SECURITY BEST PRACTICES FOR THE MOBILE ENTERPRISE

Securing the Enterprise Network as the Number and Types of Mobile Devices Proliferate

Table of Contents

Table of Figures

Executive Summary

The proliferation of mobile devices can undermine enterprise network security in the absence of effective security policies and provisions. The biggest challenge related to mobile device proliferation is how to accommodate all of the different users, many of whom have multiple devices running a variety of different applications as they connect to the enterprise network in different ways at different times and from different locations. Some of these devices will be secured with up-to-date antivirus software and other endpoint security provisions, but because many of these devices are personally owned, a growing number will not be secure. It is necessary, therefore, to treat different devices differently, even for the same user, while giving all users a consistent and seamless connectivity experience, regardless of the device and access method used.

Security Best Practices

The proliferation of mobile devices, including wildly popular tablets, is having a profound impact on network security. This white paper highlights how changes in user mobility are affecting security requirements in the *Unwired Enterprise*, and outlines a set of best practices for controlling access within and beyond perimeter protections.

For these reasons, enterprise network security today needs the flexibility to deliver different levels of service to different types of devices depending on how secure they are, while maintaining the granular manageability needed to ensure that policies are applied continuously and consistently based on user identity and profile. These provisions will need to leverage existing security infrastructure such as directory services for authentication and access control, while also incorporating industry standards for wireless encryption.

This proliferation of wireless devices will cause significant changes throughout the enterprise. The good news is there should be a dramatic increase in productivity from and for mobile users, including those in IT departments. The bad news for IT is that the proliferation of wireless devices involves many new challenges to the status quo in the enterprise network. How will the WLAN accommodate so many new and different devices? What impact will pervasive WLAN access (the Unwired Enterprise) have on network reliability? Will the integration of voice and data communications on tablets and smartphones be the catalyst for increased adoption of unified communications and fixed-mobile convergence in the enterprise? Or does the integration of voice over IP (VoIP) into 3G/4G smartphones somewhat obviate the need for the latter? What changes will be needed to existing enterprise security policies and procedures to enable the use of personal devices? Should the organization even permit the use of personal tablets and smartphones at or for work? (Good luck trying to stop that!)

This white paper highlights how changes in user mobility are affecting security in the enterprise and for the IT staff that manages network access. It outlines a set of best practices for controlling access within and beyond perimeter protections. Most affected are network access control and related provisions, where market changes and trends warrant new capabilities and procedures. To ensure comprehensive treatment of this critical topic, those standard security provisions that are assumed to be in place, including the basic authentication, authorization, and accounting (AAA) framework, and which are relatively unaffected by the proliferation of mobile devices, are covered briefly in an appendix at the end of this paper.

Introduction

It is important to note that virtually all users are now mobile to one extent or another. Even those bound to a desk are likely to have a personal smartphone or tablet. And they will want to use those devices to do work while at work, as well as to access the Internet or other resources for personal needs, also while at work. And many will want to use their personal devices for work-related activities while traveling or at home. This means that the previous network-centric security framework will need to become more user-, device- and mobility-aware. In short, it must become mobility-centric. This is where wireless local area networks (WLANs) have excelled, because they were built from the start with mobility in mind.

The proliferation of mobile devices is both accelerating and shaping the migration to what many are calling the *Unwired Enterprise*, which has long been a goal in most organizations. To be sure, the Unwired Enterprise is a bit of a misnomer, because the backbone and core of the network will remain wired. Rather, the focus here is on the fact that wired access is fast becoming a thing of the past. Perhaps the term *Wired-less Enterprise* would be more accurate. The proliferation of mobile devices, combined with the advent of 802.11n, now makes achieving the goal of fully wireless access as urgent as it is practical. One needs look no further than one's own mobile phone use to realize the advantages of anywhere, anytime voice communications. And just as it makes sense to untether voice, so too does it make sense to untether data—provided, of course, that the entire network, wireless and wired, remains solidly secure.

The best approach to securing the multitude of different devices roaming about the wireless LAN infrastructure is central control with distributed enforcement. Wireless access security must, therefore, coexist seamlessly with the well-established centralized policy management infrastructure enterprises use today. Wireless access security must expand upon this foundation to address new security considerations brought about by user mobilty. And, distributed enforcement of network policy must be both efficient and effective to ensure enterprise-wide scalability.

Another important aspect of the required flexibility is granularity. With the growing number of permutations and combinations of different users, devices, and situations, greater granularity becomes critical when determining both pre- and post-admission access control policies. Simply put, security that relies solely on detailed information about users stored in traditional AAA or directory servers, while necessary, is no longer sufficient. Access control must now also take into account the user's device, its endpoint security, and potentially the user's changing location and other criteria, in addition to the user's identity and role.

Mobility Madness

Perhaps madness is too strong a word, but the excitement surrounding tablets makes the debut of these devices more than just another adjustment to business as usual. This is indeed a disruptive trend that is accelerating the demand for ubiquitous WLAN access that is campus-wide. Perhaps the reason is that the tablet appears to be a fulfillment of the capabilities long portrayed in science fiction movies like Star Trek and Star Wars: A powerful, handheld device with a touchscreen, integrated global positioning, even voice recognition—and hundreds of thousands of applications. It seems to be what everyone has always wanted. And now it is finally here.

Sales of laptops overtook desktops in 2007, according to Frost & Sullivan. Frost & Sullivan, along with other analyst firms like iSuppli and IDC, all now concur that laptops (not including tablets) will dominate the PC market going forward. IDC expects that portable PCs will account for 70% of the market by 2012—a number that again excludes tablets, as IDC does not consider a tablet to be a PC. Gartner also predicts that mobile PCs will have a 70% share of the market by 2012.

Enter the tablet. *PC World* magazine declared 2010 as "The Year of the Tablet Computer"—in the year Apple's iPad debuted. Some 20 million units later, the iPad has been joined by a growing number of tablets, and a platform battle is currently raging among industry titans. Microsoft, Hewlett-Packard, Google, Motorola, Dell, Samsung, Research In Motion, and others are all trying to unseat Apple's early lead by grabbing more marketshare. But this is only the tip of the iceberg that lies ahead. In 2010, Forrester Research predicted that the number of tablets in the U.S. alone would grow from 10 million that year, to 82 million in 2015. Of course, that was before the growing multitude of other vendors had even announced their tablet offerings. Forrester also estimates that at least half of the 15 million Apple iPads sold to date are already being used daily in a business setting.

