

Delivering Zero-Day Defenses with Symantec Endpoint Security

Applying a Zero-Trust Model, with Application Isolation and Antimalware

WHITE PAPER



OCTOBER 15, 2019

Acknowledgements

Initiated and released by Symantec, this document was developed with support from the organization and in direct collaboration with the following:

Authors: Sheetal Venkatesh, Ashok Banerjee, Torry Campbell

Contributors: Deb Banerjee, Susan Hassall

Contents

When it Comes to Antimalware and Application Isolation, 1+1 = 3	3
Why Trusted Applications Are the Riskiest	3
Preventing Attacks from All Angles	3
Introducing the Powerful Combination of Blacklisting, Whitelisting & Application Isolation	4
Case Studies – Jails and Castles in Action	6
Case Study 1: Using Application Isolation’s Jail Mode to Prevent a Weaponized System Utility	6
Case Study 2: Using Application Isolation’s Castle Mode to Prevent a File-less Attack Using MS Excel	7

When it Comes to Antimalware and Application Isolation, 1+1 = 3

Increasing layers of defense have forced the attackers to change their approach, using file-less attack techniques that **'live off the land'** and are difficult to detect by traditional means. More and more attackers are taking advantage of what already exists on a device to carry out their objectives. They aren't downloading any new executables or files, so file-based detection methods aren't effective. They are exploiting vulnerabilities in common applications and using scripted content in document files to disguise their activity behind trusted applications. They are running attacks directly in memory and persisting in registry keys, using common scripting languages, such as Powershell, WMI, Javascript, WScript, etc., to avoid raising any suspicions.

To defend against these new 'living off the land' tactics, Symantec has combined application isolation and antimalware to give you a single solution that can keep up with emerging, zero-day threats. While antimalware detects and prevents malware from executing, application isolation assumes malware has somehow evaded detection and is already executing on your endpoints. When combined, application isolation compliments the benefits of antimalware, by proactively blocking malicious behavior using a zero-trust model for any application, both well known and possibly suspicious applications.

This paper explores in more detail how you can use these complementary application isolation and antimalware capabilities, delivered by Symantec Endpoint Security, to achieve a more holistic, layered approach to your security, without sacrificing productivity.

Why Trusted Applications Are the Riskiest

There are a number of applications that are critical for your productivity and day-to-day business operations. These include browsers, email clients, productivity applications (Microsoft Office), platforms tools (Java) and common development tools (Visual Studio), among others.

Most of these applications contain vulnerabilities that can be exploited by an attacker to take control over that application and get a foothold into your endpoint, and ultimately your network.

Some documents, such as Microsoft Word, Excel and Adobe PDF files, also allow scripted content, which means an attacker can run malicious code from seemingly innocent documents. Many of these applications use higher privileges to operate normally, particularly those that run in the context of privileged admin users. As a result, when attackers take control over these applications they have unrestricted access to the endpoint. An attacker can use a compromised application to download malware payload and then execute it within the context of that trusted application.

It is possible to mitigate some of these issues by patching applications and upgrading them to the latest versions. However, rolling out patches or updates can be challenging in large environments. Plus, given how frequently zero-day vulnerabilities crop up, even the most disciplined software update process has trouble keeping up. It is therefore critical to adopt defensive measures that prevent exploitation or tampering of applications to prevent attackers from using vulnerable applications to infect the environment.

Preventing Attacks from All Angles

The threat continuum paradigm is a useful way to analyze the right protection strategy for different threat vectors. If you classify all the files and applications on an endpoint, they will roughly fall into 5 broad buckets – 'threats,' 'potential threats,' 'unknown,' 'potentially good' and 'known good.'



Figure 1. The Threat Continuum.

An effective endpoint security strategy will deliver protection along the entire threat continuum. Known bad applications, 'threats,' should be blocked and eliminated immediately. 'potential threats' and 'unknown' applications should be monitored to identify and then stop behavior that could harm the OS or other applications. 'Potentially good' and 'known good' applications should be protected from exploits and monitored to prevent file-less attacks.

