Intel Security, McAfee Labs White Paper

Understanding McAfee Next Generation Performance Technology

The next-generation of McAfee Anti-Malware technology is part of an architecture that provides significant new capabilities to counter the newest malware threats with next-generation speed and efficacy. Understanding the basic operation of the new architecture will help provide clarity on how the technology accomplishes this.

What Problems Does the New Technology Solve?

A market problem exists:

- The capacity and capabilities of endpoints has increased dramatically and multi-terabyte endpoints are now the norm in the Enterprise.
- The number and kinds of malware attacking endpoints has increased exponentially over the last few years.
- Previous generations of AV which involves scanning every individual file is not optimized for this new environment.

One of the key design goals of the next-generation McAfee Anti-Malware Engine Core (Code-named AMCore) is to provide top-tier performance and Enterprise-level protection by introducing an intelligent strategy to only scan items that really need to be scanned, instead of scanning all items equally.

Authors: Graham Strachan, Senior Director of Engineering, John Teddy, VP Engineering, Robert H. Leong, Director of Product Management, McAfee Labs
May 2015, Version 1.3

AMCore Performance Model: The Actors in the Drama

In the new AMCore performance model, the strategy is based around the notion of an 'actor'. An actor is defined as a running process and its state of trust (as seen below). The state can be one of three values:

- Suspicious: The actor has done something or has come from someplace on the web that is not trusted, e.g. a website that is known to host malware downloads.
- Normal: This is another way of saying "Unknown". The actor is in a state that we have not determined to be suspicious or trusted yet.
- Trusted: The actor is trusted, e.g. is a known part of the OS, has come from a trusted origin (e.g. Microsoft Signed Installer), or is directly related a trusted package (e.g. a DLL that is part of a trusted application package)

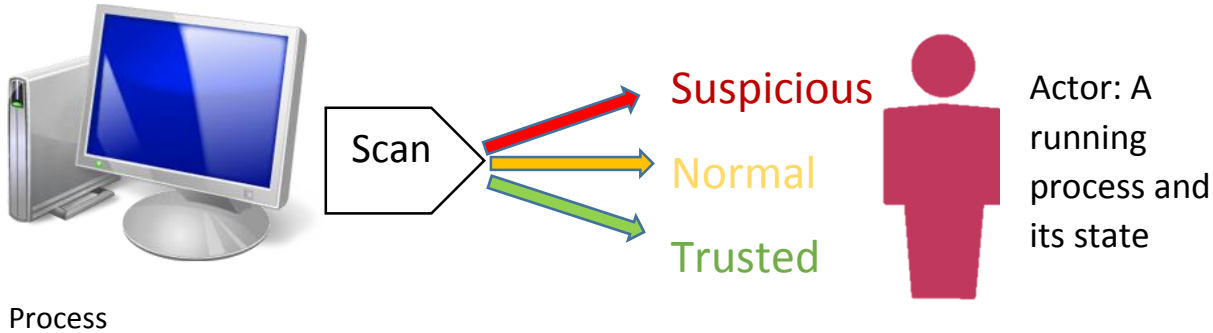

Why is the process referred to as an "Actor"? When doing behavioral observations, the "actor" is the process that is "performing" events. It is nomenclature that is used to reflect the process that is "acting", that is being observed, hence the name "Actor".

AMCore does an initial scan of the Actors in an endpoint, and does an initial classification of its state, based on the body of knowledge that AMCore has available to it. This initial scan is extremely low-impact (i.e. it is very fast) as it is not the same as doing signature scanning – rather it is performing a rapid initial state classification.

The Process is a combination of base image, loadable images, and what the process has being doing – events and/or read/write operations. The base image and loadable images are objects that are potentially hash-able, based on whether we have a priori knowledge of whether that process is malicious or not, whether they have a signed certificate, or a hash value that is

Authors: Graham Strachan, Senior Director of Engineering, John Teddy, VP Engineering, Robert H. Leong, Director of Product Management, McAfee Labs
May 2015, Version 1.3

known to be trusted. Base images and loadable images are used to help classify actors as being trusted.

