2.0 Defensive Information Warfare in the 21st Century

Dr. Larry Druffel

This paper assumes, and provides supporting motivation for, the proposition that an AF goal should be to achieve information dominance to enable the execution of its missions through unconstrained, but protected, use of cyberspace, including systems the AF does not control.

The author acknowledges the contribution of Tom Longstaff of the Software Engineering Institute to the paper, in particular to the discussion of malicious code and bounded vs. unbounded considerations.

The Importance of Protecting Cyberspace[18]

Cyberspace is essential to AF mission execution.

Successful execution of all AF missions will depend on AF ability to exploit information. Consistent with the trends in our society, use of information and the supporting information systems technology by the Air Force has become ubiquitous. The success of both combat systems and support systems relies on access and ability to process information. These capabilities, including people, information and supporting systems are geographically and organizationally distributed, reflecting the AF global mission. The Air Force depends on cyberspace. The information available to the commander is not local and is often not under his direct control, but it is accessible. This trend will continue or expand as the Air Force seeks to reduce its decision time to operate within an enemy decision cycle in order to achieve its goal of information dominance.

AF systems will include commercial products and use commercial infrastructure that the AF does not control.

Although there will always be a need for military unique capabilities, the AF simply cannot avoid using commercial products to ensure that the best technology is available in a timely and affordable manner. Likewise, although the AF will own and control some infrastructure, its anywhere/anytime mission requires that the AF also use commercial infrastructure such as communications, networks and information services that might be available.

Air Force must protect its cyberspace.

With this increasing dependence on information and on commercial applications and infrastructure, it is increasingly important that the AF protect its cyberspace. This challenge is much broader than the normal security considerations. Protection must not only include the AF assets, but also its access to commercial infrastructure and in some cases protect the infrastructure itself.

18. *Cyberspace*, "that consensually imagined universe where information reigns supreme," is synonymous with the phrase *infosphere* in our Panel's report.

Air Force must use commercial solutions for protection, but must not depend on those solutions solely.

Commercial owners of information will develop technology solutions to protect their interests. The AF should use those products and approaches but, since the risk tradeoffs may be different, should not depend on them solely. The AF should lead in the development and application of technology for protection. The Advanced Research Projects Agency has an exciting vision and is committing substantial funding in this area. The AF should work with ARPA to participate in the technology development and lead in the application of that technology.

While it may be desirable to embed protection mechanisms into systems, from both a communications perspective and a commercial products perspective, protection must be considered an overlay in much the same way that STU-3 is an overlay on the telephone system. The important point is that the protection mechanisms allow integration of available software products and access to local communications.

The AF cannot assume technical superiority.

Clearly the AF can depend on access to the best technology. However, so can potential adversaries. For the foreseeable future, the AF will rely more heavily on information than potential enemies will. In addition, the time it takes to acquire and field new technologies and train people in their use will put the AF at a disadvantage against a non-traditional adversary, terrorist, or protester. This is particularly true with respect to a technology that is changing as rapidly as information technology.

For less than a million dollars, a drug lord or terrorist group can acquire highly competent people, trained at the best US. universities, and equip them with the very latest technology. A small team of less than a dozen such people can easily conduct attacks on AF cyberspace, from outside the US., and often at no personal (physical, legal or social) risk. The average AF user will be powerless against such attacks and may not even recognize (s)he is being attacked. This does not imply that the AF will be at a net disadvantage with respect to all potential adversaries. On balance, we have more experience and greater access to technology. Our defensive systems can be better than their defensive systems and our offensive systems can be better than their offensive systems. But we must plan for the situation in which the systems we want to protect are not as sophisticated as an adversary's offensive capability. The implication of this is that the AF must plan for the possibility that an adversary can get inside our Observe, Orient, Decide, Act (OODA) loop.

The AF should train and equip information warriors.

Consequently, the AF should be prepared to train and equip highly competent teams of information warriors to monitor, detect and thwart such attacks. A sophisticated attack on AF systems will involve numerous preliminary probes to find vulnerabilities, test the ability to modify, leave backdoor traps for later entry, and assess the ability and kinds of responses to such actions. A trained team of information warriors with sophisticated monitoring capabilities could detect such probes, recognize patterns and help users take precautionary and preemptive defensive actions.

