The VAP | Very Attacked People

Dashboard Starter Kit

Companion Guide to ***How to Find Your Very Attacked People and Protect Them***



graylog

What Is a VAP, Again?

A VAP, or Very Attacked Person, is an employee whose access, not their title, makes them a high-value target.

A marketing director who signs off on vendor contracts. An HR manager holding sensitive employee data. A facilities coordinator with building access control privileges. None of them show up on a typical risk register, and all of them are exactly what an attacker is looking for.

The 2025 Verizon Data Breach Investigations Report (DBIR) tells us that 74% of breaches involve the human element, including social engineering, phishing, and errors. The days of indiscriminate phishing blasts are over. Attackers are narrowing their sights on specific individuals. They're tailoring their tactics. And they're succeeding.

74% of breaches involve the human element.

Spear-phishing aimed at a specific person is 60% more effective than a generic attempt.

Source: 2025 Verizon Data Breach Investigations Report

Attackers already know precision beats volume. This guide picks up where the VAP eBook left off and walks through how to actually build what it describes: a risk-scored, people-focused dashboard inside Graylog Security, using identification, scoring, and automation the platform already has.

What Needs to Be Flowing Into Graylog First?

Five data types need to be flowing into Graylog before you build anything: identity and authentication logs, email security logs, endpoint and EDR logs, vulnerability scan data, and user and group activity logs.



Identity and authentication logs: VPN, Azure AD, Okta, EntraID, or on-prem AD. This is where account takeover and unusual login patterns show up first.



Email security logs: Targeted phishing attempts start here.



Endpoint / EDR logs: Graylog's existing integrations with CrowdStrike, Microsoft Defender, Bitdefender, and Carbon Black cover this.



Vulnerability scan data: Tenable Nessus or Microsoft Defender feeds let Graylog factor real exploitability into a person's risk score, not just activity.



User and group activity logs: Active Directory, Azure AD / Entra ID, or file server and SharePoint audit logs. This is where privilege escalation and data staging show up: a user added to an admin group, an unusual volume of files accessed in a short window, or permissions changed on a sensitive folder.



Skip one of these and the risk score in Step 1 will have a blind spot. A dashboard built only on login logs misses the phishing angle. One built only on email security misses account takeover. One built only on endpoint data misses the privilege escalation that happens entirely inside the directory, no malware required. Correlate across sources, don't lean on one signal.

Step 1: How Do You Tag and Score Your VAPs?

Tag VAP-worthy users as an asset category, then let Graylog Security's risk scoring do the ongoing work.

Start where the eBook does: create a VAP asset category in Graylog's Assets section, then assign it to the specific users who carry outsized risk, regardless of title.

What's an Asset, Exactly?

