Security Assessment

Chronicle by Google Cloud Security

August 17, 2020

World Wide Technology

PRESENTED BY

Anthony Glackmeyer
Consulting Security Architect
World Wide Technology
706-442-4268
anthony.glackmeyer@wwt.com

www.wwt.com

Executive Summary

Chronicle is a capable and well thought out security and event monitoring tool. It is not a SIEM, in the traditional sense, but can be used to segment security events out of tools like Splunk and ArcSight. This type of segmentation can not only improve security posture but can also reduce storage and SIEM licensing cost significantly as security event data is by and far the largest amount of data stored in a SIEM.

Organizations are typically double-taxed on security events due to the licensing model of many SIEMs: paying by the amount of data ingested as well as paying for the underlying storage. Chronicle licenses its service by the number of employees in an organization. A license covering unlimited data ingest and a one-year retention.

With the traditional logging platform free of SIEM data, additional ops and application tracing data can be stored. Increasing user experience through more relevant application performance monitoring.

Chronicle does have a few barriers to adoption, one being the perception of its parent company Google. This assessment will step through this and other barriers with pro's and con's and possible risk mitigations for each.

For organizations already leveraging cloud platforms, Chronicle is a real alternative to traditional logging. Decreasing spend, increasing observability and threat relevance all without changing your risk posture. The enriched data from Chronicle is a catalyst for more meaningful use of an existing SOAR platform, increasing the confidence of executed actions and freeing up SOC staff to hunt real threats.

This assessment is meant to educate and shorten research time. Further architecture discussion is warranted when any change to security event monitoring is considered.

Table of Contents

Request Statement

Companies are looking to expand their threat detection and monitoring capabilities, especially around remote working and endpoint security. These solutions can generate a large number of events that should translate into relevant logs within a SIEM tool. This exponential increase in event generation will require additional compute in log forwarders and correlation engines along with additional storage for the resulting logs.

The prospect of additional spend is driving more customers to research alternatives to the traditional SIEM and log correlation/aggregation tools. Chronicle Security is one such tool.

Scope of Assessment

Classification of Data: data going into Chronicle is assumed to have an elevated classification due to the nature of the content and risk associated with its exposure. This will vary based on the customers regulatory and contractual compliance burden, as well as the use-cases targeting the solution. i.e. a Tiered approach based on classification.
Service(s)/Application(s): Chronicle Security platform, not specific to any one service or feature
Users: Admins, SOC Analysts, Threat Hunters, and Application Security Engineers
Assets: any system that can generate logs
Connectivity: Internet

Artifacts

The following artifacts were used and/or referenced in this assessment effort:

Artifact	Source	Date Received
Platform security discussion with Joseph Lee, Chronicle Engineer	Discussion	2020-08-05
Encryption at Rest in Google Cloud **URL**	Google Public Information	2020-08-05
Assessment questionnaire provided by WWT and completed by Chronicle team	Document	2020-08-07
Platform deep-dive with Joseph Lee, Chronicle Engineer	video conference, live demo	2020-08-12

Terms and Definitions

The following acronyms and terms are used in this document:

Term	Definition
ACL	Access Control List
API	Application Programming Interface – a set of communication methods between software. Typically, API is used when referring to web-services, but can be used between two pieces of software to enhance the capability of each.
BYOK	Bring Your Own Key – an industry term for a cloud customer generating an encryption key outside of the cloud service and importing it into the cloud service with the goal of making the customers data unreadable by the service provider.
CASB	Cloud Access Security Broker – a system that integrates with a customer's cloud services with to enforce the customers security policies. The level of enforcement is directly related to the compatibility between the CASB and the cloud service.
CD	Cardholder Data, also shortened to CHD, is the combination of a PAN and any PII associated with the account holder.
CSP	Cloud Service Provider
DEK	Data Encryption Key
DLP	Data Loss Prevention – the function or process that scans data for key words and phrases deemed sensitive to the business. Data loss prevention is largely intended to prevent exfiltration of sensitive data. The term DLP has come to represent the product that performs the scanning rather than the process.
DRM	Digital Rights Management – often used with music and movie files, DRM refers to the ability to restrict who does what with a file. With DRM, a file can be "marked" to only be read by a specific person or persons. If the file is forwarded to others, it is unreadable.
FedRAMP	Federal Risk and Authorization Management Program – an assessment that certifies a provider meets the minimum-security requirements to process data for a government entity. The security requirements are based directly on the moderate security controls defined in NIST 800-53 with specific clarifications and enhancements.
GDPR	General Data Protection Regulation – an EU regulation that governs the lifecycle of personal data of EU citizens. It went into effect May 25, 2018.
GRC	Governance Risk Management and Compliance

