance calculations to check elapsed time for specific operations. A noticeable lag exposes the virtual environment.

Polymorphic files and URLs

Malware files can morph and mutate like an infectious virus to escape signature-based detection. Using automated systems, hackers can change a letter, insert a few extra bits, pack (compress) the code, reverse some non-essential instruction, or add some junk data, and recompile to generate tens of thousands of variants. Every time the file presents itself, it looks different. There is no vaccine for this. Signature-based security systems can't cope with this flood of viruses; some of the files are bound to penetrate and begin to operate.

Attackers do similar things with URLs, using domain-generating algorithms (DGAs) to mathematically compute new domains. The malware has access to these algorithms; hackers may communicate with an URL for a set number of hours or days, moving on to another URL so blacklisting can't keep up.

Multi-stage, multi-vector attacks

Cybercrime can tailor its activities to the target. Hackers can select among web-based, email, and file-based intrusions, coordinating them and staging the timing of events to achieve a focused result.

Encrypted communication

Because most network security systems are unable to scan encrypted data to detect malware, hackers find it effective to utilize SSL to build communication tunnels between embedded malware and remote command and control (C&C) servers.

Misleading file types

Malware may masquerade as harmless files. Executable files may pretend to be JPEGs. There may be executable files inside a media file, an Excel file, or a PowerPoint file. A malware file can initially be named (name).jpg, then renamed (name).exe and run by a second malware file. So if your defenses are set to block all executables, the JPEG file may make it through – and stay until another file arrives that turns it into an executable and runs it.

Sleeping malware

Malware may be programmed to lie inactive until a specified date, exemplified by New Year's Day and April Fool's viruses. It may be analyzed but not considered malicious because it is dormant.

User interaction triggers

Malware requires a response to be activated. Because it often pretends to be legitimate, it may display a dialog box that asks users to install some software. It will probably be accompanied by certification and a Microsoft look that seems familiar and friendly. The user says, 'Yes – install it', and the malware can go into operation. Without the interaction, without the mouse click, it remains passive. Legitimate software may be packaged inside malicious software; download a free version of something popular, and you may get more than you bargained for.

Unique and targeted malware

Some malware can be incorporated in a targeted "spearfishing" attack. If it's aimed at you, it will trick you into opening a file by utilizing information specific to you. The hackers may know your environment well, and the specific assets they're looking for. They may look for a specific system state – for example, the presence of a custom application – that indicates a desired target. Malicious insiders may look for the directory path that leads to the target. They leave no signature, and they never get caught.

The Solution: Blue Coat Next-Generation Targeted Analysis

Blue Coat technology development has not been sleeping while malware authors have been evolving new attack techniques. We have created next-generation analysis techniques that identify and neutralize malware designed to evade detection technology. They block known threats; they analyze anything new and not known; and they stifle evolved attacks. They grade the risk of each threat, and the car-alarm/cry-wolf syndrome goes away.

BLUE COAT®

Security Empowers Business

Dual-detection methodologies

The Blue Coat Malware Analysis Appliance uses the powerful combination of emulation and virtualization to identify both types of malicious code –VM-aware and non-VM-aware. Virtualization takes place in a virtual machine that is a full licensed version of Windows in which the user can install any application – Office, Adobe, Quicken, or custom applications that are built in-house. We call it Intelligent VM (iVM).

The emulative sandbox environment is not Windows software; it's a fully re-created computing environment based on a Windows-like API. In this completely controlled artificial space, users can exercise the malware to make it think it's interacting with a real computer. If malware is set to sleep until a specific date, for example, the Malware Analysis Appliance can make it believe that that day is here.

It's extremely difficult for any malware author to evade both the virtual and emulative environments. With this unique combination, Blue Coat is approaching 100 percent success in identifying malware.

Kernel-level detection

Blue Coat analysis detects behaviors deep in the kernel – not in the user space. This makes detection very difficult for malware to evade, because evasion techniques are programmed in the userland space. The intention to perform malicious action remains discoverable in the kernel, and the Malware Analysis Appliance can report on these low-level events. So even if the malware author is able to display harmless characteristics, Blue Coat analysis sees the truth down deep in the kernel.

Anti-VM environment settings

Blue Coat has spent years and taken great pains to make its virtual machine environment undetectable to malware. Malware authors know how to look for indicators of a virtual environment, such as virtual devices and tell-tale registry settings and keys. The Blue Coat Malware Analysis Appliance applies multiple settings that include changing virtual device names, removing registry entries and disabling guest additions. If the malware tests the waters by offering up a dialog box or asking questions, the Malware Analysis Appliance responds like the real system.

Anti-VM pattern-matching and risk-rating

The Malware Analysis Appliance replaces signature-based detection with behavior detection patterns. Blue Coat has developed hundreds of behavior patterns, and a subset of them is specifically designed to catch malware that is looking for the presence of a virtual environment. They analyze behaviors deep in the kernel, looking for specific actions.

