

A Federated Model for Cyber Security

Scott C. Pinkerton

Abstract—We are becoming increasingly reliant on a digital infrastructure at a personal, corporate, and national level. With that reliance comes a corresponding increase in our cyber risk profile, while simultaneously the sophistication, skill, and organization of the *bad guys* is increasing rapidly as well. These conditions together are creating new challenges for managing and mitigating cyber security threats. We are proposing a new approach to dealing with these problems. Our federated model uses a message passing architecture that has the potential to dramatically improve our ability to observe, orient, decide, and act on cyber security related events.

Index Terms— Federation, message passing, near real time, OODA loop, cyber security.

I. BACKGROUND

EFFECTIVE cyber security continues to be a challenging problem for federal, commercial, and educational organizations alike. The number and sophistication of cyber attacks, and the financial gains to be had from them are increasing at an alarming rate as indicated in [1]. Consider the following:

- Our personal and corporate lives are becoming increasingly digital. The younger generation is seemingly “always on, always connected”, and obviously deeply immersed in a digital age. The older generations might not knowingly be active in a digital sense, but with our current banking and credit based economy – people are indeed immersed in a digital economy whether they knew it or not. People are just now starting to understand the concept of a digital identity; are learning how to manage that identity, and are coming to terms with the risks associated with this paradigm shift. Deftly illustrated in [2].
- This transition to a digital world came about with the adoption of standardized technologies, one in which the Internet Protocol (IP) serves as the *lingua franca* for communication, data exchange, and in many ways commerce, as detailed in [3]. This adoption of standardized technologies has enabled a faster, and wider, deployment of digital systems, but unfortunately

brings with it a host of known vulnerabilities, and a growing collection of readily available exploit tools.

- As our personal and corporate lives are increasingly bound to a digital environment – our risk exposure is increasing at a rapid, possibly exponential, rate. Often without our awareness of it. How many corporate databases, or backup tapes, currently store our credit card numbers, our social security numbers, or our bank account information? How rapidly is that number growing – weekly, monthly, or yearly? Unfortunately today we don’t have any mechanism that enables us to track our digital risk profile. These problems are summarized in [4].
- Lastly, we are seeing a marked increase in the sophistication, skill, and organization of the *bad guys* – as described in [5]. Since cyber crime obviates the need for physical proximity – many elements (assets) of our life are accessible from anywhere in the digital realm. Commercialization and other economic forces are developing within the black hat community – point of sale style payment on malicious web site click through, purchasing botnet resource time, rootkit and other malware toolkits available for purchase, etc.

With these dynamics as a backdrop it is more important than ever to develop effective methods for managing and responding to cyber security threats.

II. INTRODUCTION

Today we have numerous sources for acquiring information for cyber security purposes – two common techniques being host and network based sensors. Host based analysis utilizes software loaded directly onto a computer to monitor actions and interactions occurring. This approach has some definite advantages but doesn’t always scale well, and requires cooperation, or at least some level of administrative control, for installing the software onto the computer. Host based security is well suited for environments with a strong centralized Information Technology (IT) function.

Network based analysis is more opaque; it does not require any access or interaction with each computer, nor cooperation from a user/owner. Instead it relies on access to the switch or router infrastructure that the computer(s) connect to. This approach also has advantages and disadvantages. Today, network sensors commonly take the form of either flow data analysis (E.g. Netflow, Sflow, or IPFIX – reference [6]), or deep packet inspection (analyzing data content within each packet traveling on the network). In the near future, we will

Manuscript received October 14, 2007; revised November 18, 2007. The material in this paper was presented in part at the Cyberspace Research Workshop, Shreveport, LA, November 2007. This work was supported in part by the U. S. Department of Energy under Prime Contract No. DE-AC02-06CH11357.

S. C. Pinkerton is with Argonne National Laboratory, Argonne, IL 60439 USA, (phone: 630-252-9770; e-mail: pinkerton@anl.gov).

likely see a wider array of sensors developed and deployed that will provide new insight into cyber-centric activities.

However, a challenge with the sensors currently in use and more so with the sensors coming in the future, is that we are faced with the dilemma of data glut and information famine. By this I mean that we are inundated with large volumes of raw data from sensors (often with discouraging signal/noise ratios), and very little actionable information.

At Argonne National Laboratory (ANL) we collect flow data from our network switches and routers, signature fires from Cisco Intrusion Detection System (IDS) sensors spread across the campus infrastructure, as well as syslog data from our firewalls. Given this raw data from three different sensor sources we initiated a number of “active response” events to block IP addresses from communicating with the Lab based on the hostile behavior detected – the actionable information. In August and September, 2007 ANL blocked 1,994 and 1,679 IP addresses respectively of which 668 and 592 of the addresses were unique. Repeat offenders are common for us.

