Mobile Threat Protection: A Holistic Approach to Securing Mobile Data and Devices

A SANS Whitepaper

Written by Lee Neely

February 2016

Sponsored by

Skycure

Introduction

The ubiquitous use of mobile devices results in a mixture of corporate and personal data stored on devices that are online continuously, seamlessly connecting to the closest available network, downloading and uploading data whenever possible, and carried with users continuously. This trend has radically changed the landscape of data protection.

Complicating that proposition is an abundance of applications for mobile devices designed to empower productivity from a small footprint.

Regrettably, not every application is what it seems, and it is difficult for users to detect a well-crafted forgery or application that secretly exfiltrates data in addition to the displayed functions. Additionally, not every network is what it seems. Users stumble across networks that impersonate legitimate networks but which actually intercept or even change intended mobile device communications. Lastly, operating systems and genuine applications have numerous vulnerabilities that can be exploited.

Tools have emerged to support and secure mobile devices, and corporations have deployed many solutions. This paper reviews the current and emerging services and practices designed to help secure and protect the data on these devices, and it identifies areas where solutions are needed to fill the remaining gaps. We also look at the intersection of security tools such as analytics and their role in managing mobile devices holistically.

Compared to laptops, mobile devices have some behaviors that increase their risk of potential compromise.

First and foremost, they are always on and always connected (or seeking a path) to the Internet. Without user intervention, they connect to networks that match the characteristics of previously known networks, which means that if someone can impersonate a known network, the device will connect to it. Figure 1 shows the variety of risks posed by mobile devices.

Figure 1. Mobile Device Risks

Generally, three mechanisms are used to impersonate Wi-Fi networks. In the first, a misconfigured router can be compromised by attackers. In the second, a malicious participant in a genuine network can attack other participants. Finally, attackers can fake real Wi-Fi networks using tools such as Karma and Pineapple. Because Wi-Fi clients broadcast for known networks 10 times per second, (re)connecting to the strongest network in their list of preferred networks, a nearby strong signal for an insecure but known network is enough for clients to connect to it, leaving their possibly more secure network connection behind. Users also might react to a "free" network by attaching to it more frequently, making them targets.

When a device connects to the fake network, the traffic can be spoofed, changed and recorded, and the user may unknowingly accept fake service certificates, exposing encrypted "secure" sessions to man-in-the-middle (MITM) scenarios.

The data network security that mobile devices use is a big deal due to the mixture of applications that are continuously reaching for updates and passing credentials to servers and services. Many have been implemented to favor speed over security, making the device data streams a rich source of data to intercept or manipulate.

Sadly, the likelihood of a user installing a mobile application or a profile that has nefarious purposes is high because most users can't determine which are from legitimate versus questionable sources. By default, iOS and Android devices are configured to install only applications from Apple and Google Play, respectively. iOS devices generally have to be jailbroken to install applications from non-Apple sources, but Android allows users to install applications from other sources, either continuously or for one-time use by checking a box under Settings. The application source is a concern for those worried about the introduction of malware or repackaged applications because the largest sources for malicious applications are application stores outside Apple and Google Play.

Unpatched vulnerabilities in the device OS and/or applications also contribute to the risk factors. Not only should applications be kept updated, but good sources also have to be used to ensure that genuine updates are applied. In addition, mobile devices need to be physically protected—just having the device may be enough to expose it to infiltration.

What Are We Protecting and Why?

Mobile devices contain applications, documents, stored credentials, photographs, preferences and email, and most of those items are not intended for unrestricted access. Therefore, the devices need to be protected against access to, use of and modification of data. Although data comes in many forms, it can generally be broken into two broad categories: corporate and personal.

Corporate data is information relating to the business. Unauthorized use of this data can harm the business or give a competitor an advantage. Most often, this data is intellectual property, customer information or trade secrets in the form of documents, diagrams and spreadsheets.

More subtly, corporate data can include information or means necessary to access other information systems. Sometimes this data is in the form of credentials, sometimes readily usable, such as stored application credentials or a VPN profile. Other times, access information is in the form of a password database that cannot be used directly but is of high value in the event a toehold is made into the corporate network. These password stores take all forms—from a password management application to plain text notes or even voice memos.

Personal data is similar to corporate data, except that the unauthorized use of this data directly impacts the individual. This includes contacts, identity information, health information, financial data, addresses and information about family members, the aggregation of which can lead to impersonation or even identity theft. Also, access information is present, including password stores, and applications that auto-login as the user, such as email, calendar and social media. Password databases can provide suggestions to passwords used in other contexts. An added complication is that personal information may include information about shared credentials, so the compromise of one data source may lead to impact on other family members or acquaintances.

