Global Information Assurance Certification Paper

A Review of Cybersecurity Risk Factors

David F. Beck

In a nutshell, the Internet is all about connecting a client to a remote resource server and exchanging information. In the existing environment, covert sensors (sniffers) and forged network data exist that can comprise the confidentiality, integrity and availability of cyber systems,[1] including both the operation of the Internet as a system (the combination of networks and hosts) as well as the information contained thereon. Various threat models can be developed that describe the relationship between the threats and vulnerabilities that can result in a system compromise. One approach is to relate the important factors in what is known as a risk equation. One instantiation of such a model, shown in the figure below,[2] is useful because it is easy to see how the various factors relate to each other. For example, for a given threat and set of vulnerabilities and countermeasures, increasing the value of the target (the impact) will increase the chance for loss (risk). As another example, consider that for a given threat, target (impact), and system vulnerabilities, increasing the countermeasures will reduce the risk of loss.

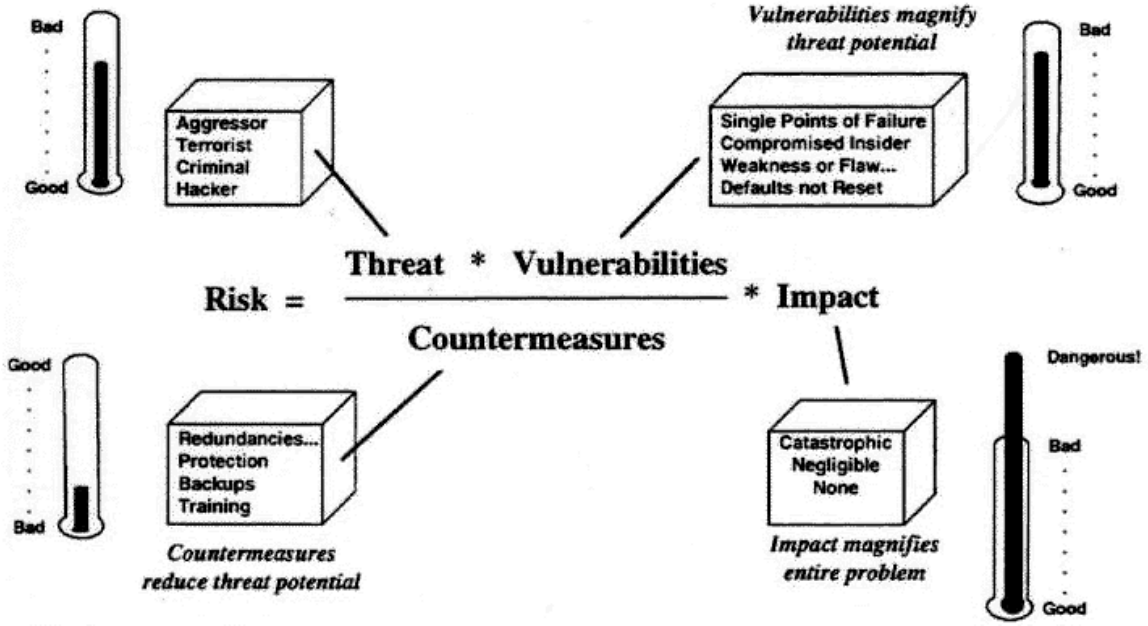

Another form used for the risk equation involves probabilities, and can be something as simple as R = Pa(1 – Pi)C. Here R is Risk, Pa is the probability of attack, Pi is a measure of system effectiveness, and C is the consequence of loss of the asset being considered. The actual form used depends on its use and on the form of the data that is available.[3] It is this last issue, what data related to risk assessment is readily available, that is reviewed in this paper.

---

[1] Sometimes called the "Three Bedrock Principles." E.g., Northcutt, Stephen, 2000, :"Core issues and challenges," SANS LevelOne Information Assurance Foundations, The SANS Institute, v 1.3, 6/28/00, p.3.

[2] Defense Science Board (DSB), 1996, Report of the DSB Task Force on Information Warfare (Defense), Office of the Secretary of Defense, Washington, 25 November. Electronic version found at http://cryptome.org/iwdmain.htm

[3] E.g., Carroll, John M., 1984, Managing Risk: A Computer-Aided Strategy, Butterworth Publishers, Boston.

