

# Developing a Distributed System for Infrastructure Protection

George Cybenko and Guofei Jiang

*Last February's distributed denial of service attacks revealed that we need a distributed monitoring and management system to make our national critical infrastructure secure.*

**Y**our business increasingly relies on computer-controlled systems vulnerable to intrusion and destruction. The recent distributed denial of service attacks (DDoS) against e-commerce companies showed that this vulnerability extends beyond your own corporate networks: The very infrastructure of the Internet is at risk. When infoterrorists use the networks' high connectivity and low security to launch

attacks against critical information infrastructure systems, they can not only disrupt global e-commerce and communications, but can also adversely affect other critical infrastructure services such as energy, transportation, healthcare, finance, and water supply.

How can organizations protect these systems from infoterrorism? They must leverage modern information technologies to create an infrastructure protection process that can operate quickly and seamlessly. We propose a six-stage protection process that involves intelligence gathering, analysis, interdiction, detection, response, and recovery. To implement this process, we've designed an underlying Web-like architecture that will serve as a platform for the decentralized monitoring and management of critical infrastructures.

## THE PROTECTION PROCESS

The recent spate of DDoS attacks revealed two disturbing trends:

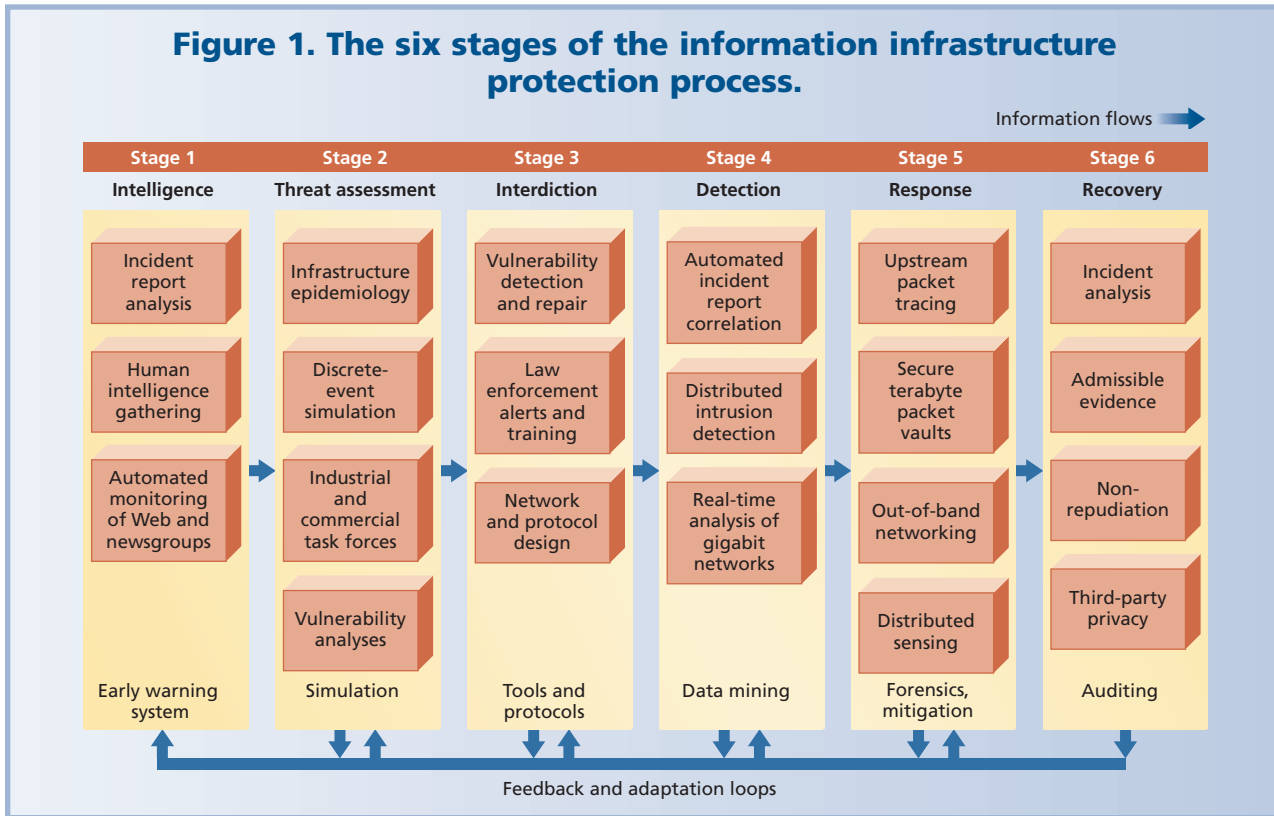
- The time intervals between when a new threat is identified, when it manifests itself, and when it mutates into different forms are relatively short and appear to be shrinking.
- Threats within one sector, such as the Internet, can easily spill over into other sectors, such as human services and the financial system.

