Review of New Security Paradigms 2002 Workshop Papers

By Christina Serban and Hilary Hosmer

The New Security Paradigms Workshop (NSPW) offers researchers a safe, constructive environment to explore radical rather than evolutionary approaches to information assurance. This year's workshop took place at the Founders' Inn in Virginia Beach, VA from Sept. 23-26, 2002. Thirty-eight researchers participated, including Ph.D. students, faculty members, and information scientists. Most were from the USA, but participants came from Ireland, Switzerland, Japan, Russia/Israel, Czechoslovakia, Egypt, China, and India. Proceedings, which are published after the workshop, will be available from ACM and from the ACM Digital Library in 2003.

In NSPW's highly interactive environment, each author presents a new paradigm for 20 minutes, but with discussion the session usually lasts about an hour. This year there were many interesting new paradigms. Below we summarize each one briefly.

Session 1. Intrusion Detection and Response

An Experimental System for Malicious Email Tracking

M. Bhattacharyya, M. Schultz, E. Eskin, S. Hershkop, S. Stolpho, Columbia U.

Commercial virus scanners find known viruses, but can't detect new ones. They also don't provide data about the propagation of viruses across the network to warn untouched users. Malicious Email Tracking (MET) limits propagation of malicious email attachments and tracks points of entry and initial distribution.

This research proposes a MET server (trusted, central location) plus MET clients (at mail servers) to monitor the behavior of email attachments across all your mail domains, then detect and contain email-based attacks. Each email attachment entering the domain is assigned a unique identifier (MD5 hash / "signature"). The ID, timestamp, sender, and receiver are logged. Two key statistics are obtained for attachments:

1) Prevalence: # of times attachment observed by MET client;
2) Birth rate: average number of copies sent from same user.

Because rapidly self-replicating viruses have extremely high birth rates, an attachment with a very high birth rate is a potential self-propagating virus.

In practice, the MET server collects data on malicious activity, stores them in a DB, and calculates derived stats. It also keeps a list of IDs for known malicious viruses, updates it, and propagates to MET clients for automatic updates.

The system can detect self-replicating viruses, even previously unknown ones, if the birth rate is > threshold t, sent to > u users. For emails over the threshold, blocking (discard or sideline) is used. Detection is done at MET clients with alerts to MET server which propagates to other MET clients.

Future work includes IDs for polymorphic viruses, mailbox "fingerprinting" tool, and early spam categorization based upon cliques of senders.

Predators: Good Will Mobile Codes Combat Against Computer Viruses

H. Toyoizumi and A. Kara

Just as white blood cells in animals replicate to attack invading organisms like viruses, bacteria, poisons, and foreign matter, in information technology good will mobile codes will self-replicate to attack invading malignant code (viruses, worms, etc.). The paper models the interactions between computer predators and viruses using the Lotka-Volterra equations widely used on mathematical biology. In nature predators are kept in check by the disappearance of their prey (food), but other techniques are needed to dampen the number of mobile code predators so they don't degrade the network.

An Empirical Analysis of NATE- Network Traffic Analysis of Anomalous Traffic Events

Carol Taylor and Jim Alves-Foss, University of Idaho, Moscow, Idaho

NATE was presented at NSPW 2001 as a low-cost approach to intrusion-detection, detecting attacks from packet header information. It detects probes, scans, and DOS type attacks from normal traffic. Unlike most statistics-based anomaly-detection programs, NATE can self-configure, so does not require the system administrator to know a system's normal parameters in order to configure the system. Because it only looks at headers, it can handle encrypted information. Anomaly-based detection allows it to pick up new attacks, unlike firewalls and filters whose rules can be by-passed. NATE can operate inside or outside a firewall, providing additional filtering capabilities and monitoring compromised machines inside the firewall.

Carol Taylor reported results this year from using NATE on a real operational non-academic data set from a small network with web, email, and firewalls. She found that the real data was much more variable than the constructed test data, and had to refine anomaly tests to eliminate large numbers of false negatives. She also had to include some of the constructed test data to cover possibilities not represented in the real data. She found that sampling by attribute distribution was a good alternative to sampling by TCP type, and that various measures of distance from the norm each had advantages and disadvantages. She recommends using a distance measure that captures relationships between TCP session attributes.