HSM	Hardware Security Module – a hardened physical security device designed to securely store and process encryption keys. Traditionally, HSM's are thought of as a server or add-on module. In today's cloud-centric world, the term HSM is often used to refer to a multi-tenant key repository that is backed by an actual hardware device.
IaaS	Infrastructure as a Service – in comparison to other cloud delivery models, IaaS requires the most management but give the customer the most customization and control. The CSP provides the foundation on which the customer can build any number of servers, networks, etc. The CSP is responsible for the data center, the foundation or hypervisor and the management interface used by the customer. The customer is responsible for everything else.
KEK	Key Encrypting Key – a key used to wrap another key, usually a data encryption key, in an encryption envelope. Typically, increases the security of the data encrypted by the DEK because the KEK is generated and managed by another system with a different key hierarchy. Accessing the data would require breaking the encryption chain on two systems.
MTTR	Mean Time To Respond (or Remediate) is the amount of time it takes an organization to act upon a threat that has been identified.
PAN	Primary Account Number – credit card number.
PII	Personally Identifiable Information – any information or combination of information that can be used to identify a person.
SaaS	Software as a Service – in comparison to other cloud delivery models, SaaS requires the least amount of management and skill. The CSP provides everything including the data center, servers, network, application, patching, and security controls. The customer is responsible for configuring their preferences within the application and nothing more.
SDK	Software Development Kit – a set of instructions and tools to guide developers in creating applications designed to interface with a specific software package.

Elements of Circumstance

Who

Google's parent company, Alphabet, created a cyber security company named Chronicle. Chronicle was acquired by Google and the company is now part of Google Cloud Security.

What

Formerly called Backstory, the platform is now referred to as Chronicle Security or Google Chronicle. Chronicle Security is a SaaS platform that ingests unstructured security event data and structures it. Indexing and organizing the data in a way cybersecurity researchers and remediation experts can find and respond to threats.

Where

Google core (private) datacenters located in US, EU, and AP. Google's IaaS platform is available in more locations than Google core. GCP customers may be able to run cloud workloads in a geographic area where Google Core, and thus Chronicle, is not physically available.

When

Alphabet started Chronicle in 2018 with the first Chronicle security service (Backstory) released spring of 2019. Soon after, Google acquired Chronicle as part of their Cloud Security practice and rebranded Backstory as Chronicle in 2020.

With

The customer needs security tools with log generation capabilities, a docker host, cloud security tools with API's, internet access, SAML 2.0 IDP.

Why

Security tools and event sources like AD, Tanium, proxy, firewalls, Carbon Black, Crowdstrike, etc. are chatty, producing a large number of events. Most organizations cannot keep up with the amount of data ingest required to detect all threats, requiring decisions to filter on what is deemed critical at the time of setup. Chronicle allows ingestion of all events over long periods of time to allow sec ops team to understand what is normal vs abnormal.

How

Chronicle ingests security telemetry data from different solutions and uses data modeling to automatically stitch related information together at high-speed to enrich incident investigation. Reducing time to detect and respond/remediate.

Security Analysis

The Service

Google Chronicle is classified as a SaaS offering hosted on Google's core infrastructure. The same infrastructure as Google Search, YouTube, G-Suite, etc. This infrastructure has a different risk posture and utilizes a different security architecture than GCP. Google core geo-locations differ from GCP available regions. Chronicle is a multi-tenant solution and customers can specify geo-based tenancy to fit their compliance needs within Google Core locations in North America, Europe and Asia.

Encryption

Every customer receives their own instance within the platform. All data uploaded into and processed within Chronicle tenants/instances is broken up into chunks (or shards), indexed, encrypted with their own key, then distributed and replicated across Google's core infrastructure.

Google uses FIPS 140-2 encryption modules for encrypting data.

Google's key hierarchy is optimized for security and speed. A DEK is generated for each chunk of data by a key vault in close physical proximity to where the chunk will be stored, to optimize cryptographic performance. The chunk is encrypted with the DEK and stored while the DEK is encrypted by a KEK specific to Chronicle and stored with/near the chunk. The KEK used to encrypt the DEK may be part of a hierarchy, or chain, depending on geographic or density distribution of the service.

DEKs and KEKs are generated and managed by different systems; DEKs by the storage systems encryption engine and KEKs by Google's KMS. DEKs are only used once, if a chunk is decrypted and re-encrypted a new DEK is used. KEKs generated by KMS, on the other hand, can be reused (but not exported) and each encrypt/decrypt action is logged for audit purposes. KMS keys are automatically rotated at regular intervals and are encrypted by Root KMS keys. The Root keys are encrypted with a Root KMS master key stored in memory on a global P2P network. These keys are backed by physical safes.

