

Why Leaks Matter

Leak Detection and Mitigation as a
Critical Element of Network Assurance

Table of Contents

Executive Summary

Defining a Leak

How Leaks Happen

Inbound vs. Outbound Leaks

Prioritized Patch Management and Data Leak Prevention

Powerful Compliance Advantages

Leaks in the Broader Context of Network Assurance

Executive Summary

Network Assurance assesses the impact of network change on security, availability and compliance. Modern organizations are discovering that only with a comprehensive set of network facts can true global visibility be achieved and its benefits of reduced network risk be realized. Network leaks or unauthorized connectivity to or from the public Internet are a classic example of the problems that Network Assurance programs uncover. This paper concludes that leaks can pose serious risk to the integrity of systems and data and are a risky exposure that organizations must mitigate on an ongoing basis. The best way to do this is through the comprehensive visibility of connectivity that a discipline in Network Assurance inherently provides.

Defining a Leak

Perimeter defenses are well-tested protective elements that have been used for thousands of years. Instead of protecting each house in a city against invaders, walls were built around the city, and well-guarded gates controlled access to the city.

Often, there were lesser entry points through the walls, for convenience or special uses. These included “postern gates,” which were small entrances far from the main gates. There are numerous tales of cities that fell because their perimeter defenses were subverted by these little known entry points.

Applied to today’s Internet-enabled networks, these unknown or unauthorized entry points are leaks – a means to malicious or unauthorized entry across the network perimeter. Firewalls and intrusion detection systems serve as gatekeepers to defend the network; nevertheless, circumvention can and does happen. Not to be confused with data leaks, which represent the egress of sensitive information from an organization’s control, Internet leaks are unrestricted pathways into and/or out of an organization’s network perimeter.

Reported losses of revenue from worms, viruses, spyware, and other security breaches are in the millions of dollars per incident in large organizations. If one includes additional recovery costs, that figure grows even larger. The full scope of the problem is almost certainly understated, since, for obvious reasons, network exploitations of this nature are kept private, and many go unreported.

Continual changes to the network landscape, including infrastructure, operating systems, and applications can cause organizational security policy and network defense configuration to become

misaligned, contributing to a proliferation of leaks – and it only takes one leak to allow malicious intrusion into a network.

Proactive identification of leaks and exposed network zones allows effective prioritization of remedial resources to prevent network subversions. When combined with the other aspects of a comprehensive Network Assurance program, leak discovery can be a powerful mechanism for comprehensively protecting an organization's network.

How Leaks Happen

The root of some of the most dangerous leaks comes courtesy of manufacturers, according to SANS.¹ The SANS 'Top 20'² identifies the most critical and/or common vulnerabilities in Windows, UNIX, and Linux environments, many of which are likely to create leaking ports. For example, vulnerable ports may be open by default, creating holes that are the easiest and most convenient route for attackers to exploit. The Blaster, Code Red, and Slammer worms can be attributed to these vulnerabilities. In addition, vendor products may be defaulted to enable routing and remote access. If these are not properly secured, the impact of these vulnerabilities can include denial of service attacks, exposure or compromise of sensitive files and data, random command execution on the server, or complete server intrusion. Until all software is one hundred percent hardened by default, it is essential to evaluate the configurations of servers and enable only the features and services required.

At the same time, outside manufacturers cannot take all of the blame. Inconsistent or ineffective internal administration also contributes to leaks in the network. Improperly configured VPNs, routers, and switches all leave networks susceptible to unauthorized access to host or end-point devices. In addition, host computers may be enabled to allow IP forwarding without proper packet filtering configurations. All of the cases reflect the ease at which defense configuration can become misaligned with security policy. In modern decentralized organizations, control over business units and the management of IT assets often becomes disconnected. As an increasing number of network configuration and connectivity decisions are made by more people, disparate management controls combine with inconsistent enforcement of security policies to compromise access control lists that govern connectivity both to and from the Internet.

¹ The SANS™ Institute (SysAdmin, Audit, Network, Security). Bethesda, Maryland. www.sans.org.

² SANS Top-20 Internet Security AttackTargets (2006 Annual Update). The SANS Institute.

A third issue is outsourcing. Outsourcing, whether an organization's entire IT infrastructure or just security services, has become a widely adopted, cost-conscious business practice. However, outsourcing configuration errors, where the perimeter of one or more networks are not secure, can allow "bleeding" of one outsource customer's network into another. This kind of mis-configuration can allow outsiders to find pathways into a network. Not monitoring the security of outsourced connections can have dire consequences and ultimately decrease security, compliance and availability of core business infrastructure – the very areas that outsourcing is designed to optimize.

Lastly, the proliferation of unknown and rogue connections creates a breeding ground for leaks. Common sources of rogue connections include acquired business units that do not have accurate network documentation, divested business units where connections were not terminated in accordance with the new business structure, unsanctioned wireless network connections, and remote network connections established by even the least technically-savvy end users such as a wireless access point that can be easily acquired at any electronics super store.