The events are behaviors, and reflect "what the file/process has been doing", and can contribute to whether we classify the file as suspicious or trusted, as can the read/write activities of the process. The events and read/write operations help AMCore to classify actors are being malicious.

AMCore has built-in behavioral classification technology that can help to quickly classify an actor as potentially malicious or not. A key point in understanding how this methodology can increase performance is realizing that AMCore can classify large groups of actors as trusted (e.g. through inheritance, canonical relationship, certificate signing, etc.) and that avoiding signature scanning on large groups of actors increases the performance while protecting the endpoint at the same high level.

AMCore Performance Model: To Scan or Not to Scan, that is the Question

The core idea is this: If it is necessary to signature scan a file, then it takes time to scan that file. If you can avoid signature scanning the file you will save time (i.e. increase performance).

The reader might rightly question this strategy, as doing a full certificate validation of a file takes longer than to signature scan a file, so checking the 'trust level' of a file through certificate validation might seem counter-intuitive. However, when you realize that doing a full certificate validation of a file will also impute trust to inherited items that are directly related to the file whose certificate was just validated, then the value of this approach becomes clear.

For example, if there is an installer or executable file that has a valid Microsoft certificate, then yes, it does take time to retrieve the certificate and validate it. However, that installer may spawn many other files that inherit the trust of the installer without having to empirically check the trust of every file that the installer places. This dramatically reduces the signature scanning burden and illustrates the core of the AMCore performance model.

Of course, there are obvious exceptions. For example, there are certain applications like browsers, where the browser binary is trusted, but the items the browser loads cannot be, as it is known that the browser can be instructed to load things that are malicious in nature. Microsoft Office is another good example, where the components and libraries of "Office proper" can be trusted, but documents that are ingested into Microsoft Office cannot all be trusted.

However, the performance model still has validity in that most of the items that trusted applications load can be trusted in the same way as the application that is ingesting the item. The bottom line is that this approach can be used to dramatically reduce the number of items that require signature scanning and certificate validation.

Therefore, increased performance is directly related to "scan avoidance", and having trusted certificates for items like trusted installers is key to providing good performance. Conversely, if

Authors: Graham Strachan, Senior Director of Engineering, John Teddy, VP Engineering, Robert H. Leong, Director of Product Management, McAfee Labs
May 2015, Version 1.3

the system is missing trusted certificates for an installer that should be trusted, then the AMCore performance model will drop back to the default mode of certificate scanning every file that is placed by the installer in question. Knowing what to trust is key to performance. Maintaining an up-to-date whitelist is therefore important to this performance model.

AMCore Performance Model: The "Fastest Path"

Key to understanding the AMCore performance model is the notion of "Choosing the Fastest Path":

