

A Strategy for Information Survivability

Joon S. Park and Judith N. Froscher
{jpark, froscher}@itd.nrl.navy.mil

Center for High Assurance Computer Systems
Naval Research Laboratory

Abstract

The need for survivability is more pressing for mission-critical systems. In this paper, we introduce a three-pronged survivability strategy for mission-critical systems: protection, detection and response, and recovery and/or reconstitution. Additionally, we provide possible technical solutions for the survivability of mission-critical systems, categorized by the different thrusts of our survivability strategy (i.e., protection, detection and response, and recovery and/or reconstitution).

Introduction

Many current mission critical systems are distributed and depend on commercially available products and services, which are heterogeneous and autonomous. No longer are they only embedded systems with well-defined boundaries under the control of a single administrative authority. Instead, these systems must withstand the documented and undocumented flaws inherent in commercial off-the-shelf (COTS) products and information operations (IO) attacks, whose perpetrators range from internet hackers to organized, well-financed terrorist groups or nation states. Traditional schemes for ensuring that they can tolerate a set of anticipated failures do not scale to meet the challenges of providing assured survivability for current distributed mission critical systems that must rely on commercial services and products.

A survivable system must enable mission critical functions to continue operation, perhaps in a degraded mode, even though the system has sustained damage from successful IO attacks or other types of failures. For these systems, it is not always possible to anticipate every type of failure because they are usually large, and attempting to predict and protect against every conceivable failure soon becomes too cumbersome and too costly. Additionally, some damage will be the result of novel, well-orchestrated, malicious attacks, which are simply beyond the abilities of most system developers to predict. Under these conditions, even correctly implemented systems do not ensure that the system is survivable. To be successful in this environment, our survivability strategy must enable dynamic recovery of critical functionality and provide intrusion tolerance. This position paper proposes a possible approach for making mission critical systems survivable and identifies some technical measures that can be used.

Survivability Strategy

The damage suffered and whether a system can recover from this damage will determine the survivability characteristics of a system. Our approach for building such systems begins with defining what survivability means for the system [KS00]. We must decide which functions are critical for the success of the mission [MEL00] and identify the dependencies and priorities among these critical functions. Then the system must be structured so that it is very difficult to attack these critical functions, but in case of attack, either the system tolerates the intrusion or it is easy to recover or reconstitute the critical functionality. Research on how such systems can be developed has only just begun [FK99]. Once this initial analysis has been done, we need a systematic approach for building a survivable system and for determining whether we have succeeded. In this paper, we identify our survivability strategy, which includes three different aspects: protection, detection and response, and recovery and/or reconstitution. These measures must work together to provide assured survivability.

First, we must attempt to protect the target system against possible damage using conventional and novel security technology, environmental measures, operational protection measures, and policies. Cryptographic technologies, access control mechanisms, and security devices can be used for this purpose. Ensuring that security patches have been installed and that approved configurations of COTS products and services are used protects the system from known cyber attacks. For example, denying the attacker intelligence about the system configuration and its assets protects the target system against a well-planned and executed cyber attack. It will also entail obtaining the various authorizations and approvals for operation.

Protection measures alone will not be sufficient for ensuring survivability since mission critical systems use COTS products and commercial services that are not under the control of the system owners. Knowing that intruders discover new and novel ways to attack systems, we must recognize that protection measures have a limited lifetime and must anticipate that some attacks will succeed. For example, we cannot guarantee that a program from a known source (e.g., checked by the digital signature) is free of malicious code or that systems, not in the same administrative domain, are configured correctly. The second aspect of our strategy is to detect anomalous behavior and respond to attacks. Each layer of the computational infrastructure must be monitored; detection events will be correlated to obtain a better understanding of the threat and failure environment. Log analyses programs, intrusion detection systems, and anti-virus programs can be used for this purpose. Intrusion tolerant techniques, such as redundant data service with voting, provide for the detection of a compromised server. If the application itself is monitored, unexpected behavior can be detected and recovery procedures can begin. When an attack is detected, a proper response must be mounted to protect system resources, such as isolating the host under attack or a partition of the system, redirecting the attacker to a honey pot, or simply alerting security managers of the impending threat. Providing a capability to strengthen protection mechanisms in response to observed threat behavior could prove to be quite effective. However, the response measures themselves may cause the application to fail by making resources unavailable.

