

“Aurora” Response Recommendations

February 17th, 2010
Prepared by Alex Stamos, Partner
Version 1.0

Background

On January 12th, 2010 Google publicly revealed that they were the victim of a “highly sophisticated and targeted attack on our corporate infrastructure originating from China that resulted in the theft of intellectual property from Google.”¹ Google was not the only company affected by this attack; at the time Google notified over 30 other companies of infection by this malware. In the time since then, further investigations have uncovered that over one hundred companies may have been targeted, although it’s difficult to ascertain how closely related these attackers are to Google’s assailants.

iSEC Partners has been investigating this attack with several victims, and has found a number of common oversights and vulnerabilities that enabled these attackers to be successful. There has been a great deal of discussion around the Aurora² malware suite due to the large amount of information released by the anti-virus vendors³. While finding and investigating an infection by Aurora is an important component of responding to this incident, we believe it is important to take into account the overall actions of the **people** behind these attacks when considering how to respond.

Despite the diversity of victims in these attacks, we have seen a common pattern in the attacks, which generally proceed like this:

1. The attacker socially engineers a victim, often in an overseas office, to visit a malicious website.
2. This website uses a browser vulnerability to load custom malware on the initial victim’s machine.
3. The malware calls out to a control server, likely identified by a dynamic DNS address.
4. The attacker escalates his privilege on the corporate Windows network, using cached or local administrator credentials.
5. The attacker attempts to access an Active Directory server to obtain the password database, which can be cracked onsite or offsite.
6. The attacker uses cracked credentials to obtain VPN access, or creates a fake user in the VPN access server.
7. At this point, the attack varies based upon the victim. The attacker may steal administrator credentials to access production systems, obtain source code from a source repository, access data hosted at the victim, or explore Intranet sites for valuable intellectual property.

In this document we outline our recommendations for organizations that have **not** been contacted or found evidence of an Aurora infection. Known affected organizations can contact us for help putting together a more aggressive incident response plan.

¹ <http://googleblog.blogspot.com/2010/01/new-approach-to-china.html>

² AKA Trojan.Hydraq

³ <http://www.symantec.com/connect/blogs/trojanhydraq-incident>

Tactical Recommendations

These are recommended steps IT and security teams can take to detect the type of malicious activity exemplified by the Aurora incident.

1. **Log and inspect DNS traffic.** The most effective surveillance mechanism for tracking the automated and manual attacks has been monitoring of internal DNS queries. DNS is a common control channel for malware, and these attackers have shown an affinity for using dynamic DNS providers to provide control information in DNS responses allowing the attackers to quickly move command and control servers ahead of investigators. In response to this incident we are recommending that potential victims:
 - a. Enable logging on their local DNS responders and centrally aggregate the request and response bodies
 - b. Monitor and audit all requests for dynamic DNS names
 - c. Investigate devices making suspicious or repeated queries to dynamic DNS names
2. **Establish internal network surveillance capability.** Without the ability to inspect and capture internal network traffic it is extremely difficult to trace this type of attack. The quickest solution that could provide this capability would be the deployment of software-based Intrusion Detection System (IDS) sensors inside the corporate firewalls, possibly using a free platform such as Snort. In the long run this capability could be replaced by a commercial product or managed service, although some of our clients have had difficulty effecting rule changes in their managed IDS sensors quick enough to be effective.
3. **Control inbound and outbound network traffic.** The initial route of infection for all of the known attacks has been through exploiting flaws in Internet Explorer or Adobe Acrobat using content hosted on external servers. Although these exploits were “0-days” and undetectable by ruleset based malware scanners, victims with extensive controls around Internet access by workstations, laptops and servers have been able to respond more quickly to these attacks, control outbound control channels and use log data to track down infected hosts. We strongly recommend that all companies preparing for this level of adversary deploy content-scanning web proxies with extensive, long-lived logging of requests. DNS resolution should happen on these proxies, and advanced clients should utilize SSL interception capability by deploying a corporate CA.
4. **Expand log aggregation.** The companies with the most effective response to these attacks have utilized their central log aggregation mechanisms to track the actions of these attackers. The storage of log data in multiple systems can seriously hamper investigation efforts and we are strongly recommending that clients create a central logging capability dedicated to security monitoring and incident response. We recommend that the security log aggregation system contain at least these logs:
 - a. Security and System event logs from Windows member servers.
 - b. Security, System, and Application event logs from Domain Controllers.
 - c. Login and SSH logs from UNIX servers.
 - d. DNS request and response records from all local resolvers.
 - e. Alerts from internal IDS sensors.
 - f. Request logs from web proxy devices.
 - g. Web access logs from Intranet sites and applications.