The point is there is valuable information on mobile devices that is attractive to adversaries. Holders of those devices may not be fully aware of the consequences of loss or modification of that data.

How Is Data Typically Stolen?

Direct exfiltration of data, possibly bypassing data loss prevention (DLP) systems and violation of administrative controls (restrictions on user behavior) are typically the realm of the insider. Tricking users to install malicious profiles, malware or repackaged apps that then could be used to transmit or relay data, establishing a path where network traffic is sniffed or, better yet, allows an attacker to act as a MITM, are actions an outsider would leverage to access data.

Regardless of the actor, physical possession is still the most viable path to data collection. Two simple practices reduce the viability of this scenario: adding a good password and enabling on-device encryption. A 2014 study by Consumer Reports[1] found only 47 percent of users had a password on their devices, and among those with passwords, 77 percent used a four-digit PIN. This means 53 percent of devices can be accessed just by physical possession. And if they have a password, it is likely a four-digit PIN.

Device encryption has been provided by default for iOS devices since iOS 4, but it must be user-enabled for most Android devices. Worldwide smartphone studies from International Data Corporation (IDC)[2] in Q2 of 2015 show iOS and Android make up 97 percent of devices, 85 percent of which are Android, indicating a large target of opportunity in this area.

Even if the device is encrypted and has a strong password, what avenues remain for accessing the valuable data contained in these devices? Vulnerabilities in the device OS, malicious Wi-Fi networks or installed applications may still leave opportunities for compromise. For example, a recent Pwn2Own exercise[3] was able to install arbitrary applications on a fully patched Android using weaknesses in Chrome, and lock screen bypass exploits continue to be found and patched on iOS. Having reputable applications, a secure network and a fully patched OS are the frontline steps to reducing this risk. In the next section, we also look at network attacks and malware as attack vectors.

> Two simple practices reduce the viability of unwanted data collection: adding a good password and enabling on-device encryption.

---

[1] "Smart phone thefts rose to 3.1 million in 2013," Consumer Reports, May 28, 2014,
www.consumerreports.org/cro/news/2014/04/smart-phone-thefts-rose-to-3-1-million-last-year/index.htm

[2] Smartphone OS Market Share, 2015 Q2, IDC, www.idc.com/prodserv/smartphone-os-market-share.jsp

[3] "Latest Android phones hijacked with tidy one-stop-Chrome-pop," The Register, Nov. 15, 2015,
www.theregister.co.uk/2015/11/12/mobile_pwn2own

Common Attack Vectors and Possible Mitigation

Mobile devices are susceptible to various types of attacks. These attacks can be physical, over the network, a vulnerability exploit or via malware. Let's look at some common scenarios and possible mitigation methods for those attacks. Figure 2 represents the various attack vectors for mobile devices.

Figure 2. Mobile Device Attack Vectors

Physical Security

When considering attack vectors, location is a factor. Device management solutions allow modified security policies to be applied or removed based on location. When location-based policy changes were first introduced, the thinking was that they could be used to enable/disable features to meet business restrictions. For a more exciting perspective, imagine dynamically disabling or changing settings based on real-time threat data. The idea is that when your device is taken into a hostile environment, policies are applied to strengthen the security settings, and when you leave the hostile environment, the settings revert to normal.

Physical access allows devices such as an IP-Box or USB Rubber Ducky to be connected to the mobile device in order to attempt to bypass its passcode. These devices use brute force methods to successfully guess a four-digit PIN in fewer than 17 hours by taking advantage of the OS response to an incorrect passcode. Android forces a 30-second pause after 10 failed passcode attempts, and the iOS bad password counter, prior to iOS 9, could be reset by power-cycling the device after five attempts.

Additionally, because there is a relatively small number of four-digit PINs, trying every possible combination is practical, making this a realistic attack vector. Having users select a six-digit PIN raises the bar to one in 1,000,000 rather than one in 10,000 possible choices, making this attack much harder and less viable. If the device supports biometric authentication, consider using that rather than a passcode that can be observed. Although biometrics also can be used to offset the inconvenience of requiring a complex passcode, personally owned devices may have fingerprints from multiple users, which is a risk to corporate applications that use device biometric authentication.

In addition to strengthening the passcode, configure the device to wipe after a set number of passcode failures (typically 10).

