Global Information Assurance Certification Paper

Internet Service Provider Insecurity

Or

You Are the Weakest Link… Hello!

James Joyce
GSEC Practical Version 1.4b
Option 2

Abstract

This practical covers issues pertaining to the lack of security often encountered when conducting vulnerability assessment/penetration testing, especially as it pertains to Internet Service Providers (ISP). Specifically, a client site can be solidly locked down, but security weaknesses at the ISP can make the tightest site security a moot point. Moreover, in the last section the argument is made, in the absence of relevant legislation or regulation, for the formation of Secure ISPs. In that vein, the case is presented for holding ISPs responsible for damages incurred from malicious exploitation of said ISP, their client sites, and/or other interconnected sites as a result of the failure of the ISP to implement at least a minimum of industry standard security best practices. Moreover, any organization involved with United States Critical Infrastructure [1] should be required to use ISPs that adhere to security best practices, in the best interests of protecting not only themselves, but also our nation's way of life.

The subjects, or relevant businesses referred to in the technical sections of this paper are a financial institution and a tier-two ISP. The financial institution has multiple branch locations and supports in-bank as well online banking transactions. They host their own Internet transaction servers and they control their entire infrastructure, with the exception of their primary edge router and their Internet pipe. The ISP is (or rather was) a sizeable regional provider of Internet connectivity supplying service to small to large businesses, as well as providing redundant backup connections for the largest financial information network in the U.S. The penetration testing of the bank, proper, revealed that the network staff at the bank had put a lot of thought into securing their perimeter – as is often the case, the interior of the network was quite chewy (they have "complete trust" in all of their employees, but that's a topic for another practical), but the perimeter was solid. Fortunately, for the bank, their contract with the ISP contained a provision that allowed the bank, or a bank approved third party, to test the security of the provider. The ISP was wide open. Via simple techniques such as a DNS zone transfer, the judicious use of (or abuse of) SNMP, and social engineering, it was possible to take over the ISP within a matter of minutes, and to "own" all of the data entering or leaving the bank's network. The end result of this particular test was that the ISP turned its' customer base over (essentially sold the business) to their tier one provider and went into the computer consulting business, and the financial institution ended up with more secure service from a different provider.

Before Snapshot

The bank described in the above abstract needed to have a security assessment in order to help to fulfill Gramm Leach Bliley (GLB) [2] requirements for financial institutions. This was their first security assessment from an external third party. The Information Technology (IT) staff had been with the bank for several years, and had a good understanding of their systems, their environment, and security as it pertained to locking down their external perimeter through the use of, and appropriate configuration of, firewalls and routers. They had a good antivirus implementation, and they understood and utilized a well thought out set of Access Control Lists (ACL) throughout the enterprise. They even isolated internal from external email through the use of separate servers, external email being provided by the ISP, and the twain did not meet. To the best of their knowledge, aside from a couple of minor non-destructive computer viruses, they had never experienced a cyber security incident from either the Internet or from their network interior.

As is often the case, the inside of the network was considerably less secure. There were no Intrusion Detection Systems (IDS) in place, neither host nor network based. Core routers had default passwords (i.e. cisco), development NT systems had administrator accounts with blank passwords, and core switches had blank passwords. Security auditing was not enabled on the Windows servers, and RISC/6000 syslog files were never reviewed. User passwords were weak and included the password "password", passwords that matched the user ID, user's first name followed by a number, the user's birthday, the name of the city's professional football team, and, naturally, the blank password, to list a few. The passwords were validated through the use of L0ftcrack. In this instance, the program "PwDump3" was used to pull the password hash file (Security Accounts Manager – SAM database) from the Windows NT Primary Domain Controller. The SAM database was then fed into L0ftcrack for cracking. Examples follow:

```
DOMAIN      USERNAME    LANMAN PASSWORD   LESS THAN EIGHT   NTLM
PASSWORD    LANMAN HASH NTLM HASH   CHALLENGE   CRACK TIME  CRACK
METHOD
      Administrator     * missing *
      00000000000000000000000000000000
      D4A495112F81884CD45012FC0F18632D
.
.
      James             44444 x           44444
      C34DBA145F6B05AAAAD3B435B51404EE
      E7C4D9EDAEF615786E761330B5037636          0d 0h 0m 2s
      Hybrid
.
.
      Jennifer          122873      x     122873
      C5049EFDAFCE7F01AAD3B435B51404EE
      AE5E08BC100DA137551324C41C8A6DC9          0d 0h 17m 28s
      Brute Force
```