Delivering on this strategy requires a solution that combines multiple controls:



1. **Blacklisting** - follows a default allow model, allowing everything to execute freely, unless it has been characterized as 'bad.' This model enumerates the applications, processes, scripts, etc. that are known to be malicious and eliminates them as soon as they are identified, with high fidelity. If it's not characterized as bad, it's assumed to be 'innocent' and permitted. This has been the preferred security model for endpoints, such as end-user desktops and laptops that change frequently, because it is easy to deploy, with simple, effective policies. It is very effective at eliminating known threats, but it is not optimized to catch well known applications that have been compromised and are performing malicious actions.
2. **Whitelisting** - follows a default deny model, allowing only the applications that are whitelisted (per a policy) to run. Nothing else is trusted or allowed, which makes it very effective at reducing the overall attack surface. This model typically works well for environments that don't change frequently, but if it's well managed, can also be used for 'potential threats' that are changing frequently. It will catch and prevent zero-day threats, as long as they fall outside of what's allowed. It won't stop threat activity being performed by a whitelisted application.
3. **Application Isolation** - follows a zero-trust model, building on whitelist security to allow not just approved applications but also restrict the behavior of approved applications. For example, an application isolation policy could define the acceptable network connections, file activity, registry activity, etc. of an application, so it is restricted to known good behavior. Non-whitelisted applications could be allowed to run, but with very severe restrictions. This model is highly effective for reducing the overall attack surface. It can also mitigate zero-day attacks, by restricting an allowed application from doing something malicious, such as making changes to protected system settings or applications.

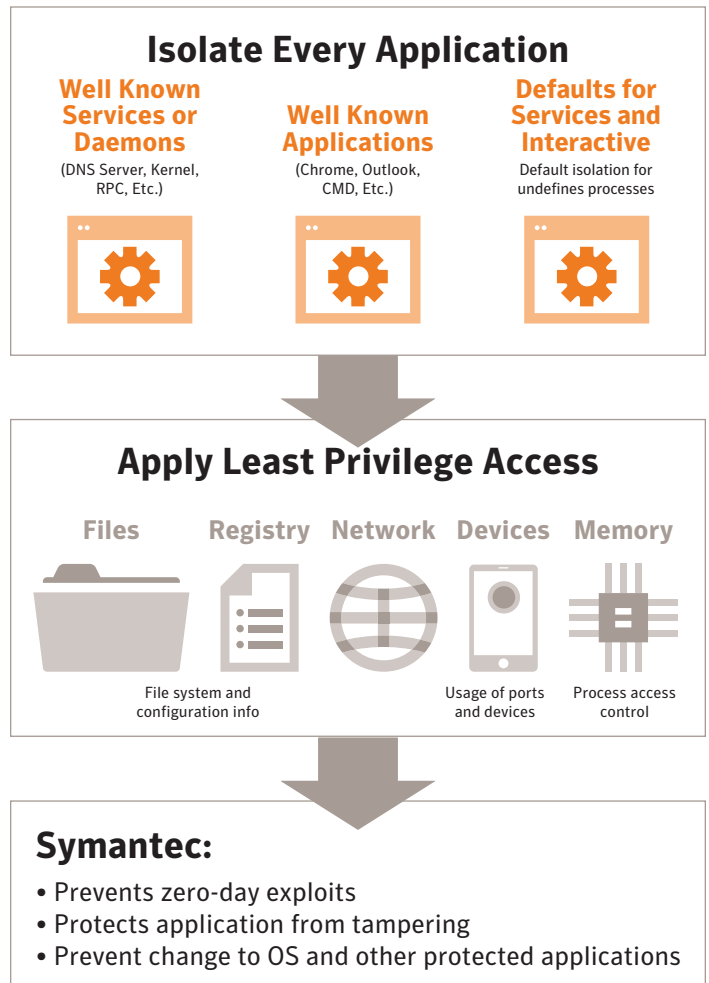


Figure 2. Zero Trust Model with Application Isolation.

Introducing the Powerful Combination of Blacklisting, Whitelisting & Application Isolation

Symantec has combined blacklist, whitelist, and application isolation security models for the first time in the industry to provide comprehensive multilayered protection against evolving threats targeting your endpoints. Until now, these complementary technologies have been used to protect non-overlapping endpoints. By bringing them together, protection coverage is maximized and a wide variety of attack vectors can be shut down.

