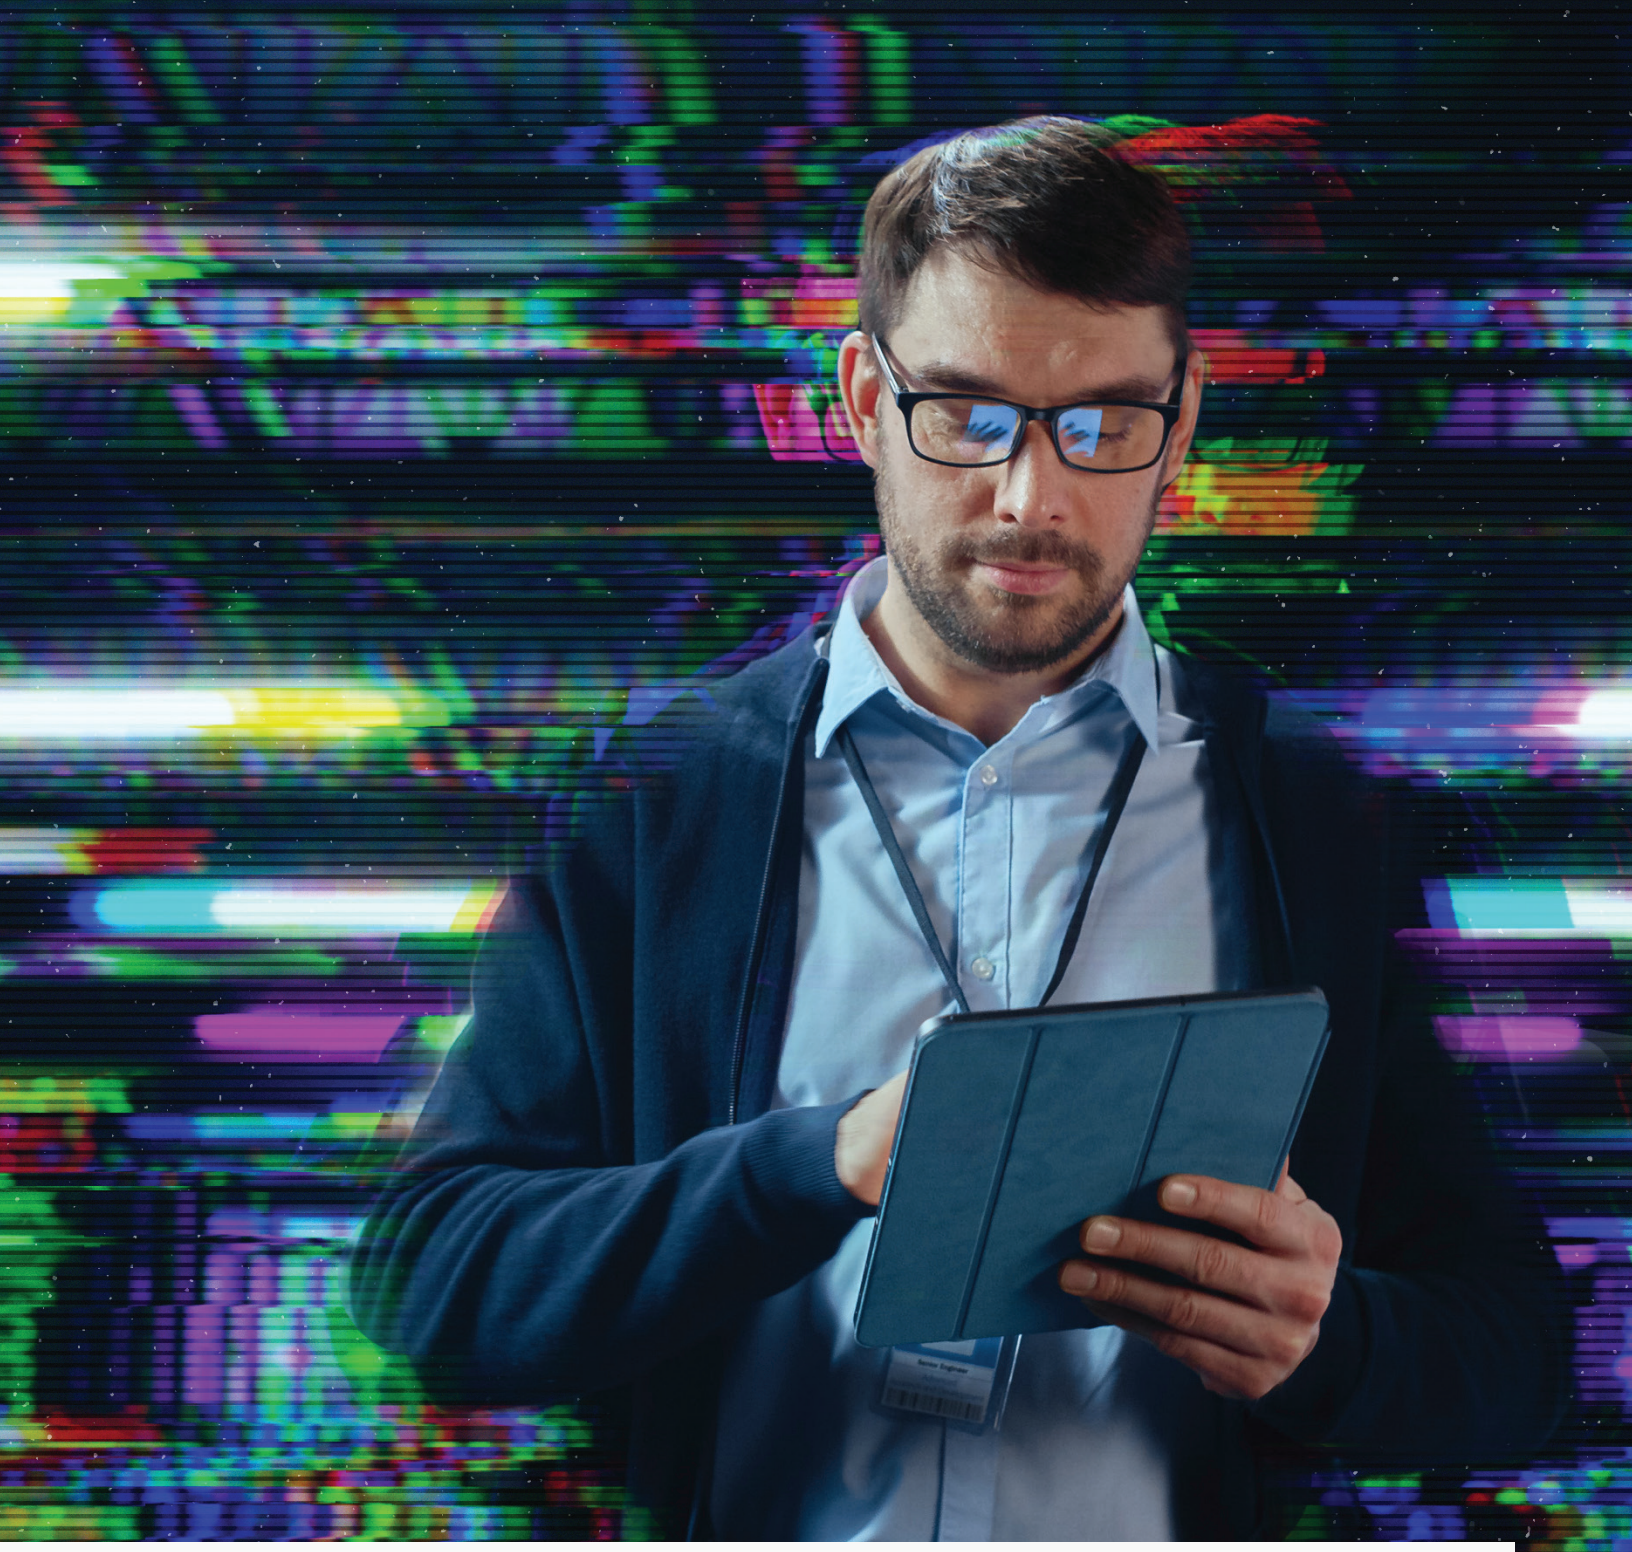

The 2023 Buyer's Guide to Next-Gen SIEM

EBOOK

DEVO

Table of Contents

Introduction

SIEMs have been deployed in security operations centers (SOC) for 15 years. The technology was created to take in data and events from security sources, usually the perimeter of the network, and bubble up critical events that required action. But times—and security technologies—have changed, and the demands placed on SIEMs have changed as well. The perimeter has disappeared as services and infrastructure have moved to hybrid cloud and multi-cloud environments, and users have moved to mobile devices and work-from-home scenarios.

Organizations of all types and sizes need to protect more attack surfaces than ever before, in a more connected world, with more data being generated than at any time in history. And the stakes have never been higher. The spoils for attackers have increased dramatically, leading to an exponential increase in the number and sophistication of adversaries. For these reasons, in the last few years a new type of SIEM has emerged: the next-generation SIEM.

THIS GUIDE IS MEANT TO EDUCATE THE READER ON THREE KEY CONCEPTS:

1. **How to distinguish** a next-gen SIEM from its older, less sophisticated predecessors
2. **How to recognize** the signs that it's time to move toward a next-gen SIEM
3. **How to compare and evaluate** next-gen SIEM solutions to choose the correct one for your needs. As examples, we'll compare four market leaders in the SIEM space: Splunk, Microsoft Azure Sentinel, Google Chronicle and Devo

WHAT IS A SIEM?

According to Gartner, security information and event management (SIEM) technology is used for threat detection, investigation, compliance and security incident management by collecting and analyzing (both near-real time and historical) security events, along with many other event and contextual data sources.

NEXT-GEN SIEM VS. XDR?

SIEM and XDR (extended detection and response) provide value in two different but potentially complementary ways. SIEM had its genesis in compliance and has evolved to serve as a broader threat and operational risk platform. XDR evolved with a specific focus on endpoint threats and provides a platform for deep threat detection and response.

NEXT-GEN SIEM VS. SOAR?

While a SIEM will ingest various log and event data from on-prem and cloud data sources, a SOAR (security orchestration, automation and response) automates response path workflows to reduce the time required to handle alerts and investigations. The SIEM is the brain, the SOAR is the muscle.

LEGACY SIEM IS DEAD. LONG LIVE NEXT-GEN SIEM

The question on the table is, what exactly is a next-gen SIEM? Many vendors, including legacy SIEM providers, lay claim to the "next-gen" label. How can you tell the difference between a legacy SIEM and a true next-gen SIEM? And what criteria should you use to evaluate vendors?

Here are six core criteria that distinguish legacy and next-gen SIEM vendors so you can accurately evaluate which solution is best for your organization.

Deployment Models

Legacy SIEMs that run on premises just don't scale well because of the on-prem hardware limitations of compute, memory and storage.

