Use of Diversity as a Defense Mechanism

Panelist's statement

Roy A. Maxion
Dependable Systems Laboratory
Computer Science Department
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213 / USA

maxion@cs.cmu.edu

ABSTRACT

Diversity, a concept suggestive of a composition of distinct or unlike elements or qualities, has served to mitigate error in modern computer systems for decades, going back at least as far as the 1971 JPL STAR (self testing and repairing) system, designed and built in the Spacecraft Computers Section of the Jet Propulsion Laboratory Astrionics Division [2]. In that context the concept of diversity was termed *redundancy*. In computer security, diversity is being contemplated as an approach toward mitigating security breaches, or what might be regarded as errors in security. The panel contemplates various issues regarding diversity and security, and this panelist in particular raises a number of questions whose answers may prove valuable at such time as they become available. Until then, perhaps these questions will serve to provoke thoughtful research directions.

Categories and Subject Descriptors

D.2.4 [**Performance of Systems**]: Fault tolerance, reliability, availability and serviceability; K.6.5 [**Security and Protection**]: Management of Computing and Information Systems

Keywords

Diversity, intrusion tolerance, security

1. INTRODUCTION

Diversity has been a consistent theme in building computer systems whose logical engines must continue to operate even when their physical engines fail [7]. The hardware notion of "masking redundancy" is similar, for example, to how we think of diversity's being helpful in today's security systems [10]. In the evolution of fault tolerance, diversity took on software in the form of N-version programming; different, independently-developed versions of software were executed concurrently, based on the idea that the same fault would not affect all the versions at once [1].

Participants in this workshop are not the first to consider diversity as an asset to security. There are too many contributions to review in the short space available, so only a few papers are selected for mention here, simply to provide a place to start. Yves Deswarte and his colleagues have specifically discussed diversity in the context of deliberate (malicious) faults [4]. Bev Littlewood (a fellow panelist) and Lorenzo Strigini have written about the roles and limits of diversity in security, and the extent to which research in security can draw on the lessons from diversity in fault-tolerant systems [5]. Eric Totel and his colleagues address an approach to anomaly detection (a frequently-used technique in intrusion-detection systems), based on design diversity [9]. Finally, James Reynolds and his colleagues show how they have built an intrusion-detection system that avails itself of redundancy, diversity and N-version programming [6]. These are just a few examples. Much of the history of redundancy and diversity in computing is reflected in the taxonomy of concepts for dependable and secure computing recently produced by Avizienis and his colleagues [3].

2. DIVERSITY IN ANOMALY DETECTION

A common assumption in anomaly-based intrusion detection is that one size fits all: a single anomaly detector should detect all anomalies. Experience has shown this assumption to be too broad, so compensation for performance shortcomings is sometimes effected by resorting to "correlation" techniques that combine the results from different detectors; this could be seen as making use of detector diversity. Such diversity is intuitively based on the proposition that detector coverage is different – perhaps widely so – for different detectors, each covering some disparate portion of the anomaly space. Diversity, then, enhances detection coverage by combining the coverages of individual detectors across multiple sub-regions of the anomaly space, resulting in an overall detection coverage that is superior to the coverage of any one detector. However, there are few studies which have portrayed the measured effects of diversity amongst anomaly detectors. One recent investigation indicated that detection coverage due to diversity does indeed improve, but the improvement is due less to broad disparities in coverage, and due more to small differences at the edges of the detector space [8]. Despite the interesting results of this study, much remains to be done before the concept of diversity in anomaly detection can be usefully employed.

NSPW 2005 Lake Arrowhead CA USA

It's important to note, incidentally, that different detection algorithms, such as neural nets and Markov models, may seem intuitively quite diverse, but their effective diversity is very small. Because their detection performances overlap significantly, little is gained through their combination.

3. LOOKING FORWARD

Various questions and issues arise when one thinks about diversity in anomaly-detection software and anomaly-detection programs, as well as in security settings in general. A few of these, which might be interesting for discussion, are:

