



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

**Kenneth Butler
KABUTLE001
GSEC Practical V1.4B
Mentor Led Training
Hartford, CT
August 1, 2003**

Untrustworthy Laptops

Protecting the Infrastructure from Worms and Unauthorized Access

© SANS Institute 2004, Author retains full rights.

Summary

Organizations face two problems caused by the proliferation of laptops and worms. A compromised laptop belonging to a visitor may be brought onsite and release a worm, or an unpatched corporate laptop may be compromised and release a worm. Controlling access to the network itself will block unknown laptops, but not known, compromised ones. Some solutions exist to check a laptop's patch level before permitting network access, which can include patch level, but it may not integrate with a solution to control access from unknown devices. Existing solutions, while not perfect, can bring improvements in network security.

Introduction

Early this September, a consultant from our ad agency brought a laptop on site and attached it to our network. This laptop carried the Nachi worm¹ and released it into our network. (Nachi is a "beneficial" worm that exploits the RPC vulnerability in Windows operating systems.

The situation could have been much worse. First, our most of our systems patched; the only machines affected were a few in the process of being built. We now keep such systems offline until they are up to date on antivirus and patches. Also, we have been considering an upgrade of our Automatic Teller Machines to Windows based systems, but we were still using OS/2 based teller machines at the time. Banks that had made the switch discovered that their teller machine visitors were too lax about security patches.²

In fairness to the vendor, our first Windows based teller machine came in with patches for MS03-026³ and MS03-039⁴ applied, as determined by Microsoft's tool for testing for the vulnerability.⁵

At our "Lessons Learned" meeting, we discussed the procedural and policy problems that lead to the incident. However, as we examined the problems, it became clear that all laptops present unique risks.

1. Departments misplace laptops. Some missing laptops resurface after being having been borrowed for an extended time, and come back with missing patches, old anti-virus software, viruses, spyware or worms.

¹ "Network Associates Virus Information Library", http://vil.nai.com/vil/content/v_100559.htm

² Poulsen, <http://www.securityfocus.com/news/7517>

³ TechNet, <http://www.microsoft.com/technet/security/bulletin/MS03-026.asp>

⁴ TechNet, <http://www.microsoft.com/technet/security/bulletin/MS03-039.asp>

⁵ Knowledge Base, <http://support.microsoft.com/support/misc/kblookup.asp?ID=827363>

2. Departments will bring their visitors with laptops onsite and attempt to help them to the internal network without consulting IT first.
3. Since laptops may or may not be connected to the network at any given time, it is difficult to confirm that patches have been applied, anti-virus is up to date and firewall software is functioning correctly. We are dependent upon users to update AV signatures, bring in laptops for patches and not tamper with firewall software.

Our organization is not unique. "Vince Tuesday", who writes the Security Manager's Journal for Computerworld, describes how a laptop with old anti-virus signatures and a wide-open personal firewall brought a worm into the network that he supports.⁶

The security model for many organizations consists of a well-considered, layered defense of the Internet presence, including a firewall or two, a DMZ, intrusion detection both host-based and network and hardened hosts running a minimum complement of services. The devices inside this perimeter are implicitly trusted.

Laptops that may or may not have the latest patches, that may or may not be authorized for connection to the internal network and that may have just come from a network with standards very different from yours cannot be considered to be trustworthy. A new layer of defense is needed to prevent unauthorized laptops from connecting to confirm that laptops have current anti-virus protection, the latest hot fixes and a functioning personal firewall before they have unfettered access to network resources.

A note on usage and scope

For brevity's sake, I will be using the term *patch* to refer to service packs, hot fixes and current anti-virus definitions and correctly configured personal firewall. *Patch control* will refer to assuring current versions of anti-virus files service packs and hot fixes. *Patched* will mean a laptop with the latest hotfixes and correctly configured anti-virus and personal firewall, and *unpatched* will refer to laptops that lack patches, anti-virus updates or appropriate firewall policies.

