

F-Secure DeepGuard

Proactive on-host protection against new and emerging threats

1. The case for proactive behavioral analysis

One of the most demanding challenges security programs have had to address in the last few years has been the increasing diversification of attack vectors through which malware can arrive onto a host machine, especially as more applications, networks and services become hosted on or accessible over the Internet. This has been of particular concern with the growing popularity of online-based attacks that exploit vulnerabilities in applications installed on a machine in order to run malicious code.

Some of the difficulties involved in dealing with modern attacks stem from major changes in the threat landscape that have taken place in the last ten years or so, including:

Exponential growth in malware

Since the mid-2000s, when malware creation kits that automated the process of producing malicious programs first became widely available, the numbers of malware samples seen by antivirus labs have grown exponentially, with hundreds of thousands of new or variant strains being created and propagated every month. In addition to the overwhelming numbers, many of these variants are designed to live only for a short time, sometimes only days or hours, in a deliberate attempt to overwhelm antivirus programs by sheer volume.

Attacks move online

The days when malware was most commonly distributed via e-mail attachments are long gone. Today, the most common attack vector is through a silent drive-by download during a visit to a compromised legitimate site or a malicious website that hijacks traffic from search engines or compromised sites. By moving distribution from direct delivery to the target machines to the nebulous online world, malware distributors and attackers not only increase their target audience but also make it much harder to prevent infections. Without a mechanism to identify the attack site and prevent users from visiting it, the user's machine can be successfully exploited without any overt sign that an attack has occurred.

Malware becomes a cybercrime tool

The consequences of an infection have also changed as organized criminals increasingly engage in cybercrime. Data and identity theft and monetary fraud are all criminal activities that have in recent years been facilitated by malware, in some cases in staggering amounts. For example, the United States Federal Bureau of Investigation (FBI) reported in a 2012 Senate hearing ^[1] that \$14 million in "illegal fees" were generated in the 2011 Ghost Click click-bot operation. With most real-world authorities lacking the resources or political will to prosecute cybercrimes, there is strong monetary incentive for cybercriminals to continue and improve their online activities.

Overview

This whitepaper explains the trends and developments in computing that have made host-based behavioral analysis and exploit interception necessary elements of computer security and provides an overview of the technology and methodology used by DeepGuard, the Host-based Intrusion Prevention System (HIPS) of F-Secure's security products.

DeepGuard introduces dynamic proactive behavioral analysis technology that efficiently identifies and intercepts malicious behavior. In 2013, an exploit interception module is being introduced that recognizes and blocks attempts to exploit vulnerabilities in installed programs, preventing malware infection. DeepGuard provides lightweight and comprehensive endpoint protection with minimal impact to the user experience.

Key Features

- Updatable scanning engine uses the latest detections to protect against emerging threats
- Continued application monitoring protects against delayed malicious actions
- Exploit interception module recognizes and blocks exploit attempts, including document-based attacks

Benefits

- Provides immediate on-host protection against known and new threats, even before signature databases are updated
- Intercepts exploit attacks against programs installed on the machine
- Recognizes and blocks suspicious activity
- Reduces potential loss of sensitive data or privacy due to malware infection

“MALWARE IS CONSTANTLY EVOLVING, WITH NEW TRICKS AND FEATURES. BUT ONE THING REMAINS CONSTANT - MALWARE WILL ALWAYS EXHIBIT MALICIOUS BEHAVIOR.”

Mika Stahlberg
Chief Technology Officer, F-Secure Labs

Popular software is heavily targeted

Although almost any software can contain vulnerabilities, of particular interest to cybercriminals and other attackers are vulnerabilities in popular applications, such as Java Runtime Environment (JRE), Adobe Reader, Microsoft Office and web browsers. These programs typically have millions of users, making them prime targets for attack.

