



Vaccination Against Worms and Viruses

Peter J. Welcher and Carole Warner Reece

Introduction

Many networks were severely affected by the recent computer worm/virus outbreaks. If you were afflicted, you have my sympathy.

Now that most staff have their life back, it may be appropriate to do "lessons learned", before the next wave hits. I'm planning on doing a couple of articles on this topic.

The topic rather naturally falls into **Before** and **After** categories. Before is Best Practices and things to do before you have a problem. After is the things you can do to facilitate repairs and mitigate nasty side-effects after an attack. I plan to cover this topic backwards, with the After part first, since it may be helpful to some, and there seem to be more new ideas floating around here. As far as the Before part, I'm a bit hesitant to claim sufficient expert (ego?) status to be putting out my own list of Best Practices. However, I do have some thoughts I haven't seen elsewhere, I've got lists of things from various sites, and I can certainly provide a set of links to other compendiums of Best Practices that you may find useful.

This article is going to be biased a little bit towards the network side. That's because the network side is where I feel most of my expertise lies. I'm pretty darn good on UNIX/Linux, and can hold my own with Windows -- but I cannot claim to be current with what system administrators are doing on the security front. The other justification for any network bias: too often viruses are seen as a systems problem, but there is a valuable role for the network team in helping fight the virus attacks as well.

New Factors

Sometimes events cause changes in how we think about things. (I'm very carefully **not** saying "paradigm shift".) The recent viruses were qualitatively a bit different from their predecessors in several ways:

- 1 Rate
- 1 Scale
- 1 Impact
- 1 Other

Rate

Concerning Rate, viruses are spreading faster and faster. This was discussed in the article *How to Own the Internet in Your Spare Time*, which can be found at the URL <http://www.icir.org/vern/papers/cdc-usenix-sec02/>. (Love that catchy title. Wish the predictions had taken longer to arrive.)

What we're seeing now is that the viruses hit so fast and hard that many networks were effectively down, sites couldn't download patches or fixes, etc. The other impact here is that perhaps the virus scanner signature file mechanism is getting overtaken by the bad guys. If all your computers are infected before the vendors have signature updates, or before your hourly/daily refresh of host signature files, you caught the virus.

The partial good news is that the vulnerabilities were known, so those who had kept up on patches weren't as badly hit. There are always home users and others who get missed in patching, which is why automation and tracking of who's

got what patches is so crucial. (Note to self: stop preaching to the choir!)

Cisco seems to be in the right place at the right time with the Okena acquisition. Their Host Intrusion Prevention System (HIPS), or CSA, doesn't use signatures and supposedly would have blocked these attacks. See also Cisco Security Agent, <http://www.cisco.com/en/US/products/sw/secursw/ps5057/index.html>. I like the idea of attempting to protect against the unknown, but there are obvious limits to how far that can effectively be done.

Rate leads to another thought, one that I've now heard from several sources: Is there a way to slow or contain the spread of the virus? The analogy is perhaps SARS and face masks.

Cisco PVLAN's (Private VLAN's) in switches can help keep one host from infecting another. They have the virtue of being unlikely to disrupt services, if implemented with reasonable care.

Another idea is what I'm calling **network lockdown**. In network lockdown, switch VACL's (VLAN Access Lists) allow end-user hosts to only talk to the server subnets (not too hard if you have server farms). You do need to be careful with this approach: what about networked printers or print servers? Other services some staff may need, such as local file servers? If printers and services are all local, then do ACL's (Access Lists) controlling inter-VLAN traffic help? I've noticed recently (see the previous DSNIFF article) that smaller VLAN's or subnets to help mitigate the impact of Layer 2 attacks, such as MAC flooding or ARP spoofing. I'm not sure any of the above are complete answers, but they may work in your environment.

What I've just said can also be thought of as at least having some internal firewalling, be it PIX or ACL or whatever. You may wish to put some extra protection in front of the server farm even. For quite a while we've been doing the hard exterior / soft interior firewalling model. The security experts have been saying that's not enough. One way to think about this: do you have to wear a badge in the building, or is showing it to the guard at the front door enough? Do you have security on just airport passengers, or do you also have some internal controls on airport employees?

