

Microsoft 365 Threat Hunting

SEC402

Go beyond alerts and become a proactive hunter. Learn hypothesis-driven threat hunting across Microsoft 365, Entra ID, Defender XDR, and Sentinel to find stealthy attackers, uncover hidden compromises, and hunt like a professional before threats become breaches.

Modules	26
Lessons	264
Phases	3
Free preview	Course Orientation (9 lessons)

The method

A hunt is a hypothesis you can be wrong about, run against data with a stated coverage. This course runs the same cycle for every hunt it works, sixteen modules of it.

1. Write a hypothesis that can fail. If no result would change your mind, it is not a hunt. The falsifiable version is what makes an empty outcome reportable.
2. Establish how far back you can look. Retention differs by table and by license tier, and a hypothesis about six months ago is bounded by the shortest source it touches.
3. Query for the behavior, then baseline it. A result is not a finding until it has been compared against what normal looks like in your estate.
4. Check the entity before escalating. An event inside your window is not evidence your hypothesis was right. The entity check is what separates a finding from a coincidence.
5. Convert the finding into a rule, and the miss into coverage. A hunt that found something should end as a detection. A hunt that found nothing should end as a coverage statement naming what was searched.

The course closes on building the hunt program: a backlog with an order, documentation that outlives the hunter, and hunts that run without one.

What this course assumes

No minimum experience and no prerequisite course. KQL, the ATT&CK model and the M365 telemetry surface are introduced where they first matter.

What makes it go faster: a tenant with real telemetry and any prior KQL. Neither is required. Every hunt in the course runs against the Practice Hub, which holds populated tables and known answers.

What this course does not cover: incident response process, forensics and detection engineering as a program. A hunt hands off to those, and the handoff is taught rather than the destination.

Full syllabus

Phase 1: Hunt Methodology & Advanced Toolcraft

Course Orientation [FREE PREVIEW](#)

What Microsoft 365 Threat Hunting teaches: run the hypothesis-driven hunt cycle to find the stealthy attacker no alert fired on, across Entra ID, Defender XDR, and Sentinel, through ten complete hunt campaigns, and turn every hunt into a new detection.

1. 0.1 The Detection Gap
2. 0.2 Why Detection Engineering Cannot Close the Gap
3. 0.3 The M365 Threat Landscape
4. 0.4 Where Hunting Fits
5. 0.5 The Business Case for Hunting
6. 0.6 Organizational Readiness and Data Sources
7. 0.7 The Hunter's Skillset and Maturity
8. 0.8 Your First 90 Days
9. Module Summary

The Hunt Cycle, A Structured Methodology

The six-step methodology that every campaign module follows.

1. 1.1 Formulating Hunt Hypotheses
2. 1.2 Scoping the Hunt
3. 1.3 Collection: Iterative Querying
4. 1.4 Analysis: Separating Signal from Noise
5. 1.5 Concluding the Hunt
6. 1.6 Converting Hunts to Detection Rules
7. 1.7 The Hunt Documentation Standard
8. 1.8 The Hunt-to-Detection Pipeline: Worked End-to-End
9. 1.9 Hunt Cadence and Scheduling
10. 1.10 Hunt Quality Assurance and Metrics
11. Module Summary
12. Check My Knowledge

Advanced KQL for Hunting

The KQL patterns that separate hunting queries from detection rules.

1. 2.1 Troubleshooting and Debugging KQL Queries
2. 2.2 Percentile and Statistical Deviation for Outlier Detection
3. 2.3 The make-series Operator
4. 2.4 series_decompose_anomalies()
5. 2.5 series_fir() and Time Series Smoothing
6. 2.6 top-nested for Frequency Analysis
7. 2.7 autoclassify() for Behavioral Grouping
8. 2.8 arg_max and arg_min for Entity Investigation
9. 2.9 Dynamic Column Parsing for M365 Logs
10. 2.10 materialize() and Query Performance
11. 2.11 Entity Pivoting Across Tables
12. Module Summary
13. Check My Knowledge

ATT&CK Coverage Analysis: Prioritizing Your Hunt Backlog

Map your current detection rules to MITRE ATT&CK.