KMS manages access and usage of KEKs through per-key ACLs. ACLs programmatically define what services and users can access the keys for encrypt/decrypt operations. Services are granted access on-demand, per action. Users are granted access through programmatic, multi-step process. Customers are not required to be involved in this process, it is internal, heavily logged, audited and logged. The length of access is determined by the request, there is a max and a timeout. Since the access requests are programmatic, they cannot be extended.

Customers do not have access to keys nor can customer keys be introduced into the key hierarchy.

Logging

All service activities within the customers control plane, including key management operations and key usage activities, are logged but are not directly available for consumption by the customer.

Authentication logs would be available from the IDP when SSO is configured and can be consumed by Chronicle.

Ingesting events from outside sources, including IDP's, can occur in three ways: a Docker appliance forwarder, API, out-of-band pull. The Docker appliance can be hosted on Windows or Linux and will send data to the customers Chronicle instance. Chronicle can use API's from 3rd party cloud security solutions to pull data using API keys and passwords. Out-of-band involves retrieving flat file data from cloud object stores, like S3 buckets, as well as SFTP sources. This utilizes a Google Service outside of Chronicle for retrieval. That service is not assessed here.

For base "coverage", Chronicle needs seven core sets/sources of logs to begin stitching together a complete picture: Web proxy, DNS, DHCP, AD, EDR, Threat Intelligence and Email. The more sources, the better the picture and the more layers of the threat landscape can be modeled.

Once the raw logs from these disparate sources are ingested, they are parsed into a Unified Data Model. Parsers are built for specific security solutions to ensure data is properly understood. Parsers for most security solutions are in-place. Additional parsers can be built to ingest data from newly discovered or acquired sources that may not be already defined.

Once parsed and stitched together, data is kept and continuously used for one year. Customers can pay for keeping data longer than a year

Chronicle is meant to be the final resting place for security event data. Unique circumstances may call for data to be in multiple places, for these use-cases, APIs are available to pull metadata out of Chronicle and the Chronicle forwarder can also send data to multiple places. Big Data querying capabilities is in development, should be available in the first half of 2021. Use-cases requiring Chronicle analysis to trigger an action or response, can rely on Chronicle's integrations with SOAR platforms.

Authentication

SSO is achieved through AD or SAML integration through IDPs like Okta, Ping, or Azure AD. Chronicle does not have any built-in login restriction capabilities, this can be achieved through the customers IDP configuration. Administrative RBAC is not currently available in Chronicle.

Chronicle supports API key and password for https authentication/connectivity to data sources.

Chronicle generates a unique API key for each customer to use with SAML 2.0 authentication for pulling metadata out of the platform via the Chronicle API.

Service Level Agreements

Chronicle has SLA's in place that govern the customer data stored in the service. Google does not sell, mine, extract to google-analytics or use customer data for sales or marketing.
As of writing this assessment, there was no SLA information available regarding uptime or access guarantees to modeled and stitched data.

Data Governance

As data flows into Chronicle it passes through a number of logical tiers before being encrypted and stored. These tiers perform different actions upon the data, including the stitching and enrichment. When the event data is stored, a complete and in-time representation of all supporting data and metadata are stored, including IP address and DNS lookups.

The final logical step before storage involves what Google calls the "parser." As explained in the Logging section, Parsers take raw event data, enrich and store the resulting logs in a Unified Data Model; however, Parsers can be used to manipulate data in other ways as it flows into Chronicle. With respect to data governance, the parser can be configured to identify restricted or classified data, as determined by the customers corporate security policy. Logic can be added to the parser to identify any pattern of data in any field, such as PAN, SSN, DOB, and IP addresses, and perform an action. Actions can be anything from sending an alert to obfuscating or redacting the data in some programmatic way. Since the parser is run before storage and executed in-memory, data at-rest should remain compliant with governance policies.

The parser can be used to perform custom actions based on the event source, such as performing hashes, lookups and field creation to enrich modeled data and provide SOAR tools additional actionable criteria.

Parser actions are powerful and if configured incorrectly or with malicious intent, data could be lost. It is for this reason Google does not give customers access to the parsers. Changes must be requested before Google engineers will work with customers on creating custom parser actions. Before the actions "go-live," an internal review is performed to ensure data is not adversely modified to the point it impacts data modeling.

Attestations of Compliance

Chronicle has achieved ISO-27001, SOC-2 and SOC-3.
Chronicle uses Google Cloud's storage encryption platform, which has achieved FIPS 140-2 certification.
FedRAMP, HITRUST and PCI-DSS certifications are in-progress for Chronicle.