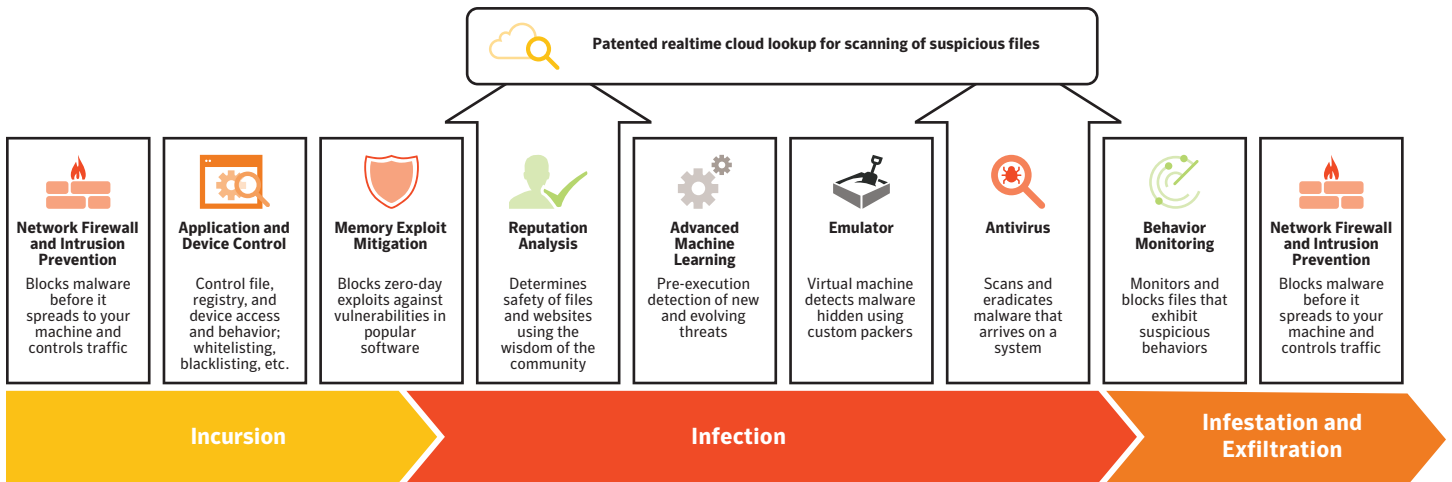


Figure 3. Layered defense strategy in Symantec Endpoint Security – Collection of antimalware engines.

Symantec Endpoint Security delivers the layered controls needed to protect endpoints. With a powerful combination of antimalware, device control, exploit mitigation, advanced machine learning, and behavior monitoring engines, Endpoint Security delivers the best threat prevention efficacy in the industry. As a result, Endpoint Security gives you the enhanced visibility you need to identify and appropriately defend against known threats and activities that are even remotely suspicious.

Endpoint Security’s new intensive protection feature allows you to tune threat detection engines collectively with a single knob, so you can block against all ‘threats’ and detect ‘potential threats’ in the continuum. For a file that is identified as a threat, intensive protection will delete or quarantine the file to stop it from doing any damage.

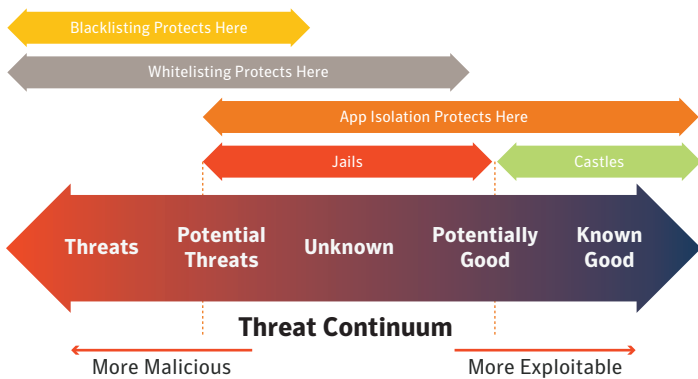


Figure 4. Complete coverage for the Threat Continuum.

For an application that is a ‘potential threat’ or ‘unknown’, intensive protection will flag the file as suspicious and hand it off to Application Isolation, a key feature of Symantec Endpoint Security Complete.

Application isolation can run these suspicious apps in ‘jail’ mode, which allows an untrusted application to execute within a Jail-like environment. It will allow the application to run with limited privileges to protect the OS and other good applications from any harm or tampering. It can contain items opened from an untrusted source (email or web, by example) to mitigate any risk they may pose and restrict these applications to only ‘good’ behavior. The same treatment can be applied to ‘potentially good’ applications which are often applications that may not exhibit suspicious behavior, but don’t have a credible reputation yet to trust them entirely. Application isolation can run these application in an ‘ankle bracelet’ type of jail where their behavior is largely unrestricted but some privileged operations, like modifying the OS or installing new applications, will be prevented.