To deal with disaster situations, the emergency management, public health, and computer security communities have decomposed their management processes into smaller, logically concise stages. For example, the DARPA Information Assurance program uses the three-stage "Protect-Detect-React" paradigm to organize work within that area (<http://www.darpa.mil/iso/ia/>). Figure 1 shows the six stages we propose for information infrastructure protection. These stages roughly correspond to the stages used in other emergency management areas.

### Stage 1: Intelligence

Infrastructure management begins with gathering intelligence about emerging threats. Figure 2 identifies the methods for early threat identification. These methods include intelligence reports, unusual-incident analysis, and automated information harvesting from the Web and news services. By drawing upon these and related open sources, organizations and governments can create a warning system that identifies new threats early in the process. Red teaming—the use of selected experts for proactive analysis through

**Figure 1. The six stages of the information infrastructure protection process.**



scenario building and threat design—also plays an important role in intelligence gathering.

Ideally, a new incident report could be quickly and automatically matched against an online database of previous threats. Currently, experts—who rely on their own memories, networks of colleagues, and ad hoc archive searches of previous attacks—perform this early warning stage.

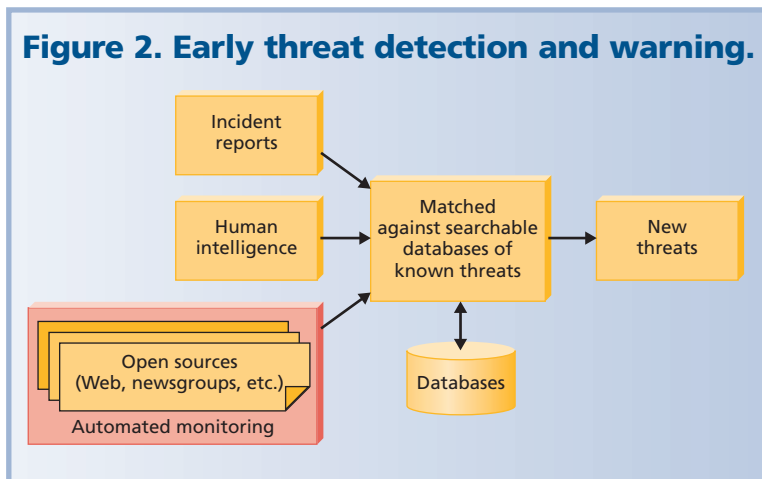
Although technically possible, automated monitoring of the Web and various newsgroups to identify threats early (<http://informant.dartmouth.edu>) has yet to be implemented.

### Stage 2: Threat assessment

Once we identify a new threat, we must perform a risk assessment and some sort of “cost-benefit” analysis for responding to the threat. As Figure 3 shows, this process involves modeling how an attack or failure based on the threat will manifest itself, and how it will affect other infrastructure systems.

Assessing a threat’s potential consequences presents the challenge of developing quantitative models. These models reveal the national and even global distribution of infrastructure vulnerabilities, and how failures based on those vulnerabilities can cascade through the overall infrastructure.

**Figure 2. Early threat detection and warning.**



### Stage 3: Interdiction

In the interdiction stage, organizations and governments attempt to proactively prevent or prohibit failures based on known threats. Typical interdiction agents that private organizations can employ include virus scanners, software patches, and improved network designs and protocols. Companies should also consider training responders—system operators that form the first line of defense—in appropriate interdiction actions.

Unfortunately, the rate at which new threats arise outstrips the ability of such personnel to

attend training meetings and courses that explain how to combat them. Thus, this information must be transmitted using remotely accessible distance training supported by networked interactive material.

