snare
A PROPHECY SOLUTION

THE SIEM INDEPENDENCE CHECKLIST

How locked in is your security-data architecture? Ten domains, fifty tests, one score and a prioritised path to fixing what the score exposes.

WHO THIS IS FOR
CISOs, security architects, SOC leaders and MSSPs approaching a SIEM renewal, evaluating a platform change, planning a consolidation programme, or trying to work out why ingestion costs keep rising without visibility improving.

How to Use This Checklist

Vendor lock-in is rarely a decision anybody made. It accumulates, one agent rollout, one parser, one detection library, one archive format at a time, until the question "should we change our SIEM?" has quietly been replaced by "could we?"

This worksheet is designed to make that accumulation visible and measurable before a renewal conversation rather than during a migration. It covers ten domains. Each domain contains five test statements, a description of what a strong answer looks like, and the specific evidence to go and gather.

Scoring

Response	Score	Use when
Yes	2	The statement is true today, and you could demonstrate it, not merely assert it
Partial	1	True for some sources, some environments or some tenants, but not consistently
No	0	Not true today, or nobody in the organisation can confirm either way

Each domain scores out of 10. The full worksheet scores out of 100. Transfer your domain totals to the summary table near the end.

ONE RULE THAT MAKES THE RESULT USEFUL

Score "Yes" only where you could produce the evidence within a working day. "We think so" and "the vendor says so" are both "Partial".
The value of this exercise is not the number. It is the list of things nobody could answer, because that list is an accurate map of where your dependency actually lives.

RUNNING IT AS A WORKSHOP

The checklist works best with the security architect, the SOC lead, the detection engineering lead and someone from procurement in the same room for ninety minutes. Procurement matters more than people expect, three of the ten domains turn on contract language rather than technology.

For MSSPs, run it once per tenant archetype rather than once globally. The answers usually diverge sharply between customers you onboarded onto your own stack and customers who arrived with their own.

1.Collection Independence

The collection layer is the part of a logging architecture that touches the most machines and takes the longest to change. If it is owned by the analytics platform, every other decision downstream inherits that dependency, including the decision to leave.

	Test statement	Yes (2)	Partial (1)	No (0)
1	Our log collection agents are supplied and controlled independently of our SIEM vendor, or are standards-based.	□	□	□
2	We can change what an agent collects centrally, without touching each host individually.	□	□	□
3	Agent configuration and collection policy are held and version-controlled in our own records, not only exported from a vendor console.	□	□	□
4	Our collection policy is defined by investigative requirement, not by what the SIEM licence makes affordable.	□	□	□
5	We could inventory every host running a collection agent, and confirm it is currently reporting, within one working day.	□	□	□

WHAT A STRONG ANSWER LOOKS LIKE
Collection policy is a document you own, applied by a layer you control, and it would read the same regardless of which analytics platform sits downstream.

Evidence to gather: Agent inventory and version report; collection policy document; the change process for modifying what is collected.

Domain 1 score: _____ **/10**

2.Agent Dependency

This is the question that turns a SIEM migration into an endpoint project. A destination change should be a routing decision. If it requires software to be redeployed across thousands of hosts, the analytics platform has been given a veto over its own replacement.

	Test statement	Yes (2)	Partial (1)	No (0)
1	Changing our SIEM destination is a configuration change, not a software redeployment.	□	□	□
2	We can add a second destination without an agent upgrade or a restart across the estate.	□	□	□
3	We know the exact host count, who would perform a redeployment, and what change window it would require.	□	□	□
4	No collection agent in our estate stops functioning when the SIEM contract ends.	□	□	□
5	We have pointed at least one production collection group at an alternate destination within the last 12 months.	□	□	□

WHAT A STRONG ANSWER LOOKS LIKE
Adding a destination is a config push measured in minutes, tested regularly, and requires no coordination with endpoint or desktop teams.

Evidence to gather: The documented procedure for adding a destination; the date of the last time it was exercised; CMDB host counts by agent type.

Domain 2 score: _____ **/10**

3.Log Ownership

Ownership is a contractual property before it is a technical one. Many organisations discover only at termination that access to their own historical evidence was a function of the licence rather than a right, and that the export format offered is not one another platform can consume.

