CENTRIPETAL
NETWORKS

Threat Surface Reduction: Prevent known cyber criminals from attacking your networks

A Technical White Paper
Centripetal Networks, Inc.

Threat Surface Reduction:
Prevent known cyber criminals from attacking your networks

A Technical White Paper
Centripetal Networks, Inc.

The Problem

Your organization's network is continually attacked by cyber criminals that steal sensitive company data or employee and/or customer login credentials and personal information, then sell it to other criminals. They attack e-commerce web sites, causing customers to experience poor service, which may cause them to flip to other business providers. If intrusions are discovered, it is after damage is done. Most organizations, even the largest ones that invest heavily in cyber security, are not fully aware of the threat, have little ability to stop attacks before they cause damage and cannot directly measure the economic impact of the attacks on their businesses. Conventional cyber defenses are mostly ineffective, inefficient and expensive. Is there any hope that this problem can be solved?

Yes. Critical intelligence information about many cyber criminals is well known, namely, the Internet addresses of the cyber criminals' computers, aka "threat hosts." These Internet addresses are included in every Internet packet that attacks a network. A straightforward defense would include the following actions:

1. **Inspecting (filtering) all Internet packets entering and exiting the network;**
2. **Examining packets for Internet addresses of known threat hosts; and**
3. **Identifying and isolating suspicious packets before they can enter or exit the network.**

This defense method is known as threat surface reduction. Threat surface reduction efficiently and effectively stops known cyber criminals from attacking networks. This raises the following questions:

Q: Why isn't threat surface reduction used to protect networks?

A: Conventional cyber defenses cannot scale and adapt to the size and dynamics of the threat.

Q: Are there any emerging solutions that can handle the cyber threat?

A: Centripetal Networks' RuleGate® packet filters and Advanced Cyber Threat (ACT) intelligence service use a highly scalable, adaptive threat surface reduction defense to shut down attacks before they damage your network.

Keep reading to find out how cyber criminals attack networks and how cyber defenders can stop them.

Threat Surface Reduction

The collection of Internet connections between the Internet addresses for all threat hosts – both well known and unknown – and the Internet addresses for your network is called the "threat surface." The collection of Internet addresses at the perimeter of your network is called the "attack surface," because attack packets can be sent between your network's Internet addresses and the Internet addresses of threat hosts, i.e., your Internet addresses compose the surface through which threats can attack you.

Your network's threat surface may be visualized as a colored disk (Figure 1A below), constructed as follows: Internet addresses can be represented as all the integers from 0 to approximately 4 billion (for IP Version 4 [IPv4]). Draw a number line with endpoints of 0 and 4 billion and then connect the two endpoints of the number line to form a circle.

In Figure 1A, the circle – the disk's perimeter – represents the Internet address surface. Place a dot in the center of the large circle. This dot represents your network's attack surface. Now, draw a straight red line between each threat host's Internet address, a point on the large circle, and the dot in the center. Each red line represents a potential communication – an attack – between a threat host and your network. Lines drawn between the contiguous threat host addresses, which may represent a network operated by a cyber crime organization or a hostile government, create red sectors in the disk. Fill out the entire disk by drawing a white line between all other (non-threat) Internet addresses and the center dot. The white portions of the disk represent safe, allowed communications between your network and all non-threat Internet addresses. The red portions of the disk represent the threat surface for your network.

Threat Surface Reduction:
Prevent known cyber criminals from attacking your networks

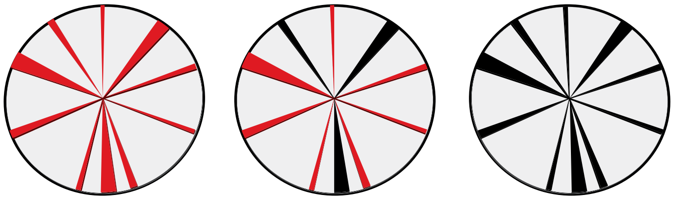

Figure 1A Figure 1B Figure 1C

Figure 1A: A communications disk for a network. The disk's perimeter represents the Internet address space; the disk's center represents the Internet addresses of your network, which is your attack surface. Lines between the center and the perimeter represent all possible communications between the Internet and your network. Red lines and sectors are the threat surface of potential attacks on your network; white lines and sectors are safe communications.

Figure 1B: Blocking communications between threat hosts and your network reduces the threat surface. Blocking a threat communication – an attack – is shown in the diagram as a color change from red to black. In Figure 1B, some, but not all, of the potential attacks have been eliminated.

Figure 1C: When the threat surface is fully reduced by blocking all threat communications, all red lines and sectors are colored black, and your network is fully secured from Internet attacks.

Conceptually, one way to protect your network is to reduce your network's threat surface by eliminating the red from the disk. Suppose we have some method for eliminating a threat by blocking communications between a threat host and your network. Visually (Figure 1B), we represent a blocked communication by drawing a black line over the red line. Each blocked communication, therefore, reduces the red in the disk. Ideally, all threat communications are blocked, i.e., all red is changed to black, and your network's threat surface has been cleared of all known threats (Figure 1C), i.e., the disk is colored white and black only.

