KQL for Detection and Threat Hunting

SEC201

Write powerful, efficient KQL queries that turn massive amounts of security data into fast, actionable insights. Master Kusto Query Language to hunt threats, build high-fidelity detections, investigate incidents, and automate response inside Microsoft Sentinel, Defender XDR, and Log Analytics.

Modules	22
Lessons	142
Phases	5
Free preview	Course Orientation (8 lessons)

The method

KQL is a query language, and most people learn it as a set of examples to copy. This course teaches the execution model first, because every performance problem and every wrong answer downstream is a model problem.

1. Understand how the query is processed. Data flows left to right through operators. Where you filter decides both what you get and what it costs.

2. Filter before you shape. The cheapest query is the one that discards rows early. This is the single habit that separates a query that runs at scale from one that times out.

3. Aggregate to answer a question, not to summarize. A summarize with no hypothesis behind it produces a table nobody reads. Know what the number is for before computing it.

4. Join only when the answer needs two sources. Joins are the most expensive thing in the language and the most reached for. Most correlation problems are a filter and a lookup.

5. Turn the query into something that runs without you. A detection rule, a workbook or a scheduled hunt. A query that lives in somebody's notes is a query that runs once.

What this course assumes

No minimum experience and no prerequisite course. The language is built from the first query onward, and the course assumes no prior SQL.

What makes it go faster: access to a Sentinel workspace or Defender advanced hunting with real data. Not required, because every query in the course runs against the Practice Hub, which holds populated tables and known answers.

What this course does not cover: incident response process, detection engineering as a program, and SIEM administration. This is the query language, taught to the depth where you stop copying examples.

Full syllabus

Phase 1: Anatomy of KQL

Course Orientation FREE PREVIEW

What KQL for Detection and Threat Hunting teaches: the query language behind every hunt, detection, and investigation in Microsoft Sentinel, Defender XDR, and Log Analytics.

1. 0.1 Three Questions You Cannot Answer Without KQL
2. 0.2 The Microsoft Security Data Model
3. 0.3 The Eight Tables You Will Query Every Day
4. 0.4 Your First Security Query
5. 0.5 What This Course Builds
6. 0.6 Setting Up Your Lab Environment
7. Module Summary
8. Check My Knowledge

How KQL Processes Data

The tabular data model, the query pipeline, every data type, type conversion, and null handling.

1. 1.1 The Tabular Data Model. Everything Is a Table
2. 1.2 The Query Pipeline. How Operators Chain Left to Right
3. 1.3 Data Types in Depth
4. 1.4 Type Conversion and Coercion
5. 1.5 Null Handling. The Silent Query Killer
6. Module Summary
7. Check My Knowledge

Filtering and Shaping Data

Master every where operator, string comparison, datetime filter, extend pattern, and project variant.

1. 2.1 Where. Comparison Operators, Logical Operators, and Compound Conditions
2. 2.2 Where. String Operators Deep Dive
3. 2.3 Where. Datetime Filtering
4. 2.4 Extend. Calculated Columns
5. 2.5 Project, Project-Away, Project-Rename, and Project-Reorder
6. Module Summary
7. Check My Knowledge

Aggregation and Statistical Analysis

Every summarize function, time-based grouping with bin, set and list aggregations, ordering and limiting, and pivot patterns.

1. 3.1 Summarize. Counting, Distinct Counts, and Conditional Aggregation
2. 3.2 Summarize. Set, List, and Bag Aggregations
3. 3.3 Bin and Time-Based Grouping
4. 3.4 Top, Sort, Take, and Limit
5. 3.5 Pivot Patterns. Turning Rows into Columns
6. Module Summary
7. Check My Knowledge

Phase 2: Intermediate Techniques

Joining and Correlating Data

Every join flavour with performance implications, union for combining tables, lookup for enrichment, and cross-table investigation patterns that trace attacks across data sources.

