

System Dynamics Modeling for Information Security: Invitational Group Modeling Workshop

Preliminary Description of Thread on the Outsider Threat Problem

Proposal contributors

Agder University College, Norway: Jose J Gonzalez, Johannes Wiik

SEI/CERT, USA (PA): Andrew Moore, Howard Lipson, Tim Shimeall

Proposal history¹

Proposed: November 17, 2003

Revised: December 10, 2003, January 16 and 27-29, 2004

Objectives of the Outsider Threat Problem Thread

- to develop a preliminary System Dynamics model of important dynamic aspects of the outsider threat problem,
- to determine whether it is feasible to develop a generic system dynamics model for outsider threat that will be useful to all organizations, or at least all organizations within a single critical infrastructure sector
- to identify additional data on the outsider threat problem that is unknown or unavailable, but needed for future progress on this problem
- to investigate possible collaborations for longer-term work to propose to prospective sponsors

Objectives of this Proposal

- to scope the aspect of the outsider threat problem to be dealt with in the workshop
- to describe the problem in sufficient detail so as to provide a good starting point for the workshop and to ensure a good chance of success
- to provide the opportunity for other workshop participants to review and provide feedback for further refinement of this proposal prior to the actual workshop
- to help the reader make a decision regarding which workshop thread to choose (for an alternative thread, see “Proposal for a Thread on the Insider Threat Problem”)

Outsider Threat Thread Problem Outline

Issue

Attacks on information networks typically originate with the discovery of a ‘system vulnerability’ by advanced intruders (Schneier 2000, Ch. 13; see also Ellison and Moore 2002, p. 43ff). Crude exploit tools are developed and distributed, leading to the first attacks. Next, advanced automated scanning/exploit tools are used to roll out more numerous and more sophisticated attacks. On the ‘victims’ side, a process of developing defenses ensues (detection

¹ In previous stages the problem proposal had the working title “The System Vulnerability Threat to Information Security.”

of the source of the vulnerability, development of a software patch, distribution and installation of the patch). Ultimately, the attack wave subsides (Arbaugh, Fithen, and McHugh 2000; Lipson 2000). The whole process is referred to as a ‘vulnerability exploit cycle’ or ‘vulnerability life cycle’.²

Traditionally, most attacks on information networks have been performed by the hacker community. Although nasty enough, hacker attacks have mostly caused transient problems. However, the hacker threat might become an increasing nuisance as more advanced attack tools are developed and an exponentially growing number of poorly maintained units belonging to single individuals expand the number of vulnerable sites to unprecedented levels.

Schultz (2002) discusses the claim that insider attacks are more numerous than outsider ones. He concludes that it is wrong to claim that. But Schultz adds: «At the same time, however, to say both (sic!) that more *successful* attacks come from the inside (especially considering that so many organizations’ network security amounts to a “hard outer coating, but a soft-chewy middle”) is more likely to be true. Additionally, there is no debate that insider attacks pose a far greater level of risk than do outsider attacks.»

Careful reading of this statement reveals several issues:

1. There does not seem to exist data (at least not by 2002) to substantiate a hard claim about the relative prevalence of successful insider vs outsider attacks.³
2. No data is provided for the *dynamics* of point 1 (e.g. whether the relative prevalence of successful insider vs outsider attacks is changing over time and how).
3. The strong statement «... there is no debate that insider attacks pose a far greater level of risk than do outsider attacks» should be qualified, given that e.g. Code Red is claimed to have cost an estimated *total* of US\$ 2.6 billion (Ghosh 2002) and Love Letter of US\$ 8 billion (Ghosh 2000). Further, no data is provided for *dynamics* of the threat (e.g. whether the risk pattern is changing over time).

Indeed, a good question is whether increasing sophistication and automation of attacks might lead to more effectiveness of outsider attacks. (Although such automation also feeds into insider attacks, but up to this point subsets of the insiders have been too unsophisticated to be aware of such attack automation or already possess the access provided by such automation. This COULD change rapidly, however.)