Risk

Here we will jump to the end of the story first and ask just what are the risks of being connected to the Internet? If we represent the impact as some non-dimensional number with a value of one, risk, in essence, becomes a measure of the probability of that consequence occurring. One obvious data set that might be used to develop a measure of this probability is to analyze the incident data from the CERT[4] (although data from the first year, being a startup year, was not used below). Since the Internet is growing, both in terms of users, hosts, and, (presumably at this point) threats, the raw incident data must be normalized if a correct perspective is to be achieved, with the number of Internet hosts being the obvious choice due to the availability of reasonable data.[5] For the purposes of this paper, mid-year estimates were made for the number Internet hosts. The resulting, normalized CERT data is shown in the figure below, with the incident count being in terms of the number of reported incidents per 10,000 Internet hosts.

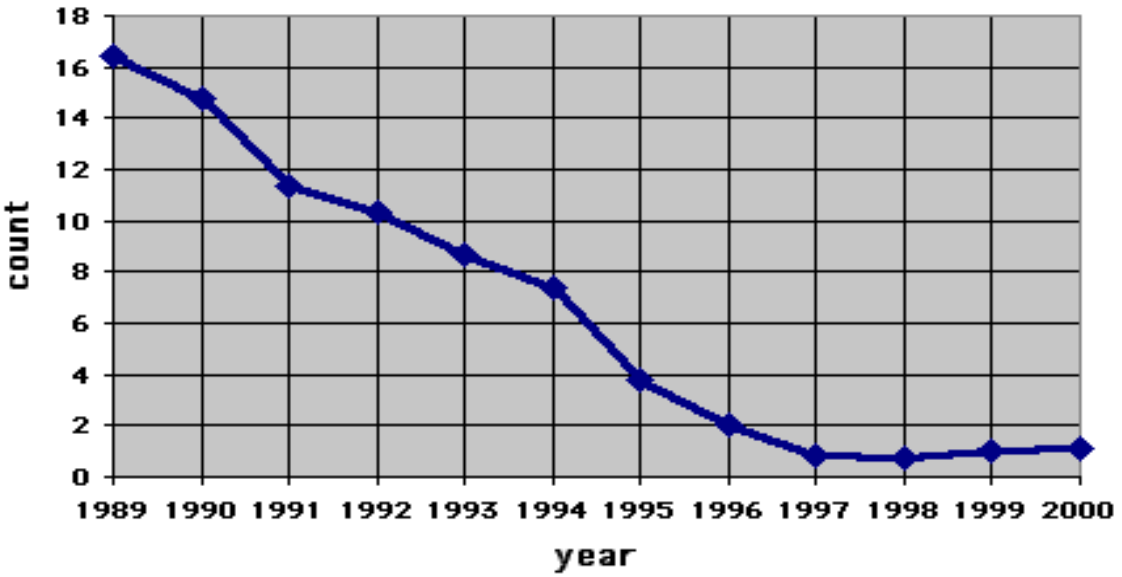

It must be recognized that this data is based not only on estimates, but on the reported incidents. An excellent study that tries to work through the limitations of such a "quick and dirty" look at the data was conducted by John Howard for his Ph.D. thesis.[6] Based on the CERT data, but with considerations for the implications of other studies (especially those conducted by DISA and the AIFWC), he concluded[7] that, in 1995, and individual host would have been involved in an incident at a rate of 1 time in 45 to 850 years. The range in this estimate reflects both uncertainties in the number of actual incidents, the type of incident (reporting is not thought to be uniform across all incident types), as well as differences in site security (e.g., the differences observed in the DISA and AIFWC studies). Based on the drop in relative incident rate from 1995 to 1997, Howard's rates should be lowered by a factor of ~4 for current threat levels.

---

[4] http://www.cert.org/stats/cert_stats.html
[5] http://www.isc.org/ds/host-count-history.html
[6] Howard, John D., 1997, An Analysis Of Security Incidents On The Internet 1989 - 1995, Ph.D. Thesis, Carnegie Mellon University, Pittsburgh, PA, April 7. Electronic version located at: http://www.cert.org/research/JHThesis/table_of_contents.html
[7] Howard, Table 14.1.