Many of these applications have multiple known vulnerabilities, and though most are fixed by security patches released from the vendors, the time needed to develop and deploy these fixes to all affected machines still leaves an interval in which the users are vulnerable. Additionally, new or zero-day vulnerabilities are periodically found for which no patches are yet available, leaving the users wide open for exploitation.

Exploit kits make attacking easier

The advent of commercial-grade exploit kits such as BlackHole, Cool Exploit or Sweet Orange, which automate the process of scanning and exploiting a user's machine within seconds of a

visit to an attack website, have significantly lowered the level of technical expertise needed for cybercriminals to successfully infect new victims with malware.

Exploit kits have transformed vulnerability exploitation from a niche activity into a common attack vector. The increasing number of malware being distributed using exploit-

based methods have in turn led to a need for on-host security solutions that are able to identify and block attempts to exploit vulnerabilities in installed programs, before malware can be successfully dropped onto the machine.

Targeted attacks make detection harder

More focused targeted attacks can involve more obscure exploits and delivery mechanisms. These attacks typically use document or executable files carefully crafted to fit the profile of the intended victim, taking into account their topics of interest, preferred operating system and any security programs they may be using. The highly specific nature of these attacks makes them particularly difficult to detect using traditional signature-based detections.

Identifying clean programs becomes more critical


The number of clean or non-malicious applications globally available today runs into the millions, far more than the normal user is likely to be familiar with at any one time. The abundance of programs, their easy accessibility over the Internet and the need to stay abreast of constant program updates all makes it cumbersome for security solutions to depend solely on local user-driven white- and black- listing to provide adequate protection.

The majority of programs seen on a typical machine are clean, so correctly identifying non-malicious software is a significant step towards pinpointing truly harmful programs for further attention. Eliminating false positives on clean files is also critical in optimizing a security program's performance and of course, minimizing interference with the user's experience.

Given the various challenges presented by today's more complex computing realities and more fluid threat landscape, traditional signature-based scanning is now just one layer of a multi-tiered approach to endpoint security. Cloud-based file and web reputation checking, HIPS (Host-based Intrusion Prevention System) and behavior analysis have all become integral components of the modern proactive protection system.

2. Multi-layered protection


F-Secure's multi-layered approach to security is comprised of the following modules, each designed to address a particular aspect of the threat landscape and work together to provide a complete solution:



As mentioned before, most attacks and malware downloads today take place online. Ideally, protection should begin even before the machine environment is reached, by preventing exposure to possible infection points - and so, enter Browsing Protection.

To prevent users from inadvertently visiting compromised legitimate or outrightly malicious sites, Browsing Protection provides critical assessment of a website's security. If the site is known to be

CHART 1: MOST PREVALENT EXPLOIT KITS ONLINE, Q1 2013^[2]



malicious, or contains features that render it suspect, the user is cautioned against entering it. To deal efficiently with the millions of sites available on the Internet and their constantly fluctuating changes in security, Browsing Protection's functionality is based on lookup queries to F-Secure's Security Cloud (see page 4), which includes a database of known safe and malicious files and websites. The entries are updated automatically in real-time based on rules maintained by response analysts.

Though Browsing Protection is able to prevent most visits to known malicious sites, it's always possible to stumble onto an unrated or newly compromised or malicious site, or for malware to be introduced onto the host machine some other way, perhaps on removable media. If a suspect file does successfully arrive on the machine, it is then subjected to multiple layers of security checks.

Whenever a file arrives on a machine, is installed or modified, it is first scanned using a traditional signature detection engine to determine if it is a known threat. The scanning engine uses custom, family, generic and heuristic detections, which respectively identify specific malware, families of malware with similar features, and broad ranges of malicious physical features and behavior patterns. If the file's characteristics match those of previously seen malware, it is blocked.

Though often overlooked in favor of more sophisticated technology, signature-based scanning is still an effective method of identifying and blocking the vast majority of malware seen to date, protecting users against lingering threats such as Downadup or Melissa, which debuted and peaked years ago but are still present in the wild, where they continue to infect new victims. The effectiveness of this check depends on keeping the signature database updated with the latest detections.

