

Network Response and Remediation

In the last few articles of my column, I've alluded to various forms of remediation that can be employed by security practitioners in accordance with organizational policies. I've received a number of requests to explore this in more detail with a dedicated piece using example case studies. Therefore, per those requests, this article will use hypothetical case studies that I hope will increase awareness around remediation and illustrate some of the benefits it brings.

I'm defining network response and remediation generally as responding to a security event by leveraging automated tools with or without human intervention at a network layer. I've purposely titled this article "Network Response and Remediation" as opposed to "Network Access Control or Network Admission Control (NAC)"; and I've opted to not focus on the value propositions and shortcomings between pre- or post-network access solutions. This was done in order to keep the concepts at a higher, more vendor-agnostic level and to avoid simply focusing on ensuring clients meet policies before getting network access. I believe that response and remediation is a larger subject in which NAC can play a role.

Background

Automated tools for remediation have been around for many years. My first experience was in the late 1990s when I was part of a security team deploying intrusion detection systems (IDS) in conjunction with network-based firewalls. The idea was that when the IDS detected malicious activity — such as an attack coming from an exter-

nal IP address — that the offending IP would be blocked at all firewalls, protecting not just the target IP, but all the organization's assets from that attacker. This sounded really great at the time; what a time saver and security would undoubtedly be enhanced.

Then reality rolled in and the idealistic notion of intelligent network IDS devices working diligently modifying firewall rules on-the-fly to protect the network from malicious attackers was quickly overshadowed by a flood of helpdesk tickets from legitimate users complaining that they couldn't access the organization's resources from the Internet. The VPN clients used by employees were being interpreted by the IDS as malformed packets (an attack), and traffic from those employees was subsequently blocked; they couldn't even browse the company's public web site. We had effectively turned our own security safeguards into a weapon that was being used against us. It didn't take long for the plug to be pulled on the remediation project and for the security team to return to more manual techniques.

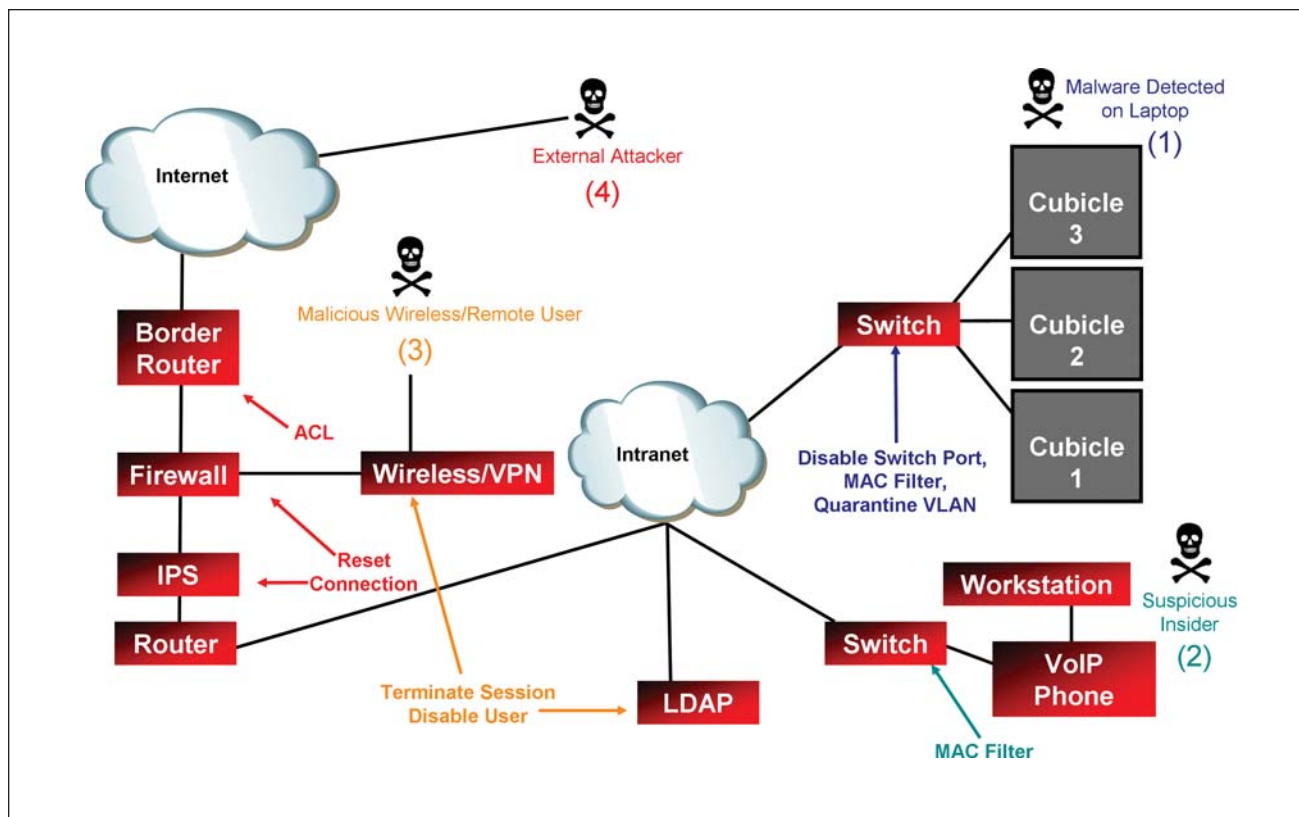
As time has progressed, organizations have become increasingly aware of the need to not only protect sensitive data, but to be able to do so rapidly. While discovering an incident post-mortem is certainly helpful, stopping an incident is almost always the gold standard. But with organizational security focus not purely on protecting the inside from the outside, but rather protecting every side from every other side, it is easy to get overwhelmed. Organizations today are well aware of insider threats, identity thieves, cyber