Another factor that must be kept in mind is that this data is global in nature. Not all sites attract the same level of "attention" or even the same type of threat agents. A good example of this can be found in the results of the 1998 InformationWeek/Price-Waterhouse-Coopers Global Security Survey[8] that was conducted in 50 countries and completed by 1,600 IT and security professionals. The key findings: (1) organizations engaged in Web commerce, electronic supply chains, and enterprise resource planning experience three times the incidents of information loss and theft of trade secrets than everybody else; and (2) revenue loss, though not prevalent, is seven times more likely to strike Web commerce sites compared with non-commerce sites. That E-commerce sites are likelier targets was also reported[9] in 1997; based on the Wheelgroup and NetSolve customer base, serious attacks occur 0.5 to 5.0 times per month per customer, with E-commerce sites falling at the upper end of the range. Another view can be formed by considering the computer crime cases that have been prosecuted under the computer crime statute, 18 U.S.C. §1030. From the USDOJ case summary[10] for March 18, 1998 through December 6, 2000, public (government) targets were involved in 10 cases, while private (commercial) targets were involved in 15 cases.

Rather than using probability or frequency data, it is sometimes possible[11] to measure risk in terms of cost impacts from security incidents (e.g., theft of proprietary information or financial fraud). Various organizations conduct periodic security surveys of organizations that include requests for cost impacts. It should be noted that such data are probably much fuzzier than that for incidents discussed above. This is because the organizations responding no doubt have widely differing impact potentials as well as differing definitions and practices in how they might calculate or estimate incident cost impacts. Nevertheless, the data can provide a "feel" for the loss potential due to cyber incidents. The figure below was created from data taken from a 1996 WarRoom survey,[12] and represents responses from 236 organizations spanning a wide range of industries.

[8] Dalton, Gregory, 1998, "Acceptable Risks," August 31, Copyright 1999, CMP Media Inc. Article found at http://www.informationweek.com/shared/printArticle?article=infoweek/698/98prrsk.htm&pub=iwk

[9] Roland, Craig, 1997, ProWatch Secure Network Security Survey, Wheelgroup, Nov 19, http://www.securityfocus.com/templates/archive.pike?list=1&msg=Pine.SOL.3.92.971119170616.13826A-100000@falcon.wheelgroup.com&_ref=1508664547

[10] USDOJ, Computer Crime and Intellectual Property Section (CCIPS), Computer Intrusion Cases, found at http://www.usdoj.gov/criminal/cybercrime/cccases.html

[11] After reviewing a number of such surveys, it would appear that roughly one-third of respondents are able to quantify their losses.

[12] Gembicki, Mark, 1996 Information Systems Security Survey, WarRoom Research, LLC, 23 November, found at http://www.infowar.com/sample/results.html

