CLEARVECTOR

Identity Intelligence Report 2026

Table of Contents

1 Executive Summary

Report overview

This report explores why production environments are uniquely exposed; what identity activity reveals about AI, human, non-human, and third-party behavior; where existing security approaches fall short; and how predictive models close the identity intelligence gap.

Recognizing the identity intelligence gap

Production environments have become both the primary target of the adversary and the least understood from a security perspective. While most security controls effectively govern authentication and access, these controls provide limited visibility into what identities actually do after access is granted. This blind spot creates a persistent **identity intelligence gap**: organizations can verify who (or what) has access but lack the ability to assess whether subsequent activity is risky or not.

Significance of identity

Our data shows that **91% of identities in production are non-human**, and activity frequently occurs outside traditional operating hours in complex, automated patterns. At the same time, identity-driven breaches are on the rise, and the adversary increasingly relies on valid credentials, making access alone an insufficient signal of trust. For example, the June 2026 Red Hat "Miasma" compromise leveraged stolen credentials and trusted CI/CD workflows to distribute malicious packages through legitimate channels, demonstrating how authorized identities can be abused to perform malicious actions while appearing legitimate. As a result, traditional security models - built to govern human access, check configurations every so often, or examine activity for adversary TTPs - struggle to distinguish legitimate behavior from malicious activity once authentication has occurred.

91%

of identities in production are non-human

Key report findings

A key insight of this report is that existing security tools are not broken; they were designed to answer different questions. Identity providers, corporate endpoint tools, and cloud security platforms collectively confirm access, posture, and configuration. Other security tools examine activity for adversary TTPs - without the context of identity. None of these tools can determine whether a specific identity's behavior is legitimate. The disconnect between identity context and activity context creates a critical visibility gap that the adversary can exploit. Without contextual understanding of identity behavior and risk, high-risk activity (especially when performed with valid credentials) can persist undetected.

Recommended approaches

Closing this gap requires a shift from access-centric, adversary-focused security to **identity intelligence**. By continuously modeling each identity's pattern of life based on observed activity, meaningful risk signals are surfaced in real time. This approach enables instant identification of misuse, reduces time to detection, and provides a foundation for security that is aligned with how modern production environments actually operate.

Why is this important now? The need for identity intelligence is becoming increasingly urgent. AI is ushering in an unprecedented expansion of identities - this enables the adversary to act with greater speed, automation, and sophistication.

2 Introduction

This ClearVector Identity Intelligence Report (CVIIR) is the first in a planned series designed to address a fundamental challenge in modern cybersecurity: Achieving a comprehensive understanding of production environments and the intricate way identities operate within them. The ultimate, strategic objective of this work is to provide the data and insights necessary to prevent and decisively stop the adversary.

While a significant portion of the cybersecurity industry has historically focused on the Tactics, Techniques, and Procedures (TTPs) of known adversaries - a practice that involves continuous chasing and reaction - ClearVector (CV) advocates for a more foundational and strategically advantageous approach.

We believe that true security mastery and the ability to gain the strategic upper hand comes not just from studying the adversary, but from **intimately understanding and monitoring your own environment.** By leveraging your unique, proprietary data, you can build a defensive posture that stops an adversary, even a novel one, without the prior knowledge or extensive study of their specific methods. This report serves as a starting point for utilizing this strategic data advantage.

Report scope and context

This report focuses on presenting key statistics, trends, and analytical insights concerning identities and identity-related activity as observed within live production environments.

Defining the production environment

For the purposes of ClearVector's analysis in this report, a "production environment" is defined broadly as any infrastructure, source code management system, or connected service that hosts or processes live application data. This includes, but is not limited to:

- Amazon Web Services (AWS)
- Google Cloud Platform (GCP)
- Microsoft Azure
- GitHub
- Okta
- Any other system intrinsically connected to hyperscalers or development platforms.

Data inclusion and future expansion

The data presented within this inaugural CVIIR primarily includes activity observed by ClearVector within **Amazon Web Services (AWS) and Google Cloud Platform (GCP)** environments. These platforms represent a significant portion of the global production environment footprint and provide a robust foundation for initial analysis.

As part of our commitment to providing comprehensive coverage, future iterations of the CVIIR are planned to expand the scope of analysis. These subsequent reports will incorporate identity activity and trends from:

- Microsoft Azure
- GitHub
- Okta
- AI identity activity
- Other emerging or critical components of the production environment ecosystem.

Artificial intelligence (AI), generative AI identities, and AI-driven activity

The topic of AI and AI activity is slated for a dedicated, future report in this series. AI-driven activity exhibits a spectrum of credential usage and identity inheritance, which often involves a combination of human, non-human, 3rd party, or AI-specific credentials.

