

Assessing Adaptation in the Context of Security and Survivability *

Paul Rubel, Partha Pal
BBN Technologies
{prubel,ppal}@bbn.com

May 4, 2001

1 Introduction

Adaptation is not generally considered while assessing the security and survivability [1] of a system. Typical security and assurance attributes (loosely defined as the features of the system that are interesting from a security and survivability point of view) either ignore or overlook the dynamism and the ability to cope with change necessary for survival. For instance, one often cited attribute is “vulnerability.” On one hand, it has a static flavor: i.e. a system is either judged to have a certain vulnerability or not. On the other hand, it says nothing about the system’s survival: i.e. what the system does or should do, if this vulnerability is exploited. We argue that this limits the usefulness of traditional attributes in assessing the survivability characteristics of a system. We also argue that adaptation should be an important factor in the assessment of security and survivability. It complements other attributes that are already being proposed and used.

2 Problems with Measures that Focus on Static Attributes

Most current measures of a system’s security are based on attributes that focus on static aspects of the system, such as the design and implementation of the

*This work is supported by DARPA under the contract No. F30602-00-C-0172

system, the nature and number of preventive barriers to thwart the attackers, etc. To better assess the survivability of a system, these need to be augmented by measurements of the dynamic aspects of the system. By survivability of a system we mean its ability "... to fulfill its mission, in a timely manner, in the presence of attacks, failures, or accidents." [1] These events change the operating environment of the system and the availability and quality of resources that are essential for the system to run. How a system *adapts* to cope with these changes determines the likelihood of its survival. Consequently, in order to be effective and accurate, survivability metrics must also pay attention to dynamic aspects of the system such as the capability to adapt.

Static measures are also inadequate for dealing with the dynamic nature of many attacks. Unlike classic fault-tolerance, which deals with non-malicious faults, survivable systems need to be robust against malicious attackers. The nature of the attacks that a system needs to defend against will change as new attacks are created and circulated. The publishing of an implementation flaw or weakness can make a server more vulnerable overnight. The assumption that the system is secure can turn into the certainty that it is not. An adaptive survivable system takes into account that bad things may happen, and attempts to mitigate their ability to take down the whole system. We can deal with this issue by factoring the role of adaptation into our survivability metrics.

This paper is intended to point out the need for and stimulate discussion about including and evaluating the role of adaptation in survivability metrics.

3 Role of Adaption in Survival

Capabilities that are associated with static measures generally provide detection of attacks and protection from these attacks. Adaptive capability provides flexibility that can effectively fit between detection and protection, potentially complementing their deficiencies. If total protection is unachievable as a practical matter, perhaps the system can adapt itself to reduce its services to a small set of "safer" features temporarily. As a simple example, consider suspending administrator privileges or disabling write operations for a while. If detection is not certain, an adaptive system can take a posture that can be upgraded as conclusions become more certain. This is analogous to what people do in their everyday lives. They change their defensive posture when they perceive threats are more probable. They do not want to pay for vigilance all the time, but sometimes it is the correct choice.

Adaptation can significantly impact various proposed IA metrics such as

adversary work factor, effort to break systems, and expected minutes to achieve target. Without any formal and systematic way to measure adaptation these metrics become less practical and less accurate. On the other hand, there are situations where adaptation is not useful and these should impact the IA and survivability metrics accordingly. We argue that assessing the role and value of adaptation is an important and difficult task that the metrics community ought to pay attention to.

4 Issues Involved with Assessing Adaptation

One simple way to look at adaptation is to see whether or not the system has adaptive capabilities. This is easy to measure but says nothing about the utility of the adaptive behavior. A more desirable measure would indicate whether adaptation makes a system more or less survivable, which leads to a more composite view of the value of adaptation. Perhaps the utility of adaptation is measurable in the context of multiple dimensions like responsiveness, effectiveness, and the riskiness of the adaptation. Responsiveness measures how quickly a survivable system can react to a fault. Effectiveness measures the success of the adaptive response. Riskiness measures how the adaptive response to one attack affects the system's ability to react to another.