.
.

```
kimberley          PASSWORD          password
E52CAC67419A9A224A3B108F3FA6CB6D
8846F7EAEE8FB117AD06BDD830B7586C          0d 0h 0m 1s Dictionary
```

Given the strong network perimeter, it is understandable that the inside of their network had not been hacked from the Internet; however, given the security posture of their interior network, only divine intervention can account the good cyber-behavior of the several hundred some odd employees of the bank. To be sure, there were probably a few inappropriate or unauthorized access incidents, but there were never any indications to the IT staff or bank management of anything of the kind, either technically or interpersonally.

During Snapshot

Vulnerability Assessment (VA) on the interior of the network consisted of a typical combination of techniques including, but not limited to, the use of Nessus, SNMP discovery, Nmap, manual infrastructure hacking (i.e. Cisco HTTP exploits – e.g. http://192.168.x.x/level/16/exec) [3], L0ftcrack, brute force telnet, inappropriate SMTP authentications, etc., and everything was thrown at the firewall. Results of the testing were interpreted and a suggested remediation plan was conveyed to the client with respect to three areas: management controls, operational controls, and technical controls. Management controls were broken down into the following categories: change management, risk management, and security controls. Operational controls were evaluated according to these categories: Personnel security, Business continuity, Data integrity, Security awareness training, and Incident response capability. Similarly, Technical controls were divided as follows: Identification and authentication, Logical access controls, and Audit trails. The results of the VA were conclusive, and the steps presented in the client report established a framework that allowed the bank to solidify their internal security posture.

Penetration Testing (PT) from the outside of the network consisted of the gambit of typical techniques: scans of all types, brute force techniques, war dialing, DNS, SNMP, and SMTP exploit attempts, throwing everything at the firewall and edge router, etc. However, the details of this part of the PT are not essential to the gist of this practical. What is important is that in spite of all attempts, the exterior of the network was not penetrated from the Internet and no sensitive information was gleaned. This result was, from this engineer's perspective, quite frustrating, and represented the first network that I had not been able to crack during a PT. Naturally, the logical end to this frustration was to find another way to compromise the bank.

The next step was to see if it was, first, legal, and second, possible to compromise the ISP. Upon review of the bank's contract, it was determined that

it was permissible to test the security of the ISP. Going into this phase, it was anticipated that it would be possible to take advantage of something that the ISP might have overlooked, and to, at least, make a small dent in the bank's armor. [*And with dramatic flair*] Nothing could have prepared this unsuspecting penetration tester for what was about to be discovered next.

This is where it started to unravel very quickly. Step one was to traceroute to the bank's network (FYI – all IP addresses and domain names in this document have been sanitized to protect the privacy of the various respective organizations. As well, the ISP will be hereafter referred to as "sadISP" and the bank will be referred to as "sadBank".). Provided that the ISP was not blocking responses, the traceroute would allow me to identify IP addresses of systems within the ISP that were routing Internet traffic to and from the bank. To clarify the traceroute shown below, 192.168.195.235 is my (sanitized) IP address, the other 192.168.x.x and 10.117.x.x addresses represent (sanitized) IP addresses of systems in between my system and the ISP in question, 172.23.52.78 represents a router within the ISP, and A.B.C.D represents a target host ("www.sadBank.com" – [note: this is not a real link]) at the bank. The command *tracert A.B.C.D* yielded the following results:

hop	name	time0	time1	time2
192.168.195.235		197 ms	201 ms	200 ms
192.168.195.252		197 ms	180 ms	181 ms
192.168.140.1		195 ms	190 ms	191 ms
192.168.144.253		183 ms	191 ms	191 ms
192.168.152.150		182 ms	191 ms	190 ms
192.168.139.195		189 ms	191 ms	190 ms
192.168.136.10		189 ms	191 ms	190 ms
10.117.242.197		188 ms	180 ms	191 ms
10.117.243.21		187 ms	190 ms	191 ms
10.117.243.193		217 ms	210 ms	221 ms
10.117.240.210		216 ms	210 ms	221 ms
10.117.240.245		215 ms	221 ms	220 ms
10.117.240.153		224 ms	231 ms	230 ms
10.117.14.231		212 ms	221 ms	231 ms
172.23.52.78		230 ms	231 ms	241 ms
A.B.C.D		250 ms	241 ms	240 ms