Application isolation can also protect ‘known good’ (whitelisted) applications, running them in ‘castle’ mode to fortify these trusted applications and protect them from exploitation and tampering through a layered security approach. First, Symantec Endpoint Security’s Memory Exploit Mitigation engine protects an application’s process from a wide spectrum of exploit techniques against known and unknown vulnerabilities. Next, in the extreme event that an attacker gains control of an application’s processes, the attacker won’t be able to use its process privileges to install new software, change the system settings, or modify other application processes or resources. All the operations that the application doesn’t typically need to perform are blocked by the isolation policy. Note, the end user doesn’t perceive any change when using the application, unless the application engages in malicious behavior. This is a critical requirement for effective application isolation and ensures security does not come at the cost of productivity.

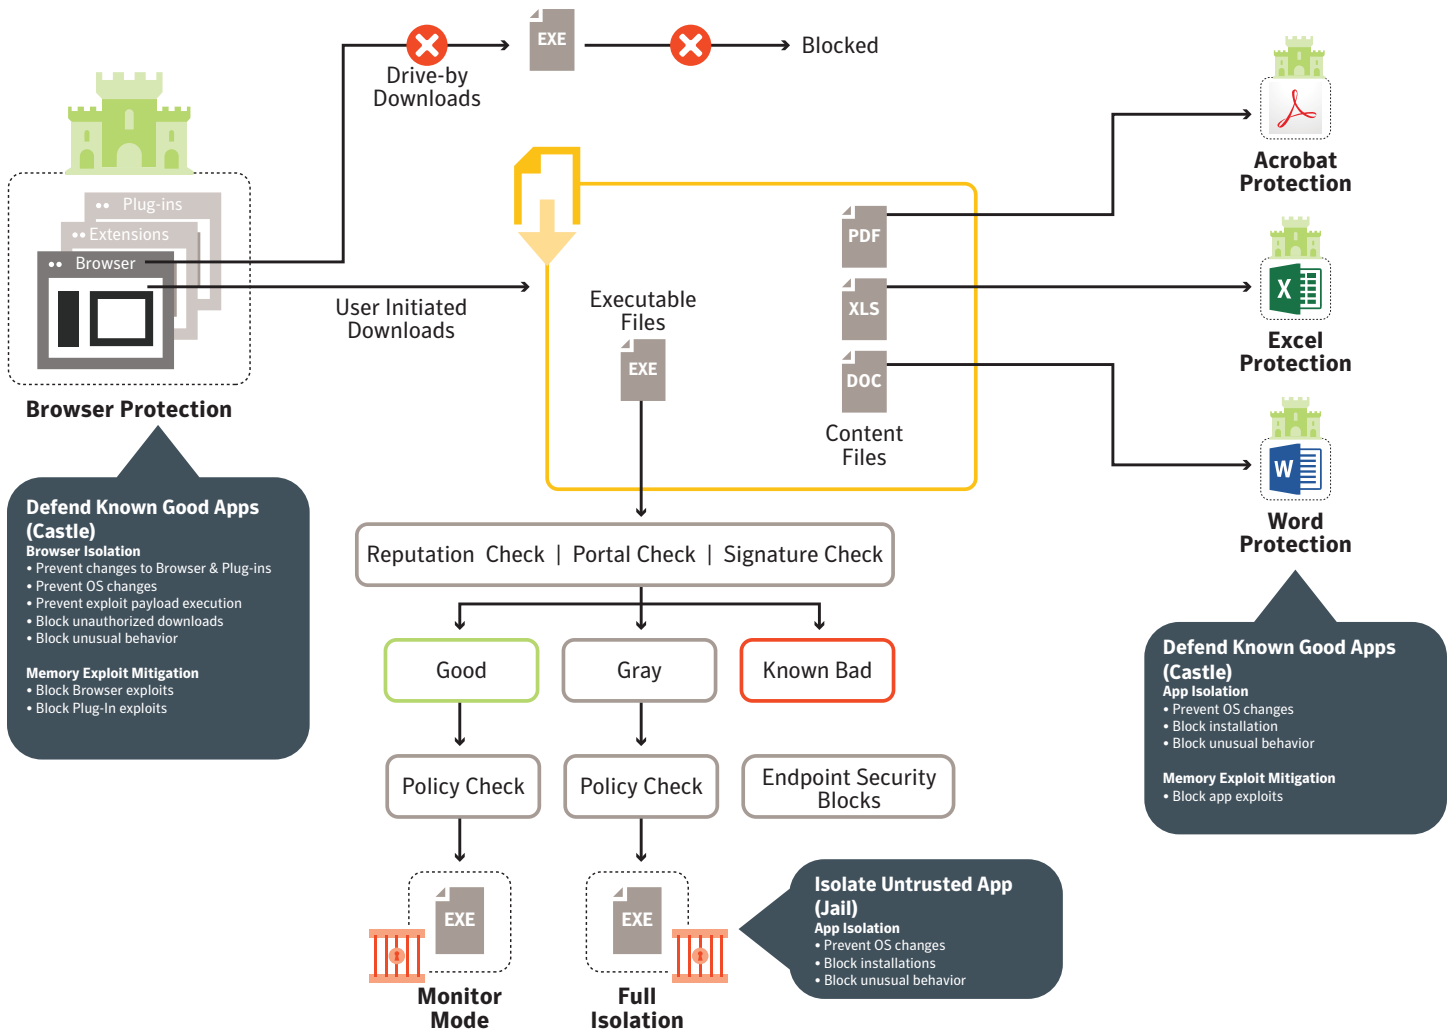


Figure 5. How 'jails' and 'castles' work.

This powerful combination of blacklisting, whitelisting and application isolation controls ensures that, even if an attacker gets past the formidable defenses of intensive protection, they will be isolated and prevented from propagating, rendering the entire effort fruitless.

Ease-of-use

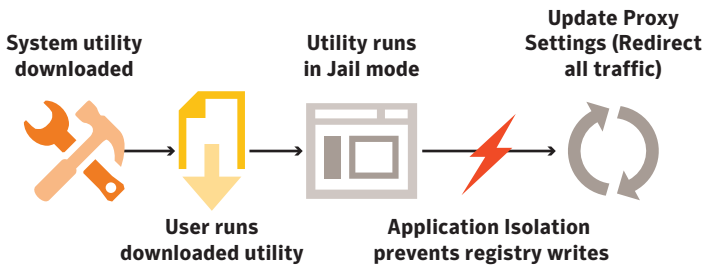
While security efficacy is critical, it does not have to come at the cost of productivity. This unique application isolation technology poses minimal performance impact on the endpoint and has no special hardware requirements. End users will not even realize that application isolation is in effect until an application engages in suspicious behavior. And because application isolation is able to assess ones and selectively jails only suspicious applications, the time spent by security teams in dealing with false positives is very low. Application isolation is easy-to-use with minimal performance and productivity impact.