Report series and contact information

The ClearVector Identity Intelligence Report (CVIIR) is the **first publication in a continuous series.** As the adversary evolves, this series will provide ongoing, in-depth analysis to help security professionals stay ahead of the curve.

We value the input of the security community. If you find the data and insights within this report useful, or if you have specific ideas, questions, or requests for data to be included in future reports, please contact us directly.

For more detailed information about ClearVector, our product, and our approach to identity-driven security, please visit our website at www.clearvector.com.

3 Identity: The centerpiece of modern cybersecurity

At ClearVector, we operate from the core conviction that identity is the key to stopping the adversary. By shifting the focus onto what an organization can inherently control - an organization's own universe of identities - it becomes possible to more accurately predict, and more importantly, decisively stop malicious activity. Crucially, a strong identity focus allows for near-real-time detection (seconds to minutes) in the event of credential theft, abuse, or misuse.

This focus on the origin and history of an action is essential for accurate detection, response, and prevention.

The ClearVector identity taxonomy

We categorize all identities into three top-level types, based on their inherent nature and relationship to the organization's control plane:

- **Human Identities (HI):** These are the simplest to define. Human identities represent real, physical people - employees, contractors, or any individual - who exist in the real world and interact with systems directly or indirectly.
- **Non-Human or Machine Identities (NHI):** This category represents any non-person entity that performs an action. This includes but is not limited to service accounts, API keys, managed identities, execution roles (e.g., AWS IAM Roles), container identities, and cryptographic keys. For the purposes of this report, we use the all-encompassing term **"non-human" (NHI)** to refer to these identities.
- **Third-Party Identities (3P):** This category is a critical delineation, representing a company, external vendor, or an individual from an external organization that requires access to the customer's production or enterprise environment. This includes vendor-specific accounts, partner federation roles, or outsourced developer credentials.

The fuzziness of identity: Beyond a simple categorization

Even with these three categories, the lines between Human, Non-Human, and Third-Party can quickly become fuzzy. A natural question that arises is: **Where do things like Generative AI or large language models (LLMs) fit into this structure?** The answer is complex, as we'll show later in this section. Furthermore, many people assume a strict, linear hierarchy or a simple, binary decision model (e.g., *is it Human OR is it Non-Human*). This perspective fails in modern distributed environments.

The reality is that a Human can and often does use a Non-Human identity in a variety of complex ways, making a simple binary decision impossible.

Consider these common scenarios:

- **Standard federation:** A human (HI) successfully authenticates to a central Identity Provider (IdP) and then federates access to a production environment, assuming a specific execution role (NHI) to perform necessary tasks. The human is the initiator, but the role is the active identity in the production environment.
- **Adversary exploitation:** An adversary (HI) successfully exploits a vulnerability in a containerized workload, gains code execution, and then utilizes the container's pre-assigned execution role (NHI) to access the underlying hyperscaler control plane. The activity looks like a machine, but the provenance traces back to the human adversary.

In conclusion, you cannot simply assign an identity to a single, permanent category like "it's human" OR "it's non-human." This interaction is fluid. **Human identities can leverage non-human identities, and vice versa, at various points within a single workflow.** The relationship is better visualized as a dynamic, interconnected network, not a strict hierarchy.

Why delineate third-parties (3P)?

At ClearVector, we place significant emphasis on clearly delineating third-party identities for at least two major security and compliance reasons:

- **Inherent risk and monitoring:** Third-parties inherently introduce a higher level of risk because their security posture is controlled by an external entity, making it a common vector for large-scale breaches (e.g., SolarWinds). Consequently, most organizations are required by compliance mandates (e.g., SOC 2, HIPAA, PCI-DSS) to specifically track and elevate the monitoring of all third-party access.
- **Data model integrity:** From a data model perspective, the Third-Party category explicitly captures the context that an external company, or a human from that external company, is accessing the organization's production environment. This allows for specific policies and detection tailored to these higher-risk interactions. For example, in a typical AWS environment, you may have multiple third-parties with cross-account access to various AWS accounts.