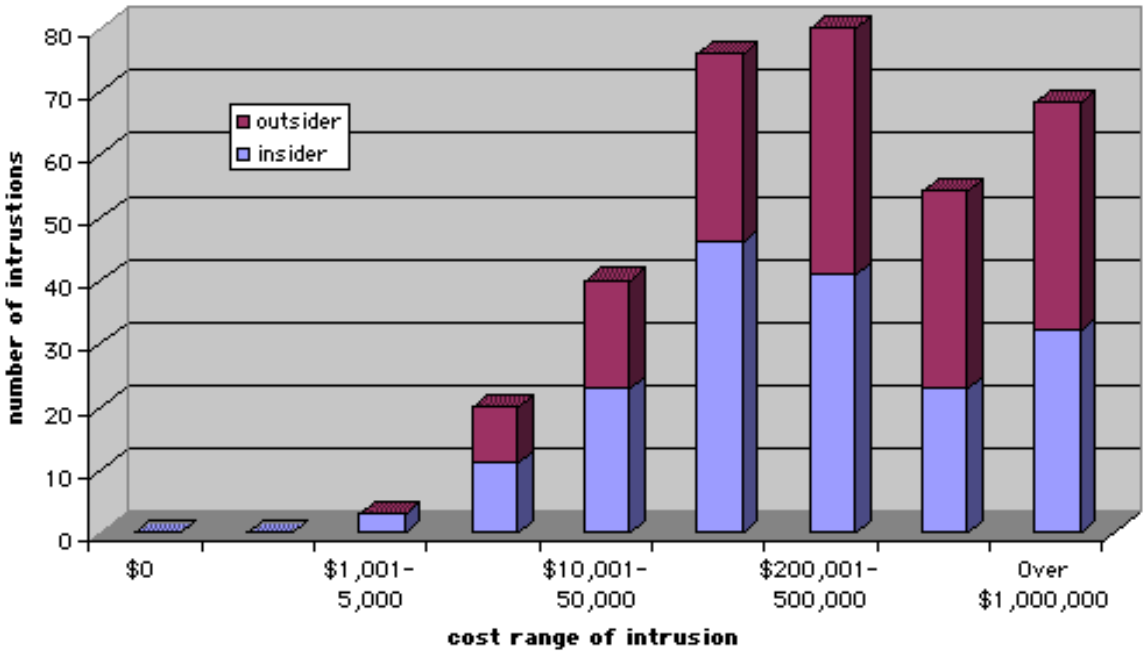

The data published by the Computer Security Institute (CSI) would seem to agree with these data reasonably well. For example, the summary data[13] for incidents with quantifiable losses published in 1997, 1998 and 1999 give an average loss in the range of $190K to $240K, a minimum reported loss of $50, and a maximum loss of $25M. (These figures were not corrected for inflation because it was felt the numbers were too soft for an adjustment to add any value.) In contrast, the data collected by Information Week[14] gives lower values, with 84% of the respondents that were able to quantify their losses placing them in the $1000 to $100,000 range, and only 16% in the over $100,000 category. Two things should be noted: (1) this data does not capture the effect of incidents that do not have a direct cost impact on an organization; and (2), it does not capture all aspects of computer crime (e.g., how do you measure the impact on society of someone using a site as a temporary but unauthorized staging area for distributing pornographic material or that from losses to investors due to stock manipulation via Internet fraud schemes that are currently estimated[15] to be running at something like $10,000,000,000 per year?).

Threat

So what gives rise to these risks? People. They range from the incompetent to the highly skilled. From the curious to the highly motivated. The recent DSB report included a

[13] http://www.gocsi.com/gifs/9903_loss.gif
http://courses.cs.vt.edu/~cs3604/lib/Crime/CSI.FBI.Report.1998.html
http://www.gocsi.com/testify.htm
http://www.gocsi.com/prelea_000321.htm
[14] Dalton
[15] Freeh, Louis J., 2000, Statement on Cybercrime for the Record before the Senate Committee on Judiciary, Subcommittee for the Technology, Terrorism, and Government Information, Washington, D.C., March 28. Found at:http://www.usdoj.gov/criminal/cybercrime/freeh328.htm

threat assessment,[16] part of which is shown below, that is important if for nothing else than because it illustrates one aspect of the changing nature of the threat: different players are coming onto the scene.

	Validated Existence[17]	Existence Likely but not Validated	Likely by 2005
Incompetent	W		
Hacker	W		
Disgruntled Employee	W		
Crook	W		
Organized Crime	L		W
Political Dissident		W	
Terrorist Group		L	W
Foreign Espionage	L		W

Another breakdown of the threat agents has been generated by the Global Annual Information Security Survey[18] conducted by Ernst and Young, as shown in the figure below.