- Apart from the intuitive understanding of diversity, is there a concrete definition of diversity from which we can work, and for which we can obtain sensible and reliable measures?
- Can diversity be used to the benefit of computer security by virtue of the same practices with which diversity has been employed in dependable and fault-tolerant systems? What are the similarities and differences, and at which points do they matter?
- What is the precise definition of diversity as one uses it in a specific context (e.g., anomaly detection)? How does it differ, if at all, from definitions for alternative contexts or for a very general usage? Does the definition extend beyond software or detection, and into other areas of diversity, or is it restricted just to specific domains?
- Diverse means different. How different do two programs, detectors, behaviors, etc. need to be before we call them diverse? How would this be measured? Would these measures be in terms of differences among the diverse entities themselves (e.g., among the programs, etc.), or would they be among their behaviors?
- In terms of using diversity in a security setting, what is the goal? What is it that diversity will make better?
- It appears that diversity in the security context is a concept that has been borrowed from early studies in hardware (and software) diversity. Hardware diversity was able to handle static faults, such as stuck-at-zero logic, not dynamic faults, such as malicious faults. How can diversity in any guise be made to handle (measurably) the dynamic nature of malicious faults whose characteristics change continually?
- If diversity works best when faults (intrusions, etc.) are independent from one another, how should that independence be determined and measured? We'll need these measurements in order to justify the expense of diversity. Less independence means less effectiveness, but not necessarily less effort/cost.
- Must "faults" be independent, if diversity is to help?
- What are the dimensions of diversity that we care about? Common dimensions are time, space (extra hardware or software), and data (replications). What more, if any, are needed for a security setting?
- What is the quantifiable benefit of adding one more diverse component? How is that determined?
- What is the cost of adding one more diverse component?
- What would a taxonomy of diverse detectors look like?
- One benefit of diversity is that it may contribute to what is called "reliability growth," the phenomenon of increasing (growing) reliability by having diverse alternatives to fall back on in fault conditions. A related concept may be security growth. Does this concept work in security?

4. CONCLUSIONS

These questions are intended to provoke thought and discussion. No widely accepted answers are known, but this workshop may provide a place to start.

5. ACKNOWLEDGEMENTS

The author is grateful for thoughtful comments contributed by Kymie Tan.

6. REFERENCES

[1] A. Avizienis and L. Chen. On the implementation of N-version programming for software fault-tolerance during program execution. In *Proceedings of IEEE Computer Society's First International Computer Software and Applications Conference (COMPSAC '77)*, pages 149–155, Chicago, IL, 08-11 November 1977. IEEE Computer Society Press.

[2] A. Avizienis, G. C. Gilley, F. P. Mathur, D. A. Rennels, J. A. Rohr, and D. K. Rubin. The STAR (Self-Testing and Repairing) computer: An investigation of the theory and practice of fault-tolerant computer design. *IEEE Transactions on Computers*, C-20(11):1312–1321, November 1971.

[3] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1):11–33, January-March 2004.

[4] Y. Deswarte, K. Kanoun, and J.-C. Laprie. Diversity against accidental and deliberate faults. In *Proceedings of the Conference on Computer Security, Dependability, and Assurance: From Needs to Solutions (CSDA '98)*, pages 171–181, Washington, DC, 07-09 July 1998. IEEE Computer Society Press.

[5] B. Littlewood and L. Strigini. Redundancy and diversity in security. In *Proceedings of the 9th European Symposium on Research in Computer Security (ESORICS 2004)*, P. Samarati, P. Ryan, D. Gollmann, and R. Molva (Eds.), pages 423–438, Sophia Antipolis, France, 13-15 September 2004. Lecture Notes in Computer Science, Vol. 3193, Springer-Verlag, Berlin.

[6] J. C. Reynolds, J. Just, L. Clough, and R. Maglich. On-line intrusion detection and attack prevention using diversity, generate-and-test, and generalization. In *Proceedings of the 36th Annual Hawaii International Conference on System Sciences (HICSS'03)*, CD-ROM Track No. 335.b, Wailea, Maui, Hawaii, 06-09 January 2003. IEEE Computer Society Press.

[7] D. P. Siewiorek and R. S. Swarz. *Reliable Computer Systems: Design and Evaluation.* A. K. Peters, Natick, Massachusetts, third edition, 1998.

[8] K. M. C. Tan and R. A. Maxion. The effects of algorithmic diversity on anomaly detector performance. In *International Conference on Dependable Systems and Networks (DSN-05)*, pages 216–255, Yokohama, Japan, 28 June - 01 July 2005. IEEE Computer Society Press, Los Alamitos, California.

[9] E. Totel, F. Majorczyk, and L. Me. COTS diversity based intrusion detection and application to web servers. In *Eighth International Symposium on Recent Advances in Intrusion Detection (RAID'2005)*, Seattle, Washington, 07-09 September 2005. (Preliminary proceedings.)

[10] J. H. Wensley, L. Lamport, J. Goldberg, M. W. Green, K. N. Levitt, P. M. Melliar-Smith, R. E. Shostak, and C. B. Weinstock. SIFT: Design and analysis of a fault-tolerant computer for aircraft control. *Proceedings of the IEEE*, 66(10):1240–1255, October 1978.