By focusing on laptops, I am ignoring PDAs and other devices that could cause problems. Laptops present a larger immediate threat, and since any technique that proves useful in dealing with untrustworthy laptops could probably be adapted to deal with untrustworthy PDAs.

⁶ Tuesday,

<http://www.computerworld.com/securitytopics/security/holes/story/0,10801,85369,00.html?f=x584>

Finally, I am ignoring issues specific to hackers posing as legitimate visitors to gain inside the perimeter access. Solutions discussed below will mitigate risks associated with malicious unauthorized access.

© SANS Institute 2004, Author retains full rights.

Controlling Network Access

Any visitor to an organization, a sales or service representative, a consultant or an auditor, is likely to bring a laptop with them. In the last few years, visitors seem more likely to want to have access to resources on the local network for their laptops. Perhaps they want to get to the Internet to demonstrate a product or check their email, or they want to print to a network printer.

In an ideal world, visitors with laptops would check in with IT. We could help them get the access they need, and confirm that their laptop appropriately patched and did not represent a threat to our infrastructure. Instead, we find out about people who want to connect to our network when they can't get the resources that they want.

The End User Community

The road to network security problems is paved with good intentions. Our corporate culture places a high value on service, and our employees pride themselves on being friendly and helpful. To our end users, helping visitor get hooked up to the network seems like common courtesy and treating a visitor's laptop as suspect seems rude.

Our organization has put a computer security curriculum in place, and all are end users are being trained. Once our people understand that the best way to help a visitor is to refer them to IT, they will become the network's first line of defense.

Physical Access to the Network

The most basic approach to controlling physical access to the network is to disconnect unused switch ports and lock wiring closet doors. This means that a visitor must ask for a place to plug in their laptop. Hopefully, the visitor will ask IT and this will afford us a chance to examine the laptop in question, rather than just disconnecting a workstation and taking its network drop. This approach does not protect the network from an unpatched user laptop, since the user probably already has a live drop available to them.

Advantages

- Low cost
- Easy to implement and understand

Disadvantages

- Easily defeated
- Does not address the problem of unpatched laptops directly

Conclusion

Disconnecting unused drops is an easy way to add some protection against unauthorized access to the network.

Network Access Control via MAC Address

Many switch manufacturers provide the ability to limit access to switch ports based on MAC address. Cisco (our switch of choice) not only provides this capability, but also makes configuration somewhat painless, through dynamic MAC assignment of port security.⁷

For example, the command:

```
Set port security module/port enable
```

will limit the specified port on the specified module to traffic from the MAC addresses that were in the MAC Address Table at the time of configuration. Other traffic will be dropped.

It should be noted that a determined visitor could reconfigure their laptop with a permitted MAC address *if they knew which addresses were permitted*. This information could be obtained from the workstation they intend to impersonate, or from the ARP cache of a device in the same segment as the device they intended to impersonate.

Advantages

- Low capital cost – if existing equipment supports port security
- Additional security compared to disconnected drops.
- Administrators can choose to block a laptop if it isn't known to be patched


Disadvantages

- Additional labor cost
- Complicated moves, adds and changes
- MAC addresses can be spoofed
- Protection against untrustworthy laptops is indirect, and must be manually administered.

Conclusion

Port security can be used to improve security in a relatively static environment. The configuration below could be used to manage a group of mobile users.

⁷ Odom et al., Pgs 410-412



All of the mobile users have desks in an area serviced by the same workgroup switch. Each port permits the list of MAC addresses for each of the mobile users, as does the port on the backbone switch. Each mobile user can use any desk.

If a new security patch is released, and a mobile user doesn't bring in their laptop in a timely fashion, the laptop's MAC address can be removed from the list of permitted addresses. This will force the user to bring their system in for maintenance and protect other devices in the process.

Network Access Control via Access Control Lists

Some routers and Layer 3 switches can be configured to permit or deny access based on network source and destination and on protocol, functioning as a packet filter firewall. This could be used to set up network segments or VLANs for limited access by untrustworthy devices.⁸

For example, some laptop users might require access to an intranet on ports 80 and 443, and email on ports 110 and 25. Providing a designated area for these laptops, and permitting only those ports to the specific addresses of the email and intranet server could provide a level of protection against NetBIOS based worms.