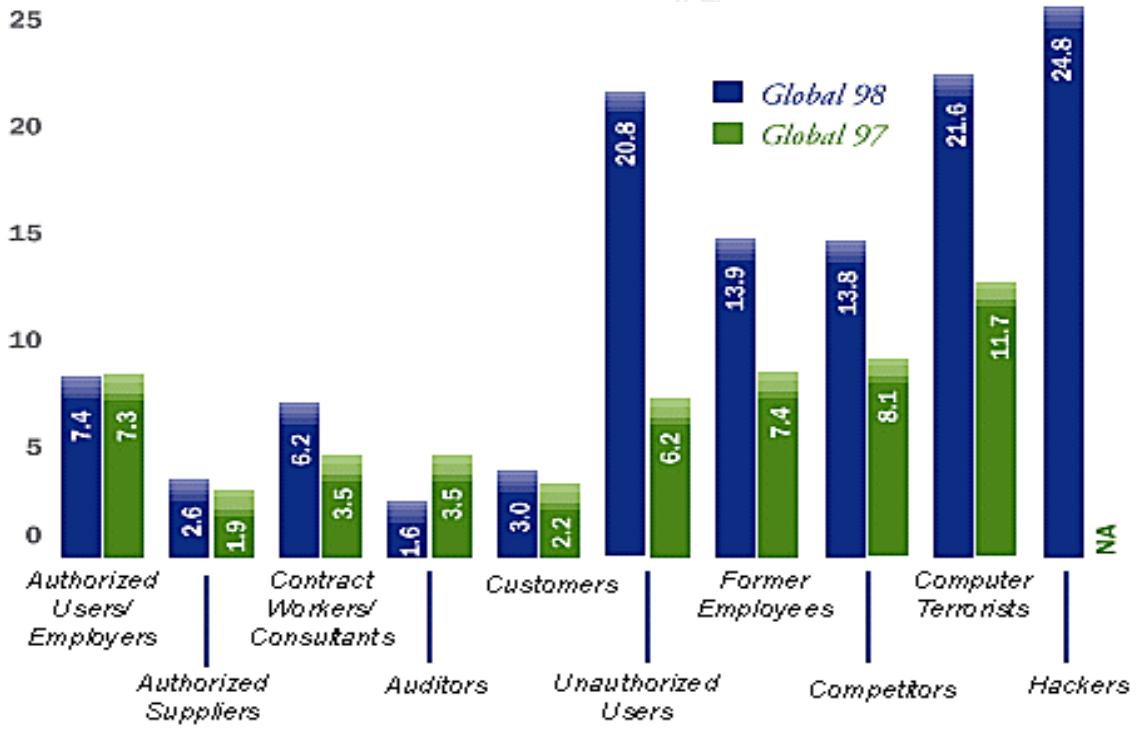

Considering potential targets against this or some other list of threat agents can lead to an understanding of intent. That is, understanding threat agents will help illuminate the questions of "are my systems a target?" and "to whom?" This, of course, is a function of what the capabilities of your system are and what is stored on them. (Recall this idea was

[16] DSB, Exhibit 2-6.
[17] Validated by DIA; W = widespread; L = limited
[18] Ernst and Young, 2001, 2nd Annual Global Information Security Survey. Found at http://www.ey.com/global/gcr.nsf/US/The_Current_State_of_the_IT_Enterprise_-_GIS_-_Information_Systems_Assurance_and_Advisory_Services_-_Ernst_&_Young_LLP

discussed in the Risk section above.) Such people can also be grouped by another category: whether they are "insiders" or "outsiders." Insiders are threats that access systems either directly (physical access) or via a company network or intranet. Outsiders rely on external connections (e.g., Internet or modem connections). While it is generally held that insiders pose the biggest threat (e.g., the 1998 InfoWeek survey[19] indicates that 58% of the respondents hold this to be true), the trend in the CSI data indicates that the external or outsider threat is now at least as great, if not greater:

> "Survey results illustrate that computer crime threats to large corporations and government agencies come from both inside and outside their electronic perimeters, confirming the trend in previous years. Seventy-one percent of respondents detected unauthorized access by insiders. But for the third year in a row, more respondents (59%) cited their Internet connection as a frequent point of attack than cited their internal systems as a frequent point of attack (38%)."[20]

Given a threat agent and intent, there is still the question of ability. One approximate breakdown[21] gives the following:

- 30% are accidental or curious people
- 50% are joyriders or vandals who have some networking knowledge but mostly use kiddy scripts
- 16 to 18% are IT professionals, network security experts, and programmers
- 2 to 4% are the true creators and innovators

Another view[22] suggests:

Sophistication	Population
Very!!	Hundreds
Aggressive	Thousands
Moderate	Tens of thousands
Script or browser users	Millions

While it is of some comfort to know that serious threats are limited in numbers, the Internet culture is such that significant sharing takes place—from a threat perspective this includes attack tools, target databases, and techniques. This gives rise to another aspect of the changing threat. Not only are new tools being created at the top, they tend to "flow down" to the less capable. This idea is best illustrated by the following two figures:[23]

---

[19] Dalton

[20] Rapalus, Patrice, 2000, Computer Security Institute press release, March 22. Found at: http://www.gocsi.com/prelea_000321.htm

[21] GTO Federal Network Systems, 2000, Introduction to Network Security & Intrusion Detection, v 1.3, p.10

[22] CERT, 2000, CERT/CC Overview Incident and Vulnerability Trends, CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, p.90. Found at: http://www.cert.org/present/cert-overview-trends/cert-history-trends-2000-08-17.pdf

[23] CERT, p. 88 and p. 93.