1. TH3.1 Exporting Your Analytics Rule Inventory
2. TH3.2 Mapping Rules to MITRE ATT&CK Techniques
3. TH3.3 Visualizing Coverage with ATT&CK Navigator
4. TH3.4 Identifying Coverage Gaps by Tactic
5. TH3.5 Coverage Clustering. Where Rules Concentrate
6. TH3.6 Threat Intelligence–Driven Gap Prioritization
7. TH3.7 Data Availability Scoring
8. TH3.8 Impact Severity Scoring
9. TH3.9 Building the Composite Score Matrix
10. TH3.10 Assembling the Hunt Backlog
11. TH3.11 Selecting Your First Three Campaigns
12. TH3.12 Validating Hypotheses Before Hunting
13. TH3.13 Backlog Maintenance and Quarterly Review
14. TH3.14 Common Coverage Analysis Mistakes
15. TH3.15 From Analysis to Action. Your 90-Day Plan
16. Module Summary
17. Check My Knowledge

Phase 2: Hunt Campaigns

Hunting Identity Compromise

The first campaign module. Six identity attack techniques hunted across interactive and non-interactive sign-in tables: AiTM session hijacking, credential stuffing, password spray, MFA fatigue, custom impossible travel, and session token anomalies. Every query runs against your environment....

1. TH4.1 Building the Authentication Baseline
2. TH4.2 Hunting AiTM Session Hijacking
3. TH4.3 Hunting Credential Stuffing
4. TH4.4 Hunting Password Spray
5. TH4.5 Hunting MFA Fatigue Attacks
6. TH4.6 Hunting MFA Bypass Techniques
7. TH4.7 Building Custom Impossible Travel Detection
8. TH4.8 Session Hijacking via Non-Interactive Tokens
9. TH4.9 Hunting Compromised Service Principals
10. TH4.10 Authentication Anomalies with Time-Series Analysis
11. TH4.11 Correlating Identity Compromise Indicators
12. TH4.12 Building the Identity Hunt Report
13. TH4.13 Converting Identity Hunts to Detection Rules
14. TH4.14 Common Identity Hunting Mistakes
15. TH4.15 Identity Hunt Operational Playbook
16. Module Summary
17. Check My Knowledge

Hunting Cloud Persistence Mechanisms

The second campaign module. Seven cloud persistence techniques hunted across CloudAppEvents, AuditLogs, and email tables: inbox rules via Graph API, mail forwarding, OAuth consent persistence, MFA method registration, conditional access manipulation, federated trust abuse, and temporal...

1. TH5.1 The Cloud Persistence Landscape
2. TH5.2 Hunting Inbox Rules Created via Graph API
3. TH5.3 Hunting Mail Forwarding and Redirect Rules
4. TH5.4 Hunting OAuth Consent Persistence
5. TH5.5 Hunting Admin Consent and Tenant-Wide App Access
6. TH5.6 Hunting MFA Method Registration as Persistence
7. TH5.7 Hunting Conditional Access Policy Manipulation
8. TH5.8 Hunting Federated Trust Abuse
9. TH5.9 Temporal Correlation with Initial Compromise
10. TH5.10 Building the Persistence Hunt Report
11. TH5.11 Converting Persistence Hunts to Detection Rules
12. TH5.12 Common Persistence Hunting Mistakes
13. TH5.13 Persistence Hunt Operational Playbook
14. Module Summary
15. Check My Knowledge

Hunting Privilege Escalation and Abuse

The third campaign module. Seven privilege escalation techniques hunted across AuditLogs, PIM logs, and directory change records: direct role assignments outside PIM, Global Admin activation anomalies, emergency access account misuse, service principal credential abuse, admin consent for...