Advantages

- Low capital cost if current equipment will support configuration
- Additional security compared to disconnected drops
- Can provide some protection from a compromised laptop

Disadvantages

- Devices in the same segment or VLAN aren't protected from one another.

⁸ Odoms et al., pg 410

- Network addresses can be spoofed.
- Does not address the problem of unpatched laptops completely.

Conclusion

If there are mobile users that need limited access to network resources, this could be a good way to improve the safety of network resources. For example, if a group of mobile users connect to the Internet for access to an Intranet server, the following configuration could be used. This configuration would allow the laptops access to the Intranet on ports 80 and 443 only, while protecting other network assets

IP addresses

Laptops will need either static addresses or DHCP services. If a DHCP server were used, it would be best to place it on the far side of the router and configure it with reservations for the mobile users. Additional rules would have to be added to router's access list to permit the DHCP negotiation, and the DHCP Helper would have to be activated on the router to propagate DHCP requests.

Either static or reserved IP addresses could be exploited if someone were so inclined, but they would only have the access permitted to a mobile user.

In this specific configuration, requiring authentication for access to the Intranet server would further tighten security.

802.1x Authentication

802.1x, while usually associated with wireless applications, can be used to improve security in wired networks. 802.1x provides a mechanism for authenticating a device and permitting or blocking access at a port level. A "supplicant", a device with an 802.1x client, sends an EAP request to the "authenticator", a switch, router, RAS server or other device that services the request. The authenticator passes the request to a RADIUS server, which sends a challenge back to the supplicant. If the supplicant sends the appropriate response, it is granted access to the network.⁹

⁹ Kelley, "The X Factor", pg 61

The supplicant can be authenticated by PAP, CHAP, MS-CHAP or certificates depending on the specific implementation.

802.1x provides robust control of network access, but still doesn't have the ability to address the trustworthiness of a configuration.

Advantages

- Low capital cost if current equipment will support configuration. This assumes relatively new infrastructure, the presence of a RADIUS server, and possibly additional infrastructure, such as a certificate server.
- Highest level of network access control
- Good protection from unknown laptops
- Administrators can block corporate laptops if their patch status is uncertain

Disadvantages

- Complicated implementation
- Cost of infrastructure upgrades – switches or routers that can function as authenticators, RADIUS server, PKI infrastructure if desired, etc.
- Failure of the RADIUS server can lock out large groups of users
- New standard. Support and compatibility may be a problem
- Does not address the problem of unpatched laptops directly
- Visitors with laptops may not be able to connect at all, or may have to be configured with an 802.1x client. Administrators may find it easier to provide a workstation to visitors

Conclusion

802.1x authentication will improve security. Only users with an 802.1x client configured for the correct authentication can gain access. The primary limiting factors are likely to be time and money. Visitors who wish to connect to your network are apt to be inconvenienced.

A non-strategy – disabling computer accounts in Active Directory

Disabling computer accounts in Active Directory cannot be considered a valid strategy for protecting the network from untrustworthy laptops, since a compromised laptop doesn't need AD authentication to propagate a worm.

Disabling computer accounts is useful for encouraging users to bring in a laptop to be patched.

Controlling the laptop configuration

Network access control provides one time patch control. Once a laptop's patch level is validated, the laptop has access until an administrator revokes access.

Suppose an administrator examines a visitor's laptop and grants access, but doesn't revoke it. Six months later, the visitor returns and plugs directly into the network without checking in with IT. During those six months, their laptop hasn't been patched, and they've picked up the latest worm. The worm is released and pandemonium ensues.

The same scenario could play out with a corporate laptop that isn't very carefully managed.

Manual Patching

At this time, we depend on a manual system for patch control. We use a locked down configuration that precludes users from altering the contents of system directories. This can slow down a worm or virus that operates in the users security context, but means that the user cannot apply operating system patches. Therefore, IT has to patch all laptops.