Cost-benefit analysis performs the essential service of identifying the threats and vulnerabilities most likely to have a severe impact. Lacking unlimited resources, first responders cannot prepare for all possible failures and thus focus on interdicting high-cost and high-probability events.

#### Stage 4: Detection

By monitoring distributed “sensors” positioned throughout the infrastructure itself, we can detect actual failures or attacks. Such sensors include computer network monitors based on, for example, SNMP (Simple Network Management Protocol) agents or packet analyzers. Governments could monitor public health records, medical laboratory results, environmental monitoring stations, financial-market trend monitors, and so on. Raw sensor data must be harvested, mined, correlated, and otherwise analyzed.

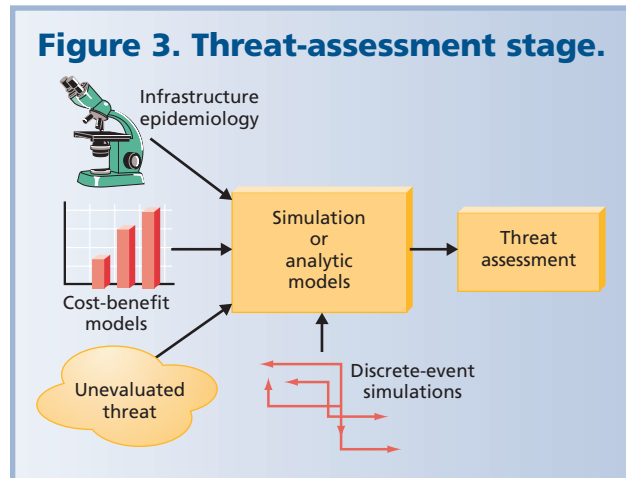
Whereas early warning systems anticipate attacks and failures through proactive intelligence gathering and analysis, detection responds to mature attacks and imminent failures. Ideally, threat assessment and interdiction have prepared the community for these events.

The challenge in automating this stage lies in flagging previously unseen anomalous events without generating large numbers of false positives. Meeting this challenge requires that a trained, automated system recognize “normal” and known behaviors so that it can flag those that fall outside this regime. Complicating this process is that many new behaviors emerge in the course of natural, non-threatening operating modes.

#### Stage 5: Response

Once the process detects an attack or failure, system security must respond appropriately. Our work focuses on law enforcement or internal auditing responses to information infrastructure events. Infrastructure attacks challenge responders to swiftly identify the problem’s source. Doing so requires forensic techniques that can build a trail of legal evidence for future investigation while respecting third parties’ privacy. These considerations require fast and reliable upstream packet tracing, something that currently requires time-consuming and relatively slow operator intervention. Moreover, given that many Internet links now operate in the multiple-megabit- and even gigabit-per-second range, archiving network traffic for forensic analysis presents a major technical challenge.

Another challenge arises from the telecommunications networks, which responders must use to coordinate a response. Such systems are themselves part of the infrastructure and thus highly vulnerable to failure. To address this vulnerability, any future infrastructure Web architec-



ture must provide for out-of-band and otherwise redundant communications.

Multiple communication channels based on different protocols, implemented by different vendors, offer one solution. They ensure that a single vulnerability does not compromise the entire system. Thus, from a security perspective, standardized systems compromise survival and should be replaced or supplemented with heterogeneous ones. Radio and satellite networking can provide additional out-of-band communications capability.

