



Dr. Anup Ghosh
Program Manager
Advanced Technology Office

Defending Warfighter Networks

Imagine, for a moment, a network that behaves like a living organism: It is fully cognizant of its environment. It can recognize attacks and failures and adapt to its new environment. It can monitor the behavior of its users and its software and detect when it is being misused. It knows how to repair its own faults. And, above all else, it can defend itself, allowing it to provide sustained service even during continuous attacks.

This is the level of sophistication in network design we must achieve to realize DoD's goals for network centric warfare. We're not there yet, not by a long shot.

The network is the most important weapons platform for the military of the future. My ATO colleagues have spent much of the last hour describing the potential of network centric warfare. They have discussed the projects they are managing to bring this potential on line within our military. The technologies they're working to deploy have the ability to powerfully expand our warfighting capabilities. They will allow our forces to operate with greater accuracy and lethality, while putting fewer of our soldiers in harm's way.

Overcoming the technological challenges they've discussed will be an awesome job. But that is just the start. Because of all the potential of network centric warfare, all the future capabilities DoD is counting on to fight and win our wars hinge on a crucial, self-evident fact: The networks must work.

That is my area of research within ATO and what I'd like to talk about today: How do we design networks that are self-defending and self-sustaining through attack?

As my colleagues made clear: DoD is counting on the development of the Global Information Grid (GIG) to fight its wars. Our military leaders envision a GIG that provides reliable access to a rich stream of data and information for every DoD user, from war planners to individual Soldiers in the most forward deployed units.

In order to fulfill this critical mission, the GIG must offer a reliable, secure and robust computing network. Yet there are many technological Everests to climb and conquer before the scientific community can make good on the GIG's promise. As it stands today, because of fundamental problems our current computing systems, if we don't develop robust, self-defending, and self-sustaining networks, the GIG is likely to fail in delivering on its promise.

Today's networked command and control systems (C4ISR) are predicated on the ability to share information timely and securely. Our military expects to be able to have a continuous picture of the battlefield, a picture they must be able to relay through networks and data links both up and down the entire chain of command and laterally within units in an ever wider circle of information sharing. This network must be secure against all natures of attack, even though the network is based on infrastructure that is sometimes mobile, ad hoc, and always under attack.

Networks enable our C4ISR capabilities, but in an era where networks are constantly under attack, we really need to develop C4ISR capabilities for the network itself. That means a network that can conduct its own surveillance, a network that can process intelligence about the threats it faces, a network that can command and control itself.

Defending Warfighter Networks

This requires not merely evolutionary, but revolutionary changes to the way we build networking and computing technology.

Our current approach to network defense is medieval. Network defenses are designed with a fortress mentality. We build the toughest possible shell in order to repel as many attacks as possible. We use tools like firewalls, anti-viral systems, intrusion detection systems and patch compliance tools to strengthen the fortress walls in a valiant, but ultimately futile, effort to keep intruders outside the walls of the network. We treat network defense like a war of attrition, hoping we can somehow outlast the enemy's siege.

The flaw is, once the fortress walls are breached, attacks can systematically undermine the network one node at a time, and often the entire enterprise, within a matter of seconds. One Trojan horse can defeat a carefully constructed fortress. And while attackers only need to exploit one vulnerability, we have to close every hole, including holes we don't know about. It is ironic how brittle our systems are, often crashing or compromising when presented with unexpected input. Meanwhile, viruses and spyware are infamously robust, often able to withstand a barrage of detection and clean-up tools and keep on ticking. Talk about asymmetric warfare.

Obviously, we need to radically upgrade our cyber defense mentality to the 21st century. We must wage the equivalent of a war of network maneuver, rather than hope to survive a war of attrition.

To secure our goals for network-centric warfare, our networks must be designed to be self-defending and self-sustaining. What are the attributes of a self-defending network?

First, it has to be cognizant of its own behavior. Our networks must be their own doctors, with the ability to develop a baseline of health and recognize when they are sick.

Second, self-defending networks need their own command and control function. This would allow

the network to recognize attacks and failures, distinguish between malicious and benign users, determine when software is misbehaving, and provide traceback and attribution of attacks. Based on this intelligence, the network would require the ability to adaptively reconfigure in the face of attacks and failures. In other words, an autonomic command and control system for networks.

Third, self-defending networks must be self-correcting. After sensing and evading attacks, our networks must have the ability to adapt, developing new immunities so that they are no longer vulnerable to the same attacks.

What are the technology challenges that need to be overcome to achieve this vision? What are some of the approaches that might prove effective in meeting these challenges? Without intending to limit your creative input, let me sketch out a few ideas of where we'd like to go.

1. Software assurance controllers. We need on- or off-board devices that execute control algorithms for monitoring and controlling the dependability and security of essential software systems. These devices not only monitor applications for runtime failures or security violations, but also apply appropriate correctives in case of failure or compromise. Any corrective actions autonomously applied must have high precision and be cost-optimized in order to preserve as much of normal system operation as possible while isolating and correcting the problem. And we must accomplish all of this without the aid of a human in the loop.
2. Dynamic measures of system health. We must define what constitutes system health and then use these as inputs to our assurance control models. Examples include uptime, network latencies, available memory, hung processes, system restarts, and intrusion alerts.
3. Real-time, large-scale network health status. To make this possible, we require scalable real-



time measures of the operational impact of degraded system services. For example, systems functions may be lost or degraded due to attacks, malfunctions, or autonomous corrective actions. When this happens, we must be able to alert human operators and decision-makers to the operational impact of those lost services. We also need to give real-time mission health assessments to global network operations centers.

4. Out-of-band network defenses. Host defenses must run on separate hardware from the applications they are defending. Ideally, we should use a different instruction set or operating system, and a separate command and control channel for our defenses. We must reverse the long trend of building our defenses on a house with a broken foundation.

5. Trust-based Credentials. We must develop a credentialing system that can discriminate between trustworthy and untrustworthy patterns of system resource use. Individual users' credentials, system permissions and accesses, must be continually reassessed in light of each user's current behavior. And user credentials must be automatically degraded, revoked, or restored with fine-grained controls when untrustworthy behavior is detected.

The initial groundwork for self-defending networks has been laid by two ongoing DARPA projects, Dynamic Quarantine of Computer-Based Worms and Defense Against Cyber Attacks on the MANET Systems programs. While these programs have begun to show the promise of autonomic defense technology, the challenges described today remain.

The DoD vision for the GIG and future military operations requires all DoD and Intelligence Community users to have timely assured access to information. To realize this vision, to be able to achieve victory in the future, we need you to solve these technical challenges in network centric warfare.

We look forward to solving these problems together in the future.