Network Security

Bad actors are looking for ways to get between the device's communication and back-end services. One mitigation approach is to leave Wi-Fi disabled so network communication is passed over the cellular network, which is harder to impersonate. Another way to mitigate the risk is to run all communications over a full tunnel VPN, which would further protect the communication from interception or modification. There are many disadvantages to the full tunneling approach, such as battery life, privacy concerns and server maintenance, if it is always on. A much better approach is to turn on full tunneling only when a threat is detected.

When choosing to route all the device traffic over a VPN, consider the kind of traffic to be included and the effect of that traffic on the network, including bandwidth and content/destination. For example, corporate users streaming content may affect off-site business data flows.

Consider the risks of an application-specific VPN. Although the impact to the device is lessened, a compromised device may allow other applications to access that VPN and consequently the corporate network. Root or jailbreak detection software may be able to mitigate this risk by preventing execution, raising an alert or even wiping the device.

Malware Security

Malware comes in many forms that need to be detected and stopped. Detection is a challenge because signature-based detection software can be easily bypassed simply by providing an updated/repackaged mobile application that will have a new signature.

Fortunately, new malware detection techniques are emerging. Instead of just relying on the malware signature, behavior-based algorithms are used to detect inappropriate actions by applications, possibly using this information to build a reputation database for otherwise uncategorized applications.

The mobile OS itself may also provide barriers to detecting malware. On iOS, the sandboxing of applications makes detection of malicious activities from another application impossible. The first defense is to limit where applications can be installed from; both Google Play and Apple's App Store have screening mechanisms designed to prevent the introduction of malware into these distribution channels. Beyond that, whitelisting and blacklisting solutions are needed to block the execution or possible installation of malware. To be successful, these need to be fed with real-time threat data, which is challenging to do in-house.

New Paradigm: Distributed Data Gathering/Aggregating

Mobile devices are small computers, and as such, they can collect and process data about threats in their vicinity. Mobile devices have a lot of compute power and storage, and they doggedly work to preserve a network connection, which makes using them to process threat data viable so long as this doesn't impede end user processing or experience. Once the mobile devices are acting like sensors and sending the results to a repository, it becomes possible to aggregate that data and provide it to device management systems for real-time threat analysis.

Keeping OS and Applications Updated

Vulnerabilities in mobile device operating systems and applications drive the need to keep the OS and applications updated. When a device needs an update, ask the following questions:

• Has the update been regression tested?

• What will the requirement be for applying that update?

• Who is responsible for updating the items and how the update will be applied?

• What are the consequences of not applying updates?

• What is your communication plan to affected parties?

Threat Intel and Analytics

Threat intelligence can mean many things. In this case, we are talking about data that has been collected and been through some analysis process that adds relevance. There are many great sources of threat intelligence from the vendor community, which provide evaluated data that yields insight into threat actors, vulnerabilities, exploits, indicators or compromise from many data sources.

Threat analytics considers these data feeds and turns them into actionable information by discovering threats and patterns. In addition to the data feeds above, organizations incorporate their local security monitors so that externally reported actions or trends can be made relevant to the current operational state. Have you considered that mobile devices could be data sources in this equation?

Threat and response information from mobile devices should be fed into existing security tools with two desired outcomes. First, to provide in-depth information for improved situational awareness. Too often mobile devices are blind spots to traditional security operation centers. Second, to learn which threats affect other IT assets because without information about these threats, appropriate actions cannot be taken. If many mobile devices feed threat data to a common source, that data can be aggregated into a valuable threat intelligence source.

Using the device to collect data has to have nominal impact on the users. Device monitoring and data collection services on traditional computers often have a negative impact, either by resource (disk, memory, CPU, network) consumption or by anecdotal (water cooler story) information. The net effect is an erosion of the trust relationship between IT and the user. Mobile device data collection has the potential to be even more disruptive. Because of these changes to how users work, selecting a technology and process that users are already familiar with, such as an app download, might be effective.

Careful evaluation of the tools selected to collect data from devices and clear communication of both the return and impact to users in a context of collaboration are key. Users will want to know if they can disable, remove or otherwise mitigate impacts. Working with representative user groups ahead of time to develop responses, mutually acceptable settings and documentation is worth its weight in gold.