Case Studies – Jails and Castles in Action

Case Study 1: Using Application Isolation’s Jail Mode to Prevent a Weaponized System Utility

Problem: Jen, an end user, is looking for a utility to ‘speed up’ her laptop’s performance. She visits a popular downloader website to check out a few utilities and downloads the ‘SpeedUp.exe’ app. Jen is obviously unaware that SpeedUp.exe is weaponized and will modify the Window’s proxy settings to redirect Web traffic to an attacker’s server.

Solution: Application isolation neutralizes this threat and prevents the attacker from gaining a foothold on the endpoint.



How it Works: As soon as SpeedUp.exe is downloaded, Symantec Endpoint Security performs an exhaustive analysis of the file. The intensive protection engine determines that SpeedUp.exe is suspicious ('potential threat'), but not confirmed to be malware. As a result, when the user launches the SpeedUp.exe application, Application Isolation automatically runs SpeedUp.exe in 'jail' mode. While running in the jail, SpeedUp.exe can render its user interface and examine system performance, as long as it is not attempting to read or modify any protected OS resources. However, when SpeedUp.exe attempts to modify the Windows proxy settings, by accessing the system registry, the jail will block the operation instantly and generate a security event. Jen will get a notification indicating that SpeedUp.exe was blocked from modifying the OS, but can continue to run the utility for its system performance functions. Note, Endpoint Security could be configured to prevent the download of any executable files from the internet. However, in this example, we have allowed all files to be downloaded to demonstrate how jailing addresses the threat.

Benefits: Application Isolation dynamically 'jails' suspicious applications and prevents any malicious changes to the endpoint, so Jen can continue to work, without putting her endpoint or the network it connects to at risk.