ON PREMISES VS. SAAS

The easiest way to identify a legacy SIEM is that it only runs on premises. This forces users to make several compromises.

The biggest compromise is scalability. Legacy SIEMs that run on prem just don't scale well because of the on-prem hardware limitations of compute, memory and storage. This makes it difficult for a legacy SIEM to grow with your business. Legacy SIEM users cannot collect as much data as they want because these solutions lack the compute power to search or the storage capacity to retain the data. Inevitably, legacy SIEMs force organizations to make tough decisions about what data is "must have" vs "nice to have," or invest significant time, money and effort to manage a complex SIEM infrastructure.

These organizations also must make tough decisions on what data they store and how long they retain it. This usually results in organizations only being able to search 90 days (or less) of hot data. This is problematic because to be effective, many threat investigations need to determine the "Day-1 event"of when an attack first appeared in the environment. Going back in time more than 90 days to find that Day-1 event is a slow and tedious process when the data is in cold storage. Headline-making attacks, such as Sunburst and others, clearly show the amount of time an attacker can lay dormant in an environment. To put it simply: historical perspective matters.

Finally, legacy SIEMs are constrained by the static resources of on-prem compute and memory. Since hardware is typically refreshed every 3 to 5 years, this greatly limits the resources accessible to the on-prem, legacy SIEM. This resource limitation often slows down search performance and dashboard rendering, forcing security teams to limit the number of searches and users due to performance concerns. Also, on- prem SIEMs have significant administrative overhead that adds significant additional costs measured in both dollars and time.

Conversely, the next-gen SIEM is delivered through SaaS. It can take full advantage of the elasticity of the cloud to deliver compute, memory and storage resources on demand. This liberates owners from many of the compromises of a legacy SIEM. Owners of a SaaS-delivered next-gen SIEM can collect all the data they deem necessary, store it longer, and search it more often by more users. These advantages mean greater visibility into more data sources by more people, which will result in a greater security posture. And all this value is gained with less administrative overhead.

Don't make the mistake of thinking that because a SIEM is offered in the cloud, it's SaaS. Running an on-prem version of SIEM in the cloud but managing it yourself is not the same as SaaS. Self-managing a solution in the cloud is simply trading the architecting and administering of on-prem infrastructure for the architecting and administering of cloud infrastructure, which can be just as problematic if not done correctly. A true next-gen SIEM is completely managed by the vendor so the user can focus on securing their business—not keep your SIEM infrastructure up and running.