5. **Expand Windows endpoint control.** Networks with loosely managed Windows desktops and laptops are especially vulnerable to attacks against end users. We recommend that security sensitive organizations:
 - a. Disable the use of local administrator access for day-to-day activities. In most organizations local administrator access should be eliminated for almost all employees on their corporate-managed systems.
 - b. Utilize a patch solution that can update non-Microsoft products. Managing the patch levels of popular products such as Adobe Flash, Reader, Mozilla Firefox and Sun Java is key to preventing common exploits from working on your users.
 - c. Consider application whitelisting technologies.
6. **Audit VPN access and enrollment.** We have found evidence that these attackers prefer to establish remote access via “legitimate” VPN portals instead of always relying on the control channels for their custom malware. It is advisable to begin auditing your VPN access servers for unexpected users and to consider protections that prevent attackers from easily issuing themselves VPN access once they compromise a small number of user accounts.
7. **Test malware scanning against known rootkits.** Our investigations have shown that the attackers are willing to use variants of standard rootkits in their attacks alongside their custom code, such as TLD3. We recommend that organizations test the effectiveness of their malware control mechanisms by intentionally infecting a managed system with common rootkits and running a “full scan”.

Strategic Recommendations

These are longer term recommendations to help prepare for this type of advanced, targeted threat. This list is in no way exhaustive, but represents several issues that we commonly see in client networks.

1. **Build a security operations team.** Many of our suggestions will result in the creation of much larger datasets of use for quickly finding and responding to attackers inside of a company's perimeter. Unfortunately, this data is useless if there is no staff to monitor or respond to anomalies. A proper security operations team is related to but stands apart from security architecture and audit functions. It is important that security operations employees are granted access to production systems and logging facilities and are able to quickly retrieve data from monitoring systems.
2. **Secure your overseas offices.** Many of the companies involved have found that a large portion of their infected workstations originated in an overseas office. The preferred route of attack seems to be via social networks and chat channels, which are effective in socially engineering these employees into clicking on a malicious link or opening an infected file. It is important for central IT security teams to evaluate what IT functions are absolutely required for the functioning of their remote offices and to restrict internal network access to those resources. More restrictive companies should consider which social networks and chat mechanisms should be provided for use on corporate systems.
3. **Classify and catalog sensitive data.** Victims of these types of attacks who do not understand the disposition of their critical data and systems start at a disadvantage. Having this list in hand at the start of an incident response can guide the responders to the systems of greatest import to the victim and attacker.
4. **Secure your Active Directory network.** A theme of many of these intrusions has been a concerted effort by the attackers to utilize the corporate Active Directory forest to gain privilege across the network. The attackers are quite patient, and have utilized software that waits in the background for a use of a NTLM credential by a domain admin for one of several services on the machine (RDP, SMB, DCOM) that then captures the NTLM hash for cracking. Once they have this hash, the attacker is able to pull the entire Active Directory database for offline pre-computed dictionary cracking. There are several steps to securing your AD forest from these attackers:
 - a. **Make Domain Administrator account logins smartcard-only.** This configuration greatly reduces the risk of theft and misuse of administrator credentials. If backup logins with passwords are required, they should be distinct accounts limited to console login only.
 - b. **Do not use shared local accounts.** The use of local (non-Domain) accounts with a shared password is a common technique in corporate IT environments, especially when machines are imaged from a common corporate build. Unfortunately shared accounts can easily be cracked by offline computing resources or the NTLM hash used as a password equivalent.
 - c. **Beware of using Domain Admin accounts in automated processes.** Many IT departments use system management and security devices that regularly log into Windows clients and servers to check for patch levels, configuration compliance or to perform log aggregation. Whether these logins are performed by dedicated devices or software installed on a domain member server, it is often possible for an attacker to downgrade the authentication handshake to NTLM and perform a brute-force attack to retrieve the domain admin password. These processes should be configured to use non-Administrator users whenever possible. Consider upgrading such systems to Windows 7 or Windows Server 2008 R2 and using GPO to audit or restrict NTLM traffic to remote servers.