1. TH6.1 The Privilege Escalation Landscape
2. TH6.2 Hunting Role Assignments Outside PIM
3. TH6.3 Hunting Global Admin Activation Anomalies
4. TH6.4 Hunting Emergency Access Account Misuse
5. TH6.5 Hunting Service Principal Credential Abuse for Escalation
6. TH6.6 Hunting Admin Consent for Privilege Escalation
7. TH6.7 Hunting Security Group Manipulation
8. TH6.8 Hunting CA Weakening for Elevated Access
9. TH6.9 Temporal Correlation. Escalation Timeline
10. TH6.10 Building the Privilege Escalation Hunt Report
11. TH6.11 Converting to Detection Rules
12. TH6.12 Common Mistakes
13. TH6.13 Privilege Escalation Hunt Operational Playbook
14. Module Summary
15. Check My Knowledge

Hunting Email-Based Threats

The fourth campaign module. Seven email threat techniques hunted across EmailEvents, CloudAppEvents, and authentication tables: BEC from compromised internal accounts, internal phishing from legitimate mailboxes, vendor email compromise, mail flow rule manipulation, auto-forwarding to external...

1. TH7.1 The Email Threat Landscape in M365
2. TH7.2 Hunting BEC from Compromised Internal Accounts
3. TH7.3 Hunting Internal Phishing from Legitimate Mailboxes
4. TH7.4 Hunting Vendor Email Compromise
5. TH7.5 Hunting Mail Flow Rule Manipulation
6. TH7.6 Hunting Financial Keyword Interception
7. TH7.7 Consolidated External Forwarding Audit
8. TH7.8 Email Campaign Correlation with Authentication Anomalies
9. TH7.9 Building the Email Threat Hunt Report
10. TH7.10 Converting to Detection Rules
11. TH7.11 Common Mistakes
12. TH7.12 Email Threat Hunt Operational Playbook
13. Module Summary
14. Check My Knowledge

Hunting Data Exfiltration

The fifth campaign module. Six data exfiltration channels hunted across CloudAppEvents, SharePoint audit logs, and device telemetry: SharePoint and OneDrive bulk download anomalies, external sharing link creation, email-based data exfiltration, Teams file exfiltration, browser downloads to...

1. TH8.1 The Data Exfiltration Landscape
2. TH8.2 Hunting SharePoint and OneDrive Bulk Downloads
3. TH8.3 Hunting External Sharing Link Creation
4. TH8.4 Hunting Email-Based Data Exfiltration
5. TH8.5 Hunting Teams File Exfiltration
6. TH8.6 Hunting Downloads to Unmanaged Devices
7. TH8.7 Multi-Channel Exfiltration Correlation
8. TH8.8 Building the Data Exfiltration Hunt Report
9. TH8.9 Converting to Detection Rules
10. TH8.10 Common Mistakes
11. TH8.11 Data Exfiltration Hunt Operational Playbook
12. Module Summary
13. Check My Knowledge

Hunting Endpoint Threats

The sixth campaign module. Seven endpoint threat techniques hunted across DeviceProcessEvents, DeviceNetworkEvents, DeviceFileEvents, and DeviceRegistryEvents: LOLBin abuse with network connections, process injection indicators, registry and scheduled task persistence, defense evasion...

1. TH9.1 The Endpoint Threat Landscape
2. TH9.2 Hunting LOLBin Abuse with Network Connections
3. TH9.3 Hunting Process Injection Indicators
4. TH9.4 Hunting Registry and Scheduled Task Persistence
5. TH9.5 Hunting Defense Evasion
6. TH9.6 Hunting C2 Beaconing
7. TH9.7 Hunting Fileless Execution
8. TH9.8 Process Tree Analysis
9. TH9.9 Building the Endpoint Threat Hunt Report
10. TH9.10 Converting to Detection Rules
11. TH9.11 Common Mistakes
12. TH9.12 Endpoint Threat Hunt Operational Playbook
13. Module Summary
14. Check My Knowledge

Hunting Lateral Movement

The seventh campaign module. Seven lateral movement techniques hunted across SigninLogs, DeviceNetworkEvents, DeviceLogonEvents, and hybrid identity tables: cloud-to-cloud token reuse across applications, cloud-to-endpoint pivot correlation, RDP/SMB/WMI/PowerShell Remoting detection, service...

