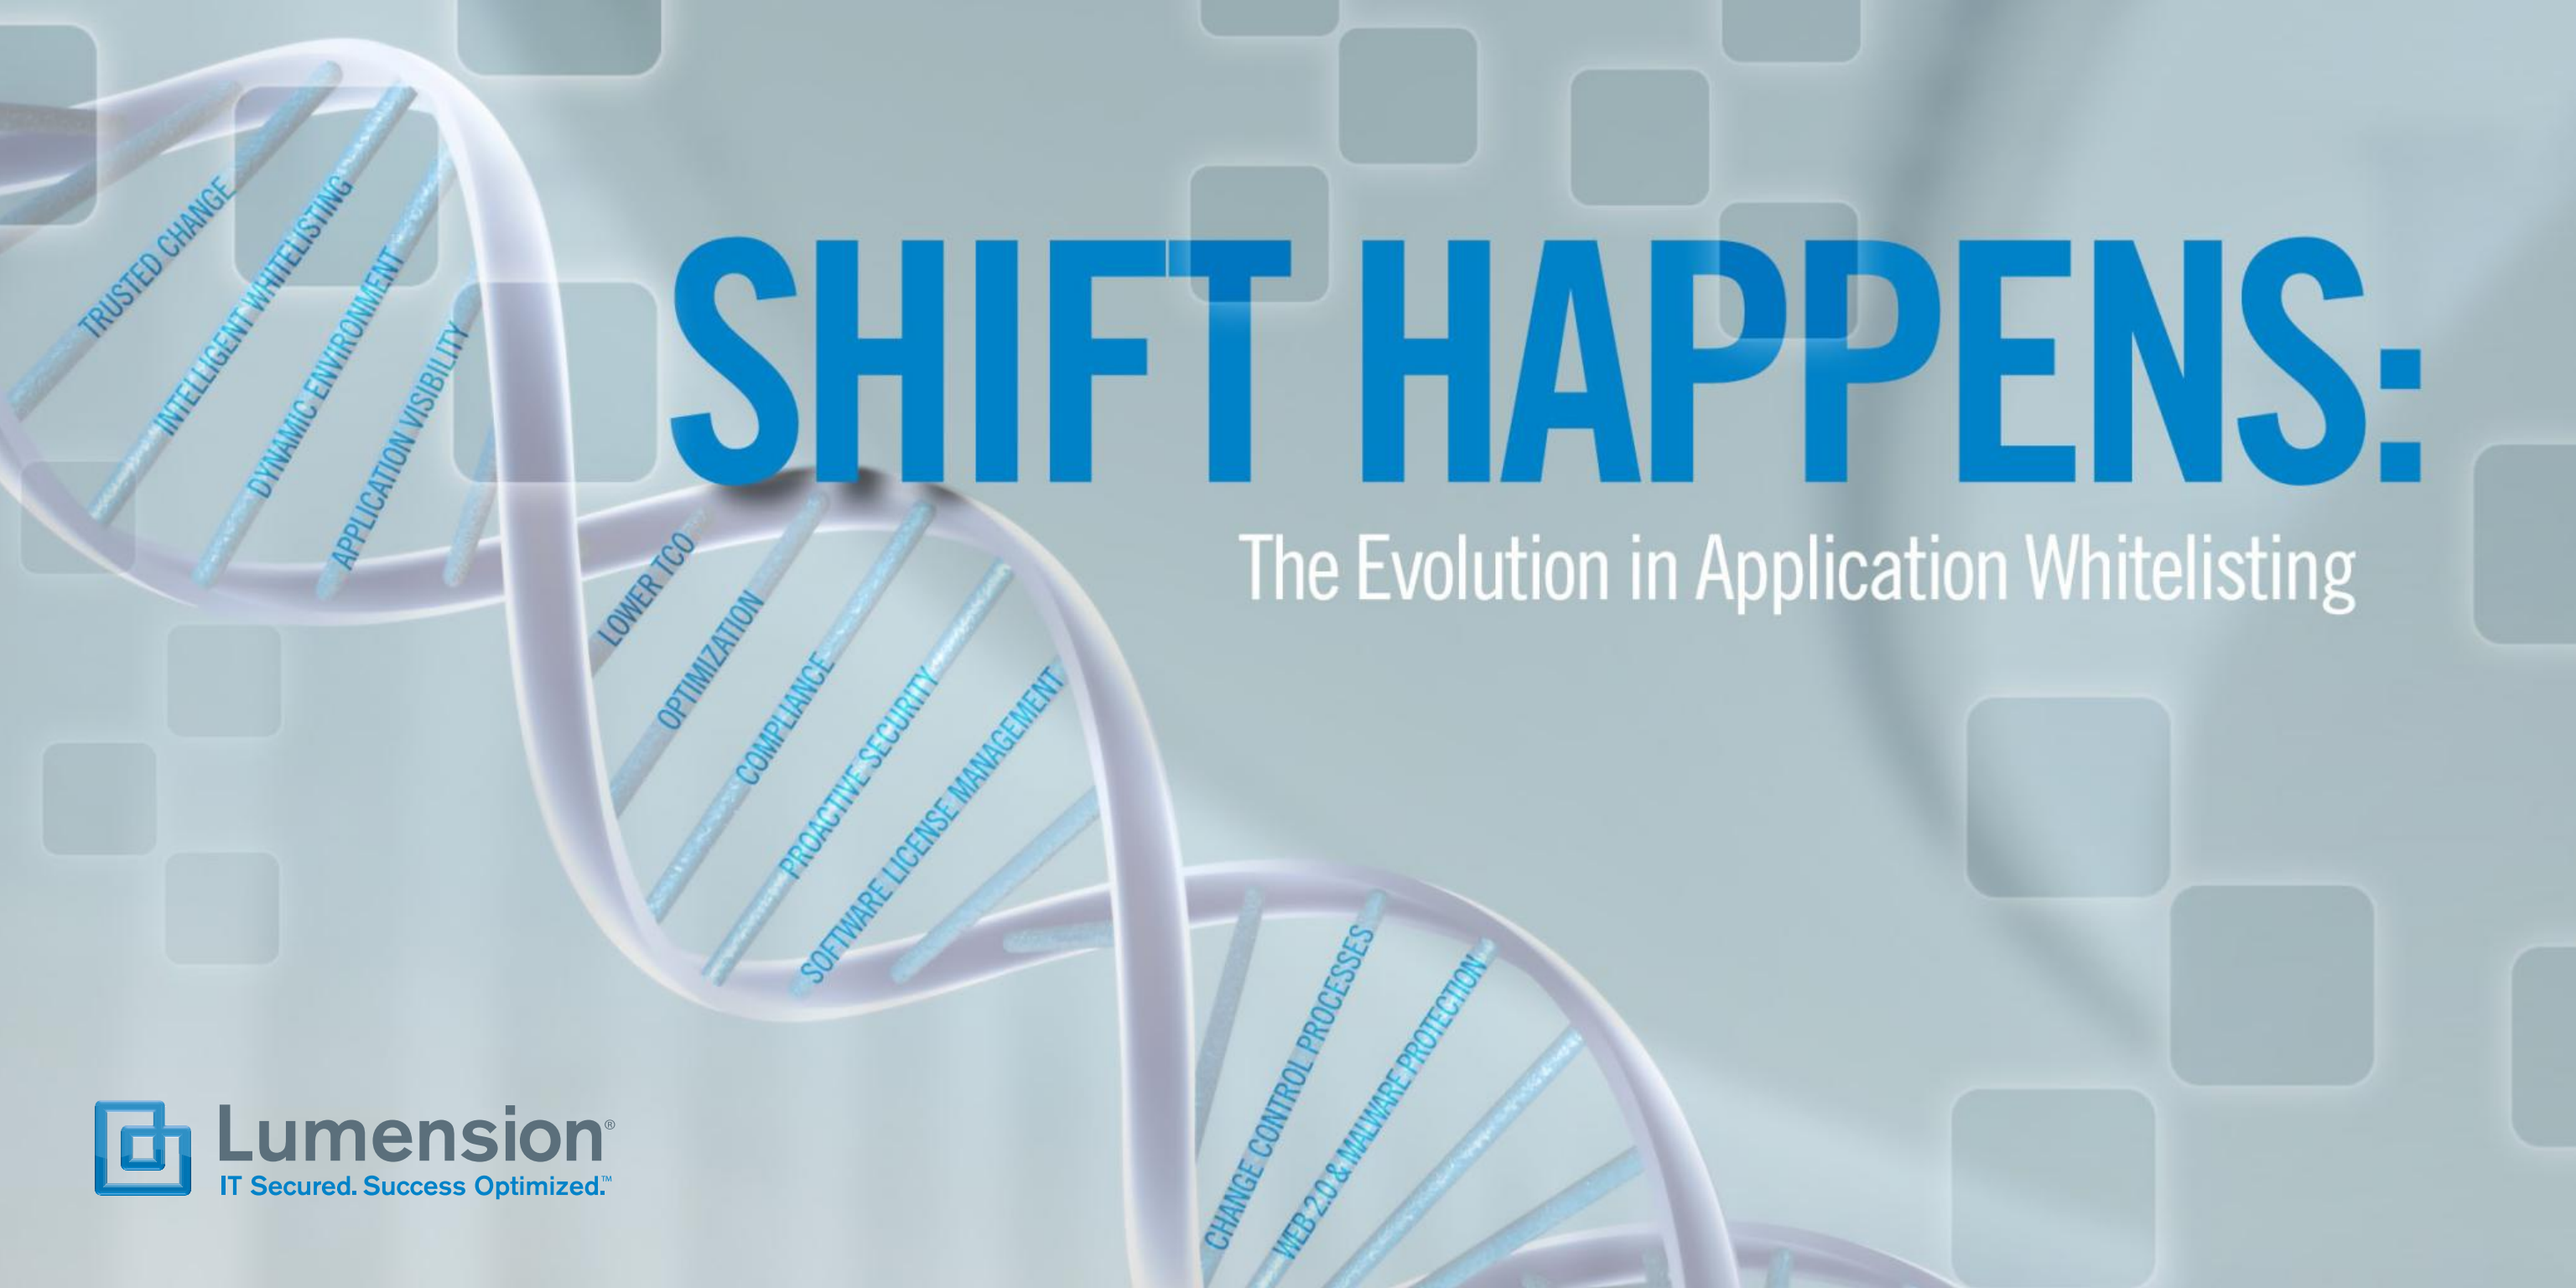

SHIFT HAPPENS:

The Evolution in Application Whitelisting

Lumension®
IT Secured. Success Optimized.™

The Past, Present and Future of Whitelisting

The explosion in malware and vulnerabilities over the last several years has narrowed the usefulness of blacklisting technology.

While still important, it can no longer remain the mainstay of the modern security program. A shift is needed, and whitelisting seems to be the answer. By allowing only trusted applications to execute on a system, organizations can beat hackers at the malware game by taking the ball and glove and quitting the game altogether.

To better understand how whitelisting has evolved over the years, how it can offer immediate benefits within static environments and how future advances will help make whitelisting operationally feasible within dynamic environments, read on.

