Fast Abstracts: Suppl. Proc. International Conference on Dependable Systems and Networks, DSN04, Florence, Italy, June 2004

# Reputation-based Security in Routed Networks

Scott E. Coull
*Rensselaer Polytechnic Institute*
*110 Eighth Street*
*Troy, NY 12180*
*coulls@cs.rpi.edu*

Boleslaw K. Szymanski
*Rensselaer Polytechnic Institute*
*110 Eighth Street*
*Troy, NY 12180*
*szymansk@cs.rpi.edu*

With the increasing popularity of the Internet, and the continued reliance on its communications abilities, comes the need to protect this infrastructure against increasingly complex and destructive attacks. While it was once almost inconceivable that malicious users could render the modern form of the Internet useless, it is now quite possible that the amount of Internet traffic produced by some worms and viruses could do just that. The Code Red worm, at its peak number of infections, left parts of the Internet totally flooded with traffic while it was trying to infect new hosts. Even months after the first appearance of the Code Red worm, there were still hundreds of thousands of infected hosts sending traffic to random parts of the Internet, attempting to infect new hosts. These new types of widespread worm and virus attacks are a clear and present danger to the Internet infrastructure on which we rely so heavily. There are currently no measures in place to protect this vital resource in the case of a major attack or worm outbreak.

Unfortunately, the Internet is built upon protocols that require implicit trust among all parties involved. For instance, any given host must trust the intermediary routers to properly route packets and not to perform malicious activities in order for that host to communicate with other hosts on the Internet. Currently, the burden of securing the Internet has fallen upon the hosts, and not the Internet architecture itself. Various add-on protocols and programs have been used to ensure privacy and secure communications between parties across the Internet, as well as to protect these hosts against potentially malicious attacks. While this may work to some extent, the amount of knowledge and time required of the average end-user is increasing rapidly with the addition of newer security features like virtual private networks, virus scanners, personal firewalls, and home networking. In most cases, these home users, who often have bandwidth capabilities comparable to most small or medium sized businesses, do not want to or know how to update their systems, or correctly configure their hardware. Virus writers and crackers are continually targeting end user systems for these very reasons.

All the above reasons lead us to believe that the only way to secure the Internet against wide scale attacks is to add security to the Internet architecture itself. This would allow the Internet to stop most of the major attacks early and to do it at points in the network that would significantly reduce the impact of these attacks on the network at large. To the best of our knowledge, there has been only one other attempt to develop a protection mechanism for computer networks that operates within the network architecture. The Pushback method, developed by John Ioannidis and Steven Bellovin, is a mechanism that allows routers to cooperate in the control of excessive congestion caused by denial of service and distributed denial of service attacks [1]. In the Pushback method, the router that is being inundated by denial of service traffic sends a message to upstream routers telling them to begin limiting the rate of traffic that has been determined to be a denial of service attack. This spreads the responsibility of traffic rate limiting among a number of high bandwidth connections to enable the router to continue to service legitimate clients. Unfortunately, this solution allows for a Byzantine General attack in which the router initiating the Pushback can discriminately limit the rate of traffic through upstream routers in order to degrade the quality of service not only to that specific router, but also to others that may receive traffic through those upstream routers.

To overcome this Byzantine General problem, our system will extend the concept of Pushback to include a reputation management and propagation system to inform neighboring routers of misbehavior, and to allow each router to take a course of action commensurate with the reputation of the misbehaving host or network. In order to provide this functionality, we propagate the reputation from the node being attacked to all neighboring nodes within a certain radius by using a gossip-based protocol. These gossip-based protocols require each node that receives the reputation to send it to a fraction of its neighbors, who will then in turn send it to their neighbors, until the radius of the reputation propagation is reached. This reputation should be stored at each node separately, and computed based on the distance from the source of the

reputation. The reputation should also be allowed to decay gracefully over time. Once the reputation has reached a certain cutoff value, communications from the host or network whose reputation is poor can be blocked or degraded accordingly.

