BLOODHOUND ENTERPRISE

OVERVIEW

Defend AWS Environments with BloodHound Enterprise

Adversaries reach AWS through identities that already have a way in. A directory account synced into the cloud, an over-trusted federated role, a token left on a machine one SaaS platform away: each artifact grants access that looks legitimate to the implemented controls built to catch it. Following those trust relationships across account boundaries, an adversary reaches critical resources without a single failed login.

The Problem

Most security tools that monitor AWS are evaluating one account at a time. They inventory misconfigurations and flag toxic combinations of permissions, then leave teams to guess which ones an adversary could reach and chain together. A trust policy that looks acceptable on its own can hand a low-privileged identity a path to administrative control three accounts away, and a permission audit reports all three as compliant.

Defenders think in lists. Attackers think in graphs. As long as this is true, attackers win. As identity boundaries and access needs expand, this adage proves truer every year. Attackers follow role assumption, cross-account trust, and federated login to wherever those grants ultimately reach. The unknown exposure that matters is the reachable path across accounts and identity providers, yet it stays invisible to any review that stops at the account boundary.

The Solution

BloodHound Enterprise gives defenders a map of what adversaries see, the chained relationships that combine into exploitable paths to critical assets. It models on-premises Active Directory, Entra ID, and the cloud and SaaS platforms that power your business as one graph, so a path that starts with a synced identity and ends in an AWS production account stays visible the whole way. OpenGraph is the engine that brings AWS into the attack graph.

For AWS, BloodHound Enterprise evaluates authorization the way AWS evaluates it. Rather than pattern-matching permissions into toxic combinations, it reasons across identity policies, resource policies, trust policies, permission boundaries, service control policies, and resource control policies to determine who could perform an action, then traces where that access leads across accounts. The AWS OpenGraph Extension models more than 20 AWS resource types and over 150 edges spanning IAM, AWS Organizations, STS, S3, KMS, Lambda, EC2, CloudFormation, EKS, Systems Manager, and more.

Taken together, that identity security coverage answers the question no configuration audit can:
Given everything in place across your environment, where can an adversary go from any particular point in my identity and application network?

Cross-Account Reachability

Resolves role trust and assumption across account boundaries to show where a foothold in one account reaches administrative control in another. It also identifies inbound trust from external AWS accounts, vendors, and contractors to expose third-party attack paths into your environment.

Hybrid and Federated Identity

Maps federation, web identity, OIDC, and directory relationships that carry identity into AWS, revealing how compromised Active Directory, Entra ID, Okta, GitHub, or workload identities ultimately reach cloud roles.

Privilege Escalation

Identifies hidden administrative paths created through actions such as policy attachment, PassRole abuse, trust policy modification, and access key creation, exposing how low-privileged identities can ultimately reach administrator access.

Credentials and Secrets

Shows who can reach sensitive resources including S3 buckets, KMS keys, Lambda functions, Systems Manager parameters, EC2 workloads, and the credentials and secrets they contain. Rather than identifying only direct access, BloodHound Enterprise reveals who can ultimately reach those resources through role chaining, privilege escalation, and cross-account trust.

Ready to see your AWS environment the way adversaries do?

Request a demo of BloodHound Enterprise:

Visit specterops.io/get-a-demo ›

What This Looks Like in Practice

More than 700 organizations had data stolen in the 2025 Salesloft Drift breach, and the path to each of them ran through AWS. The threat actor Google tracks as UNC6395 first gained access to Salesloft's GitHub environment, then reached the AWS environment where the Drift integration stored OAuth tokens for its customers' platforms. Those tokens unlocked hundreds of downstream Salesforce tenants, which the adversary exported and mined for further credentials, including plaintext AWS access keys staged for the next pivot. One GitHub foothold reached AWS, and AWS held the trust that reached everyone else.

BloodHound Enterprise models the same trust relationships the Salesloft adversary walked: how an account in one platform reaches an AWS role, and where that role reaches next. The path that took the adversary months to walk quietly is the kind it surfaces as a single traversable chain.

AWS Is One Piece of a Larger Attack Surface

The paths that reach AWS rarely start there and can often be the jumping off point to broader breaches. One may begin on a Jamf-managed Mac, move through Okta, and pass through GitHub before it ever assumes an AWS role. From there, an attacker can leverage trusted relationships to reach production workloads, sensitive data, CI/CD pipelines, and other connected cloud and enterprise resources. Each step looks legitimate to the platform that owns it, and the full attack path appears only when they're exposed in one graph.

Previous BloodHound Enterprise OpenGraph extensions for Okta, GitHub, and Jamf bring the same attack path analysis to identity provider configurations, source code and CI/CD pipelines, and macOS device management. Together with AWS, they give defenders visibility into the cross-platform paths that form between systems, across administrative boundaries, and outside the reach of any single platform review.