This approach gives us the chance to test patches before they are applied, and to confirm that laptops have been patched, but means more work.

Other organizations permit users administrative access to their operating systems, and rely on users to connect to patch sites.¹⁰

This is explained in greater detail below.

¹⁰ Eustice et al., pg 3

In order to have any success with a manual system, you need:

- A VERY accurate inventory of laptops
- A policy requiring that users co-operate with patching
- A co-operative end user community.

Advantages

- Low *capital* cost
- Little development required
- Patches can be tested prior to application.

Disadvantages

- High labor cost
- Inconvenient for end users
- Slow deployment of patches
- Possibility of missing machines, laptops remain somewhat untrustworthy

Conclusion

Manual patching is imperfect, but unavoidable.

Web based updates

Some organization use Windows Update, Ximian's Linux update site, anti-virus FTP sites and other web based update services to keep machines patched. This is a low cost approach, but administrators cannot be sure that updates have taken place and some sites may require administrative access to the operating system to perform updates.

Having users update their own machines can have two complications. A very diligent user can download and apply a patch before it has been tested, possibly doing more harm than good, and a lax user can delay the application of a critical patch until it is too late.

Advantages

- Low capital and labor costs

Disadvantages

- Reliance on end users, who may not apply patches in a timely manner
- End users may apply patches before testing.
- End users must be able to update OS, reduced security as a result.
- Possibility of missing machines, laptops remain somewhat untrustworthy

Conclusion

This approach is best suited for organizations where users are highly autonomous and IT resources are limited.

Patch Distribution Software

Microsoft has addressed the problem of patch distribution with the System Update Service. SUS provides a central repository of Microsoft patches with client software to retrieve and apply them. The SUS client can be configured to provide administrative credentials, so the users don't need to have administrative access to their machines.¹¹

Microsoft's Systems Management Server can be used to distribute patches, but laptops present the same problem to both SUS and SMS: if the laptop isn't connected, it can't be patched.

If SMS has a pending patch installation for a laptop, it will attempt to install it. The user can inadvertently interrupt this process by disconnecting the laptop. SMS will then have to start from scratch the next time the laptop is attached. We have to instruct our users to bring in laptops and leave the connected but logged off until SMS can install the patch.

Microsoft has recognized this problem, and their latest version of SMS contains background transfer technology that is more forgiving of slow or interrupted connections.¹² It should be noted that other products, such as CA Unicenter, have already handled this problem.¹³

Advantages

- Automation of patching process results in less labor and more consistency
- Some systems, such as SMS will record success or failure of patching process

Disadvantages

- Laptops' unpredictable connectivity complicates patching process
- Patches must be packaged and tested

Conclusion

Software based patch delivery will make patching laptops a little easier. It is worthwhile if the infrastructure is already in place for patching desktops.

Policy Based Access Control

Several companies have developed policy based configuration confirmation. An agent running on a laptop compares the laptop's configuration to a policy and permits the laptop to connect if the policy criteria are met.

¹¹ Smith, URL: <http://techrepublic.com.com/5100-6268-1050973.html>

¹² Sturdevant, URL: http://www.eweek.com/print_article/0,3048,a=110206,00.asp

¹³ Sturdevant, URL: <http://www.eweek.com/article2/0,4149,1134220,00.asp>

Depending on the specific product, this policy can specify OS version and patch level, anti virus software versions, the presence or absence of a specific process, or other criteria.

Although many of these products are intended to be used with a VPN, they could be used in an internal network to isolate incorrectly configured laptops. The discussion of Checkpoint's Secure Client serves to illustrate this. Other VPN products, such as Confidence Online¹⁴, or Zone Labs Integrity¹⁵ could be used in the same general way.

Checkpoint Secure Configuration Verification

Checkpoint's Secure Client, the VPN client that provides connectivity with Firewall-1, provides authentication, encryption and a centrally configurable personal firewall.¹⁶

At startup, Secure Client has a default firewall policy of Any↔Any Drop. Once the laptop is authenticated, it fetches the current policy from the Policy Server. This means that all Secure Client laptops are running a centrally configurable firewall.

