OFIR ARKIN

demystifying passive network discovery and monitoring systems

Ofir Arkin is the CTO and co-founder of Insightix (http://www.insightix.com), conducts research in the information security field, and has published research papers, advisories, and articles in the fields of information warfare, VoIP security, and network discovery & management. He is a member of the honeynet project (http://www.honeynet.org) and is a director and chairs the "security research" committee at VoIPSA (http://www.voipsa.org).

- *ofir@sys-security.com*

THE QUESTIONS OF WHAT AND WHO is on the enterprise network and what is being done on and over the network has captured the attention of many researchers interested in finding appropriate network discovery technology. Such technology would not only allow these questions to be answered accurately, completely, and in a granular fashion but would allow this information to be maintained in real time.

In the past several months a number of commercial companies have hyped a new technological solution for network discovery: passive network discovery.

This article sheds light on the weaknesses of passive network discovery and monitoring systems. While acknowledging the advantages of this technology, the article explains its shortcomings, weakness by weaknesses, and demonstrates why it is unable to deliver complete, accurate, and granular network discovery and monitoring.

Passive Network Discovery

Passive network discovery and monitoring is a technology that processes captured packets from a monitored network in order to gather information about the network, its active elements, and their properties. It is usually installed at a network chokepoint. The roots of passive network discovery and monitoring technology go back to the mid-1990s, where references regarding use of the technology can be found [1].

The kind of information collected through passive network discovery and monitoring might include the following:

- Active network elements and their properties (e.g.., underlying operating system)
- Active network services and their versions
- The distances between active network elements and the monitoring point on the network
- Active client-based software and their versions
- Network utilization information
- Vulnerabilities found for network elements residing on the monitored network

Such information can be used for the following purposes:

- Building the layer 3–based topology of a monitored network

- Auditing
- Providing network utilization information
- Performing network forensics
- Performing vulnerability discovery
- Enhancing the operation of other security and/or network management systems by providing context regarding the network they operate in (information about the network, the active elements found on the network, and their properties)

Strengths

Passive network discovery and monitoring systems have important advantages related to their mode of operation.

Real-time operation: The operation (i.e., processing received network traffic and providing relevant information) is performed in real-time.

Zero performance impact: A passive network discovery and monitoring system has zero impact on the performance of the monitored network [2]. This is because the monitored network's traffic is copied and fed into the system, the operation of which involves no active querying. This all means that passive monitoring poses no risk to the stability of a monitored network and can theoretically be installed on any network.

Data processing: Passive network discovery and monitoring systems have the ability to gather information from all TCP/IP layers of network traffic processed.

Detection of active network elements and their properties: A passive network discovery and monitoring system is able to detect network elements along with some of their properties, by observing network activity related to the network element, provided that it is receiving and responding to network traffic. This means a passive system can:

- Detect active network elements that transmit and/or receive data over the monitored network
- Detect network elements as they become active and transmit and/or receive data over the monitored network

The ability to detect active network elements based on their network activity allows passive network discovery and monitoring systems to:

- Detect network elements that have low uptime
- Detect network elements that may transmit and/or receive data only for short time periods
- Detect which network elements on the monitored network are operational and serving requests coming from network elements on other networks
- Detect active network services running on non-default ports
- Detect active client-based network software operating on network elements on the monitored network

Detection of elements behind network obstacles: A passive system can detect active network elements that operate behind network obstacles and send and/or receive network traffic over the monitored network. A network obstacle is a network element that connects multiple networking elements to a network while filtering traffic from that network to these network elements (which are logically hidden behind it). Network obstacles include a network firewall, a NAT device, and a load balancer.

Granular network utilization information: A passive solution can provide information regarding the network utilization of its monitored network link. Unlike active monitoring solutions, which only provide basic network utilization information regarding the amount of traffic observed over a certain amount of time through SNMP [3], a passive network discovery and monitoring system supplies network utilization information by observing actual network traffic. A passive system has the ability to supply more granular and detailed network utilization information (i.e., per network element, per service, etc.) than active solutions.