	Test statement	Yes (2)	Partial (1)	No (0)
1	Our contract explicitly grants us access to our own raw log data throughout the term.	☐	☐	☐
2	Our contract specifies data access, export format and timeframe after termination.	☐	☐	☐
3	We hold a copy of our security evidence outside the SIEM vendor's environment.	☐	☐	☐
4	We know which jurisdiction our retained evidence sits in, and we control that placement.	☐	☐	☐
5	Where a third party collects on our behalf, an MSSP, or a SaaS provider exposing audit APIs, we receive or can retrieve our own copy.	☐	☐	☐

WHAT A STRONG ANSWER LOOKS LIKE
Someone in the security team can name the contract clause, the export format and the post-termination window without asking procurement.

Evidence to gather: The relevant contract clauses; a sample export in the specified format; the data-residency statement for retained evidence.

Domain 3 score: _____ **/10**

4.Historical Retention

IBM puts the mean time to identify and contain a breach at 247 days. Thomson Reuters discovered on 30 June 2026 activity that occurred in March. Retention set to the alerting window rather than the detection lag guarantees that the most important investigations start blind.

	Test statement	Yes (2)	Partial (1)	No (0)
1	We can state, per source, how far back our evidence actually reaches, and it matches our stated retention policy.	□	□	□
2	Our retention horizon for authentication, access and egress evidence exceeds the current 247-day detection average.	□	□	□
3	Retained evidence is searchable within a working incident timeline, not restore-only.	□	□	□
4	Retention is set per source and per data class, not as a single platform-wide setting.	□	□	□
5	We have run a query against data older than six months within the last 12 months, and we know how long it took.	□	□	□

WHAT A STRONG ANSWER LOOKS LIKE
Retention is deliberately differentiated by data class, documented against both the investigation clock and the regulatory clock, and periodically tested by querying old data.

Evidence to gather: Retention settings by source; a timed query result against data older than six months; the mapping of retention periods to notification obligations.

Domain 4 score: _____ **/10**

5.Routing Portability

Routing is where cost control and independence meet. A layer that can send one stream to several destinations, each with its own treatment, lets you optimise ingestion spend and run a migration in parallel, using exactly the same mechanism.

	Test statement	Yes (2)	Partial (1)	No (0)
1	We can send one collected event stream to more than one destination simultaneously.	□	□	□
2	Each destination can have its own filtering, transformation and output format.	□	□	□
3	Destinations have independent delivery queues, so a slow or unavailable destination does not back-pressure the others.	□	□	□
4	We can add or remove a destination without any change at the endpoint.	□	□	□
5	Routing rules live in a layer we control, not inside the analytics platform.	□	□	□

WHAT A STRONG ANSWER LOOKS LIKE

Routing is a first-class, documented capability with per-destination policy, and a second destination has been run in production, not just in a lab.

Evidence to gather: The routing configuration; evidence of a second destination in use; queue and backpressure behaviour under a destination outage.

Domain 5 score: _____ **/10**

6.Schema and Format Portability

Normalisation is optimised for correlation, not forensics. Fields that do not map to a schema class get dropped or flattened, and in a real investigation the decisive detail is often one of the awkward vendor-specific fields no common schema anticipated. If the only surviving copy is the parsed one, the evidence has already been edited, invisibly.

	Test statement	Yes (2)	Partial (1)	No (0)
1	We retain the raw event alongside any normalised version of it.	□	□	□
2	We can export evidence in an open format, Syslog RFC 5424, JSON or CEF, that another platform can parse without bespoke work.	□	□	□
3	Normalisation is performed in a layer we control, not only at SIEM ingest.	□	□	□
4	We have assessed our major sources against an open schema such as OCSF.	□	□	□
5	We could hand twelve months of evidence to external forensic counsel without providing a licence to our SIEM.	□	□	□

WHAT A STRONG ANSWER LOOKS LIKE
Raw and normalised are both retained and both retrievable, and the export path to open formats has been exercised at least once against real data.

Evidence to gather: A raw and a normalised copy of the same event; a sample export in an open format; any OCSF mapping work completed to date.