An early capability is evolving within the Internet. Small teams of highly trained people, called incident response teams, one variant of which is a Computer Emergency Response Teams (CERT), support various networks. The AF has a CERT team. In general these teams are more like volunteer fireman in that they provide a response service but have no authority to take preventive action.

Training specialists is not enough. Any AF person may be involved in the information war. As a user of information systems, each AF person is a potential target and must be trained to understand her/his role in protecting cyberspace.

The AF must make core systems impenetrable.

The AF must also analyze its core systems (such as those that provide coordinates to weapons) and ensure that they are impenetrable. Protection schemes involve assessment of risk. They often include assessment of the cost of penetration vs. the damage to the AF. For those core systems, such as the link between an aircraft and a weapon, AF will want to ensure that the communications is assured and the system fully protected.

In IW, the AF must be biased toward protection.

Since the AF places greater reliance on information technology than any potential adversary for the foreseeable future, information dominance will depend on the ability to protect its systems. Also, since the AF will make increasing use of commercial systems upon which much of the US infrastructure is also based, vulnerabilities in AF systems are likely to also be present in US communications, power, medical, financial and other systems. Potential adversaries will also be using the same technology and products. When vulnerabilities are discovered, there will be a natural tendency to keep that knowledge for offensive exploitation purposes. The process for making such decisions must include advocates for protection and the process must be biased toward protection.

The AF should consider providing leadership in protecting cyberspace.

No US Agency has clearly established a capability to protect the US infrastructure that increasingly depends on cyberspace. Many of the systems our society depends on are vulnerable. Our power, transportation, financial, airline control, individual airplane safety, and healthcare, to name just a few, are all systems that could be attacked.

The AF must develop a capability to protect its systems and its use of cyberspace. In doing so, it will be developing the technology and a capability to accept a broader role as a US champion in much the same way the AF has established itself as a champion for air and space.

The Dimensions of Cyberspace Protection

Unconstrained use of cyberspace implies protection in multiple dimensions.

These dimension include encryption, protection from malicious code such as viruses and worms, use of agents that cannot be corrupted as double agents, protection from intrusion, detection of viruses and trojan horses in incorporated software, operating system and infrastructure control, and mechanisms for recovery and alternate operation in the face of failure, disruption, and denial.

Both data and control must be protected

Data must be protected both from unauthorized disclosure and from corruption or loss.

Data is a sequence of bits to which meaning may be assigned. Data is generally well understood as consisting of the elements that are input to computers and the output produced by them. It is obvious that data must be protected from disclosure to an enemy. It is also important that it be protected from corruption or loss. Data represents both an enormous investment and an important resource that reflects the current state of many systems.

Control must be protected both from unauthorized users and from automated attacks.

Control refers to the process (computer program) that has execution authority of a computer system. Although many programs may be resident on a system and be invoked from a variety of levels, only one program has control of the execution of a given processor at any one time. Normally, the issues of control are managed by the operating system for resource allocation and performance reasons. For an intruding process to have an effect on a computer system, it must gain control. There are a variety of ways in which control may be passed to a program, including human action, local system action, or action by a remote process. Any of these methods may be used for legitimate purposes, but they may also be the means of relinquishing control to malicious code.

The distinction between data and control is becoming less clear in many modern systems.

Most traditional applications treat data only as input to the program. In such applications, data could be corrupted to produce inaccurate results, even crash the system if certain circumstances or sequences were not properly accommodated, but could not assume control. Increasingly, applications accommodate sequences of code as part of the data. This approach offers additional power, but it also introduces additional sources of risk. Programs such as Powerpoint, a popular application for creating viewgraph presentations, accepts code as input and will pass control to that code. In this way, malicious code can be introduced through data, providing yet another reason why data must be protected.