There is a key difference between insider and outsider attacks: Insider attacks target normally single organizations – there are hardly instances of two or more insiders targeting several distinct organizations.⁴ In strong contrast, outsider attacks can be automated to target virtually any

² And, indeed, such vulnerability exploit cycle resembles a ‘product life cycle’ (see e.g. Ch. 9 in Sterman 2000).

³ Tim Shimeall added the following comment: «All of the objectively citable data I’m aware of shows much fewer insider vs. outsider attacks – but there is a lot of reason to believe the data is incomplete. So your statement here is true.»

⁴ A lot of the truth of this statement depends on the word “organization” – in the Harn case, http://www.baselinemag.com/print_article/0,3668,a=34708,00.asp, more than three different organizations were targeted (more than three different racetracks/pools of betters). They were all attacked via the same methodologies and access at a single location, but the loss was felt by several different organizations.

number of organizations. This aspect of *number* – if paired with large severity of attack – can make outsider attacks an issue of national or even international importance.

On the other hand, several authors argue that the distinction between insider and outsider threat is becoming more and more fuzzy (Lipson and Fisher 1999; Schultz 2002). Lipson and Fisher write: «In the highly distributed applications and Internet-based systems of today there is little distinction between insiders and outsiders.» Schultz states: «With all the outsourcing that is occurring, it is becoming difficult and hard distinction between insiders and outsiders... Additionally, many so-called “insider jobs” have turned out to be the result of complicity between an insider and an outsider.» With this in mind, this workshop should look for possible synergies between its two threads (insider and outsider threat).

There are reasons for considering the possibility that outsider attacks – rather than originating in dispersed hacker communities and without other major purpose than beating the defenses – could be coordinated and launched by organized groups (e.g. mafia or terrorists).

What would happen if a billionaire fundamentalist were to fund an organized activity involving hundreds of brilliant computer scientists for the purpose of identifying dormant system vulnerabilities, develop advanced exploit tools and – without advanced notice – roll out a series of devastating attacks towards some neuralgic point of the global information network? Could such roll out be too fast for the defenses to cope with, i.e. could the attacks outnumber the defenses? Could the financial game of chess be in a state of check for days and weeks – could it even be set mate? Could a well-timed attack cause the energy supply to break down at the worst thinkable moment, e.g. in the middle of an unprecedented ‘cold wave’ striking North America? (The issues here are very complex; it is hoped that the best answer is “probably not”, but the parameters on this “probably” are a very complex mix of defensive measures, business processes, redundancies and emergency planning as opposed to malicious access, malicious processing, attack resource requirements and other attack planning complexities. A nice challenge to system dynamics and other tools that have been developed to manage complexity to tack!)

A preliminary study of available literature has documented that such concerns about concerted and planned attacks are shared by colleagues (see e.g. Johnson, Guttman, and Woodward 1997; Schneier 2000; Shrobe 2002). We take such concern as indicative that it is worth looking closer at the issue, without necessarily assuming worst case scenarios.


Problem milestones

The problem sketched in the previous section is vast. For convenience we suggest approaching the problem in several stages, each of intrinsic interest, all of them connected by logical threads.

The Single Vulnerability Problem

Briefly stated, the problem is to understand the determinants of the life cycle of a single vulnerability (Arbaugh, Fithen, and McHugh 2000). For the purpose of illustrating the issue, the problem can be simplified to understand a generic life cycle such as given in Fig. 1 of Arbaugh,

Fithen, and McHugh's paper. (A qualitative rendering is presented in the figure below, using arbitrary units for both axes).⁵



A successful model would throw light on aspects of the hacker community that are difficult to assess directly (i.e. the ability of the model to render actual reference behavior from several well-studied vulnerabilities – such as Figs. 3-5 of Arbaugh, Fithen, and McHugh's paper – would increase our faith in the model's assumption about how knowledge spreads in the hacker community, etc.)

Further, the model could be used to:

- Experiment with policies to prevent or at least reduce the impact of attacks.
- Help identify trends and explain the determinants of trends
- Assess the extent to what current data can be used for numerical simulation analysis; suggest further empirical studies; suggest further modeling studies; all these would be steps in developing an application for funding.
- Build a basis for the next two problem stages

We do have a simple System Dynamic model for the Single Vulnerability Problem that is able to render the idealized reference behavior shown in the figure above. Arguably, this model can be used as a basic for further discussions at the workshop. (A detailed description of the model will be distributed one week ahead of the workshop.)

