Three Paradigms in Computer Security

Catherine Meadows
Naval Research Laboratory
Center for High Assurance Computer Systems
Code 5543
Washington, DC 20375
meadows@itd.nrl.navy.mil

Abstract

This paper describes three paradigms in computer security in terms of how they relate to the existing infrastructure : by existing within it, replacing it, or by extending it or replacing only small portions. We identify the third as the most desirable, and discuss some of the implications of this approach.

1 Introduction

We are in the middle of an explosion of new application of computers to making communication between people and organizations faster and more convenient. Unfortunately, this burgeoning new functionality has brought with it a host of new and potential security problems. As the older approaches to security prove to be inadequate, we are actively searching for new paradigms and approaches to address the new problems we face. The thesis of this paper is that we should not only be looking at new paradigms for solutions, but to take a step back and rethink how we should be looking at the security problem as a whole. In particular, we should be looking, not only at different solutions, but the context in which they are applied.

This became very clear to me when, last year, I was invited to take part in a panel on "High Assurance Systems: the Good, the Bad, and the Ugly," in which each panelist was asked to present examples in his or hew own field in each category [8]. The "ugly" category consisted of practical but messy solutions of doubtful assurance, the "bad" category consisted of sound but impractical solutions, while the "good" consisted of solutions that were both sound and practical. As I began to answer these questions, it became clear to me that which category a solution fell into was usually not so much a result of the techniques used, as the way in which the problem itself was formulated, in other words, the kind of paradigms that were used. As a matter of fact, the three different categories tended to fall into three different paradigms that can be characterized in terms of the way in which the existing infrastructure is used. I will call these the Live With It, Replace It, and Extend It Paradigms. The Live With It paradigm uses the infrastructure much as it is and tends to produce messy and ad hoc, that is to say ugly solutions. The Replace It paradigm attempts to replace significant portions of the infrastructure, and tends to produce sound but impractical, that is to say bad solutions. The Extend It Paradigm either extends the infrastructure or replaces underutilize parts of it. It is my contention that the Extend It paradigm tends to produce good solutions. Below we consider these three paradigms in more detail.

2 The Live With It Paradigm

By the Live With It paradigm we mean the approach of applying patches to a system as it is in order to make it more secure. In this approach, the system and the environment in which it operates is taken as a given. Patches are applied without making any attempt to modify or extend the underlying structure of the system. There are numerous examples of this approach; firewalls and virus checkers are what comes most readily to mind. They provide a coarse-grained, imperfect security. Firewalls allow traffic through based on such information as unauthenticated source and destination addresses or format, not security attributes, so they may both disallow desirable traffic and allow undesirable traffic. Virus checkers check only for viruses they know about, and thus are no protection against new, unknown viruses. Moreover in many systems the virus checker itself is not protected, and so is itself vulnerable to attack.

1997 New Security Paradigms Workshop Langdale, Cumbria UK
0-89791-986--6/97/ 9

It is easy to think of more effective solutions than firewalls and virus checkers. A network with strong authentication so that it is possible to know the provenance of any message traffic would do a lot more to make a network secure than any improvement in firewall technology. Likewise, operating systems that do not allow modification of their data without permission would be more effective than virus detectors. But both solutions would involve extensive modification to the systems they were intended to protect. On the other hand, firewalls and virus checkers can be installed with minimal modifications to the systems they run on. This makes them cheap, popular alternatives that provide reasonable, if not perfect, security.

3 The Replace It Paradigm

The Replace It paradigm can be summed up in the phrase "Replace X with a secure X." The Replace It paradigm is attractive from a theoretical point of view; it allows one to design the best possible X as far as security goes. However, it overlooks the possibility of the entrenchment of X, that is that X may have a lot of loyal, or at least resigned, users, and that it may interface with a number of Y's and Z's that would now have to be made compatible with the secure X.

All of us who have worked in computer security in the Eighties have seen the Replace It paradigm in action. This is the paradigm behind the Orange Book [3], which set out a set of criteria for secure operating systems. Most existing operating systems did not satisfy these criteria, and thus would have to be replaced with new ones. Moreover, the new secure operating systems often broke interfaces with existing software, especially at the higher evaluation levels. Thus code would have to be ported to the new systems. This contributed to the cost of using the systems, and increased the time needed for development. Although these were probably not the only reasons Orange Book evaluated systems have not seen nearly as much use as was originally hoped for, they were certainly a contributing factor.

As computer systems become ever more widespread and interlinked, the Replace It paradigm becomes harder and harder to apply. Clearly, whatever security is introduced must be compatible with existing systems as much as possible, and it should be possible to introduce it without requiring people to replace too much of their existing environment.

4 The Extend It Paradigm

This brings up our last paradigm, the Extend It Paradigm. The idea behind the Extend It paradigm is to add components that extend a system's capabilities to operate securely, and to do this in such a way that very little of the system or the infrastructure supporting it has to be replaced, if any.

The examples of the Extend It paradigm that I present here all come from multilevel security, where the approach is being applied to the controlled sharing of data at different security levels. In particular, we can think of the work that has been done on replicating data from a machine operating at a low security level to a machine operating at a high level as an example of this paradigm in action. Such approaches usually require the addition of a relatively small hardware component to the system that governs the transmission of data from the low machine to the high machine. Since all this device has to do is make sure that information does not get transmitted back from high to low, it can be extremely simple, although it may be more complex if we allow some controlled information to go from high to low. Also, such use of a device like this makes few requirements on the systems that are being hooked up; instead it concentrates on the communications between them. This greatly reduces the amount of interfaces that are broken and the amount of software that must be replaced. A prominent example of this type of approach is the SINTRA database management system [7], but the approach is also being applied to other types of multilevel systems [4], and a sizable body of work exists on producing reliable one-way flow devices [6, 5, 2, 10].