Practical examples of identity flows

The following paths illustrate how all three identity types interact within common, real-world security flows, using HI for Human Identity, NHI for Non-Human/Machine Identities, and 3P for Third-Party Identities:

Standard employee workflow:

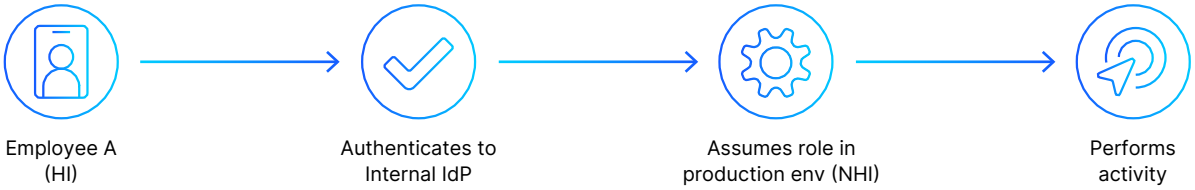

Credential theft by the adversary:

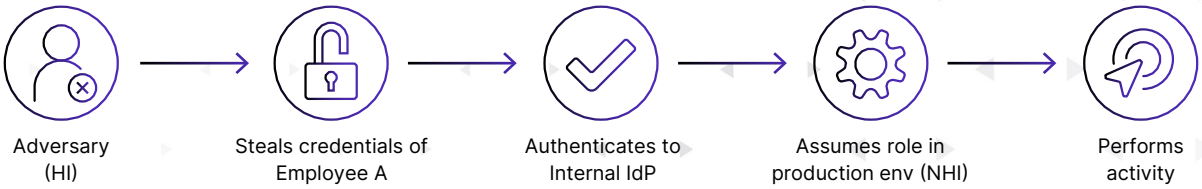

Vendor access and chaining: (The most complex, involving all three types)

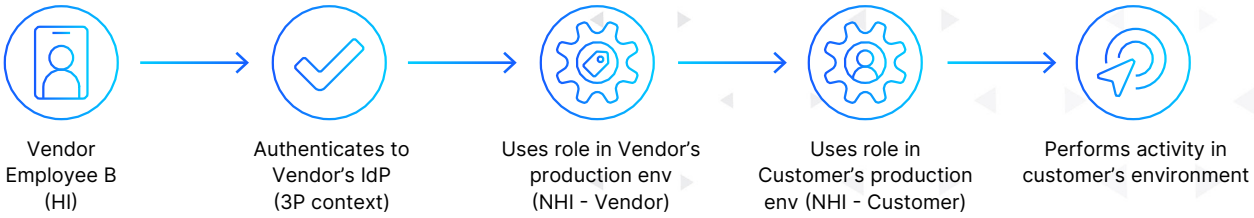

As demonstrated in example three, all three core identity types are utilized in a relatively simple, common vendor access flow. Understanding this complex interplay of identities at runtime is the foundation for effective identity-driven security.

4 Identity trends and activity in production environments

Identity-driven attacks are evolving

As identity activity becomes increasingly automated and distributed across human, non-human, and third-party identities, the adversary is adapting accordingly. Recent incidents show that identity compromise is no longer limited to traditional credential theft - now spanning software supply chains, cloud-native infrastructure, and AI-assisted operations.

Supply chain attack

In June 2026, the Red Hat "Miasma" compromise demonstrated how stolen credentials and trusted publishing workflows could be abused to distribute malicious packages through legitimate channels.

Cloud-native infrastructure breach

In March 2026, the Mercor breach was linked to a compromise of the widely used LiteLLM project, highlighting how trusted third-party software dependencies can create downstream identity and credential exposure across large portions of the ecosystem.

AI-assisted attack

In May 2026, researchers observed an intrusion in which an LLM-driven attack workflow used harvested cloud credentials to retrieve secrets from AWS Secrets Manager and further access protected infrastructure - illustrating how AI can accelerate identity-driven post-exploitation activity.

While these incidents differ in execution, they share a common characteristic. Each relied on trusted identities, credentials, or automated systems to achieve the objectives of the adversary. This reinforces a core theme of this report: Understanding how identities operate in production is as or more important than governing who or what is granted access.

Recent attacks and evolving cybersecurity priorities point to the same conclusion: access alone is no longer a sufficient measure of security. From CISA's BOD 25-01 directive for cloud visibility and monitoring to the White House's March 2026 Cyber Strategy focus on AI-driven detection and measurable security outcomes, the message is consistent: organizations must move beyond determining who has access and toward understanding how identities operate in production. This report showcases, based on data, why identity intelligence is a foundational component of modern cybersecurity strategy.

Identity activity analysis of production environments

ClearVector (CV) analyzed identity activity within observed production environments, yielding significant insights into the composition and behavior of various identity types. The key findings are summarized below, categorized by identity composition and activity observations.

Identity composition by type

CV's observations reveal a significant skew in the types of identities present in production environments, indicating a departure from traditional, human-centric security models.