The Multiple Vulnerability Problem

Hackers roll out attacks on vulnerabilities so that exploitation of several vulnerabilities might overlap and interact. At the outset one would expect that the ideal shape of the vulnerability life cycle depicted above will not hold when several vulnerabilities compete for the hacker

⁵ At this stage the figure should not be taken as too literally. This is one of many curve shapes (basically, one can vary the amplitude, skew, kurtosis, and variance of the above curve pretty arbitrarily and find examples in the real world; there are numerous examples where the tail is much steeper than the head). Rather, the purpose of the figure is to serve as generic reference behavior mode for a simple model prototype.

community's attention. In fact, Figs. 3-5 in Arbaugh, Fithen, and McHugh's paper do suggest the existence of two or more humps. Indeed, there are lots of reasons for multiple humps other than competition – such as bundling of attack tools, “attractiveness” of the target, new exploit discovery, etc.

We have not investigated in depth the issue of reference behavior for the Multiple Vulnerability Problem. It would appear that there is at least partial information in terms of the aggregated number of incidents reported (the CERT/CC Statistics 1988-2003, http://www.cert.org/stats/cert_stats.html). There are some curves showing incident reports against services and vulnerabilities that may be relevant – see <http://www.cert.org/archive/ppt/IntelDataExamp.ppt>, although there is more data than is shown there.

Again, a successful model would throw light on aspects of the hacker community that are difficult to assess directly (i.e. recruitment to the hacker community, how the hackers pick various vulnerability types, the R&D pipeline in the hacker community, etc)

Further, the model could be used to:

- Experiment with policies to prevent or at least reduce the impact of attacks. An intriguing issue would be to model traceability (Lipson 2002).
- Devise better ways to monitor and measure hacker activity including spread of information
- Impact of growth of IT systems including larger system integration
- Help test controversial issues (e.g., whether full disclosure of vulnerabilities is beneficial or detrimental).
- Assess the impact of the R&D competition between defenders and attackers. Again, traceability – if feasible – could provide ways to explore this issue.
- Help identify trends and explain the determinants of trends
- Help explore the robustness of internet sectors (i.e. the energy or the health care sector) in various scenarios of increasing frequency of attack or even of greater attack sophistication and automation.
- Help estimate damages of potential attacks and assess cost-benefit ratios of defensive policies.
- Assess the extent to what current data can be used for numerical simulation analysis; suggest further empirical studies; suggest further modeling studies; all these would be steps in developing an application for funding.
- Build a basis for the next problem stage

The Concerted Attack Problem

Here the issue would be the ability of internet sectors (i.e. the energy or the health care sector or even the .mil domain) to perform critical functions in case of a concerted attack.

At first sight, modeling such challenge would involve the identification of a target sector within the information network for the study (such as the financial system, the energy sector, possible the .mil domain). One would need to estimate the amount of dormant system vulnerabilities (e.g.

buffer overflow bugs in Internet Explorer, etc),⁶ the rate of discovery of vulnerabilities depending on the attacker's resources, the rate of development of defenses once vulnerabilities are exploited by attackers, etc.

Some, may be most of the data needed might be known, or at least easily ascertained, other data might be uncertain, but critical (thus demanding further studies).

The model could be used to:

- Explore the robustness of various internet sectors to concerted attacks
- Help test controversial issues (e.g., degree of causation by software developers vs system managers).
- Devise better ways to monitor hacker activity
- Help identify trends and explain the determinants of trends
- Help explore the robustness of internet sectors (i.e. the energy or the health care sector) in various scenarios of increasing frequency of attack or even of greater attack sophistication and automation.
- Help estimate damages of potential attacks and assess cost-benefit ratios of defensive policies.
- Assess the extent to what current data can be used for numerical simulation analysis; suggest further empirical studies; suggest further modeling studies; all these would be steps in developing an application for funding.

We have implemented a generic model developed by Rudolph and Repenning (Rudolph and Repenning 2002b, 2002a) using Powersim Studio. The purpose of the exercise was to keep the model in store for adaptation to information security.

