


THE CROWDSTRIKE SECURITY CLOUD

Transforming security for today's modern cloud business

TABLE OF CONTENTS

4	CHAPTER 1: SECURITY ISSUES WITH CLOUD INFRASTRUCTURE
4	HUMAN ERRORS
4	RUNTIME THREATS
6	SHADOW IT
6	LACK OF CLOUD SECURITY STRATEGY AND SKILLS
8	CHAPTER 2: ADVERSARIES HAVE THEIR HEADS IN THE CLOUD
11	TIPS FROM THE CLOUD SECURITY EXPERTS
12	CHAPTER 3: PERCEPTION VS REALITY
12	THE CLOUD IS DYNAMIC
13	MULTI-CLOUD ENVIRONMENTS AND VARIED WORKLOADS
14	SECURITY CONTROL DIFFERENCES CAN LEAD TO MISCONFIGURATIONS
15	CHAPTER 4: THE CROWDSTRIKE WAY
15	FOCUS ON THE ADVERSARY
16	REDUCE THE RISK OF EXPOSURE
17	MONITOR THE ATTACK SURFACE
20	CHAPTER 5: STOP ADVERSARIES WITH CLOUD ANALYTICS — CROWDSTRIKE THREAT GRAPH BREACH PREVENTION ENGINE
21	BUILDING BLOCKS FOR BREACH PREVENTION
21	FEEDING THE GRAPH
22	FULL VISIBILITY, INSTANTLY
23	VISUALIZATION
24	HISTORICAL AND RETROSPECTIVE SEARCHES
25	SECURITY AUTOMATION
25	FINDING HIDDEN THREATS
26	MANAGED THREAT HUNTING FOR THE CLOUD AND BEYOND
27	HOW THREAT GRAPH WORKS
28	CHAPTER 6: THINK IT, BUILD IT, SECURE IT ...WITH CROWDSTRIKE CLOUD SECURITY FOR WORKLOADS AND CONTAINERS
29	ABOUT CROWDSTRIKE



THE CROWDSTRIKE SECURITY CLOUD TRANSFORMING SECURITY FOR TODAY'S MODERN CLOUD BUSINESS



5T+

FALCON PLATFORM PROCESSES
OVER 5 TRILLION EVENTS PER WEEK



140M+

IOA DECISIONS MADE
EVERY SECOND



15PB+

CROWDSTRIKE STORES 15+
PETABYTES OF DATA IN THE CLOUD



75K+

CROWDSTRIKE STOPS OVER
75K BREACHES ANNUALLY

STATISTICS AS OF MAY 2021, AND GROWING

Upheaval is arguably one of the best words there is to describe the effects of recent events on enterprises in 2020. As quarantine restrictions forced workers remote, it accelerated the growth of both telecommuting and cloud service adoption. The result was a new proving point for the power of cloud computing and the criticality of securing and supporting both cloud applications and a remote workforce.

Gartner has already predicted that end-user spending on cloud services will grow more than 18% globally in 2021. As enterprises shift more of their resources and attention to the cloud, having a cloud-native security platform that can properly secure cloud-native architectures and empower the development practices of DevOps culture is crucial.

As a cybersecurity company that has built one of the biggest cloud architectures in the world, CrowdStrike has gained an exceptional vantage point and garnered unique experience on what it takes to secure cloud workloads and containers. The CrowdStrike Security Cloud processes over 5 trillion events per week, 140+ million indicator of attack (IOA) decisions made every second—and that only represents the streaming data. In addition, CrowdStrike stores 15+ petabytes of data in the cloud, protects 1+ billion containers every day, and stops over 75,000 breaches annually. All that data is combined with a diverse set of security solutions designed to shrink the attack surface for enterprises and obtain visibility into the events taking place across the environment.

As both a cloud customer and a security company, CrowdStrike has a deep understanding of the complexities and risks of protecting corporate data and the cloud infrastructure that holds it. Embracing the cloud is critical to digital transformation initiatives, but for them to be successful, security must transform alongside the business. Quite simply, it is time for enterprises to rethink security to keep pace with an evolving landscape of risks.

CHAPTER 1

SECURITY ISSUES WITH CLOUD INFRASTRUCTURE

As more enterprises adopt private, public and hybrid cloud infrastructure, the need for comprehensive security across cloud workloads and containers increases exponentially. This need has become more urgent as adversaries have quickly turned their attention to the cloud, and cloud breaches continue to rise.

The reasons behind cloud breaches run the gamut, but can be broadly classified into four categories: human errors, runtime threats, shadow IT and poor strategic planning.

HUMAN ERRORS

Because of the nature of cloud environments, the majority of breaches in the cloud are caused by human error. In fact, according to Gartner, through 2025, 99% of all cloud security failures will be the customer's fault. In the cloud, the absence of perimeter security can make those mistakes very costly. These errors can include misconfigured S3 buckets, leaving ports open to the public, or the use of insecure accounts or APIs. Sometimes, organizations are not even aware of what APIs are being used, let alone understanding whether or not they are secure.

Those errors transform cloud workloads into obvious targets that can be easily discovered with a simple web crawler. Multiple publicly reported breaches started with misconfigured S3 buckets that were used as the entry point. Other examples of misconfiguration leading to a breach involve servers in the DMZ that have ports wide open to the world. These configuration issues continue to happen, often leaving workloads and containers publicly exposed.