Identity Type	Percentage of total identities	Key observation
Human (Individual users, administrators)	**9%**	Human users account for only ten-percent of the total identities, suggesting that the primary volume of access and activity is driven by automation.
Non-human (Service accounts, automated tools, etc.)	**87%**	Non-human identities constitute the vast majority of active identities, underscoring the critical need for robust machine identity management and monitoring.
Third-party (External integrations, vendor access)	**4%**	Third-party identities, while the smallest group, represent a distinct risk vector and require dedicated oversight due to their external origin.

Figure 1: *Breakdown of the types of active identities observed in production environments.*

The breakdown in Figure 1 illustrates that more than two-thirds of all identities interacting with production resources are automated or non-human, cementing the necessity for security strategies to prioritize non-human identity governance.

Identity pattern of life (PoL)

Analyzing when identities are active provides crucial context for defining "normal" behavior and detecting anomalies. CV focused on comparing activity during standard business hours (e.g., 9:00 AM to 5:00 PM local time) versus off-hours.

Identity Type	Activity during normal business hours	Implication
Human	**44%**	Human activity is predictably concentrated around the working day, with almost half of the activity occurring outside of standard hours, potentially reflecting administrative tasks, globally distributed teams, or out-of-hours maintenance.
Non-human	**20%**	Less than a quarter of non-human activity occurs during business hours, meaning **more than 75% of non-human activity occurs during off-hours.** This highlights that most automated processes operate continuously or are scheduled for periods of low human interaction.
Third-party	**35%**	Similar to non-human identities, **almost three-quarters of third-party activity occurs outside of the primary business day**, suggesting these integrations are often used for scheduled, batch, or background processes.

Figure 2: *Breakdown of the pattern of life (PoL) of active identities observed in production environments.*

Figure 2 shows that the low percentage of non-human and third-party activity during business hours reinforces that continuous monitoring, rather than simple "working hours" checks, is essential for these identity types.

Identity activity distribution and periodicity

To better characterize the nature of the observed activity, CV further broke down the pattern into three categories: uniform (consistent 24×7 operation), periodic/bursty (scheduled tasks or high-volume short periods), and minimal weekend activity.

Identity Type	Uniform 24×7 activity	Periodic/Bursty activity
Human	3%	51%
Non-human	22%	54%
Third-party	18%	42%

Figure 3: *Detailed analysis of the periodicity of activity of active identities observed in production environments.*

Figure 3 shows three key observations on activity distribution:

1

Non-human and third-party activity dominance: Both non-human and third-party identities exhibit an almost identical activity distribution. For both groups, approximately a quarter of the activity is constant (24×7 uniform), indicating persistent processes like monitoring or data syncing. The vast majority, nearly half of the activity, is periodic or bursty, reflecting scheduled jobs, automated deployments, or sudden bursts of high-volume processing.

2

Human activity characteristics: Human activity is rarely uniform 24×7 (only 3%). The 51% of activity classified as periodic or bursty for human users likely corresponds to specific work sessions, large administrative tasks, or development sprints, as opposed to the predictable scheduling seen in non-human identities.

3

Minimal weekend activity: Across all identity types, CV observed minimal activity during the weekend, with less than 5% of the total activity occurring on Saturday or Sunday. This provides a strong baseline for defining an activity anomaly, as nearly any significant activity surge on the weekend would be immediately suspicious.

A final noteworthy observation is the **fairly uniform similarity** between the identity composition and activity patterns across the various production environments observed by CV. This consistency suggests that the reported statistics represent a general trend in modern, automated production infrastructure rather than an isolated finding specific to a handful of unique environments.

5 In depth: Human identities in production environments

ClearVector's (CV) analysis of production environments during the period of observation provides significant insights into the nature and behavior of human identities.

Human identity presence and activity

- **Human identity share:** Only **9%** of all active identities within the observed production environments were identified as human. The remaining 91% are associated with service accounts, machines, third-parties, or other non-human entities.
- **Activity timing: normal business hours:** The vast majority of human activity adheres to traditional working schedules, with **44%** of all observed human identity activity occurring during standard business hours (e.g., 9:00 AM - 5:00 PM, Monday - Friday, in the relevant time zones). This suggests that for most human roles, production environment interaction is a daytime, workday function.
- **Activity timing: 24×7 uniformity:** A very small subset, only **3%,** of human identity activity demonstrated a uniform, 24×7 pattern. This group likely represents roles with global operational responsibilities, automated maintenance scripts tied to human accounts, or high-priority, on-call support personnel.
- **Activity timing: weekend only:** Activity strictly confined to the weekend was minimal, accounting for only **2%** of the total human identity activity. This low percentage suggests that most non-scheduled work and maintenance requiring human access is either handled during business hours or is critical enough to be categorized in the 24×7 group.
- **Activity periodicity:** Almost half of all human identity activity, specifically **51%,** exhibited a distinct periodic nature. This strong periodicity is a key indicator of scheduled tasks, recurring administrative checks, or regular development deployment cycles.