Most tablets and nearly half of all smartphones are now dual-mode, with both cellular (GSM/3G/4G) and Wi-Fi communications built-in, according to ABI Research. ABI predicts that by 2014, a full 90% of smartphones will incorporate Wi-Fi, and that sales will nearly quadruple from 144 million Wi-Fi equipped smartphones in 2009 to 520 million by 2014. This means that smartphones that might have previously been isolated from the enterprise network will now be able to access it (with permission granted, of course).

All for Naught Without Holistic Security

While there are many such issues, the focus here is on security and the impact is equally profound. If not properly handled, the use of personal devices presents a serious security threat. Most of these devices will lack robust security provisions such as antivirus and personal firewall protection. Many will contain hundreds of apps, including downloaded music and video files, games, and other software. Many will access social networking sites. Some users will even have multiple devices and will want to synch their personal data (including their calendar and contacts) among, for example, the corporate issued laptop and the personal tablet and smartphone.

There are direct threats to the wireless LAN, as well. Consider the dual-mode device. Because it has both Wi-Fi and cellular communications, some also offer the ability to create a local Wi-Fi hot spot that is backhauled via the cellular network. Imagine, a feature that can turn a smartphone into a rogue access point. (Thanks vendors, that's exactly what corporate IT needed!) Will new policies prohibit such features from being used? (They probably should.) And if so, how will these rogue access points be detected and isolated?

There are many more issues involved in securing the Unwired Enterprise, to be sure, and these cannot be addressed without some control over, or at least a better understanding of, the security posture of the multitude of mobile devices themselves.

Mobile Device Security

Mobile devices pose a potential risk to the enterprise network owing to their very mobility. They routinely travel outside enterprise perimeter protections, including the firewall and any intrusion prevention systems (IPS). They are loaded with potential vulnerabilities (much content and many applications or apps), and are inevitably used for personal reasons (e.g., Web browsing, checking email, downloading applications, conducting online banking, connecting to social networking sites, and so on), even when issued for business use by the organization. And as mobile devices, they can be lost or stolen, exposing their confidential content to unauthorized users—potentially in violation of applicable security regulations.

Of all the mobile devices, laptops are the easiest to secure, and smartphones are the most challenging, with tablets falling somewhere in between. Laptops are relatively easy to secure for three reasons. The first is that they are often issued by the organization, which can then readily control both the operating system and the security provisions. The second is the maturity and broad assortment of security applications such as personal firewalls, and antivirus and anti-spam software, that are readily available. Also, the disk drive of a laptop can easily be encrypted because this capability is built right into the operating system. The third is that users have long recognized the need for securing PCs—both desktop and laptop—and cooperate (for the most part) with the IT department's enforcement of these security provisions.

Smartphones are far more difficult to secure, however. The reasons are essentially the exact opposites of why laptops are so easy to secure. Many smartphones are the personal property of their users, giving the enterprise no direct control over the choice of platform or feature sets. Users often fail to see the need to secure a phone and may even resent the attempt, especially when the mobile phone is their own. For the IT department, the security provisions available might be immature or might only work with one of the many mobile operating system platforms being used, thereby requiring the use of different mobile security software and management systems for each mobile OS.

Fortunately, there is an emerging solution to this problem—purpose-built mobile security applications. An enterprise-class mobile security solution should contain a complete suite of mobile protections, including personal firewall, anti-malware, and spam filtering, and should also support critical mobile device management (MDM) capabilities such as remote wiping or deletion of sensitive data for lost or stolen mobile devices. All of these mobile security protections and device management capabilities should, ideally, support all of the smartphone and tablet platforms being used, enabling them to be configured through a single management console.

Centralized Network Access Control

There is good news and bad news in the effort to unify and centralize access control provisions for the mobile enterprise. The good news is that proven authentication, authorization, and accounting systems and protocols are all battle-tested through years of use. Most organizations implemented AAA long ago using RADIUS or some other directory service. The bad news is that many AAA systems were implemented from a network-centric perspective with a dependency on the port being used for access, which makes sense for a purely wired Ethernet environment. This traditional approach does not make sense, however, in a world where every user is potentially a mobile user. Mobility recognizes no such dependency and indeed, the port changes constantly as users roam from one access point to another.

Fortunately, all existing AAA systems can serve as the foundation for a far more robust and *mobilized* form of network access control—one that is based on user identity; that is, who is attempting to access and use the network rather than which port a user is plugged into (either via an access point or an alternate remote access method). This is the most fundamental difference between wired and wireless networks. In wireless networks, the notion of user mobility is at the core of the security model, which is why all WLAN vendors have implemented some form of identity-based networking. Wired networks simply have not had the same pressing need for a user-centric approach. Nevertheless, most Ethernet switches are now adding support for user identity to unify access control across wired and wireless domains.

But identity-based networking on its own still is not quite sufficient, because the very act of roaming about while connected exposes the network to new security risks. Consider an authenticated user who has access to the accounting system. If the user steps outside the building with the laptop still connected, should access to the accounting resources still be permitted? Or is it more prudent to assume that the device is now in a high risk area and immediately shut down access to corporate resources? Rather than simply attempting to preserve existing user security profiles that give all users consistent access wherever they connect—whether stationary or nomadic or mobile—a more granular approach is needed to modify access privileges in real time depending on what users are doing on which devices, the security posture of those devices, where they are, and perhaps even when they are doing it.

To be fully effective, mobile authentication and access control must be dynamic—both pre- and post-admission—and must be sufficiently granular to accommodate all possible situations and user behaviors. Dynamic and granular control enables network managers to, for example, lock down bandwidth abusers, restrict guest access to controlled areas such as meeting rooms, prevent certain network access based on location or the time of the day or the day of the week, or upon changes in the security posture of the device. But it all begins by authenticating each and every user.

Authentication—By the Book

To ensure that the authentication process works well, the networking industry has adopted the Institute of Electrical and Electronics Engineers (IEEE) 802.1X standard for port-based network access control (NAC), which defines a quite rigorous and solidly secure means of authenticating users. It also affords interoperability with existing RADIUS and other back-end identity stores such as the Lightweight Directory Access Protocol (LDAP) for user and device authentication. The combination of 802.1X and the Extensible Authentication Protocol (EAP) standard from the Internet Engineering Task Force (IETF) creates a solid authentication framework that is depicted in Figure 1.