Running an on-prem version of SIEM in the cloud but managing it yourself is not the same as SaaS. Self-managing a solution in the cloud is simply trading the architecting and administering of on-prem infrastructure for the architecting and administering of cloud infrastructure, which can be just as problematic if not done correctly.

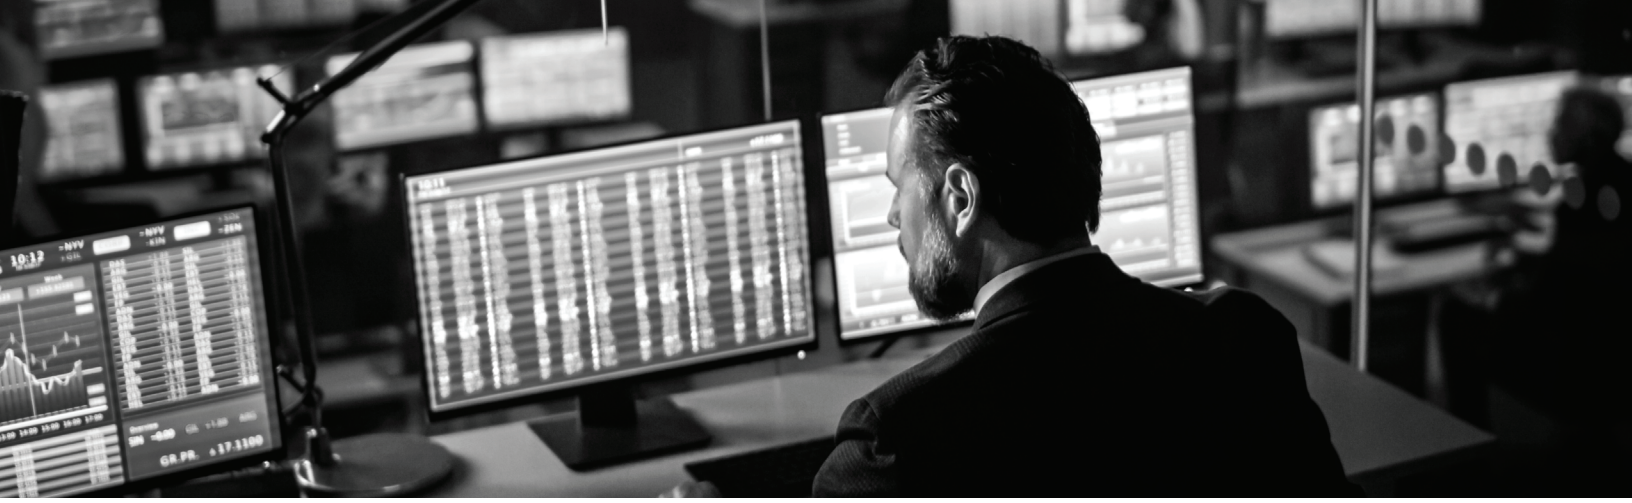

EVALUATION CRITERIA: ARE YOU TRULY A CLOUD-NATIVE, SAAS SOLUTION?

It's important to state that simply offering a SIEM as a SaaS solution does not make it a next-gen SIEM. As discussed above, many legacy SIEM vendors have just lifted and shifted their old, inefficient architectures to the cloud and are delivering it as a SaaS offering. Running a legacy SIEM in the cloud yields high cloud infrastructure costs that organizations must bear. Beyond cost, security teams forced to use a legacy SIEM as a SaaS solution also will have to deal with a dearth of hot, searchable data and slow search performance at scale—not exactly a recipe for success.

Usually any self-hosted solution, even in your own managed Amazon Web Services (AWS) or another cloud environment, is not a next-gen solution because it relies on you to make sure it scales out appropriately.

A good sniff test of whether a SIEM is truly cloud-native and SaaS is to see if it is offered as an on-prem, self-hosted solution. If it is offered for on-prem operation, it is almost certainly not a next-gen SIEM. Even if you can self-manage it in your private cloud, buyer beware. Usually any self-hosted solution, even in your own managed Amazon Web Services (AWS) or another cloud environment, is not a next-gen solution because it relies on you to make sure it scales out appropriately. Like most modern, cutting-edge technologies, a next-gen SIEM is a completely SaaS offering.

VENDOR COMPARISON: WHICH VENDORS DELIVER A CLOUD-NATIVE SAAS SOLUTION?

Splunk

Splunk is not a cloud-native SaaS solution. It was designed as an on-prem solution that the company later moved to the cloud. While its SaaS solution—Splunk Cloud—is growing, it's a lifted-and-shifted architecture.

Microsoft Sentinel

Sentinel is a cloud-native SaaS solution. Built on Azure, Sentinel was designed to live in the cloud and overcome many of the challenges of on-prem solutions. But Sentinel can only be deployed in Azure.

Google Chronicle

Chronicle is also a cloud-native SaaS solution, built on top of Google's public cloud. It's the least mature offering in this list, but does scale well. It can only be deployed in GCP.

Devo

Devo is a fully managed SaaS solution born in the cloud to handle the multi-terabyte needs of today's data age. It fully supports ingesting data from multi-cloud and hybrid cloud environments.

Next-Gen SIEM Evaluation Criteria	Splunk	MS Sentinel	Chronicle	Devo
Truly Cloud-Native SaaS	● ● ● ●	● ● ● ●	● ● ● ●	● ● ● ●

Integrated Capabilities

A next-gen SIEM correlates data regardless of data source (network, server, security, endpoint, application) and regardless of the domain (on prem, private cloud, public cloud).

MODULAR SIEM VS. A COMPLETE AND OPEN SIEM

Another hallmark of a legacy SIEM is its modular architecture. Common examples of add-on modules (and license costs) involve certain data sources, such as NetFlow, or add-on modules for functionality, such as machine learning (ML) or reports. The curse of the modular SIEM is usually a result of two factors: maximizing revenue and technology acquisition.

As their market share grew, on-prem SIEM vendors quickly realized they needed to sell more products to maintain revenue growth. Instead of developing new products, it was easier to create add-on functionality for their existing SIEM and charge extra for it. This gave their salesforce something new to sell—and add on to revenue from existing maintenance charges. Sometimes these add-on modules worked great, other times not so much.

Often, the legacy SIEM vendor didn't even create the add-on technology. Instead, they acquired another vendor and integrated its separate technology as a new chargeable "module." Technology acquired via the acquisition of niche vendors typically integrates much more poorly with the core product than if the legacy SIEM vendor had developed it.

A necessity of this modular approach is that each module can be used without the others since the vendor can't know which modules a customer will buy. Unfortunately, this results in a suite of modules with a disjointed workflow. Each module has its own screen or UI that doesn't work seamlessly with the others. Invariably this results in a frustrating, inefficient user experience where analysts need to have multiple windows open across multiple monitors and then manually correlate a lot of disparate data.