Human identity resource interactions

To understand the scope and impact of human identities, CV analyzed the median number of unique resources and categories each identity interacted with during the observation period. As shown in Figure 4, the ranges provided (the median value +/- the standard deviation) illustrate a high degree of variability in the breadth of access, even for the “average” identity.

Description	Median Interactions (+/- std dev)	Interpretation
AWS accounts or GCP projects	**1 +/- 20**	Most human identities interact with only 1 unique accounts/ projects, but the standard deviation of 20 indicates that a small, highly privileged group interacts with a significantly larger number of accounts.
Regions	**1 +/- 1**	The typical human identity is focused on a single region. The deviation suggests that global or multi-regional administrators represent a smaller but more complex segment.
Services	**3 +/- 6**	A typical human identity utilizes 3 unique services (e.g., EC2 and S3, or Compute Engine and Cloud Storage). The standard deviation implies that a few “super-users” or high-level architects interact with a six or more unique services.
CV Risk Categories	**2 +/- 1**	The typical identity performs actions categorized into a two CV Risk Categories. This suggests a balanced role profile. The variance indicates that the most privileged or widely deployed identities engage in a broader spectrum of high-risk activities.
User Agents	**3 +/- 4**	The majority of human identities use multiple types of access method (e.g., AWS CLI, browser console, specific IDE). The high deviation highlights identities that utilize multiple access tools, potentially complicating identity tracking and risk assessment.

Figure 4: *Breakdown of the resource interaction of active human identities observed in production environments.*

Human identity activity by CV risk category

A granular breakdown of the nature of the API calls and resource interactions made by human identities, categorized by CV's internal risk model, reveals a heavily skewed distribution towards passive, informational actions:

1

Read-only activity: 99% of all observed human identity activity was classified as Read-only. This includes actions such as Describe, List, Get, and other purely observational functions, suggesting that the primary role of most human identities is monitoring, debugging, and information retrieval.

2

Privileged activity: ~1% of the activity involved privileged actions. This includes configuration changes, policy modifications, and environment settings that do not immediately result in data destruction or service outage but require elevated permissions. This segment represents the core administrative and development work of maintaining the environment.

3

Destructive activity: The riskiest actions, classified as Destructive (e.g., Delete, Terminate, or major security group changes), accounted for less than 1% of the total activity. Although not contributory on a percentage basis, the low volume of destructive activity is a positive indicator of proper role segregation and the infrequent need for highly destructive operations in day-to-day work.

6 In depth: Non-human identities in production environments

ClearVector's (CV) analysis of production environments during the observation period reveals a significant portion of production environments are dominated by automated and programmatic identities.

Non-human identity presence and activity

The data underscores the pervasive role of non-human identities in modern production infrastructure:

- **Dominance of Non-Human Identities:** A substantial 87% of all active identities observed within production environments were classified as non-human (e.g., service accounts, execution roles, managed identities). This highlights the critical need for a security model centered on machine-to-machine interactions.
- **Temporal distribution of activity:** The activity of non-human identities exhibits complex, non-uniform patterns that challenge traditional 9-to-5 security monitoring paradigms:
  - **Normal business hours (9-to-5): 20%** of non-human activity occurred during standard business operating hours. This is often linked to deployment pipelines, automated configuration changes, or developer tool execution.
  - **Continuous uniform activity (24×7): 22%** of activity was uniformly distributed across all hours and days. This signature is typical of persistent monitoring agents, logging services, or baseline infrastructure maintenance tasks.
  - **Weekend activity: 4%** of activity is concentrated during the weekend. While the lowest percentage, this is often the window for large, resource-intensive maintenance, patching, or long-running analytics jobs.
  - **Periodic activity (Scheduled/Triggered): 54%** of all non-human activity is periodic, meaning it occurs at regular, scheduled intervals (e.g., hourly backups, daily vulnerability scans, weekly data synchronization). The high proportion of periodic activity emphasizes the importance of understanding and profiling expected automation schedules.

Non-human identity resource interactions

To quantify the operational scope and complexity of non-human identities, CV analyzed the median number of unique resources and categories each identity interacted with during the observation period. As shown in Figure 5, the variance (represented by the +/- value) reflects the diversity and dynamism of these interactions across the population of non-human identities.