criminals, espionage and the like. But with so many angles to watch, new vulnerabilities, new exploits and the flood of increasingly broader daily responsibilities (e.g., compliance, risk, privacy, governance, incident prevention, detection and management), most security teams are doing everything they can just to keep their heads above water — let alone being able to actually stop an attack in near-real-time. Fortunately, times have changed, and while technology alone is never the sole solution, in this space, it has matured to the point where rapid remediation isn't only interesting from a theoretical perspective, but offers real, practical and usable capabilities.

Some remediation solutions are vendor specific, only allowing remediation within their own product family. Since most organizations have heterogeneous environments, this is rarely a desirable solution.

Other remediation solutions work in a synergistic manner with broader security management solutions. Enterprise Security Management (ESM) solutions, for example, provide incident detection capabilities, but only offer limited features in the form of comprehensive network response. However, ESM is complementary to remediation solutions. In fact, many remediation solutions depend on the capabilities of an ESM to make a decision or prompt a security analyst to make a decision as to whether or not remediation is necessary. The remediation solution simply becomes an extension of the ESM.

History has shown us that reactive, manual responses can be more complex,



costly, slow, and ineffective. Additionally, following a response, the essential task of documenting what was done is sometimes overlooked. This makes post-mortem reviews, change rollbacks, and addressing future events more difficult to manage. Most organizations also require change management procedures to be followed – even for an incident that requires rapid response. If this step is also manual, valuable time will be lost between the detection and remediation of an incident.

In the next section, four case studies will be explored to help illustrate some practical uses for network response and

remediation. The cases will reference the diagram in Figure 1.

Case Studies

(1) In the first case, malware is detected on a user's laptop. The user isn't doing anything intentionally malicious, but nonetheless, their laptop is carrying malware that is capable of infecting other assets within the organization's network. This malware could be detected in a number of ways:

- The organization may scan every system that attaches to the network to determine if it meets minimum requirements such as anti-virus, oper-

ating system and application patches, as well as personal firewalls and other relevant parameters.

- After being connected to the network for a few minutes, the malware may have attempted to propagate, thus generating alerts.

Regardless of the how the malware was detected, a log collection and analysis solution such as an ESM would process the information and, based on pre-determined policies, instruct the remediation tools to respond.

The actual response may be automatic – with or without human intervention.

It may require a security analyst to approve the change, or most likely, it would require an approval process that could be a series of automated notifications and approvals, such as e-mails and responses from individuals in the escalation chain. These responses could then trigger the actual remediation, while keeping in accordance with change management procedures and still allowing an efficient response.