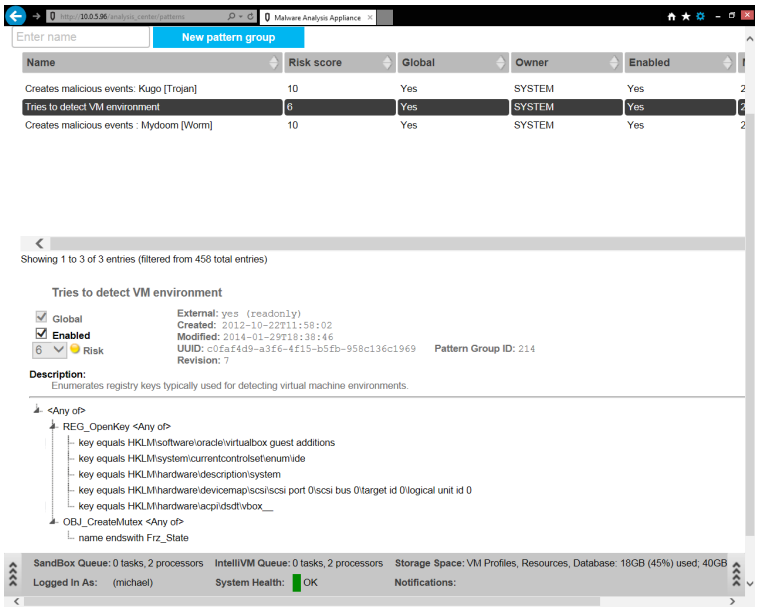

In the screen shot above from a Malware Analysis Appliance, the highlighted malware is trying to detect a VM environment by checking registry keys. The keys can be hidden or removed to make them undetectable.

Instead of giving malware a binary Good or Bad label, Blue Coat custom pattern matching lets you grade the seriousness of each risk when it's identified. Scoring the threats lets you eliminate false alarms.

Customizable environments

The virtual environment offered by most security vendors is generic, preloaded with configurations designed to represent a typical setup. But generic environments only detect generic threats. Blue Coat takes it many steps further.

BLUE COAT®

Security Empowers Business

Only a custom environment will catch threats specifically directed at you. The Malware Analysis Appliance enables you to customize a virtual environment that is extremely close to your corporate gold image. You can build virtually the same system your work force uses – the same version of Windows, the patch level revisions, the applications and their versions, the updates, and the custom applications. Blue Coat will soon be offering the capability to simply clone your system for this purpose.

In this customizable, realized environment, you can add plug-ins that recognize whether the malware is replicating a human being with mouse clicks, cursor moves, or interaction with dialog boxes. You can set up a specific directory tree or false financial records to nail malware that targets you or your company.

Most importantly, with all these advanced Blue Coat threat defenses, any malware you identify you can block at the web gateway and integrate into the Blue Coat Global Intelligence Network – a collaborative defense that will block it worldwide.

Deploying the Blue Coat Advanced Threat Protection Lifecycle Defense

Done right, malware analysis gives you enhanced protection against the evolved techniques of malware authors. If it's not deployed effectively, you may be swamped by responses to trivial malware and pushed into the car-alarm syndrome: too many alarms, followed by overreaction, followed by the dangers of benumbed under reaction.

The Blue Coat Advanced Threat Lifecycle Defense lets you overcome these challenges by combining next-generation malware analysis in a complete security solution to perform the following:

Block all known web threats	BLUE COAT PROXYSG SECURE WEB GATEWAY
Allow known good traffic with application white listing, and block known bad traffic with malware scanning	BLUE COAT CONTENT ANALYSIS SYSTEM
Analyze unknown threats	BLUE COAT MALWARE ANALYSIS APPLIANCE

Building trust in the system by scoring and pre-filtering

The filtering power of this intelligent defense-in-depth lets you reduce the necessity of analysis, grade threats to eliminate false alarms, and build trust in the reliability of the system. The Content Analysis System draws on a database of more than a million records to identify applications and files in real time. It enables you to pre-filter threats by scoring them on a scale of 1 to 10. Is it malware, goodware, or unknownware?

A file or application that has been coming into the system for years without a problem would rate a 10. Another that is unknown but comes without negative reports could be rated a 2 or a 3. A retail organization, like the one we described above, would rate malware that targets cash registers and point-of-sale card-swiping as a 1. A company running Linux will register the entry of Windows-targeted malware, but it will not be rated high enough to create an alarm that requires attention. Some industries – banks, for example – may choose to restrict downloads to high-rated known-good files.

With Content Analysis System filtering in place, fewer files will be sent to the anti-malware engines and to sandboxing systems. Alarms will be fewer and more meaningful. To put this in perspective, let's look at a typical business day at a financial organization – a Blue Coat customer with 250,000 employees:

- Employees will make 660 million attempts to contact websites.
- They will make 2.2 million attempts to access known malicious sites that are blocked by Blue Coat WebPulse, using input from 75 million users worldwide as a part of the Global Intelligence Network.
- Network perimeter anti-malware will block 244 malicious files.

Keep in mind that the blocked sites referred to here are truly malicious, not simply undesirable. They could include popular and trusted sites that have been infiltrated and corrupted. But the point here is the tremendous potential for the generation of alarms at any large corporation, and the time-consuming challenges they present to IT security.