As part of Secure Client, Checkpoint provides the Secure Client Verification function. The firewall administrator creates a policy to check for patches, processes, registry keys, or custom verifications based on the execution of a batch file or executable. If criteria aren't met, the laptop cannot connect to any network resource.

Secure Client can be used to block an untrustworthy laptop from local access, but not perfectly. Secure Client permits a laptop to obtain a DHCP lease, register with WINS, and communicate with a timeserver before the default policy is in place. Also, Secure Client will permit communication with the Policy Server for authentication and policy update. It is conceivable that a compromised laptop could release a worm before Secure Client is in control.

The network trace below shows the communication between a Secure Client laptop and a local network before Secure Client is in full control of the network interface.

Software based patch delivery will make patching laptops a little easier. It is worthwhile if the infrastructure is already in place for patching desktops.

¹⁴ URL: <http://www.nwfusion.com/reviews/2003/0811prodpeek.html>

¹⁵ URL: <http://www.zonelabs.com/store/content/company/corpsales/intOverview.jsp>

¹⁶ URL: http://www.checkpoint.com/products/connect/vpn-1_clients_scv.html

Network Trace of a Secure Client Laptop Starting Up

1	[0.0.0.0]	225.255.255.255	CHCP Request, Source type: DHCP Broadcast
2	[0.0.0.0]	225.255.255.255	CHCP Request, Source type: DHCP Broadcast
3	[0.0.0.0]	225.255.255.255	CHCP Request, Source type: DHCP Broadcast
4	[0.0.0.0]	225.255.255.255	CHCP Request, Source type: DHCP Broadcast
5	[0.0.0.0]	225.255.255.255	CHCP Request, Source type: DHCP Broadcast
6	Destination: Broadcast		ARP C FA [0.0.0.0]
7	[0.0.0.0]	0.0.0.0	CHCP Request, Source type: DHCP Broadcast
8	Destination: Broadcast		ARP C FA [0.0.0.0]
9	Destination: Broadcast		ARP C FA [0.0.0.0]
10	Destination: Broadcast		ARP C FA [0.0.0.0]
11	Destination: Broadcast		ARP C FA [0.0.0.0]
12	Destination: Broadcast		ARP C FA [0.0.0.0]
13	Destination: Broadcast		ARP C FA [0.0.0.0]
14	Destination: Broadcast		ARP C FA [0.0.0.0]
15	Destination: Broadcast		ARP C FA [0.0.0.0]
16	Destination: Broadcast		ARP C FA [0.0.0.0]
17	Destination: Broadcast		ARP C FA [0.0.0.0]
18	Destination: Broadcast		ARP C FA [0.0.0.0]
19	Destination: Broadcast		ARP C FA [0.0.0.0]
20	Destination: Broadcast		ARP C FA [0.0.0.0]
21	Destination: Broadcast		ARP C FA [0.0.0.0]
22	Destination: Broadcast		ARP C FA [0.0.0.0]
23	Destination: Broadcast		ARP C FA [0.0.0.0]
24	Destination: Broadcast		ARP C FA [0.0.0.0]
25	Destination: Broadcast		ARP C FA [0.0.0.0]
26	Destination: Broadcast		ARP C FA [0.0.0.0]
27	Destination: Broadcast		ARP C FA [0.0.0.0]
28	Destination: Broadcast		ARP C FA [0.0.0.0]
29	Destination: Broadcast		ARP C FA [0.0.0.0]
30	Destination: Broadcast		ARP C FA [0.0.0.0]
31	Destination: Broadcast		ARP C FA [0.0.0.0]
32	Destination: Broadcast		ARP C FA [0.0.0.0]

Secure Client presents one additional limitation to connecting inside a network. The user logs onto the laptop after Secure Client has implemented a policy of Any↔Any Drop and before authentication and communication with a policy server. Therefore, the only way to use a domain or Active Directory account is to permit cached logon credentials. It should be noted that SANS does recommend cached logon credentials for laptops.¹⁷

Advantages

- Central management of personal firewall
- Central management of policy including minimum OS version and patch level, anti-virus version applications present and applications absent
- Access to network resources if and only if criteria are met.