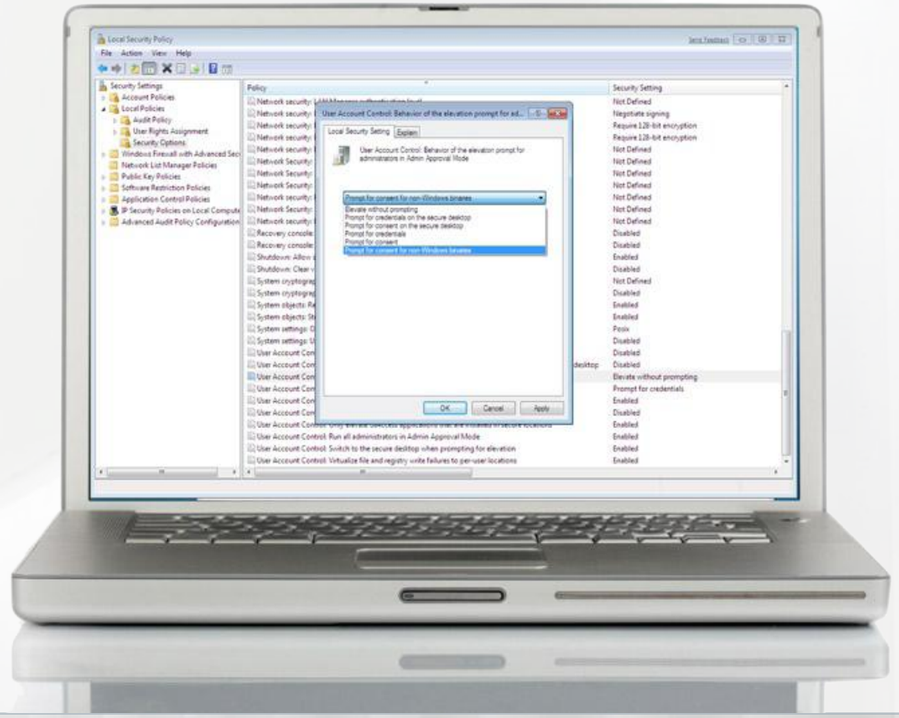

TABLE OF CONTENTS

Lumension
IT Secured. Success Optimized.

Part 1: The Past

Whitelisting may be a novel idea to some, but the old oak trees of security will tell you that it is hardly a new approach to shoring up systems. In fact, whitelisting was the very first computer security model. In the days before PC ubiquity and Web application mania, corporate mainframes and other systems were run in the iron-clad lockdown - running only the known good applications. Anything not on the approved list just didn't run. Period.

As computer use blossomed outside the server room and more users required a growing panoply of new applications, IT managers found that this level of inflexibility could no longer fly within the modern enterprise. There was no dynamic trust model built into that first iteration of whitelisting—either an application was trusted or not trusted with no wiggle room built into the system.

And so began the era of antivirus. Whitelisting was cast off and forgotten in favor of the much more accommodating blacklist method, which allowed everything to run but what was known to be bad for the system. At the time, blacklisting as the major mode of defense made a lot of sense. There were very few vulnerabilities being reported, and the amount of malware online was limited.

Fast-forward to today. The numbers of vulnerabilities and the malware being pushed out to exploit those vulnerabilities have exploded. Hackers are not only taking advantage of more than 7,000 published vulnerabilities per year, but they're layering obfuscation methods that subtly alter single pieces of malware to make them look like hundreds of different applications to the blacklist signature engines.

The end result is that the typical antivirus engines and other blacklist technologies must contend with more than 16 million different malware signatures and still might not be able to detect the thousands more bad programs that emerge daily.

It's time to re-evaluate the situation. Now, let's be real. Antivirus still plays a major part of any company's security program—after all, it is a great first line of defense against all of the bad programs out there that we already know to exist. So there's no reason to throw it out the window.

At the same time, the economics and sustainability of the blacklist security model alone can no longer sustain IT needs of today and tomorrow. Blacklisting technology solves a very narrow field of risks compared to what we face today which is why it just might be time to revisit security's very first model, which has evolved since the first go-around way back when.

THE PAST, PRESENT, AND FUTURE OF WHITELISTING

Vice President of Solution Marketing Paul Zimski provides an in-depth look at the fast changing IT landscape and how intelligent whitelisting will play a key role in mitigating risk while enabling the business.

Part 2: The Present

When we go back to security's roots and examine whitelisting today, we find that what's old is new again. As it stands, whitelisting is not a bolt-on, set-it-and-forget-it type of technology. Whitelisting still does pose some operational challenges while vendors work on ways to hone the technology and processes to address the issue of heterogeneous and constantly changing IT environments. It's a great solution for static environments that have strict change-control processes in place; however, turning on whitelisting in order to lock down systems in a dynamic environment more or less requires a change in the way an organization manages its systems as a whole. Clearly, that's a long-term project.