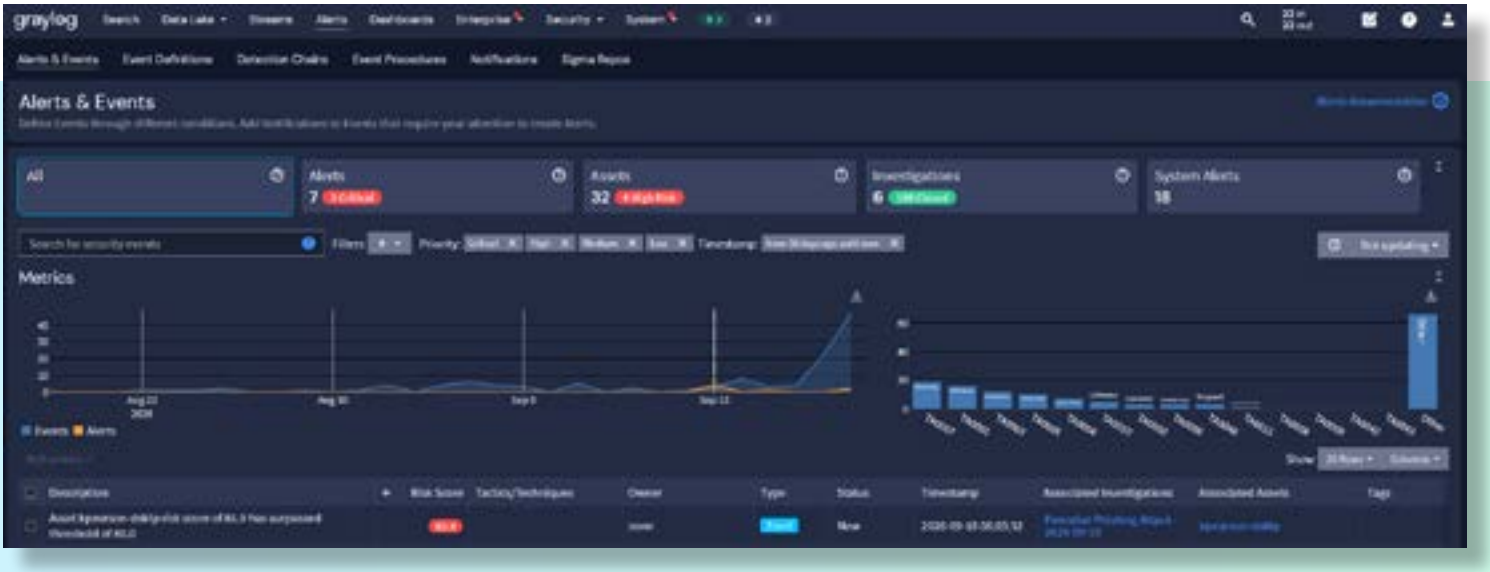
An asset isn't a login or a device. It's an identity construct: the mechanism that lets Graylog recognize the same user or system even when different data sources describe them differently. A VPN log might call someone "jsmith." An email security log might call the same person "john.smith@company.com." A Windows event log might call them "COMPANY\JSmith." Graylog's asset enrichment resolves all three into one asset, so a VAP tag or a risk score follows that person everywhere they show up, not just in the source where you applied it.

Build your first list

Cross-reference:

- Who has admin, finance, HR, or facilities-level access
- Who's been flagged by email security or endpoint tools in the past 90 days
- Whose credentials have shown up in a threat intel or dark web feed
- Who signs off on vendor payments, wire transfers, or contracts

That's your first VAP list. Assign the category. This step is manual, and that's fine. It's a starting point, not the finish line.