Scale

Scale comes in because so many hosts were infected. Sites needed

- 1 Automatic scan or quarantine
- 1 Auto-detection of viruses
- 1 Automated cleanup

I break them out this way because some sites took a mostly host-oriented approach, whereas others combined network-centric and host-centric tools in an effective way.

Universities were being hit with the virus outbreak just as students were returning to campus. Some have been building toolsets that lend themselves to helping with this situation. Eric Gauthier was kind enough to type up a draft document and tell the NANOG list about it. It can be found at <http://www.roxanne.org/~eric/blaster.html>. He references a University of Connecticut writeup, which can be found at http://www.security.uconn.edu/uconn_response.html.

Solving a slightly different problem, University of Florida apparently has a traffic level scanning program called Icarus that does scans for peer file sharing programs like Kazaa, and turns down student links if usage is detected. See also http://www.mae.ufl.edu/sysinfo/uf_takes_action.htm.

The basic idea behind these is that many universities welcome back students with a registration page, perhaps with acknowledgement of the Acceptable Use Policy. (I have a son and daughter at University of Maryland College Park; they do this). That also gives the university a chance to run a vulnerability check on a computer before providing it with connectivity. Eric's article discusses using recorded MAC addresses, DHCP, and DNS to trap users in a quarantine / "jail" until they pass a scan.

I'm mentioning this since enterprises may not be aware of this approach. The enterprise equivalent is software that verifies the user virus scanner and/or personal firewall is running before allowing connection to the network. Zone Integrity, McAfee e-Policy Orchestrator go part-way here. Right now, vendors (including Cisco) are more focused on good hygiene before connecting via IPSec VPN. Reading about Zone Integrity, I see it has features resembling the network lockdown I mentioned above. Is that something network devices should be doing, or is it something you want your personal firewall software doing on each and every computer?

I keep hearing that most recent worm/virus outbreaks came from laptops that caught the virus at home and then

imported it to the enterprise network. At most sites, the firewall, email or web proxies, and any Enterprise virus tools did fine at keeping the viruses at bay, as far as infection coming in directly off the Internet. So what can we do about laptops coming in? Unless they are scanned before connecting (slow!), or the agent can provide assurance it wasn't turned off, do you really want that laptop connecting up?

As far as Auto-Detection of viruses, see below. There are some nice network-centric techniques that you may not have thought of. If you can figure out who has the virus (without slow scanning), then the systems folks can be more effective at virus removal.

Concerning automated cleanup, that's pretty firmly on the systems side of things. People seem to use some combination of PERL scripts, Web CGI scripts, Windows SMS, virus vendor tools, etc.

Impact

The recent worms spread so fast and to so many systems that they became a Distributed Denial of Service (DDoS) attack, in effect. Since external DDoS is also a concern these days, it's good to know how to mitigate a DDoS attack. This is somewhat related to the detection of infected hosts, so we'll go into this in more detail below as well.

Other

Staffing Levels

I've been noticing that staffing levels are a problem. If you're barely keeping your head above water, you don't have time for network or performance management, let alone reading IDS logs, etc. In which case, you're to some extent driving blind. An analogy comes to mind: how long will your car work without the engine oil sensor gauge? So why do you have one? If all your time is going into scrambling to keep up with new sites, swapping out gear, and resolving implementation issues, then how are you supposed to be finding time to use all that network management and security software? Yet that's what can alert you to problems as they begin to hit.

Lean times and budget crunch may explain some of this. I am beginning to see that as network devices proliferate, and as responsibilities and complexity grow, perhaps staffing has not kept pace. I've seen this phenomenon elsewhere, sometimes somebody notices and fixes this, sometimes people get stressed or the network gets more and more fragile until something breaks (or somebody quits).

I'm also hearing some folks say "we know we should have done X, but we haven't had time". Resolving your business operations after getting hit by a hacker, virus, or DDoS attack can certainly also consume a good bit of time. And you don't get the chance to schedule **those** activities!