As attackers grew more subtle and sophisticated, the defenders on security teams realized they desperately needed to bring all data and functionality together in a seamless workflow with one UI. This is exactly what a next-gen SIEM does. It correlates data regardless of data source (network, server, security, endpoint, application) and regardless of the domain (on prem, private cloud, public cloud). And it also brings all the functionality, such as ML, data visualization, enrichment and analytics together in the same interface to expedite the analyst workflow.

EVALUATION CRITERIA: ARE ALL FEATURES INCLUDED AS PART OF THE SAAS PRICE?

Beware the curse of the modular SIEM and its many negative aspects. The first is a costly, complicated and unpredictable cost structure. Remember that you are going to live with your new SIEM for at least 2 or 3 years, so you don't want to go through a complex pricing and licensing process every time you renew. And you don't want to be surprised with costs you didn't anticipate for a function or feature you needed all along. Another negative effect of the curse is it causes swivel-chair analysis and disjointed workflows as analysts must bounce back and forth between modules and screens. This slows down your SOC analysts and increases the likelihood of failing to detect and stop a serious threat before it's too late.

Next-gen SIEMs avoid the modular curse by including everything you need as part of the SaaS offering. The cost structure is simple to understand and, more importantly, easy to predict year after year. Everything in a true next-gen SIEM is built into the core product, right at the fingertips of your SOC analysts.

When evaluating next-gen SIEMs, it's critically important to obtain an itemized cost breakdown of all included features and functions. If you're handed a list with lots of line items and a mix of different licensing costs, you're probably looking at a legacy SIEM in next-gen-SIEM clothing. When test-driving or going through a proof of concept of a prospective SIEM, pay close attention if you're required to log into a new interface. That could indicate you are entering another module that requires a separate license.

Next-gen SIEMs avoid the modular curse by including everything you need as part of the SaaS offering. The cost structure is simple to understand and, more importantly, easy to predict year after year.

VENDOR COMPARISON: ARE ALL FEATURES INCLUDED IN THE SAAS PRICE?

Splunk

Splunk's pricing model is very complex and does not include all features. The company charges extra for storage, extra for encrypting data at rest, and the SIEM itself is an additional cost on top of its core product (which you must buy). In fact, of the three vendors in this evaluation, Splunk is probably the worst when it comes to surprising customers with additional costs.

MS Sentinel

Sentinel comes with all features enabled. But Sentinel alone isn't all you need to purchase. First, you must pay for the data to be ingested into a Log Analytics workspace, which has its own pricing: **https://azure.microsoft.com/en-us/pricing/details/monitor**.

After you pay for data ingestion and storage of your Log Analytics workspace, you're not finished. You then must pay Sentinel's ingest pricing and storage costs. You'll find Sentinel costs (separate and additive to the Log Analytics costs) here: **https://azure.microsoft.com/en-us/pricing/details/azure-sentinel**.

While Sentinel's license includes all features, it does have some pricing pitfalls you need to consider. The biggest charge to watch for is the additional cost associated with exceeding your reserve pricing. Since Sentinel pricing is reserve-based, exceeding your reserve puts you into an "on-demand" pricing structure, which can quickly escalate if you significantly exceed your reserve. This model presents a challenge for customers with bursty data needs—either you over-provision for the majority of the time, or you pay occasional penalties for exceeding your reserve. It's challenging for Sentinel users to plan an annual budget for SOC costs.

Google Chronicle

Chronicle's pricing model is based either on the number of employees in the customer organization, or by amount of data ingested. Pricing includes a year's worth of hot data by default. However, since Chronicle has no native dashboard capability, you absolutely need additional GCP products like Looker and BigQuery to do custom dashboards and these are priced separately. Google also has a SOAR (formerly Siemplify) but this is also priced separately. Finally, other Google products such as VirusTotal and Mandiant professional services used in conjunction with Chronicle are also priced separately.

Devo

Devo's pricing model is refreshingly simple and includes all features. You pay only for data ingested, averaged over a 30-day period. The price includes support. All SIEM, SecOps, ITOps, machine learning, and automation functionality are included in the price as well. Devo also includes 400 days of hot storage in the default license cost—the most of any vendor.

> **Devo's pricing model is refreshingly simple and includes all features. You pay only for data ingested, averaged over a 30-day period.**

Next-Gen SIEM Evaluation Criteria	Splunk	MS Sentinel	Chronicle	Devo
All Features Included in SaaS Price	● ○ ○ ○	● ● ○ ○	● ● ○ ○	● ● ● ●

Playing Well with Others

A next-gen SIEM makes use of a well-documented, open API to integrate with solutions from any vendor. In today's world of hybrid cloud, multi-vendor technology, next-gen SIEMs don't develop their solutions in a walled garden.

CLOSED ECOSYSTEM VS. OPEN ARCHITECTURE

Since legacy SIEMs have been around for more than a decade, most legacy vendors have had a chance to integrate their offering into a suite of tools that include capabilities such as ticketing systems, SOAR platforms, threat intelligence, etc. To try and wring more revenue out of their customers, SIEM vendors started limiting integrations with outside vendors. For example, your SIEM would work great with the same vendor's SOAR platform but wouldn't work nearly as well with SOAR platforms from other vendors.

Such a walled-garden approach ensures that a single vendor maximizes the revenue generated by each customer. Interestingly, this approach wasn't solely based on greed; it is easier to develop and test new integrations when one company has complete control over the entire codebase of all the solutions. Of course, while this is advantageous for the vendor, it prevents your organization from easily integrating new cutting-edge technologies into the heart of your security stack.

Fortunately for your security team—and your budget—next-gen SIEMs take a more modern, open approach to integration. A next-gen SIEM makes use of a well-documented, open API to integrate with solutions from any vendor. In today's world of hybrid cloud, multi-vendor technology, next-gen SIEMs don't develop their solutions in a walled garden. This more modern, open architecture gives customers the ability to change a piece of their security portfolio with minimal impact on the rest of the organization's existing solutions.

EVALUATION CRITERIA: DO YOU PLAY WELL WITH OTHERS?

Your SIEM does'nt operate in a vacuum. It must play well with the rest of your security ecosystem. And that security ecosystem is a two-way street. Yes, your SIEM must be able to ingest data from firewalls, EDR, IPS/IDS, and your cloud environments. But it also must work seamlessly with the SOAR platform you use today and any other solutions you might adopt down the road.

Therefore, part of your evaluation criteria should be how easy it is to ingest data sources and send data out via an open API. The more open and feature-rich the API, the better. You can't allow your organization to get locked into a situation where your SIEM platform limits what solutions you can integrate with others. Make sure that in addition to thoroughly testing the SIEM's open API you also ask the vendor about specific integrations it supports with other vendors' technologies.