If the file isn't identified as a known threat, a query is sent to F-Secure's cloud infrastructure to gather the latest metadata available for the file. Analysis is subsequently handled by **DeepGuard**, which collectively handles all the behavioral analysis, process monitoring and exploit interception of suspect files, both at the point of application launch and during execution.

3. More about DeepGuard

Put simply, DeepGuard observes an application's behavior and prevents any potentially harmful action from successfully completing. The apparently simple nature of this task belies its importance however, as this proactive, on-the-fly monitoring and interception serves as the final and most critical line of defense against new threats, even those targeting previously unknown vulnerabilities.

Behavior-based analysis addresses the Achilles' heel of signature-based scanning: the need for analysts to have an actual sample of the malware in order to create the signature to identify it. Given the huge numbers of malware constantly being created and distributed, new threats will often be able to successfully infect at least one victim in the wild before most antivirus labs are able to acquire a sample, analyze it and issue a detection.

Behavior-based detection covers that crucial gap between the first appearance of new malware and the first signature detection being issued for the threat. By moving the focus from unique physical characteristics to patterns of malicious behavior, DeepGuard can identify and block programs performing harmful actions, even before an actual sample has been acquired and examined.

THE ROAD TO DEEPCLOUD

2006

Heuristic analysis technology introduced

DeepGuard 1.0 introduces behavioral analysis to complement existing signature-based detection technology. When a program is launched, DeepGuard performs two tests - a static check for features commonly found in malware and emulation of the program in a virtual sandbox to evaluate its behavior. Programs that show no features or behavior matching known malware are allowed to execute as normal; those with tell-tale characteristics or malicious routines are blocked from execution

2008

First AV product to incorporate cloud lookups

In addition to signature scanning and emulation, DeepGuard 2.0 queries the Security Cloud for an almost instantaneous check of a suspect file's reputation. Response Labs analysts constantly monitor and update file reputation information, providing crucial human intelligence to the automated process.

2010

File metadata used in DeepGuard detection logic

In addition to signature detection and behavioral analysis layers, DeepGuard 3.0 includes a component that uses a file's metadata - e.g., the file's rarity, when it was first seen, related objects, and more - to gauge its threat potential. This feature allows malware to be identified using reputation-based factors such as whether the file was downloaded from a known malicious site, without needing further examination of its features or behavior

2011


Prevalence logic increases effectiveness against rare files

DeepGuard 4.0 revises the scanning engine to use updateable detections and beta detections for false alarms reduction. It also improves the prevalence logic used to identify files that are both rare and malicious, a feature that proves decisive in winning both AV-Comparative's *2011 Product of the Year* award and AV-Test's *2012 Best Protection Award* ^[3].

2013

Enhanced protection against exploit-based attacks

Malware infections facilitated by exploits targeting vulnerabilities in common applications have become a favored attack vector. DeepGuard 5.0 introduces enhanced behavior-based detection logic, including a module that monitors the runtime behavior of commonly targeted programs and potential attack files. This broad behavioral analysis approach allows DeepGuard to identify and intercept exploit-based attacks, regardless of the specific vulnerability targeted



For example, out of all Zeus crimeware infection attempts reported in April 2013, 80% involved previously unseen variants. In those cases, DeepGuard successfully prevented infection by recognizing the file's malicious behavior and blocking the attack. Subsequently, signature databases were updated to identify these samples, but for users facing new threats, DeepGuard's proactive analysis provides immediate protection against infection.

In 2011, an entirely rewritten DeepGuard engine was introduced that included (among numerous other improvements) a switch from using hard-coded scanning logic to an updateable detections database. Response Labs analysts constantly monitor the threat landscape and analyze the latest threats in order to determine the best way to identify malicious behavior. Being able to update the scanning engine with the results of this research keeps DeepGuard consistently effective against the latest threats.