Case Study 2: Using application isolation's castle mode to prevent a file-less attack using MS Excel

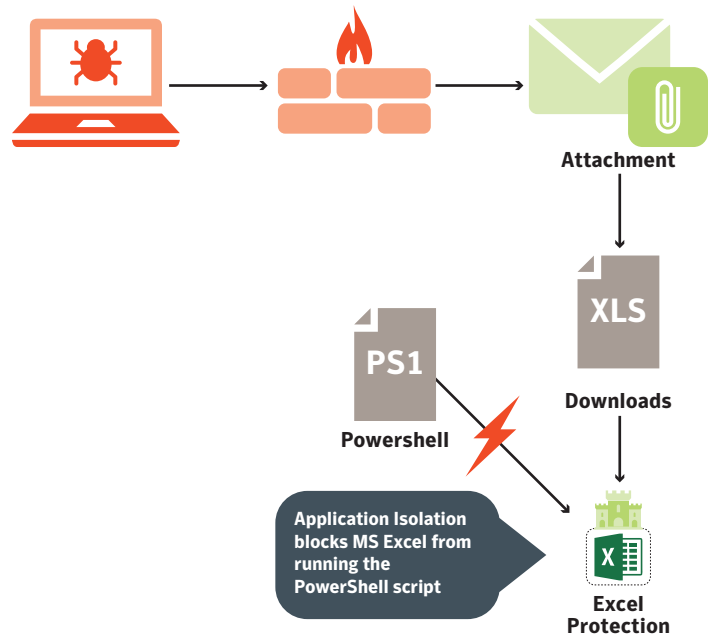
Problem: Sam receives a spear phishing email with a project quote from an attacker pretending to be a vendor with whom he regularly works. The email contains an attachment, 'QuoteForReview.xlsx,' which is a weaponized document. When the document is opened in MS Excel, the spreadsheet looks

deceptively similar to the quotes he normally receives. Sam gets tricked into disabling MS Excel's 'protected view mode' with an instruction that macros should be enabled for the quote's numbers to show up accurately. The sophisticated macro then runs a Visual Basic script that constructs a PowerShell script using snippets of code from the XLS file and runs the script using Windows PowerShell (a commonly used system tool on Windows). This initiates the attack sequence – the PowerShell script communicates with the attacker's hosted Command & Control (C2) server to download additional tools that are used to extract credentials, conduct network reconnaissance, and move laterally to other endpoints.

Solution: Application isolation will prevent a 'known good' application from introducing a threat to the environment.

How it Works: With application isolation, Microsoft Excel runs in a castle that protects it from being tampered with and restricts its behavior. When Sam opens the weaponized XLS document, Excel will be able to render the document's content. However, when the XLS file attempts to construct a PowerShell script file or run the script using the PowerShell utility, the castle will block the operation because MS Excel is neither allowed to create files of certain types nor launch any executables, particularly scripting tools like PowerShell.

Benefits: Application isolation renders the 'fileless' attack useless, protecting Sam's device and eliminating the progression of the attack.



Conclusion

Today's advanced attacks are increasingly focused on exploiting existing applications and tools on an endpoint. An effective endpoint security strategy has to deliver protection that covers the entire threat continuum. With the combination of blacklisting, whitelisting, and application isolation, Symantec now offers the visibility and broad range of controls needed to deliver unprecedented protection across the threat continuum.

With Symantec, threats can be identified and stopped before they can get a foothold and do any damage. Suspicious applications brought into the environment by end users, via browsers, email or other collaboration tools, get automatically jailed, so they pose no harm to the endpoint. Vulnerable applications that are core to end users' productivity can now be run in castles, so they can be used without fear of exploit or file-less attacks. The result is a true zero-day defense for endpoints, that allows you to confidently run any application you need to run your business.

About Symantec

Symantec Corporation (NASDAQ: SYMC), the world's leading cyber security company, helps organizations, governments and people secure their most important data wherever it lives. Organizations across the world look to Symantec for strategic, integrated solutions to defend against sophisticated attacks across endpoints, cloud and infrastructure. Likewise, a global community of more than 50 million people and families rely on Symantec's Norton and LifeLock product suites to protect their digital lives at home and across their devices. Symantec operates one of the world's largest civilian cyber intelligence networks, allowing it to see and protect against the most advanced threats. For additional information, please visit www.symantec.com, subscribe to our [blogs](#), or connect with us on [Facebook](#), [Twitter](#), and [LinkedIn](#).



350 Ellis St., Mountain View, CA 94043 USA | +1 (650) 527 8000 | 1 (800) 721 3934 | www.symantec.com