Aggregating threat information from many devices across multiple customers can be risky because it might reveal the mobile device details of one organization to another, providing a possible competitive advantage to one or the other of those organizations. Because the goal of this information sharing is to protect all of the devices from discovered or predictive threats, solutions must be tailored to only provide the core information needed. A specific risk to a given application doesn't need to include detailed information about the device on which it was detected nor does it need to be attributed to the organization that device is associated with. In some cases, tokenized data can provide enough information while obfuscating sensitive specific information. To be successful, both service providers and consumers have to have a common understanding and agreement of what is being collected and shared.

How Mobile Devices (and Data) Are Often Protected and the Effectiveness of Those Protections

Traditionally, mobile devices have been managed by a mobile device management (MDM) solution. These systems put configuration policies onto mobile devices using either built-in or added device management APIs.

Prior to 2007, the "gold standard" of secure device configuration was the Research In Motion (RIM) BlackBerry Enterprise Server (BES). This system was used to manage almost every aspect of the device, in short because the device and server were both RIM/BlackBerry products.

In 2007, we saw the introduction of the iPhone and in 2008, the introduction of the Android. These two devices, along with their use and application paradigm, transformed the smartphone market and presented the consumer with a new model for how smartphones could be used. User adoption became widespread and the variety, availability and use of mobile applications exploded.

These changes—iPhone, Android, adoption and mobile application explosion—made it necessary for corporations to incorporate these devices into their formerly homogeneous mobile device environment. However, the management capabilities of the different devices vary greatly, both in manageable features and how that management is achieved, making it harder to ensure equal protections across different platforms (iOS, Android, Windows Phone and possibly BlackBerry).

Mobile devices are now pervasive and no longer the tool of the elite. Users install applications whenever they need another tool to help them achieve desired results, and they expect to use them wherever they happen to be, connecting to cellular and Wi-Fi networks with less regard for security and more concern for getting online and connected. Figure 3 shows the various aspects of enterprise mobility management (EMM).
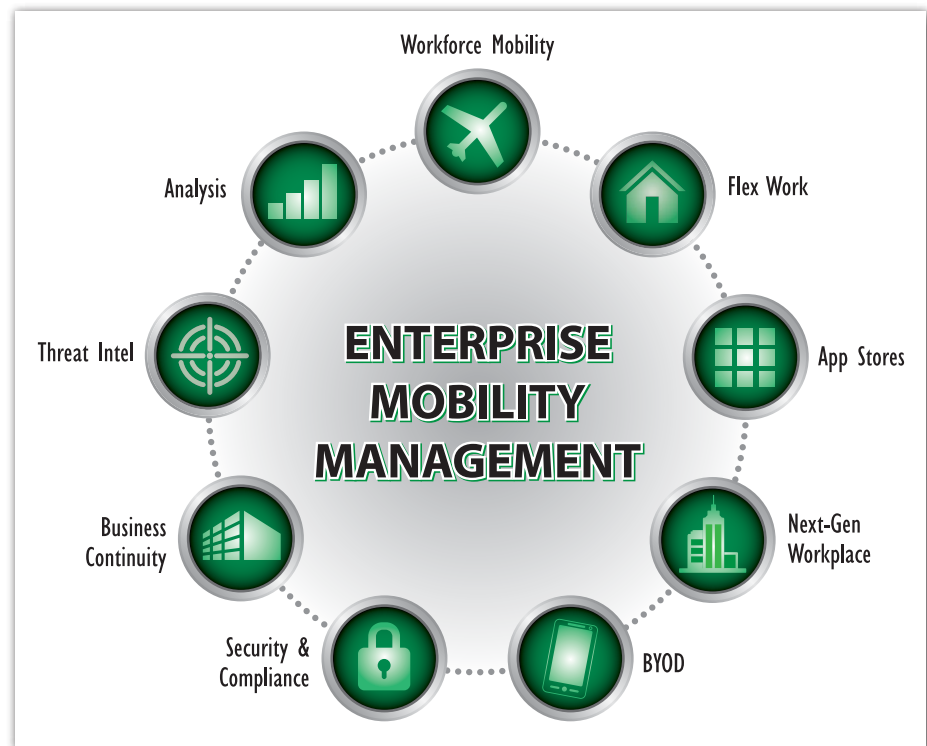
Some new approaches have evolved in device management to attempt to manage enterprise mobility. Some organizations whitelist or blacklist applications, largely manually, in an attempt to reduce the introduction of mobile malware. Additional efforts have been focused on limiting the sources of applications that can be installed, reducing the possible sources from which a malicious application can be installed. To address network layer risks, device-level VPNs and application-level VPNs have been used to try to route traffic securely to back-end systems. The problem is that to be effective, the applications need a connection that establishes itself transparently. That means the credentials for these connections must be stored on the device, which increases the possibility they can be compromised.