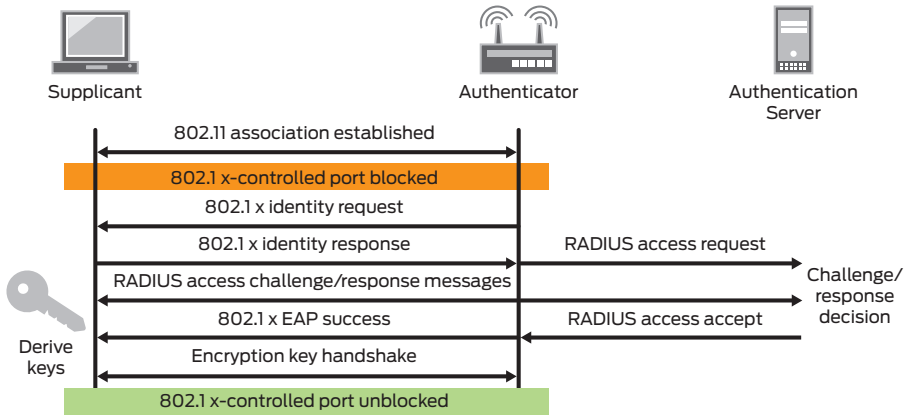

Source: Core Competence

Figure 1: Typical authentication process

Before admitting an endpoint to the network, the user, and ideally the user's device (see *Performing an Endpoint Integrity Check* below), is first authenticated. Ethernet switches and WLAN access points block all traffic from 802.1X clients (which are commonly referred to as *supplicants*), except for EAP traffic (specifically EAP over LAN or EAPOL, and EAP over RADIUS) until the authentication is successful. Once the tethered or mobile user enters the necessary authentication data (e.g., a user name and password or other authentication credentials, including multifactor authentication methods such as token-based one-time passwords, or biometric scans), the endpoint device communicates an authentication request to the switch or access point, which relays it to an authentication server such as a RADIUS server. If the authentication process succeeds, the user is granted access. If authentication fails, network access is denied and the client can be quarantined, depending on the endpoint's security posture.

This simple "yes/no" permission process provides a strong foundation for the preadmission control portion of network access control. The 802.lX standard ensures interoperability among 802.1X compliant clients running on endpoints, 802.1X-enabled switches and access points, and network access control servers, allowing the enterprise to leverage the existing infrastructure. (For more detailed information on how this works, read the Juniper Networks white paper "802.1X: Port-Based Authentication Standard for Network Access Control (NAC)" at **www.juniper.net/us/en/local/pdf/whitepapers/2000216-en.pdf**.

Performing an Endpoint Integrity Check

In a mobile environment with numerous different types of devices—both corporate issued and user owned—NAC is now often dependent upon the results of endpoint integrity checks. While the endpoint can mean both the user and his or her device, the integrity check focuses on the device and is a test of its security posture. Devices with robust security provisions, such as those supported by a mobile security solution, may be permitted greater network access. It is not uncommon (and is becoming increasingly prudent) for devices lacking security provisions, or not meeting enterprise security policies such as for up-to-date antivirus software, to be isolated or quarantined until their lack of compliance has been remediated.

Endpoint integrity checking requires cooperation among multiple applications and networking systems to ensure that the check itself, along with possible quarantining, remediation, and alerting, all function properly. This interoperability has been made possible through both existing industry standards and the work of the Trusted Network Connect (TNC) workgroup of the Trusted Computing Group, which has developed a set of open industry standards for network security and access control. Network elements that comply with the TNC standards ensure that each endpoint accessing the network has passed an integrity check before it can be granted any access. And if the endpoint is determined to pose a security risk, a variety of TNC protocols may be employed to subsequently ensure that either a device's attempted network access is denied, or the device is quarantined to a secure segment of the network.

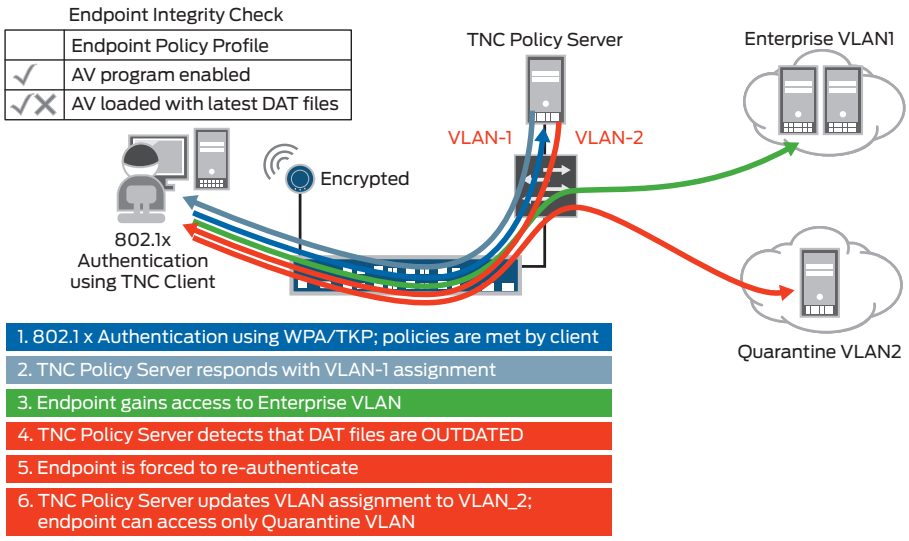

Figure 2: 802.1X authentication using TNC client

Authorization to Access Resources

As explained above, authentication in an 802.1X environment is a yes/no process. Either the user's access is accepted or not, and if it is, several different EAP types are made available to provide the keys for encryption during the user's network session. Authorization is where—based on the user's identity or role, and ideally, the security posture of the device—a rich set of access and security polices can be established and conditional restrictions can be applied. Sometimes referred to as access control rules, authorization is the "We trust you are who you claim to be, now here's what we will allow you to do" portion of NAC. Authorization is where the real power of distinguishing among different types of users and devices, and levels or nuances of access, come into play. Some users, based on identity, role, location, etc., may be granted broad access privileges; others, such as guests, may be granted restricted access, perhaps significantly to just one or a few resources.