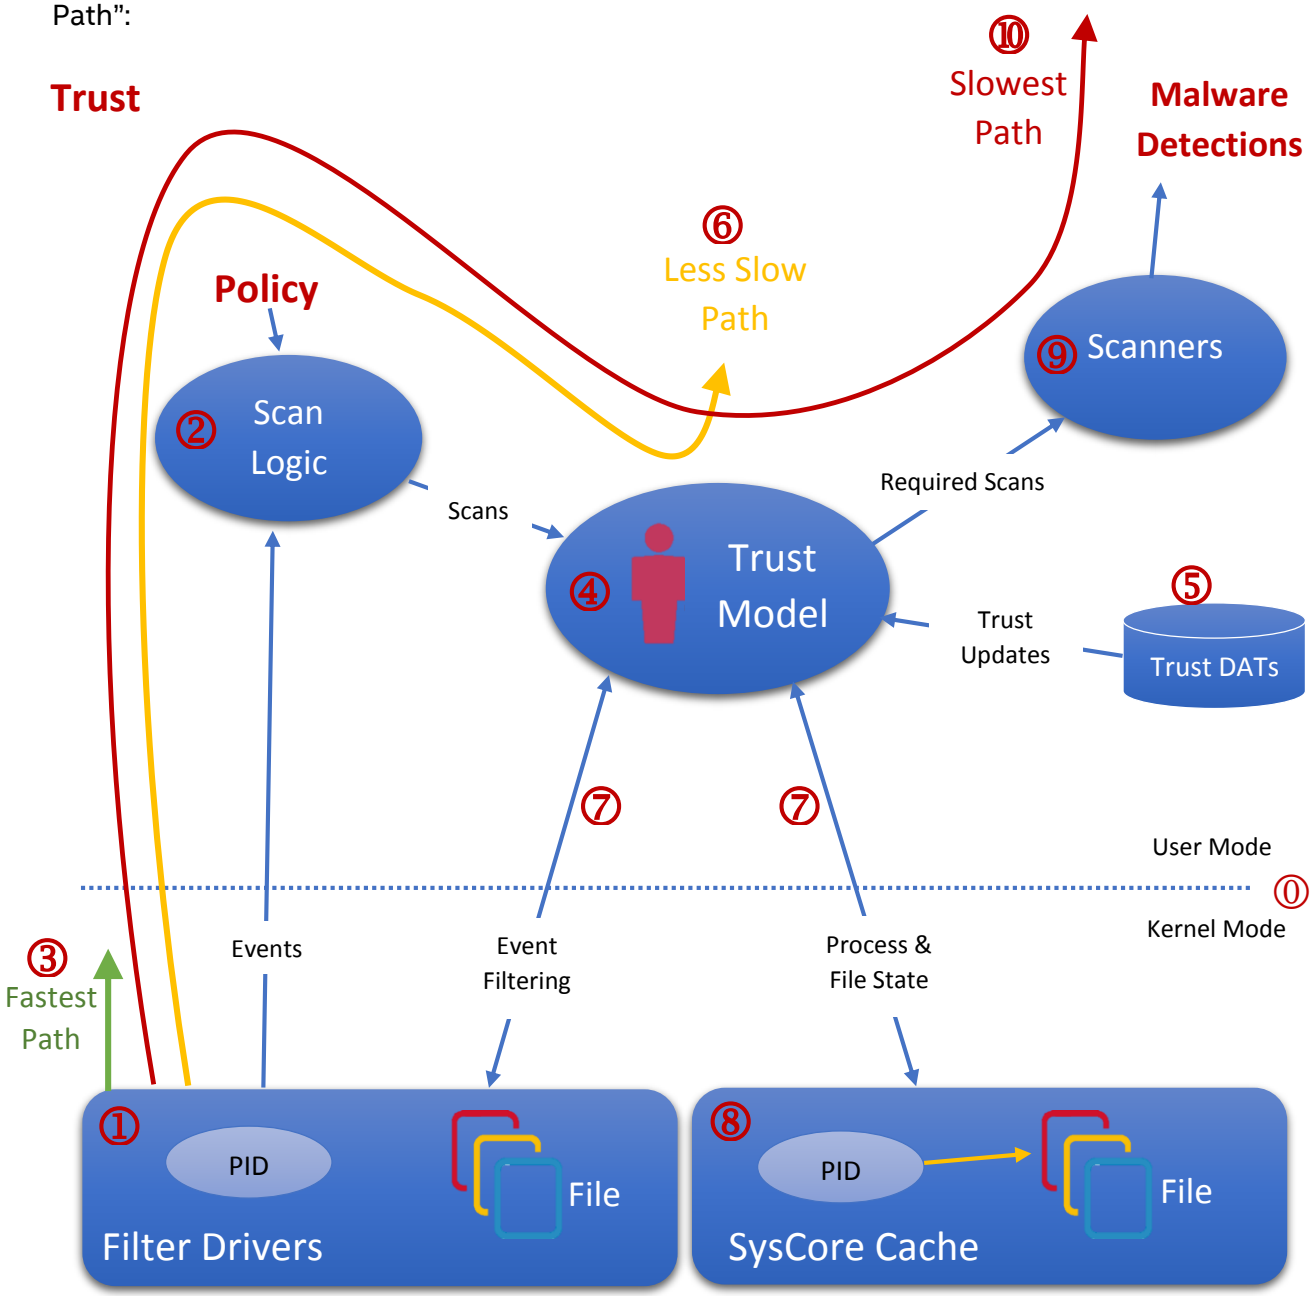

Authors: Graham Strachan, Senior Director of Engineering, John Teddy, VP Engineering, Robert H. Leong, Director of Product Management, McAfee Labs
May 2015, Version 1.3

