

Survivability-Over-Security: Providing Whole System Assurance

William Yurcik¹ David Doss

Hans Kruse

Dept. of Applied Computer Science
Illinois State University
{wjyurci,dldoss}@ilstu.edu

McClure School²
Ohio University
hkruse1@ohiou.edu

***Abstract:** Whole system assurance is necessary since over-reliance on protection solutions for system components has actually contributed to the fragility of information systems when viewed as a whole. For instance, the use of authentication and encryption to protect networked systems may actually add more vulnerabilities to the system as a whole than they eliminate. The goal of this research is to increase the survivability of information systems to underlying failures by focusing on reducing net system vulnerabilities and increasing net system restoration capabilities. Specifically we give examples of current projects that are in varying states of development.*

1.0 Introduction

Research into reliable and survivable systems is important due to the increasing number of information systems and society's increasing dependence on these systems for service. Disruption of service can be very expensive to businesses, life-threatening for emergency services, and ultimately threaten the defense and economic security of the U.S. In addition to the increasing number and dependence upon information systems is also the increasing interconnectedness and interdependence between large and complex systems; a failure in one sector can easily affect other sectors. The question has been raised that given the fragility of U.S. infrastructures to small random failures (such as hardware failures, bad software design, innocent human errors, and environmental events) which have caused large scale regional or national outages at an increasing rate over the last two decades, what is the vulnerability of U.S. infrastructures to a malicious, intelligent, well-funded and determined attack?

Current security approaches to protect information systems have focused on preventing attacks from being successful by hardening defenses with authentication, encryption, and a variety of layer-violating network devices (i.e., firewalls, network address translators, intrusion detection systems). What is not being captured is the survivability of an entire system to failures or attack. While security approaches may protect one layer of a networked system they often introduce vulnerabilities in other layers. There are many examples: authentication schemes that introduce a single-point-of-failure (Certificate Authority); encryption systems that rely on single-point-of-failure secure bastion hosts; action-response intrusion detection systems that disconnect from the Internet given an attack signature (self-inflicted denial-of-service); and firewalls which break end-to-end protocols for network management while allowing trivial insider attacks once circumvented. After considering the net vulnerability of information systems, we find that security protection may actually add more vulnerabilities to the system as a whole than they eliminate.

Thus our research group has coined the term “**Survivability-Over-Security**” (SOS) as a descriptive focus of our over-reaching goal: to increase the survivability of information systems using innovative techniques which simultaneously reduce net vulnerabilities and increase restoration flexibility. While security is one technique to protect system components, we feel that survivability is a higher goal over security since survivability encompasses the functionality of an entire information system and not individual components.

¹ author for correspondence; additional contact information: telephone/fax (309) 438-8016/5113 ; hardcopy (postal) mail: 202 Old Union, Campus Box 5150, Normal IL. 61790 USA

² J. Warren McClure School of Communication Systems Management

2.0 The Concept of Survivability

Previous research in the application of reliability research to large complex systems has focused on characterizing failure distributions to answer questions such as why systems fail, are systems failing more or less often, and what can be done to make systems fail less. Reliability analysis assumes failure events are independent (for mathematical analysis) and emphasizes preventing the most probable failure events from occurring.

Reliability is defined as the ability of a system to perform its required functions under given conditions for a given time interval.

Survivability is defined as the ability of a system to perform required functions at a given instant in time after a subset of components of a system becomes unavailable.

Survivability research builds upon reliability research to focus on recovery after a failure has left part of a system unavailable. Survivability analysis allows failure events to be correlated (as in an intelligent attack) and emphasizes recovery from failure events to the most critical system components (not necessarily the most probable failure events). The goal of survivability research is to maintain continuous service performance to users despite underlying system failures. The best examples of implemented survivability research are combat aircraft that can still fly despite extensive underlying system damage. While reliability research assumes failures may be eliminated, survivability assumes failures will continue to occur but mission functionality can be maintained despite underlying system degradation. Survivability is the appropriate risk-avoidance approach given an unbounded failure set and the potential a failure could cause a large-scale outage to a critical national infrastructure if no recovery procedures exist.