Another important area to investigate is if the SIEM you're considering works with multiple cloud providers and in multiple regions. If the SIEM only runs in AWS, or only in a handful of AWS regions, that could cause problems as your organization grows. If your organization already operates globally, compliance rules may require it to store data from certain regions in that region, but your security team, wherever they are based, still needs to be able to access and search it. So the SIEM you choose should not only be cloud-agnostic, but multitenant as well. A multitenant SIEM that can store data in a region but also make it searchable across regions, which will give you global visibility while maintaining compliance.

Splunk

Splunk plays well with others on the data ingestion side, but not on the integration side.

Splunk does ingest data from just about any source, on prem or in the cloud. But Splunk wants you to use everything in its ecosystem. Splunk has its own SOAR and wants you to use it. It can be difficult for users of Splunk Enterprise Security to integrate with a different SOAR platform.

Microsoft Sentinel

Sentinel plays well with anything inside the Azure stack. Microsoft includes a SOAR as part of the solution and uses playbooks to automate tasks and responses to alerts and detections. This is done using Azure Logic Apps as the connectors between Sentinel and other components or services.

However, automating tasks for anything outside of Azure, such as AWS or Google Cloud Platform (GCP), will be much more difficult and require a great deal of effort and coding. Sentinel uses the common event format (CEF) as a schema for all data. So everything must be parsed into that format, and the original format of the event is permanently lost. For customers that are 100% Azure, Sentinel has a great deal of flexibility, but for customers with a multi-cloud environment, it may not be the best fit.

Google Chronicle

Like Microsoft's Sentinel, Chronicle plays well with anything from the GCP ecosystem, but doesn't integrate well with anything else. Just as Sentinel uses the ASIM as a schema to parse all data, Chronicle uses its own Unified Data Model, or UDM, as a schema to parse all data. This means the original message format is lost, and anything not parsed into this schema at ingestion time is also permanently lost. Like Microsoft, Google also has its own SOAR and Google isn't interested in integrating with anything outside of its own ecosystem. This means automating actions outside of the GCP platform is going to be difficult at best.

Devo

Devo is the most agnostic solution of the three SIEMs profiled, thus it works well with most other technologies. Devo has a fully extensible API and can work with the SOAR platform of your choice, regardless of provider. Devo can ingest data from virtually any source, in structured or unstructured formats. Unlike Sentinel and Chronicle, Devo parses data at query time and NOT on ingestion, which preserves the original event in case you want to parse it differently in the future. This also makes Devo the most change tolerant solution since changing data format does not break ingestion.

A multitenant SIEM that can store data in a region but also make it searchable across regions will give you global visibility while maintaining compliance.

Next-Gen SIEM Evaluation Criteria	Splunk	MS Sentinel	Chronicle	Devo
Open Architecture	●●○○	●○○○	●○○○	●●●●

Architecture

The parse-on-ingest approach has several drawbacks. The biggest drawback is the delay between when the raw data is received and when it's parsed and indexed.

LEGACY VS. NEXT-GEN ARCHITECTURE

One of the easiest ways to distinguish a legacy SIEM from a modern next-gen SIEM is to look at the differences in how they parse and store data.

How a SIEM parses data is one of the most subtle yet most important factors in differentiating a legacy SIEM from a next-gen SIEM. It may seem like a technical detail, but this subtle distinction has vast implications across many aspects of the SIEM.

Legacy SIEMs parse data on ingest. This entails receiving the raw log or event, putting it in a queue, breaking the data into a set of predefined fields, and finally indexing the data to make it searchable.

This parse-on-ingest approach has several drawbacks. The biggest drawback is the delay between when the raw data is received and when it is parsed and indexed. During this "dead time," no alerts or searches can occur on the data until the parsing and indexing are complete. Another drawback of parsing on ingest is that large spikes of incoming data cause CPU contention and slow down search performance—usually at the exact time you need to know what's going on.

Another drawback is that changing the data format can cause problems with data ingestion. Migrating to a new firewall vendor or even upgrading to a new software version can affect the format of the data, which breaks the parser and results in it missing new data. Even when you do parse the data correctly at ingest, it alters the raw data forever.

This leaves you with only the processed version of the data—not the original raw data. When that occurs, you can never go back and parse the data differently or look for something your parser may have ignored. From that point forward, you can only parse what you already know to look for, which means you will miss everything you don't know to look for in the data. This can lead to dangerous oversights in your security posture.

For all the reasons above, next-gen SIEMs don't parse on ingest. Instead, they parse on query. Take a minute to think about all the advantages of the next-gen SIEM's parse-on-query approach. First, the raw data is stored as soon as it hits the platform. This means it's immediately searchable and alertable. The lag time can last between 15 and 30 minutes. But in today's security world, even a 15-minute delay can make a critical difference between stopping an intrusion or letting it spread. That's why next-gen SIEMs store data raw, so you can always go back and parse it differently to look for a piece of data you might have missed previously. You also don't need to worry about a new data source or a change in data format breaking your ingestion and creating gaps in data. Finally, since you store the data raw as soon as it comes in, your compute resources on ingestion are far more efficient and enable you to handle large spikes in incoming data with ease. For all these reasons, parsing on query is far superior to parsing on ingest and is a key indicator of a modern, next-gen SIEM.

Another area of contrast between a legacy and next-gen SIEM involves the way data is stored. Legacy SIEMs store data in a normalized fashion—usually in multiple databases or data storage systems. That's because legacy SIEMs are modular, and each module sometimes has a different data store. For example, network or NetFlow data would be stored separately from log files. There are a few reasons for this, one being that legacy storage systems simply can't scale well with multiple data sources, causing queries to take an unacceptably long time. Another reason is the data in a legacy SIEM needs to be normalized into a set of fixed fields to match a data structure such as that of a database table.

The result of this legacy SIEM approach to data storage is that you typically have multiple places where you store the data and thus must run multiple queries. This approach has two inefficiencies at scale. First, it requires more compute power since you must run multiple queries against multiple data sources. Second, it requires more storage space since you must allocate storage multiple times for different data sets. These inefficiencies can cripple a legacy SIEM at today's cloud-scale of data and security threats.

Fortunately, a next-gen SIEM has a completely different architecture to handle these problems. A true next-gen SIEM stores all data in a single place and compresses it to keep the data as small as possible. Since all data is in one place, you don't need to run queries against multiple data stores to see a correlated result. This makes much more efficient use of compute power.

> **A true next-gen SIEM stores all data in a single place and compresses it to keep the data as small as possible**

EVALUATION CRITERIA: HOW DO YOU APPROACH DATA PARSING AND STORAGE?