Disadvantages

- Per laptop cost.
- Only laptops with Checkpoint software are managed. A device WITHOUT the Secure Client software can still connect and cause problems.

Conclusion

¹⁷ Shawgo, pg. 27

Worthwhile if a Checkpoint VPN is already in use. The ability to configure a policy centrally is a real plus. For example, if a worm is using TFTP to spread, the policy could be configured to deny access to any device running a TFTP server.

Confidence Online

Wholesecure's Confidence Online, provides a firewall independent way to check the integrity of laptops. A fat client, Java Plug-in or Active-X plug in checks the laptop to see if it is compromised by a worm, spyware or other Trojan. The product is intended for additional security for web-based email and other SSL applications.

Confidence Online could be used to set up a quarantine area that would provide limited access to network resources for laptops that were determined to be clean.

Controlling network access based on laptop configuration

If it were possible to combine these two solutions, controlling access to the network and confirming patch status, a network could be configured to deny connectivity to all untrustworthy laptops. We are tantalizingly close to this goal.

Security to the port level

In June of 2003, Information Security Magazine published responses to an RFP for network security to the port level. Respondents provided strategies to provide port level authentication and security.¹⁸

While this stops short of controlling network access based on a laptop's patch level, integrating authentication and firewall functionality could go a long way toward protecting the network from worms, viruses and malware.

Vendors were asked to secure a main campus of 3000 users, smaller branch offices, and provide secure wireless services for laptops and integrated management. 3Com, Cisco, HP, F5 Networks and Netscreen Technologies responded. Below is a very brief summary of each approach.

3Com

3Com relied on firewall technology embedded in the newest 3Com NICs. Servers could be protected and, if the budget permitted, every workstation

¹⁸ Snyder, Joel, pg. 28

and laptop would be behind a firewall-on-a-NIC. Authentication was performed by 802.1x and RADIUS, but the two functions remained separate.

F5

F5 proposed using switches that enforce authentication and encryption. This approach could be used to block unknown laptops, but not untrustworthy ones.

Netscreen

Netscreen proposed placing firewalls between VLANs. Administrators would have to configure the VLANs manually. Authentication would be provided by RADIUS and a web based client, so only authenticated users could cross into the servers' VLAN.

HP

Hewlett Packard combined authentication and policy-based security. A device authenticates using 802.1x and RADIUS. Based on that identity, the RADIUS server pushes an access list down to the switch port and places that port in the appropriate VLAN. The VLANs are connected to a firewall, which protects servers and enforces policy.

Cisco

Cisco's functionality is largely similar to HP's, except that they did not propose to place packet filters at each switch port.

Conclusion

None of the proposals fully answered the RFP's requirements, although some came close. The proposed technologies are certainly worthy of consideration.

Zone Labs Integrity

Zone Labs, well known for its Zone Alarm firewall, developed Integrity, a policy enforcement product. The Integrity client fetches a policy from the Integrity server, and allows the client device to connect to the network if the policy is met. The policy can specify OS patch levels, check to see if anti-virus is functioning and current and confirm the presences of applications.

Integrity is intended for mobile systems when connected to the corporate network and when functioning as standalone devices. Authentication can be provided by LDAP (including Active Directory), RADIUS or NT Domain. Integrity has just added 802.1x functionality, suggesting that network access can be tied to adherence to policy and to patch level.

On December 15th, 2003 Checkpoint and Zone Labs announced their intention to merge. In the future, we may see an integration of Secure Client and Integrity, providing both VPN and improved policy enforcement.¹⁹

Cisco's Network Admission Control Initiative

Cisco is working with a group of visitors to develop Network Admission Control (NAC).²⁰ Endpoint devices - laptops, PDAs and PCs – would run “trust agents,” clients that could confirm the configuration of anti-virus software, and update anti-virus DAT file. Several Vendors have announced an intention to develop trust agent software.