With an IP address within the ISP's internal address space, reconnaissance on the ISP continued with an attempt to discover their primary DNS server. The following *nslookup* commands revealed that "sadISP" was resolving the IP addresses for "sadBank". Note: explanatory comments have been inserted to the right of the commands and responses.

```
C:\>nslookup
```

```
Default Server:  dns.proxy.xyzISP.com   - xyzISP.com is fictional
Address:  192.168.195.134               - my DNS server

> www.sadBank.com
Server:  dns.proxy.xyzISP.com
Address:  192.168.195.134               - my DNS server

Non-authoritative answer:
Name:    sadISP.com
Address:  A.B.C.D                       - survey says…

> A.B.C.D                               - reverse check for good measure
Server:  dns.proxy.xyzISP.com
Address:  192.168.195.134

Name:    www.sadBank.com                - and it checks out
Address:  A.B.C.D
```

Many organizations often follow a common domain naming convention for their core systems. In the case of DNS servers, these systems are very often named ns1.*domain-name*.com, ns2.*domain-name*.com, etc. This ISP was no exception and an educated guess revealed:

```
> ns1.sadISP.com
Server:  dns.proxy.xyzISP.com
Address:  192.168.195.134               - my DNS server

Non-authoritative answer:
Name:    ns1.sadISP.com
Address:  A.B.C.X                       - sadISP's DNS server
```

The next step was to attempt to switch my primary DNS resolver over to sadISP's primary DNS server. From an *nslookup* prompt, the following command succeeded in performing this task:

```
> server A.B.C.X
Default Server:  ns1.sadISP.com
Address:  A.B.C.X          - I am now resolving with sadISP's DNS server
```

Then ensure that it works:

```
> www.sadbank.com
Server:  ns1.sadISP.com    - shows that ns1.sadISP.com is the resolver
Address:  A.B.C.D          - yields correct results
```

Given that I could successfully use the ISP's DNS server to directly resolve IP addresses, the next step was to see if it was possible to perform a DNS zone transfer of the ISP's internal IP addresses. If successful, the DNS zone transfer would reveal sensitive information about the internal hosts and structure of the ISP itself. The following shows the results (partial and sanitized) of the zone transfer attempt, which was initiated with the command "ls –d *domain-name*". Note: system names, IP addresses, and domain names have been sanitized.

```
> ls -d sadISP.com
[ns1.sadISP.com]
 sadISP.com.                SOA        ns1.sadISP.com
webmaster.sadISP.com
. (2001100101 3600 3600 604800 86400)
 sadISP.com.                NS         ns1.sadISP.com
 sadISP.com.                NS         ns1.fictionalTier1ISP.net
 sadISP.com.                A          192.168.16.7
 sadISP.com.                MX   10    mail2.sadISP.com
 ns3                        A          172.16.1.7
 pri                        A          192.168.16.1
 mail2                      A          192.168.16.2
 ssssssss                   A          192.168.16.147
 home3                      A          192.168.16.249
 news                       CNAME      news.fictionalNewsSvc.net
 stage                      A          192.168.16.247
 home4                      A          192.168.16.248
 aaaaaaaa                   A          192.168.16.6
 bbbbbbbb                   A          192.168.16.184
 firewall.sadISP.com        A          192.168.16.251
 signup                     A          192.168.16.8
 home                       A          192.168.16.3
 mail                       A          192.168.16.2
 www                        A          192.168.16.7
 pri2                       A          192.168.16.252
 ns1                        A          A.B.C.X
 dev                        A          192.168.16.250
 cccccccc                   A          192.168.16.37
 ns2                        A          172.16.14.201
 sadBank.com                A          A.B.C.D
 cust2.com                  A          A.B.C.E
 cust3.com                  A          A.B.C.F
.
 cust100.com                A          A.B.F.Z
.
 custN.com                  A          A.B.X.N
.
```

Presto… The above results represents significant sensitive information about the internal structure of the ISP, including host names and addresses of the systems that comprise the core of the ISP as well as IP addresses and names of the ISP's customer base. This DNS zone transfer not only yielded the information needed to attack the ISP and the bank, but also contained information about other companies that could be attacked through this ISP (i.e. cust2.com, cust3.com, etc.).

