

Peer-to-Peer Enclaves for Improving Network Defence

David W. Archer and Adam Wick

“If we were sincerely looking for a place of safety, for real security and success, then we would begin to turn to our communities.”

Wendell Barry

Author, critic, and farmer

Information about cyberthreats within networks spreads slowly relative to the speed at which those threats spread. Typical "threat feeds" that are commercially available also disseminate information slowly relative to the propagation speed of attacks, and they often convey irrelevant information about imminent threats. As a result, hosts sharing a network may miss opportunities to improve their defence postures against imminent attack because needed information arrives too late or is lost in irrelevant noise. We envision timely, relevant peer-to-peer sharing of threat information – based on current technologies – as a solution to these problems and as a useful design pattern for defensive cyberwarfare. In our setting, network nodes form communities that we call enclaves, where each node defends itself while sharing information on imminent threats with peers that have similar threat exposure. In this article, we present our vision for this solution. We sketch the architecture of a typical node in such a network and how it might interact with a framework for sharing threat information; we explain why certain defensive countermeasures may work better in our setting; we discuss current tools that could be used as components in our vision; and we describe opportunities for future research and development.

Introduction

Current approaches to network defence are limited in scale and speed by the limited availability of skilled human operators and the inability of these operators to share information and work at cyberspeeds. We believe that network defence can be scaled out and enhanced through timely sharing of relevant information about common threats that are reasonably expected to affect a host in the near term. The goal is for critical information about relevant threats to be shared rapidly enough that the information is useful to a recipient in preparing a timely defence that adapts to current threat conditions. Unfortunately, in current practice, such timely and relevant sharing does not typically occur.

Current approaches to sharing threat information make use of Internet-wide threat feeds that are commercially available. Such feeds typically propagate threat inform-

ation with delays on the order of minutes to hours, though Microsoft's Cyber Threat Intelligence Program under Windows Azure promises updates as often as every 30 seconds (tinyurl.com/mslnv2h). In contrast, the attacks being reported may move from one host to another in milliseconds. In addition, commercial threat feeds report on a wide variety of threats, requiring that consumers filter and prioritize threat information before acting on it. Thus, a rapid, autonomous improvement in defensive posture against imminent threats is currently prevented both by delay in the availability of threat information and by delays due to the filtering and prioritizing of that information.

In this article, we articulate a novel design pattern for defensive cyberwarfare: *enclaves* of cooperating hosts that use autonomous, timely, peer-to-peer sharing and exploitation of relevant threat information to solve these (and other) network defence problems. We begin

Peer-to-Peer Enclaves for Improving Network Defence

David W. Archer and Adam Wick

Central to our approach is the timely sharing of threat information among hosts. There is a significant prevalence of cyberattacks in which hosts sharing a network or other resource are subject to the same attack in rapid succession. This prevalence may result from a frequent structural choice in the current Internet, where sub-nets tend to contain hosts running similar operating systems with similar application loads (Chen et al., 2003: tinyurl.com/kw56ckh; Abu Rajab et al., 2005: tinyurl.com/l5yvezem). As adoption of IPv6 continues to expand, these problems may get worse because system administrators organize their machines into logical sub-nets that are globally addressable. A variety of advanced persistent threats (APTs) typify this iterative attack pattern. Hutchins and colleagues (2011; tinyurl.com/8qhsj5u) note that, “APT actors, by their nature, attempt intrusion after intrusion.” For example, RSA Security's well-known network breach in 2011 (tinyurl.com/mvk2yjhl), which started with a phishing campaign targeting two groups of employees, subsequently targeted many hosts on the RSA corporate network. Similar behaviour is seen in attacks that affect multiple networks that share characteristics of interest to an attacker. For example, Operation Aurora attacked several technology and defence corporations in 2009 (tinyurl.com/np89339), methodically exploiting the software configuration management subsystems on hosts throughout target networks. The literature shows that rapid sharing of information about such threats can be an effective enabler of improved defensive posture. For example, Weaver, Staniford, and Paxson (2004; tinyurl.com/kaxwhu3) show that defence against scanning worms can be improved by rapid communication of threat information among autonomous defensive elements.

Enclaves use a *peer-to-peer* paradigm rather than a client-server approach for sharing threat information. By peer-to-peer sharing, we mean sharing performed autonomously by participating hosts, thereby avoiding human intervention or use of a central network resource. Our motivation for this choice is that centralized resources such as server-deployed enterprise applications are attractive targets for attack (tinyurl.com/q6xyuhw), and security applications are especially attractive targets (tinyurl.com/m7zk9dn). In addition, the client-server approach requires an explicit build-out of server resources as well as investment in system administration effort, while peer-to-peer resources scale naturally as new peers are added, and they require no central system-administration resources. In addition, peer-to-peer architectures are more robust than client-server architectures due to lack of single points of failure (Lua

et al., 2005; tinyurl.com/kygjjen). Conversely, peer-to-peer systems have inherent security weaknesses, because each peer is controlled by the host on which it runs. Thus, “bad actors” – peers providing irrelevant or distracting threat information to peers in an enclave – can adversely affect peer-to-peer networks more easily than client-server arrangements. We recognize that our proposal of a peer-to-peer approach requires care in authenticating and trusting peers, and we address this problem in the next section.

Our approach limits timely sharing to threat information that is likely to be immediately relevant to peers, because we expect the reasoning capability in autonomous cyberdefence elements to be limited. Our goal is that shared information should be actionable without substantial filtering, interpretation, or prioritization. For example, if a threat manifests a port scan, relevant information shared among peers might specify the ports scanned, the operating system and version of the attacked host, and the applications installed at the scanned ports on that host. A peer host might exploit this information for example by applying simple rules to block the reported ports for a specified time period if the host was running the same operating system as reported.

Component Technologies for Peer Agents

In this section, we break down our envisioned system into five concrete, manageable components: a detection system, a communication language, opt-in communication channels, secure authentication and trust mechanisms, and dynamic countermeasures. In each subsection, we describe the existing technologies that may begin to meet the needs of these components.

Detection

We expect enclaves to leverage existing host and network-intrusion-detection systems, as shown in Figure 1. Host Intrusion-Detection Systems (HIDSs) look for internal changes to a system; examples include Tripwire (tinyurl.com/d4pty), which monitors file changes, and OSSEC (ossec.net), which checks system logs and registries, and looks for rootkits. More traditional anti-virus tools, such as Norton Internet Security (tinyurl.com/23shn7p), also may be considered in the HIDS category. Network Intrusion-Detection Systems (NIDSs) such as Snort (snort.org) detect bad behaviour by sniffing packets on attached networks. Other technologies such as firewalls may also detect and report threats in a timely way.

Peer-to-Peer Enclaves for Improving Network Defence

David W. Archer and Adam Wick

Communication language

Communicating threat information among peers requires that both sender and receiver use the same language. The semantics of such a language can be captured in one or more ontologies, whereas syntax can be captured in a language specification. An ontology in this context is a machine-usable specification of the entities, concepts, and relationships in a domain of discourse. Orbst, Chase, and Markeloff (2012; tinyurl.com/kbrhrf) describe the development of ontologies for cybersecurity at the MITRE Corporation (mitre.org) as part of an effort called Structured Threat Information eXpression (Barnum, 2013; tinyurl.com/kdov4c8). Assured Information Security (ainfosec.com) is developing an ontology for describing malware behaviour and cyberenvironments (Taylor and Hall, 2013; tinyurl.com/m5yplkz).

Communication channels

The channel for transmitting timely, relevant cyber-threat information must be: *decentralized*, to make it difficult to attack and more robust than a single point of failure; *reliable*, to ensure that threat information is delivered; *timely*, to enable peers to react at cyberspeed; and *efficient*, to minimize impact to normal business logic. Publish-subscribe middleware, such as implementations of the Data Distribution Service (DDS; portals.omg.org/dds/), are designed with such properties in mind, and thus may be suitable choices for communication among enclave members. DDS family members are fully distributed without need for brokering of mediation between publishers and subscribers. Reliability and timeliness have been demonstrated in several DDS implementations such as OpenSplice Community (tinyurl.com/p8pw24g) and OpenDDS (opendds.org).

At least one communication channel is in development specifically for transport of cyberthreat information: the TAXII sharing service (taxii.mitre.org) being developed in conjunction with MITRE's STIX language. DDS or TAXII are existing technologies that demonstrate how the content-distribution mechanisms we envision are both feasible and practical.

Authentication

A fundamental issue in communicating threat information is the degree to which a consumer of the information should trust what is communicated. Establishing trust requires action on at least two levels: authentication of transmissions, and trust in their contents. Communication and authentication standards for data transmission are well understood in general. We expect that typical protocols such as the Secure Sockets Layer

(SSL; tinyurl.com/c9jdg), or similar protocols that achieve efficient data transmission and encryption may be sufficient. Message authentication and other techniques may also be applied to authenticate threat data.

Enclave peers will need to guard against malicious or broken peers, which may correctly implement data-transmission policies but may also transmit information counter to enclave interests. This problem is the subject of ongoing research in the general case, but mechanisms based on reputation systems seem a likely solution to the problem (Resnick et al., 2000; tinyurl.com/km43orc). In a reputation system, a node keeps track of reputation data from its peers. As an example, node A may keep track of threat information provided by each of its peers. If a threat reported by one peer, B, is correlated by another peer, or system-countermeasures report stating that the threat became reality, then A may increase its "opinion" of B. If a threat never materializes and no other peer mentions it, A may decrease its opinion of B. Once generated, this reputation data can be used to quickly and easily weight threat information introduced to a node. In the long run, such reputation systems may also be used to remove peers that do not provide good, relevant data to the node and to find new peers that can provide such information.

Countermeasures

Enclaves offer a unique opportunity for dynamic adjustment in defensive posture. The timely exchange of relevant threat information allows hosts to take dynamic defensive action, and then revert to less aggressive defensive postures when threats pass. In contrast, current network defence techniques rely on static defensive postures that may impose hardships on users and system administrators. For example, countermeasures that automatically block network access (in part or in full), restrict account privileges, back up or obscure sensitive data, or temporarily disable ports can disrupt business processes and reduce utilization of computing resources if used consistently. However, if deployed for short time periods surrounding an attack, such disruption can be minimized.

Additional countermeasures may be available that are suitable for short-term, dynamic deployment, but might impose too much disruption for static deployment. Through recent research at Galois, Inc. (galois.com), we demonstrated the use of virtual-machine creation on-the-fly as a network defence technique called CyberChaff. Upon detection of an imminent threat, a CyberChaff device deploys a significant num-

Peer-to-Peer Enclaves for Improving Network Defence

David W. Archer and Adam Wick

ber of lightweight virtual machines onto a network, with network configurations that can be tuned to appear as particular operating systems running standard sets of services. By doing so, CyberChaff has the potential to obfuscate the network structure in order to confuse attackers. In addition, CyberChaff's virtual machines can serve as honeypots (tinyurl.com/37scmk), gathering information about patterns of cyberthreats to provide greater insight into the attackers' identities, goals, and preferred attack patterns. The Chaff Controller, shown in Figure 1, illustrates how CyberChaff fits into our notional enclave peer architecture.

Other recent research at Galois demonstrated a network stack called Ditto, which can allow a host to falsely display its configuration to external network scans. Using Ditto, a host can appear to be running a different operating system than actually used by the host. Ditto is intended to solicit attackers to waste time by applying exploits that are less likely to succeed because they target incorrect operating systems. The Ditto Obscuration Controller, shown in Figure 1, illustrates how Ditto fits into our enclave peer.

There is increasing interest in using software-defined network routing such as that provided by OpenFlow (openflow.org) for intrusion response. OpenFlow allows hosts to specify policies that classify traffic as belonging to specific network flows and thus enables redirecting of that traffic upon detection. For example, OpenFlow policies might re-direct port scanning traffic from its intended destination to a honeypot. The FRESKO framework (Shin et al., 2013; tinyurl.com/n2z24wv) is a recent system that employs a related approach. Software-defined networking might be included as part of the Programmable Network Adapter shown in Figure 1.

Conclusion

Current approaches to network defence rely on static end-point defensive postures taken by individual hosts that lack timely and relevant information about threats they may soon face; or actions orchestrated by centralized command-and-control systems that receive threat information and adjust postures slowly relative to attacks. Our vision is to change this defensive landscape by enabling the creation of enclaves that are responsive, informed, and armed. In such enclaves, each host dynamically adjusts its own defence at cyberspeed, and all hosts share information about emerging threats with their peers in a timely way. In doing so, hosts can reduce disruption to users and system administrators be-

cause some countermeasures can be deployed dynamically in response to such information instead of statically, and hosts gain the advantage of access to new countermeasures specifically designed for such dynamic deployment. Such enclaves may be localized to a single network or may include hosts from distinct networks owned by organizations that face common cyberthreats. For example, as the Internet of Things (tinyurl.com/5qr2nq) emerges and home networks grow to be more attractive targets, home networks in a physical neighbourhood may face common threats such as drive-by network hacking, and these networks may form enclaves in response.

In this article, we presented a notional architecture for hosts capable of operating in the enclaves we describe, as well as a notional means for these hosts to communicate timely, relevant threat data. For the most part, the key technologies required to create a first generation of such enclaves already exist. However, some key technologies still require advancement, and the pieces must be combined into an integrated whole. We note, in particular, the need for practical, rapid methods for describing and communicating threat information, as well as the need to develop advanced-decisions engines capable of receiving, analyzing, and acting on network threats.

Peer-to-Peer Enclaves for Improving Network Defence

David W. Archer and Adam Wick

About the Authors

David Archer is a Research Program Lead at Galois, Inc., where he directs research into high-assurance methods for large-scale cyberconflict. He holds a PhD in Computer Science from Portland State University in the United States as well as an MS in Electrical Engineering from the University of Illinois at Urbana-Champaign. Dr. Archer's research interests also include efficient methods for computing on encrypted data, and information integration, assurance, and provenance. At Intel Corporation, Dr. Archer was instrumental in the development of the communication network for the ASCI Red TeraFLOPS system at Sandia, and in the development of multiple generations of high-performance server and workstation memory and I/O systems.

Adam Wick directs the Systems and Networking Group at Galois, Inc., where he has worked with DARPA to create advanced network-defence techniques, including CyberChaff and Ditto. He holds a PhD in Computer Science from the University of Utah in the United States, as well as a BS in Computer Science from Indiana University Bloomington. Dr. Wick also has been collaborating with SRI, LG, and others to build secure mobile devices for the United States Marine Corps. Prior to this work, he developed the HaLVM, a lightweight machine for running custom, single-purpose applications in the cloud. In all of this work, he maintains a focus on using next-generation operating system and networking technology to create practical tools for critical systems.

Citation: Archer, D.W. and A. Wick. 2013. Peer-to-Peer Enclaves for Improving Network Defence. *Technology Innovation Management Review*. July 2013: 19–24.



Keywords: network defence, cybersecurity, enclave computing, dynamic cyberdefence, cyber countermeasures, peer-to-peer