The NAC infrastructure also includes a policy server, which could be used to specify minimum acceptable patch levels.

Finally, Cisco switches and routers would provide the enforcement. Based on policy, a non-compliant device could be shut out, quarantined or provided minimal access to network resources.

Cisco routers will be able to implement NAC in the first half of 2004. Switches will follow at a later time.

QED – Quarantine, Evaluation and Decontamination

Eustice et al.²¹ describe a security model for wireless clients that involves an initial connection of an untrustworthy device to a wireless network with limited access (quarantine), evaluation of the device for integrity (evaluation) and, if necessary, remediation (decontamination). Like 802.1x, the concepts could be made to work in a wired environment.

This paper is interesting because of the QED model's reliance on open source products.

Quarantine phase

A wireless device is configured with a QED client, including a certificate. The client connects to an 802.11 gateway, which functions as a DNS, DHCP and router for the client. Along with the IP address, the local client receives an IPPackets rule to DENY all traffic from any source except the gateway. The gateway uses the certificate to verify the identity of the client.

¹⁹ Greene, URL <http://www.nwfusion.com/net.worker/news/2003/1215checkzone.html>

²⁰ http://www.cisco.com/warp/public/cc/so/neso/sqso/adcon_wp.pdf

²¹ Eustice et al. URL: <http://lasr.cs.ucla.edu/reiher/papers/qed.pdf>

Examination phase

The client is then scanned with NMAP for anomalous activity. The client software would also present information regarding the functionality of the anti-virus software. The authors recognize this as a weak link, since a compromised machine might “lie” about its state.

Decontamination phase

Decontamination, if possible, would be achieved with the RedHat Package Manager. The user would be prompted to accept packages for updates to anti-virus software or OS fixes.

Our Chosen Course of Action

Our organization is a community bank. We have about 550 employees, with 350 at a main location and the rest in branches within a fifty-mile radius. Our fourteen mortgage representatives use laptops and a Secure Client based VPN to support a loan origination application and to get email over Outlook Web Assistant. Our investment specialists and commercial loan officers use laptops as mobile network clients, but only connect to the Internet through our infrastructure. (We have disabled their modems, but the possibility still exists for a laptop to go home and be hooked up to a cable modem or DSL, if the end user can figure a way to configure PPPoE without administrative rights.)

We have another twenty or so laptops with the ability to connect to the Internet via modem or to our local network. Some of these use the VPN and Checkpoint's Secure Client; others are configured with Zone Alarm or other personal firewalls.

All laptops are at a minimum of NT4.0; most are at W2K or XP. Users do not have administrative access to their machines. Users do not have write access to SYSTEM32, or critical parts of the registry.

During the last year, we have replaced older 3Com and Synoptics network gear with new Cisco switches and routers.

Our Goals and Constraints

No additional capital expenditures

The decrease in interest rates has been hard on financial institutions, and ours is no exception. Aside from a Citrix server, which we intend to make available via VPN, no additional capital are likely to be approved.

Better control over laptop configuration

At this time, we rely on our users to bring laptops in for updates and, in some cases, rely on our users to bring the laptops to us before connecting to the network. We also rely on our departments to bring third parties with laptops to us, rather than finding a live drop, hooking into the network and attempting to connect to resources.

Regulatory requirements

GLBA requires very tight controls on customer data. If a laptop is stolen it is much better for us and for our customers if there is no customer data on the hard drive, particularly now that Passware has announced way to crack EFS.²² Our standards now call for the SAM to be password protected on laptops, as SANS recommends.²³ Lost Password says that their product needs access to the SAM to break EFS; password protection of the SAM may be enough to stop Passware, although we have not tested this yet.

Our next steps

Controlling Network Access – Without Spending Money

Enforcing a Policy

We have enough Secure Client licenses to cover all our laptops. Once they are all converted we will have a way to control their access to our network based on policy.