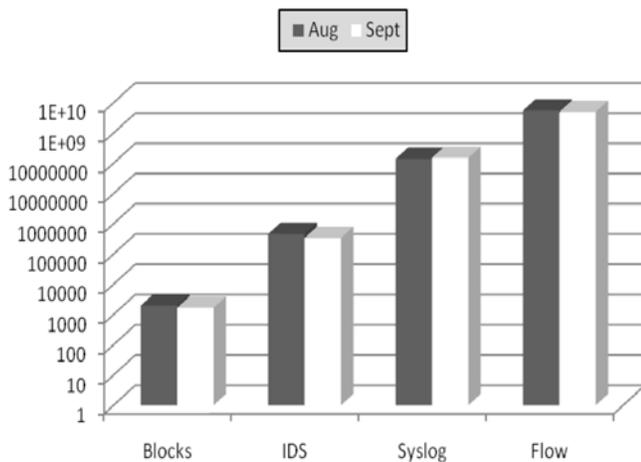


Fig. 1 – Active response events (blocks) relative to raw data records from different sources. Data was collected at Argonne National Laboratory.

The premise of our initial federated model for cyber security was that we could work smarter to mitigate the cyber threats facing us by sharing information about the IP addresses that have been hostile to computers at our site. By sharing this information we hoped to increase our responsiveness; improve our data correlation methods; enable the acquisition of pertinent background information in near real-time; and to improve our reaction and response to cyber security incidents.

To that end we built a framework to share unclassified cyber security related information amongst a trusted community via a secure message passing infrastructure. This included a limited function web service that allowed for the upload and download of formatted data files. The file formats were XML, and included a Bachus-Naur Form (BNF) type definition to allow sites to effectively share information. The message formats were developed based on the Intrusion Detection Message Exchange Format (IDMEF) – RFC 4765.

III. TECHNICAL APPROACH

ANL has been leading a grass roots effort to create a framework for sharing information including a repository, a communication scheme, an XML/BNF definition of messages to be transferred, and the initial formation of federations (a trust community).

A. Repository Web Server & Communication Scheme

The framework we developed is comprised of two repositories and an arbitrary number of participating sites. The two repositories are functionally equivalent and are intended to hold identical copies of the data being uploaded from the sites. Two were used merely to increase availability and minimize disruption based on localized network outages. Participating sites (organizations) would collect information about hostile IP addresses (their actionable information based on local detection and analysis methods), format the information into an XML file based on the BNF template definition, encrypt the XML file according to which sites they chose to share their information with, and then upload the encrypted data file to both repositories. The local sites retain complete control of what information they share, and whom they share it with – they encrypt the file specifically for the recipients they choose. The repositories will accept file uploads, and downloads, from at most two unique IP addresses from each of the participating sites. The frequency at which data is uploaded to the repository, or downloaded from, is up to the discretion of the each site.

The repository which is actually a web server with customized scripts supports an RSS model so that sites can detect when new data is available from another site. Argonne currently uploads information on hostile IP addresses hourly; this time could be reduced or alternatively we could upload a single file for each IP address that we block – in very near real time. Other federation members have chosen to upload their information on a daily basis. The repository is agnostic relative to the trust relationships established between sites. We encourage organizations to develop trust relationships and to exchange PGP keys independently. Sites can share data with one other organization, or they can share with numerous organizations.

B. XML/BNF Definition

We have a working template definition, referenced as version 1, which allows for the exchange of information on hostile IP addresses that were detected at the local sites. The initial template definition was developed in conjunction with other participating Research & Education organizations, and a commercial IDS vendor. We used RFC 4765 – IDMEF as a reference for creating the template. The current definition supports both IPv4 and IPv6 addressing models. It is organized in two parts – one for current information (why we blocked an IP address today), and one for historical information (what the IP address did last week, last month, or last year). For instance, the very first time we block an IP address we would report it, citing the offense or behavior that caused the site to take action. If a month later we blocked the

same IP address again we would report on the new current offense, as well as provide information about the past behavior. Internally at Argonne we refer to this as our anti-host database. Since 1999, we have blocked IP addresses over 1.2 million times, of which there have been ~550,000 unique IP addresses. Acknowledging that each site will have varying local thresholds and criteria for defining hostile behavior, the framework provides for each site to define, and share, their local analysis schemes and local thresholds for detecting hostile behavior.

Today, Argonne has chosen to not block an IP address based solely on malicious behavior at another site. We are not ruling this out, but have not implemented this capability at this time. Instead we have chosen to incorporate the knowledge of hostile behavior at other sites in our *history proportional blocking* concept. Under this scheme we increase the minimum block time for hostile IP's based on their historical behavior.

C. Fostering Federations

Forming alliances or federations is not the most technically challenging part of this effort, but it is certainly one of the biggest political/cultural challenges. Obstacles for participation range from regulatory restrictions to a simple reluctance on the part of an organization to acknowledge, much less actively report on, hostile behavior at their site. However, one on one personal trust relationships frequently exist, and people will often call upon a trusted colleague to support forensic or other cyber related tasks. We encourage organizations to participate even if they are in a listen only mode. It is also worth noting that sites can simultaneously be a member of one or more federations. Different information can be shared, at the sites discretion, within different federations.