From the perspective of unwitting users, they simply notice that their laptop isn't working in cube one, so they try cube two, three and so on. Some response solutions only focus on disabling a switch port; if that is the response to this action, it would have to happen every time the user moves to a different cube. Alternatively, a more elegant solution based on this scenario would be a MAC address filter, blocking the laptop at layer-2 regardless of where it plugs in. Also, the logical network

connection could be automatically moved, positioning the laptop on a quarantined VLAN and redirecting that user/laptop to a scanning and cleaning solution before it can be moved back onto the desired network.

(2) In the second case, we see another example of the value of MAC filtering. Here we have a device/user that appears to be generating suspicious events. It isn't something the AV scanners or policy checkers detected, but the traffic appears to represent anomalies that are consistent with malicious activity based on the determination of the event collection and analysis solution. Maybe it's an attack, a zero-day exploit, a worm, or perhaps it's simply a false positive.

MAC filters can be extremely beneficial when you have multiple endpoints connected to the same switch port. Consider a lab with several servers connected to a hub and up-linked to a

single physical port on a switch. Disabling the port on the switch disables all the servers connected to the hub from communicating over that network connection. VoIP networks are another example of this type of connectivity. Some VoIP networks work by having a user's workstation connected directly to the IP phone, which in turn is the actual device connected to the switch as illustrated in Figure 1.

During a suspicious incident, and based on organizational policies, the last thing that you may want to do is disable the switch port, thus not just disabling their computer's network connectivity, but their phone as well. Instead, you may want to configure your remediation solution to implement a MAC filter at that switch port, thus disabling the connectivity from that end workstation, but still allowing the user to communicate via phone to investigate why access has been disabled.



As time has progressed, organizations have become increasingly aware of the need to not only protect sensitive data, but to be able to do so rapidly. While discovering an incident postmortem is certainly helpful, stopping an incident is almost always the gold standard.

(3) In the third case, a mobile user enters the mix. Perhaps the reason is because he is a contractor. This contractor seems to be trying to access sensitive systems and conducting large file downloads as detected by various operating system and application logs.

For mobile users, remediation tools can terminate the session at the access point, whether it's a wireless controller or a VPN for example. Additionally, the remediation tools can communicate directly with an LDAP system, such as active directory, and disable that user on the directory server or on the remote connectivity point itself. This is an extremely useful feature, since a user could always try to reinitialize a terminated session and regain access to the network if his accounts are still enabled. Thus terminating the session at the access point and disabling the user on the directory server are both necessary to stop the live session and future connection attempts.

While it isn't illustrated in Figure 1, should a common authentication mechanism such as active directory be used for authenticating both physical and logical access, the user's physical access to the organization's facilities could also be revoked, increasing efficiencies in de-provisioning.

(4) In the fourth and final case, an external attacker is attempting to gain access to the organization's network. The possible solutions are blocking and resetting connections. The remediation solution can communicate natively with the preventative controls such as network devices, firewalls and IPS devices to initiate connection resets on those systems or make rule/ACL changes. Since the attacker is coming from the Internet, the best solution is to reset his active session and block the attack at the outermost ingress, which in this case would be ACL changes on the border router.

Summary

The following is a short summary of some of the features that can be found in some of today's effective network response and remediation solutions.

- Integrate with real-time, comprehensive threat detection solutions that analyze data across the existing detection and prevention infrastructure as well as hosts, applications, databases, physical security solutions, and the like.
- Minimal infrastructure changes/upgrades, minimal deployment and maintenance overhead, no inline devices and no client-side software.
- Work with a breadth and depth of

vendor agnostic devices for response such as routers, switches (layers 2 and 3), IPS, VPNs, wireless access points, firewalls, LDAP, and physical access control systems.

- Automatically discover network devices and topology, and during a response, use that information to intelligently respond to attacks with optimal mitigation strategies within the constraints of organizational policies.
- Simulate the remediation impacts to render potentially unknown issues, and thus allowing an analyst to see any problems that may arise from remediation efforts before any changes are made.
- Easily reverse any changes made, if necessary.
- Automatically and consistently document all changes and create audit trails and reports.
- Integrate with the organization's approval-based response policies, for example, following change management, approval and escalation procedures. ■■

Brian T. Contos, CISSP, is the Chief Security Officer of ArcSight