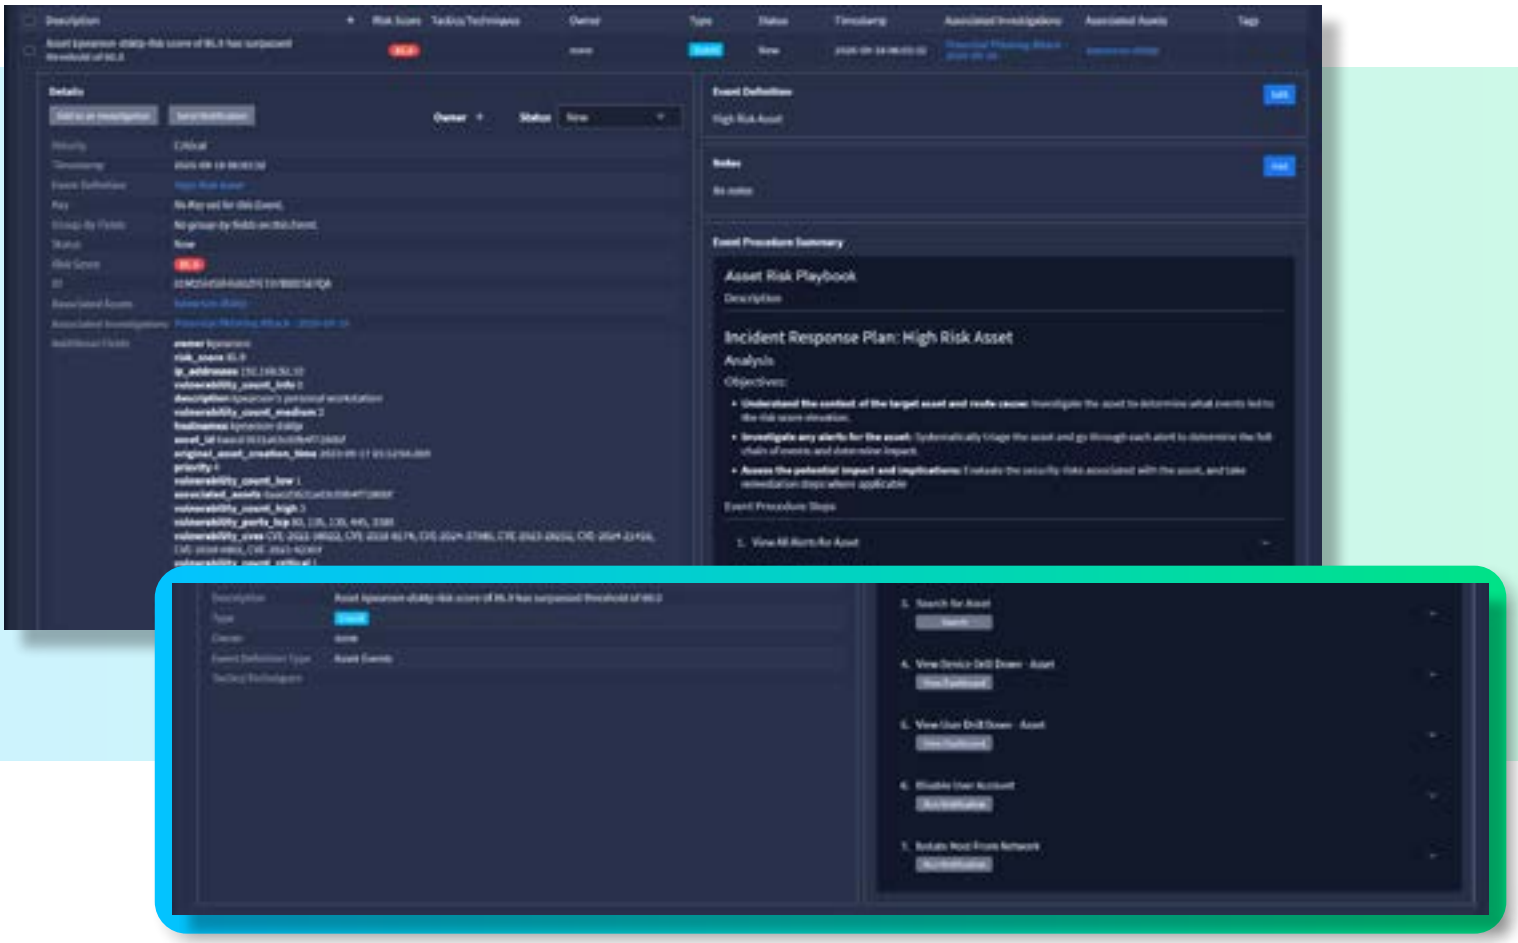


Let risk scoring take over from there

Two Graylog Security mechanisms do the heavy lifting once the category is assigned.

Asset-Based Risk Scoring assigns every asset, including VAP-tagged users, a running score based on security posture. For a person, that includes any linked vulnerability or exposure data.

Events & Alerts Risk Scoring calculates a score for every event using log severity, event priority, and asset priority together. Set a VAP-tagged user's asset priority high, and an event that would otherwise sit in a queue of low alerts now scores high the moment it touches that person. This is the mechanic that solves the eBook's core problem: a "low" alert tied to a high-value person isn't actually low.




```
category:VAP AND risk_score:>=70
```

Illustrative syntax only. Confirm exact field names against your own Graylog schema before use.

Risk-based scoring can turn 10,000 alerts into a shortlist of 10 that actually matter. Teams using it report cutting noise by up to 99.9% and responding up to 60% faster.

Vulnerability Scan Ingestion pulls in data from Tenable Nessus or Microsoft Defender. That data only populates on machine assets, never on the person. A VAP's own score doesn't move because their device is vulnerable. The device's score does. Check both: a VAP working from an unpatched, exploitable machine is a different risk than one who isn't, even before any attack activity shows up.