IV. RESULTS AND DISCUSSION

The technology aspects of this project were not terribly challenging; we had a working system for sharing data in a fairly short period of time. Reaching a consensus on the definition (template) for the data to be shared took slightly longer. Getting around the cultural and political reservations associated with sharing this type of data has been the slowest part. We are gaining traction within the Department of Energy Labs as well as with members of the REN-ISAC community – a collection of colleges and universities served by Internet 2.

A. Adapting for other Cyber Functions

Once we had a working system that successfully transferred data between sites on a regular (site determined) schedule, it did not take us long to ask *what if* we used this message passing framework in a query response manner. Given a secure near real time communication infrastructure – we only needed to define some new message templates in order to support a wider range of cyber security functions.

Consider questions like the following that could be encoded and used within this framework:

- Have you had outbound port {80|25|22|...} traffic to IP

address XYZ ?

- Have you had inbound traffic from IP address ABC ?
- Have you had inbound udp traffic on port {#} from IP address xyz ?
- Is the following network {a.b.c.d/mask} active (being routed) at your site ? [This is a question that I would like many of the CERT organizations to ask me, *before* telling me about a problem machine at my site.]

Or more complex interactions could be created, such as:

- Could you send me any flow records associated with outbound port {80} traffic to IP address {xyz} on date {mm/dd/yyyy} ?
- Could you add the following IP address {abc} to a watchlist for in|out|both-bound traffic on port {#} ?
- Advise that you block in|out|both-bound traffic to IP address {xyz} – reference {id} for further information

We have been developing these concepts in our version 2 and 3 template definitions. For those wishing to participate in this effort additional information about this project can be found here [8].

B. Promoting Commercial Collaboration

In addition to encouraging organizations to join a federation and share information, we have also been working with commercial vendors to explain and justify support for the creation of the customized XML data files “out of the box”. One approach to this would require a template editor – so that the output files could be tailored to match the desired definition template. Again the definition templates are derivations of RFC 4765.

An additional capability that we are lobbying for is the ability to automatically create the formatted output files – either periodically or based on an event, E.g. during the block/shun of an IP address for hostile behavior, or on an hourly basis. So far we have been working with Cisco Systems and Arbor Networks to develop this capability.

V. BENEFITS

The federated model for cyber security and the underlying capability for secure near real time message passing, offers numerous benefits to those working in the area of cyber security. Specifically, we hope to significantly speed up our OODA (observe, orient, decide, and act) loop, which is well defined here [7], for responding to a wide range of cyber related events.

- Observation and orientation improves with information that can be confirmed and correlated across multiple organizations or sites.
- Deciding on a next step improves with the ability to request, and receive, information from another organization in near real time. Even when we have good personal contacts between organizations – synchronizing schedules across time zones and with busy calendars is not terribly efficient.

- Our ability to act or react is improved dramatically using a secure near real time message passing infrastructure. Being able to send a flash advisory that could be implemented without human intervention is a powerful capability – within certain communities.

Today we share information with the US-CERT in our federated model initiative, and early feedback indicates that our data is a good complement to the Einstein program – described here [9]. With some future algorithmic development and enhancements with additional message templates – we see numerous new capabilities becoming available to meet the cyber challenges that lay ahead.

ACKNOWLEDGMENT

The author wishes to thank Tami Martin and Aashish Sharma for their assistance in defining and developing the concept for this framework; and to Tami Martin for her diligence in bringing this system to operation, and her on-going support.

REFERENCES

- [1] “The National Strategy to Secure Cyberspace,” *The White House*, Washington D.C., February 2003.
- [2] D. Hardt, Identity 2.0, O’Reilly Open Source Convention Keynote Speaker, Portland, OR, August 2005.
- [3] P. Burrows, “Cisco’s Evolving IP Pitch,” *Business Week*, October 24, 2005.
- [4] R. Tehan, “Data Security Breaches: Context and Incident Summaries,” *Congressional Research Services*, May 7, 2007.
- [5] L. Tung, “Infamous Russian ISP behind Bank of India hack,” ZDNet Australia, <http://news.zdnet.co.uk>, September 4, 2007.
- [6] B. Trammell, E. Boschi, “From NetFlow to IPFIX,” presented at Nanog 41, Albuquerque, NM, October 15, 2007.
- [7] http://en.wikipedia.org/wiki/OODA_Loop.
- [8] <http://www.anl.gov/it/federated>.
- [9] US-CERT, “Privacy Impact assessment EINSTEIN Program – Collecting, Analyzing, and Sharing Computer Security Information Across Federal Civilian Government,” September 2004.