

A Submission to Phoenix Challenge 2005 Award Paper Program

Advanced Security Reporting Systems for Large Network Situational Awareness

Presenters

Michael Collins, CERT/NetSA (mcollins@cert.org)
Gregory Virgin, NSA (gsvirgi@missi.ncsc.mil)

Executive Summary

In collaboration with NSA and JTF-GNO, the Network Situational Awareness (NetSA) group at CERT/CC, Software Engineering Institute, Carnegie Mellon University, has developed large-scale network traffic reporting systems that provide analysts with the capacity to dynamically query large summaries of network traffic over time. These systems are deployed on the NIPRNet as part of the JTF Centaur capability. In this presentation we describe the sensor and analysis technologies that support an asset inventory system, and serve as a foundation for a flexible, ad-hoc intrusion detection capability. These facilities have greatly increased our ability to respond strategically to information security challenges, and to detect novel threats to the NIPRNet, in an environment where both attacks and normal traffic are changing continuously.

Contacts

Andrew Kompanek
CERT/NetSA Technical Lead, DoD Projects
ajk@cert.org
(412) 268-9744

Rex Brinker
SEI/CERT Project Manager, Information Assurance
rbrinker@sei.cmu.edu
(412) 268-7722

Discussion

Network security analysis faces a constant logistical conflict due to the highly skewed nature of network traffic. The overwhelming majority of network traffic is, if not innocent, at least innocuous, while attacks can be very small and well hidden. Any security system must infer behaviors from an enormous volume of packets passing through a network at any time. As networks get larger, these problems become exponentially more difficult; the NIPRNET, by virtue of being one of the world's largest computer networks in terms of volume, IP real estate, and geographic coverage, may be the most complex IP network known.

Traditionally, security analysis systems manage this problem by filtering data: an IDS matches the attributes of packets against known signatures and sends an alarm; a security audit takes a snapshot of the system's behavior at one time and evaluates the whole environment. However, both of these responses are limited; IDS recognize traffic they have been trained to recognize, and audits are disruptive events.

In collaboration with NSA and JTF-GNO, CERT/NetSA has been developing systems that approach this problem from a different angle. Instead of filtering traffic, the SiLK toolset and data collection system used by CENTAUR summarizes traffic using an abstraction standard from CISCO called Netflow.

Because flows are compact, agnostic to protocol, and collected by the routing infrastructure, flow provides a 'sky-eye' view of the network. Daily, hundreds of terabytes flow through the NIPRNet. This enormous volume of data is summarized by approximately 8 billion flows/day, yielding around 170 GB of summarized flow information. Implementing the SiLK tools on cluster architecture, the current implementation of CENTAUR can generate daily summaries of traffic in less than 5 minutes. The result is a near-complete low-resolution image of the traffic crossing to and through the NIPRNet on a regular basis.

We are now expanding these facilities by replacing the CISCO Netflow collectors with system jointly designed by CERT and NSA. These revised systems implement a flow abstraction that, like Netflow, provides complete coverage of the network using small traffic summarizations, but focuses on security aspects. Since Netflow was originally designed for traffic accounting, it is primarily intended to provide an aggregate measurement; Netflow as implemented by CISCO is both relatively low resolution and tolerant to a degree of report loss.

These attributes are being rectified by implementing new flows and by revising the collection system using two new tools: AMP (Analytic Metadata Producer), which designed by NSA to provide more security-relevant flow data, and FlowCap (Flow Capacitor), which CERT designed to replace the UDP-based, lossy Netflow reporting structure with a compact and robust TCP implementation.

AMP collectors are connected to routers and passively monitor the traffic passing through them, thus freeing the routers from the additional processing overhead for Netflow. Using a DAG card to capture and process the traffic, AMP is currently robust at gigabit rates and is currently being tested for OC-48 bandwidth using commodity hardware. AMP generates enhanced flow records that expand on the traditional flow definition by including protocol and payload information. Attached to each AMP collector is a FlowCap, which replaces the fast-but-unreliable UDP with an intelligently moderated TCP connection; this capability lets FlowCap transmit information over lower bandwidths and compensate for network disruptions.

Attack activity is hidden, flexible, and constantly changing. Networks themselves are incessantly reconfigured, as new applications become standard tools. Consequently, one of flow's great strengths is that it is nonjudgmental – it records all data, rather than assuming what any particular datum is. Our new approach adheres to that ideal, while adding additional fields for protocol-specific information. This new information is stored in expanded flow records specified by CERT.

The new capacities added by AMP include more in-depth TCP information, particularly concerning how TCP sessions are terminated, and a more thorough examination of TCP flags. Our experience has been that the activity that takes place in the first few packets is more relevant than what happens at the end of the session. Consequently, we now include records of the flags for the first packet and a summary of how the session terminated. We have added limited payload information as well; ICMP, which is a very compact message passing protocol, now gets a specific payload field in AMP flows. We represent other protocols, where the payload is both larger and less well defined, by hashing their payloads. This allows us to recognize when the same payload is sent multiple times (such as with a worm).

While we record abstractions of the payloads in flow data, packet information is also used to maintain a dynamic inventory of NIPR. Packet payloads examined by Trickler, an AMP facility developed by NSA that compares payloads with known services to generate an inventory of applications running on a machine. This information is stored in a central database that provides a constantly running security audit of NIPRNet. NIPR is not only threatened by attackers, but by the constant and necessary changes to its architecture; the Trickler capability ensures that even as NIPR is changed, these changes are recorded automatically. In addition, passive monitoring facilities like Trickler identify and quantify the difference between stated policy and the tools actually used to run NIPR.

These combined facilities provide a significant strategic advantage to NIPR. Already, we have used the SiLK, AMP, and Trickler capacities to identify large gaps in the routing architecture, sources of wasted bandwidth and infiltrations, and damage to the NIPRNet from hostile sources. By virtue of solving these problems for large networks, and by focusing on making the collectors and architecture commoditized, our work is being adopted by other security organizations, notably AFCERT.