RUNTIME THREATS

In public clouds, much of the underlying infrastructure is already secured by the cloud service provider (CSP). However, everything from the operating system to applications and data are the responsibility of the user. This is what is referred to as the "shared responsibility model." Unfortunately, this model can be misunderstood, leading to the assumption that cloud workloads are fully protected by the CSP. This results in users

“ CLOUD SECURITY IS A SHARED RESPONSIBILITY.”

unknowingly running workloads that are not fully protected, meaning adversaries can target the operating system and the applications to obtain access. Attackers use zero-day exploits to gain a foothold, then establish persistence by planting advanced persistent threats (APT) and moving laterally within the data center.

Any available attack surface will be leveraged by adversaries. Even securely configured workloads can become a target at runtime, as they are vulnerable to zero-day exploits and unpatched vulnerabilities. In addition, the cloud provides more than just compute power. It has also become a storage facility for intellectual property and confidential documents, making cloud workloads and containers an increasingly attractive target for attackers. This is a trend observed by the CrowdStrike Services team across numerous breaches it investigated this year that originated in cloud workloads, which many adversaries seem to be targeting specifically.

SHARED RESPONSIBILITY MODEL

CUSTOMER	RESPONSIBLE FOR SECURITY "IN" THE CLOUD	CUSTOMER DATA			
		PLATFORM, APPLICATIONS, IDENTITY AND ACCESS MANAGEMENT			
		OPERATING SYSTEM, NETWORK AND FIREWALL CONFIGURATION			
		NETWORK TRAFFIC ENCRYPTION, SERVER-SIDE ENCRYPTION AND DATA INTEGRITY			
CLOUD SERVICE PROVIDER	RESPONSIBLE FOR SECURITY "OF" THE CLOUD (INFRASTRUCTURE)	COMPUTE	DATABASE	STORAGE	NETWORKING
		REGIONS		EDGE LOCATIONS	
		AVAILABILITY ZONES			

“

EVEN IF CUSTOMERS
OWN CLOUD
SECURITY PRODUCTS,
THEY OFTEN FAIL
TO SECURE THEIR
CLOUD WORKLOADS.
IN FACT, ALL LARGE
COMPANIES THAT
EXPERIENCED A
BREACH LAST YEAR
HAD SOME CLOUD
SECURITY SOLUTION
IN PLACE.”


SHADOW IT

Shadow IT is another major issue for enterprise cloud environments. By its very nature, shadow IT challenges security because it circumvents the normal IT approval and management process. The reason behind shadow IT's existence is not normally malicious. It is typically the result of employees adopting cloud services to do their jobs. The ease with which cloud resources can be spun up and down makes controlling its growth difficult. Developers can easily spawn workloads using their personal accounts. These unauthorized assets are a threat to the environment, as they often are not properly secured and are accessible via default passwords and misconfigurations.

Cloud and DevOps teams like to run fast and without friction. However, obtaining the visibility and management levels that the security teams require is difficult without hampering DevOps activities. As DevOps becomes more mainstream, both security and IT teams need to adapt. DevOps needs a frictionless way to ensure that they deploy secure applications and directly integrate with their continuous integration/continuous delivery (CI/CD) pipeline. There needs to be a unified approach for security teams to get the information they need without slowing down DevOps. IT and security need to find solutions that will work for the cloud — at DevOps' velocity.

LACK OF CLOUD SECURITY STRATEGY AND SKILLS

As workloads move to the cloud, administrators continue to try and secure these workloads the same way they secure servers in a private or on-premises data center. Unfortunately, traditional data center security models are not suitable for the cloud. Cloud may give organizations agility, but it can also open up vulnerabilities for organizations that lack the internal knowledge and skills to effectively understand security needs in the cloud. Poor planning can manifest itself in misunderstanding the implications of the shared responsibility model, which lays out the security duties of the cloud provider and the user. It can also manifest in mistakes by DevOps, which may find itself playing a much larger role in security as part of a shift left approach without the necessary skills or knowledge.



OF CYBERSECURITY PROFESSIONALS
BELIEVE THEIR CYBERSECURITY
PROGRAM NEEDS TO EVOLVE TO SECURE
CLOUD APPLICATIONS AND
PUBLIC CLOUD INFRASTRUCTURE

These issues reveal that a good cloud strategy must include education. Educating teams on security best practices such as how to store secrets, how to rotate keys and how to practice good IT hygiene during software development is critical, but it is often lacking. DevOps may be happening, but DevSecOps may not — which is hampering the industry's ability to make the cloud secure.

CHAPTER 2

ADVERSARIES HAVE THEIR HEADS IN THE CLOUD

In many organizations, the adoption of cloud infrastructure had happened to some degree prior to the global pandemic. While the pandemic certainly accelerated the move to the cloud for many organizations, with some electing to go all-in on cloud, other organizations continue to test the waters and gradually move certain services or capabilities into various cloud platforms. CrowdStrike is even seeing some early cloud adopters moving from legacy cloud deployments to new architectures in the hope of gaining improvements in scalability, maintenance and security.

No matter where you are in your cloud journey, managing the security posture of cloud environments can play a critical role in preventing a breach.

One trend CrowdStrike saw in its 2020 CrowdStrike Cyber Front Lines Report involved threat actors targeting cloud infrastructure slated for retirement or simply neglected for various reasons, including:

- Adversaries target neglected cloud infrastructure slated for retirement that still contains sensitive data.
- Adversaries use a lack of outbound restrictions and workload protection to exfiltrate your data.
- Adversaries leverage common cloud services as a way to obfuscate malicious activity.