- d. **Set GPO to reduce authentication attack danger.** There are several Active Directory group policy settings that can increase the difficulty of attacking local and domain credentials. Although the details of these settings go beyond the scope of this document, there are several excellent public resources for Windows hardening⁴. Important settings to harden include the **NoLMHash**⁵, **LDAP Signing**, **DCOM Packet Security** and **LAN Manager Authentication Level**⁶ policies.
- e. **Audit Domain users for unusual permissions.** Attackers who gain escalated privileges on a domain often create user accounts and join them to unusual groups, such as Backup Admins, or directly grant themselves privileges unnecessary for normal use. There are several commercial tools for auditing an Active Directory for permission issues, which can be useful in finding outlier accounts.
- f. **Deploy read-only domain controllers in overseas offices.** Active Directory forests with the Windows 2008 functional level have the option of deploying read-only Active Directory servers. This is a critical security enhancement to Windows networks that is appropriate for use in semi-trusted physical environments. It should be noted, however, that password database cracking techniques are also available against read-only domain controllers.
- g. **Consider an external Forest with a two-way trust for remote offices.** Administrative security boundaries in AD are enforced at the Forest, not the Domain level. Creating distinct AD Forests with two-way trusts and SID Filter Quarantine⁷ enabled (the default) allows seamless Kerberos authentication between all systems while preventing elevation of privilege attacks that are possible between systems in the same Forest. Alternatively, Active Directory Federation Services can be used to enable user-transparent, cross-Forest resource access with even more limited interconnectivity and trust.

⁴ http://www.nsa.gov/ia/guidance/security_configuration_guides/current_guides.shtml

⁵ <http://support.microsoft.com/kb/299656>

⁶ [http://technet.microsoft.com/en-us/library/dd560653\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd560653(WS.10).aspx)

⁷ [http://technet.microsoft.com/en-us/library/cc772633\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc772633(WS.10).aspx)

Lessons Learned

Although none of the attacks or technique used in this series of attacks are particularly novel, the skill set, patience and tenacity of the attackers is much greater than most enterprises are equipped to deal with. Here are some general lessons that we can take from these attacks:

- **Attackers are ignoring the front door.** Despite the focus of the security industry and enterprise security teams on production networks and applications, attackers have learned that “back door” attacks against end users are much more effective at gaining access to major corporations. It is generally much harder to secure internal corporate than production networks, and internal and external security teams will need to refocus on designing trustworthy corporate IT environments.
- **Current Anti-Virus solutions are not working.** All of the victims we have worked with had already deployed enterprise-wide anti-virus solutions, none of which prevented the initial attacks or the escalation of privilege within the network. Anti-virus tests are mostly rule-based, and the majority of heuristic detection mechanisms can be easily bypassed if an attacker is customizing his malware for that product. Enterprise IT departments should not rely on the promises of their vendors or third parties, and should verify the effectiveness of their AV solution by testing it against known advanced malware and rootkits.
- **Patching sometimes is not enough.** The vulnerabilities exploited during this attack were 0-days, meaning no patch or mitigation directions were available to correct these flaws. Prompt and verifiable patching is still a key component of any information security plan, but IT groups should keep in mind that advanced attackers will often be able to find new flaws in complicated end-user products like web browsers, office suites and document readers.
- **You might be playing in the big leagues.** The most interesting aspect of this incident is that a number of small to medium sized companies now join the ranks of major defense contractors, utilities and major software vendors as potential victims of extremely advanced attackers. This is concerning for many reasons, not the least of which is that even most Fortune-500 companies will not be able to assemble security teams with the diversity of skills necessary to respond to this type of incident. It is extremely unlikely that SMBs will be able to properly prepare for these threats alone, opening the door for more effective product and managed service solutions than what is currently available.

Version History

0.9– Initial Draft (Feb 9, 2010)

0.91 – Added lessons learned (Feb 12, 2010)

0.92 – Added additional Win hardening tips (Feb 15, 2010)

1.0 – Small fixes, first public release (Feb 17,2010)