The third aspect of the strategy is to be able to recover and/or reconstitute resources to support the survivability requirements (i.e., ensure that critical functions continue to operate) of the system. The recovery and/or reconstitution approach will rely on conventional recovery techniques (e.g., transaction management systems), redundant data sources, multiple critical paths through the distributed application,

the ability to use undamaged resources, software repositories, etc. Recovery will also depend on capturing state information as well as having an accurate picture of the computing infrastructure so that undamaged resources can be exploited to ensure that critical functions are able to continue from an undamaged state. As a result of the operational and of the IO situations, the mission itself may need to change. To adapt to a changing computing as well as operational environment, the system must be dynamic so that it can be configured differently depending on the mission status and objective. Redundancy and reconfiguration of critical resources play key roles in ensuring survivability. However, to be effective, these techniques must be applied systematically with knowledge about the extent of the damage (e.g., whether the attack has been thwarted) and the resources available for continuing operations.

Possible Techniques for Survivable Mission-Critical Systems

To develop a survivable mission-critical system, we can use a variety of survivability techniques. We classify some possible technical measures that can be used to develop survivable systems. In Table 1, these measures are organized into the categories, which correspond to the different aspects of the survivability strategy (i.e., protection, detection and response, and recovery and/or reconstitution). We can use some of these techniques to build a survivable mission-critical system.

Possible Technical Measures		
Protection	Detection & Response	Recovery
Access Control	Log Analyses	Advanced Middleware [PWS01]
Data Fragmentation and Scattering [WBS00]	Virus Check	Dynamic Reconfiguration
Dynamic IP Addresses	Host Monitoring	Redundancy
Strong Authentication	Application Monitoring	Degraded Services
PKI	Mobile Agents	Transaction Recovery
DNSSEC	Network Monitoring	State Capture
IPSec, SSL, VPN	Situation Awareness of Computing Environment	Dynamic Routing
Proxy Services	SNMPv3	Backup Server
Service Registration	IDS	Recovery Controller
Trusted OS	Knowledge Base	Mobile Agents
Authentication Server	Dynamic Routing	
Directory Server		
Firewalls		
Policy Server		

[Table 1] Possible Survivability Measures for Mission-Critical Systems

In this paper, we do not describe each technique in detail. There are probably many other techniques that are not listed in the above table.

Future Work

We have introduced a three-pronged survivability strategy: protect, detect and respond, and recover and/or reconstitute the critical functions of the system. However, the strategy does not describe an approach for determining what the optimum combination of countermeasures from each prong of the strategy is. It does not provide a methodology for deciding how to allocate limited survivability resources. A framework for deciding whether protection techniques should be extensively applied or whether it is more effective to be able to repair damage as it occurs is required. Of course, the optimum solution is a combination of countermeasures. What is needed is a comprehensive approach for reasoning about how the survivability requirements for a mission can be refined, how different countermeasures complement and depend on each other, and how to make trade-offs for the system as a whole. NRL has developed a methodology for reasoning about the comprehensive security posture of a system [PF01] by providing an assurance map, which is an outline of the assurance argument for the system. By refining the survivability objectives into survivability requirements, we believe that it is possible to develop a survivability methodology, with its roots in the survivability strategy, for analyzing the composition and dependencies of countermeasures from the three-pronged survivability strategy and for deciding whether the survivability objectives have been satisfied.

Another issue that results from the survivability strategy and intrusion/fault tolerant systems, generally, is the management complexity introduced into every level of the computing infrastructure. Survivability of these management functions presents another critical survivability requirement. Recent enterprise application integration and management middleware has become available and could possibly be quite applicable not only to the management complexity of survivable systems but also as a technology for designing systems that must satisfy strict survivability requirements. The need for research into the larger issue of how to develop and comprehensively provide assurance for survivable, mission-critical systems becomes extremely crucial as intelligence about the emerging threat environment becomes more concrete.

References

- [FK99] J. Froscher and M. Kang. *Secure, Survivable Distributed Computing*. NATO Information Systems Technology Symposium, Washington DC, 1999.
- [KS00] J. Knight and K. Sullivan. *Towards a Definition of Survivability*. Proceedings of the 3rd Information Survivability Workshop (ISW), Boston, MA, October 2000.
- [MEL00] N. Mead, R. Ellison, R. Linger, et al. *Survivability Network Analysis Method*, SEI Technical Report: CMU/SEI-00-TR-013, September 2000.
- [PF01] J. Park and J. Froscher. *Tools for Information Security Assurance Arguments*. Proceedings of the 2nd DARPA Information Survivability Conference and Exposition (DISCEX II), Anaheim, California, June 2001.
- [PWS01] P. Pal, F. Webber, R. Schantz, M. Atighetchi, and J. Loyall. *Defense-Enabling Using Advanced Middleware- An Example*. Proceedings of MILCOM 2001, McLean, VA, October 2001.
- [WBS00] J. Wylie, M. Bigrigg, J. Strunk, G. Ganger, H. Kiliccote, and P. Khosla. *Survivable Information Storage Systems*. IEEE Computer, August 2000.