The most efficient way to block communications between threat hosts and your network is to apply packet filters to network links – such as your network's Internet access links or your ISP's ingress points and peering points – that transport the packets composing the communications to be blocked. Packet filters (a) examine each packet's headers for values of source and destination IP addresses, and source and destination ports (applications, e.g., port 80 for web applications), and (b) compare the header values to filter rule databases (called network security policies). If a match is found between a rule in the policy and the packet header values, then (c) apply an action – BLOCK (drop) or ALLOW (forward) – to the packet that is specified by the matching rule.

Threat surface reduction is a conceptually simple method for eliminating threats to your network. In practice, threat surface reduction may be achieved by using packet filters to block threat communications. So, why isn't threat surface reduction a standard and ubiquitous method for cyber defense?

Conventional Cyber Defenses Are Not Effective at Threat Surface Reduction.

There are several obstacles that cause conventional packet filtering devices – network firewalls, routers (access control lists), intrusion prevention systems (IPS), and web proxies – to be ineffective at reducing threat surfaces:

Lack of scalability is the primary obstacle. Well-known threat hosts – spam servers, botnet controllers, malware servers, bogons, spoofers, scanners, hostile governments, cyber criminal networks, onion routers, etc. – span several hundred thousand Internet addresses and subnet prefixes. The size of the complete threat, both well known and unknown/suspected, is certainly larger. In any case, a network security policy that covers the well-known threats will consist of several hundred thousand packet-filtering rules. Conventional filters, however, are limited to applying policies composed of about ten thousand rules or less, which will reduce your network's threat surface by a small fraction of the total threat (e.g., 1 percent or less reduction of red in the disk).

A conventional filter may apply policies somewhat larger than ten thousand rules, but network performance, measured by latency and packet loss, degrades to unacceptable levels that users will not tolerate. Therefore,

Threat Surface Reduction:
Prevent known cyber criminals from attacking your networks

the narrow security provided must be balanced against network performance and the resultant quality of experience for the users. In practice, security is almost always compromised to maintain acceptable network performance.

Lack of adaptability to the rapidly evolving threat is another obstacle. New threats are identified continually. Filter rule bases should be updated several times per day to afford maximum protection. Conventional filters typically require several minutes to update rules, during which time the filter stops functioning, resulting in loss of service and/or loss of security during an update. Accordingly, even the largest organizations with the most critical network security requirements limit updates to a few times per week, while incurring significant operational costs to do so.

Lack of availability of known threat surface information is another obstacle for most organizations. Conventional cyber security solutions assume a dedicated team of cyber engineers will gather intelligence information, monitor their networks for attacks, and continually execute defensive actions. The size of the threat, however, makes it difficult for the team to manually identify a significant portion of the threat. Even if complete threat intelligence information were available, it is impossible to manually manage the hundreds of thousands or millions of associated filtering rules. The threat is evolving continually and rapidly, i.e., the threat surface is highly dynamic. Most organizations cannot afford to host a cyber defense team. Organizations that can afford a cyber defense team have difficulty staffing the team. There is a critical lack of cyber defense engineers (recognized as a significant national security issue by the US Government). Cyber defense needs to be automated to make it efficient and cost-effective to implement. Security is compromised when cyber defense information is not automatically available, accessible and manageable.

A Complete Threat Surface Reduction Solution: RuleGate® Network Protection System

A comprehensive, threat surface reduction solution is now available from Centripetal Networks, Inc. This solution combines Centripetal's RuleGate® packet filters and ACT intelligence service to attain the necessary scalability, adaptability and availability properties for defending networks against current and future cyber threats. Centripetal's Global RuleGate Manager automates the management of large filter rule sets and the management of the multiple RuleGate devices that are defending your network's security borders. Centripetal's List Agent SDK integrates the ACT service information with other threat intelligence information sources provided by your internal resources or third parties. The RuleGate device uses real-time attack logging capability and the QuickThreat™ visual dashboard application to allow you to know, at a glance, which threats are attempting, but failing, to attack your network **while these attacks are occurring.** CAPEX and OPEX costs are very low compared to the costs of conventional cyber defenses and of the cyber defense team that which operates them. The security value received is greater by 1000X or more.

A RuleGate can apply millions of packet filtering rules – the scale of the Internet threat surface – to the highest speed (100M/1G/10G) network links even when those links are heavily loaded with small packets. The RuleGate devices' performance improves on conventional filters' performance by several orders of magnitude (1000X or more). Latency is insignificant (measured in microseconds) and packet loss is zero; thus, RuleGate devices have no effect on the users' quality of experience. A RuleGate also has the unique capability to instantaneously update million-rule policies without loss of service and without loss of security during a policy update event.

The ACT service fully automates the end-to-end process of collecting threat intelligence in near real-time, translating that intelligence into packet filtering rules, packaging the rules into policies that you can customize to your needs (using the Global RuleGate Manager), and downloading the policies directly into your RuleGate devices. The RuleGate devices immediately enforce the new policies. Your policies may consist of several hundred thousand or even millions of rules, which cover the entire known threat surface and represent an improvement of several orders of magnitude (1000X or more) in cyber security effectiveness. The ACT can deliver updated policies at an interval cycle that you specify.