#### Stage 6: Recovery

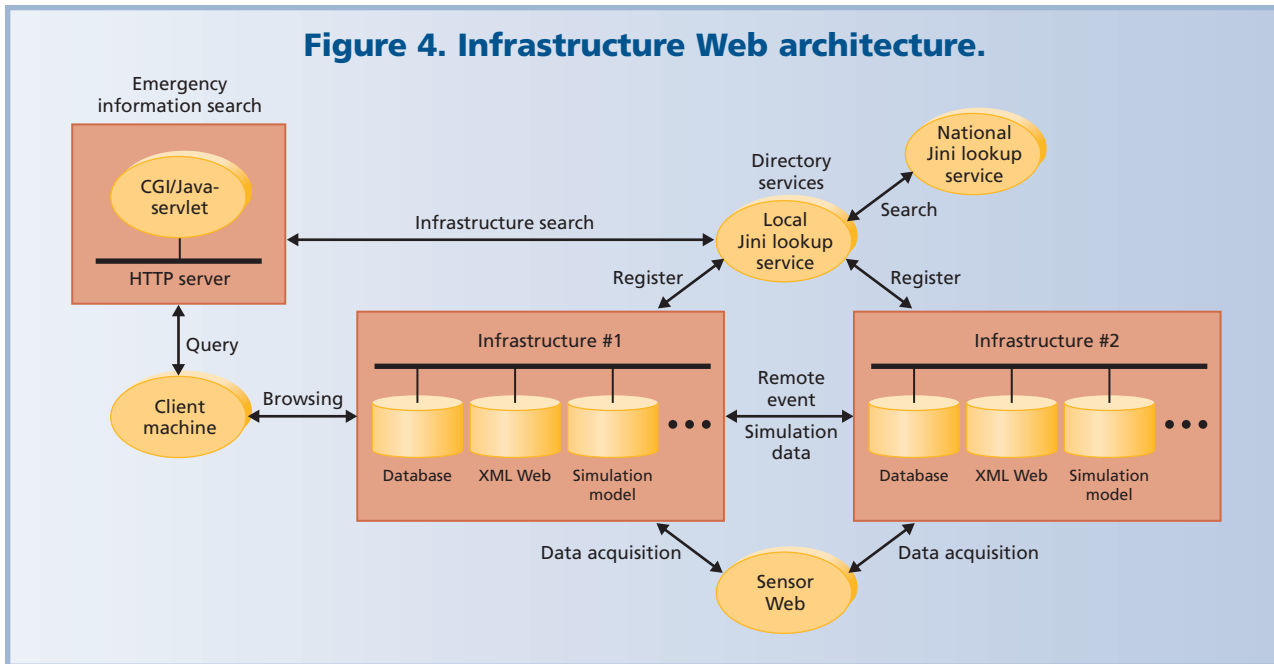
Recovery from an attack on a critical infrastructure involves collecting evidence pertinent to it without violating privacy laws and standards. Investigators also require a complete analysis of the incident to learn from it and to archive its characteristics for future detection and training. First responders must be trained in the appropriate forensic techniques to accomplish these goals.

### BUILDING A SYSTEM ARCHITECTURE

According to the report of the President’s Commission on Critical Infrastructure Protection (<http://www.ciao.gov>), the infrastructure networks of greatest importance to national security and stability include telecommunications, electrical power systems, gas and oil, banking and finance, transportation, water supply systems, government services, and emergency services. To effectively protect these critical infrastructures, it will be necessary to have a system in place that can monitor and manage these very large and complex, dynamic networks. Our proposed infrastructure Web provides such a basic architecture, as well as the underlying paradigms by which a problem of this scale and scope can be addressed.

How do we integrate and implement the protection process into a functional monitoring and management system? We propose that national infrastructure Web networks be built with four types of basic distributed compo-

**Figure 4. Infrastructure Web architecture.**



nents: directory services, infrastructure servers, sensor webs, and emergency information search servers. We are using Sun's Jini system (W. Keith Edwards, *Core Jini*, Prentice Hall, Upper Saddle River, N.J., 1999) to organize and integrate these distributed components throughout the nationwide networks. Jini, designed for deploying and using services in a network, enables the construction of dynamic, flexible, and robust systems built with independent distributed components. Further, access to Jini's source code has fostered a thriving community of developers who continue to enhance and expand Jini's capabilities (Jim Waldo, "Alive and Well: Jini Technology Today," *Computer*, June 2000, pp. 107-109).

Figure 4 shows the framework for the Infrastructure Web architecture. Using this kind of architecture, we believe that the Infrastructure Web can be exploited as a platform to implement a scalable distributed infrastructure assurance vision.

### Infrastructure server

In the Infrastructure Web system, one infrastructure network server represents one critical infrastructure in the physical world. The server's IP (Internet Protocol) address functions as the unique identification for the infrastructure. The infrastructure server will have a real-time database, an XML-based Web interface, a simulation model, and possible other services running on some ports with the server's IP.