Examples of different options during authorization include:

- Requiring the use of encryption or even a particular encryption type
- Restricting access to certain servers based on role or departmental association
- Allowing broader access from only those systems that are up-to-date with enterprise approved antivirus or anti-malware software (See the previous section on *Performing an Endpoint Integrity Check.*)
- Limiting access to only the Internet
- Allowing/disallowing access for different users or groups based on time-of-day or day-of-week, or even from what location access is requested

There exists a wide variety of means to enforce authorized access privileges throughout the enterprise IT infrastructure, including in the servers or applications, and in the network. Listed here are just some of the means available in the network itself. These techniques can also be combined and customized to create a strong authorization model that provides solid security, yet requires minimal administrative intervention.

- Virtual LAN membership—Users on the same service set identifiers (SSIDs) can be joined to different VLANs based on their identities. This is done while maintaining full separation between the VLANs.
- Time of day/day of week—Guests, contractors, employees, or other users may be restricted from accessing certain resources on weekends or after hours.
- Simultaneous logins—Perhaps a visiting vendor has no business being on two machines at the same time, while employees are allowed to have their laptops, smartphones, and wireless VoIP phones running simultaneously.
- Quality of service (QoS) profile—This maps specific bandwidth control settings to a user's identity or role. It can also be used for dividing bandwidth between SSIDs as well (e.g., when contention occurs, guests get 10% and employees get 90%).
- Bandwidth usage—With Dynamic Authorization (RFC 5176), it is possible to change any authorization attribute on the fly, even after the network session is up and running (referred to as post-admission access control). This makes it possible, for example, to quench bandwidth abusers by terminating their sessions or by throttling down their bandwidth after they have exceeded a certain threshold within a specified timeframe.
- Stateful inspection—Another example of Dynamic Authorization involves stateful inspection of application traffic to make certain desired changes. For example, an employee using an application known to have vulnerabilities can be diverted to a quarantined network area.
- Firewall filters—These filters can be applied to users based on their identity or group membership. This is very helpful for isolating or separating certain traffic such as Session Initiation Protocol (SIP) voice traffic coming from a laptop rather than a known voice device, and treating it differently than other data coming from the same device.

Providing Secure Guest Access

Most organizations have guests, whether customers, contractors, or partners, and these routine visitors want (and deserve) at least some network access privileges. Access to the network enables these guests to conduct business more productively, which also benefits the host organization. Just like employees, these guests desire access through a myriad of mobile devices over which IT has no control.

But could guest access present a threat to the enterprise network? Unless guest access and activity is restricted somehow and monitored continuously, the threat can be very real. Uncontrolled guest access could potentially use the network to send spam or maliciously attack other network users. Guests could secretly capture data from other users or from the data center, or simply drain available Internet bandwidth. In other words, even if the guest network is completely isolated from corporate data, there is still no reason to allow open access for all guest users.

Here are three suggestions for satisfying the needs of both guest users and the host IT security manager:

- The guests should be welcomed. They should be real guests, and not some hacker in the parking lot. Just like issuing a visitor badge, it is best to have some modicum of authentication through a receptionist to grant the minimal authorized access required. This could be done via a designated SSID or by creating a special guest account in the AAA server and giving this information only to legitimate guests needing access. Of course, the SSID or guest ID and password should be changed regularly, even daily.
- Guest access should be easy and compatible for everyone. Guests will bring in a variety of devices, and because endpoint integrity checking is an invasive process, the best way to ensure seamless compatibility is to restrict guests to a secure portion of the network. For most users, having access limited only to the Internet is sufficient. Optionally, some guests might be granted access to a special portal or walled garden that permits some additional capabilities. KISS is the best approach here: Keep It Simple and Secure.
- There should be automated mechanisms to purge expired guest records and easy provisioning tools to allow non-IT personnel to grant access to different classes of guests within the parameters established by IT—but without bothering them for maintenance or day-to-day guest access provisioning. And, ideally, any guest access mechanism should track who enabled the guest's access to ensure regulatory compliance and corporate accountability.

Other Network Access Control Best Practices

Here are some additional examples of how mobility-centric dynamic NAC provisions can be made more granular to accommodate different users, devices, and applications:

- A combination of user, device, time of day, and location can be used to grant anything from full access to no access at all. Note how time of day and location require post-admission controls.
- A more restrictive guest access account can be created for certain devices, such as those with both Wi-Fi and cellular communications (3G/4G) capabilities.
- Groups of users can be isolated in ways that prevent eavesdropping or tampering. This can be done by using different SSIDs and optionally segmenting the traffic for these groups onto different VLANs.
- Certain devices, especially those with an unknown security posture, can be assigned to a specific SSID or VLAN to segment traffic on the network.
- Different firewall rules and filters can be applied to different combinations of user groups or devices.
- Wireless intrusion detection/prevention systems (WIDS/WIPS) can be employed to provide another layer of security for users with endpoints that are not secure. While most wireless LAN solutions may be able to detect some of the more common threat types, such as rogue access points and denial-of-service (DoS) attacks, most institutions, especially those in the financial and government sectors, will prefer more advanced WIDS/WIPS solutions capable of uncovering hundreds of different attack types on wireless LANs.

The ultimate in flexible and granular access control requires one additional capability—deep inspection. DI is the most accurate and often the only way to determine the specific applications and protocols being used in any session. Normally such awareness is used only to apply an appropriate QoS profile during a network session. But this same deep awareness can and should have a more fundamental role in the mix of criteria used to determine the appropriate level of network access.

Distributed Enforcement

Distributed enforcement is achieved by propagating security credentials out to the furthest edge of the network where access originates. For the wireless LAN, this normally happens via the controllers, thereby enabling each controller's access points to instantly recognize existing authenticated and authorized users who roam into range. This approach has the advantage of delivering the fast, secure roaming needed to avoid disrupting Voice over Wireless LAN (VoWLAN) calls and latency sensitive data applications, even when crossing major network boundaries such as roaming across controllers, or from indoors to outside. Although fast roaming is not a security consideration per se, centralizing enforcement at the controllers inevitably creates a bottleneck that does have a negative impact on quality of experience (QoE) for all users.

