Active Trust Management for Adaptive Survivable Systems (ATM for ASS's)

Howard Shrobe & Jon Doyle

Artificial Intelligence Laboratory
Massachusetts Institue Of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: Computer systems are under attack as they have never been before. Every day computer systems around the world are compromised: data is stolen, security is breached, services are denied. These attacks are enabled by flaws in our current operating system, application code and middleware. It would be nice to replace all this with new software that is based on provably impenetrable trusted computing base, but of course this isn't practicle. Furthermore, even if we could afford to make such a wholesale upgrade, compromised insiders would invalidate any such guarantee of security. We instead need a way to make systems behave in a more survivable fashion and to complete their critical tasks even when compromises have happened.

Motivation: There are two ways to look at this problem of information survivability. One can either, as is traditional, try to build impenetrable walls around computer systems, or one can try to build highly adaptive computer systems. Such adaptive systems will find alternative ways of providing useful services that avoid using compromised resources as much as is possible. We believe that this unexpored approach is vital for two reasons: First, there is an enormous base of useful, installed software which is known to have vulnerabilities and new vulnerabilities are revealed more quickly than the old ones are fixed. Secondly, even if all of the holes were closed, there is still the problem that even a highly secured system can be corrupted by compromised insiders.

Previous Work: Most work that has looked at computer security has focussed on either preventing intrusions or on rapidly detecting them so as to prevent damage. To date, there has been very little work on using adaptivity as the organizing principle behind survivability.

Approach: Our project aims to build Adaptive Survivable Systems that are capable of performing their intended function even when underlying computational resources have been successfully compromised. In particular, we wish to build systems that model the trustworthiness of computational resources and that make rational choices about how best to achieve their goals in light of the risks and benefits involved in using alternative computational resources.

Our project will focus on four major topics:

- Trust Models: An Adaptive Survivable System must know what resources are trustable and for what purposes they may be trusted. This in turn depends on what components have been compromised and on the form of the compromise. Finally, this depends on what attacks have been conducted, which have succeeded, and with what intent they have been conducted. Our trust model will therefore have three levels, each with its own ontology and inference techniques. The Trustability level will center on properties of significance to applications (e.g. privacy, quality of service). The compromise level will focus on computational components that provide these properties and on the ways in which they may be compomised. The attack level will focus on the types of attacks and on how they enable compromise of critical resources.

- Perpetual Analytic Monitoring: The trust model is constructed and kept current by constant monitoring of information streams arising from multiple soruces such as intrusion detection systems and the self-monitoring of application systems. We collate and analyze these reports, looking for temporal trends that are indicative of coordinated attacks or of particular compromises. Thus, our goal is not as much to spot attacks as to assesss the degree of compromise already present. This part of our effort will be based on our MAITA monitoring system.

- Self-Adaptive Survivable Systems: Trust models influence the way a Self-Adaptivie system attempts to perform its computation. Self-Adaptive systems are structured so that each sub-task has many methods available for achieving its goal. Each of these methods requires specific types of resources and each of these resources is assessed for its trustworthiness; each method also promises a certain quality of answer. A self-adaptive system makes the rational choice of using that method which is most likely to achieve maximum net benefit. Self-adaptive systems also inform the trust model.

The goals and invariants of each computation are explicitly represented and checked as the computation proceeds. The failure of a computation to behave as expected provides evidence that the resources used by that computation have been compromised. This evidence is reported to the monitoring system and is used to help assess the degree of compromise and the trustworhiness of the resourrces.

- Rational, Trust Driven Resource Allocation. Trend detection, self-monitoring and trust assessment all consume resources which might otherwise be used by applications to perform their critical services. Dedicating too many resources to house-keeping functions would prevent the applications from rendering their functions (i.e. a self-inflicted denial of service); dedicating too few resources to the house-keeping functions necessary for an accurate trust-model can lead to the use of compromised resources in tasks for which they are not trustworthy. Similarly, application systems themselves constantly make decisions about how to achieve their goals and which resources to use. Each of these decisions can be viewed as a rational decision making problem, that is assessing how best to achieve maximum expected net benefit, given the trustability of the resources, the political situation and the likelihood of coordinated, malicious intention.

Difficulty: Each of the major components of our approach involves significant research into uncharted waters. There has to date been no systematic analysis of how trust, compromise and attack are related to one another nor are there practical ontologies of each of these areas. Our whole approach involves making rational decisions in real-time; but rational decision making is inherently unbounded. Thus we need to find more tractable, rule-based approaches to this decision making that adequately approximate the outcomes of decision theory.

Impact: This project is attempting to create an entirely new approach to the management of survivability in critical computational systems. We believe that current approaches are inherently limited and that approaches like ours are critically necessary.

Future Work: We are just at the beginning of this project and it is difficult to project future work beyond our current research plan. Our current plan involves the following efforts: We plan to construct a preliminary ontology underlying the trust model and to distribute it for discussion with other projects working on information survivability. We plan to enhance our MAITA monitoring system to understand the information provided by a variety of intrusion detection systems and by self-monitoring applications. We also plan to construct a library of ”trend-templates” that describe the temporal pattern of behavior that characterize successful attacks and compromises. We plan to develop techniques for instrumenting an application system so that it checks its own progress towards achieving its goals and generates reports in the event of failure. Finally, we plan to develop initial models for rational resource management that take into account the information in the trust model.

Research Support: This work is funded by a contract from DARPA ISO and administered by SPAWAR.