Ms. Taylor recommends more testing with NATE, to see if expanding the number of attributes, such as time since DOS attacks, results in better attack detection. She also recommends an actual prototype deployed on a high band width network

to assess real-time performance under actual working conditions. More testing with a wider range of attack and normal data needs to be done. False positives and negatives need to be researched.

Session 2: Large Systems

Small Worlds In Security Systems: An Analysis of the PGP Certificate Graph

Srdjan Capkun, Levente Buttyan, and Jean-Pierre Hubaux, Swiss Federal Institute of Technology Lausanne, Switzerland

The problem of securing fully self-organized mobile ad hoc networks motivates this work. Mobile ad hoc networks have no fixed infrastructure; all networking functions are performed by the nodes themselves in a self-organizing manner. Many of Milgram's small world phenomena apply, and PGP certificate graphs (directed graphs G(V,E) where V is a set of vertices representing users' public keys and E is a set of edges that represent public key certificates) are an inspiration.

In a small world, the average number of acquaintance links between any two people is five or six. The equivalent of an acquaintance link in secure computing is a public key certificate. These are distributed among nodes based upon sociological relationships between users in the network, so small world principles apply. For example, to authenticate a public key, each user keeps a local certificate repository of certificates and processes them to develop a chain of trust. If the user's own certificates can't produce the chain of trust, the two users wishing to communicate can merge their certificates to get a chain of trust.

Since existing small world models do not correctly model certificate graphs, the authors propose a new certificate graph model with irregular vertices and an irregular lattice. For future work, the authors propose to study in detail mechanisms by which trust is likely to emerge in fully self-organized systems.

Breaking the Barriers: High Performance Security for High Performance Computing

Kay Connelly, Indiana University and Andrew Chien, UC San Diego

High performance workstation clusters are insecure computing environments. The standard practice is to have no security beyond simple logins and access rights, so that nothing interferes with the performance of the search engine, the reservations system, or command and control system. All data is sent in plaintext, since encryption requires too much overhead. Attackers can (1) send remote procedure calls (RPC) to various components to change the execution of an application; (2) eavesdrop and attack when the system is in a vulnerable state.

Security mechanisms for the HPC environment must have low overhead and protect data long enough to change state. The authors' approach includes streamlining the encryption process when data is put onto the wire, and precomputing during idle time to reduce communication latency.

The authors define three metrics and describe an initial prototype.

From Privacy Promises to Privacy Management—A New Approach for Enforcing Privacy Throughout an Enterprise

Ashley, M. Schunter, Powers, IBM Research Labs, Switzerland

Privacy is the right of individuals to determine for themselves when, how, and to what extent information about them is communicated to others. The paper's focus is on Personally Identifiable Info (PII) privacy as managed by a service provider. Looking at OECD privacy principles and usage phases (Notice, Collection, Cataloguing, Control, Release, Recording, Response), there aren't any tools to address the phases beyond Notice and Collection throughout an enterprise. Most privacy policies are unimplemented throughout the enterprise.

The proposed framework:

Define enterprise privacy policy
Deploy policy to IT systems containing PII
Record user consent to advertised privacy policy when submitting PII
Enforce privacy policy, create audit trail of access to PII
Generate enterprise-wide and individualized reports of accesses to PII and conformance to governing privacy policy.

The privacy policy consists of:

Elements: data users, operations, data types, purposes, conditions.
Rules: ALLOW [data user] to perform [Operation] on [Data Type] for [Purpose] provided [Condition]. CARRY OUT [Obligation].
Deploy policy: Map data, users, tasks into policy elements.
Record consent: Record collected data plus PII infor, timestamp of consent and applicable version of privacy policy.
Enforce policy, create audit trail of access: Can be real-time or near-time.
Report: Respond to individual inquiries as well as enterprise level inquiries (auditors, outside agencies).

