Rapid Enterprise Triaging (RETRI)

How to run a compromised network and keep your data safe

By
Michael Murphy and Aaron LeMasters

June 2009

Imagine this scenario

Routine log analysis uncovers suspicious activity dating back several months, and active beaconing reveals a backdoor channel in an exploited piece of production software on your network. Anti-Virus did not catch it - updated IDS signatures reveal dozens of compromised machines, all buried beneath a hierarchy of domain controllers and NAT'ed subnets across different autonomous organizations throughout a globally distributed network. What do you do without the necessary infrastructure and tools to respond?

Once the initial feelings of panic subside, you realize that you have countless questions and decisions which must be made in a matter of moments. What are you going to tell the boss? How are you going to "fix" this issue? What competitive advantage was lost due to data which was stolen? What will your network users do while you scope the problem and begin the long and agonizing process of rebuilding the network? How do you save your job…

Network administrators need to come to grips with reality. If you operate any large scale network which is home to valuable data you will be broken into at some point. Security professionals agree, it is only a matter of time and preparing for the inevitable will ensure that when disaster strikes you are ready to react and limit the impact to your organization. Cyber COOP has become a vital part of any mature IT disaster recovery plan.

Large-scale response efforts are expensive, inefficient and dangerous

What is worse, most companies are not typically resourced or prepared to respond. As a result, they are forced to either purchase a service contract or acquire the tools to respond internally. Neither response is a complete solution nor makes the network any less prone to infection. More importantly, these approaches ignore the persistent threat at hand, and proprietary company data and resources remain exposed throughout the effort. To add insult to injury, you just spent your entire IT security budget for a one-time fix.

Traditional Incident Response

The current industry best practice incident response process includes four major phases:

1. Discovery
2. Containment
3. Eradication
4. Recovery

While there are many variations of this theme, including iterative components, analysis phases, and "always-on" phases such as network monitoring, these are the four parts most companies care about.

In our experience, implementations of this incident response model are linear, resulting in lingering infections; or worse, parallelized but not synchronized, causing chaos, wasting money and producing conflicting results.

In either case, it is fruitless to model your response process around the perceived threat. In recent years, common malware has increased dramatically in sophistication. Your infected network may contain deeply entrenched malware

with multiple layers of infection and advanced evasion and persistence capabilities. You must assume the worst.

Thus, we submit that the most critical phase of any incident response process is the **containment** phase. Prior to this paper, there were two schools of thought for containment of enterprise-wide infections.

Traditional Option One: Offline Rebuild

Option one involved the network disconnecting from the internet and doing a complete rebuild from trusted media. Once the rebuild is completed (several weeks later) the network is reconnected to the internet. This is optimal if a high confidence of security is required in the cleanup efforts. Users are left with no access to email and web services. Not many business models operate effectively under this scenario. Option one is extremely costly due to extreme end user productivity loss.

Traditional Option Two: Online Rebuild

Option two is the most commonly used method today. Network and Host level scans are employed to identify, locate and remediate infected host. This method's major benefit is the nearly non-existent impact on end users. They continue to work on the compromised network while the IT Security teams work in the background. This method suffers from some major issues. How do you scan for what you don't know? If the malware is polymorphic both in how it communicates and installs on the host, it will be nearly impossible to reliably locate all instances of the malware. There are numerous documented cases of this methodology SPECTACULARLY failing. In one such case, a year later it was discovered that the attackers had changed backdoors and moved to recently "cleaned" sections of the network and stayed hidden for well over nine months. That is nine months of data being quietly exfiltrated from a "cleaned" network. In the end, they had to start all over at a cost of several millions of dollars. This is no longer the "cheap" way to respond to large scale network intrusions.

A New Method is Required: Enter the RETRI Methodology

Our methodology of rapid enterprise triaging (RETRI) embraces a macro-approach to the containment phase of the Incident Response process. Rather than focusing on individual network segments or hosts, our approach prioritizes broad network isolation to contain the threat and ensures core business functions remain operable. The result is less strain on your IT staff and nearly no downtime for your users. The immediate benefit is that the infection is *isolated* and *contained* from further spread or damage, allowing the appropriate teams to clean the network at a reasonable pace.

The cornerstone of this process is isolating the affected network in a quarantined network zone (Q-net) with no logical access to unaffected network segments. This Q-net provides isolation of the affected network and prevents any further intrusion, command and control, or exfiltration of data. This isolation further provides a secure environment for analysis and remediation.

The second fundamental element in this methodology is the delivery of basic office resources such as email, internet access, and document production over an Application Streaming Infrastructure (ASI) to the affected users. The ASI is locked down and its users are forced to

use strong two factor authentication to gain access to its resources. All services and access is delivered to the Q-Net through secure connections limited to the ports and protocols needed for the remote access. Compromise of these services on the backend provides no risk as command and control and exfiltration routes have been blocked.

The final element consists of the remaining components of the traditional response process and restoration of compromised machines to a known good state. It is the Incident Response Team's responsibility to discover what that attacker was doing on the network, what files were accessed and stolen, and finally, which hosts were compromised. This is done at whatever pace the Security / IT team desires, since the malware no longer poses any threat to the network and its valuable data.

Admittedly, the remainder of the traditional response process (discovery, analysis, eradication and recovery) is fraught with inefficiencies, including limited tools and training. This is why we have decided to release a free tool to assist forensic investigators during the analysis and recovery phases.

Codeword, a free tool to assist in the RETRI methodology

Codeword was designed in the spirit of the RETRI methodology: to create options for incident responders when selecting response tools and to help alleviate enormous costs of enterprise Incident Response.

Codeword is similar in functionality to commercial agent-based forensic tools, but focuses on the needs of management and analysts during rapid triaging. These needs typically include an estimate of the spread of infection; an ability to search hosts for malware with constantly evolving indicators; a thorough evidence-collecting capability; and advanced, seamless reporting. These capabilities are implemented in a server-agent framework.

Codeword has two components:

1. **Administrator Console:** allows administrators to generate custom agent scanners packaged in an easy-to-use MSI installer and to connect/interact with the agent on the host. The administrator can issue a variety of commands to the agent to investigate the host, including extracting binaries from memory, enumerating devices installed on the system, and debugging a process.
2. **Host Agent:** contained in the deployed MSI installer, the Agent is composed of a user-mode executable and a kernel-mode driver, both of which can optionally install as a persistent service. For rapid analysis, the agent can also execute a single time, remove its footprints, and automatically report results over FIPS 140 compliant encrypted channels.

Codeword also boasts advanced rootkit detection and kernel integrity checking capabilities. It is also capable of stealth installation and self-protection.

Conclusion

Large enterprises now have a third option for dealing with massive compromises. RETRI address many of the limitations which are presented by the previous schools of thought, additionally it does so at a greatly reduced cost. Not only is the financial burden reduced, but so is the stress level of the IT staff doing the cleanup. IT staff can work at their normal

pace, greatly reducing stress and fatigue which lead to mistakes.

Lastly, the non-IT staff is happy because they are only slightly impacted as the network is being cleaned and rebuilt around them.

Contact Us

Michael Murphy –
mike.a.murphy@gmail.com

Aaron Lemasters –
aaronlemsaters@yahoo.com