Inbound vs. Outbound Leaks

Leaks can allow traffic to bypass security provisions into and out of a company's enterprise network. Hence, they are generally referred to as "inbound" and "outbound" leaks, respectively. These unsafe passageways provide hackers, worms, viruses, and other unwanted web flotsam free access to enterprise resources.

Inbound Leaks

Most companies restrict inbound traffic to minimize the company's vulnerability to worms, real-time hacking, and other attacks. Inbound leaks almost always pose a significant threat. Detection of an inbound leak is an indicator that security controls have been bypassed. Every security policy will have guidance to block unauthorized connectivity since the advent of the firewall and DMZ, however these solutions are only as effective as their individual configurations.

Outbound Leaks

Companies vary in how stringent they are about what exists inside their networks. While some feel that if they control inbound traffic, they are sufficiently protected, others endeavor to maintain higher control of outbound traffic to help prevent internal infections from leaving the company and attacking elsewhere, particularly in today's environment of increasing insider threats. These are scenarios that could result in negative publicity or legal consequences. Reducing the number of outbound leaks significantly minimizes the likelihood that an intruder can compromise a system

within a network and launch further attacks on the outside – resulting in a situation in which the company’s system appears to be at fault.

Outbound leaks are often seen as less critical than inbound leaks. Nevertheless, organizations often use the lists of outbound leaks to validate stated corporate security policies used by divisions or sub-groups within the organization. Outbound leaks also show the Network Address Translation (NAT) devices in the network. If multiple rows in the table have a corresponding IP address, then that is probably the IP address of the device that is performing NAT for a portion of the network.

Methodology

Network Assurance scans should test for both inbound and outbound connectivity to and from the Internet via multiple protocols including User Datagram Protocol (UDP), Internet Control Message Protocol (ICMP), Simple Network Management Protocol (SNMP), Domain Name System (DNS), and user-specified Transmission Control Protocol (TCP) and UDP ports. From there, reports should be generated to clearly and accurately present the gathered data for informed task prioritization.

Only through collection of a comprehensive set of network facts, including tests from every port using multiple protocols, can organizations discover and ultimately remediate leaks.

Inbound Leak Discovery

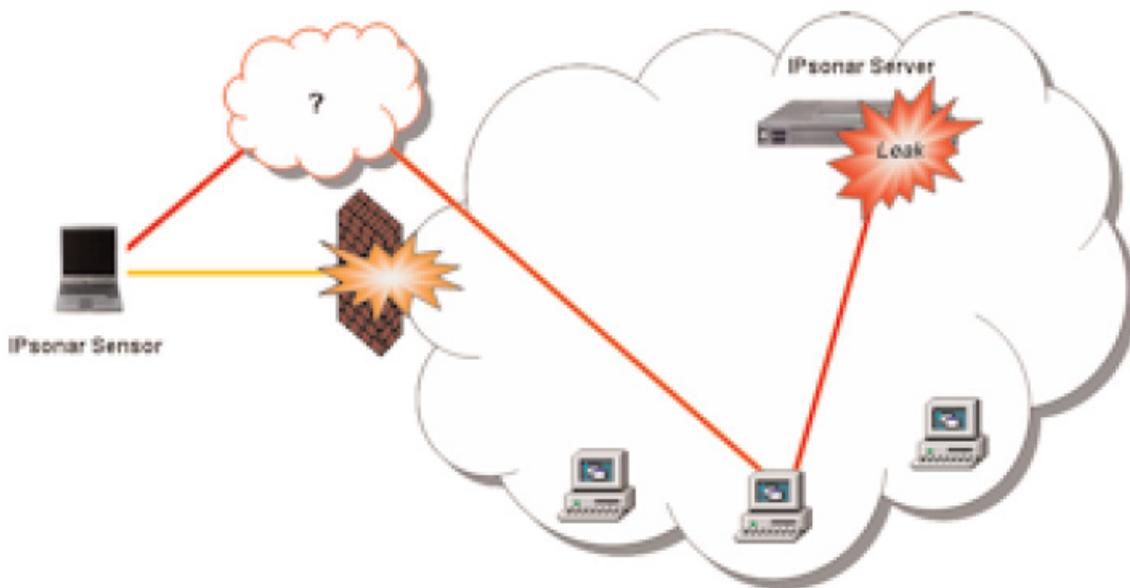


Figure 1: Discovering inbound leaks – A request sent from outside the network bypasses a firewall and is received, indicating an insecure port of entry – and the possibility that an unauthorized source can view the network from outside.