Another example of this notion of data being interpreted by the client is illustrated in the World Wide Web. The World Wide Web is a collection of programs that operate on files available on the Internet. Through a convenient point and click user interface, a human may look for information. The user interface is an interpretive browse. When the user clicks on a link, the page pointed to by that link is brought over to the user's machine. The interface program then interprets the information on that page. It may be asked to perform a variety of functions including displaying graphics. It is interpreting commands to run specific programs. The data could be structured to change the programs to be run and thereby have an unintended effect.

The World Wide Web also offers another kind of opportunity for confounding the user. Organizations place information that they wish to make available in a specified format. Each item for which they wish to make additional information available is provided a pointer. This pointer may be to similarly formatted data on the same computer or may be to another computer. If a malicious user gains unauthorized access to these files, (s)he can change the links and the unsuspecting user would follow the wrong trail - a wild goose chase through cyberspace.

Data and control must be protected in bounded and unbounded systems

A bounded system has common administrative control.

In a bounded system, the Air Force (or other DoD agent) has either central or distributed authority over all components of the system, and can conform to a defined set of policies and procedures. Isolated classified networks and layered encrypted networks are examples of bounded networks.

An unbounded system has no common administrative control.

The control and data takes place in an environment where the interconnected systems and sites are not under a common administrative control. In this case, the boundaries of such a system cannot be defined, since there is no central or distributed authority that "keeps track" of all the components of the unbounded system. The Internet—a non-hierarchical network of systems under local administrative control only—is one current example of an unbounded system.

At a bare minimum, a bounded system can be understood and all of its various parts identified. In an unbounded system, the various parts cannot be identified, their actions cannot be predicted, and there is no unified administrative control over the parts of the system. There are conventions that allow the parts of the Internet to work together, but there is no global administrative control to assure that these parts are behaving according to these conventions.

The architecture of secure, bounded systems is built upon the notion of a security policy with the existence and enforcement, or lack thereof, imposed by the exercise of administrative control. In contrast, an unbounded system can impose no global security policy. For instance, on the Internet today the backbone architecture is independent of security policy considerations because there is no global administrative control on an unbounded system.

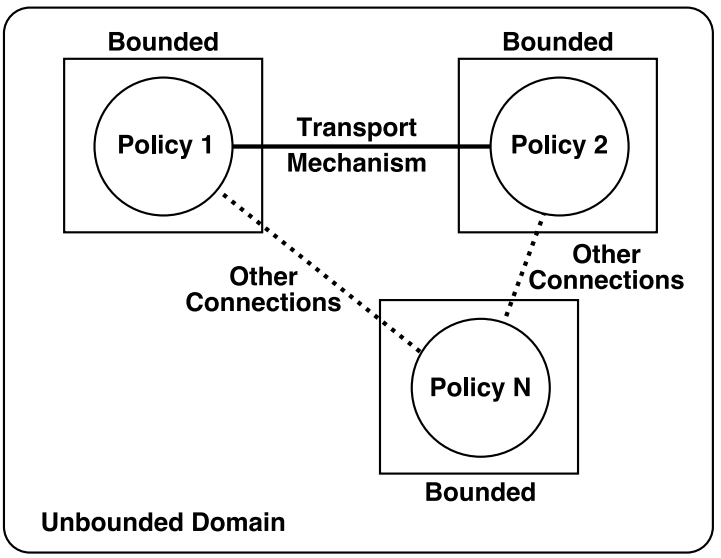

Figure 10. An unbounded domain viewed as a collection of bounded systems

Figure 10 illustrates an unbounded domain consisting of a collection of bounded systems, where each bounded system is under separate administrative control. If each bounded system were completely disconnected from the other systems, it would be possible to fully characterize the security state of each system. Note that the notion of boundedness does not make any presumption about geographic constraints.

It is not sufficient to focus only on the bounded systems. Protection for unbounded systems is also important because the AF needs the information derived from unbounded systems and needs to use unbounded systems as part of the infrastructure.

Threats and Countermeasures

There are a variety of threats and each threat requires its own countermeasure.