So, how are devices managed and what are the limitations?

Traditional Device Management Solutions

Traditionally, devices are managed by an EMM solution, which is good at installing policies/settings and detecting compliance issues. EMM solutions may not manage application whitelists/blacklists, nor may they introduce protections intended to secure the network path to the mobile device.

Next, in an attempt to segregate corporate data and keep it from exfiltration, containers have been deployed to create software and policy boundaries around data. The container secures the data, typically using encryption to separate container information from the other data stored on the device, and provides partner applications within the container to share data and possibly credentials. Some container approaches also add a secured/trusted communication link to ensure the legitimacy of communication between the applications processing corporate data and the back-end information systems. This approach can have significant impact on the user experience, though, because users often need to process their data with applications outside the container. There also have been a few cases where containerization has failed because the underlying device was compromised.

Another data isolation approach is virtualization, where the application and data are not actually present on the mobile device or are executing as a guest OS on the device. Both options have limitations similar to containers, with the possible added complication of needing a network connection to the remote environment, as well as the added risk of compromise of the credentials used to access the virtual app/environment. This approach, particularly in the form of a Virtual Desktop Interface (VDI), has been successfully used in bring-your-own-device (BYOD) or remote office computing environments, and there is some attraction to using it with mobile devices. In addition to the limits mentioned previously, there are also challenges with how the user interface is mapped to the mobile device, reducing the effectiveness of the solution.

The last frontier of device management is user behavior. Having users agree to behavior restrictions that cannot be technically enforced both mitigates and carries risks. If you are counting on user acceptance and execution, you need to provide a user-friendly solution.

Modern Solutions

As mobile device technology advances, so do threat management solutions. Application analysis, location-based defense and distributed threat intelligence gathering represent major steps toward mitigation.

New and exciting services are emerging to analyze applications from a wide variety of sources. Some even use the mobile devices themselves as data collection points. The idea is that a distributed set of analysis tools will work together to examine whether an app is risky based on the origin, static analysis, dynamic analysis, behavior patterns and other parameters. The downside here is the IT department must find a service that integrates with its company's present and future EMM systems. That EMM solution must also have sufficient access to the device application store so that accurate inventory can be collected and disallowed applications can be stopped prior to installation or removed/quarantined when detected post-installation. Also, a process has to be in place to address categorization issues that may not match user/business needs. This gets particularly tricky with BYOD.

Another evolving space is location-based policies. As the name implies, they enable or disable services and actions based on the location of the device. For example, such a policy could disable Wi-Fi in a facility that doesn't permit it. If a location-based update is missed, however, policies aren't updated. The creation of new policies is not terribly dynamic, so creating and deploying a response based on rapidly changing threat conditions—such as detecting a bogus Wi-Fi hotspot and configuring the device not to use it—is problematic.

A shift is occurring from reactive threat detection and response to proactive and even predictive threat detection. This shift is analogous to discovering an attack that happened after reviewing log files versus real-time detection, alerting and blocking of an attack as it is attempted. Most current EMM solutions are not designed to be reactive or to implement "real-time" responses based on dynamically changing threat scenarios.

Another focus is on situational awareness, which provides continuous visibility into devices to learn the threats they are facing and hopefully defend against them. This is

not just a case of having a secure configuration, but also of having information to show it is working. For example, if a Wi-Fi hotspot is deemed dangerous, having information that shows defenses were deployed as well as how many devices saw that hotspot, how many connected, etc. would provide better information about the environment mobile devices are being used in and could drive risk management decisions.

An opportunity exists to create the ability to push app/OS updates to mobile devices. IT organizations have spent years maturing centralized OS, patch and application management solutions for desktops/laptops and servers. Mobile devices are dependent on end users for these updates, which can be difficult because users are accustomed to this being handled for them. These updates need to be applied without user interaction, or by visiting a help desk, and aligned with risk and business impact.

Checklist

Selecting the right mobile security solution makes holistic management and monitoring of a mobile device fleet much easier. Keep the following guidelines in mind when it's time to choose one or to purchase additional products to integrate with your EMM to fill any gaps that expose you to risk.