Following the logic that many organizations use a common domain naming convention, the host named "pri" at 192.168.16.1 suggested that it could possibly be the core (primary) WAN switch for the ISP. By running an Nmap scan at that host, the following was revealed:

```
Starting nmap V. 3.00 ( www.insecure.org/nmap )
Interesting ports on pri.sadISP.com (192.168.16.1):
(The 1598 ports scanned but not shown below are in state: closed)
Port       State       Service
23/tcp     open        telnet
25/tcp     open        smtp
79/tcp     open        finger
Remote operating system guess: Ascend/Lucent Max (HP,X000-Z000)
version 6.1.3 - 7.0.2+
Nmap run completed -- 1 IP address (1 host up) scanned in 66 seconds
```

Oh my, does that say Ascend/Lucent Max? It does. This indicates that the initial guess was correct. The core switch has been discovered. Surely, SNMP was not enabled on the device… Haven't their engineers read the SANS Top 20 [4]? Regardless, the next step was to attempt to see if the Ascend could be compromised via SNMP. Before testing to see if SNMP was enabled, I loaded the appropriate Ascend Enterprise Management Information Base (MIB) into an SNMP tool (I obtained the correct MIB from an engineer who worked at the manufacturer of said device) in order to ensure that the maximum amount of information would be pulled from the device (in an understandable format), and that I would be able to correctly write to the device if SNMP was, in fact, enabled – my assumption was that SNMP would be enabled; after all, everything else had worked on this ISP. I then implemented a series of SNMP inquiries to the host, pri.sadISP.com. Partial (sanitized) results of this scan follow:

```
General parameters
------------------
Host Name : 192.168.16.1
IP address : 192.168.16.1
DNS Host Name : <not in DNS>
Read Community : public
Write Community : private
SNMP timeout : 2000
```

SNMP retries : 3
System Description

Note that default community strings are in place ("public" and "private").

SysName : PRI.Sadbank
SysDescr : 'Ascend Max-HP T1/PRI S/N: ABCDEFG Software +7.0.26+'
SysContact : Sad Guy 555-555-1234
SysLocation : SadNOC, Anytown, USA
SysObjectID : enterprises.529.1.2.5
SysServices : 14
SysUpTime : 140:10:25:12
IfNumber : 179

Reachability parameters

FTP : not reachable
HTTP : not reachable
Netbios : not reachable
NNTP : not reachable
POP3 : not reachable
Print : not reachable
SMTP : OK
SNMP : OK
Tcp Echo : not reachable
Telnet : OK

Note above that Telnet is enabled and accessible from the Internet. This represents another security hole that could have been "brute force" compromised to directly reconfigure the switch. If SNMP had not been so accessible, I would have attempted to exploit this vulnerability.

admin	oper	type	MTU	descr.	speed	ip address	mask	phys	Vendor
up	up	33	0	Console 1	9600			""	
up	up	ds1	0	T1 Slot 1 Line 1	1544000				""
up	up	ds1	0	T1 Slot 1 Line 2	1544000				""
down	down	ds1	0	T1 Slot 2 Line 1	1544000				""
down	down	ds1	0	T1 Slot 2 Line 2	1544000				""
up	up	45	0	Serial WAN Slot 11 Port 1	0				""
up	up	ethernet-csmacd	1500	ie0	10000000				
	00C0XXXXXXXX	Ascend Communications ISDN bridges/routers							
down	down	other	1500	wan0		0		000000000000	
down	down	other	1500	wan1		0		000000000000	
up	up	ppp	1524	wan2		56000		000000000000	

```
.
.
.
up      up      ppp     1524    wan33           64000                   000000000000
.
.
.
up      up      ppp     1524    wan69           64000                   000000000000
.
.
.
up      up      ppp     1524    wan139          64000                   000000000000
.
.
.
```