Domain 6 score: _____ **/10**

7.Multi-SIEM Capability

Parallel operation is what makes a competitive evaluation real. Without it, a proof of concept runs on synthetic or partial data, the comparison is unconvincing, and the renewal defaults to the incumbent, not because it won, but because a fair contest was never possible.

	Test statement	Yes (2)	Partial (1)	No (0)
1	We could run two SIEMs in parallel on the same live data for a validation period.	☐	☐	☐
2	We could route different classes of data to different analytics platforms according to purpose.	☐	☐	☐
3	We could run different SIEMs in different regions or business units without duplicating collection.	☐	☐	☐
4	We have modelled the ingestion cost of a parallel-run period.	☐	☐	☐
5	A competitive proof of concept could be fed live production data without touching a single endpoint.	☐	☐	☐

WHAT A STRONG ANSWER LOOKS LIKE

A parallel run is a planned, costed and technically routine activity, something the architecture supports by design rather than something that would need a project.

Evidence to gather: The parallel-run cost model; the routing change required to enable it; any prior evaluation fed with live data.

Domain 7 score: _____ /10

8.Migration Readiness

The hardest part of most SIEM migrations is not the new platform. It is history, proving the new system sees what the old one saw, and keeping years of evidence usable after the old licence lapses.

	Test statement	Yes (2)	Partial (1)	No (0)
1	We can replay historical evidence into a new platform without re-collecting it from source.	□	□	□
2	We have a documented cutover sequence that does not require an endpoint change window.	□	□	□
3	We know, by count, how many detections, dashboards and integrations would need rebuilding.	□	□	□
4	We have a validation method to prove a new platform sees what the current one sees.	□	□	□
5	Decommissioning the current platform would not remove our access to its historical data.	□	□	□

WHAT A STRONG ANSWER LOOKS LIKE

Replay is a supported capability that has been tested; the rebuild inventory is a real number someone has counted, not an estimate.

Evidence to gather: Replay test results; the detection, dashboard and integration inventory; the documented cutover and validation plan.

Domain 8 score: _____ **/10**

9.Ingestion-Cost Control

Cost programmes that only reduce what is collected are visibility-reduction programmes wearing a finance label. Filtering, transformation and routing are three different operations, and only separating them makes cost reduction compatible with better investigation coverage.

	Test statement	Yes (2)	Partial (1)	No (0)
1	We filter, transform and de-duplicate upstream of the point at which billing is calculated.	□	□	□
2	We know our current duplication rate, the same event arriving via more than one collector.	□	□	□
3	We can measure volume by source and by destination independently of the SIEM's own reporting.	□	□	□
4	We have classified our top twenty sources by detection value and investigative value separately.	□	□	□
5	Reducing SIEM cost does not require reducing what we collect.	□	□	□

WHAT A STRONG ANSWER LOOKS LIKE
Volume is measured on your side of the meter, duplication is a known number, and the levers available include transformation and routing, not only dropping sources.

Evidence to gather: Volume by source from your own instrumentation; the measured duplication rate; the source classification by detection versus investigative value.

Domain 9 score: _____ **/10**

10.AI Investigation Portability

AI is the largest step-change in investigation productivity the SOC has seen in a decade, and it should be adopted aggressively. The question is only where the capability sits. CrowdStrike observed AI agent-triggered detection leads growing at 2.5 times the rate of human-triggered leads, an investigation layer that ends with a contract is a capability you will have to rebuild at the worst possible moment.

	Test statement	Yes (2)	Partial (1)	No (0)
1	Our natural-language or AI investigation capability queries data we control, not only the analytics platform's own store.	☐	☐	☐
2	Investigation history, saved questions and established methodology would survive a platform change.	☐	☐	☐
3	Analysts could continue investigating during a migration cutover.	☐	☐	☐
4	Our AI investigation capability is not priced or licensed as an inseparable part of the SIEM contract.	☐	☐	☐
5	We could evaluate an alternative AI investigation layer against our own retained evidence.	☐	☐	☐

WHAT A STRONG ANSWER LOOKS LIKE

The investigation layer is scoped to your evidence rather than to a platform, so adopting, changing or adding AI capability is a choice rather than a migration consequence.

Evidence to gather: Where the AI capability reads from; the licensing basis; whether investigation history is exportable.