Your data sources and data formats are going to change over time—probably sooner and more often than you expect. When you change vendors or even upgrade software from one version to the next on the same product, your log format can change. If your SIEM doesn't have a parser for your new data format, you can lose some or all of that data. This is a major administration problem for most organizations.

The easiest way to avoid this potential problem is to deploy a next-gen SIEM that parses on query, instead of on ingest. That way, even if data sources and formats change, you still ingest the raw data. You never have gaps in your data due to parsing errors. Make sure that all the replacement SIEMs you're considering support a parse-on-query rather than a parse-on-ingest approach to data ingestion.

Splunk

Splunk uses a variety of standard ingestion methods, most of which are fairly straightforward. However, Splunk does need to index data before it can be queried or alerted on. In addition, changes in data format can negatively affect data indexing. This can cause gaps in data and break alerts until data is re-indexed. This dramatically impacts Splunk's agility and makes changes in data format a common problem.

Splunk's approach to storage is a common "hot, warm, cold" approach. Hot storage is typically 90 days, with additional hot storage available at a significant additional cost. Splunk uses multiple large indexes to speed up search times. As a result, its data compression ratio is not very good—usually 2:1.

It's also important to look at how your replacement SIEM stores data and the compression ratio of ingest to storage. One of the biggest advantages of next-gen SIEMs over their legacy predecessors is data storage efficiency. Thanks to this approach, you should get more hot, searchable data with a next-gen SIEM. You should be able to search data older than 4 months as quickly and easily as data from last month. This is critical when your SOC analysts are doing an investigation and need to see the first instance of an attack. It's also important for compliance use cases. Finally, being able to compare data from today or this month to data from a year ago is critical to understanding how your environment is changing over time—for better or worse.

For all these reasons, one of your evaluation criteria should be the difference in search performance for data 30, 90, 120 and 365 days old. If the older data is in cold storage and takes significantly longer to search, that should factor prominently into your decision. Slower search performance means slower investigations, something a next-gen SIEM enables you to avoid.

MS Sentinel

Data ingestion for on-prem Microsoft sources or Azure cloud sources is relatively easy. Ingestion from all other 3rd-party sources requires sending data to Log Analytics via Syslog in the common event format (CEF).

This means any data that is not in the CEF, such as custom application logs, is not ingestible by Sentinel unless it is first converted to CEF. And since custom application logs are always changing, this is a difficult maintenance task. If your custom application logs, are important to you then Sentinel is not a good choice.

Sentinel comes with 90 days of storage included in the price. Extra storage is available at an additional cost. The maximum retention time for storage inside of Azure is 730 days.

Google Chronicle

Again, there are a lot of similarities between Chronicle and Sentinel in the way they parse data. Chronicle uses the Unified Data Model to parse all data. Most GCP data sources use this, but non-GCP data sources do not. If custom application logs or data sources outside GCP are important, then Chronicle may not be the best choice.

Unlike Sentinel, Chronicle comes with 1 year of hot storage included. Search speed for data older than 90 days is good for standard fields like IP address, hostname, etc. But custom searches using regex are much slower.

Devo

Not only is data ingestion in Devo easy but it is also the most flexible solution when it comes to changes in data sources and format. Unlike Splunk and Sentinel, Devo does not parse and index data on ingest. It stores data raw and never changes it. Instead of parsing on ingest, Devo uses tags, then stores data in a nested file structure based on those tags. This method gives Devo a few key advantages, the first being that a change in format does not impact ingestion in any way.

Another Devo advantage is your data is immediately searchable on ingest because you don't need to wait for it to be indexed. On the architecture side, Devo's nested file storage enables a 10x data compression ratio, which uses less disk space and makes searching much faster. These advantages enable Devo to include 400 days of always-hot searchable storage in its base price—the most of any vendor. This makes Devo one of the most cost-effective vendors in the SIEM space when it comes to the cost per day of hot searchable storage. Devo offers up to 5 years of storage at an additional cost.

Next-Gen SIEM Evaluation Criteria	Splunk	MS Sentinel	Chronicle	Devo
Parsing & Storage	●●○○	●○○○	●○○○	●●●●

The Ability to Enrich

Enrichments are a force multiplier for SOC analysts, enabling them to make critical decisions about the nature of the data they see.

FLEXIBLE DATA ENRICHMENT AND THREAT INTELLIGENCE

Enrichment means adding useful context to your data. A few common examples of enrichment are using DNS to add machine names to IP addresses in a table, correlating usernames to people names, and geolocating IP addresses to a physical location. Enrichments are a force multiplier for SOC analysts, enabling them to make critical decisions about the nature of the data they see. For example, an analyst can see that a user was logged in from a machine geolocated in California in the afternoon, and the same user was logged in from a machine geolocated in China later that day. This "impossible traveler" scenario is easy to spot when the IP address of the machines used are automatically enriched with geolocation data. But it would be very difficult for an analyst to perform manually for all logins.

Legacy SIEMs, to be blunt, are not good at enrichment. Either they require a time-consuming manual process of cutting and pasting information across various spreadsheets, or have bolt-on modules (which usually cost extra) that use a fixed approach (such as DNS and LDAP lookups). While they provide limited enrichment, they're far less flexible than a next-gen SIEM.

Next-gen SIEMs deliver all these fixed enrichment capabilities and much more. This typically works by enabling users to upload business context data and write custom queries that cross reference the newly collected data with uploaded contextual data. The main advantage of the enrichment capabilities of next-gen SIEMs is they are flexible (many types of data can be used for enrichment), they are programmatically driven (new data is enriched automatically), and they are dynamic (users can add or change them anytime, based on the needs of the business). This gives the customer the ability to add and update enrichments as needed, without being held captive by the vendor's release cycle.

Cyberthreat intelligence is an important enrichment type that requires special attention. Legacy SIEMs typically don't include threat intelligence. In fact, until a few years ago legacy SIEM vendors viewed threat intelligence products as competitors. You can still find web pages that describe the pros and cons of choosing a SIEM vs. a TIP (threat intelligence platform). In hindsight, this is just another example of the poor response by legacy SIEM vendors to the needs of a modern SOC. If the SIEM you're evaluating does not include an integrated threat intelligence platform or wants you to pay extra for it, you're looking at a legacy SIEM.

Next-gen SIEMs come standard with an integrated threat intelligence platform, and the best vendors don't charge extra for it because they know the vital role it plays in an effective cybersecurity program. Since the purpose of the SIEM is to enable security teams to hunt, detect and respond to threats, it's in everyone's interest to share indicators of compromise (IOC) of targeted attacks. And that's exactly what a threat intelligence platform enables you to do.