	Requirement	Priority	Additional Info
Deployment process	Support app download from public stores	High	Official app should be available on Apple's App Store and Google Play
	Overall ease of deployment	High	Considering required actions by the end user and the admin
End user experience	Low impact on device battery usage	High	Usage should be under 3%
	Low data usage	Medium	Both on cellular network and Wi-Fi
	App maintains end user's privacy	High	Not exposing sensitive user information
	Clear display of detected threats and mitigation options	High	Provide a clear and simple display of detected threats with an advisory for mitigating them
	Provide automatic mitigation options for most threats	High	For minimizing actions required from the end user
Threat detection	Network threats		
	• Secure communication downgrading (SSL stripping) attack detection	High	Man-in-the-middle attack in which the device communication is downgraded from SSL to plain text
	• Secure traffic decryption (SSL decryption) attack detection	High	Man-in-the-middle attack in which traffic from the end user's device is decrypted by the attacker
	• Content manipulation attack detection	Medium	Attack in which the content of a web page is altered in order to manipulate the end user
	• Rogue networks detection	High	Identify anomalies in public hotspots to identify rogue networks
	• Ability to perform automatic mitigation on detected network threats	High	Mitigate network threats without end user intervention, keeping traffic secure without losing connectivity
	Malware		
	• Detection of malicious apps based on different app properties	High	For instance, app source, requested permissions, certificate, etc.
	• Detection of repackaged/fake apps	High	Detection of malicious apps that impersonate legitimate apps
	• Detection of malicious apps based on signatures/known exploits	Medium	Using standard antivirus capabilities
	• Ability to block malicious app installation	High	Intervene in real time to stop installation in case the app is risky
	• Detection of iOS malware	Medium	Ability to detect new and existing iOS malware such as XcodeGhost and YiSpecter
	• Detection of malicious profiles on iOS devices	High	Malicious profiles can be used for monitoring/controlling activity on an iOS device
	Device vulnerabilities		
	• Ability to identify jailbroken or rooted devices	Medium	Detection and policy enforcement on these non-compliant devices
	• Ability to identify device OS vulnerabilities	High	Present vulnerability details and risk clearly for each device
	• Ability to prompt end users to upgrade their device OS version	Medium	Ability to do this as soon as the update is available (sometimes even before the formal vendor announcement arrives)

	Requirement	Priority	Additional Info
Management and administration	Provide visibility on detected threats and vulnerabilities	High	Present a clear, detailed description of each threat (including network and malware) and vulnerability (OS/device configuration)
	Provide an overall risk estimate per device	High	Risk calculation should take into account current threat, device history, vulnerabilities, etc.
	Provide forensic capabilities on identified threats	Medium	Present details about the impact of each detected threat
	Provide the option to define an organization-level compliance policy	High	Devices that do not comply with the organizational policy can be blocked from using organizational resources
	Reporting	High	Provide reporting capabilities, including scheduled email reports, support for different data formats (tables, graphs) and document formats (PDF, CSV)
Other	EMM integration	High	Work with or without an existing EMM solution such as AirWatch, MobileIron and XenMobile
	SIEM integration	High	Support integration with different SIEM systems (ArcSight, McAfee ESM, Splunk, etc.) for exporting detected threats
	Provide a third-party API	Low	Provide a third-party API for retrieving device security information

Summary

Mobile devices are more than just small computers in continuous use with perpetual connections to the Internet. They are key business and productivity tools. As such, they need to be identified, secured and managed as you would any business IT asset. The operating paradigm of these devices calls for new approaches to ensure the data processed by them remains secure while maintaining productivity. The ecosystem to manage these devices must include both technical and operational controls, and it must integrate into the overall operational awareness for the business.

Before selecting a management suite, you need deep visibility into not only the configuration of the devices, but also the environments in which they operate to continuously and appropriately update their security posture. Consider the mobile device fleet as an extension of your existing security sensor network. After leveraging the provided checklist to achieve an optimal solution, fill any gaps with user training and guidance. Users can either be your greatest security risk or they can be your greatest security asset; you hire them to solve problems rather than contribute to or create them.

About the Author

Lee Neely, a SANS mentor instructor, teaches cybersecurity courses, including the new cybersecurity management training, and Information System Security Officer training. He worked with the SANS SCORE project to develop the iOS Step-by-Step configuration guide as well as the Mobile Device Configuration Checklist included in the SEC 575 course. A senior IT and security professional at Lawrence Livermore National Laboratory (LLNL), Lee has been involved in many aspects of IT. He currently leads LLNL's new technology group, working to develop secure implementations of new technology, including developing its secure configurations, risk assessments and policy updates required for corporate and BYOD mobile devices.

Sponsor

SANS would like to thank this paper's sponsor: