Resilient and Self-Organizing Overlay of Collaborative Security Monitors

Jun Li, Virginia Lo, Xun Kang, Dayi Zhou, Dan Rao
Department of Computer Science
University of Oregon
{*lijun, lo, dayizhou, kangxun, rao*}*@cs.uoregon.edu*

Abstract

As attacks toward the Internet become more sophisticated, destructive and widespread, especially those distributed attacks that exploit multiple launching sites, security systems that are deployed in isolation within administrative domains and do not exchange information across those boundaries fail to address those attacks successfully. To overcome these limitations, current distributed monitoring systems have made significant advances in integrating data gathered by monitors distributed throughout the Internet [1]. However, these systems are also far from satisfactory because they often rely on a single central site to analyze the data and use a non-scalable method (unicast) to issue alerts.

Driven by realistic security applications, including our current research in worm defense and open proxy blacklists, we design, evaluate, and deploy a distributed robust communications architecture for Internet-scale security monitoring systems, called *Sequoia*. It consists of a set of security monitors distributed throughout the Internet and supports a rich array of security information communication patterns among them, including gathering, sharing and delivering security-related information in such a large-scale system. In particular, an effective security communication architecture must be: fast in order to quickly react to attacks, scalable to accommodate thousands of participants, adaptive to react to dynamic changes, deployable to be really useful, robust to continue to function in the face of failures, and secure to protect both data in transit and the communication structure itself.

There are three key protocols of Sequoia: Monitor Neighbor Discovery Protocol (**MND**), Distributed Dominator Selection Protocol (**DDS**), and Communication Path Discovery Protocol (**CPD**). Using these protocols, monitors self-organize into a two-level hierarchy on which scalable, fast and trustworthy message delivery can be achieved. We have finished the design for MND and DDS and currently working on CPD.

The goal of the MND protocol is to form a topology-aware flat overlay among monitors, on which every monitor is connected to nearby nodes as its neighbors. A monitor node joins the monitor overlay by contacting known landmark nodes to obtain its coordinates, which are then used to query a directory server for a recommended list of nearby nodes. The monitor then chooses the closest neighbors based on round-trip measurements. Each node can further optimize and maintain its neighborhood relations through local gossiping.

The DDS protocol is used to form a two-level communications hierarchy from the flat neighbor overlay constructed by MND. A monitor in the higher level of this hierarchy (*dominators*) must meet minimum requirements regarding trustworthiness and routing performance. A monitor can choose to apply for a Sequoia-certificate, or *S-certificate*,

from a registry service, certifying this monitor's service type, trust level, public key, and other information. Each monitor in the lower level of the hierarchy (*dominees*) eventually selects one or more higher level monitors; thus, each dominator acts as a hub for a group of dominee nodes to reach the rest of monitors. While improving the scalability, the two-level structure ensures that untrusted nodes will not be able to forward security information for others, providing a robust communication structure. To form the two-level hierarchy, a dominator periodically advertises itself to its x-hop neighborhood, and presents its S-certificate and other qualifications to dominees. As needed, a dominee node can search in its y-hop neighborhood for dominators, selecting those it wishes to utilize based on the dominator's attributes. A caching mechanism is used to reduce message overhead.

The CPD protocol is used to discover multiple delivery paths from one or more senders to one or more destinations, in order to support a rich set of communication modes among monitors, including a publisher-subscriber paradigm. CPD considers both efficiency and security constraints. Every dominee node can reach or hear from other nodes through its dominators, so the path discovery among dominators is the key for CPD. In our ongoing design, we are investigating the possibility of using structured overlays (such as CAN [2], Tapestry [3], or Chord [4]) to organize dominator nodes and handle routing issues among them, with two particular goals: (1) How can disjoint paths be discovered between sources and destinations, and (2) How can paths with high trust values be discovered.

In addition, security of Sequoia and applications of Sequoia are the other two important research issues. The security of Sequoia is critical for its success. Not only must Sequoia protect the security updates exchanged between monitors, but it also must carefully protect itself, including its functions, messages and elements. Sequoia must secure all its functions, including neighbor discovery, monitor overlay construction, and communication paths discovery. Applications that use Sequoia are also important in order to demonstrate the usefulness of Sequoia. For instance, if Sequoia can support multiple worm monitors to share the information about suspicious traffic patterns or validate each other's discoveries, a more effective worm detection system can be developed.

In summary, the need for an architecture in order for security monitoring systems to gather, share, and deliver information in a large-scale system without centralized control has never been more compelling. Sequoia's use of self-organized topology-aware structure to support rich, fault-tolerant communication is an important step towards this goal. We expect to further advance this research by applying Sequoia to specific monitoring applications.

References

[1] R. Robbins. Distributed intrusion detection systems: An introduction and review. January 2002.

[2] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content addressable network. In *Proc. ACM SIGCOMM*, Aug. 2001.

[3] B. Zhao, L. Huang, J. Stribling, S. Rhea, A. Joseph, and J. Kubiatowicz. Tapestry: A resilient global-scale overlay for service deployment. *IEEE Journal on Selected Areas in Communications, Special Issue on Service Overlay Networks*, 22(1), January 2004.

[4] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan. Chord: a scalable peer-to-peer lookup service for internet applications. In *Proc. ACM SIGCOMM*, pages 149–160, Aug. 2001.