In the meantime, though, whitelisting as we know it has quite a bit to offer in the here and now. In specific-use cases involving static environments, whitelisting can actually be an extremely simple and cost-effective means toward implementing air-tight security within systems. The type of systems that can offer the most immediate return and benefit are those that you've identified as invariable and will likely not have new applications installed on them on a regular basis. Some textbook examples of whitelist-ready environments and systems include kiosks, point-of-sale (POS) terminals, servers and even SCADA systems.

For example, take the retail environment, which makes liberal use of both kiosks and POS terminals. Both types of systems rarely need new applications rolled out on them, and when they are needed, these deployments are done centrally from IT.

Utilizing whitelisting tools would be fairly painless for these types of systems and greatly improve their security posture. Think about this: One of the largest breaches on record, the TJX debacle was actually brought about through malware insertion on an in-store kiosk. That breach would have been entirely preventable through whitelisting methodology.

As organizations start to find ways to take advantage of the immediate benefits of whitelisting, a very clear business case for the technology begins to appear.

Whitelisting should not only be of interest to CSOs from a security standpoint, but also to the CFO and CEO for its very real impact on the bottom line. The implementation of these tools can offer a number of ways to not only reduce the cost of security expenses, but also to bring down the total cost of ownership (TCO) for IT's hardware and software investments.

The most obvious TCO reduction is, of course, the elimination of costs revolving around the fight against malware. IT organizations spend a great deal of employee resources reacting to malware infections and incidents. It takes manpower to clean up systems and money shipping them back and forth between employees as administrators go through the process of reimaging completely corrupted systems. Because whitelisting technology controls unwanted or untrusted change on endpoints, it prevents malware from ever running on a system, and therefore eliminates the cost of reacting to it, including the potential impact from compromised data.

HOW A TRUST-CENTRIC MODEL IS KEY TO ENDPOINT MANAGEMENT

Matt Mosher, senior vice president of the Americas at Lumension, outlines why blended trust-centric approach to endpoint protection provides better security and compliance posture and flexibility as well as efficiencies across IT operations and security.

Part 2: The Present

Slightly less obvious is the reduction of TCO surrounding other support activities. By introducing a known and consistent set of whitelisted applications to enterprise environments, IT can drastically reduce the help desk overhead required to support all applications.

When the help desk works with a finite number of applications, they no longer have to waste time on obscure questions about unknown applications. Similarly, a finite list of applications cuts down on the number of unwanted applications, such as games and peer-to-peer networking, that can be a drain on support resources when they malfunction, that can drag down system and network performance, and that can cut down on general employee productivity.

Finally, one of the least known TCO benefits—and perhaps most intriguing to the CFO—may well be the ability to leverage whitelisting as a utility to aid in software license management. The in-depth auditing of application use can help IT departments proactively prevent licensing headaches and efficiently make licensing decisions. By getting a handle on who is executing which applications and how often they run those pieces of software, IT can get good visibility into licensing situations before they run amok. For example, IT will be able to spot when more people are using a specific application than a license allows and adjust the licensing agreement before the Business Software Alliance comes knocking on the door. And on the flip side, your organization will also find out for which applications it has overbought software licenses. If an enterprise has purchased a site-wide license for a title and only five people are regularly using that software, the organization stands to save a lot of money through renegotiation with the vendor.

So if all of these business cases for reduced TCO and improved security appeal to you, but you're running a dynamic environment, what then? Well, as we speak, the whitelisting community is making great strides to adjust to meet those needs. The dialogue has already begun about finding ways to define ecosystems and ensure that even in dynamic environments, we can understand change control and enforce policies in order to achieve a more comprehensive visibility into the endpoint beyond when AV can tell us something bad landed on them.