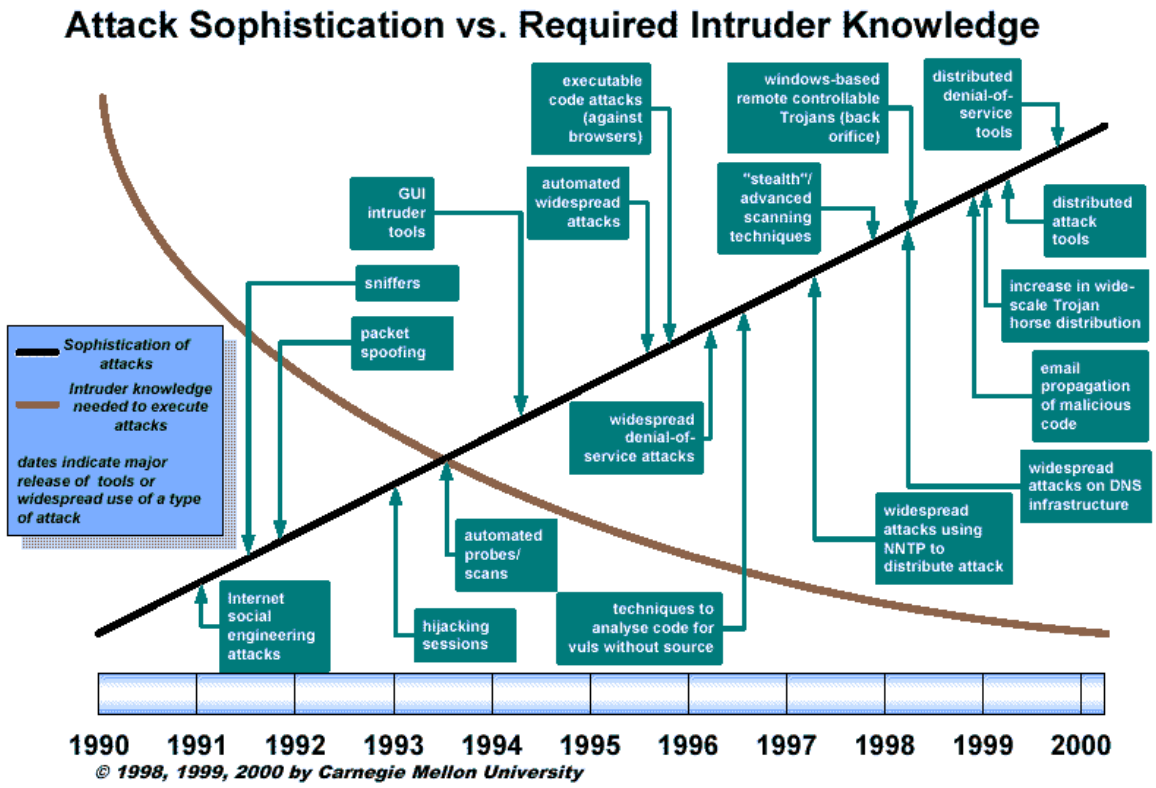

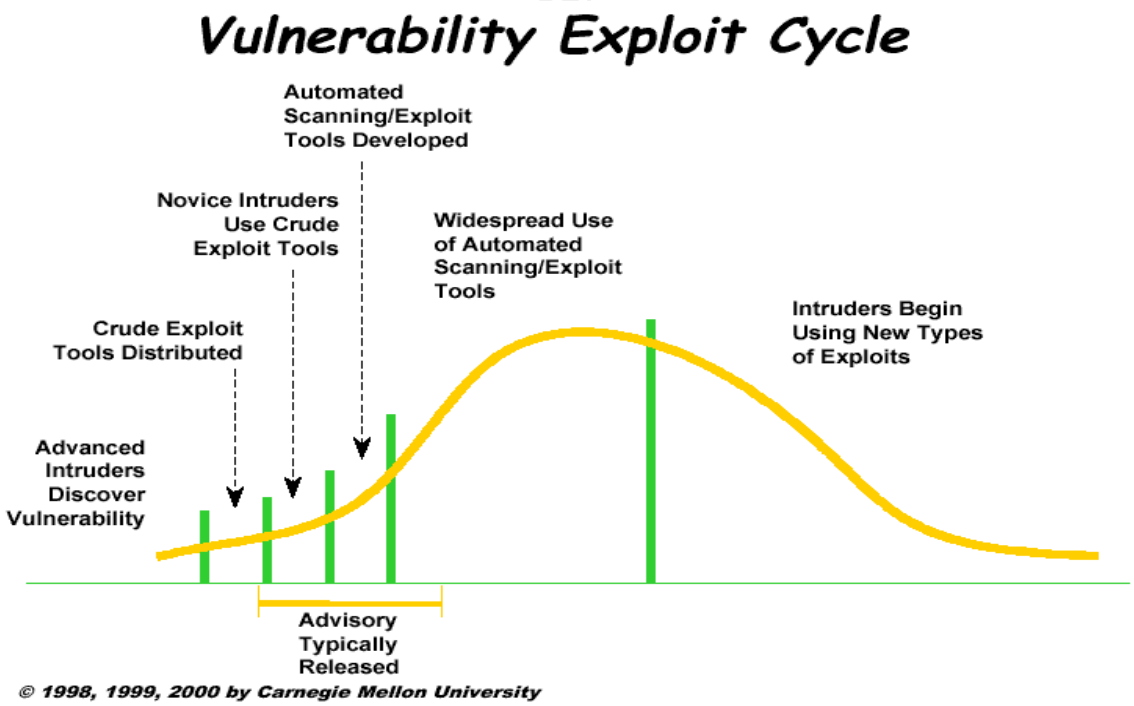

Note that the length of the exploit cycle might be measured in days or weeks, or it can be measured in years.[24]