The above is a partial listing of device interfaces, interface states (i.e. is the interface up or down), and other relevant information. But wait – there's more:

```
192.168.017.080     192.168.016.252             8       255.255.255.248
        indirect        local   9419695     .ccitt.nullOID
.
.
.
192.168.017.153     192.168.017.153             1       255.255.255.255
        direct  other   645468      .ccitt.nullOID
.
.
.
192.168.018.032     192.168.018.033             1       255.255.255.224
        indirect        other   2424374     .ccitt.nullOID
192.168.018.033     192.168.018.033             1       255.255.255.255
        direct  other   2424375     .ccitt.nullOID
.
.
.
192.168.018.184     192.168.018.184             1       255.255.255.255
        direct  other   7634    .ccitt.nullOID
```

The above represents a partial sanitized listing of the routes that were configured on the device, and is what was ultimately used to reroute traffic to compromise the ISP and the bank.

```
net address   name  phys address      Vendor
192.168.016.002           00A0XXXXXXXX    Intel (PRO100B cards)
```

192.168.016.003	0060YYYYYYYY	3Com
192.168.016.252	00C0ZZZZZZZZ	Ascend Communications ISDN bridges/routers
192.168.019.250	0040AAAAAAAA	Sonic Mac Ethernet interfaces
192.168.19.254	02BBBBBBBBBB	

The above represents sanitized IP addresses and sanitized MAC address mappings found on the device. This could potentially have been used in an exploit centered on ARP cache poisoning if that had been necessary or desired.

Within a couple of minutes, the configuration and all other SNMP information had been pulled from the Ascend Max. It filled up several hundred pages in MS Word, so I'll save a few trees in the interest of brevity. Note that the ISP's entire customer base could easily be mapped out from the SNMP scan results. In many instances, the SNMP trace also contained the names of individual customers and customer contacts. This information could easily have been used to socially engineer and technically compromise each of these organizations. After going through the SNMP information and validating that it was possible to remotely write to the device (i.e. modify the configuration of the routes from over the Internet) with an SNMP tool obtained from the Internet, it was time to move on to the next step.

Full exploitation of the ISP required both technical components as well as human components. Social engineering of this ISP resulted in the "Game Over" phase of this assessment. This began with a phone call to the ISP along the lines of:

> "Hi, this is Wile E. Coyote of Acme Products and I would like to purchase an 8-bank of IP address space for my company. How much will that cost?… Really?… When can you turn the service up?… Really?… What's the address space?… Thanks, and send the invoice to…."

Once the "Acme Products" address space was "active", it took less than five minutes to actually modify the Ascend Max's routes so that all data passing into, or out of, the bank passed through my new "Acme Products" network. Naturally, there was some preparation up front to ensure that the switch's configuration would be changed correctly the first time, and so that services would not be interrupted for other clients of the ISP. With the new routes in place, all Internet traffic flowing into or out of the bank first passed through my network.

(*Note: the technical description of the rerouting of Internet data at this ISP has been greatly simplified (i.e. modification of other devices, aside from the Ascend, was required), so that this practical can not be used by the malicious element as a step-by-step guide for hacking ISPs. Also, please be aware that this vulnerability is not specific to Ascend/Lucent products – any vendor products that*

support SNMP and have SNMP enabled in a like manner would have been equally vulnerable to this type of attack.)

With liberal use of sniffers (my sniffer of choice, in this instance, was made by Shomiti – it is a commercial product and worked perfectly. Note: almost any commercial or freeware sniffer could have been used for this part of the testing), all unencrypted traffic was captured, yielding expected results. As well, all encrypted traffic was also captured – while not immediately readable, given enough computer horsepower, it would have been possible to brute force the encryption keys to un-encrypt the bank's sensitive communications. Since the bank conducted on-line transactions, and since the servers that housed the on-line banking applications are on the bank's internal network, the encrypted traffic contained user IDs, passwords, account numbers, and transaction details. Cracking this encryption could easily have resulted in a catastrophic compromise of the bank, and a large increase in the amount of money in a malicious hacker's own bank account. Completion of this scenario would have resulted in a "Game Over" condition. SadISP – You are the weakest link… HELLO!