Since we have already budgeted for the Citrix Server, we will use this to limit laptop interaction with the local network and remove data from laptops. Secure Client will permit laptops to connect to Citrix via the local network, but nothing else. Laptops will be able to access the Citrix server if their laptop complies with the Secure Client policy.

Customer data remains on application servers, where it is backed up regularly and where it is more secure.

Network Access Control

We do disconnect idle network drops, both for security and to conserve switch ports. However, it is easy to miss the move of a PC and leave a network drop live and unused.

²² URL: <http://www.lostpassword.com/news/pnl31.htm#efs>

²³ Shwago, pg. 16

There has been a fair amount of resistance internally to implementing port level security. Our desktop support personnel are concerned about the complications that it will present, when compared to an incremental increase in security. Implementation of 802.1x may prove more palatable, and we intend to prototype a configuration. We already have a RADIUS server, and our infrastructure will support 802.1x

We hope our security training (which is also already budgeted) will lead to better control of our network perimeter. Once our end users understand the importance for involving us if a visitor wants to connect to the network, we will have decreased the risk from visitors' laptops.

We also believe that a clearly articulated policy requiring that all visitors check their laptops with us before attempting to connect will help, since bankers are true believers when it comes to policy.

Conclusion

After a year of fast breaking worms and companies laid low by laptops, network security vendors see an opportunity to help. Unfortunately, many of the most promising solutions are either brand new or not yet available.

We can, however, improve our security by taking full advantage of the solutions immediately available to us.

References

Poulsen, Kevin, "Nachi Worm Infected Diebold ATMs", SecurityFocus, 2003
URL: <http://www.securityfocus.com/news/7517>

Microsoft TechNet
"Buffer Overrun in RPC Interface Could Allow Code Execution"
URL: <http://www.microsoft.com/technet/security/bulletin/MS03-026.asp>

Microsoft TechNet
"Buffer Overrun in RPCSS Service Could Allow Code Execution"
URL: <http://www.microsoft.com/technet/security/bulletin/MS03-039.asp>

Microsoft Knowledge Base Article 827636
"How to Use the KB 824146 Scanning Tool to Identify Host Computers That Do Not Have the 823980 (MS03-026) and the 824146 (MS03-039) Security Patches Installed"
URL: <http://support.microsoft.com/support/misc/kblookup.asp?ID=827363>

Tuesday, Vince "Layered Defense Falls to Worm Attack", Computerworld, September 29th, 2003
URL: <http://www.computerworld.com/securitytopics/security/holes/story/0,10801,85369,00.html?f=x584>

Seam Odom and Hanson Nottingham, Cisco Switching Black Book, Coriolis, 2001

Kelley, Diana, "The X Factor", Information Security, Vol. 6, No. 8, August 2003

Eustice, Kevin et al. "Securing WiFi Nomads: The case for Quarantine, Examination and Decontamination", Laboratory for Advanced Systems Research, Computer Science Department, University of California
URL: <http://lasr.cs.ucla.edu/reiher/papers/qed.pdf>

Smith, Jeremy, "Take Control of Patch Deploys with Software Update Services", Tech Republic
URL: <http://techrepublic.com.com/5100-6268-1050973.html>

Sturdevant, Cameron, "SMS 2003 Goes From Desktop to Laptop", Enterprise News and Reviews, October 20, 2003
URL: http://www.eweek.com/print_article/0,3048,a=110206,00.asp

Sturdevant, Cameron, "Unicenter Tools Deliver" Enterprise News and Reviews, June 23, 2003

URL: <http://www.eweek.com/article2/0,4149,1134220,00.asp>

Shawgo, Jeff, Editor, Securing Windows 2000 Step by Step, v1.5, SANS Institute, July 1, 2001

URL: <http://www.sans.org>

Greene, Tim, "Checkpoint Buys Zone to Bolster Endpoint Security", Network World Fusion, December 15, 2003

URL: <http://www.nwfusion.com/net.worker/news/2003/1215checkzone.html>

Snyder, Joel, "Turning the Network Inside Out", Information Security Magazine, Vol. 6, No. 6, June 2003

© SANS Institute 2004, Author retains full rights