The database acquires real-time data from the sensor Web or other sources, such as host-based detection systems. This data consists of the infrastructures' security status, internal states, and so on. Some data will be displayed on the Web in

real time to show the infrastructure's current status, and some will be used in simulations such as those for threat assessment. We will investigate Open Database Connectivity (ODBC) and Java Database Connectivity (JDBC) technologies to make database access transparent.

By browsing the infrastructure's XML-based Web, clients can directly check some of the infrastructure's security statuses and internal states, such as the packet traffic throughput of an important LAN infrastructure. The relationships between related infrastructures will be described by XML's X-Link and X-Pointers (Elliotte R. Harold, *XML Bible*. IDG Books Worldwide, Foster City, Calif., 1999). For every infrastructure category listed in the President's Commission on Critical Infrastructure Protection, a Resource Description Language (RDL), created with XML document type definitions (DTDs), will let us clearly describe an infrastructure's attributes of this category in standard formats. We'll also use other Web technologies, such as Java applets and JavaScript, to describe the infrastructures.

These infrastructure servers will need suitable analytical and simulation models of the infrastructures, together with a description of their behavior under dynamically changing interconnections. Just as we wire the pins of chips in a printed circuit board design, we can "wire" the inputs and outputs of infrastructure blocks to implement large-scale discrete-event simulations for threat assessments. We believe that, by applying adaptive-learning technologies such as neurocomputing and evolutionary computing (Dimitri P. Bertsekas and John N. Tsitsiklis, *Neuro-Dynamic Programming*, Athena Scientific, Belmont, Mass., 1996), we can derive better planning, control, and coordi-

nation strategies from the simulations. These strategies and policies will help the related infrastructures cooperate efficiently once a disaster or attack occurs.

Other services can also be implemented on the infrastructure server, such as intrusion detection. Moreover, some systems, like the early warning system itself, can be implemented as an information infrastructure in our architecture to offer automated intelligence collection service.

### Directory service

The Infrastructure Web system's directory service will have two tiers, one each at the state and national levels. We intend to implement directory services with Jini's lookup service. Infrastructures must register themselves in the local Jini lookup services, and all local lookup services must register themselves in the national-level Jini lookup services. An infrastructure's registration describes its attributes, such as its category and location, server IP and URL, proxy interface program, and so on. Jini's attribute mechanisms support both type- and content-based search styles, which makes searching for particular attributes simple, quick, and effective.

Jini has five basic concepts: discovery, lookup, leasing, remote events, and transactions. These concepts form the basis for Jini's ability to support spontaneously created, self-healing communities of distributed components. Further, Java's remote method invocation (RMI) and object serialization techniques (Joseph L. Weber, *Using Java 1.2*, Que Publishing, Indianapolis, Ind., 1998) make the implementations of these concepts possible. The Infrastructure Web system will be organized and integrated with the Jini system and will inherit these concepts' advantages from Jini.

For example, with Jini's leasing concept, all infrastructures must sign a lease with the lookup service in the registrations. Once the leasing time expires, the infrastructure will automatically be removed from the lookup service. This process provides Jini's lookup service with a self-healing ability that prevents its directory management and clients from receiving outdated or nonexistent infrastructure information. Moreover, the remote-events concept lets Jini better describe the relationship between related infrastructures. For example, Infrastructure 1 can tell related Infrastructure 2 which of Infrastructure 2's statuses it cares about. Once these statuses change, remote events from Infrastructure 2 will automatically notify Infrastructure 1, just as if they were events local to Infrastructure 1. This methodology lets geographically distributed infrastructures cooperate efficiently to detect, respond, and recover from possible intrusion and attack.

Based on these Jini concepts, we believe that the Infrastructure Web can be implemented easily with the

capabilities required to swiftly and decisively respond to infrastructure attacks and failures.

### Sensor web

Another essential part of our proposed system is the ability to monitor and detect stimuli or the states of distributed infrastructures. When assessing DDoS attacks, an analyst or upstream packet tracing system needs not only the packet information from the local machines, but also those from remote routers or firewalls. So our sensor web system will actually function as a large-scale distributed smart-sensor network that collects distributed sensor information from intelligent software sensors and smart hardware sensors. Just like the infrastructure, the sensor web will register its sensors in the directory service, and all these sensors can offer distributed data-sensing services.