A next-gen SIEM not only gives you threat intelligence enrichments as part of the SaaS solution and price, but it also gives you the flexibility to import multiple types of threat intelligence from multiple sources and providers. This adds an element of crowdsourcing to security teams and helps SIEM users stay up to date on threat indicators so they can spend more time looking for and responding to those threats instead of updating the product or spending time integrating their SIEM with a TIP. It just makes sense to bundle threat intelligence with a SIEM.

If the SIEM you're evaluating does not include an integrated threat intelligence platform or wants you to pay extra for it, you're looking at a legacy SIEM.

EVALUATION CRITERIA: ARE YOU ABLE TO AUTOMATICALLY ENRICH WITH THREAT INTELLIGENCE?

A next-gen SIEM must be able to enrich your log data with data from other sources to add context that accelerates the ability of analysts to make decisions. When evaluating SIEMs, choose one with many options for enriching log data from other sources. Think about how you'd like to enrich your log and security data with other data sources and make sure the next-gen SIEM you choose supports as many of these as possible. If there is a hard limit on the number of enrichments or a limited number of data sources that can be enriched— you should broaden your search. A true next-gen SIEM provides deep and flexible enrichment capabilities.

One of these enrichment capabilities is threat intelligence. Make sure your replacement SIEM comes standard with threat Intelligence feeds. This will help your analysts identify attacks more quickly and respond efficiently. Ideally, you don't want your threat intelligence platform to be from the same vendor that provides your SIEM. It's unwise to put all your eggs in one basket. You want an open integration that enables your team to bring in data from multiple threat intelligence sources.

The FBI InfraGard Portal, the Department of Homeland Security, MISP, and the SANS Internet Storm Center are just a few examples of open-source threat intelligence feeds. As you evaluate SIEMs, make sure you choose one that supports open integration with one or more threat intelligence feeds, and that the integration isn't difficult. If it takes a lot of work and coding to integrate threat intel feeds, that should be a red flag to keep looking for a different SIEM.

Most next-gen SIEMs not only support multiple threat intelligence enrichments but they also are aligned with the MITRE ATT&CK framework. ATT&CK provides context about the individual parts of an attack to help teams predict an adversary's behavior and next move. You don't want your SOC analysts to have to manually cross reference detections and IOCs. Having your next-gen SIEM fully aligned with the MITRE ATT&CK framework's tactics, techniques and procedures will enable your team to identify threats faster and respond quickly and effectively.

Splunk

Splunk does not offer threat intelligence enrichments out of the box. It does offer the ability to integrate with a TIP, but that integration must be set up manually. The integration process is described in Splunk's documentation under the section "Threat Intelligence Framework."

MS Sentinel

Sentinel does not offer an integrated TIP out of the box. But it does support several connectors with threat intelligence feeds. These must be manually configured in your Sentinel environment. You'll find documentation on how to manually set up data from TIPs in Sentinel's documentation **here**.

Google Chronicle

Chronicle does have threat intelligence built in, and uses information gathered by Google's monitoring of threats from the internet. Additionally, Chronicle has integration with VirusTotal for malware detection, but this comes at an extra cost. However, a major shortcoming in Chronicle's approach to Threat Intel is that you can't enrich it from any other sources.

Devo

Devo comes integrated with the MISP threat intelligence storage platform. This is operational on day one and doesn't require any manual setup, scripting or coding. Other threat intelligence platform integrations, including Recorded Future, also are available.

Devo also has an incredibly flexible capability for other types of enrichment. You can load any type of data into a table and create a lookup that enriches data in one table from data in another. This robust ability to enrich data in any table from any source includes the ability to add business-specific context to the raw log data collected.

Next-Gen SIEM Evaluation Criteria	Splunk	MS Sentinel	Chronicle	Devo
Data Enrichment & Threat Intelligence	● ● ○ ○	● ● ○ ○	● ● ● ○	● ● ● ●

Purpose-built for Analysts

A next-gen SIEM is significantly easier to use and more powerful than its predecessors. And while it requires a bit of time and effort to master, the benefits of a next-gen SIEM should be easy to recognize.

ACCELERATING THE ANALYST WORKFLOW

As mentioned above, legacy SIEMs are usually modular products and require multiple UIs to be open and require a lot of cutting and pasting between modules and UI. This kind of swivel-chair analysis is painful and time-consuming for the SOC analyst. It is also very difficult to collaborate with other analysts or hand off an investigation.

Next-gen SIEMs—for all the reasons already presented in this eBook (and more)—accelerate the analyst workflow by putting all the information and tools they need in a single UI. Next-gen SIEMs also make it easier for analysts to collaborate by providing built-in tools, such as evidence lockers, where analysts can store information so their colleagues can contribute to or take over an investigation during a shift change.

As should be the case with any modern technology, a next-gen SIEM is significantly easier to use and more powerful than its predecessors. And while it requires a bit of time and effort to master, the benefits of a next-gen SIEM should be easy to recognize. At first glance, you may be struck by the many ways a next-gen SIEM is different from the legacy solution you've been using. But before long, it should be obvious that the many benefits your organization will obtain make it worth the investment in time and money to transition to such innovative, feature-rich technology.

But once you recognize that it's time to break up with your legacy SIEM, how should you evaluate the different next-gen SIEMs on the market to find the right one for your organization?

EVALUATION CRITERIA: WILL IT ENABLE MY ANALYSTS TO WORK MORE EFFECTIVELY AND FASTER?

The ultimate test for any prospective next-gen SIEM should be an evaluation where you bring multiple data sources into the SIEM, enrich them for context, add threat intelligence, and use the SIEM for specific use cases. Don't just use it for detection: put its threat hunting, threat investigations, and incident response capabilities to the test, as well. You should come away from the evaluation fully confident that the capabilities of the next-gen SIEM will make your analysts better and enable them to accomplish their critical tasks faster and more effectively.

Test out the breadth and depth of the API with a few automation use cases. Look for features such as case management that enable you to assign an investigation to a SOC analyst. Next, evaluate how easy it is for analysts to share information while collaborating on an investigation.

As with any new piece of technology, it may be unfamiliar to use at first. But after spending a bit of time in the cockpit, the benefits should become crystal clear.

You should come away from the evaluation fully confident that the capabilities of the next-gen SIEM will make your analysts better and enable them to accomplish their critical tasks faster and more effectively.

Splunk

For an experienced Splunk ninja, it could improve analyst performance, but most SOC analysts are not Splunk experts. It takes a tremendous amount of time and training to become proficient with the Splunk platform. Since Splunk uses a proprietary query language (SPL), it's not easy for general security analysts to use. Many SOC analysts struggle with Splunk Enterprise Security.

If you have a team of dedicated Splunk experts who can perform the configuration, set up the dashboards, and build the queries for your SOC analysts, then you might obtain a lot of value from using it. However, if your security team is on its own to do all the setup, configuration, and dashboard and query building, it may be difficult to realize value from such an expensive platform.