1. 4.1 Join Fundamentals. Combining Rows from Two Tables
2. 4.2 Join Flavours. Choosing the Right Join for Every Investigation
3. 4.3 Join Performance. Why Joins Are Expensive and How to Control the Cost
4. 4.4 Union. Combining Rows from Multiple Tables
5. 4.5 Lookup. Lightweight Enrichment Without Full Join Overhead
6. 4.6 Cross-Table Investigation Patterns. Tracing Attacks Across Data Sources
7. 4.7 Materialize and Let. Reusing Results Across Complex Queries
8. Module Summary
9. Check My Knowledge

String Parsing and Data Extraction

Parse, extract, parse_json, string manipulation, URL decomposition, and building log parsers.

1. 5.1 The parse Operator. Extracting Structure from Unstructured Strings
2. 5.2 extract and extract_all. Regex-Based Extraction
3. 5.3 parse_json and Dynamic Nested Data
4. 5.4 String Functions, split, substring, replace, and Transformation
5. 5.5 URL and Email Decomposition for Threat Analysis
6. 5.6 Building Custom Log Parsers
7. Module Summary
8. Check My Knowledge

Advanced Filtering and Pattern Matching

Regular expressions, dynamic arrays with mv-expand and mv-apply, the let statement, user-defined functions, and externaldata.

1. 6.1 mv-expand and mv-apply. Working with Dynamic Arrays
2. 6.2 Advanced Regex Patterns for Security Analysis
3. 6.3 let Statements and User-Defined Functions
4. 6.4 externaldata. Querying External Reference Files
5. 6.5 IP Range Functions and CIDR Matching
6. 6.6 Watchlist-Driven Detection Patterns
7. Module Summary
8. Check My Knowledge

Phase 3: Advanced Patterns

Time-Series Analysis and Anomaly Detection

Make-series, series decomposition, automated anomaly detection, statistical analysis, baseline comparison patterns, and building anomaly detection rules from raw logs.

1. 7.1 make_series. Converting Events into Time-Series Data
2. 7.2 series_decompose. Separating Trend, Seasonal, and Residual Components
3. 7.3 series_decompose_anomalies. Automated Anomaly Detection
4. 7.4 Baseline Comparison Patterns. Per-Entity Behavioral Profiles
5. 7.5 Rate-of-Change Detection and Sliding Windows
6. 7.6 Building Anomaly Detection Rules for Production
7. Module Summary
8. Check My Knowledge

Graph and Relationship Analysis

Entity relationships, attack path reconstruction, network graph patterns, identity mapping, and recursive multi-hop investigation queries.

1. 8.1 Entity Relationship Modeling in Security Logs
2. 8.2 Attack Path Reconstruction
3. 8.3 Multi-Hop Investigation Queries
4. 8.4 Identity Mapping Across Tables
5. 8.5 Network Connection Graph Patterns
6. 8.6 Process Tree Reconstruction
7. Module Summary
8. Check My Knowledge

Performance Optimization and Scale

Query engine internals, performance analysis, materialized views, optimizing for large datasets, common anti-patterns, and writing queries that scale from development to production.

1. 9.1 Query Engine Internals. How KQL Executes Your Queries
2. 9.2 Common Anti-Patterns That Kill Performance
3. 9.3 Partition Pruning and Term Indexing
4. 9.4 Materialized Views and Stored Functions
5. 9.5 Optimizing Joins and Time-Series Queries at Scale
6. 9.6 Query Limits, Diagnostics, and Performance Monitoring
7. Module Summary
8. Check My Knowledge

Phase 4: Mastery

Detection Rule Engineering with KQL

Converting investigation queries into production detection rules.

1. 10.1 From Investigation Query to Detection Rule
2. 10.2 Scheduled vs NRT Rules. Choosing the Right Rule Type
3. 10.3 Entity Mapping and Alert Enrichment
4. 10.4 Threshold Tuning and False Positive Management
5. 10.5 Multi-Condition Correlated Detection
6. 10.6 Detection Rule Testing and Validation
7. Module Summary
8. Check My Knowledge

Threat Hunting with KQL

Hypothesis-driven hunting methodology, MITRE ATT&CK hunting, behavioral hunting with baseline deviation and peer group analysis, IOC sweeps, and building a reusable hunting query library.