A critical piece to making whitelisting easier to implement has been the development of foundational libraries that allow organizations to identify the majority of the applications running in an environment based on the prevalent applications that run at most enterprises. The idea is not brand new. Some whitelist vendors have been using libraries populated through the scraping of billions of enterprise files. They use AV engines to weed out the known bad files and applications found through the scrape, and slap what's left into the library. But at the end of the day the finished product is a very blunt tool with so-so security capabilities: essentially a very promiscuous library comprised of suspect information.

Part 2: The Present

The most recent refinement of this library idea is vendors are partnering with ISVs who can provide reliable application information via an ISV-sponsored repository. For example, a whitelisting technology leveraging a Microsoft-sponsored repository would then have a standard schema, nomenclature and anatomical understanding of how Microsoft applications manifest themselves within any given system. For example, such a repository could tell the whitelisting utility that Microsoft Word 2007 consists of a certain list of files, hashes, patches, timestamps and default install directories. Then, if the whitelisting technology finds any of that information installed on a system, it would be able to say, "Aha! This is Word 2007!" It takes much of the legwork and the guesswork out of deciding what is and is not acceptable on a large number of dynamic endpoints.

This is a big improvement over the scrape method of library formation because the provenance of the source information offers a much higher degree of integrity and accuracy.

After all, the real goal of this exercise is not necessarily to build the biggest whitelisting library, but to build the most relevant and accurate library. Once an organization has that "Rosetta Stone," it can then do some really interesting things to leverage other apparatuses within the security infrastructure.

CUSTOMER SUCCESS STORY: EC SUITE.COM PROACTIVELY MANAGES ENDPOINT RISK

ec suite
E-COMMERCE SUITE

EC Suite.com, a major processor of credit card transactions for e-commerce companies, saved considerable time and costs as a result of their preventative security measures and procedures using Lumension solutions. EC Suite.com has reduced complexity and TCO, achieved stronger security and compliance posture, lower TCO, and higher level of protection against malicious attacks. Read more on how EC Suite.com has achieved 258.3% rate of return in the first year.

» EC Suite.com ROI Case Study

Part 3: The Future

Alone, these ISV-sponsored repositories can help organizations take a great step toward evolving into a whitelisting security adoption model, and they will provide much of the software identification necessary to reconcile lists of files and hashes into human-consumable policy management. However, this will not be the only mechanism required to achieve an intelligent approach to whitelisting in dynamic environments.

If we look fundamentally at the root cause of the majority of security and configuration issues at the endpoint, most of what we see can be distilled into a single core issue: a breakdown in change control.

Whether that manifests itself through malware or end users installing applications against corporate policy, the lack of an established change-control policy or enforcement thereof is causing quite a bit of security and operational overhead in today's IT environments. (That's a loss felt on the bottom line, folks).

Whitelisting's mission is really about establishing rules and enforcement around how change can be introduced into the environment. Organic whitelisting, or lockdown, does a really good job at stopping change.

So much so, that if it's not used in static, purpose-built systems, it has an ugly habit of impeding the ability to allow enough change to support the business. So how are we going to extend whitelisting's change control enforcement while delivering the flexibility required for dynamic environments and the overall business needs? How do we make whitelisting intelligent? Intelligent whitelisting will be delivered via a rules-based trust engine that can define what types of change are acceptable.

By setting up rules around how change can be introduced, rather than focusing solely on what kinds of change should be stopped, a balanced and effective model of endpoint management can be achieved. It's simpler than it may sound at first.

The trust engine will be easy to configure and based on some straightforward rules. For instance, we will identify which users can introduce change, what applications can introduce change, what software authors (vendors) we may allow, what digital signatures are permitted, etc. So if a new application is being invoked, a series of simple questions will be asked (some of these might look familiar, while others are an extension of how organizations currently reconcile trust):

1. Is this a known bad application, such as malware?
2. Is this an unwanted application, such as games, unsupported software or unlicensed software?
3. Is this a known good application?
4. Do I trust the application that introduced this application?
5. Do I trust the vendor that signed the application?
6. Do I trust the user who is installing this application?
7. Do I trust where this application came from?