After Snapshot

Post game wrap-up – The ISP was completely unaware that they had been compromised. The ISP did not have technical personnel that could handle even the barest minimum of security precautions. It would have been equally as simple to reroute any of the ISP's client's networks without anyone being the wiser. When this situation was pointed out to the ISP, they realized that groups with malicious intent might very possibly have enacted this scenario in the past. That, in conjunction with the fact that they provided a redundant link into the nation's largest financial information network, prompted the ISP's owner to say that he had been considering getting out of the ISP business anyway. Within a week, the ISP was in negotiations with their provider to convert their customer base over to another provider. Additionally, the bank had no idea that they could be compromised in this manner, and they were quite dismayed. As a result, they were integral in convincing the ISP that they should either tool up for security immediately or get out of the business. The ISP realized that there could be potential civil litigation centered on their complete lack of any security precautions for their clients, especially those in the financial sector, and has since shut down the ISP component of their business – they are now a computer network consulting firm (they still have the same engineers on staff – buyer beware!).

The above scenario leads to a much larger issue: especially with respect to clients that are controlled by either GLB, Healthcare Information Portability and Accountability Act (HIPAA) requirements [5], or owners of United States Critical Infrastructure, in the absence of regulation or legislation for Internet Service Providers, there is a strong need for ISPs to establish a level of service that meets the security and privacy needs of their clients, and, furthermore,

contributes to the security of the United States. ISPs should be considered to be a part of the U.S. Critical Infrastructure and should secure themselves accordingly. The current thinking in Washington D.C., based upon my conversations and correspondences with relevant decision makers in the White House, the U.S. Congress, and numerous other government agencies, is that a voluntary compliance model is preferred over a regulatory model, as it pertains to getting "security buy-in" from the ISPs. Unfortunately, in my experience, the unsecured ISP example above is the rule and not the exception, and unless some significant incentives are presented to the ISP owners, I do not see this condition changing.

The Computer Fraud and Abuse Act, 18 U.S.C. §1030 does provide for the prosecution of perpetrators of certain computer crimes (i.e. unauthorized intrusion into computer systems whether by an outsider, an insider, or due to administrator/authorized user impropriety – actually seven types of activities are defined as being criminal)[6]; however, there is no current legislation that protects organizations against an ISP that is negligent with respect to securing their own systems/network links/etc. In the absence of this type of legislation, clients are, whether they know it or not, essentially at the mercy of their ISPs, with respect to the security of their Internet data traffic. In light of the current threat that this condition poses to our nation, my feeling is that the U.S. Congress should seriously, and immediately, address this issue. Additionally, this issue, again in my opinion, should not be limited to ISPs. Any and all organizations operating a computer network that are in any way associated with the U.S. Critical Infrastructure should be required to take measures to ensure, at least, a minimum level of compliance with industry best security practices. Furthermore, they should take measures to ensure that their ISPs also comply with the same standards (i.e. that they should function as, and be certified as, a "Secure ISP" [development of certification standards for a "Secure ISP" could be the subject of another SANS practical]) Personally, I believe that Section (a)(5) of the Code …

> *"prohibits anyone from knowingly causing the transmission of a program, information, code, or command, and as a result of such conduct, intentionally causes damage without authorization, to a protected computer"* [7]

…could be tested in court against a significant breach of a protected system that was the result of flagrant neglect with respect to security on the part of an ISP; however, this will remain to be seen, and I am not an attorney – just a wishful thinker. Currently, the analysis of the latest version of this section [Section (a)(5)] lists negligence on the part of authorized users as "no crime" [8]; however, in my opinion and in the interests of national security, this needs to be tested. The reasoning for this is that (giving the ISPs the benefit of the doubt that they did not previously realize that turning up an ISP without any security precautions could endanger the nation) due to the rash of destructive attacks against the Internet over the past three years, and the myriad of well publicized threats against our

infrastructure, ISP owners cannot possibly be under the same misconceptions (re: the aforementioned threat to the nation); therefore, complete negligence with respect to security on the part of an ISP should be considered criminal. While the need for legislative reform was argued by the U.S. Department of Justice [8], and this argument did lead to reforms which now protect businesses against rogue administrators/authorized users, this topic must be revisited, and fresh arguments need to be made in order to protect our infrastructure, since, in this case, the refusal to act (i.e. to willingly not implement security precautions for, specifically, reasons of cost containment) facilitates criminal activities against protected computers. The definition of a "protected computer" is as follows:

> the term "protected computer" means a computer
>
> > (A) exclusively for the use of a financial institution or the United States Government, or, in the case of a computer not exclusively for such use, used by or for a financial institution or the United States Government and the conduct constituting the offense affects that use by or for the financial institution or the Government; or
> >
> > (B) which is used in interstate or foreign commerce or communications, including a computer located outside the United States that is used in a manner that affects interstate or foreign commerce or communication of the United States; [9]

This statute should be revisited so that the above definition of a "protected computer" is expanded to include computers that are used by or for organizations that are associated with the U.S. Critical Infrastructure.

An alternative side to this is that if ISPs are not willing to improve their security for these types of clients, they should be required to assume the liability for security breaches that are a direct result of their lack of security. Legislation would not necessarily be required to bring this situation into effect. Larger business clients could demand that their ISPs assume this liability. Were this to happen, we would, more than likely, quickly see ISPs working towards improving their security posture. Unfortunately, one problem with the "transfer of risk (liability)" solution is that is does not take into account the incalculable cost associated the compromise of certain information contained on some computer networks within the U.S. Critical Infrastructure. For example, the loss of specific information pertaining to nuclear weapons could result in the destruction of our country – for losses such as this, there is no reasonable or realistic quantification of the risk; therefore, there is no way that an ISP could assume this liability. Consequently, the notion of transferring risk (i.e. the ISP assuming the liability for

exploits against the U.S. Critical Infrastructure) does not, in and of itself, make our nation secure. It simply allows ISPs to continue operating insecure networks, and permits them to keep "playing the odds", allows insurance companies to develop an additional source of revenue (via policy riders), and allows the legal profession to pursue new avenues of litigation (i.e. cases against ISPs AND their insurance companies).

These opinions may not be popular, especially throughout the ISP community; however, again with the current absence of regulation, there may be no other alternatives for security conscience organizations and companies to ensure that they truly are secure. Whether accomplished via regulatory legislation, through tax incentives, as a result of customer mandates, or through the sudden altruistic or enlightened realization that they are hurting the security of the United States (and thereby putting our nation, their income source, and, moreover, the very fabric of our society, at great risk), the security posture of our nation's ISPs MUST BE IMPROVED. With respect to the possibility of a newfound altruism growing among the owners and controllers of our nation's ISPs, I am not going to hold my breath. I have actually had success in converting a mid-sized ISP over to a more secured posture, and they are currently marketing the fact that they are a "secure" ISP; however, they are only one in a very large number of ISP. As well, their owners are retired military and extremely focused on the security of the U.S. - again, another exception to the rule.

In the very probable absence of the ISP owners suddenly beginning to care as much about the security of the U.S. as they do about money, it seems that either a regulatory model or an incentive-based model are the only options to building a more secure Internet. This is somewhat sad and reflects a serious lack of insight on the part of these business owners – they fail to realize that, if they do not shore up their defenses, it is only a matter of time before their infrastructures are compromised to the point that the Internet crashes, our power grids are brought down, ATMs fail, our financial sector is "e-raided", our emergency responders lose the ability to quickly respond to incidents, and, given our current and growing reliance on the Internet to conduct business, our economy grinds to a screeching and immediate halt. In the quest for immediate and larger profits, there does not seem to be a belief that our nation could be compromised in this manner. If I am wrong in the assumption that ISP owners are profit-centric to the point that they would sacrifice U.S. national security in order to realize an extra half of a percentage point in net revenues, someone please show me that I am wrong, step up to the plate, and promote "responsible capitalism" within the ISP space. On the other hand, if I am right, it may well require that the government step in to change an industry that holds the key to our security and future well being.