Network utilization abnormality detection: The ability to provide statistical information regarding network utilization information, per network element, per network service, and the ability to gather information from all TCP/IP layers, enables a passive solution to build usage profiles for any element using the network and for any service used over the monitored network. These usage profiles can later be used to detect network-related abnormalities.

Detection of NAT-enabled devices: A passive system might be able to discover network address translation (NAT)–enabled devices that operate on the monitored network and to guess the number of network devices they might hide behind them [4].

Weaknesses

Although associated with important advantages, passive network discovery and monitoring systems have a number of critical weaknesses that affect their discovery and monitoring capabilities.

What you see is only what you get: By definition, a passive system will analyze and draw conclusions about

a monitored network, its elements, and their properties from network traffic observed at a monitoring location on the network. Consequently, a passive solution cannot draw conclusions about an element and/or its properties if the related network traffic does not go through the monitoring point. Moreover, information that needs to be collected by a passive system might never be gathered, if there is no network activity to disclose the information. A passive solution cannot detect idle elements, services, and applications.

The discovery performed by a passive system will be partial and incomplete, since it is unable, technologically, to detect all network assets and their respective properties. Finally, A passive system is blind when it comes to encrypted network traffic.

No control over the pace of discovery: A passive system has no control over the type of information that passes through its monitoring point and its initiation. Statistically, certain packets might not pass through the monitoring point for extended periods of time.

Limited IP address space coverage: Lacking control over the type of information that passes through its monitoring point, a passive network discovery system can generically cover only a limited IP address space.

Not everything can be passively determined: In some cases, information cannot be discovered by using passive network discovery. Passive vulnerability discovery is a good example: not all vulnerabilities can be determined passively, e.g., the vulnerabilities abused by the Code Red worm [5], the Blaster worm [6], and the Sasser worm [7].

Incomplete and partial network topology: A passive network discovery and monitoring system gathers network topology information based on the distances discovered between network elements and the monitoring point on the network, by relying on the time-to-live field value in the IP header of observed network traffic. The time-to-live field value is decremented from its default value by each routing-enabled device that processes the IP header of the packet on its way from the sender to its destination. Some passive network discovery and monitoring systems first determine the underlying operating system of a certain network element before relying on the time-to-live field value found with network traffic initiated by this network element.

The network topology information provided by a passive system relates only to layer 3–based information, i.e., routing-based information. A passive network discovery system cannot detect the physical network topology of a network it is monitoring, for several key reasons:

- It cannot detect the network switches that operate on the network. Usually a network switch will not generate network traffic other than the spanning tree protocol, sent only to its adjunct switches.
- A passive system cannot query switches for their CAM tables, detecting which network element (or elements) are connected to which switch port.

Additionally, a passive system would supply an incomplete and inaccurate network topology map, because:

- It cannot uncover routing that does not pass through its monitoring point.
- It cannot detect other routers operating on the monitored network.
- It is unable to uncover all of the network assets operating on the monitored network.

Deployment location and the number of sensors needed: The deployment location of a passive solution determines the data quality of the network traffic it receives. Network traffic data quality is relevant to the information collection process and is maximized when the deployment location is as close as possible to the access layer (i.e., between layer 2 and layer 3). A passive system loses some of its information collection abilities when it does not observe layer 2–based traffic of its monitored network elements.

A number of passive systems must be deployed in an enterprise implementation in order to have complete coverage, with the highest quality data collection, of the enterprise networks.

Network utilization–related issue: Although it is able to receive network traffic from multiple monitoring points passively, a passive system is unable to supply per-link utilization information. Furthermore, a passive system cannot uncover communications between network elements found on the same switch on the monitoring network.

Limited service monitoring: A passive network discovery system cannot monitor service condition state transitions or uncover idle services. For example, a network service might shut down soon after serving network traffic observed by a passive system, which will remain in the dark regarding this operational state transition.

Lesser-Known and More Important Weaknesses

Some weaknesses have not had widespread publicity. Here are details about some of them, showing why they are so very important.