APPLICATION WHITELISTING: KEEP THE BAD GUYS OUT – LET THE GOOD GUYS IN

Brent Rickels, senior vice president of First National Bank of Bosque County, provides a look at how his whitelisting approach has been critical to protecting their systems from potential unknown threats. Take a listen to this podcast as Rickels outlines how they are balancing security and end user productivity with whitelisting.

» Listen to the Podcast

Part 3: The Future

Answering these questions greatly extends the user's ability to reconcile trust, verify if change should occur and to prioritize risk incurred by these changes. We must remember that while whitelisting can offer iron-clad security in a lockdown mode, there is an optimization to be achieved that balances the appetite for risk against the needs of the business. This new approach can offer a quantum leap of protection and reduced operational costs without stifling productivity—and is easy to implement. There will also be a spectrum of policy control that can be applied depending on the criticality of the system or the user. Servers and purpose-built systems can be locked down in a pure whitelist and only allowed to change during prescribed maintenance windows, while executive laptops can have far greater flexibility—all provided by the same whitelisting solution. Regardless of how lenient or strict the trust engine policy is that will be implemented, there will be far greater operational stability and security protection from incidents over today's approach of "if I can't prove that it's bad, it can run."

Even further into the future, the truly optimistic realization for whitelisting comes about when technologists are able to tie that rich fabric of information from the application library into the information provided by other security systems such as vulnerability management and patch management systems.

This is the special bonus that can be put atop the library to create a relational awareness about the environment.

For example, say the IT department was getting ready to deploy a patch and the patch management system knows that it will affect a specific DLL that is shared by five applications contained within the library. At that point, the relational awareness will be able to tell the department to test those five applications to make sure they don't break within the environment.

Or, perhaps an IT department is preparing to deploy new software. Altogether, the new relationally aware supersystem would be able to tell you that the software will change a certain set of files or that there are a certain set of vulnerabilities within that software. It would alert the administrator or user that the software is about to introduce vulnerabilities into the environment and then direct that administrator to patches and update levels necessary to remediate those vulnerabilities before the software is even installed on the system.

What this is going to enable users to do is to create an ecosystem where IT management can automate better risk decisions about who can run and what can run on systems even in constantly changing environments. It will allow dynamic trust so that administrators can create a continuum of application permissions based on risk and convenience.

Though it might be in the distant future, the inevitable goal for whitelisting is to ultimately provide a platform for change management. Currently we only talk about whitelisting in the context of applications or devices. But there is no reason why the industry can't eventually expand the concept to non-executable realms.

One of the most eloquent ways to describe the principle behind whitelisting is that it determines what new elements are introduced to the environment.

The theoretical expansion of the whitelisting ideal would be to offer technological controls that enable IT to ensure that nothing is changed without its explicit knowledge or that if it is being changed, it is based on some level of trust, whether it be applications, configurations or even files.

Lumension
IT Secured. Success Optimized.™

A PRACTICAL APPROACH TO SIMPLIFYING WHITELISTING

There has been a lot of talk about what role whitelisting will play in the endpoint protection suites of the future. Opinions dissent about what it will take for whitelisting to become easily implementable and whether it will replace or augment the traditional antivirus approach.

With a fundamental shift taking place in the security environment, organizations need to move away from a threat-centric model and adopt a trust centric approach to achieve higher level of security, visibility, and control.

» Read the Optimal Security Blog
by Vice President of Solution Marketing Paul Zimski

Conclusion: What You Can Do Right Now

Clearly, the latest generation of whitelisting technology offers a spate of security and TCO-reduction opportunities for most organizations.

Though there are currently some operational challenges organizations face in implementing whitelisting as it stands, new innovations, such as ISV-sponsored application libraries and relational awareness, are evolving this security model in such a way that it will soon be able to offer a better way to determine what applications are trusted within a dynamic environment.

As that evolution occurs, there are still plenty of ways organizations can currently take advantage of whitelisting benefits. First and foremost is to identify static systems and environments and start implementing now.

And secondly, organizations can start prepping their dynamic environments by engaging whitelisting utilities in audit-only mode to gather more information about these systems. Doing so will aid your organization in developing a whitelist for future use and can also provide the means to learn more about what applications are running within the environment.