Given the short-lived nature of most malware variants, signature detections tend to have narrow windows of effectiveness before the malware they detect 'expire'. In contrast, DeepGuard detec-

tions can effectively identify malware over a much longer time period, as malware behavior is much less mutable. For example, on 12 July 2012, DeepGuard was updated with one new detection, while the signature database received 600 new additions. Nine months on in March 2013, tests run using the same database set against a random collection of more recent malicious samples showed the DeepGuard detection blocking 12 times more infections of the newer malware than the 'aged' set of signature detections.

The proactiveness and longevity of DeepGuard detections is illustrated in Chart 2 (above), which is based on detection statistics from F-Secure's internal systems for Urausy ransomware variants. The DeepGuard detection was able to identify variants (and therefore block attempted infections) earlier and continued to do so for longer, while the equivalent signature detection peaked and then declined rapidly, as newer Urausy variants appeared. (The reason for the signature detection's higher peak is due to it being a previous defense layer to DeepGuard. Had those signature detections been missed, it would have been DeepGuard with the high peak.)

Security Cloud

In operation since 2008, the Security Cloud (formerly known as the Real-Time Protection Network) is F-Secure's cloud network, housing the various databases and automated analysis systems that support and enhance the performance of F-Secure security products installed on client machines. The infrastructure for this network is hosted on servers in multiple data centers around the world.

Client machines that connect to the Security Cloud are able to retrieve the most up-to-date details of threats seen in the wild by other protected machines, making response far more efficient and effective. When a new object, such as a file or URL, is encountered on one client, the product communicates with the Security Cloud using the strongly encrypted Object Reputation Service Protocol (ORSP) to query for the object's reputation details. Anonymous metadata about the object, such as file size and anonymized path, are sent to the Security Cloud. These queries are completely anonymous and the IP address is not stored, maintaining the client's privacy.

By evaluating the metadata sent, together with information drawn from the in-house databases and various other sources, the Security Cloud's automated analysis systems (which make up to 8 million decisions per day) can provide a fully-informed, up-to-date risk assessment for the object during DeepGuard's pre-launch security evaluation stage, immediately blocking a threat that has been previously seen by any other machine connected to the Security Cloud. This also removes the need to perform further analysis of the object on the client, reducing impact on the user's experience.

The Security Cloud also allows Response Labs analysts to provide critical human intelligence and judgment to complement the automated systems and on-host scanning technology. In addition to creating and maintaining the rules that underpin the databases and automated analysis systems, analysts actively monitor the threat landscape and research malware characteristics and behavior patterns to find the most effective ways to identify truly malicious programs. Once a threat has been confirmed (or a known file's reputation is modified), the updated details take 60 seconds to replicate across all products connected to the Security Cloud, ensuring up-to-date protection.

DeepGuard's updateable detection logic is especially useful in countering attacks that exploit vulnerabilities in installed programs in order to run malware on a machine. In such cases, the dropped malware itself can be spotted and blocked by signature or behavior-based scanning. To halt the attack at an even earlier stage however - that is, at the point of exploitation - Response Labs analysts examine the exploit mechanism for tell-tale actions or behavior patterns, and then incorporate the research results into DeepGuard's scanning engine. It is then able to pinpoint and block suspicious actions that bear the hallmarks of a vulnerability exploit attempt, preventing malware from being dropped on the machine at all.

By taking into account characteristic exploitation mechanisms as well as the features and behavior of malware being dropped on the system, DeepGuard can effectively identify and block threats on the fly, even when faced with totally new malware targeting zero-day vulnerabilities.

4. How DeepGuard works

DeepGuard's behavioral analysis is activated by two events. When a program is launched for the first time, DeepGuard analyses it to determine if it is safe to run. Subsequently, DeepGuard continues to monitor the program while running.