Electronic Warfare provides a more familiar analogy**.** There are a variety of threats and countermeasures. The data and control threats are different in bounded and unbounded systems. Likewise the responses to these threats are different. In addition, the maturity of the technology available to respond varies considerably. To a large extent the DoD has invested heavily in the technology and mechanisms to protect data, but has done very little to deal with control.

It is not possible to project all of the potential future threats. However, it is reasonable to begin by understanding the current threats and the countermeasures for those threats. Then as each new technology is introduced, the potential vulnerabilities of that technology can be assessed in the context of a deep understanding of these threats. While it may seem that the description of the threats and responses is irrelevant for the 2025 projection, threats introduced into new technology often follow the concepts used in previous technologies. There is some evidence that this is the case even though the mechanisms are often very different. The AF should not limit its thinking to a simple projection of known threats, but that is certainly a reasonable place to start. Table 2 captures many of the known or predictable threats and their possible countermeasures.

	Threat	Countermeasure
Data/bounded	Disclosure	Data encryption Access control
	Loss of Integrity	Crypto Checksums
Data/unbounded	Disclosure	Authorization Authentication
	Disclosure in Transit	Data Encryption
	Integrity	Data Encryption
	Traffic Analysis	New Future Technology
Control/bounded	Trojan Horse	Strong Policy & Procedure
	Viruses	Limited Detection Prevention
	Internal excess of auth	Actg. & logging
Control/unbounded	Worms	Limited Detection Prevention
	Corrupted Agents	Docking Protocols
	Intrusions	User Proxy Firewalls

Table 2. Threats and Countermeasures

Data Threats - Bounded Systems

Disclosure

One of the principal reasons for creating a bounded system is to protect the data from disclosure to unauthorized people. This the most traditional concern of military intelligence and predates the use of computers. The typical response is to encrypt the data when it is stored. Unauthorized access to the data would be useless without the ability to decode.

Current capabilities for the prevention of disclosure include the use of both symmetric or asymmetric key encryption. The primary measure of the strength of encryption is in the time it takes to decrypt information without knowledge of the key. This time is related to the available performance and availability of computers used to "guess" or mathematically deduce the key. The constant increase in speed and availability of computers results in longer and more complicated key and encryption algorithms.

In the next 25-30 years it is likely that strong encryption algorithms will be available to the general population world-wide and it is equally likely that the ability to break current encryption techniques will be commonplace. At the same time, encryption technology is a commonly used technology to layer a bounded network on a larger unbounded network. Thus the government, and the AF in particular, should continue to develop strong encryption techniques for specific military use.

A second means of protecting data within a bounded system is by controlling access to the system so that only those who are authorized may access the data. These methods include physical identification schemes, passwords, and technology such as personal identification cards (like credit cards). This has traditionally depended on people and been labor intensive.

The technology for verifying that an individual is who (s)he claims to be is advancing rapidly. By 2025, a combination of voiceprint, use of uniquely coded identification cards, and possibly even human chemical analyses will be available to ensure that the person is who (s)he claims to be. It will not eliminate detection of an authorized user operating under duress or otherwise being "turned" to operate for contrary purposes.

Loss of Integrity.

A second consideration for data is that it not be changed by an unauthorized agent or process. In addition to preventing unauthorized access, techniques such as cryptographic checksums enable detection of changed data. They are also useful in the face of certain errors that may not be caused by a malicious actor.

Integrity has had little comparative attention as much of the research in security has focused on the prevention of disclosure. However, integrity will become more important to the AF mission as interactions with contractors and vendors migrates to the public-access networks. In this instance, the protection and integrity controls for all members of the public also benefits the AF. Older integrity models such the Biba integrity model and others may be adequate for use in bounded systems (Biba, K.J., 1975, "Integrity considerations for secure computer systems", Report MTR 3153, MITRE Corp.) The primary need for the development of integrity models is in unbounded systems and to protect against internal, unauthorized modification of critical data.

Data Threats - Unbounded Systems

Disclosure.

Disclosure of information to unauthorized parties is a distinct threat in unbounded systems. The threat of disclosure is similar to the case of bounded networks, but requires additional countermeasures as all end points may not be within a single administrative control.