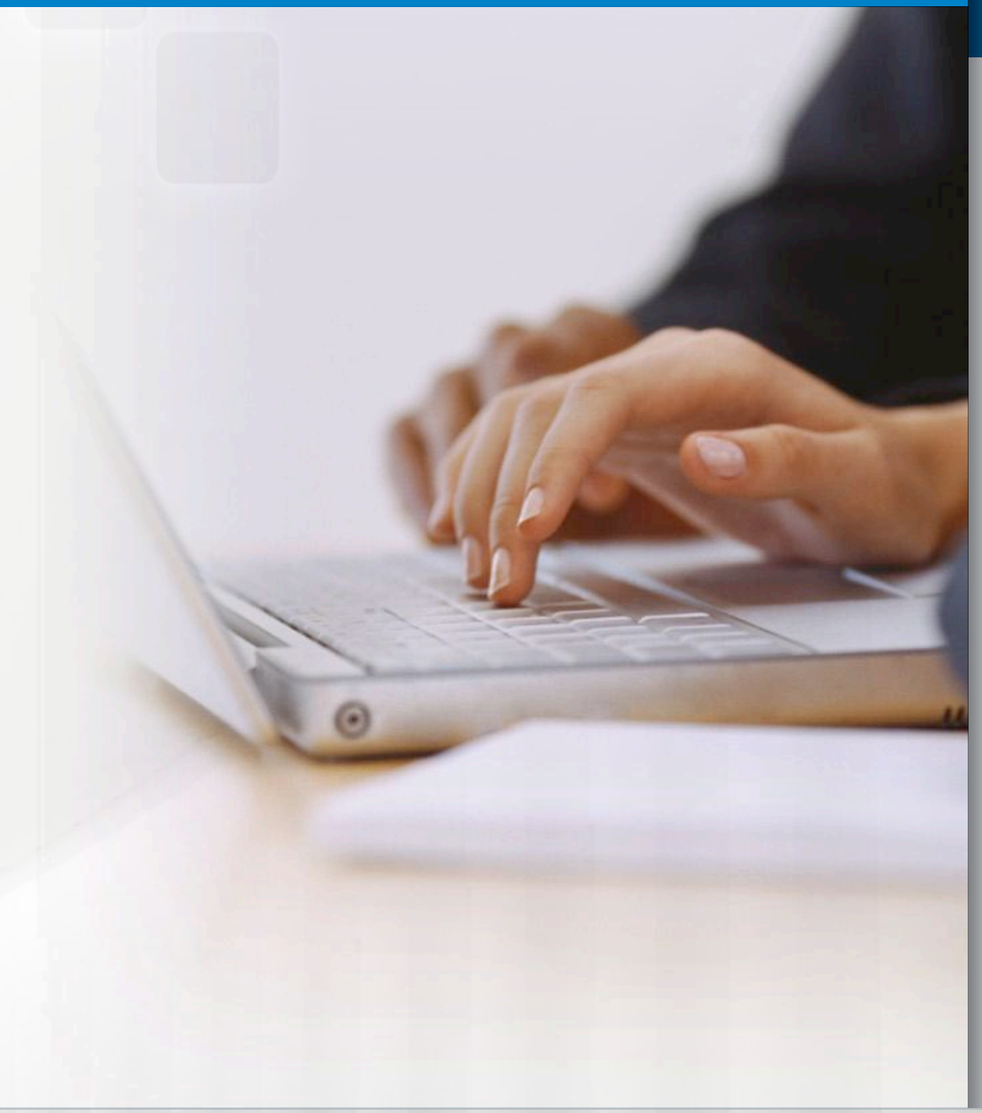

A PRACTICAL APPROACH TO SIMPLIFYING WHITELISTING

Vice President of Solution Marketing Paul Zimski outlines key steps to help organizations overcome obstacles around whitelisting and successfully implement and manage a trusted environment.

Lumension®
IT Secured. Success Optimized.™