There are tradeoffs between reliability and survivability: if a system has high reliability (large MTBF or slow growth distribution) then survivability may not be important (especially if the MTBF is beyond the expected lifetime of the system); if a system has high survivability (can transparently recover from all possible fault scenarios) then reliability may not be important. However, when considering a malicious attacker, the assumptions of reliability analysis no longer hold since failures become correlated and designed for maximum disruption and damage.

Survivability is an especially appropriate concept in the context of computer network infrastructure because, as most network managers will attest, under normal operating conditions some part of the network is almost always broken during any given time period. Thus providing survivability to computer networks can serve as an illustrative model for any infrastructure based upon underlying networked computer systems. Another reason for choosing computer networks is that while emphasis on end-user computer system security is vital, a network model can exhibit the dynamics of system interconnectivity where a failure in one system can affect many other systems due to interdependencies and multiplier effects. While it can be argued that end-system computer security encompasses network devices (i.e., routers, switches, etc.), the reality is that network components have additional vulnerabilities due to layered time-based coordination.

3.0 Selected SOS Research Projects

3.1 VPN Deployment Issues [3, 5, 6]

Deploying Virtual Private Networks (VPNs), especially in a large-scale environment, requires a large amount of planning. We are developing a framework for organizations to use in planning for VPN deployment based on *layered* security issues. It is designed as a high-level process that can be used to distribute VPN deliverables among different workgroups. Specifically, we describe security issues corresponding to five distinct layers of analysis: (1) VPN technology; (2) Firewalls; (3) Legacy Networks; (4) Survivability; and (5) Legacy Applications. Despite the common perception that a VPN is not a customizable solution, we conclude that there is a broad spectrum of VPN options available with each having its own strengths, weaknesses, and vulnerabilities. It is anticipated that no single VPN solution will

supplant others but instead a diversity of choices will continue to emerge. The goal is to coordinate VPN deployment without introducing additional vulnerabilities due to distributed weaknesses to attack.

3.2 Unintended Protocol Interactions Caused by Layer-Violating Devices [2,7]

In 1981, Saltzer, Reed, and Clark identified “end-to-end” principles related to the design of modern *layered* protocols. The Internet today is not as transparent as envisioned by Saltzer et. al. While most of the intelligence remains concentrated in end systems, users are now deploying more sophisticated processing within the network for a variety of reasons including security, network management, E-commerce, and survivability. Applications and application-layer protocols have been found to interact in unexpected ways with these new *layer-violating* (LV) network devices (which break the end-to-end model) such as network address translators, firewalls, proxies, intrusion detection, and differentiated service functions. We survey examples of problems caused by the introduction of this new processing within the network, which is counter to the end-to-end Internet model. LV devices can take on both positive and negative roles with respect to survivability. On the positive side firewalls, traffic shapers, and intrusion detection devices keep the network running but break some application deployment. On the negative side, some LVs require the storage of state information, and therefore become a single point of failure. Transition to IPv6 is one solution, but by the time IPv6 is eventually implemented the nature of the Internet may be drastically different. Future work is obviously needed to create a consistent environment for protocol development that preserves the transparency provided by the end-to-end Internet model. The goal is to reduce net vulnerabilities by preserving network transparency while coexisting with LV devices.

3.3 Open Source versus Commercial Security Solutions [4]

We are participants in the current debate between “open source” versus “security-by-obscurity” by having compared two specific firewall solutions in detail highlighting corresponding security risks: (1) a firewall using only open source software available for the LINUX operating system and (2) a commercial firewall solution using the CISCO IOS firewall feature set. Open source follows the “many eyes” principle which states that the more developers work on shared code, the fewer secrets and the harder to compromise since more people will detect errors and fix them. Security-by-obscurity argues for hiding the code as a deterrent to breaking the code. While open source code is less expensive and offers more control to system administrators, commercial solutions are a proven commodity that have held up relatively well over time despite being more expensive and less flexible. Which approach is better is not a simple question. We found that the most powerful present firewall solution may be a hybrid of an open source application-level gateway system with commercial network-level packet filtering. The goal is reduce net vulnerabilities introduced by both open source and commercial products.