1. TH10.1 The Lateral Movement Landscape
2. TH10.2 Hunting Cloud Token Reuse Across Applications
3. TH10.3 Hunting Cloud-to-Endpoint Pivot
4. TH10.4 Hunting RDP, SMB, WMI, and PowerShell Remoting
5. TH10.5 Hunting Service Account Lateral Abuse
6. TH10.6 Hunting VPN Pivot. Population Analysis
7. TH10.7 Hunting Credential Harvesting. NTLM and Kerberos Anomalies
8. TH10.8 Hunting Azure AD Connect Abuse
9. TH10.9 Building the Lateral Movement Hunt Report
10. TH10.10 Converting to Detection Rules
11. TH10.11 Common Mistakes
12. TH10.12 Lateral Movement Hunt Operational Playbook
13. Module Summary
14. Check My Knowledge

Hunting Application and API Abuse

The eighth campaign module. Seven application and API threat techniques: shadow IT discovery via Cloud App Events, OAuth application inventory and risk scoring, Graph API abuse detection, third-party app excessive permissions, dormant high-privilege applications, AI tool usage with corporate...

1. TH11.1 The Application Threat Landscape
2. TH11.2 Shadow IT Discovery
3. TH11.3 OAuth Application Risk Scoring
4. TH11.4 Hunting Graph API Abuse
5. TH11.5 Hunting Excessive Application Permissions
6. TH11.6 Hunting Dormant Application Reactivation
7. TH11.7 Hunting AI Tool Usage
8. TH11.8 Hunting Unsanctioned File Uploads
9. TH11.9 Building the Application Abuse Hunt Report
10. TH11.10 Converting to Detection Rules
11. TH11.11 Common Mistakes
12. TH11.12 Application Abuse Hunt Operational Playbook
13. Module Summary
14. Check My Knowledge

Hunting Pre-Ransomware Activity

The ninth campaign module. Six pre-ransomware indicators hunted across DeviceProcessEvents, DeviceFileEvents, DeviceRegistryEvents, and DeviceNetworkEvents: reconnaissance tool execution sequences, volume shadow copy and backup disruption, credential harvesting indicators, staging directory...

1. TH12.1 The Pre-Ransomware Landscape
2. TH12.2 Hunting Ransomware-Specific Reconnaissance
3. TH12.3 Hunting Backup Disruption
4. TH12.4 Hunting Credential Harvesting for Domain-Wide Deployment
5. TH12.5 Hunting Ransomware Staging
6. TH12.6 Hunting Known Ransomware Tooling
7. TH12.7 Temporal Chain Correlation. The Complete Pre-Ransomware Timeline
8. TH12.8 The Pre-Ransomware Hunt Report
9. TH12.9 Converting Pre-Ransomware Findings to Detection Rules
10. TH12.10 Common Pre-Ransomware Hunting Mistakes
11. TH12.11 Pre-Ransomware Hunt Playbook
12. Module Summary
13. Check My Knowledge

Hunting Insider Threats

The tenth and final campaign module. Insider threats are entirely different from external attacks, the adversary has legitimate access, their individual actions are authorized, and the detection surface is behavioral deviation rather than signature or anomaly. This module teaches you to detect...

1. TH13.1 The Insider Threat Landscape
2. TH13.2 Hunting Data Hoarding and Bulk Access
3. TH13.3 Hunting Access Pattern Deviation
4. TH13.4 Resignation and Termination Correlation
5. TH13.5 Hunting Privilege Abuse by Authorized Users
6. TH13.6 Hunting After-Hours Bulk Activity
7. TH13.7 Hunting External Account Communication
8. TH13.8 Multi-Source Behavioral Assessment
9. TH13.9 The Insider Threat Hunt Report
10. TH13.10 Converting Findings to Detection Rules
11. TH13.11 Common Insider Threat Hunting Mistakes
12. TH13.12 Insider Threat Hunt Playbook
13. Module Summary
14. Check My Knowledge

Phase 3: Hunt Operations

Building a Hunt Program

Phase 3 begins. TH0–TH13 taught you how to hunt. TH14–TH16 teach you how to build, sustain, and scale a hunt program. This module covers cadence selection, prioritization frameworks, staffing models, the hunt-to-detection pipeline, SOC integration, and the hunt program charter, the document that...