One aspect of distributed enforcement that does have a direct effect on security involves the endpoints of encrypted communications. The encryption function provides the privacy enforcement behind authentication and authorization and is, therefore, a critical part of building a secure network.

There continues to be a debate about where the encryption and decryption boundaries should exist in a wireless LAN for client data. The most obvious answer is over the air between the client and the access point, which is where it was originally intended (the distributed encryption model). Some, however, contend that it is better to extend the encrypted session of the user back through the wired network and terminate it on a wireless LAN controller located in the data center or core (the centralized encryption model). Traffic would then be decrypted at the controller, where it is subsequently placed back onto the same wired network in order to reach its destination.

The centralized approach is manifest from the early days of WLANs when the old and ill-fated Wired Equivalent Privacy (WEP) protocol was broken. With WEP, the only sure way to secure the network over the air was to run a VPN on every client, which all terminated on a centralized VPN gateway. The data traffic was then sent unencrypted from the gateway over the wired network. This approach has numerous problems, not the least of which is to limit capacity and scalability. With the new, improved Wi-Fi Protected Access protocols (WPA and WPA2) replacing WEP, the distributed encryption model is solidly secure, and is far superior at scaling the WLAN to accommodate the proliferation of mobile devices.

From a pure security perspective, it is also worthwhile considering the U.S. Department of Defense's view on centralized vs. distributed encryption as defined in DoD Directive 8100.2[1]. The DoD, when directing how to secure WLANs, defines three fundamental elements: End-to-End, Assured Channel, and Security Border. These all refer to where the network boundary of control ends, at which point the traffic then traverses a network potentially subject to eavesdropping. From the DoD's point of view, there is no incremental security value in performing encryption/decryption somewhere behind the Security Border. In fact, doing so would only serve to obfuscate where the Security Border is actually located. In a wireless LAN, the DoD considers the wireless access point as the Security Border.

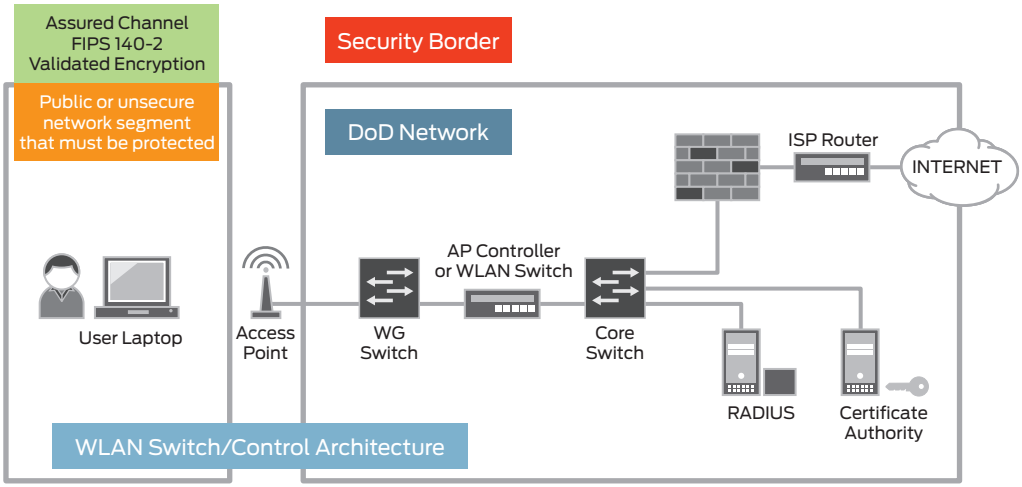

Figure 3: Overview of DoD Directive 8100.2 security architecture

But what about the possible trade-offs in a more typical enterprise wireless LAN? Here the additional load and overhead placed on the controller in the centralized encryption model can create a bottleneck, especially in 802.11n deployments. The distributed encryption model, by contrast, leverages the built-in, additive processing power of all access points, where encryption/decryption capabilities are now included in the chipset. Table 1 provides a summary comparison between the two alternatives.

Table 1: Centralized vs. Distributed Encryption

	CENTRALIZED ENCRYPTION MODEL	DISTRIBUTED ENCRYPTION MODEL
Where is the Security Border?	The controlled network begins and ends at the WLAN controller.	The controlled network begins and ends at the access point.
Does this architecture comply with DoDD 8100.2 guidelines?	Yes	Yes
Does this architecture comply with FIPS 140-2 requirements?	Yes	Yes
How does performance scale as more access points are added?	Because all data forwarding and encryption/decryption must flow through the controller, capacity can be increased only by adding more controllers.	With encryption/decryption distributed to access points, sufficient cryptographic computing power already exists to match the additional traffic load.
Does the architecture support distributed forwarding?	No. With all encrypted traffic terminating at the controller, the controller must also forward the traffic to its destination.	Yes. Solutions that support distributed forwarding can route traffic directly among access points, relieving controllers of the forwarding load.
Is it possible to preserve sessions in redundant, failover scenarios?	No. If a controller fails, encrypted sessions will be lost, and clients will need to re-authenticate to the backup controller.	Yes. If a controller fails, all active sessions can be preserved in real time by switching them over to another active controller.

[1] DoD Directive 8100.2 - Use of Commercial Wireless Devices, Services, and Technologies in the Department of Defense (DoD) Global Information Grid (GiG) **https://acc.dau mil/CommunityBrowser.aspx?id=153484&lang=en-US**.

Location, Location, Location

Location matters in today's mobile world. This consideration in the context of security can be divided into two situations—on and off premises. The latter situation is fairly straightforward to handle because it does not represent a significant change for the IT department. Organizations have long used VPN technology to grant users access to private network resources via public network infrastructures. VPN technologies are fully standardized, and both SSL and IPsec are fully interoperable with network access control provisions such as RADIUS. Virtually every device now supports at least some form of virtual private networking. As a result, proven VPN policies and practices already exist in virtually all organizations for employees working from home or public hot spots.