\$3.86
million

THE AVERAGE COST IN USD OF A
DATA BREACH IN 2020

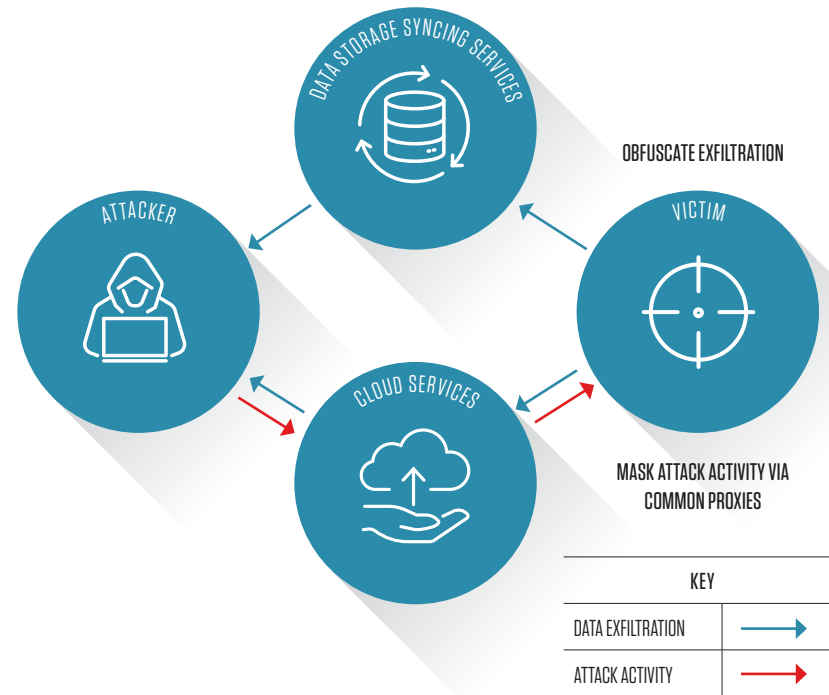
This vulnerability likely stemmed from infrastructure no longer receiving security configuration updates and regular maintenance. Unfortunately, security controls such as monitoring, expanded logging, security architecture/planning and posture remediation no longer occurred in these environments.

Unfortunately, CrowdStrike encountered cases where neglected cloud infrastructure still contained critical business data and systems. As such, attacks led to sensitive data leaks requiring costly investigation and reporting obligations. Additionally, some attacks on abandoned cloud environments resulted in impactful service outages, since they still provided critical services that hadn't been fully transitioned to new infrastructure. Moreover, the triage, containment and recovery from the incident in these environments had a tremendous negative impact on some organizations as these activities disrupted the release of a key feature launch in one case and delayed mergers and acquisitions (M&A) activities in another.

Not only did the CrowdStrike team see cloud infrastructure as a target of attacks in 2020, the cloud also served as a vehicle to launch attacks. Over the past year, threat actors leveraged common cloud services, such as Microsoft Azure, and data storage syncing services, such as MEGA, to exfiltrate data and proxy network traffic. A lack of outbound restrictions coupled with a lack of workload protection enabled threat actors to interact with local services over proxies to IP addresses in the cloud.

This gave attackers additional time to interrogate systems and exfiltrate data from services ranging from partner-operated, web-based APIs to databases to custom network services — all while appearing to originate from inside the victim's network and barely leaving a trace on local file systems.

ATTACKERS LEVERAGE CLOUD SERVICES FOR MALICIOUS ACTIVITY



This demonstrates that attackers are adapting their tactics to target the cloud. Enterprises must adapt their security approach to protect their environment. That evolution should start with a security platform that is purpose-built in the cloud, for the cloud.

TIPS FROM THE CLOUD SECURITY EXPERTS



WHAT CAN I DO TO PROTECT MY CLOUD ENVIRONMENT?

The cloud introduces new wrinkles to proper protection that don't all translate exactly from a traditional on-premises data center model. Security teams should keep the following firmly in mind as they strive to remain grounded in best practices.

1. ENABLE RUNTIME PROTECTION AND OBTAIN REAL-TIME VISIBILITY

You can't protect what you don't have visibility into — even if you have plans to decommission the infrastructure. Central to securing your cloud infrastructure to prevent a breach is runtime protection and visibility provided by solutions like CrowdStrike Falcon® Cloud Workload Protection (CWP). It remains critical to protect your workloads and containers with next-generation endpoint, workload and container protection, including servers, workstations and mobile devices, regardless of whether they reside in an on-premises data center or virtual cluster, or are hosted in a private, public or hybrid cloud.

2. ELIMINATE CONFIGURATION ERRORS

The most common root cause of cloud intrusions continues to be human errors and omissions introduced during common administrative activities. It's important to set up new infrastructure with default patterns that make secure operations easy to adopt. One way to do this is to use a cloud account factory to create new sub-accounts and subscriptions easily. This strategy ensures that new accounts are set up in a predictable manner, eliminating common sources of human error. Also, make sure to set up roles and network security groups that keep developers and operators from needing to build their own security profiles and accidentally doing it poorly.

3. LEVERAGE A CLOUD SECURITY POSTURE MANAGEMENT (CSPM) SOLUTION

Ensure your cloud account factory includes enabling detailed logging and a CSPM — such as CrowdStrike Falcon Horizon™ — with alerting to responsible parties including cloud operations and security operations center (SOC) teams. Actively seek out unmanaged cloud subscriptions, and when found, don't assume it's managed by someone else. Instead, ensure that responsible parties are identified and motivated to either decommission any shadow IT cloud environments or bring them under full management along with your CSPM. Then use your CSPM on all infrastructure up until the day the account or subscription is fully decommissioned to ensure that operations teams have continuous visibility.