[24] E.g., see Howard section 8.1.2.

Two points should be made here. First, even though ever increasing threat agents are becoming evermore increasingly sophisticated, Internet growth and improvements in cybersecurity have managed to keep the yearly rate of normalized incidents fairly level for several years (see the Risk section above[25]). However, if the combined effect of the rate of Internet growth and security improvements tapers off (which may happen soon—at least for the Internet growth factor), the bad guys might "catch up," and the normalized incident rate would be expected to increase. Secondly, in order to reduce the threat level of a new exploit, rapid incident reporting and quick responses to security advisories is called for. That is, while some sites will be initially affected by new exploit tools, in general there is sufficient response time to implement protective features. Said in another way, things can happen very quick in the cyberworld; so quick, in fact, that human response times are often inadequate (response being limited to recovery and corrective actions). However, from the perspective of the entire Internet, the widespread deployment, training and use of exploit tools does involve a human time scale, thus enabling the opportunity for human response.

Vulnerabilities

Some computer systems may be vulnerable simply because a conscious decision was made not to provide or enable any protection mechanisms. However, the concept of vulnerabilities is generally applied to systems where protections have been provided, but where such protections mechanisms contain faults, known or unknown, that can lead to a security incident. These breaches result from operational faults (configuration and policy errors), coding faults (includes programming logic errors, and faults of omission or commission), or environment faults (errors related to the execution environment). Not all faults lead to a security vulnerability. Generally speaking, known coding vulnerabilities are well documented by the security community. Given a good understanding of the hardware and software configurations of a particular system, it is possible to search such documentation for potential security problems that should be addressed. One place to start is with the ICAT Metabase,[26] a searchable index of information on computer vulnerabilities sponsored by the National Institute of Standards and Technology (NIST); as of 10 January 2001, this index included 2124 vulnerabilities. Other organizations that provide relevant information include: CERIAS,[27] FedCIRC,[28] ISS X-Force,[29] NIAP,[30] SANS Institute,[31] and Security Focus.[32]

Based on ICAT, NIST has generated some statistical data[33] on vulnerabilities deemed "important," from which several, general observations can be made (note: there is some overlap because some vulnerabilities affect multiple categories):

- 40% are exploitable locally (launched on system being attacked) while 62% are exploited across a network.

[25] This was also noted by Howard in section 7.2.2
[26] http://icat.nist.gov
[27] http://www.cerias.purdue.edu/
[28] http://www.fedcirc.gov/
[29] http://xforce.iss.net/
[30] http://niap.nist.gov/
[31] http://www.sans.org
[32] http://www.securityfocus.com/
[33] http://icat.nist.gov/icat.taf?_function=stats

- 23% affect the operating system, 24% the network stack or some other part of the communications protocol, 53% affect applications, 5% affect hardware, and 0.4% affect the encryption implementation.
- 22% directly exploit availability, 20% confidentiality, 19% integrity, and 53% security (violates access control policy; e.g., gives attacker root privilege)
- 68% affect UNIX platforms of some variant, 61% affect Microsoft windows platforms of some variant, 5% affect Apple OS, and 7% other operating systems.

Countermeasures

Countermeasures are, of course, the safeguards put in place to mitigate risks due to the combination posed by impacts or consequences, threats, and vulnerabilities. As such, they can take on various forms. For example, threat agents can be deterred from launching an attack through tough, enforced computer crime laws, warning banners, awareness and ethics training, and the like. Consequences can be reduced by segregation of data (e.g., limit the damage due to any one machine being subverted), the use of audits or other detective tools, and the use of backups (hardware and software). Vulnerabilities can be reduced through good system design (e.g., defense in depth), use of multiple technologies (e.g., mix of hardware and software vendors such that, in combination with defense in depth, no single point of failure exists), and in keeping up with vendor updates and actions recommended, for example, by one of the computer security advisory groups.