1. TH14.1 Hunt Cadence Models
2. TH14.2 Hunt Prioritization Framework
3. TH14.3 Staffing Models
4. TH14.4 The Hunt-to-Detection Pipeline
5. TH14.5 SOC Integration and Budget Justification
6. TH14.6 The Hunt Program Charter
7. TH14.7 Hunt Tooling and Workspace Organization
8. TH14.8 Program Maturity Roadmap
9. Module Summary
10. Check My Knowledge

Hunt Documentation, Reporting, and Knowledge Management

The hunt program generates findings, negative results, and detection improvements every cycle.

1. TH15.1 The Hunt Finding Documentation Standard
2. TH15.2 The Value of Negative Findings
3. TH15.3 Reporting to Leadership in Business Language
4. TH15.4 The Technical Hunt Report
5. TH15.5 Building and Maintaining the Hunt Knowledge Base
6. TH15.6 Program Metrics and Dashboards
7. TH15.7 Report Quality Assurance and Stakeholder Communication
8. TH15.8 Finding Trend Analysis and Year-Over-Year Reporting
9. Module Summary
10. Check My Knowledge

Scaling Hunts: Automation, Notebooks, and Continuous Hunting

The final module. The program is established (TH14), documented (TH15), and producing results. This module teaches you to scale it: automate recurring queries with scheduled hunts, manage hypothesis-driven investigations with Sentinel hunt management, build reproducible analysis with Jupyter...

1. TH16.1 Scheduled Hunting Queries
2. TH16.2 Sentinel Hunt Management
3. TH16.3 Jupyter Notebooks with MSTICPy
4. TH16.4 Hunting Workbooks and Continuous Dashboards
5. TH16.5 The Maturity Continuum
6. TH16.6 Continuous Hunting Operations
7. TH16.7 Course Conclusion
8. TH16.8 Automating Hunt Response with Sentinel Automation
9. Module Summary
10. Check My Knowledge

Phase 0: Course Resources

Cheatsheets

Hunt queries by domain, each with the fields that carry the decision and what the result does not establish.

1. Hunt Query Patterns
2. Identity
3. Cloud Persistence
4. Privilege Escalation
5. Email
6. Data Exfiltration
7. Endpoint
8. Lateral Movement
9. Application and API
10. Pre-Ransomware
11. Insider
12. Tables, Codes and Coverage

Cookbooks

Six runbooks: executing a hunt, converting findings to detections, working an advisory, sprints, the backlog, and reporting.

1. Running a Hunt
2. Converting a Finding to a Detection
3. Working a Threat Advisory
4. Running a Hunt Sprint
5. Building the Hunt Backlog
6. Reporting to Leadership

Lab Setup

Building an environment you can hunt in, and the constraint that separates a hunting lab from every other kind.

1. Building the Environment

Walkthroughs

Worked hunts end to end, including the wrong turns, the null results, and the candidate that looked certain and was not.

1. The Hunt That Found Nothing
2. The Same Hunt, Five Weeks Later
3. The Baseline That Had Already Learned It
4. A Grant Nobody Reviewed
5. The Cluster That Hid It
6. The Candidate That Cleared
7. The Answer You Cannot Give
8. The Chain That Kept Going

Playbooks

One operational playbook per hunt domain: the sequence, what to record at each step, and what escalates.

1. Identity
2. Cloud Persistence
3. Privilege Escalation
4. Email
5. Data Exfiltration
6. Endpoint
7. Lateral Movement
8. Application and API
9. Pre-Ransomware
10. Insider

Playground

Resources for practicing hunting and using it at work: a corpus with history, guided investigations, the detection library, response procedures, and a DFIR toolkit.

Operational Reference

The consolidated lookup and the step-by-step procedures from this course, in one place.

1. Hunt Query Quick Reference
2. Threat Hunting Field Manual

References & Further Reading

External sources this course draws on: vendor documentation, frameworks, standards, and research.

Course Completion

Course Exam

Microsoft 365 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. Microsoft 365 Threat Hunting

Generated 2026-09-01 from the published course. The current syllabus is always at ridgelinecyber.com/training/courses/threat-hunting-m365/