4. EMPOWER DEVSECOPS

At CrowdStrike, we know enterprises need security and speed to go hand in hand. Security controls should not slow the speed of application delivery. As developers continue to adopt container technologies to increase velocity, the importance of a shift-left approach to security has only grown. To meet this need, CrowdStrike Container Security was designed to integrate frictionless security into the continuous integration/continuous delivery (CI/CD) pipeline.

This integration takes multiple forms, including continuously scanning container images for known vulnerabilities and misconfigurations, detecting malware in base images before container deployment, and integrating with developer toolchains. By automating protection, organizations can empower DevSecOps to deliver secure applications without slowing the build cycle.


CHAPTER 3

PERCEPTION vs REALITY

When it comes to cloud security, CrowdStrike lives what it preaches. The company has been building and securing its own cloud since 2011. During that time, CrowdStrike has observed the following trends.

THE CLOUD IS DYNAMIC

Providers and consumers of cloud services are moving fast with dozens of new cloud-native services being introduced each year. These services are often aimed at busy developers who are focused on keeping friction low and their velocity high. While most security teams understand their roles in the shared responsibility model, it can be difficult for them to keep up with the changing landscape. Even companies with a strong security program and demonstrated expertise can be at risk of not having sufficient security. For example, a large financial services company with sophisticated cloud security capabilities suffered a breach involving its cloud infrastructure, even though it had previously contributed to an open-source cloud security toolkit. Since the cloud is dynamic, the tools used to secure it must be dynamic and portable in order to work in multi-cloud environments.



OF CYBERSECURITY PROFESSIONALS
BELIEVE THE LACK OF ACCESS TO THE
PHYSICAL NETWORK AND THE DYNAMIC
NATURE OF CLOUD APPLICATIONS CREATES
VISIBILITY BLIND SPOTS

MULTI-CLOUD ENVIRONMENTS AND VARIED WORKLOADS

Attacks can traverse multiple planes and involve different types of workloads. The attack against the financial services company involved tactics that spanned traditional web applications, endpoints and cloud-native resources. Reports about the breach mentioned that an application flaw was exploited to pull a temporary station-to-station (STS) key from the underlying host's EC2 (Amazon Elastic Compute Cloud) metadata service. The key was then used externally to access sensitive cloud resources including S3 buckets.

CrowdStrike investigated an incident that started with an insider threat, with the perpetrator running an exploit on Amazon Web Services (AWS) resources. Leveraging a vulnerability in AWS, the attacker was able to obtain data that was stored in S3 buckets. CrowdStrike expects attacks involving multiple types of workloads and the cloud to become more common going forward.

The visibility needed to see the type of attack that traverses from an endpoint to different cloud services is not possible with siloed security products that only focus on a specific niche. Only a combination of endpoint and cloud-native security tools working together can see attacks of that nature.

Organizations often have more than one cloud to support. Companies running workloads in the public cloud look to improve reliability and availability by adopting a multi-cloud strategy. While definitely a step in the right direction, these companies are finding that not all cloud providers offer the same security features.



“
UNTIL
ORGANIZATIONS
BECOME PROFICIENT
AT SECURING ALL
OF THEIR DIFFERENT
CLOUDS, ADVERSARIES
WILL CONTINUE TO
TAKE ADVANTAGE OF
MISCONFIGURATIONS.”

SECURITY CONTROL DIFFERENCES CAN LEAD TO MISCONFIGURATIONS

Security controls differ from cloud to cloud. Even when cloud service providers offer similar security controls, their behaviors and implementations can vary. Even elements as simple as the log trails needed to support threat hunting, and the design patterns for retrieving them vary from cloud to cloud. These variations make for a steep learning curve, with each public cloud provider offering different sets of security controls. Default configuration settings also vary by provider. Even where there is some overlap, there are different implementations and nuances to deployment. Until organizations become proficient at securing all of their different clouds, adversaries will continue to take advantage of misconfigurations.



CHAPTER 4 THE CROWDSTRIKE WAY

CROWDSTRIKE
USES A THREE-
PRONGED SECURITY
STRATEGY TO GUIDE
ITS CLOUD SECURITY
INITIATIVES.

FOCUS ON THE ADVERSARY

At its core, CrowdStrike's strategy for ensuring security puts the adversary first. In all areas of security, including the cloud, it is critical to understand your adversaries and their modus operandi: who they are, what they want, what they must accomplish to get it and how that maps to an attack surface.

CrowdStrike has observed that many of the same adversaries are active in the cloud and in other parts of the IT landscape. The difference is that the cloud offers adversaries the opportunity to use a new set of tactics, techniques and procedures (TTPs). CrowdStrike continues to research these cloud-native threats and has found that TTPs are maturing for AWS users and emerging across Google Cloud Platform (GCP) and Microsoft Azure. Most current techniques involve adapting traditional attack modes for the cloud, although cloud-only techniques are likely to emerge in the hands of sophisticated adversaries.





Current state-of-the-art techniques include:

Attack tools and post-exploitation frameworks: These are now available for public cloud providers with software such as Pacu and Barq that can use IAM for privilege escalation or use lambda functions for persistence and evasion.

S3 ransomware: There is published research around S3 ransomware, which could be theoretically expanded to any cloud service that offers bring-your-own-key and easy rotation, as those could be potentially vulnerable as well.

Traffic sniffing: Some public cloud providers have recently introduced capabilities in network mirroring, which in addition to improving network monitoring can also allow new paths for packet sniffing and bulk data exfiltration.

Staying informed about these threats can be challenging. That is why having strong partners for threat and situational intelligence helps. Third-party testing, internal red-teaming and bug bounty programs are also valuable when implementing an adversary-focused approach. CrowdStrike uses the CrowdStrike Intelligence and CrowdStrike Services teams to provide ongoing, comprehensive cloud security assessments.