Within the confines of the enterprise network, however, location takes on an entirely different and far more granular meaning. Location in the context of secure mobility services for an enterprise network goes well beyond knowing that a user is currently connected to a specific VLAN or SSID. Any network should know at all times which users are currently connected via which switch port or access point. This is useful information, to be sure, but the ability to precisely pinpoint the user's (or device's) actual physical location in real time has many far more useful applications, ranging from enhanced security and asset tracking to better resource management and troubleshooting. The growing need for more granularity has led to the creation of an entirely new capability for wireless LANs—Real Time Location Services or RTLS.

RTLS employs triangulation, trilateration, or other means to pinpoint a mobile device's location, and hence its user. Although the degree of precision can vary among RTLS solutions, all are far more precise than determining location based on the access point being used. Consider the situation where a user walks into a public area outside of a building, but remains connected to an access point located just inside. While it may be acceptable to grant some users with some devices full access while outside, access for others (especially guests) should be restricted in such situations to, for example, the Internet and VoWLAN calling. Another example is using RTLS in what is commonly called geo-fencing for creating a perimeter radio frequency (RF) firewall.

In addition to this location-based access control, RTLS has many other useful applications. A critical location-based service (LBS) that benefits from RTLS is Enhanced 911 (E911) service for VoIP, which also has security implications—albeit not for the network per se. A caller with a medical emergency, for example, needs to be located precisely and without delay. Ideally, others nearby would be contacted and asked to render assistance. In certain emergencies such as a fire, locating and notifying everyone in the vicinity is just as critical.

As an integral part of the security infrastructure, a robust RTLS solution should be as effective outdoors in a campus setting as it is indoors. Unfortunately, some WLAN solutions are not so seamless, which can make precisely pinpointing user locations outdoors quite challenging. For example, some WLANs have completely different product families built on different architectures for outdoor and indoor deployment, while some have no or only limited RF planning capabilities for outdoor deployments. The reason for the lack of seamlessness is simple—indoor and outdoor networking requirements are quite different. Outdoor wireless LANs require features like bridging, meshing, filtering, special antenna systems, and more to overcome issues like the lack of access to wired Ethernet and other technical obstacles related to limited bandwidth and higher range requirements.

Users shouldn't need to care about any of this, of course; they simply expect seamless mobility (with appropriate access privileges) as they move about the campus and beyond, in and out, to different locations. The IT staff does need to care, however, which makes the need for a common indoor/outdoor architecture beneficial, particularly as the number of always-on wireless devices continues to proliferate. The common architecture should also have a common management system and common network access controls to minimize the burden on IT staff. Such seamless integration enables the service and security profiles to be extended end-to-end across the entire network. Then and only then will users enjoy the seamless, secure mobility services they expect—and deserve.

Conclusion

The proliferation of mobile devices is having a profound effect on employees and enterprises alike. Gone are the days of being able to control the devices employees use. More users with more devices running more applications from more places demands more changes to the enterprise network. The security practices outlined here will hopefully help the IT department cope successfully with these changes in the Unwired Enterprise.

The ultimate best practice is to use only those solutions proven in highly mobile environments. Such an approach will, naturally, place emphasis on the wireless LAN portion of any solution. But the same network access control policies and provisions should apply equally to both wired and wireless access, as well as mobile and local devices. This approach also ensures that security provisions have a mobility-centric architecture that possesses the granular user, device, and location awareness necessary to create and consistently apply suitable access controls.

For more information about how your organization can benefit from a comprehensive, unified network access and application control solution for wired and wireless LAN, as well as local and mobile users, visit Juniper Networks on the Web at **www.juniper.net/uac**.

Appendix

Included in this appendix are some additional topics that, while not particularly relevant to most organizations today, may nevertheless be of interest to some readers.

802.1X vs. PSK

802.1X in combination with EAP is the standard for authentication that is designed around authentication requirements of large networks, and it offers several key elements that are necessary when scaling secure network access control. These elements include, among others:

- **Per user authentication**—The ability to authenticate and preserve, for accounting purposes, the identity of the user.
- **Per session encryption**—This EAP function allows each session to have unique encryption keys and thus disallow snooping, even by other authenticated users.
- **Integration to existing NAC equipment and standards**—Using RADIUS servers that either already exist or are easily deployed.

The PSK or pre-shared key authentication technique has none of these key elements. In the PSK scenario, any device with the key (password) can gain access to the network. PSK is easy to deploy, but is not a secure authentication technique for anything larger than the smallest enterprises. Due to the nature of pre-shared keys, more users means that it is too easy for the key to get into the wrong hands and too difficult to propagate a change if the key becomes compromised.

For these reasons, an enterprise should avoid using the PSK authentication technique that is intended only for home office/small office deployments by the Wi-Fi Alliance. There are, unfortunately, many older devices and wireless phone sets that are only capable of PSK. In these cases, it is usually desirable to substitute more authorization restrictions to make up for the weak authentication. For example, if PSK must be used, it should be restricted and isolated to its own subnetwork with restrictive access control lists.

In summary, WPA's Temporal Key Integrity Protocol (TKIP) is showing its age, and though useful for devices with old chipsets, the use of WPA2 and its Advanced Encryption Standard / Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (AES/CCMP) encryption has been and remains the only choice for the U.S. government's most secure wireless networks. For authentication, there is a wide choice of robust 802.1X/EAP implementations approved as part of WPA2, which provide a scalable and secure authentication technique appropriate for the enterprise in contrast to PSK, whose focus is for simple home and small business networks.

Wireless Intrusion Detection/Protection Systems

Much discussion occurred regarding wireless intrusion detection and protection in the earlier days of Wi-Fi networks. Indeed, this made sense as robust AAA implementations were difficult to deploy and had some missing elements, and as the industry quickly learned, the *de facto* security protocols such as WEP later turned out to be not as secure as originally thought. IT managers had to rely on dedicated IDS/IPS systems to stop intruders because there was not a strong enough authentication and authorization architecture in place. But today, this market is not keeping pace with wireless LAN deployments, and in many cases it is declining. Further, the principal intrusion threats, such as rogue access points and DoS attacks, are now easily detected with the base-level IDS/IPS feature set built into most access points.