Domain 10 score: _____ **/10**

Your Score

Transfer each domain total below.

Domain	Score	Priority to address
1. Collection Independence	_____ / 10	□ High □ Medium □ Low
2. Agent Dependency	_____ / 10	□ High □ Medium □ Low
3. Log Ownership	_____ / 10	□ High □ Medium □ Low
4. Historical Retention	_____ / 10	□ High □ Medium □ Low
5. Routing Portability	_____ / 10	□ High □ Medium □ Low
6. Schema and Format Portability	_____ / 10	□ High □ Medium □ Low
7. Multi-SIEM Capability	_____ / 10	□ High □ Medium □ Low
8. Migration Readiness	_____ / 10	□ High □ Medium □ Low
9. Ingestion-Cost Control	_____ / 10	□ High □ Medium □ Low
10. AI Investigation Portability	_____ / 10	□ High □ Medium □ Low
TOTAL	**_____ / 100**	

Interpreting Your Score

Band	Position	What it means in a renewal conversation
80–100	Independent	Your evidence layer is genuinely portable. The SIEM is a destination, not the architecture. You can run a real competitive evaluation, and your renewal is negotiated on merit.
55–79	Partially independent	Exit is feasible but costly, and the cost is concentrated in one or two domains. Fix those and you move a whole band, usually with less effort than the score suggests.
30–54	Coupled	The analytics platform is materially shaping your data architecture. A migration would be a rebuild, and both you and the incumbent know it. Address collection and routing first.
0–29	Dependent	The SIEM is your data architecture. Renewal is not really a decision. Treat this as an architecture programme rather than a procurement problem, and start with getting collection out from under the platform.

A NOTE ON LOW SCORES

A low score is not a failure of the security team. It is the predictable outcome of an architecture that was assembled one sensible decision at a time, under cost and delivery pressure, over several years. Almost every organisation that has run a SIEM for more than five years scores lower than it expects on first pass.

What matters is that the dependency is now visible, sized, and attached to specific domains you can sequence.

Where to Start

The ten domains are not equally urgent, and they are not independent. Collection and routing sit upstream of everything else, fixing them raises the ceiling on six of the remaining eight. This is the sequence that recovers the most optionality for the least effort.

Phase 1, Recover the collection layer (Domains 1, 2)

- Establish collection that is owned and configured independently of the analytics platform.
- Prove that a destination change is a configuration push, not a redeployment, and record how long it took.
- Bring collection policy under your own version control, so it survives any platform change.

This is the phase that removes the endpoint estate from the migration conversation entirely. Everything after it becomes cheaper.

Phase 2, Take control of routing and cost (Domains 5, 9)

- Introduce a routing layer capable of multiple simultaneous destinations with independent delivery queues.
- Move filtering, transformation and de-duplication upstream of the billing point.
- Measure volume and duplication on your side of the meter, so the numbers in the next negotiation are yours.

Phase 2 typically pays for Phase 1. It is also what makes a parallel run technically routine rather than a project.

Phase 3, Secure the evidence itself (Domains 3, 4, 6)

- Fix the contract language on data access, export format and post-termination availability at the next renewal, not after it.
- Set retention per data class against both the investigation clock (247 days on current averages) and the regulatory clock.
- Retain the raw event alongside the normalised one, and test an open-format export against real data.

Phase 4, Prove it (Domains 7, 8, 10)

- Run a live parallel evaluation against a second platform, using production data and no endpoint change.
- est replay of historical evidence into that second platform and time it.
- Confirm your investigation and AI capability reads from evidence you control and would survive the cutover.

At the end of Phase 4 the answer to “if we changed our SIEM tomorrow, what else would we have to change?” is a short, specific and unsurprising list. That is the whole objective.

Appendix: The Cost of Exit Worksheet

The checklist tells you where dependency sits. This worksheet turns it into a number, which is the version of the argument that travels to a CFO and a board.
Estimate each line for your own environment. The comparison that matters is the total against your annual licence cost.