An effective countermeasure is authentication that a packet appearing to be coming from a trusted source actually is doing so. In the link between trusted and untrusted domain, the AF should never let a bounded service be controlled by a program in an unbounded system.

Currently it is very difficult to prevent disclosure on unbounded networks. What is required is a national or international infrastructure that will allow third-party authentication and key management between parties. This would allow the AF and vendors, contractors, or private citizens to communicate using strong authentication and cryptographic protocols without prior arrangements. The use of ad-hoc technologies and partial solutions is the current state-of-the-art practice.

Unfortunately, the solution to this problem goes beyond technology. To promote the use of cryptographic technology throughout an unbounded system requires consensus with the method, strength, and exportability of the technology to be employed. Once agreed upon, the infrastructure must be funded and created to support the networks.

By 2025 it is likely that an infrastructure of some type will exist world-wide and it will be the challenge of the AF to work within this infrastructure as it evolves to effectively and securely communicate within the unbounded networks.

Disclosure during transit.

Disclosure of data while in transit from one bounded system to another bounded system through an unbounded system is at risk of disclosure. Such data is easily protected by encryption. Likewise integrity of the data can be protected by the same encryption techniques. See above for description of the appropriate technologies.

Traffic Analysis.

In traditional message-based communications systems, the Air Force has tried to block the inference of pending action by the analysis of traffic by an adversary. The nature of computer based networks with packet switching, reduces the vulnerability to such analysis. However, traffic analysis can be used to determine location of specific items. The traditional response of flooding a channel with artificial traffic is not an effective one because it only uses up available bandwidth and does not impede the analysis.

The prevention of traffic analysis is always at odds with the performance requirements of networks. Over the next 25-30 years it is likely that the use and prevention of traffic analysis will be addressed in the variety of new communication technologies that will arise during this time. The important point here is to address the security concerns (including traffic analysis) of any new communication technology prior to wide-spread deployment and use by the AF.

Control Threats

Threats to computer systems are based on system vulnerabilities.

Most systems are designed to perform desirable actions and not permit undesirable actions. Unfortunately, most systems have weaknesses that can be exploited to violate the system's intended behavior. These vulnerabilities may be exploited by direct human guided action or by programs, which are called malicious code.

Generally, control threats may be countered either by taking countermeasures into the initial engineering of the systems or by employing countermeasures upon the delivered and deployed systems. Currently, the development of systems employing these countermeasures have been guided by the DoD Orange Book requirements for multi-level secure systems. For Commercial Off the Shelf (COTS) and general-purpose systems, there has been little advance in trustworthy engineering, especially for systems designed to be deployed in an unbounded domain.

As the integration of COTS systems continues to dominate the character of AF systems over the next 25-30 years, the need for security engineering in these products will become even more critical. To accomplish this, it will be necessary to invest in new security engineering models, techniques, and products that take into account the unbounded network environment.

The terms defined in the following paragraphs offer one model of the types of threats. There are other models. These definitions are generally consistent with "Computers at Risk," a report of the National Academy.

Malicious code is a code sequence (program) which is not intended to be part of the operational system and does damage when it executes.

Malicious code is distinguished from erroneous code "bugs" that may be in a system. Bugs are code segments that were placed in the system intentionally to perform some function but, through some error on the part of the programmer, perform an unintended action. Malicious code on the other hand is inserted into the system either when it is built or inserted after it is put into operation for the express purpose of causing damage. There are a number of ways that damage can be caused: corruption of data, modification of action, creation of a vulnerability that can be exploited later, and crashing the system.

Protection from malicious code requires detection. Since detection of malicious code, in general is undecidable, cure is difficult to impossible. In practice, once a malicious code sequence is known, it can be detected. Once detected, it can be countered. In addition, protection without detection is possible using mechanisms such as limiting transitivity of trusted programs. (A transitivity limit of distance one would imply trusting a program but not trusting a program that had been modified by another program.)[19]

Current research in malicious code has focused on viruses and micro-computer platforms where there has traditionally been a lack of traditional security protections built into the operating system. While this will likely remain a threat for some time, before the year 2025 it is likely that

19. Cohen, F.B. "Defense-in-Depth against computer viruses." Computers and Security, Vol. II, No. 6, Oct 1992, pp 363-379.

all computer systems will employ multi-process and powerful computing techniques that will necessitate the use of security features to assure proper functionality.