Description	Median	Variability (+/- std dev)	Interpretation
AWS accounts or GCP projects	**1**	**10**	The median identity touches 3 projects/accounts, but the high variability indicates a significant number of "super-connectors" that operate across a large multi-cloud or multi-account footprint, often associated with centralized security or logging roles.
Regions	**1**	**1**	Most non-human identities are typically confined to a single region, suggesting a focus on regional resilience.
Services	**1**	**2**	The median identity interacts with 1 distinct service (e.g., S3, EC2, Cloud Functions, DynamoDB). The variance highlights that complex automation pipelines or core infrastructure management roles may need permissions for a very broad array of services.
CV Risk Categories	**1**	**1**	Most non-human activity is contained within two risk categories, indicating role specialization. The deviation suggests that cross-functional roles (e.g., security auditors) can span multiple risk categories.
User Agents	**1**	**1**	The median identity uses a distinct user agent. This signifies that some automated processes are accessed or executed through a single interface, such as a CLI, SDKs, or custom internal tools.

Figure 5: *Breakdown of the resource interaction of active non-human identities observed in production environments.*

Non-human identity activity by CV risk category

A granular breakdown of the nature of the API calls and resource interactions made by non-human identities, categorized by CV's internal risk model, reveals a critical need for continuous monitoring of non-human identities.

1

Read-only activity: 93% of all non-human activity involves passive operations (e.g., Describe, List, Get). This baseline activity is essential for monitoring and auditing, but it can still be exploited for reconnaissance.

2

Privileged activity: 4% of the activity involved privileged actions that create, modify, or manage resources (e.g. Create, Update, Delete, Put).

3

Destructive activity: 3% The riskiest actions, classified as Destructive (e.g., Delete, Terminate), accounted for 3% of the total activity. While necessary for lifecycle management or automated recovery, this activity profile represents the highest risk and demands immediate, high-fidelity alerting for certain identities.

7 In depth: Third-party identities in production environments

This section provides an in-depth analysis of third-party identities observed within the production environments during the reporting period, breaking third-party identities out from other non-human identities.

Third-party identity presence and activity

ClearVector's analysis reveals that third-party identities constitute a small but critical segment of the overall active identity landscape:

- **Proportion of active identities:** 4% of all active identities observed across production environments are classified as third-parties.
- **Temporal distribution of activity:** The activity of these third-party identities exhibits distinct temporal patterns:
  - **Normal business hours (9-to-5): 31%** of third-party identity activity occurred strictly during normal business hours, suggesting human-driven or third-party responses (automation) to human activity.
  - **Continuous uniform activity (24×7): 27%** of third-party identity activity was distributed uniformly across all hours and days. This indicates continuous monitoring, logging, or background synchronization processes.
  - **Weekend activity:** A notably low **6%** of activity was specifically observed during the weekend period, suggesting most critical or high-volume tasks are either business-hour dependent or part of the uniform 24×7 pattern.
  - **Periodic activity:** A significant **37%** of third-party identity activity is classified as periodic, meaning it occurs at regular, scheduled intervals (e.g., hourly, nightly, or weekly syncs/scans). This represents the largest single activity pattern, highlighting the reliance on scheduled batch processing or regular security/compliance checks by third-party tools.

Third-party identity activity by vendor

Figure 6 provides a breakdown of the activity of the top four active third-party vendor identities observed in production environments, excluding ClearVector's own operational identities.

Description	Coverage	Min Services	Max Services	RO	Priv	Destruct	Risk
Observability Vendor	100%	2	5	100	0	0	Low
Security Vendor 1	100%	4	7	60	0	40	High
Security Vendor 2	100%	2	3	100	0	0	Low
Compliance Vendor	100%	2	2	100	0	0	Low

Figure 6: *Breakdown of the activity of the top active vendors (excluding ClearVector) in the observed production environments.*

All four vendors have 100% coverage across the production environment. Three of the vendors (Observability, Security Vendor 2, and Compliance) CV categorizes as "Low" risk because the observed activity is strictly Read-Only (RO) - however, a future report may provide a detailed breakdown of read-only activity, especially as it pertains to reconnaissance and secrets.

CV categorizes "Security Vendor 1" as "High" risk, due to the observed activity being 60% read-only and 40% destructive.

Figure 7 shows a visualization of the subgraph of Observability Vendor.