Today, most dedicated wireless IDS and IPS solutions are deployed where:

- There is no intention of supplying a wireless LAN, yet wired services do exist. In this case, the IDS/IPS system is deployed solely to prevent the unauthorized creation of a wireless LAN to keep employees or intruders from installing rogue access points or creating an ad hoc network that attaches to the internal wired network.
- The wireless devices are old legacy devices with outdated security capabilities, thus limiting the AAA architecture that can be deployed. In this case, the open path to the internal network required by the legacy devices must be augmented to prevent intruders. The most common environment for this is retail, where Payment Card Industry Data Security Standards (PCI DSS) v1.2 mandate not deploying WEP and "… using strong encryption technologies for wireless networks, for both authentication and transmission." But many retailers are loath to replace handheld terminals, barcode scanners, and other systems in order to comply. The alternative of extending their life a few more years by augmenting the existing solution with WEP cloaking schemes offered by dedicated IPS/IDS infrastructure vendors seems more attractive in the short term.

Dedicated IDS/IPS systems are typically not installed where the customer is already running or planning to run a wireless LAN with WPA/WPA2 capable devices. This is because they offer very limited incremental value. Most WLAN vendors, in addition to a strong approach to AAA using WPA2, provide much of the primary functionality for DoS and intrusion detection formerly found only in IDS/IPS systems. For example, Juniper WLAN equipment is fully capable of identifying, alerting, locating, and automatically combating rogue access points and their users. There are additionally over 40 different IDS and DoS detection functions built into the existing equipment. These include flooding detection techniques using de-authentication, disassociation, and decryption error frames, use of RF jamming, fake access point flooding, spoofed access point and SSID masquerading, detecting the use of popular sniffing/spoofing applications, presence of a wireless bridge, and use of weak encryption keys, among others.

Dedicated systems are too expensive for most enterprises, but they do have additional IDS/IPS features beyond most built-in capabilities. A strong implementation of AAA combined with included rogue detection, IDS, and DoS functions, however, often offsets the underlying need and management overhead of dealing with the many false positives that standalone IDS/IPS systems tend to produce. This and the complexity of integrating and maintaining security policies across two separate wireless systems often cause more problems than they solve.

Voice over IP over Wi-Fi Security Considerations

Voice conversations need the same attention to detail when it comes to security as data does. There is, in fact, often a higher expectation of privacy between the two parties. The same AAA architecture can and should be used for understanding the security implementation of VoWLAN to ensure proper authentication, authorization, and accounting.

Here again, using a standards-based approach like WPA2 utilizing 802.1X is the ideal. The special QoS requirements for voice can be handled in multiple ways, such as through standards-based Wi-Fi Multimedia (WMM)[2] or through QoS attributes assigned via the authorization step of AAA. Additional vendor differentiation beyond WPA2 can be useful, but these need to be understood within the context of a particular enterprise's needs. Voice traffic can also be placed on its own VLAN with wired VoIP phones, leveraging the existing QoS configuration over the wired network.

[2] Wireless Multimedia Extensions (WME), also known as Wi-Fi Multimedia (WMM), is a Wi-Fi Alliance interoperability certification, based on the IEEE 802.11e draft standard. It provides basic QoS features for IEEE 802.11 networks. WMM prioritizes traffic according to four Access Categories (AC)—voice, video, best effort, and background. It is suitable for simple applications that require QoS such as VoIP on Wi-Fi phones.

But beware of letting the limited security capability of a particular handset determine the security architecture. In order to use a high security implementation like WPA2, the clients, or in this case the Wi-Fi handsets, must also support WPA2. Not supporting this implies some other, likely less secure, approach to AAA. Too often, voice handsets can only use a weak authentication technique, such as media access control (MAC) authentication combined with a weak authorization such as WEP encryption, and this can undermine the entire security infrastructure of the enterprise network. Though additional authorization attributes can be used to restrict the access of these weaker handsets (or a device pretending to be a handset), this is at best a stopgap approach to shore up security holes that can only be resolved by using handsets with robust AAA capabilities such as WPA2 and 802.1X.

Good Encryption May Not Be Good Enough

Though encryption is important, it is only a small part of the more important AAA architecture of a secure network. The IEEE (and therefore WPA and WPA2) make specific recommendations for secure authentication which must also be followed properly. The context here is an enterprise deployment of 802.1X with an EAP type that matches the vendor compatibility and administrative needs.

The primary difference between WPA and WPA2 is WPA2's compliance with a purpose-built encryption technique, AES/CCMP, which has been tested by National Institute of Science and Technology (NIST) for the government's Robust Security Network (RSN) compliant networks.

By contrast, WPA used Temporal Key Integrity Protocol (TKIP), which was developed as a temporary work-around on the old WEP protocol—the latter having been found to be seriously flawed in 2001.[3] TKIP was designed as a software upgrade that would be hardware-compatible with the old chipsets supporting WEP. Most devices shipped after 2003 support both TKIP and the more efficient and secure AES/CCMP. TKIP did, however, inherit some known weaknesses, and in 2008 researchers discovered an inherent flaw that could allow a re-injection and spoofing of short packets to a client.[4] Though the encryption was not broken, it might then allow for subsequent spoofing attacks involving short packets like Address Resolution Protocol (ARP) or Domain Name System (DNS), which is by no means an easy exercise, but is still possible.

Because TKIP was never designed as the final, secure solution but rather as a temporary work-around, the Wi-Fi Alliance began in 2006 providing certification only for products that do both WPA and WPA2, and no longer provides certification for products that meet WPA only.[5] Additionally, the IEEE will also extricate TKIP entirely from the 802.11 base standards.

Special Requirements for Certain Organizations

Organizations in some industries have special security requirements, either as a best practice or, more likely, the result of regulatory control. Highlighted here are such requirements in the U.S. for three different industries—the federal government, healthcare, and financial services, including retail.

Federal Government

U.S. federal government agencies and contractors are required to adhere to Federal Information Processing Standards, including the FIPS 140 series specifying requirements for cryptography modules related to computer security and data communications. The current version is FIPS 140-2. FIPS 140 is intended to coordinate both hardware and software requirements and standards for use by any/all departments and agencies of the U.S. government. Being certified at a particular level of the FIPS 140 requirements is not sufficient for building a secure network, but is thought to be necessary by government entities.

In addition to the agencies themselves, private contractors to the U.S. government are often required to comply with aspects of FIPS 140 as part of carrying out a government contract. This can affect the contractor's internal and external design and operation of its network.