However, the increased use of networked and distributed techniques will likely spur the development of distributed malicious code that will be difficult to counter using traditional host-based techniques. The current emphasis on firewall technology for network-based threats will not be adequate for many threats that will arise in a fully distributed environment. As a result, new protection technologies will have to be developed to maintain the security of a wide-spread distributed network.

One line of research that appears promising is the notion of mediators. The current line of research is to use an intelligent agent to intercept data base queries and to intercept the response. The agent would apply certain rules to filter the request and the response. If both were satisfied by the rules, the query would be completed. If it does not pass all the rules, it is routed to a human. The extension to full security would enable the evolution of a rich set of rules for managing the interface between bounded and unbounded systems.

Control Threats - Bounded Systems

A trojan horse is a program whose execution causes undesired side effects, usually unanticipated by the user.

A trojan horse is usually hidden within a larger program whose execution performs normally. A trojan horse is passive until its execution is triggered. When it executes, it can perform such undesirable actions as: disclose information to the outside, destroy data, or introduce a vulnerability into the system.

Protection against trojan horses is best provided by strong policy and procedures. A trojan horse must be inserted, either manually or automatically. Manual and automated policies for control of code, including change control, are essential to preventing introduction of trojan horses. Automated techniques include longitudinal algorithms, such as checksums, for detecting changes.

A virus is a self replicating trojan horse.

The biological analysis is appropriate - a virus can infect other programs. The distinguishing characteristic is that a virus is a trojan horse that copies itself, often attaching the copy to another program. If the virus propagates fast enough, it can have the effect of using up all the available processing time and clogging a network in addition to any other damage it causes.

Once a particular virus is known, its presence can generally be detected. In addition, limited detection is possible based on a virus' characteristic of replicating itself. Limited protection without detection is feasible by preventing specific actions from certain classes of programs.

Internal excess of authority is the assumption of system privileges by a program in excess of its rights.

Operating systems normally grant levels of privilege to programs. At the basic level, an operating system reserves for itself certain root privileges such as writing to certain areas of memory and controlling tables that establish privilege levels for applications programs. If a program attempts to accomplish some task such as writing to an unauthorized area of memory,

the operating system will normally block the action. A sophisticated user (or program) can often exploit an error in the operating system code to change the privileges allotted to it and thereby cause a variety of damages, including taking control from the operating system and even closing the system down. Note that this action is dependent on the existence of some vulnerability.

Control Threats - Unbounded Systems

A worm is a program that distributes itself in multiple copies within a system or across a distributed system.

A worm is much more autonomous than a trojan horse. It exploits system vulnerabilities to gain access to distributed systems. Whereas a trojan horse is passive until control is passed to it, a worm, once invoked, distributes itself by the positive action of exploiting a vulnerability in a target system, and passing control to the newly distributed worm.

An agent is a program that performs a specialized function such as monitoring activities, filtering data or seeking out data.

An agent is often considered benign and therefore trustworthy. Although some agents perform their function within the host machine, an agent may be "sent out" to perform its function on other machines. Unfortunately, a well intended agent that is sent out, can be corrupted so that if it should return, it can cause damage. This is a difficult situation since the system to which the agent is returning can no longer trust it.

Agents are powerful mechanisms but also introduce another type of vulnerability that has not been characterized or studied to the best of our knowledge. There is a current line of research to develop docking protocols which will require agents to "register" with the system to which they are visiting. Combined with encryption techniques, agents can probably be made safe for use both within bounded systems and for use between bounded systems across unbounded space.

Intrusions are unauthorized, human-directed access to computers.

Such unauthorized access is the result of some violation of policy, whether by an insider or an outsider. It may be the violation of an access policy or of a modification policy. It may be the violation of a human directed policy or of a policy that is controlled through automation.