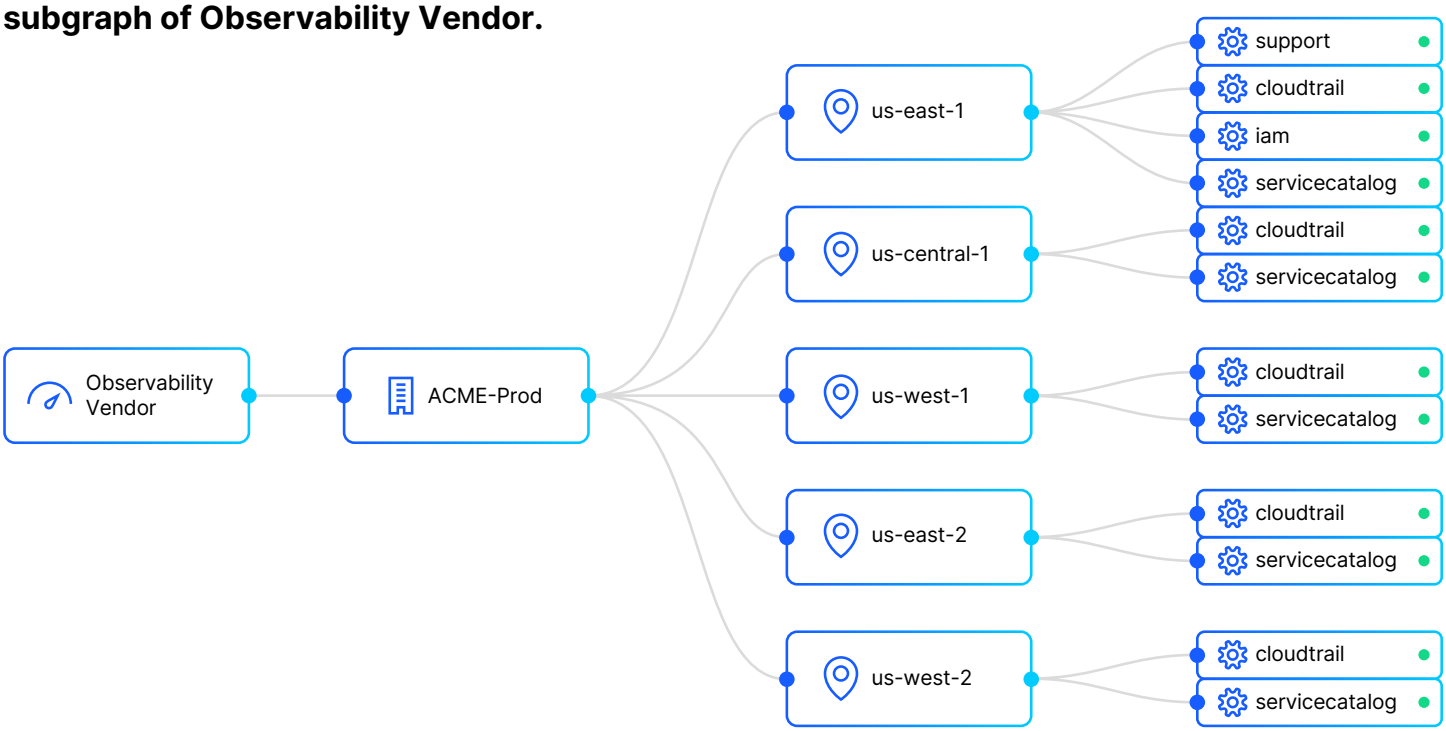

Figure 7: *Subgraph of Observability Vendor.*

Figure 8 shows a visualization of the subgraph of Security Vendor 1.