4.1 Pre-launch analysis

When a program is first executed, regardless of how it is launched (the user clicks the file icon, an e-mail attachment or program initiates it, etc.), DeepGuard temporarily delays it from executing in order to perform the following checks:

File reputation check

If an Internet connection is available, DeepGuard sends a query to the Security Cloud (see page 4) to check for the latest information on the program's reputation in the clean file database, which contains the latest security evaluations for a vast catalog of commonly used applications. This database is maintained and constantly updated by Response Labs analysts. Programs that have been rated as clean in the database are allowed to bypass additional checks and launch immediately, whereas known malicious files are blocked at once.

For the user, the clean file cloud lookup functionality offers a number of advantages. Being able to use the security verdict for a known file from the clean file database not only removes the burden of identifying unknown or unfamiliar programs as legitimate or malicious from the user, it also means unnecessary security checks on clean files can be avoided. At the same time, by reducing to a manageable level the volume of software that needs to be individually evaluated, the ability to still white- or black-list selected programs becomes more meaningful. And finally, even when the product's signature databases are outdated or rarely updated, DeepGuard can still use the most up-to-date file reputation information to fine-tune its analysis.

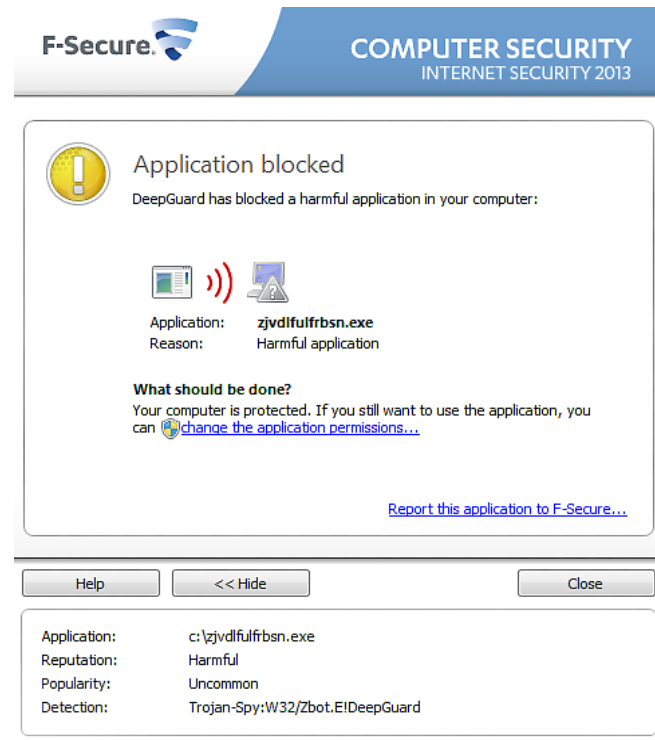


Image 1: DeepGuard blocks a harmful application

Behavioral analysis

If the program is flagged as suspicious during the file reputation check, or if Internet access is unavailable, DeepGuard executes it in a virtual environment and observes its behavior for malicious actions, such as attempting to self-replicate, edit or delete critical system files, and so on.

Response Labs analysts continually research and update DeepGuard's scanning logic with detections for the most effective behavior patterns needed to spot malware. These detections may identify specific malware families (which typically share similar features or behavior) or they may more generally identify suspect actions, such as attempting to hide from process enumeration programs, which are indicative of malicious intent. The analyst's ability to tweak DeepGuard's engine in this manner permits an element of human discretion and flexibility, to provide a more fine-grained and ultimately more accurate analysis.

Prevalence rate check

DeepGuard includes a module that focuses on a file's prevalence rate. Clean files typically have thousands or millions of users, making them highly prevalent. In contrast, malware samples are comparatively rare. According to statistics generated from F-Secure's internal systems monitoring known threats, in a random sample of malicious programs found in the first four months of 2013, 99.7% of the threats were rarely seen in our user base. Rare or new files are automatically considered more suspect and subjected to greater scrutiny during the subsequent process monitoring stage.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.