Session 3: Mobile Code

Anomaly Intrusion Detection in Dynamic Execution Environments

Hajime Inoue and Stefanie Forrest, University of New Mexico

Products such as Java are based upon dynamic compilation, profiling, and optimization technologies. The potential exists to leverage their infrastructure for anomaly intrusion detection with extremely low performance penalties and customization to a specific application. The authors propose to automate the construction of an application intrusion detection system without modifying the application by profiling information already in place for dynamic optimization. They call this "dynamic sandboxing" and demonstrate the approach.

Empowering Mobile Code Using Expressive Security Policies

V.N. Venkatakrpishnan, Ram Peri, R. Sekar, SUNY at Stonybrook

The authors aim to empower mobile code rather than disable it. Highly expressive security policies provide the basis for such empowerment while greatly mitigating the risks to the host system. Their implementation is based upon rewriting Java byte code so that security-relevant events are intercepted and forwarded to the enforcement automata before they are executed.

The Source is the Proof

Vivek Haldar, Christian Stork and Michael Franz, University of California, Irvine

There are two main approaches to mobile code security: byte code and proof-carrying code. The authors propose an alternative called WELL (Well-formed Encoding at the Language Level) which transports compressed abstract syntax trees, permitting transporting programs at a much higher level of abstraction that is closer to the source. The method provides safety by construction.

Future work includes improving performance and exploring transporting other annotations.

Session 4: Usability

An Approach to Usable Security Based on Event Monitoring and Visualization

Paul Dourish and David Redemiles, University of California, Irvine

One cause of the disparity between theoretical and effective security is the extent to which users can comprehend and make effective use of security mechanisms. The authors' thesis is that a technical infrastructure which makes available security mechanisms visible will enable users to make informed decisions, thus rendering the system more secure. They propose a layered framework for visualizing and monitoring security mechanisms, events, and sources, using probes, gauges, and alarms.

Moving from the Design of Usable Security Techniques to the Design of Useful Applications

D.K. Smetters and R.E. Grinter, PARC

The usability of security technology may be one of the largest roadblocks standing in the way of increased computer security, and it is only going to get worse as security technology undergoes radical change. The authors approach the problem from a different perspective: if you put usability first, how much security can you get?

The users look usable key management, authentication for ad hoc networks, and implicit security starting from usability.

Three engineering approaches are: Build in implicit security, Refactor security infrastructure, Build Lego Blocks for Security

Session 5: Panel Discussion on Assurance in Critical Endeavors

Session 6: Securing Information

Capacity is the Wrong Paradigm

Ira Moskwitz, LiWu Chang, Richard Newman

Capacity is the prevailing paradigm for covert channels. With respect to steganography, however, capacity is the wrong paradigm. The authors propose a new paradigm called "capability' to gauge the effectiveness of a stenographic method. Capability includes payload carrying ability, detectability, and robustness components. JPEG compressed images always have the potential to carry hidden information.

Toward Achieving Acceptable Security in Secure Multi-party Computation

Wenliang Du, Syracuse University

Secure Multi-party Computations deal with situations where two (or more) parties want to jointly perform a computation but each wants to keep the data it provides hidden from the other parties. Approaches requiring zero-information disclosure fail. The author recommends an approach where partial information disclosure is acceptable.

Guarding the Next Internet Frontier: Countering Denial of Information Attacks

Mustaque Ahamad, Wenke Lee, Ling Liu, Leo Mark, Edward Omicienski, Carlton Pu and Andre dos Santos, Georgia Institute of Technology

This position paper introduces the Quality of Information (QoI) concept and the denial of information (DoI) attack. The many dimensions of QoI include: consistency, timeliness, reliability, trustworthiness, and density/richness of information. A denial of information attack inserts noise or bogus information degrading the quality of data, in either a massive or a gradual way.

ACM and ACM SIGSAC have sponsored the workshop since its start in 1992, and several organizations, including DOD, CERT, and James Madison University provided financial support this year.