The Blue Coat Advanced Threat LifeCycle Defense also delivers actionable information about malware attacks. It provides detailed analysis of malicious behavior patterns, key indicators that malware has compromised something in the system, and a timeline for how malware works.

BLUE COAT®

Security Empowers Business

The Blue Coat Advanced Threat Prevention Lifecycle Architecture

The Blue Coat Advanced Threat Prevention LifeCycle Defense provides an elegant, high-efficiency architecture for next-generation analysis and incident resolution. It functions as follows:

1. A user downloads content from the web through the ProxySG secure web gateway, which sends it to the Content Analysis System for malware scanning via ICAP or ICAPS.
2. The Content Analysis System checks the file in real time against the known-good-file whitelist database, which is hosted in the Global Intelligence Network. If it's listed there, the file is delivered and Content Analysis System processing is finished. A temporary local cache is maintained for performance reasons.
3. If the file is not whitelisted, it's scanned by one or two anti-virus (AV) engines. If the file is known bad (rated 0) it is blocked and its URL is added to the Global Intelligence Network.
4. If the file is neither known good or known bad (rated 1), it can be sent to one or more sandboxing appliances, including the Malware Analysis Appliance or FireEye AX. (The Content Analysis System will first check the local file hash database to see if the file has already been analyzed.) When sandboxing is complete, the result goes to the Content Analysis System. If the file is malicious, the Content Analysis System updates the local cache – the file hash database – and tells the ProxySG to block all subsequent requests to the same object. It also updates the Global Intelligence Network with the object's URL, file hash, timestamp and filename.

In Summary: Next-Generation Malware Requires Next-Generation Analysis

Corporate networks are being challenged by the two problems described in this paper: too many alarms, and evolved malware with ingenious abilities to avoid detection. They are connected. The solution to both is an architecture that combines a secure web gateway, incident containment, next-generation malware analysis, incident resolution, and a threat-scoring system that puts alarms on a rational basis. The Blue Coat Advanced Threat Lifecycle Defense is designed to give you these capabilities in a complete network security solution with a simple mission: protect corporate assets by passing the known good, blocking the known bad, and analyzing the unknown.

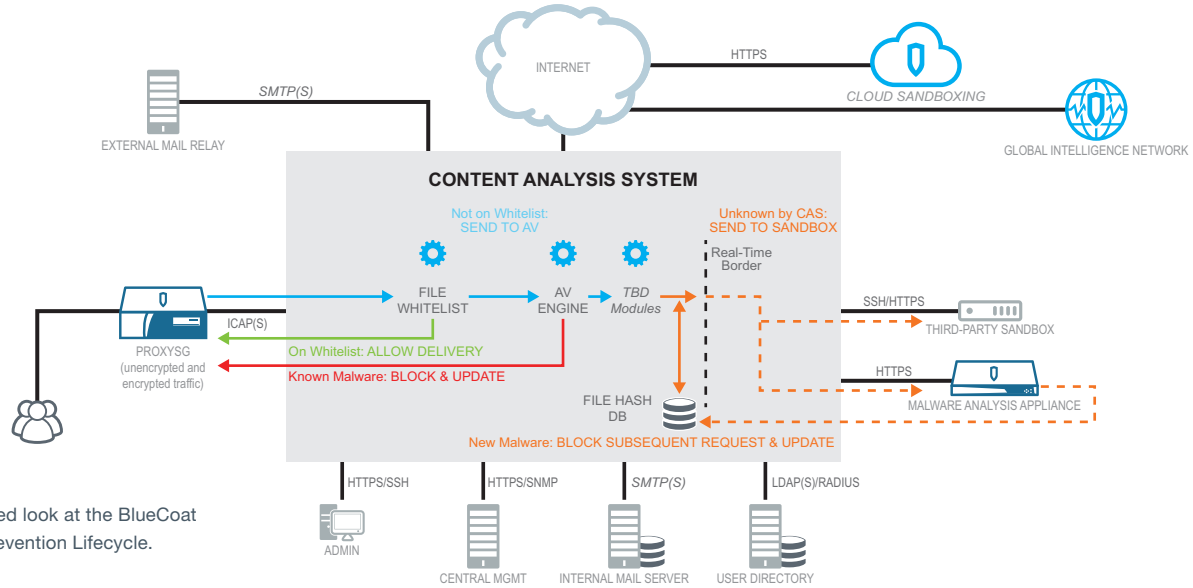

Here's a more detailed look at the BlueCoat Advanced Threat Prevention Lifecycle.

BLUE COAT®

Security Empowers Business

v.WP-DEFEATING-NEXT-GENERATION-MALWARE-WITH-NEXT-GENERATION-ANALYSIS-EN-v1a-0614

Blue Coat Systems Inc.
www.bluecoat.com

Corporate Headquarters
Sunnyvale, CA
+1.408.220.2200

EMEA Headquarters
Hampshire, UK
+44.1252.554600

APAC Headquarters
Singapore
+65.6826.7000