I don't know the complete answer here. Politely making the problem visible to management? Obtain the time or temporary staff to get things to where you want them? If management won't lighten your burden, maybe you can pass work off to someone else. We'd certainly love for you to hire us to out-task implementation projects or offsite monitoring.

Total System Approach

Another thing I've noticed is that it may be time for systems folks and network folks to work more closely together, especially in shops where that's not normally the case. That's perhaps one conclusion to be drawn from the technical part of this article: the network devices (and design) can detect and mitigate ill effects and help control the spread of viruses. Systems administrators can use anti-viral scanners and so on, and are needed to fix the infected computers. It definitely works better if you tackle the problem from both sides. Yes, you can have personal firewall policies to control spread -- but it might be easier, less of a performance impact, and more effective to do policy enforcement on network devices. Doing so does not require the presence of cooperating or controlled agent software on each and every PC.

Technical Ideas

There were two forward references above, concerning how to auto-detect viruses, and how to control their spread. I propose to provide a higher level of detail on these techniques, which means this discussion will have to extend into next month's article, for space reasons.

As mentioned above, it is tempting to do some form of network lockdown, as mentioned above. You have to be careful about cutting users off from resources they need. The attraction of this approach, where applicable, is that the access lists may be fairly simple and they can be put in place in advance or at the first sign of virus infection.

The other approach is to use specific knowledge about the virus to limit its ill effects and control its spreading. This is less likely to disrupt network services. But it does require an interval of time after the virus starts to hit before the information about how the virus spreads is available. Yes, you can watch an infected host with a Sniffer (Ethereal, TCPDUMP) and make some intelligent guesses.

So to mitigate Blaster, Cisco is recommending an inbound ACL such as the following. We've added some entries recommended by other sites, such as Microsoft.

```
access-list 101 deny udp any any eq 135
access-list 101 deny tcp any any eq 135
access-list 101 deny udp any any eq 137
access-list 101 deny tcp any any eq 137
access-list 101 deny udp any any eq 138
access-list 101 deny tcp any any eq 138
access-list 101 deny udp any any eq 139
access-list 101 deny tcp any any eq 139
access-list 101 deny udp any any eq 445
access-list 101 deny tcp any any eq 445
access-list 101 deny tcp any any eq 593
access-list 101 permit ip any any
```

This controls access to Microsoft ports 135 and 139, and to another port used by Microsoft SMB, port 445. Microsoft is recommending blocking port 593 as well.

You can then use Class Based policing to drop all such packets. Create a QoS class:

```
class-map match-all dcom-rpc
  match access-group 101
```

Then create a policy:

```
policy-map drop-dcom-rpc
  class dcom-rpc
    police 8000 1000 1000 conform-action drop exceed-action drop
```

Note that this policy ends up dropping anything. This is a minor work around to the fact that you cannot put 0 in place of the 8000 above.

Apply this inbound to the desired interface:

```
interface fast 0/0
  service-policy in drop-dcom-rpc
```

If you prefer, you can implement this approach as an inbound ACL instead.

You will also want some outbound filtering on your firewall or edge router. The worm downloads mal-ware via TFTP. Do your users really need to be able to use TFTP to servers on the Internet? I don't think that's a good idea! Hence Cisco recommends adding to your rules something like:

```
access-list 102 deny udp any any eq 69
access-list 102 deny tcp any any eq 4444
access-list 102 deny udp any any eq 135
access-list 102 deny tcp any any eq 135
access-list 102 permit any any
!
interface <your interface here>
  ip access-group 102 out
```

The port 4444 is used for control by Blaster. Port 69 is TFTP. And you really shouldn't be allowing Microsoft ports to be open to the Internet. If you do, your network is probably a security event waiting to happen. There are far more secure ways to craft an Enterprise architecture and provide services.

Now if you're being diligent, you accumulate lists of troublesome ports like this from various worm or virus attacks and add to your ACL. Since many new attacks recycle old code, this reduces any impact of such new variants. Weeks after Blaster, our customers report they are still seeing scans on Microsoft ports.

Links

The following are good links about the systems side of various recent worm/virus attacks. They might be useful if you're looking to write rulesets such as the above. I realize you've probably found your way to some of these by now, but it's helpful to know where you can find information and what's likely to be out there for the next time around.