Understanding the figure above helps to explain the AMCore performance strategies with respect to certificate checking, scanning, caches, trust and more. (Look for the indicator symbols, e.g. ⓪①②③to locate the items referenced by the descriptions)

1. The dashed line ⓪shows the separation between kernel mode and user mode code.
2. The Filter Drivers object ① is the trigger point for everything the system does for performance. It triggers on events related to file loading, file execution, and other low-level system events.
3. When the Filter Drivers object① triggers on these events, it filters them according to some specific criteria.
4. The Scan Logic object ②receives Policy instructions that it sends to the Filter Drivers Object ①. This allows the Filter Drivers Object to very quickly filter and sort according to certain filter criteria, such as
    a. File type
    b. File location
    c. Previously cached file information for malicious or trusted file
5. The use of the Filter Driver Object ①+ the Policy-based Scan Logic ②is the "Fastest Path" to resolution of the trust question and can very quickly determine the state of trust of the file that caused the trigger. This is the meaning of the green arrow③ showing the Fastest Path to state of trust of a file.
6. If the Scan Logic Object ②cannot immediately make a determination on a file (trust or block, etc.), then it passes decision-making concerning the file to the Trust Model④.
7. The Trust Model ④ performs more sophisticated analysis of the file utilizing content from the Trust DATs⑤, including:
    a. Does the file have a match in a known good whitelist?
    b. Does this file require a certificate validation?
8. Employing the Trust Model is the "Less Slow Path"⑥, as this analysis clearly takes a little more time. It is not as fast as the "Fastest Path"③.
9. There is also some logic in the Trust Model④ concerning how to handle inheritance of trust. For example, if the system determines that the file in question (that triggered the scan) is the child of a trusted parent, then we can infer trust upon the new child and the child is consequently trusted. This kind of information is then stored in the Scan Logic Object ②for quick retrieval should another "child of inheritance" is encountered by the Filter Driver Object ①.
10. As a result, the system will learn from its analysis as time goes on, and it will become faster at finding and recognizing trusted files, and the use of the "Fastest Path" ③ will become more prevalent.
11. The lines⑦ leading from the Trust Model④ to the Filter Driver Object ① and the SysCore Cache Object⑧ indicate that information on previously-filtered events

Authors: Graham Strachan, Senior Director of Engineering, John Teddy, VP Engineering, Robert H. Leong, Director of Product Management, McAfee Labs
May 2015, Version 1.3

and process and file state go from the Trust Model to the other two objects and are cached in those objects for quick retrieval in the future to speed determinations of files for which the system has prior knowledge.

12. There are two levels of cache, one in the user mode and one in the kernel mode. The User Mode cache is more flexible while the Kernel Mode cache is faster. When the system can, it will send relevant information to the SysCore Cache ⑧, but this only works for Process IDs (PIDs). Other kinds of information cannot be stored in the SysCore cache. The SysCore cache can be used in the Kernel Mode to assist in determining the state of trust of a file very quickly.
13. The Filter Driver Object ① can impose some level of up-front filtering. The Filter Driver Object receives requests from the Windows OS. The Windows OS will call into the Filter Driver Object whenever certain events occur, such as:
    a. Registry Operations
    b. File Operations
    c. Network Operations.
14. The Filter Driver Object ① will relay information to the User Mode (which houses the entire family of McAfee Security Connected applications) which houses the AMCore code. Each time information is relayed across the User / Kernel Mode boundary, a small time penalty is exacted. This is normal and occurs in all computing devices that use this architectural model. In order to optimize system performance, the system has programmed in low-level filtering into the Filter Driver Object itself. This improves performance because the filtering can occur without moving information back and forth across the User / Kernel Mode boundary. The low-level filtering occurs for things like folder filtering, sub-event filtering and more. This avoids having to send all events up to the User Mode simply to have the Objects in the User Mode layer reject the information because the object (such as a folder), cannot by definition be a malicious object.
15. If the Trust Model ④ cannot determine whether a file can be trusted through a whitelist or inheritance lookup, then the system must pass analysis to the Scanners Object ⑨ for further analysis. This defines the “Slowest Path” ⑩.
16. In general, the Fastest Path of filtering according to the SysCore cache + the basic filtering of the Filter Drivers existed in VSCore. However, the Trust Model is a completely new technology object and is a fundamental change to the Enterprise Endpoint.
17. In the previous generation of technology, in the absence of the Trust Model, the Scanners Object ⑨ would feed cache information directly to the SysCore and Filter Drivers Caches. Now, the Trust Model ④ has taken over this responsibility.
18. The Major Change to this performance model is therefore the functionality of the Trust Model. By employing the notions of inherited trust and shared certificate trust, this provides a quantum leap in performance abilities for this technology. This is especially useful for things like trusted installers that may place thousands