This approach of concentrating the security of a system in a small component that requires minimal changes to accommodate it is also being applied to other, similar, problems in multilevel security. For example, the Starlight Interactive Link [1] provides a multilevel windowing capability by linking a workstation to classified and unclassified networks and allowing the user to switch from one to another. The Interactive Link consists of three components: the Link itself, a trusted display which lets users know the current security level of the Link, and software which must be installed in the single-level systems in order for them to communicate with the Link. Only the Link and the display are trusted, and the only change required to the communicating single-level systems is installation of the software managing the communication with the Link.

5 Open Questions

We've presented a tentative outline of three paradigms in computer security. The goal of this outline is to help to organize our thinking about computer security research so that we can better identify the solutions that

are most effective. But there are still a number of questions to answer and points to clarify.

One of the most important points to make is that the identification of a solution as belonging to one of our three categories depends very much upon the context in which it is applied. For example, the examples I described as typifying the Extend It paradigm both break the usual assumption that communication is always two-way. This is not a serious problem when we consider the intended application; since the system-high systems were not communicating with each other to begin with, we are not replacing an existing two-way communication links, but instead extending the systems' capabilities by making one-way communication possible. However, in a situation in which two-way communication was in common use and even relied upon, the same techniques would probably fall into the less desirable Replace It category.

For another example of the way in which context affects the category in which a solution fall, consider the case of intrusion detection. Like firewalls and virus detectors, intrusion detectors can be though of as add-on fixes - instead of going to the source of the problem and structuring a system so that attacks are prevented, they detect attacks after they occur or while they are occurring. This would seem to put them in the Live With It category. But a good instrusion detection system can also make some serious requirements upon the systems upon which they are running: protection of the intrusion detector itself from attacks, security relevant auditing, and a security policy that is well-defined enough so that it is possible to tell attacks from legitimate use of the system. If meeting these requirements requires a massive overhaul of existing systems, then intrusion detection might fall into the Replace It category. On the other hand, if it is possible to have the system meet these requirements with minimal changes. then intrusion detection might fall into the Extend It category. Again, it will be the context that determines this as much as the nature of the solution itself.

Finally, we note that the examples of Extend It all apply to multilevel security. What kinds of solutions can we come up with for nonmilitary applications? It appears that one-way replication also has uses in the commercial sector; for example Microsoft uses one-way replication to increase the effectiveness of its firewalls, and recommends this approach in its current manuals [9]. It would also seem that, given the reliance upon small hardware components that I've described, it might seem that something like secure co-processors would also be a likely candidate for commercial security. But as a matter of fact, secure co-processors have yet to see widespread use. Why is this, and does the reason have any relation to the three paradigms that I have outlined? If secure hardware is not the answer for commercial applications, what would be?

6 Conclusion

We have characterized three different approaches to computer security in terms of their relation to existing infrastructure. The first attempts to build security mechanisms on top of the infrastructure, resulting in solutions that are relatively easy to introduce but that provide a rather limited degree of protection and little assurance. The second attempts to provide security by replacing a large part of the infrastructure; here the degree of protection is much greater, but the cost of introducing it is prohibitive. The third looks for places in which the infrastructure is nonexistent or underutilized, and builds in the security mechanisms by replacing or extending these parts. We consider this the most promising approach, but some thought must be taken in applying it.

7 Acknowedgements

I am grateful to Carl Landwehr and Judy Froscher for helpful conversations on the application of one-way communication to security, and to John McDermott for pointing out to me Microsoft's use of one-way communication. I am also grateful the participants in the New Security Paradigms Workshop for their helpful comments.

References

[1] M. Anderson, C. Nort, J. Griffin, R. Milner, J. Yesberg, and K. Yiu. Starlight: Interactive Link. In *Proceedings of the Twelfth Annual Computer Security Applications Conference*, pages 55–63. IEEE Computer Society Press, December 1996.

[2] J. Davidson. Asymmetric isolation. In *Proceedings of the 12th Annual Computer Security Applications Conference*, pages 44–54. IEEE Computer Society Press, December 1996.

[3] Department of Defense Computer Security Center. *Department of Defense Trusted System Evaluation Criteria*, August 1983. CSC-STD-001-83.

[4] J. N. Froscher, D. M. Goldschlag, M. H. Kang, C. E. Landwehr, A. P. Moore, I. S. Moskowitz, and C. N. Payne. Improving inter-enclave information flow for a secure strike planning operation. In *Proceedings of the 11th Annual Computer Security Applications Conference*, December 1995.

[5] David M. Goldschlag. Several secure store and forward devices. In *Proceedings of the Third ACM Conference on Computer and Communication Security*. ACM, March 1996.

[6] M. Kang and I. Moskowitz. A pump for rapid, reliable, secure communication. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, pages 119–129. ACM Press, November 1993.

[7] M. H. Kang, J. N. Froscher, J. McDermott, O. Costich, and R. Peyton. Achieving database security through data replication: The SINTRA prototype. In *Proc. 17th National Computer Security Conference*, pages 77–87, Baltimore, MD, Sept. 1994.

[8] Catherine Meadows. Computer security: The good, the bad, and the ugly. In *Proceedings of the 1996 High Assurance Systems Engineering Workshop*, pages 52–54. IEEE Computer Society Press, 1996.

[9] Microsoft Corporation. *Microsoft WindowsNT Server Internet Guide (For Windows NT Server Version 4.0)*, 1996. pp. 63-64.

[10] N. Ogurtsev, H. Orman, R. Schroeppel, S. O'Malley, and O. Spatschek. Experimental results of covert channel elimination in one-way communication systems. In *Proceedings of NDSS97*, 1997. To appear.