REDUCE THE RISK OF EXPOSURE

CrowdStrike strives to drive down the risk of exposure, so that it is limited to what is needed to run the business. This includes continually searching for and removing unnecessary attack surfaces. As an organization, CrowdStrike cultivates a “security-first” culture that is embraced at all levels of the company — from the C-suite to the newest engineer.

Examples of tactics CrowdStrike uses to reduce the attack surface include:

Segmenting where possible to reduce a potential attack blast radius. This entails using different cloud accounts, virtual private clouds (VPCs), subnets and roles for different types of workloads. Strive to avoid overlapping production, development and integration workloads.

“
THE CORNERSTONE
OF CROWDSTRIKE’S
INTERNAL CLOUD
VISIBILITY STRATEGY
CAN BE SUMMED
UP AS ‘FALCON
EVERYWHERE.’”

Using cloud-native encryption where available for data in flight and at rest in the cloud, and being proactive when it comes to ciphers, protocols, keys and certificates — including having a suite of internal tools to help.

Securing earlier in the process — a practice also known as “shift left” — by implementing tools, automation and standards to enable engineers to easily follow the desired security behavior. These tools reduce developer friction as well as diminish the likelihood that unsafe or default configurations in the wild will be used.

Using MFA where available and hard tokens for high-impact environments such as GovCloud deployments.

Proactively maintaining good IT hygiene by automatically discovering the cloud workload footprint.

MONITOR THE ATTACK SURFACE

Always look for ways to improve visibility into the necessary attack surface. This makes it more challenging for adversaries to hide and also drives up their attack costs.

The CrowdStrike Falcon® platform provides comprehensive visibility across CrowdStrike’s cloud infrastructure. In fact, the cornerstone of CrowdStrike’s internal cloud visibility strategy can be summed up as “Falcon everywhere.” This approach consists of deploying the Falcon agent on all cloud workloads and containers and employing the Falcon OverWatch™ team to proactively hunt for threats 24/7. In addition, CrowdStrike uses specific cloud-native indicators of attack (IOA), analyzes machine learning (ML) patterns and performs free-form threat hunting, looking for hands-on-keyboard activity by adversaries within CrowdStrike’s cloud workloads and control plane.



This level of visibility coupled with proactive threat hunting has allowed CrowdStrike to detect subtle, nearly imperceptible behaviors with uncanny accuracy, such as an incident in which an adversary was probing for the existence of certain S3 buckets. Those buckets were not publicly accessible, and they were named in a way that made using brute force impossible, which prompted CrowdStrike analysts to investigate how the adversary could have obtained a list of the S3 buckets. After considerable research, CrowdStrike intelligence sources surmised that the adversary was probably pulling S3 bucket names from sampled DNS request data they had gathered from multiple public feeds. That type of data is easily obtained by accessing resources from public Wi-Fi. The lesson here is that the adversary sometimes has more knowledge of and visibility into an organization's cloud footprint than you might think.

The expertise the CrowdStrike team has gained firsthand by defending its cloud helps the product team continue to expand the Falcon platform, fueling enhancements and creating new cloud workload and container security features and enhancements to our security posture management toolsets.



CHAPTER 5

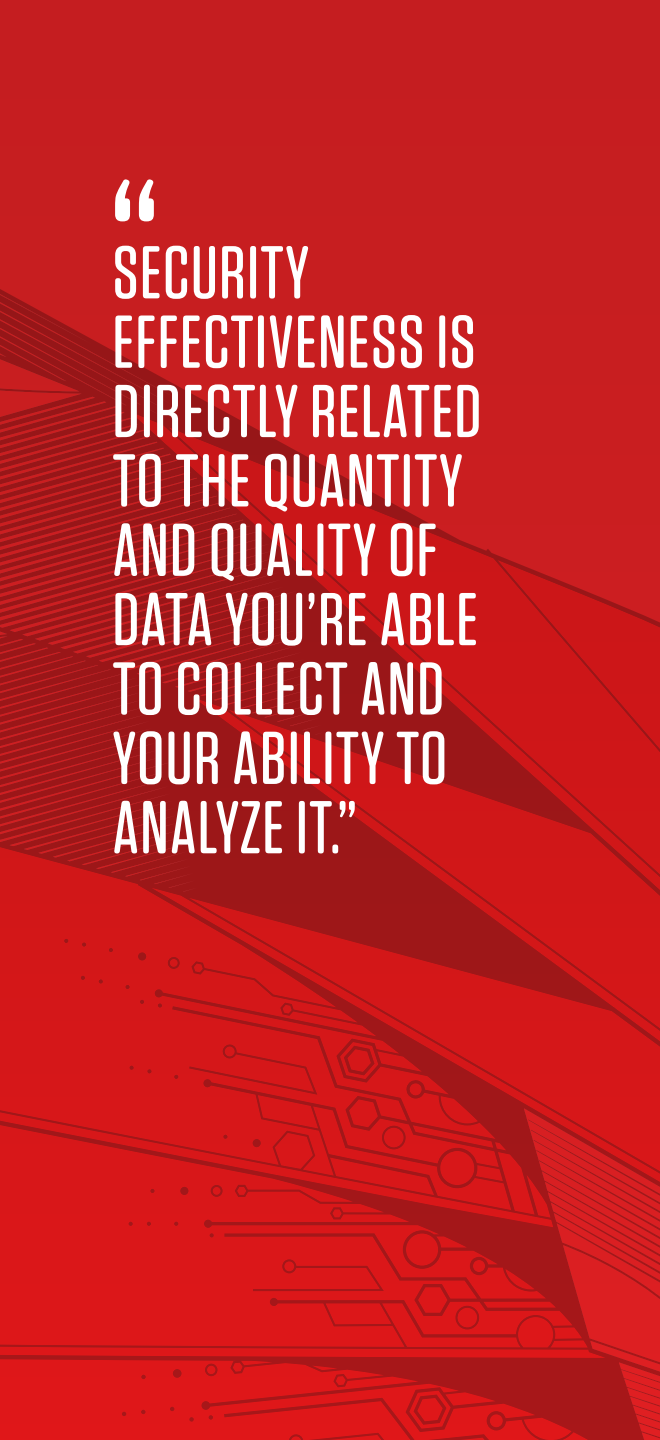
STOP ADVERSARIES WITH CLOUD ANALYTICS — CROWDSTRIKE THREAT GRAPH BREACH PREVENTION ENGINE