When a human user interacts with a computer, (s)he is interacting with a program. The user may perform only those actions that the program allows. The program with which the user interacts may be on top of several layers of other programs. Through experimentation, the user may find that if (s)he provides some unexpected input, (s)he is passed through to some other program, even the operating system. (S)he may now perform whatever actions that program will allow, including access to data or modification of tables.

Malicious Code Summary

Figure 11 shows how these threats relate. This entire field is relatively new so that new types of threats may be developed which will require new characterization.

Within the next 25-30 years it is likely that the trend will move from host-based attacks to autonomous agent-like attacks that take advantage of the interconnection of systems with no

common administrative control. To counter these new threats, the AF will need to take a leadership role in prevention, detection, and recovery from automated network attacks.

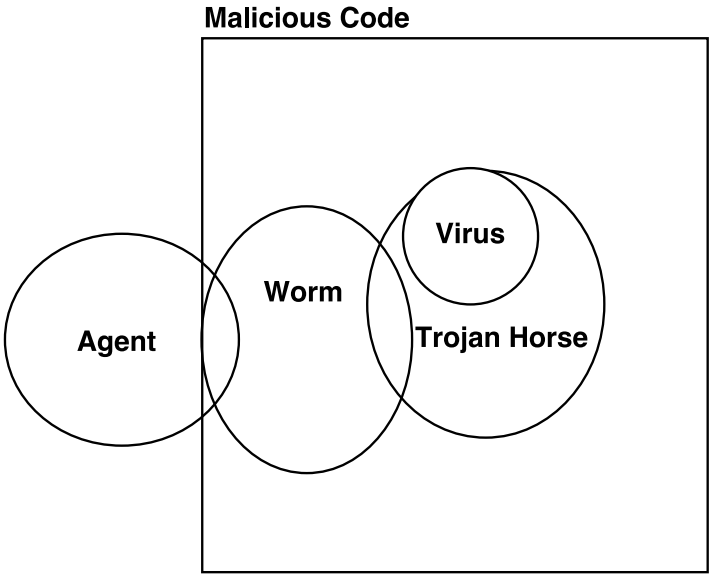

Figure 11. Relationships among threats

One trend that is likely to continue through the year 2025 is the expansion of the use of these threats by less traditional adversaries. For example, students may protest military operations through the disruption of distributed networks rather than a physical march outside an AF base. Likewise, a disgruntled employee can cause considerable disruption. The fact that the world will have access to the shared infrastructure means that an individual with motivation may be able to disrupt that infrastructure costing the AF significant time, effort, and perhaps even capability to successfully execute its mission. In addition, the expansion of the network technologies has led to more anonymous access such that the risk to an individual in performing these acts is minimized.

Recovery Technology

Despite the best protection mechanisms in automated systems, the potential always exists that some failure or disruption might occur. A single processor or communication link might fail or be denied, leaving no alternative for that element. However, other kinds of failure or disruption permit recovery.

In many cases involving human users of computing systems, the system gets into a state that is not in sync with what the user believes it is in. This can happen when a sequence of actions in which the human develops in her/his mind a model of what the computer is doing that is different from what the computer is actually doing. It can also happen due to some outside disruption such as a temporary or intermittent communications failure, action by a third party, or simply a mistake.

One example of this situation exists today. A PC user may have invoked a feature to access a remote UNIX system to acquire information via a database application. When a confusing sequence of characters appears on the screen, the user must determine which of the various systems in operation created the characters. The recovery action depends on whether it is the database application, UNIX, the communications handler, the operating system on the PC or the software controlling the remote access. A knowledgeable user can normally understand which system is creating the characters and respond appropriately to recover. As more effective user interfaces hide the underlying infrastructure software and less computer sophisticated users become more prevalent, users will be unable to recover from situations that should allow recovery.