A system may adapt reactively, i.e., in response to an event or pro-actively, i.e. on its own without any external trigger. In either case, adaptation must serve a purpose, and one would want to know how well the purpose was served. This measure of effectiveness directly contributes to the systems survivability metrics. For instance, if the purpose of the adaptation was to bypass the damage an attack caused, successful adaptation would positively impact IA metrics such as adversary work factor or expected minutes to achieve target. Future adaptive decisions may also be affected by this information.

A very slow adaptive response may lack utility. By the time it is engaged the problem it is trying to deal with may already have crippled the system. Similarly, an instantaneous reaction gives the problem no time to manifest itself and stands the risk of over- or under-reacting. Between these two extremes it is interesting to try to measure the goodness of an adaptive response. A tradeoff exists between a quick response and a thorough understanding of the extent of the problem. In the limit, both seem to be poor choices. Taken together, both are desirable attributes that must be weighed depending upon the situation.

Many adaptations change the state of the system. These changes may be advantageous for one situation but may have detrimental effects in regard to

other situations. A system needs to judge whether the results of dealing with one problem will be better than being susceptible to problems arising elsewhere. Furthermore, a sophisticated attack may launch smaller sub-attacks as a feint to put the system in a vulnerable position, and may then exploit the adaptive defense capabilities in the process. This risk should also be taken into account in the assessment of adaptation.

Based on our initial investigation of using adaptive response in self-defense and intrusion tolerance[2], the challenge in assessing the role of adaptation is in identifying and understanding these complex inter-relations and trade-offs.

Things get more complicated when planned and coordinated attacks are considered. Predictable adaptations can be planned for, making them easier to overcome. The feints discussed above are much harder to execute if the attacker is unsure what response a particular attack will elicit. Randomly choosing a response means that feints become less effective. An opponent will be unable to take advantage of patterns if there are none to be found. However, the question now is how does one go about quantifying the impact of “unpredictability”?

Intuitively, the value of choosing unpredictably from a set is likely to increase if the number of choices is large. However, this does not mean that that less useful adaptations should be fielded just to be enlarge the response space. Sometimes there will only be one good response that predictably should be taken. Other times, unpredictability can even be seen within a single adaptive response. For example, a service may be moved to a random port if it comes under attack. When there are trade-offs between responses it may be best to choose one unpredictably from among the top choices rather than always choose a single response predictably.

With multiple possible responses there are often competing goals. The risks that an adaption is balancing may pull the system in different directions, each with individual strengths and weaknesses. By having choices, a system is able to rearrange the system so that its strengths are brought forward and an attacker may need to become re-oriented before resuming its activity. An effective adaptive system may keep an attacker busy guessing and re-orienting rather than further damaging the system. This again, can be incorporated in IA metrics such as adversary work factor and expected time time to achieve target.

5 Conclusion

In conclusion, trying to measure the security and survivability of a system without taking dynamic behavior and adaptation into account may be inappropriate.

A survivable system needs to be able to respond and adapt to attacks, perhaps unpredictably, as a means to improved survivability. Measuring this adaptability, or its influences, is largely overlooked. We have been experimenting with adaptation as a tool to increase survivability[3] and believe that measuring these factors will contribute towards practical improvements in survivability techniques and more survivable systems.

6 Acknowledgements

The authors would like to thank other members of the BBN staff, Ronald Watro, Franklin Webber, Richard Schantz, David Karr, Chris Jones, and Joseph Loyall, for their help in developing ideas presented here.

References

- [1] ELLISON, R., ET AL. Survivability: Protecting your critical systems. *IEEE Internet Computing* 3, 6 (Nov.-Dec. 1999), 55–63.
- [2] PAL, P., WEBBER, F., SCHANTZ, R., AND LOYALL, J. Intrusion tolerant systems. In *Proceedings of the Third IEEE Information Survivability Workshop (ISW 2000) Boston* (October 2000).
- [3] WEBBER, F., PAL, P., SCHANTZ, R., AND LOYALL, J. Defense enabled applications. In *The 2nd DARPA Information Survivability Conference and Exposition, Anaheim* (May 2001). To appear.