Conclusion

Visibility and observability are key to any successful security strategy. Before any action can be taken, before any investigation can result in an action and before any data can be flagged for investigation, event logs must be collected and identified as relevant. Associating relevance with an event is what Chronicle does.

Many companies just point everything towards their SIEM tool, with the understanding their SOC team has created correlation rules to create relevance within all of this data. For the most mature security departments with dedicated SOC engineers and threat hunters, ensuring everyone understands the commands and rule formats necessary to properly use the SIEM is an ongoing challenge.. Even the most seasoned SOC team member with complete mastery of all the tools is always playing catch-up to the continuous addition of tools, constant configuration changes that introduce new scenarios and the sheer volume of data all of this generates. Not including cloud.

There are alternatives to keeping up and leveling the playing field. Offloading SOC operations is one, licensing product specific plugins within the SIEM tool is another. Leveraging a tools like Chronicle is somewhere in-between. Security departments that need to keep their threat data close to the vest may not be able to outsource their SOC. On the other hand, waiting for and purchasing plugins for the SIEM keeps the data in-house but doesn't really free up the SOC to hunt threats. Using a tool like Chronicle leverages Google's scale and analytics power to assign relevance to your security event data, keeping the data close to the vest while freeing up the SOC to focus on the actual threats instead of the noise.

The first hurdle to leveraging Chronicle is who owns it, Google. The first reaction by any security professional will be a knee-jerk, “No.” Just because it’s Google. After further investigation and data gathering, that will subside (reluctantly), as it did for this author. Google has SLA’s in place to ensure the data is not used for any purpose other than threat modeling for the data owner. There are also ways to ensure events ingested by Chronicle are sanitized before and after they the data center: SIEM forwarder logic and Chronicle Parsers.

Leveraging Google core also adds the capability of unlimited storage of event and telemetry data for one year. The average time to detect a threat with in the enterprise is 191 days in 2020. When stored, each event is indexed and, during its 365 day lifespan, continuously used in data modeling. This increases the likelihood that an APT or a dormant zero-day will be detected.

The second hurdle for an organization will likely be the SaaS nature of Chronicle. As outlined in the assessment, Chronicle is multi-tenant with each customer getting their own tenant and API keys. As a SaaS service; however, the customer does not have control of the encryption keys. Google has a well-documented and very mature key management process to ensure customer data remains secured, but in the end, Google handles the key material.

For organizations already leveraging cloud, especially SaaS services, storing telemetry data from these services in Chronicle will not increase the security risk, as the data already resides in a cloud service. For organizations not in the cloud or with more stringent data governance requirements, sanitizing event data will be key – as mentioned before. An API gateway or log forwarder can be used to redact data before being sent to Chronicle. This should only require enough storage to process event data and compensate for an outage that may interfere with events being sent to Chronicle. The cost of an additional forwarder or gateway will be offset by the storage and SIEM licensing savings by moving security event data to Chronicle.

The third, and possibly the toughest, hurdle to overcome when adopting Chronicle is SLAs. This assessment has covered the SLAs that govern the use of customer data ingested by Chronicle, what is unclear from research for this document are the SLAs that govern uptime and access to the data. Google’s core infrastructure has SLAs covering data resiliency, but as is typical for SaaS services, the uptime of the system that provides access to that data is not necessarily guaranteed. Such an SLA would be required for customers to ensure they can meet internal or customer-bound SLAs for threat response and remediation.

As previously stated, customer leveraging cloud services like Office 365, already accept a certain level of risk when it comes to platform uptime. The risk is often weighed against the benefits of leveraging the platform, usually in the form of cost. Chronicle is no different.

Chronicle is a self-managed security event modeling system leveraging a platform that provides unlimited scale-out without customer tuning or management overhead. Accessing the enriched threat data is achieved through any browser without the need for plugins or specialized SIEM command knowledge. When data is ingested or a query is run, Chronicles worldwide infrastructure can process petabytes of data in seconds. This type of big data analytics and threat analysis can only be achieved at a scale unattainable by even the biggest enterprises. These benefits can be realized by most organizations through the use of Chronicle, even those with a high compliance burden.

Appendix A: Example Architecture

Below is an example architecture that logically depicts the different paths event data can take when being ingested by Chronicle. Security tools, including firewalls, in your data centers can send data directly to a Chronicle forwarder. These tools can also send data to a SIEM tool, such as Splunk, that can then forward security event data to the Chronicle forwarder after sanitization on-premises (if required).

Chronicle natively ingests security event data directly from cloud platforms such as GitHub, AWS, Oracle cloud and Office 365. Organizations can configure Chronicle to leverage the same SAML 2.0 IDP as these cloud platforms.

Once events are parsed and stitched, logs can be used to trigger events in a SOAR platform, which can take action upon these events.