Yesterday's techniques for detecting and blocking threats at the endpoint are ineffective against today's modern cloud threats. Breaches can no longer be reliably prevented by monitoring and scanning files and looking for known bads.

Security effectiveness is directly related to the quantity and quality of data you're able to collect and your ability to analyze it regardless of where it comes from. Preventing breaches requires taking this data and applying the best tools, including AI, behavioral analytics and human threat hunters. It leverages this massive data to continuously predict where the next serious threat will appear, in time to act.

Harnessing the data to effectively stop breaches presents significant challenges. The first to overcome is the sheer volume of data. As endpoints, workloads and containers generate billions of events daily, the amount of data to analyze over time can reach petabytes. Second, that data is often unstructured, discrete and disconnected. Without adequate structure, determining how individual events may be connected to an impending attack becomes a tedious and time-consuming manual process. In such an environment, detecting attacks is often difficult, and sometimes impossible.

CrowdStrike sought to solve this challenge by employing a graph data model to collect and analyze extremely large volumes of security-related data. Since no commercial graph database solutions were capable of meeting the unique requirements of cybersecurity, CrowdStrike designed and built its own — the CrowdStrike Threat Graph™ — to store, query and analyze relevant security events.



“
SECURITY
EFFECTIVENESS IS
DIRECTLY RELATED
TO THE QUANTITY
AND QUALITY OF
DATA YOU’RE ABLE
TO COLLECT AND
YOUR ABILITY TO
ANALYZE IT.”

Graph theory is far from new. In fact, it has been used to solve mathematical problems for centuries. The sheer power and scalability of the graph data model has led to its adoption by some of the largest technology companies on earth, including Facebook, Microsoft and Google. CrowdStrike is the first to purposefully use a graph database for cybersecurity.

BUILDING BLOCKS FOR BREACH PREVENTION

The powerful concept of graph data modeling is at the heart of Threat Graph, a powerful and massively scalable graph database that resides in the cloud. Custom-built by CrowdStrike, Threat Graph’s capability for storing, visualizing, correlating and analyzing the vast quantity of event data generated by endpoints provides the Falcon platform with its unique ability to identify attacks in progress and actually stop breaches.

FEEDING THE GRAPH

Threat Graph is “fed” by a variety of sources. In addition to endpoint, cloud workloads and container telemetry transmitted directly from Falcon agents, it receives threat intelligence from the CrowdStrike Intelligence team and from a variety of third-party sources. Its graph data model allows Threat Graph to process billions of events daily, streaming from millions of agents, and to support more than 500,000 event writes per second. It also provides the ability to grow by orders of magnitude to accommodate petabytes of additional data. As a graph database built to fully utilize a highly elastic cloud infrastructure, Threat Graph can scale to meet any volume requirements.

This ever-growing, enriched data set creates a perfect environment for analyzing security data at an unprecedented scale, forming one of the primary pillars supporting CrowdStrike’s breach prevention capabilities.

FULL VISIBILITY, INSTANTLY

Threat Graph data becomes immediately available for viewing, visualizing and searching retrospectively. Some key use cases for taking advantage of this capability include the following:

Real-time Attack Visibility and Drill Down

Threat Graph's instant access to data allows users to see and trace process execution on any endpoint, workload or container in their environment, including rich contextual information such as the process name, arguments and exact time of execution, as well as the sequence of events following that execution (see Figures 1 and 2).

This is key to better understanding the context of code and other executables in an environment. Being able to observe the process, command-line arguments and timing allows security teams to observe suspicious and anomalous activity that could warrant action, and to trace potentially malicious activities in the full context of the affected machine.

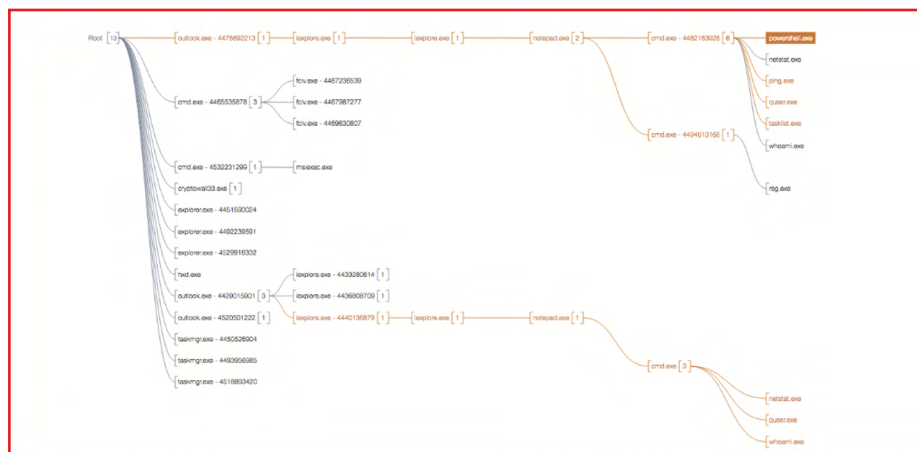


Figure 1: The full context of code execution, displayed in a process tree.

