MOAR Capability Matrix

Public methodology and candidate catalog

> *"An independent lab and capability matrix for security data engineering — so you don't have to take any vendor's word for it."*

What this is

A versioned capability matrix mapping candidate tools per security-data substrate component, against criteria that vary in weight by client-specific workload and constraints. The matrix is the standing IP of the practice; this 1-pager shows the structure, candidate tools, and criteria taxonomy. The engagement-internal version applies client-specific weights, scores each candidate, and produces a defensible recommendation.

The matrix sits underneath **MOAR (Modular Open Architecture)** — the assertion that security-data substrates are best assembled as modular, replaceable components, with each component choice informed by evidence rather than vendor narrative.

Methodology

- Each component is evaluated against weighted criteria specific to the client environment.
- Scores: **1–5** (5 = best fit for the criterion).
- Weights: sum to 100; assigned per-engagement based on workload classification.
- Final scores: weight-adjusted; ties broken by qualitative judgment with documented reasoning.
- Cross-component dependencies are explicit (e.g., catalog choice constrains engine compatibility).

Components

#	COMPONENT	CANDIDATES
0	**Substrate Pattern**	Composed · Databricks · AWS Security Lake · Snowflake
1	**Lakehouse / Storage Format**	Iceberg · Delta Lake · Hudi
2	**Catalog / Metadata**	Hive Metastore · Polaris · Nessie · Unity · Glue
3	**Query Engine**	ClickHouse · Dremio · StarRocks · Trino · DuckDB
4	**Ingestion / Route**	Tenzir · Vector · Cribl · Kafka Connect · native shippers
5	**Graph / Visualization**	Splunk SH federated · Grafana · Superset · custom · vendor SOC UIs
6	**Storage Tier**	S3 (Standard / IA / Glacier) · MinIO · Wasabi · NetApp · Dell ECS

Component 0 — Substrate Pattern criteria taxonomy

- Operational capacity (data eng / DevOps headcount available)
- Engine independence required by workload performance
- On-prem residency / data sovereignty
- Air-gap requirements (cloud→on-prem network egress allowed?)
- Security-vertical methodology fit (DetectFlow, OCSF schema work, hunting workflow)
- Cost sensitivity (egress, managed-service premium, license model)
- Time-to-value
- ML / analytics integration scope

Disclosure principles

- **No reseller margins** on any product in the matrix.
- **No vendor-paid placements.**
- **Active partnerships disclosed** — every SOW includes Appendix B listing partnership relationships at execution. Recommendations involving a partner-vendor are accompanied by alternative-candidate comparison; client may request independent review at no additional cost.
- **Reproducibility as integrity guard** — public benchmarks (securitydataworks.com/lab) make matrix scoring auditable; if a partner-vendor scored unfairly relative to independent re-runs, the divergence would be visible.
- **Annual external sanity-check** — one external practitioner reviews recommendations annually for bias creep.

> **ENGAGEMENT BUNDLE** — every paid engagement above the $25K floor includes a 6-month subscription to the engagement-internal matrix (with weighted scoring, criteria-by-criterion reasoning, vendor-claim-vs-shipped-reality deltas, recommended bundles per workload archetype) plus 2 quarterly tool-eval reports during the subscription period.

Jeremy Wiley · OCSF 2.0 co-author · ITU-T SG-17 contributor securitydataworks.com · jeremy@securitydataworks.com