Outbound Leak Discovery

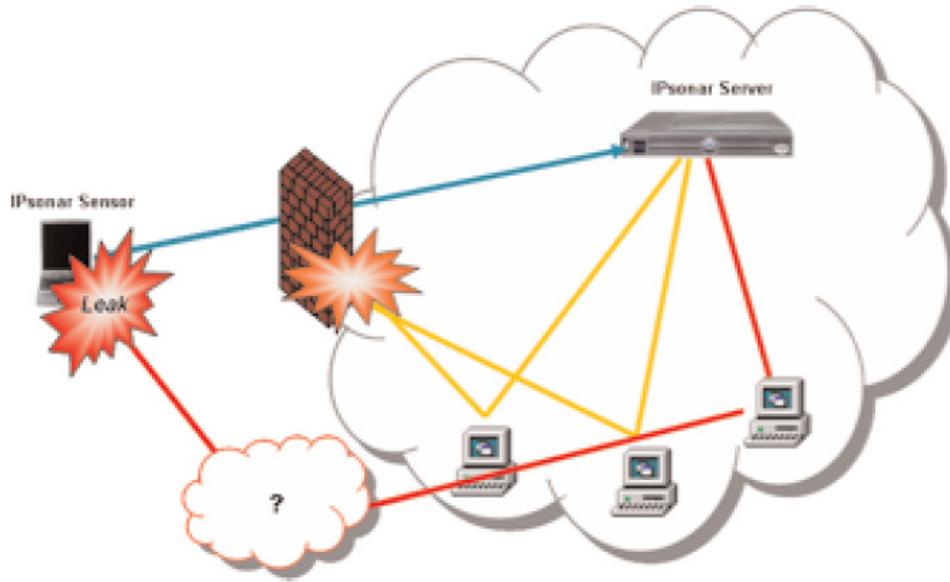


Figure 2: Discovering outbound leaks – Spoofed IP requests pass through nodes on the network. Leaking packets not blocked by the firewall are collected and reported.

Prioritized Patch Management and Data Leak Prevention with Leak Discovery

The results of inbound leak discovery enable organizations to see not only where leaks are present, but also which devices and hosts require the most critical response to control viruses that attack a particular port, such as the Slammer worm. Many organizations are reluctant to begin rolling out untested patches quickly in response to a spreading virus (often with good cause), since an untested patch could cause malfunctions in the network and may cause more harm than good. While a newly released patch is being tested, organizations can use Network Assurance scan results to rapidly identify which devices are running the port(s) vulnerable to the attack and which of those devices has leaks to and from the Internet.

Having identified the most vulnerable devices, organizations can take action to filter inbound traffic on the ports targeted by the virus. Once all of the leaking devices that use the targeted ports are protected, the organization can address the remainder of the non-leaking devices that run the targeted port. Prioritizing efforts and closing the most dangerous holes first ensures that critical devices that have a documented need to run particular services are running them safely. The

solution also provides a well-defined list of devices that will need patching once the new releases have been sufficiently tested.

Patching and access restriction, driven by leak detection, also have an additional advantage as they repair or restrict potentially significant avenues for data leakage. If sensitive data resides on a system that leaks to the Internet, for example, it may be exploited and used in unauthorized data transmission or access (a “data leak”). This results in significant risk exposure for the organization – many of which have unfortunately been exploited with the exposures documented in the media. While data leaks can happen in a number of ways, proactive network leak prevention protects a potentially high exposure channel in the defense of sensitive data.

Powerful Compliance Advantages of a Leak Discovery Solution

The structured detection, documentation, and remediation of leaking devices provides auditors with demonstrable justification that controls are in place to protect an enterprise or agency against any reasonable threats to network integrity and operations.

Leak discovery provides powerful capabilities for any public, private or government entity that must develop a compliant security policy for business risk assessment and incident response. Detecting and documenting leaks enables security organizations to get advanced level insight into vulnerabilities in the network and their potential impact. Organizations should pay special attention to inbound leaks discovered and take great care in describing the measures and technologies that are being taken to protect these leaking devices against unauthorized logical access.

Leaks in the Broader Context of Network Assurance

Discovering leaks is a vital element of a proactive Network Assurance program. The ever-changing nature of today's networks creates constant changes in connectivity, infrastructure, and defense configuration. Security organizations can no longer afford to relegate vulnerability assessments to times of convenience – leak discovery must become a continuous process.

As important as leak discovery is, handling leaks is just part of the bigger picture. To create a fully secure network, organizations must have a full range of network facts so they can:

- Make a baseline of network information
- Understand the true network perimeter
- Discover all unknown connectivity and unknown assets
- Shore up all leaks
- Handle rouge wireless devices

Network Assurance addresses dangerous gaps produced in risk management processes as rapid network change causes organizational security policy and network defense configuration to become misaligned. With the implementation of a comprehensive Network Assurance program, organizations can quantify risk from a network perspective, based on this comprehensive set of network facts. Using Network Assurance, companies can also prioritize remediation efforts based on a complete view of connectivity and risk. This is a critical requirement for validating policies and controls put into place to stop and prevent leaks, among other issues.

Continual internal and external forces keep network vulnerability in a never-ending state of flux. Although security products and policy management guidelines continue to evolve, there is ultimately no single proven solution to prevent all network leaks. However, consistent testing and assessment help bring companies as close as possible to achieving comprehensive security.