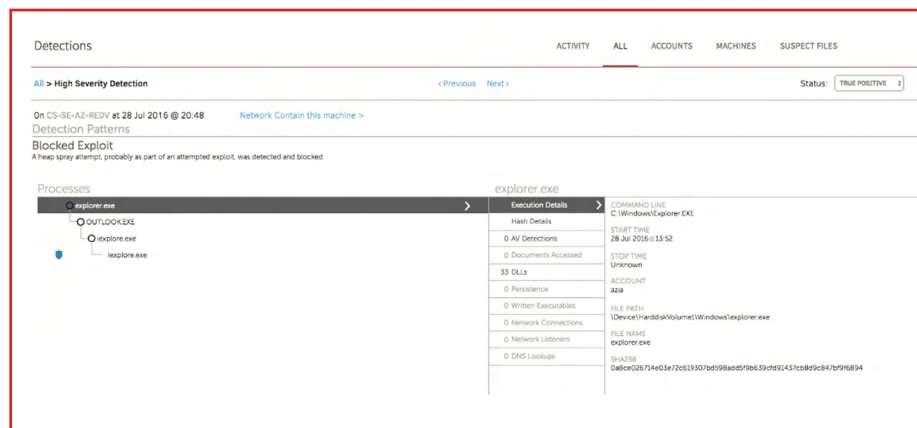


Figure 2: The same chain of events, instantly displayed in the CrowdStrike Falcon user interface.

VISUALIZATION

Another advantage of Threat Graph lies in its ability to map dependencies and present the information visually. By visualizing data in this form (see Figure 3), analysts can spot outliers, inconsistencies and variances at a glance, and identify potential security issues in seconds. Figure 3 represents a map of remote desktop connections per user account, as it is displayed in the Falcon management interface. The unusually high number of connections established by that account is a potential indication that the account may have been compromised. Threat Graph's data visualization capability allows for instant identification of this suspicious behavior which would otherwise be very difficult to detect — and could possibly go completely unnoticed.

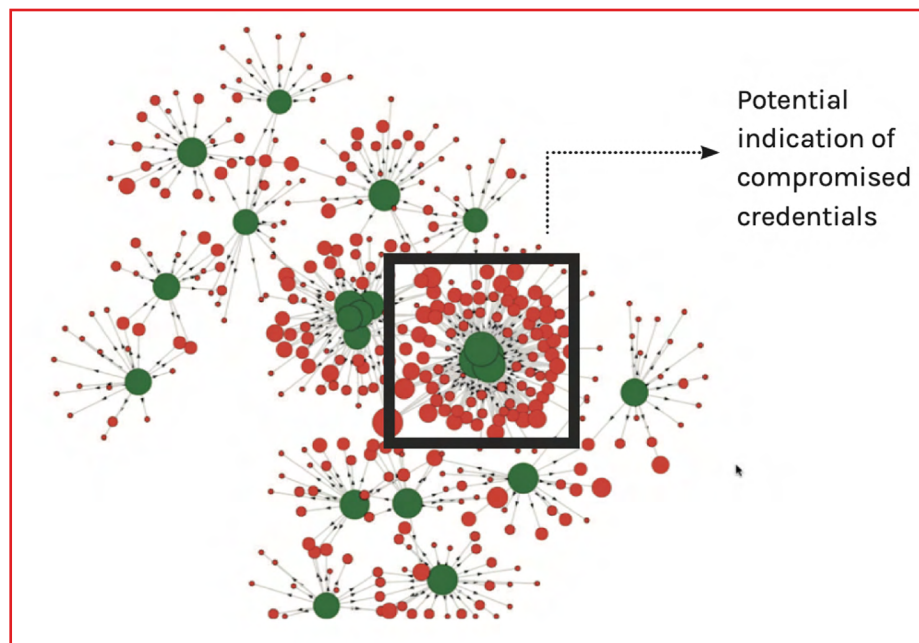


Figure 3: Visualization of remote desktop connections per user and per system.

HISTORICAL AND RETROSPECTIVE SEARCHES

Since the state of each endpoint, workload, container and the environment is kept over time, Threat Graph offers the powerful ability to look back and retrace events as they occurred. This provides analysts with the ability to trigger detections retrospectively. For example, if something is discovered today, its roots can be traced back to perform in-depth forensic analysis — whether the event took place yesterday, last week or last month. In this way, Threat Graph takes discovery and insight forward *and* backward, providing real-time and historical visibility into individual events across a customer's environment.



SECURITY AUTOMATION

The type of real-time visibility and visualization described above is possible because automated analysis is performed on the data as soon as it is written to the database. The graph nature of the Threat Graph allows multiple detection methods and algorithms to run against the data simultaneously, leading to near-instant results. These methods include, but are not limited to, checking against known malware and using machine learning algorithms for detection of unknown malware and IOAs.


IOAs are a means of detection pioneered by CrowdStrike. These indicators reflect a series of actions an adversary must perform to be successful in their attack. IOAs rely on the relationships, context and sequence of events to determine if an attack is in progress. Effective and efficient IOA creation and analysis have been enabled largely by the Threat Graph.

