Providing threat intelligence to those in the Cloud

Google Cloud is pleased to publish this "Snapshot" of its most recent *Threat Horizons* report. The information provided is based on threat intelligence observations from the Threat Analysis Group, Google Cloud Threat Intelligence for Chronicle, Trust and Safety, and other internal teams. It summarizes actionable intelligence that enables organizations to protect against ever-evolving threats. For the complete *Threat Horizon* report, please visit gcat.google.com.

Snapshot Summary

Adversaries and researchers alike are continuing to scour the web looking for vulnerable instances of Log4j. As a result, service providers have been and continue to work with their cloud customers to ensure the infrastructure is secure as well as check the status of customer-installed tools and third-party dependencies in their environments to see if they are affected. While adversaries continue to knock on this door, observations have shown that they are opting to use known open-source tools, native Cloud services, and previously established domains for persistence in their attacks.

Current observations include:

- **Vulnerable instances of Apache Log4j still sought by attackers**
  Google Cloud is seeing ongoing scans for vulnerable Log4j instances by attackers and security researchers and we recommend continued vigilance to ensure available patches and mitigations are applied and sustained. During the month following the vulnerability's disclosure there was extensive scanning across the Internet. Google Cloud and other providers had a unique vantage point over this and used this to good effect to help customers identify vulnerabilities as well as watch for the evolution of attempted exploitation to rapidly assure mitigations were effective for cloud infrastructure and customers. Google Cloud is continuing to see scanning (400K times a day) and believe similar, if not more scanning levels against all providers, and so we recommend continued vigilance in ensuring patching is effective.

- **Adversaries using an open-source platform to maintain network persistence**
  Google Cloud Threat Intelligence has observed from across the Internet that Sliver is being used by adversaries post initial compromise in attempts to ensure they maintain access to networks. Recent observations show attackers using Sliver, an open-source, cross-platform adversary emulation framework, as the second stage malware to compromise organizations, moving the possibility of attackers using Sliver for malicious actions. Network detection is possible whenever an implant uses the HTTP transport to observe the initialization session, such as with the "out-of-the-box" implementation of Sliver; however, when customized, it will be more difficult to detect.

- **Cloud Shell used for obscuring identity on the internet via reverse SSH tunneling** (revised on 3/15)
  Already compromised or fraudulent credentials used Cloud Shell to establish reverse SSH tunneling for obscuring identity on the internet or for general computing. Additionally, Google Trust and Safety has observed activity pertaining to SSH tunnels available for rent and trade, which supports the argument that this abuse is being monetized. Google Cloud allows users to create a free sandboxed environment of the Linux system ("Cloud Shell") outside customer projects or networks, with the ability to execute commands using a provided shell. While Cloud Shell inbound connections are restricted, the outbound connections are not. The availability of outbound connections to conduct a reverse SSH tunneling from Cloud Shell to any public endpoint can serve as a means for threat actors to distribute malicious campaigns or perform harmful activity. Customers who may be concerned about this potential obfuscation/abuse tactic may enable a VPC service perimeter to limit service access to services and users inside their perimeter..

- **Domain previously identified by TAG used in ongoing attacks against researchers**
  A domain previously identified by Google's Threat Analysis Group (TAG) was associated with intrusions launched by a North Korean government-backed entity against security and vulnerability researchers. These intrusions originated from a trojanized version of IDA Pro as discovered by ESET Research. The actor presumably shared the modified installer for IDA Pro, a disassembler popular among security researchers, in order to replace a specific DLL within the installer package. When executed, it would attempt to download and install a backdoor from a North Korean controlled domain. Over the past 12+ months, the actor has launched multiple campaigns against the security and vulnerability research community.

- **Lessons Learned**
  When securing an environment and ensuring exfiltration has not occurred after an incident, customers should review logs associated with file sharing and domain-wide Takeout in addition to looking for implanted malware. Threat actors have been known to use tools native to the Cloud environment rather than downloading custom malware or scripts to avoid detection. This "living off the land" technique has been used extensively in on-premise compromises and is being adopted in Cloud environments.

  In the situations where it is necessary to respond to a security incident, where there was an exploitation of third-party software in the Cloud instance, customers should review logs to determine the point of initial compromise as well as determining if additional unauthorized software was installed after the initial compromise or if unauthorized configurations were made. Adversaries have been known to close doors behind themselves, i.e., patch the vulnerability which grants initial access, and then pivot to new malware, or reverse shells to maintain presence.

Recommendations

As Google Cloud works with its customers in a "shared fate" partnership, valuable trends and lessons-learned emerge from other incidents that Google helped address:

- Google Cloud customers have Implemented Cloud Armor and Cloud IDS to mitigate threats related to Log4j.
- The Java scanning feature of Google Cloud On-Demand Scanning can be used to help identify Linux-based container images that use an impacted version of Log4j. Identifying vulnerable images can help prevent them from being deployed in production. This functionality can be used either on a locally stored container image or one stored in Artifact Registry or Google Container Registry.
- Chronicle is Google's threat hunting tool that provides extended event collection across Google Cloud and in on-premise environments. Those using Chronicle for log ingestion/SIEM and have historical event data can detect and respond to Apache Log4j.
- Google Cloud customers used Security Command Center (SCC) Premium and Cloud Logging to scan DNS calls to known malicious sites associated with Log4j vulnerability abuse and query logs.
- The Google Open Source Security Team partners with the security company Code Intelligence to provide continuous fuzzing for Log4j as part of OSS-Fuzz and Google Cloud customers leveraged resources from the Open Source Security Foundation to secure their environments.

Based on these observations there are a number of mitigating measures to counter these threats.

Risk	Countermeasures
Instance vulnerabilities	Modifying firewall rules to specifically address SSH connections. Follow password best practices and best practices for configuring Cloud environments. Update third-party software prior to a Cloud instance being exposed to the web. Avoid publishing credentials in GitHub projects. Use Container Analysis to perform vulnerability scanning and metadata storage. Leverage Web Security Scanner in the Security Command Center to identify security vulnerabilities in App Engine, Google Kubernetes Engine, and Compute Engine. Use service accounts with Compute Engine to authenticate apps instead of user credentials. Implement Policy Intelligence tools to help understand and manage policies. Use predefined configurations through Assured Workloads to reduce misconfigurations. Set up conditional alerts in the Cloud Console to send alerts upon high resource consumption. Enforce and monitor password requirements for users through the Google Admin console.
Downloading software updates	Establish a strong chain of custody by hashing and verifying downloads.

Unauthorized exfiltration	Review logs associated with file sharing and domain-wide Takeout in addition to looking for implanted malware.