Advances in measurement devices have reduced costs to the point where we can now develop affordable large-scale distributed sensing systems. Meanwhile, advances in processor technology allow for relatively low-cost, low-power, compact distributed processing. Such processing power, when integrated within these sensor devices, results in what is commonly called a *smart sensor*. Smart sensors capable of parsing and filtering only the necessary or desired information allow for efficient memory use, conserve precious wireless bandwidth, and reduce the battery power needed for transferring sensor information.

Before sending sensor information to related infrastructures, the sensor web system will preprocess the data from the distributed sensors, using techniques such as data filtering, data fusion, and data mining. More information about our distributed sensor web systems can be found in *A Study of Distributed Smart Sensor Networks* (Michael G. Corr and Clayton M. Okino, tech. report preprint, Thayer School of Engineering, Dartmouth College, Hanover, N.H., Mar. 2000).

### Emergency information search

When an infrastructure fails or suffers intentional attack, its damage must be repaired quickly. Meanwhile, system operators must adjust alternate infrastructure components to cover the damaged one. They also must check the status of related infrastructures to help coordinators determine the extent of the problem. Unfortunately, emergency information search-and-response systems of this kind have been unavailable so far. To fulfill this role, the Infrastructure Web will have a nationwide Jini-based directory services system that will register all critical infrastructures.

Like a 911 telephone emergency system, the emergency information server should have a well-known domain name. Likewise, the emergency query forms should be formatted for speed and ease of use. After clients submit a

**We intend to implement directory services with Jini's lookup service.**

query, the server will transfer the query data to the CGI (common gateway interface) or Java servlet programs. These programs will process the query data and submit a formatted attributes template to the Jini lookup services. The Jini system will then search the desired infrastructures from its lookup services and return the possible infrastructures' general information and URLs. By browsing the infrastructures' XML-based Web sites, clients can check the real-time status and internal state of any listed infrastructure.

### THE ARCHITECTURE IN ACTION

The architecture and components we've proposed offer a basic open platform for distributed infrastructure monitoring and management. We believe that after the system is implemented and deployed, it will function as follows: While sensor web systems collect data for all critical infrastructures, the services on every infrastructure server monitor these data to verify that the infrastructure is running well. Once the infrastructure server detects attacks or failures, related infrastructures will be notified and response and recovery steps will be taken automatically. Meanwhile, the coordinator can search and check real-time data from all appropriate infrastructures and choose the appropriate presimulated control strategy to respond to and recover from the attacks or failures.

However, to integrate this system with various infrastructures at different agencies, organizations, and entities we need

- more complete knowledge of the operations of the various infrastructures from different fields,
- organizations and experts to analyze the specific requirements for their areas, and

- more academic and industrial research on how to protect all these infrastructures separately.

Eventually, we'll also need state- and local-government support to deploy the system we propose. To protect the Infrastructure Web from malicious attacks, redundant communications channels and protocols will be necessary.

**T**he basic technologies and components for creating such an Infrastructure Web are currently under development at the Institute for Security Technology Studies at Dartmouth College. That effort is part of a comprehensive research program on cybersecurity and infrastructure assurance funded by the US Department of Justice's National Institute of Justice. While related national-level US-based efforts in this direction are beginning to take shape, we are not aware of any coordinated national efforts in other countries, let alone any global-scale programs. ■

*George Cybenko is the Dorothy and Walter Gramm Professor of Engineering at Dartmouth College's Thayer School of Engineering. Contact him at [gvc@dartmouth.edu](mailto:gvc@dartmouth.edu).*

*Guofei Jiang is a postdoctoral research associate at the Thayer School of Engineering. Contact him at [gjf@dartmouth.edu](mailto:gjf@dartmouth.edu).*

*An earlier version of this article was presented at the Infrastructure Protection and Emergency Management (IPEM) Symposium held as part of the 2000 Advanced Simulation Technologies Conference. The Society for Computer Simulations International sponsors ASTC.*

## Set Industry Standards

Our members write important IT standards, including IEEE 802.3, the standard for Ethernet, the most widely deployed LAN. But technology networks are not the only kind of standards developed here.

**Grow Your Career • Find Out How @**  
***[computer.org/standards/](http://computer.org/standards/)***

### Did You Know?

Right now, over 200 Working Groups are drafting IEEE standards.