MS Sentinel

For organizations that are 100% Azure cloud users (or a combination of mostly Microsoft on-prem technology and Azure cloud), Sentinel could be a very attractive solution. The ease of getting data into Sentinel from Microsoft data sources such as M365 Defender or Defender for Endpoint gives Sentinel a quick time to value curve for organizations with an entirely Microsoft ecosystem. The Logic Apps offer a way to automate tasks and responses within your Azure environment without an incredible amount of coding.

The weakest link in the Sentinel story is arguably the underlying Microsoft SQL database service, which doesn't have a great reputation for being the most performant and scalable database. This could lead to long query times, and thus long investigation times, as data ingestion scales up.

For organizations with a broad mix of Microsoft and non-Microsoft technologies, Sentinel could be more trouble than it's worth. Onboarding custom application logs will be especially difficult. And automating tasks and responses in other cloud providers such as AWS and GCP will be equally troublesome.

Google Chronicle

Of all the products listed in this guide, Chronicle is the most immature and that fact is most apparent when it comes to usability and workflows. Custom visualizations and dashboards have to be created in Google's completely separate data visualization tool Looker. But Looker doesn't query Chronicle directly - it has to query the BigQuery data lake. And while Google does a good job with detections inside Chronicle, for detailed threat hunting it seems most of that will be done by going to BigQuery directly or through visualizations in Looker. This is a result of a limited number of search types in Chronicle itself. While it is easy and fast to search for IP address, hostname, and a few other things in Chronicle, it is very difficult to do complex multi-field searches. For these kind of searches, you have to go outside Chronicle. This leads to a lot of swivel chair analysis - the last thing an analyst wants to do.

Devo

Devo makes SOC analysts more effective in many ways. Devo's 400 days of always-hot searchable data makes it easier and faster to go back and conduct investigations to see the first occurrence of a threat in your environment. Since data is immediately searchable as soon as it hits the platform, there are no delays between when something happens and when you can alert or search on it. Devo's lightning-fast query performance means shorter query times and thus, faster investigations. And Devo offers many enrichment capabilities to add context to your data, automating many investigation tasks for analysts so they can reach the right answer faster.

Next-Gen SIEM Evaluation Criteria	Splunk	MS Sentinel	Chronicle	Devo
Increased Analyst Effectiveness	● ○ ○ ○	● ● ○ ○	● ○ ○ ○	● ● ● ●

Conclusions: Scoring the Best SIEM

Now that you've learned about the differences between legacy and next-gen SIEMs, and you know which criteria to use when evaluating vendors, it's easier to apply that knowledge to help cut through the noise and determine which is the best next-gen SIEM vendor for your organization.

Of the four vendors in this guide, Splunk is the least attractive choice. Although it has a rich feature set, it is essentially a legacy SIEM. And its pricing model is complicated and expensive.

Microsoft Sentinel is a true next-gen SIEM, but it's most suitable for organizations that have predominantly Microsoft technology stacks. This bias makes it a solid niche player, but it will not work for many of today's large, multi-cloud enterprises.

Chronicle shows promise, but it's still a very immature product that relies on several poorly (or not at all) integrated solutions from Google to make it viable.

Devo has all the hallmarks of the next-gen SIEM in terms of being cloud native, massively scalable and performant, extremely interoperable due to raw data ingestion, and comes with the most impressive set of features out of the box with no additional costs.

IS SPLUNK WORTHY OF CONSIDERATION AS A NEXT-GEN SIEM? **NO.**

Although Splunk has the rich feature set you'd expect from a market leader, it's not a true next-gen SIEM. It was designed and built to be run on prem, and its lifted-and-shifted architecture doesn't benefit from the move to the cloud. And like many legacy vendors, Splunk wants to charge extra for every single feature, which quickly escalates cost, making it a very expensive solution. This is particularly true for storage, where the price jumps quite high if you want more than 90 days of hot, searchable data.

IS MICROSOFT AZURE SENTINEL A NEXT-GEN SIEM WORTHY OF CONSIDERATION?

Yes, if your organization is exclusively or predominantly in the Microsoft ecosystem. No, if your organization relies on a broad mix of technologies and cloud services.

IS GOOGLE CHRONICLE A NEXT-GEN SIEM WORTHY OF CONSIDERATION? **NO.**

Chronicle is cloud native and has scalability but it is still too immature a product. Many core capabilities such as case management and IR are in the SOAR product that was just recently acquired from Siemplify and are not well integrated into the SIEM. The fact that you have to go outside of Chronicle to other products such as Looker and BigQuery for dashboards, visualizations, and complex searching makes for cumbersome workflows. Significant development work needs to be done to integrate all the other Google solutions together to make a comprehensive solution.

IS DEVO WORTHY OF CONSIDERATION AS A NEXT-GEN SIEM? **YES.**

Devo is not only a true next-gen SIEM, but it offers the flexibility required by large enterprise accounts with multiple technology stacks across multiple cloud providers. Devo's ability to ingest data raw, with no indexing, makes it an ideal solution for customers with rapidly changing technologies. And its ability to scale out to terabytes of ingestion a day while offering 400 days of always-hot searchable storage makes it an ideal fit for very large organizations with long-term data needs. Finally, Devo's simple, all-inclusive pricing model makes understanding and predicting costs easy—now and in the future.

Devo is not only a true next-gen SIEM, but it offers the flexibility required by large enterprise accounts with multiple technology stacks across multiple cloud providers.

Next-Gen SIEM Evaluation Criteria	Splunk	MS Sentinel	Chronicle	Devo
Truly Cloud-Native SaaS	● ○ ○ ○	● ● ● ●	● ● ● ●	● ● ● ●
All Features Included in SaaS Price	● ○ ○ ○	● ● ○ ○	● ● ○ ○	● ● ● ●
Open Architecture	● ● ○ ○	● ○ ○ ○	● ○ ○ ○	● ● ● ●
Parsing & Storage	● ● ○ ○	● ○ ○ ○	● ○ ○ ○	● ● ● ●
Data Enrichment & Threat Detection	● ● ○ ○	● ● ○ ○	● ● ● ○	● ● ● ●
Increased Analyst Effectiveness	● ○ ○ ○	● ● ○ ○	● ○ ○ ○	● ● ● ●

Devo
255 Main Street
Suite 702
Cambridge, MA 02142

Devo is the only cloud-native logging and security analytics platform that releases the full potential of your data to empower bold, confident action. With unrivaled scale to collect all of your data without compromise, speed to give you immediate access and answers, and clarity to focus on the signals that matter most, Devo is your ally in protecting your organization today and tomorrow. Devo is headquartered in Cambridge, Mass. Learn more at **www.devo.com**.