Step 2: How Do You Build the Dashboard?

Turn the Step 1 search into a dashboard with four widgets: geolocation, a single-value metric, a time series graph, and a heat map.

With the category assigned and risk scoring active, the dashboard is mostly a matter of picking the right widgets. In Graylog, dashboards are just saved searches that run on a schedule, so once the Step 1 search works, turning it into a dashboard is a right-click away.



- **Geolocation:** Map where login and access attempts for your VAPs originate, so an unusual location reads as risk, not routine.
- **Single-value metric:** Current count of active VAPs above your alert threshold. A number that changes is a number people check.
- **Time series graph:** Attack activity by individual over a rolling 30-day window, the exact gap the eBook calls out in endpoint-only dashboards.
- **Heat map:** Spot which VAPs cluster around specific anomaly types, login versus data movement versus file integrity.



Make it parameterized. A click on any VAP in the ranked table should drill straight into that person's raw logs. Search and dashboards share the same engine in Graylog, so this is a query parameter passed from the widget, not a separate build.

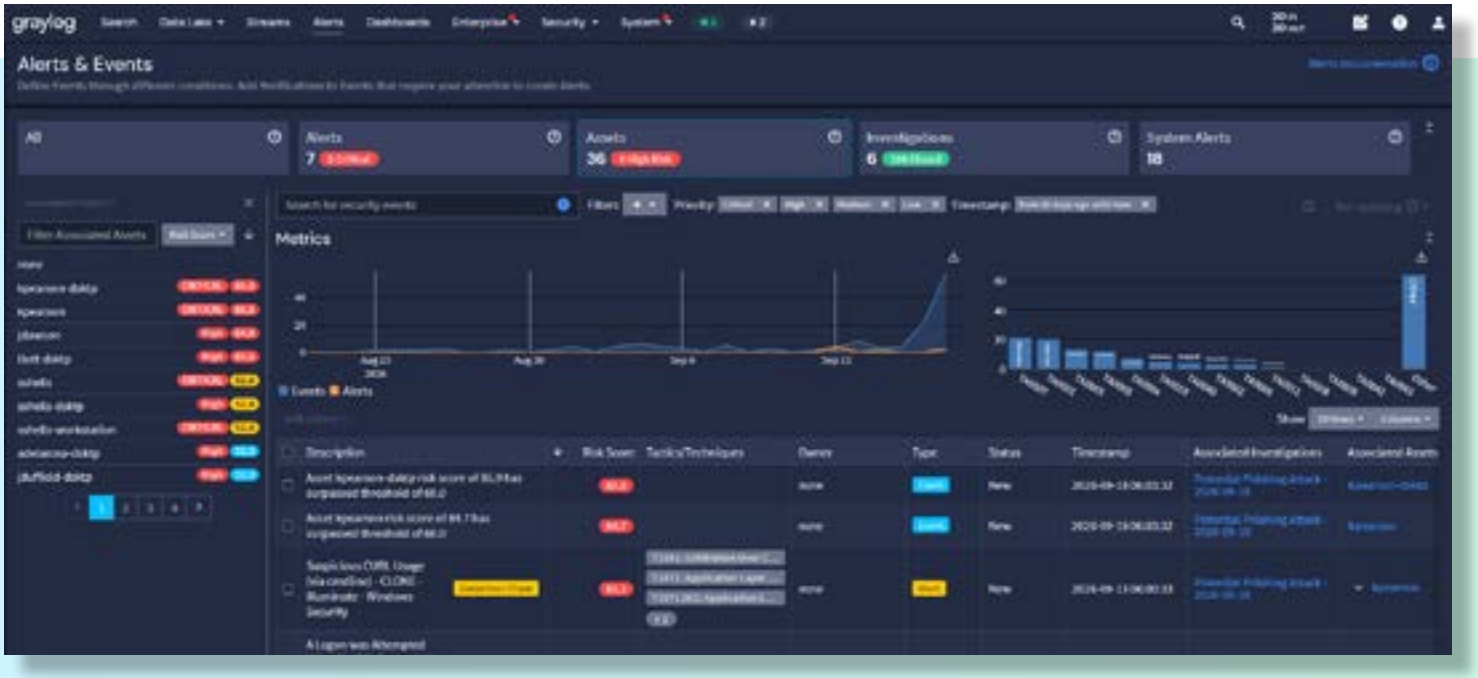
Step 3: How Do You Automate the Upkeep?