Exit cost component	Unit	Quantity	Effort / cost
Agent redeployment across the endpoint estate	Hosts	_____	_____
Parser and normalisation rebuild	Source types	_____	_____
Detection content rewrite, including validation and tuning	Detections	_____	_____
Dashboard and report rebuild	Dashboards	_____	_____
Historical data conversion, or dual-licence overlap period	Months / TB	_____	_____
Integration rework, ticketing, SOAR, enrichment, case management	Integrations	_____	_____
Analyst retraining and productivity dip during transition	Analysts / weeks	_____	_____
Parallel-run ingestion cost during validation	Months	_____	_____
TOTAL COST OF EXIT			**_____**
ANNUAL LICENCE COST			**_____**

READING THE RESULT
If total cost of exit exceeds the annual licence, the platform is no longer being renewed on merit. It is being renewed because leaving is harder than staying.

That is a legitimate commercial position, but it is a very different one from "this is the best platform for us", and it is worth knowing which conversation you are actually having before you sit down at the table.

Next Steps

Every domain in this checklist depends on the same underlying condition: that the layer collecting, retaining and routing your security evidence is one you control, and is separable from the platform analysing it.
That is the role Snare is built for.

SNARE AGENT	Collect detailed security and audit events close to the source across Windows, Linux, Unix, macOS and other environments, independently of any downstream analytics platform.
SNARE CENTRAL	Centrally manage collection policy, aggregation, de-duplication, retention and replay, so history remains usable when platforms change.
SNARE REFLECTOR	Filter, transform and route events to multiple SIEMs, data lakes, archives and analytics platforms simultaneously, with independent delivery queues and destination-specific policy.
ASKSNARE	Give analysts a natural-language investigation layer that queries Snare-managed evidence directly, independently of whichever SIEM is downstream.

PUT THE CHECKLIST INTO PRACTICE
Bring your completed scoresheet to a working session. We will walk through the domains you scored lowest, and what changing them would actually involve in your environment.
→ Request a Snare and AskSnare demonstration
→ Read Snare Insider Issue 14
→ Read the Ransomware Investigation Playbook (Issue 13)

References

IBM, Cost of a Data Breach Report 2026 (IBM and Ponemon Institute, 29 July 2026)
Mean time to identify and contain of 247 days, reversing five consecutive years of improvement; global average breach cost US$4.99M, up 12%; detection and escalation plus lost business accounting for 63% of total cost.

Verizon, 2026 Data Breach Investigations Report (19th edition, May 2026)
Vulnerability exploitation at 31% of breaches; third-party involvement in 48%, up 60% year on year; median time to fully remediate a CISA KEV entry of 43 days.

CrowdStrike, 2026 Threat Hunting Report (3 August 2026)
88% of PoC-linked exploitation occurring within 48 hours of release; AI agent-triggered detection leads growing at 2.5x the rate of human-triggered leads.

Thomson Reuters / West Publishing, C-Track cybersecurity incident (disclosed 2–3 September 2026)
Unauthorised access to C-Track files in March 2026, discovered 30 June 2026, affecting court systems in 11 US states, the US Virgin Islands and Ontario, Canada.

Open Cybersecurity Schema Framework (OCSF)
Vendor-neutral security-event schema. Launched 2022; joined the Linux Foundation November 2024; version 1.4.0 released January 2025; used natively by AWS Security Lake.

ASD Information Security Manual, event logging and retention
ISM-0859 (seven-year retention) rescinded December 2024 and replaced by ISM-1988, requiring event logs to be retained in a searchable manner for at least 12 months. ISM-1815 (logs protected from unauthorised modification and deletion) maps to Essential Eight Maturity Levels 2 and 3.

Regulatory notification clocks referenced in Domain 4
GDPR Article 33 (72 hours); NIS2 Article 23 (24-hour early warning, 72-hour notification, one-month final report); DORA Article 19 with Commission Delegated Regulation (EU) 2025/301 Article 5; SEC Form 8-K Item 1.05 (four business days from materiality determination); ISO/IEC 27001:2022 Annex A 8.15 and 8.16.

Toll Free US: 1(800) 834 1060
Asia/Pacific: +61 8 8213 1200
UK/Europe: +44 (800) 368 7423

AMERsales@prophecyinternational.com
APACsales@prophecyinternational.com
EMEAsales@prophecyinternational.com

snare
A PROPHECY SOLUTION