Authors: Graham Strachan, Senior Director of Engineering, John Teddy, VP Engineering, Robert H. Leong, Director of Product Management, McAfee Labs
May 2015, Version 1.3

of files all across the endpoint. The AMCore Trust Model can dramatically reduce the amount of time that would otherwise be spent scanning files that should be inherently trusted.

19. The Trust DAT which is shown feeding into the Trust Model is a 'slice' of the AMCore DAT that is content loaded into AMCore on a regular basis. The AMCore DAT (not shown in the figure) is a structured container for:
    a. Scanners: AV Engine, Sample Profiler, Trusted Source
    b. AV Content: MinDAT, Medium DAT
    c. Trust: Trust DAT
    d. Assessors: Sample Profiler
    e. AMCore Content: Business Logic
    f. Secure Containers: VTP.bin

    These combine the more traditional technology along with the new technology to provide a blended solution and a smooth release of new technology.

AMCore Performance Model: Performance for the Long Haul

This next generation of technology is explicitly designed to maintain top-notch performance with Enterprise-class protection for Intel Security McAfee customers. By understanding the underlying design architecture, you can better know how to take advantage of this technology and prepare for the tremendous next-generation advantages that adopting this kind of architecture can bring. You can be confident that protection and performance are with you every step of the way as you leverage new endpoints in your Enterprise.

What's coming up next?

This next generation of technology is already in use across many McAfee products, and is providing next-generation increases in performance. So, what's next? What innovations in Enterprise Security are coming next from McAfee?

Problem Statement

According to the Ponemon Institute Survey of Threat Intelligence & Incident Response (Feb 2014):

- 35% of all cyber attacks are undetected
- 86% say cyber attacks take too long to discover
- 85% say there is no prioritization of incidents
- 48% say their security products do not support importing security data from other sources
- 55% say they believe that their security team does not have the skills to investigate and remediate a security incident

In summary, even with exemplary defenses, it is clear that some attacks will inevitably get through.

Authors: Graham Strachan, Senior Director of Engineering, John Teddy, VP Engineering, Robert H. Leong, Director of Product Management, McAfee Labs
May 2015, Version 1.3

As a result, McAfee Labs is exploring customer needs to better spot security concerns, understand risk, and react. We want to enable customers to be predictive and proactive when possible, and to prioritize resources when anomalies are spotted.

Currently, we are looking at four main areas for Threat Intelligence:

1. Analytics (on premise and in the cloud): analyzing available data to spot anomalies that are indicative of attacks
2. Endpoint threat detection and response: providing the sensor data to capture the anomalous behavior via lightweight forensic tools
3. Services: producing, sharing, and consuming IOCs, IOAs, and other threat data and feeds
4. Incident response: offering capabilities to remediate and report on impact

To learn more about these and other security topics, please contact McAfee Labs.

Authors: Graham Strachan, Senior Director of Engineering, John Teddy, VP Engineering, Robert H. Leong, Director of Product Management, McAfee Labs
May 2015, Version 1.3