Cannot resist decoy and deception: Although a passive system might have some conflict resolution policies, it might be possible, although dependent on a number of parameters, to trick the system into drawing wrong conclusions about the network, its elements, and their properties, by poisoning the observed network traffic.

A passive network discovery and monitoring system's conflict resolution policies might not be effective if the monitoring location does not allow the system to receive layer 2–based traffic from the monitored network.

Influencing the accuracy of a passive network discovery and monitoring system might influence other systems, such as network intrusion detection systems (NIDS) or network intrusion prevention systems (NIPS), that rely on the data collected by the passive network discovery and monitoring system as their input.

Example 1: Changing Location Information

Discovery relies on the time-to-live field value in the IP header of observed network traffic. It is possible to trick a passive network discovery and monitoring system, under several conditions, to conclude that a certain network element is located closer to or further away from a monitoring location simply by changing the default time-to-live field value in the IP header. For example, a Microsoft Windows 2000–based networking element has the default time-to-live field value set to 128. By changing the default value to the value of 126, a passive system would identify the operating system underlying the network element as Windows, and then trust the time-to-live field value information contained within the IP header of examined packets of this network element, placing it two hops further away from the monitoring point.

Example 2: Influencing Network Traffic Utilization Information

A network element can influence network traffic utilization information by injecting bogus traffic into the network and through the monitoring location. There are many different factors that prevent a passive network discovery and monitoring system from resisting these and other more and less sophisticated types of network traffic poisoning. Among them is the inability of passive systems to validate collected information.

Denial of service & remote code execution: The need of passive systems to decode received packets passively leaves them vulnerable to DoS and remote-execution attacks, of which there have been numerous examples [8].

Conclusion

This article has examined the strengths and weaknesses of passive network discovery and monitoring technology. It has demonstrated that despite the technology's advantages, it cannot, under any circumstances, perform complete, accurate, and granular network discovery and monitoring due to limitations that directly relate to the passive nature of the technology.

REFERENCES

[1] Vern Paxson, "Automated Packet Trace Analysis of TCP Implementations," 1997.

[2] Note that it is important not to overload a network device's backplane, in case port mirroring is being used. If the network device's backplane is overloaded, the network monitored will suffer performance degradation. Another side effect would be the network device's inability to send all of the network traffic which passes through the device and needs to be monitored to the network discovery and monitoring system.

[3] For more information on active network monitoring tools, see The Multi Router Traffic Grapher (MRTG) at http://people.ee.ethz.ch/~oetiker/webtools/mrtg/.

[4] Steven M. Bellovin, "A Technique for Counting NATed Hosts," http://www.cs.columbia.edu/~smb/papers/fnat.pdf.

[5] Microsoft Security Bulletin MS01-44, Cumulative Patch for IIS, August 15, 2001, http://www.microsoft.com/technet/security/Bulletin/MS01-044.mspx.

[6] Microsoft Security Bulletin MS03-39, Buffer Overrun in RPCSS Service Could Allow Code Execution (824146), September 10, 2003, http://www.microsoft.com/technet/security/Bulletin/MS03-039.mspx.

[7] Microsoft Security Bulletin MS04-011, Security Update for Microsoft Windows (835732), April 13, 2004, http://www.microsoft.com/technet/security/Bulletin/MS04-011.mspx.

[8] For examples of DoS attacks, see "Unknown Vulnerability in the Gnutella Dissector in Ethereal 0.10.6 through 0.10.8 Allows Remote Attackers to Cause a Denial of Service (Application Crash)," CAN-2005-0009, http://cve.mitre.org/cgi-bin/cvename.cgi?name=CAN-2005-0009; Marcin Zgorecki, "Snort TCP/IP Options Bug Lets Remote Users Deny Service," post to Snort-devel mailing list, October 2004. For an example of a remote code execution, see "Buffer Overflow in the X11 Dissector in Ethereal 0.8.10 through 0.10.8 Allows Remote Attackers to Execute Arbitrary Code via a Crafted Packet," CAN-2005-0084, http://cve.mitre.org/cgi-bin/cvename.cgi?name=CAN-2005-0084.