

Is Your Security Data Architecture Costing More Than It Should?

A 5-minute self-assessment for security leaders

Most security data architectures weren't designed. They accumulated. New tools, new sources, new compliance requirements — each solved a specific problem but added cost, complexity, and fragmentation to the whole. **Check each statement that applies to your organization, then total your score at the end.**

1 Ingestion & Storage Costs

Are you paying premium rates for data you rarely use?

- More than 30% of your SIEM-ingested data is searched less than once per month
- You've excluded data sources from your SIEM specifically because of ingestion cost
- Your SIEM renewal cost has grown faster than your actual data needs
- Retention periods are driven by licensing pressure, not investigative or compliance requirements

2 Analyst Efficiency

Are your analysts investigating — or assembling data?

- Analysts routinely pivot across 4+ consoles to complete a single investigation
- Senior analysts are required for routine investigations because only they know where data lives
- Data sources exist that analysts avoid because they're too slow or difficult to access
- Copy-paste across tools is a standard part of the investigation workflow

3 Coverage & Visibility

Can you actually see what you need to see?

- High-value telemetry (EDR, identity, cloud) exists but isn't searchable from your primary analysis platform
- Detection coverage is limited by what you can afford to ingest, not what's available
- Historical data beyond 30–90 days is effectively inaccessible during active investigations
- You've had an incident where the data you needed existed somewhere — but wasn't reachable in time

4 Architecture Complexity

Is your architecture compounding — or creating drag?

- Onboarding a new data source requires a pipeline project measured in weeks or months
- You maintain multiple schemas or enrichment logic for the same entities across tools
- Your data architecture depends on 1–2 people's institutional knowledge to operate
- A SIEM migration or major tool change would require rebuilding most of your data workflows

5 AI & Automation Readiness

Is your data foundation ready for what comes next?

- You've tested AI security tools and found results inconsistent or difficult to trust
- Your data lacks a common schema or normalization layer across sources
- AI or automation tools can only access data inside your SIEM — not your full security data footprint
- There's no way to verify the evidence an AI tool used to reach its conclusion

SCORE	WHAT IT SUGGESTS
0–4	Your architecture is reasonably aligned. Look for targeted optimizations.
5–10	Cost and capability gaps are compounding. Targeted architectural changes would have measurable impact on both budget and outcomes.
11–15	Your data architecture is actively constraining your security program. Incremental tool-level fixes are unlikely to close the gap.
16–20	You're paying more and seeing less. It may be time to rethink how security data is accessed, stored, and operated as a system.

What High-Performing Teams Do Differently

- **Decouple storage from analytics.** Data lives in cost-effective storage and is queried where it sits — not duplicated into expensive platforms by default.
- **Normalize before they analyze.** A common schema across sources means detections, investigations, and AI workflows all operate on consistent inputs.
- **Design for investigations, not just alerts.** Data access patterns support the way analysts actually work: pivoting across sources, time ranges, and entities.
- **Treat security data as a product.** Assign explicit owners, define consumers, and set clear expectations across the full data lifecycle — from collection through enrichment, access, and retirement.

This assessment is based on patterns observed across enterprise security programs navigating SIEM renewals, cloud migrations, data lake initiatives, and AI-driven operations. For a deeper framework, see [Query's Security Data Operations Assessment Guide](#).