The Verascend™ Incident System keeps VAP alerts triaged automatically, correlating entity behavior, asset risk, and live threat campaigns instead of relying on manual review.

A VAP list built once and never revisited goes stale within a month. People change roles, leave the list, or should have been on it and weren't. Alert volume doesn't wait for that review to happen. This is where automation stops being a buzzword and starts doing actual work.
Verascend is Graylog's agentic triage layer, built natively into Graylog Security rather than a separate environment you have to stand up and secure yourself. It turns the flood of log lines your VAPs generate into a ranked queue of verified threats, tuning out false positives automatically as it goes. Analysts investigate. They do not filter.

For a VAP dashboard specifically, three things carry most of the weight:

- Asset risk correlation, so an alert tied to your CFO or your HR admin's account outranks a hundred lower-priority tickets automatically
- Entity and campaign correlation, so a known attack pattern hitting a VAP gets flagged as part of a campaign, not investigated as an isolated event
- Automatic false-positive tuning, so a VAP's real signal doesn't get buried under alerts that don't matter



The payoff: Alert priority for your VAPs stays current without someone manually re-triaging the queue every shift. The highest-risk, most correlated cases surface first, automatically, with the noise already tuned out. That's the difference between a VAP program that's a quarterly checklist and one that's actually operational.

What Happens When Graylog Flags a VAP?

Four steps: confirm the trigger, check correlated signals, escalate by role, and feed the outcome back into the list.

- **Confirm the trigger:** Open the parameterized view and check what changed: a login anomaly, a data movement spike, a new device.
- **Check correlated signals:** One unusual login is worth a look. That login plus an outbound data spike plus a new device is worth an immediate response.
- **Escalate by role, not raw severity:** A “low” alert tied to your CFO or HR admin is not a low-priority action item.
- **Feed the outcome back into the list:** A confirmed attempt likely keeps that person elevated for a defined review period, not just until the ticket closes.



Quick Answers

Q: What is a VAP?
A: A Very Attacked Person: an employee whose access, not title, makes them a high-value target.

Q: Do I need a dedicated analyst to run this?
A: No. Tagging is manual and quick. Graylog Security's risk scoring and UEBA anomaly detection handle the ongoing analysis.

Q: How is VAP scoring different from a standard risk score?
A: Standard models score endpoints and vulnerabilities. VAP scoring adds the person as a scored entity, tied to real attack activity.

Q: How long before anomaly detection is useful?
A: About a week of historical data before the behavioral baseline stabilizes.

Q: Does this replace existing SIEM alerts?
A: No. It re-prioritizes them. Every alert still fires; VAP tagging determines which ones jump the queue.

Quick Reference

Input	Source
Identity / auth logs	Azure AD, Okta, on-prem AD, VPN
Email security logs	Existing email security gateway
Endpoint / EDR logs	CrowdStrike, Defender, Bitdefender, Carbon Black
Vulnerability data	Tenable Nessus, Microsoft Defender
Risk score inputs	Log severity + event priority + asset priority
Recommended widgets	Ranked table, single-value metric, time series, heat map
Automation layer	UEBA Anomaly Detection, ~7 days to train a baseline

Ready to Build This in Your Environment?

A VAP program doesn't need a dedicated analyst or a six-figure add-on. It needs an asset category, a risk score that already accounts for who's tagged, a dashboard built from a saved search, and an anomaly detection layer that keeps it current without manual upkeep. All of it is already inside Graylog Security.

Get a personalized walkthrough of Graylog Security: tell Arti what you're working on, see the features that protect your VAPs.

©2026 Graylog, Inc. All rights reserved.