Of the possible options that the government could invoke, two seem to have the most potential. First, the government could legislate a required minimum-security standard within the industry along the lines of GLB and HIPAA. As

mentioned before, this option is the one that the government considers to be a last resort. The other option would be to offer tax incentives to ISPs that comply with a set of industry recognized best practices, such as those promoted by the SANS Institute and/or those being developed by other Critical Infrastructure Security-centric public/private sector organizations. In order to ensure maximum industry buy-in, however, rewards for compliance (i.e. tax incentives consistent with the security investment of the organization in question) should also be coupled with deterrents for those that do not comply (i.e. if an un-secured ISP is maliciously hacked and damage occurs as a result, said ISP should be fined twice the amount of money that it would have taken to secure the ISP [opportunity cost] PLUS the amount of loss realized by the security compromise – just a thought). Given this type of incentive, we might find that many more business owners discover a newfound sense of national pride and concern for U.S. security. This could help to realize the aforementioned "responsible capitalism" as opposed to "unbridled capitalism", and at this point in our nation's history, given the current climate of terrorist threats, we MUST come together (i.e. public AND private sectors) to ensure the security of our nation and our way of life.

Functionally speaking, the technical effort required to make an ISP more secure is not great and would not cost either an arm OR a leg. Implementation of ingress and egress filters to help to prohibit spoofing, limiting DNS zone transfers to only those authorized, locking down SNMP (either through non-default community strings or by turning SNMP off altogether), and eliminating the use of insecure protocols (i.e. Telnet), would go a long way towards securing the Internet as a whole.

An additional recommendation, if ISPs really wanted to help out, would be that they should implement IPv6 as soon as possible. While not as cost-transparent as other recommendations, the rollout of IPv6 would strongly improve Internet security. Unfortunately, in discussions with tier-1 ISPs, I have been told that they will roll out IPv6 when, and only when, their largest customers demand that they roll it out. In spite of the fact that virtually all of the infrastructure gear currently in use on the Internet fully supports IPv6, ISPs will not roll out the technology until, essentially, a significant number of Fortune 500 companies and/or the government mandate it. With all due respect to the almighty dollar, to me, this represents extreme myopia within the industry.

The above two paragraphs are not saying anything revolutionary. Those specific technical recommendations have been cited time and time again in numerous SANS publications, in many student practicals, and by security experts the world over. They are reiterated here in the hopes that perhaps some ISP owners, with a little technical savvy, are actually reading this and will finally take some action, share this with their peers, and help to secure our national infrastructure.

The leaders in this industry should take a look back on the generations of Americans that sacrificed more than a half-of-a-percentage-point (a totally unsubstantiated estimate of what it would take to implement stronger ISP security – but, hey, it's a start) to ensure the security of our nation. From the founding fathers to the veterans of WWII, those Americans did not, for the largest part, worry about how much money they were making at times of the greatest national need, and we are now at another one of those times in our nation's history. Instead, they knew that if they did not fight, our freedom was at risk… and they fought. While the logistics are different, the Internet is quickly becoming a worldwide battlefield just the same. Each ISP must be persuaded to defend its own piece of that battlefield, or we will lose more than just the war… we will lose our way of life.

To put it another way – don't let anyone say to you, "You are (your ISP is) the weakest link…goodbye."

References

1. U.S. Chamber of Commerce – Critical Infrastructure Assurance Office, White Paper – The Clinton Administration's Policy on Critical Infrastructure Protection: Presidential Decision Directive 63 May 1998, URL: http://www.ciao.gov/publicaffairs/pdd63.html
2. U.S. Federal Trade Commission, Gramm Leach Bliley Act, URL: http://www.ftc.gov/privacy/glbact
3. Cisco CERT Advisory, URL: http://www.cert.org/advisories/CA-2001-14.html
4. The SANS Institute, SANS Top 20, URL: http://www.sans.org/top20
5. U.S. Department of Health and Human Services, Healthcare Information Portability and Accountability Act, URL: http://www.hhs.gov/ocr/hipaa
6. U.S. Department of Justice - Computer Crime and Intellectual Property Section, Legislative Analysis of the 1996 National Information Infrastructure Protection Act, Electronic Information Policy & Legal Rep. 240, (1997)
7. The SANS Institute, SANS InfoSec Reading Room, URL: http://www.sans.org/rr/legal/act.php
8. U.S. Department of Justice, The National Information Infrastructure Protection Act of 1996, Legislative Analysis, URL: http://www.usdoj.gov/criminal/cybercrime/1030_anal.html
9. 18 U.S.C. § 1030, Fraud and Related Activity in Connection with Computers, URL: http://www.usdoj.gov/criminal/cybercrime/1030_new.html