I have another good set of links for Cisco TAC recommendations, etc., but I'm saving those for the next article.

CERT® Advisory CA-2003-20 W32/Blaster worm	http://www.cert.org/advisories/CA-2003-20.html
CERT® Advisory CA-2003-19 Exploitation of Vulnerabilities in Microsoft RPC Interface	http://www.cert.org/advisories/CA-2003-19.html
Vulnerability Note VU#326746: Microsoft Windows RPC service vulnerable to denial of service	http://www.kb.cert.org/vuls/id/326746
Free on-line scan of your PC	http://housecall.trendmicro.com/housecall/start_corp.asp
Basic info about Blaster from Microsoft	http://www.microsoft.com/security/incident/blast.asp
Security team info from Microsoft	http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/virus/alerts/msblaster.asp
Scanning tool to find host computers lacking the 823980 Security Patch (MS03-026)	http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=C8F04C6C-B71B-4992-91F1-AAA785E709DA
Microsoft Security Bulletin MS03-026 and patches	http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS03-026.asp
Page to check regularly for Windows updates	http://v4.windowsupdate.microsoft.com/en/default.asp
Symantec Security Response on the worm	http://securityresponse.symantec.com/avcenter/venc/data/w32.blaster.worm.html
Symantec worm removal tool	http://securityresponse.symantec.com/avcenter/venc/data/w32.blaster.worm.removal.tool.html
Symantec Webcast discussing mitigation and remediation strategies	http://enterprisesecurity.symantec.com/content/webcastinfo.cfm?webcastid=63
McAfee brief info on worm	http://us.mcafee.com/virusInfo/default.asp?id=helpCenter&hcName=sobig
McAfee more detailed virus profile	http://us.mcafee.com/virusInfo/default.asp?id=description&virus_k=100561
McAfee tool for detecting/removing some viruses, including SoBig	http://vil.nai.com/vil/avertools.asp#stinger
NTBugtraq	http://www.ntbugtraq.com/default.asp?pid=36&sid=1
NTbugtraq posting on free RPC/DCOM vulnerability scanning tool from eEye Digital Security	http://www.ntbugtraq.com/default.asp?pid=36&sid=1&A2=ind0307&L=ntbugtraq&F=P&S=&P=7719

eEye Digital Security's RPC DCOM
Vulnerability Scanner and Worm Disinfection
Utility (some free, more for a fee)

<http://www.eeye.com/html/Research/Tools/RPCDCOM.html>

Conclusion

Next month we'll continue the technical discussion of how to detect and control worm/virus attacks. I'll include Cisco links in that article.

Please contact Chesapeake Netcraftsmen (<http://www.netcraftsmen.net>) if you'd like us to come do a security risk assessment, help with virus effects mitigation, network management or security management software (CiscoWorks!), etc.

Dr. Peter J. Welcher (CCIE #1773, CCSI #94014, CCIP) is a Senior Consultant with Chesapeake NetCraftsmen. NetCraftsmen is a high-end consulting firm and Cisco Premier Partner dedicated to quality consulting and knowledge transfer. NetCraftsmen has eight CCIE's, with expertise including large network high-availability routing/switching and design, VoIP, QoS, MPLS, IPsec VPN, wireless LAN and bridging, network management, security, IP multicast, and other areas. See <http://www.netcraftsmen.net> for more information about NetCraftsmen. Pete's links start at <http://www.netcraftsmen.net/welcher>. New articles will be posted under the Articles link. Questions, suggestions for articles, etc. can be sent to [pjw <at> netcraftsmen <dot> net](mailto:pjw@netcraftsmen.net).

Carole Warner Reece (CCIE #5168) is a Senior Consultant with Chesapeake NetCraftsmen with over 15 years of experience in internetworking. Her prior assignments range from technical to sales and marketing and managing technical staff. Carole recently designed challenging new labs for version 5 of the Cisco CIT Troubleshooting course. She previously developed and worked with other authors to produce networking skills based labs at a wide variety of levels.

10/7/2003

Copyright (C) 2003 Peter J. Welcher