A FIPS 140-2 certification demonstrates that the product's relevant security features have been thoroughly validated and documented in terms of its hardware and software design. These are done to four specific levels of security that various agencies can designate as a requirement. The Juniper wireless LAN solution has been validated for compliance with FIPS 140-2.

[3] See Nikita Borisov, Ian Goldberg, David Wagner, "Intercepting Mobile Communications: The Insecurity of 802.11" **www.isaac.cs.berkeley.edu/isaac/mobicom.pdf**.
[4] "Battered, but not Broken: Understanding the WPA Crack". Ars Technica (2008-11-06).
[5] See "WPA2 Security Now Mandatory for Wi-Fi CERTIFIED™ Products" **www.wifi.org/pressroom_overview.php?newsid=16**.

Rigorous FIPS Requirements in 11 Areas

FIPS 140-2 imposes software and hardware requirements across 11 different areas, and based on capabilities, these are placed in one of four different levels. The 11 areas are:

- Cryptographic module specification and documentation
- Cryptographic module parts and interfaces (flow of sensitive information, how secure/insecure information is segregated in both hardware and software)
- Roles, services, and authentication (who can do what administratively, and how this is checked)
- Finite state model (documentation of the various states the cryptographic module can be in, and how transitions occur)
- Physical security (primarily tamper evidence and resistance)
- Operational environment (the operating system the module uses and is used by)
- Cryptographic key management (generation, entry, display, storage, and deletion of cryptographic keys)
- Electromagnetic interference/electromagnetic compatibility (EMI/EMC)
- Self-tests of cryptographic modules (what must be tested and when, and how failures are handled)
- Design assurance (documentation demonstrating good design and implementation)
- Mitigation of other attacks (if a module's function is designed to mitigate an attack, how is this done)

Four FIPS Certification Levels

- FIPS 140-2 Level 1 is the lowest, imposing very limited requirements; loosely, all components must be "production grade" and various common forms of insecurity must be absent.
- FIPS 140-2 Level 2 adds significantly more documentation requirements for cryptography and describing state models. It also adds requirements for physical tamper evidence and role-based authentication. This is the most commonly required FIPS level for government agencies, and this is the level of certification attained by the few WLAN vendors that have validated FIPS 140 solutions.
- FIPS 140-2 Level 3 adds requirements for physical tamper resistance, identity-based authentication, and for a physical or logical separation between the interfaces by which "critical security parameters" enter and leave the module, and its other interfaces.
- FIPS 140-2 Level 4 makes the physical security requirements more stringent, and requires robustness against environmental attacks.

Examples of FIPS Certifications

To see examples of the various reports involved with FIPS certifications, go to NIST's "Cryptographic Module Validation Program" at **http://csrc.nist.gov/groups/STM/cmvp/**. Available here are examples of "Security Policy" summary documents and the actual certifications for various products.

Healthcare

The Health Insurance Portability and Accountability Act (HIPAA) is a U.S. federal law creating security standards to ensure the privacy of patients' medical records and personal health information. It is intended primarily for those organizations having access to patient medical data such as hospitals, healthcare providers, and insurance companies. Any network, including a wireless LAN, used by these organizations must comply with HIPAA security standards. Also, any "business associate," which is a broad term applying to any equipment or services used in the transfer of confidential information, would also be considered subject to HIPAA requirements. As such, Juniper Networks would be considered a "business associate."

The main thrust of the security requirement is simple. Access to electronic medical information must document the individuals who accessed files, when they accessed them, and for what purpose they accessed them. HIPAA identifies five key areas for secure electronic transfer of patient records:

- **Authentication**—The first A in AAA, which ensures that the system knows who is using it.
- **Authorization**—The second A in AAA, which ensures that authenticated individuals access the network based only on a defined set of privileges.
- **Confidentiality**—Confidentiality ensures that only persons authorized to access records are allowed to do so. Conversely, users that are not properly authenticated and authorized cannot gain access to private patient data.
- **Integrity**—This ensures that data has not been manipulated en route, which, like confidentiality, can be thought of as robust authorization that can be provided through the use of session-based encryption keys (e.g., WPA2 using 802.1X).
- **Non-repudiation**—Once access is granted, neither the sender nor receiver can deny it took place. Here there is an implied requirement for accounting (the final A of AAA). Only by doing the first two A's properly (authentication and authorization), can good accounting track access to the network. In HIPAA networks, both network and database or application accounting are needed to identify what particular records are accessed and by whom.

It is noteworthy that although the thrust of HIPAA is access to patient data, the requirements can extend to discussions related to patient data that might occur. For example, using VoIP to discuss a patient should follow the same HIPAA compliance requirements. In this case, the network and call application software must account for who is talking when and to whom.

In summary, the HIPAA security requirements, like any other security requirements, can be thought of within the context of the AAA architecture for wireless LANs. Proper use of WPA2 and 802.1X normally provides compliance with all five areas of Authentication, Authorization, Confidentiality, Integrity, and Non-repudiation.

Financial Services

Banks and other financial institutions, including retail establishments, often provide Internet access for visitors and contractors. Historically this was achieved by providing a completely separate guest network that is isolated from the corporate network, and using a dedicated DSL router at each location offering such services. But this is really no longer necessary, because modern WLANs make securing such guest access an integral part of authentication and authorization.

The best way to provide secure guest access is the same as it is for providing secure employee access. This approach is more cost-effective because it uses the very same 802.1X authentication process, and the very same network access control provisions. What is different are the policies that apply. In some circumstances, the institution may want to use WIPS/WIDS, RF firewalling, VLANs, and other layers of security, which can also serve to block non-authorized access from people outside the premises.

After some disturbing early implementations of wireless networks by retailers that resulted in *war drivers* being able to pick up credit card numbers by simply being in the parking lot and using a wireless sniffer, the Payment Card Industry (PCI) extended its Data Security Standards (DSS) to include wireless networks. These standards are currently up to v1.2[6] and have notably required their members to discontinue the use of WEP. Instead, they mandate requirements for "strong encryption technologies for wireless networks, for both authentication and transmission," such as WPA2 with 802.1X. This is a further recognition that the AAA architecture using the Wi-Fi Alliance's standards matters. Juniper Networks is of course, fully compliant with PCI DSS v1.2.

[6] **www.pcisecuritystandards.org/security_standards/pci_dss.shtml**

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

2000420-001-EN July 2011 Printed on recycled paper