3.4 Network Survivability in the Multicast Context [8,9]

Multicasting is one of the fundamental technologies necessary to scale the Internet to very large sizes. Because of savings in bandwidth and switch processing, network providers consider multicasting the technology of choice for providing multimedia services to consumers. In its simplest form, multicasting allows a single stream over common links which is then replicated at branch points. In order to manage large bandwidth multimedia feeds and to slow congestion by eliminating waste, multicasting concentrates resources to gain efficiency. However, this concentration amplifies the impact of failures such that a random fault or malicious attack will affect a larger number of users and thus may introduce more serious vulnerabilities. We have pursued the use of preplanned, dedicated bandwidth, disjoint backup route techniques for multicasting on cell-switched networks (ATM) which can provide survivability guarantees. Results show that providing survivability via self-healing survivable rings at the traffic layer is both more feasible and cost-effective (in terms of resources) than either tree or mesh techniques. The goal is reduce net vulnerability to attack on multicast flows while maintaining broadcast and group communication functionality.

3.5 Application Layer Survivability [1]

We have demonstrated one project which provides survivability at the application layer. A lost essential service can be replaced by another service that supports mission fulfillment in a different but equivalent way.

4.0 Summary

The goal of Survivability-Over-Security is to coordinate different incremental research efforts such that survivability can be effectively and efficiently provided against intelligent malicious attacks. This is accomplished by reducing net system vulnerabilities and increasing net system restoration capabilities. While our research provides valuable quantitative analysis specifically for computer networks, we feel the qualitative results are equally applicable to all critical infrastructures since all critical infrastructures can be abstracted to a system of links and nodes dependent upon networked computer systems.

5.0 Acknowledgments

We would like to give special acknowledgment to the SAND (Survivability Analysis and Network Design) Research Group and the significant intellectual contributions of David Tipper/University of Pittsburgh and Deep Medhi/University of Missouri-Kansas City.

6.0 References

- [1] Grzywa, Mary S. *Resumable FTP Transfer Protocol*, M.S. Thesis, Illinois State University, Dept. of Applied Computer Science, February 2000.
 - [2] Kruse, H., Yurcik W., and L. Lessig, "The InterNAT: Policy Implications of the Internet Architecture Debate," *Telecommunications Policy Research Conference (TPRC)*, Sept. 2000.
 - [3] Patton, S., Doss D., and Yurcik W. "Distributed Weakness in Virtual Private Networks," *25th Annual IEEE Local Computer Networks (LCN) 2000 Conference*, to appear Nov. 2000.
 - [4] Patton, S., Doss D., and Yurcik W. "Open Source Versus Commercial Firewalls: Functional Comparison," *25th Annual IEEE Local Computer Networks (LCN) 2000 Conference*, to appear Nov. 2000.
 - [5] Patton, S., Smith B., Doss D., and Yurcik W. "A Virtual Private Network Deployment Framework," *25th Annual IEEE Local Computer Networks (LCN) 2000 Conference*, to appear Nov. 2000.
 - [6] Patton, S., Smith B., Doss D., and Yurcik W. "A Layered Framework Strategy for Deploying High Assurance VPNs." *Fifth IEEE Intl. Symp. on High Assurance Systems Engineering*, to appear Nov. 2000.
 - [7] Yurcik, W., Doss D., and Kruse, H. "Challenges to the End-to-End Internet Model." *Americas Conference for Information Systems (AMCIS) 2000*, Aug. 2000.
 - [8] Yurcik, W. and Tipper D., "Survivable ATM Group Communications: Issues and Techniques." *8th Intl. Conference on Telecommunication Systems*, March 2000, pp. 518-537.
 - [9] Yurcik, W., Tipper D., and Medhi D. "Scalable Survivable ATM Group Communications," *IEEE Military Communications Conference, (MILCOM'2000)*, to appear October 2000.
-