Prior to the Threat Graph, an analyst would have had to gather endpoint, workload and container telemetry — sometimes from multiple sources — then add intelligence feeds, write their own correlation rules, and finally pivot the data endlessly to determine how events might be related. This required slow, labor-intensive processes. In contrast, the Threat Graph offers one unified view of all events and intelligence known to CrowdStrike and its customer base spread across more than 175 countries. The analysis can be automated since all the data — including intelligence, events and most importantly, their relationships — are all kept in one place.

FINDING HIDDEN THREATS

We've seen that Threat Graph enhances and speeds the detection of attacks and patterns of attacks by providing critical visibility, which enables full automation of the analysis. But Threat Graph also excels in facilitating the discovery of new behaviors that have never been observed before: the so-called “unknown unknowns.”

Threat Graph continually looks for malicious activity by applying a combination of graph analytics and machine learning algorithms across its data. The algorithms not only look at file features, but also track the behaviors and sequence of code execution in the



customer's environment. More importantly, the graph data model allows those algorithms to discover relationships between events that are not directly related, but that could constitute an attack that would otherwise remain undetected.

Additionally, Threat Graph allows multiple algorithms to be run simultaneously enabling potential threats, to be discovered much faster. The discovery of new triggers can automatically spawn additional analysis. This allows the whole system to learn and build its own intelligence over time. As more data is fed into the Threat Graph, more attack patterns are discovered. Those new detections are tested and added to the analysis process, increasing Threat Graph's ability to quickly and automatically detect similar attacks. This potent combination of automation and intelligent analysis allows CrowdStrike to find these "unknown unknowns" that elude conventional security measures.

MANAGED THREAT HUNTING FOR THE CLOUD AND BEYOND

Threat Graph enables unprecedented levels of automation to eliminate many of the manual processes that response teams typically endure to detect attacks. However, the threat landscape has changed. Security professionals increasingly must face off with human adversaries, who may have the skill and ingenuity to defeat even the most sophisticated automation. That's why CrowdStrike has always advised that human security assets must be part of the defense chain. Once again, Threat Graph provides unique and exceptional help to achieve that.

A shining example of this is the way Falcon OverWatch, CrowdStrike's proactive threat-hunting team, uses Threat Graph. As noted earlier, Threat Graph first automates the process of discovering triggers, eliminating a very lengthy and tedious process that overtaxes and underutilizes most in-house security teams. Once these triggers are presented to a managed hunting team such as Falcon OverWatch, the hunters can turn back to the Threat Graph for further investigation. At that point, they may need to run ad hoc queries, something at which Threat Graph excels. Unlike traditional databases, which are only really suited to providing answers to predetermined questions on a "big data" scale, graph databases can answer these off-the-cuff questions without delay or difficulty. In addition, since analysts often don't know ahead of time what they will be asking, the ability to answer ad hoc queries is a "must-have" feature for successful and timely detections. In this way, Threat Graph enables the Falcon OverWatch team to hunt much more quickly and efficiently.

Threat Graph



Threat data is collected

Falcon Host sensor telemetry sent to Threat Graph



- 5+ trillion events per week
- From 170 countries
- From all deployed sensors

Data is analyzed

Data is immediately available



- Immediately blocks and alerts
- Process tree visibility
- Forensic analysis
- 5 second real-time and retrospective searches

Automated analysis



Runs a combination of analysis techniques including:

- Signature matching
- Static analysis
- Machine Learning
- Behavioral analysis

Managed hunting



- Continuous proactive hunting by Falcon OverWatch team
- Immediately alerts

Data is enriched

Threat Intelligence added to Threat Graph



Enriched with Falcon Intelligence and third-party threat intel, such as:

- Adversaries and attribution
- Indicators of compromise
- Known bad domains and IPs

Data is made actionable

IOAs creation



- New detections are turned into IOAs
- IOAs are validated
- All sensors updated with new IOAs

Falcon Host sensor updates



CHAPTER 6

THINK IT, BUILD IT, SECURE IT ...WITH CROWDSTRIKE CLOUD SECURITY FOR WORKLOADS AND CONTAINERS

PROTECTION WHEN, WHERE AND HOW YOU NEED IT

CROWDSTRIKE CLOUD AND ENDPOINT SOLUTIONS

ENDPOINT AND WORKLOAD PROTECTION

SINGLE LIGHTWEIGHT AGENT,
SINGLE PANE OF GLASS,
MULTI-PLATFORM SOLUTION


IT AND CLOUD HYGIENE

THREAT INTELLIGENCE AND HUNTING

CLOUD SECURITY SERVICES



CLOUD INCIDENT
RESPONSE



CLOUD COMPROMISE
ASSESSMENT



CLOUD SECURITY
ASSESSMENT



CLOUD RED/BLUE
TEAM EXERCISE



ABOUT CROWDSTRIKE

CrowdStrike, a global cybersecurity leader, is redefining security for the cloud era with an endpoint and workload protection platform built from the ground up to stop breaches. The CrowdStrike Falcon® platform's single lightweight-agent architecture leverages cloud-scale artificial intelligence (AI) and offers real-time protection and visibility across the enterprise, preventing attacks on endpoints and workloads on or off the network. Powered by the proprietary CrowdStrike Threat Graph®, CrowdStrike Falcon correlates 5 trillion endpoint-related events per week in real time from across the globe, fueling one of the world's most advanced data platforms for security.

With CrowdStrike, customers benefit from better protection, better performance and immediate time-to-value delivered by the cloud-native Falcon platform.

There's only one thing to remember about CrowdStrike:

We stop breaches.



TD SYNnex

Public Sector

Learn more at www.crowdstrike.com

© 2021 CrowdStrike, Inc. All rights reserved.