Without getting in too deep, suffice it to say that an excellent first action to take for an existing system would be to eliminate the top ten flaws.[34] Next it would probably be prudent to perform general platform configuration checks using some of the published guides and scripts.[35] Finally a system-wide look should be taken, such as by using the SANS roadmap.[36]

While it would be nice at this point to have a quantitative understanding of how security methods perform, there is little data to go on. The studies on risk presented earlier were not linked to specific systems or configurations. At best it is probably reasonable to suggest that the average risk rate would be representative of a typical installation, whatever that was at the time. One study that was contemporary with Howard's work was the 1995 Open Computing/NCSA Security Survey[37] that was based on the responses of 390 organizations. This survey suggested that at that time only half of the networks tied to the Internet used a firewall. Other security features used are depicted in the two figures below.

---

[34] SANS Institute, 2000, "How To Eliminate The Ten Most Critical Internet Security Threats: The Experts' Consensus." Found at: http://www.sans.org/topten.htm

[35] E.g., the checklists available from the Australian Computer Emergency Response Team at http://www.auscert.org.au/Information/Auscert_info/papers.html, the "Step-by-Step" guides available from http://www.sans.org, or hardening scripts such as that available for Solaris at http://www.yassp.org/ or Linux at http://bastille-linux.sourceforge.net/

[36] http://www.sans.org/newlook/publications/roadmap.htm

[37] Kelchner, Tom, 1995, Inside the Open computing/NCSA Security Survey, July. Found at http://www.wcmh.com/oc/features/previous/9507srvy.html

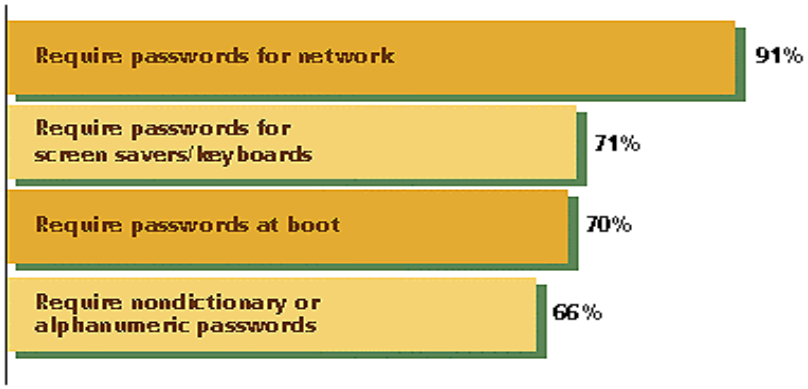

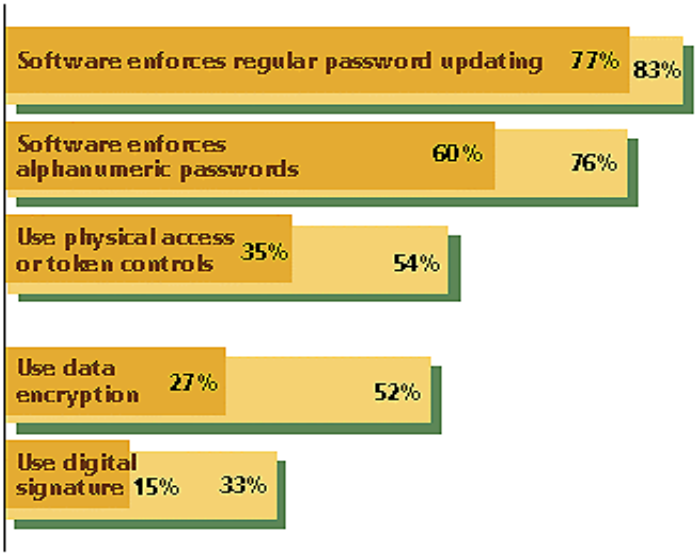

Here, in the second chart, the higher numbers reflect methods added by sites that were trying to improve their security as a result of known intrusions.

Summary

While not well correlated, sufficient data exists to provide a basic understanding of the various cybersecurity risk factors. Most importantly, the data that revealed the changing or adaptive threat pointed out the fact that security can not be implemented as a static solution. Continued improvements are required if risk is to be managed at current (or better yet, improved) levels. And because of the "logistics" cycle associated with exploit tools, organizations must coordinate their efforts if maximum reduction in threat levels is to be achieved; taking a parochial attitude will only subject one to the full brunt of the widespread use of exploit tools.