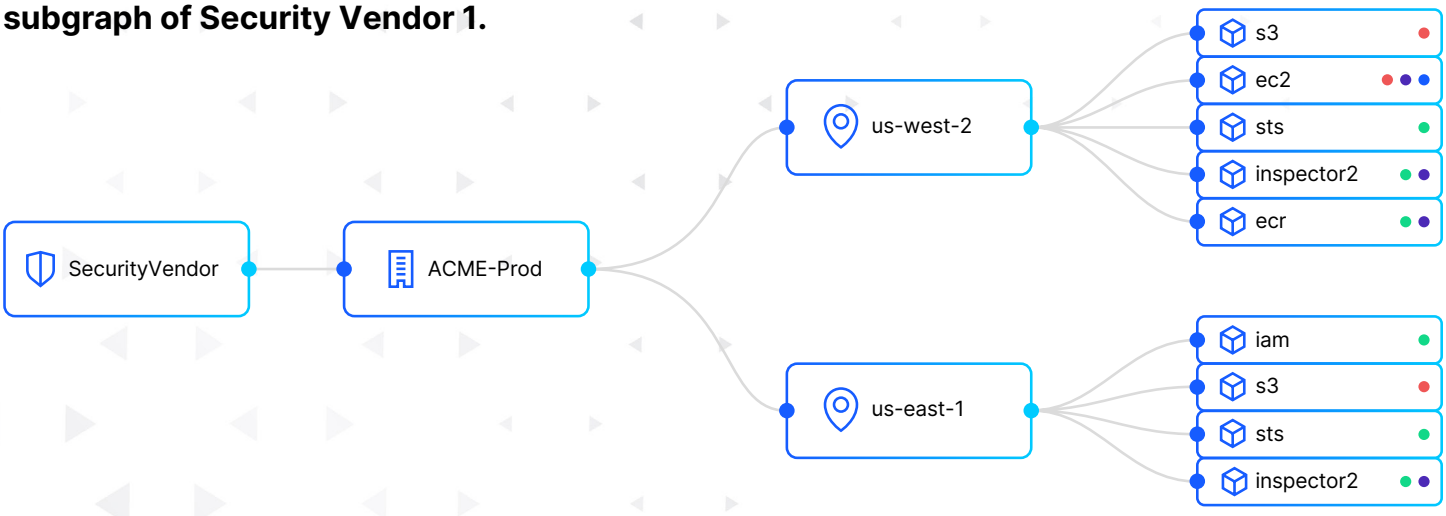

Figure 8: *Visualization of the subgraph of Security Vendor 1.*

8 Conclusion

This report established a foundational, data-driven understanding of how identities operate within modern production environments. The data highlights several critical takeaways for the security community:

1

Non-human dominance: The production environment is overwhelmingly machine-driven. With 91% of all active identities being non-human, the volume and complexity of automated activity dwarf human interaction. This mandates a shift in focus from human-centric models (e.g., VPN, user provisioning) to robust **machine identity governance** capable of continuous, high-fidelity monitoring of programmatic access.

2

The non-9-to-5 reality: A significant majority of non-human **(80%)** and third-party **(65%)** activity occurs outside of standard business hours. Furthermore, nearly half of all non-human activity (51%) is periodic or bursty. This pattern proves that security strategies relying on simple or periodic, time-based anomaly detection are insufficient; effective defense requires deep, contextual profiling of non-human **Patterns of Life (PoL)** to differentiate legitimate automated activity from potentially malicious, outside-of-hours activity.

3

Privileged and destructive activity for non-human identities: While human activity is primarily read-only **(99%)**, non-human identities, by contrast, show a higher risk profile - **5%** of activity involves privileged operations, and **3%** carries the potential for destructive impact. This substantial volume of high-impact, programmatic action means that a compromised machine identity poses an immediate and scalable threat to the environment, requiring detection and response measured in seconds, not hours.

4

The critical third-party: Although only **4%** of identities are third-party, their activity profiles - especially the **40%** destructive activity observed in “Security Vendor 1” - reaffirms that 3rd party controls are critical. Comprehensive oversight and continuous validation of third-party behavior, beyond mere contractual compliance, is essential.

The findings in this report showcase the need for organizations to move beyond reactive security models. Gaining the strategic upper hand is achieved by intimately knowing your own identity universe.

ClearVector believes the most effective defense is one that can detect and stop an adversary based solely on an understanding of **your own environment's normal identity pattern of life**, regardless of the specific TTPs used by the adversary. By focusing on the provenance, scope, and activity of all identities - human, non-human, and third-party - security teams can build a foundation for identity-driven security that is both resilient and future-proof.

About ClearVector

ClearVector is the identity-driven security company. Using Predictive Behavioral Defense, the ClearVector platform builds a unique pattern of life for every human, non-human, third-party, and AI-driven identity in production environments. By combining real-time identity discovery, activity analysis, security-relevant business data, and productized cloud expertise, ClearVector helps organizations prepare for, measure, and prevent the impact of cyberattacks and breaches.

Because each model, or Identity Intelligence Model (IIM), is built automatically for each customer environment from observed activity within each customer, the adversary cannot test or optimize attack techniques across all customers. ClearVector detects and helps stop attacks that traditional signatures and workload-level tools miss, including compromised developer credentials, supply-chain compromises, and rogue AI agents.

With ClearVector, a single click or API call empowers organizations to quickly understand, validate, mitigate, isolate, and control risk across production environments. ClearVector is defining a new era of cybersecurity by raising the bar for the adversary through the power of identity, identity intelligence, and productized expertise.

Learn more at www.clearvector.com.

CLEARVECTOR