A realistic example was posed during the AF/SAB study on integrated avionics. Using traditional federated systems, the pilot is able to build up a situational awareness based on the independent readings of a variety of sensors. When (s)he gets into overload, (s)he can focus attention on fewer sensor readings for which (s)he establishes priority. When the integrated systems begin to present a situational awareness to help with the overload, the system may present a state that is different from what (s)he expects because it has reached a state that is different from the mental model (s)he has developed. One easy remedy in that situation is to give the pilot a control which allows her/him to reset to some known earlier state even though that might be less complete than might be available.

It is essential that the AF pursue techniques that enable a user to recover from these situations. There are a number of well understood techniques and some research in this area. As commercial systems involve greater requirements for recovery, additional techniques will become available. However, the AF must ensure that available techniques are designed into the systems because they are generally not appliqués that can be added as an afterthought.

Other Threats not Covered

Denial of Access

This paper does not address denial of access or disruption through such techniques as jamming, flooding the network and physical interference such as cutting communications lines.

Detection in Communications Link

This paper does not address vulnerabilities introduced by various communications links. For instance, cellular phones use a layer of protocols that communicate, in the clear, information about the user of the phone.

Embedded Trojan Horse

With the increasing reliance on software for supporting activities such as design, the potential exists for an adversary or other agent wishing to disrupt AF capabilities, to embed a trojan horse into software or even a chip that is used in a computer aided design system. As with the general case of a trojan horse, detection is not possible for an arbitrary segment of code.

This vulnerability will be exacerbated by increased use of commercial products (COTS) in defense systems. While this is a necessary and desirable trend, it means that AF will be incorporating software that was not developed under its control. The fact that the AF uses the software establishes the vendor as a target for those who might like to install a trojan horse. The

AF will need to work closely with those vendors and in parallel will need to develop techniques for checking the software.

Summary

This paper began with the assumption that an AF goal should be to achieve information dominance to enable the execution of its missions through unconstrained, but protected, use of cyberspace, including systems the AF does not control.

To achieve this goal, the paper draws the following conclusions:

- AF systems will include commercial products and use commercial infrastructure that the AF does not control.
- Successful execution of all AF missions will depend on ability to exploit information.
- AF must protect its cyberspace using commercial solutions where appropriate, but must not depend on those solutions solely.
- The technology will exist for AF to achieve the goal, but existence of the technology is not sufficient.
  - The technology must be used
  - Employment and training will be critical
- New technologies must be introduced aggressively but each new technology must be analyzed to understand and protect against the vulnerabilities it introduces.
- The AF cannot achieve information dominance without a preeminent ability to protect its information systems.
- The AF cannot expect to have technical superiority of its defensive IW systems with respect to the offensive IW capability of every adversary.
- The AF needs to train and nurture Information Warriors.
- The AF should train and equip to (a) monitor cyberspace activity that poses a potential threat to AF systems and (b) thwart a broad range of attacks.
- When an offensive activity identifies a vulnerability, a mature process must be in place to communicate that vulnerability to those engaged in protection.
- Since the US places greater reliance on information than potential adversaries, the US (both defense and civil) systems are at risk.
- No US Government Agency has established a credible capability to protect the US against hostile activity in cyberspace. The AF should consider filling this void and be an advocate for cyberspace defense at the same level it has for air and space.
- The AF should use its technology and capability to monitor cyberspace activity and thwart a broad range of attacks to protect US cyberspace interests.

Research Threads

The paper identifies a number of technology areas that the AF should pursue. These technologies must be pursued by monitoring and motivating commercial developments and standards, as well as pursuing defense unique capabilities:

- Encryption - including key agile technology
- Monitoring and intrusion detection techniques
- Techniques for designing trustworthy software
- Technology for detecting malicious code
- Firewall and mediator technology
- Docking protocols for accepting and managing software agents
- Recovery and resyncing techniques

Potential Payoff

The payoff to the AF includes:

- Ability to achieve information dominance of the battlefield
- Leadership in a critical area of defense to which the US is enormously vulnerable and which is devoid of champions

Acknowledgment

Dr. Tom Longstaff from the Trustworthy Systems Program at the Software Engineering Institute made several technical contributions to this paper, including the notions of bounded and unbounded systems and the characterization of malicious code.