1. 11.1 Hypothesis-Driven Hunting Methodology
2. 11.2 MITRE ATT&CK-Aligned Hunting
3. 11.3 Behavioral Hunting with Baselines and Peer Groups
4. 11.4 IOC and Retroactive Sweeps
5. 11.5 Hunt Management, Bookmarks, and Evidence
6. 11.6 Building and Maintaining a Hunting Query Library
7. Module Summary
8. Check My Knowledge
9. 11.9 Mini Capstone. The Live Hunt Exercise

Workbooks, Dashboards, and Operational Reporting

KQL for Sentinel workbooks with parameterized queries, building the security operations dashboard, executive reporting queries, and automated report generation.

1. 12.1 Sentinel Workbook Fundamentals
2. 12.2 Parameterized Queries and Dynamic Filtering
3. 12.3 Building the SOC Operations Dashboard
4. 12.4 Executive Reporting Queries
5. 12.5 Detection Health Monitoring
6. 12.6 Automated Report Generation
7. Module Summary
8. Check My Knowledge

Phase 5: Capstone

Capstone, The Hunting Lab

Three complete hunting scenarios that test every skill from the course.

1. 13.1 Capstone Lab. Scenario Brief and Setup
2. 13.2 Scenario 1. Credential Spray to Business Email Compromise
3. 13.3 Scenario 2. Ransomware Kill Chain Reconstruction
4. 13.4 Scenario 3. Cloud Data Exfiltration
5. 13.5 Capstone Review and Self-Assessment

Phase 0: Course Resources

Cheatsheets

Seventy-four queries across ten subject-area sheets, each with the operators it turns on and what the result does not establish.

1. Foundations
2. Filtering and Shaping
3. Aggregation
4. Joining and Correlating
5. Parsing and Matching
6. Time Series and Anomalies
7. Graph and Relationships
8. Performance
9. Detection and Hunting
10. Reporting and Capstone
11. Operators at a Glance
12. Limits, Ceilings and Silent Failures

Cookbooks

Five procedures for the recurring problems, worked against documented platform behavior: ingestion lag, table plans, memory and truncation limits, service limits on rules, and transformation constraints.

1. A Query Returns Nothing
2. A Query Is Too Slow
3. From Query to Detection Rule
4. Parsing an Unfamiliar Log Source
5. Running a Hunt

Lab Setup

A Log Analytics workspace you build and populate yourself, free, with the activity each module's queries need and the verification that proves it holds data.

1. Building the Environment

Walkthroughs

Ten investigations worked end to end with the output at every step, each turning on a failure that produces no error message: queries that return nothing, too few rows, or exactly the right answer to the wrong question.

1. Building a Query From a Question
2. When the Join Lies
3. The Baseline That Learned the Attack
4. The Query That Timed Out
5. The Parse That Half Worked
6. Reconstructing the Process Tree
7. The Rule That Fired Four Thousand Times
8. The Hunt That Proved Nothing
9. The Dashboard That Improved
10. Capstone: Spray to Business Email Compromise

Playbooks

Five query patterns by investigative question, each built on the function KQL actually provides for it, with the constraint that decides the design and the false positives it will produce.

1. First Time Seen
2. Rare in Population
3. Volume Anomaly Per Entity
4. Paired Events in Sequence
5. Beaconing and Regularity

Playground

Resources for practicing KQL and using it at work: the detection library, a free-run query console, guided investigations, response playbooks, and an open-source DFIR toolkit.

References & Further Reading

External sources this course draws on: Microsoft KQL documentation, the language reference, and community query material.

Course Completion

Course Exam

KQL for Detection and Threat Hunting end-of-course exam: a three-phase simulation testing whether you can apply the method to a situation the course did not walk through.

1. Course Completion. KQL for Detection and Threat Hunting

Generated 2026-09-01 from the published course. The current syllabus is always at ridgelinecyber.com/training/courses/kql-detection-threat-hunting/