We discovered several minor errors in the model ("minor" in the sense that they do not affect the overall behavior of the model, i.e, the conclusions of the study stand).

The Rudolph and Repenning study might be useful for the workshop in two ways:

- It can help making concepts of System Dynamics available to a general audience
- It might be used as a starting point for the concerted attack problem, including suggesting reference behavior modes (the "system" copes with concerted attacks up to some threshold. Above that threshold it becomes "choked".)

Perspective

Several perspectives can be appropriate (to be further discussed at the workshop):

- Software vendor

⁶ Undiscovered software bugs can often provide points of attack for malicious agents. It is estimated that released professional software has between 5 and 15 undiscovered bugs per 1000 lines of code. Windows 2000 has 43 million lines of code, Windows XP even more. Interesting question: How many of the estimated half million bugs in Windows XP (and other software) can be exploited by attackers?

- System manager
- Homeland security (government, military, etc)
- Scientific (NSF, etc)

References

- Arbaugh, William A, William L Fithen, and John McHugh. 2000. Windows of Vulnerability: A Case Study Analysis. *Computer* 33 (12):52-59.
- Ellison, Robert J., and Andrew P. Moore. 2003. *Trustworthy Refinement Through Intrusion-Aware Design (TRIAD)*. CMU/SEI 2002 [cited 17.11. 2003]. Available from <http://www.cert.org/archive/pdf/03tr002.pdf>.
- Ghosh, Anup K. 2000. *Code-Driven Attacks: The Evolving Internet Threat*. Dulles, VA: CIGITAL.
- . 2002. *Challenges for Anomaly Detection of Program Exploits*. Baltimore, MA: Johns Hopkins University.
- Johnson, Dale M. , Joshua D. Guttman, and John P. L. Woodward. *Self-Analysis for Survival, in 1997 Information Survivability Workshop - ISW'97*. Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA 15213-3890 1997 [cited January 16, 2004]. Available from http://www.cert.org/research/isw/isw97/all_the_papers/no14.html.
- Lipson, Howard F. *Survivability – A new security paradigm for protecting highly distributed mission-critical system* 2000 [cited 17.11.2003]. Available from <http://www.cert.org/archive/pdf/surviv-paradigm.pdf>.
- . *Tracking and Tracing Cyber-Attacks: Technical Challenges and Global Policy Issues*. CMU/SEI 2002 [cited 29 January 2004]. Available from <http://www.cert.org/archive/pdf/02sr009.pdf>.
- Lipson, Howard F., and David A. Fisher. 1999. Survivability — A New Technical and Business Perspective on Security. Paper read at New Security Paradigm Workshop, September 22-24, 1999, at Caledon Hills, Ont. Canada.
- Rudolph, Jenny W., and Nelson P. Repenning. 2004. *Disaster Dynamics Model Documentation* 2002a [cited January 14 2004]. Available from Home page: [http://mitsloan.mit.edu/omg/people/index.php-->People-->Faculty Nelson Repenning-->\[3\]](http://mitsloan.mit.edu/omg/people/index.php-->People-->FacultyNelsonRepenning-->[3]) Rudolph, J. and N. Repenning (2002). Disaster Dynamics: Understanding the Role of Stress and Interruptions in Organizational Collapse, *Administrative Science Quarterly*, 47: 1-30. -->Link to model.
- . 2002b. Disaster Dynamics: Understanding the Role of Quantity in Organizational Collapse. *Administrative Science Quarterly* 47:1-30.
- Schneier, Bruce. 2000. *Secrets and Lies: Digital Security in a Networked World*. New York: John Wiley & Sons, Inc.
- Schultz, E. Eugene. 2002. A Framework for Understanding and Predicting Insider Attacks. *Computers and Security* 21 (6):526-31.
- Shrobe, Howard. 2002. Computational Vulnerability Analysis for Information Survivability. *AI Magazine* 23 (4):81-91.
- Sterman, John D. 2000. *Business Dynamics : Systems Thinking and Modeling for a Complex World*. Boston: Irwin/McGraw-Hill.