We developed a prototype system in simulation, which shows the usefulness of this method in stopping malicious traffic in networks such as the Internet. The simulation has been shown to stop the propagation of a worm, similar to the Code Red worm mentioned above, without inundating the network with excessive traffic from the worm. The simulation, which was created using SSFNET [2], contains six local area networks that contain four worm-susceptible file server hosts each. A seventh network contains only one host, a file server host that has been infected with the worm. These seven networks each contain a single gateway router that is used to connect the local area network to the core network of four routers connected in a full-mesh topology. The worm propagates by randomly choosing a host from the available address range and attempting to infect that host. We assume some type of virus or intrusion detection capability on the host being attacked, such that when it receives an attempt to infect it, it is able to trigger a reputation propagation to its local gateway router. We have also added a 20% false positive rate for legitimate file server traffic to be marked as an infection attempt. The simulation showed that core routers effectively block traffic from the infected file server when their reputation levels for the infected file server rose above the trustworthiness cutoff due to the reputation messages sent by the six local area networks.

There are a number of interesting implications from the use of this system in large-scale routed networks. The first, and most poignant, is that it provides a method to stop malicious communications immediately at the router that is being attacked. The second is that it provides a sort of warning to neighboring routers. Once some percentage of the routers in a given radius are attacked, communications from the attacking host or network are blocked from the entire neighborhood because of a convergence in the reputations such that the reputation of the attacker becomes greater than the cutoff value at every node in the neighborhood. When this convergence occurs, the other nodes in the neighborhood have been saved from attack due to the reputation gathered from surrounding nodes that were attacked. The localized approach makes this a lightweight system that is very resistant to Byzantine General-type attacks. Suppose a single, or small number of, routers attempt to force the blockage of traffic from a specific host at upstream routers. This will fail because there have not been enough reputation propagations for the specified host in the neighborhood of the upstream routers. This would mean that the upstream routers still have a positive reputation for the specified hosts, and therefore traffic should be unabated. Probably the best feature of this system is that it is independent of all attack detection technologies. A security administrator can easily trigger a reputation propagation for an attacking node based on logs they have reviewed, or an automated intrusion detection system can be made to trigger a reputation propagation when it is significantly confident that an attack is occurring, or has occurred.

While this system holds a lot of promise for the security of the Internet, it also holds many possibilities of misuse if it is poorly designed. We must first consider the limitations of the devices, which will be running this system, routers. Routers have very limited memory and computational capabilities, and the system must be designed such that it does not hinder the focus of a router, that is routing. Additionally, the system itself must be designed so that it could not be used maliciously to deny service to other hosts on the network. There are a number of parameters that direct the actions of the reputation system, and with the correct settings, these parameters can ensure a fair and equitable system. For instance, we can use the cutoff and the reputation determination formula to ensure that a majority of surrounding routers must concur that a given host is malicious before enacting any type of blocking or degradation of service. We can also ensure that the radius up to which a reputation propagates is of a size small enough to ensure that any one node cannot send a reputation that will propagate to every node in the network. We must also consider scenarios in which a given node will continually propagate false reputations for some other node in the hopes of flooding neighbors such that they will begin to deny service to the falsely accused node. This situation can be dealt with by setting a period of time that a node that has propagated a reputation has to wait before propagating a reputation again. We can set this time such that the reputation has decayed enough to make it impossible for any one, or small number of nodes to subvert the reputation system.

Though there are a number of challenges, we believe that they can be overcome with proper experimentation and rigorous examination. This system adds a layer of security to not only protect hosts against attacks, but also to protect the Internet infrastructure itself against large-scale attacks that could possibly lead to wide scale denial of service throughout the Internet.

[1] Ioannidis, J., Bellovin, S.M. Proceedings of Network and Distributed System Security Symposium, Catamaran Resort Hotel San Diego, California 6-8 February 2002.

[2] Scalable Simulation Framework - http://www.ssfnet.org