Adversarial Tradecraft for Defenders

SEC403

Read an intrusion as the operator's decisions: what each step needed, how it was carried out, what it left on every host and in the cloud, what was attempted against what succeeded, and which defensive changes remove the operator's options, proved against the records.

Modules	20
Lessons	166
Phases	5
Free preview	Investigative Foundations of Adversarial Tradecraft (8 lessons)

The method

Every lesson teaches one attacker procedure from documented intrusions and the tools' own documentation, then finds it in the course's records. The loop is the same throughout.

1. Start from the attacker's decision. What the operator wanted at this step, what access they held, what constrained them, and which option they chose. The decision explains the step better than the tool does.

2. Learn how it is carried out. The commands operators run, shown from published intrusions, Atomic Red Team and vendor documentation, with the conditions each one needs.

3. Read what it leaves. Each step is found in endpoint, identity, network and cloud records, in KQL in the browser with the same search in SPL beside it.

4. Keep the claims honest. An attempt is not a success and a success is not an effect. Every finding says which it is, and what the records cannot establish.

5. Put it together. Every module ends with a practical on the campaign, written up as an evidence and decision record, and the last tests a set of defensive changes against the whole week.

What this course assumes

Comfort with basic networking, Windows administration and security logs, at the level of someone who has read alerts for a living. No prerequisite course and no offensive background.

What makes it go faster: having written a few simple queries, and knowing how processes start other processes. Module 0 includes a diagnostic exercise that shows where to refresh before starting.

What this course does not cover: penetration testing as a profession, building payloads or implants, exploit development, and red team operations. Attacker procedures are taught to the depth a defender needs to recognize and stop them, and nothing is run against live systems.

Full syllabus

Phase 1: Foundations, Reporting and Initial Access

Investigative Foundations of Adversarial Tradecraft FREE PREVIEW

Module 0 of Adversarial Tradecraft for Defenders: what tradecraft is, how an operator's objectives, access and constraints shape each choice, how techniques differ from procedures and tools, and how to separate attempts from successes and facts from inferences before you reconstruct an intrusion.

1. 0.1 What Adversarial Tradecraft Means for a Defender
2. 0.2 Attacker Objectives, Access and Operational Constraints
3. 0.3 Techniques, Procedures, Tools and Observable Behavior
4. 0.4 Attempted Activity, Successful Execution and Business Impact
5. 0.5 Facts, Inferences, Hypotheses and Confidence
6. 0.6 Reading an Intrusion Without Assuming a Linear Sequence
7. 0.7 Diagnostic Exercise: Explain a Short Intrusion
8. 0.8 Practical: Build Your First Evidence and Decision Record

Turning Threat Reporting into an Operational Model

Module 1 of Adversarial Tradecraft for Defenders: read a threat report for what the operator did, separate what was observed from what was assessed, and turn it into an attack path, a campaign brief and an evidence plan for one organization.

1. 1.1 Reading Reports for Behavior Rather Than Indicators
2. 1.2 Separating Reported Facts from Analyst Interpretation
3. 1.3 Identifying Objectives and Targeted Business Services
4. 1.4 Mapping Behaviors to ATT&CK Techniques
5. 1.5 Recording Prerequisites and Environmental Assumptions
6. 1.6 Building an Attack Path with Decision Points
7. 1.7 Identifying Alternative Routes and Failed Attempts
8. 1.8 Selecting Relevant Full and Micro-Emulation Plans
9. 1.9 Predicting Evidence and Defensive Opportunities
10. 1.10 Practical: Produce a Campaign Brief and Evidence Plan

Reconnaissance and Initial Access

Module 2 of Adversarial Tradecraft for Defenders: what an operator learns before making contact, how a way in is chosen, and how to investigate two entry routes into one environment from the records each leaves.

1. 2.1 External Reconnaissance and Target Selection
2. 2.2 DNS, Certificates and Infrastructure Relationships
3. 2.3 Exposed Services and Remote-Access Interfaces
4. 2.4 Reconnaissance Visibility: Approved Scanning and Suspicious Discovery
5. 2.5 Initial Access as a Choice Between Routes
6. 2.6 Phishing, Attachments, Downloads and Origin Information
7. 2.7 Valid Accounts, Remote-Access Services and Trusted Relationships
8. 2.8 Exposed Applications and Controlled Exploitation
9. 2.9 Delivery, Execution and Access: Blocked and Partial Chains
10. 2.10 Practical: Investigate Two Entry Routes into the Same Environment

Phase 2: Execution, Discovery and Credentials

Execution and Living-off-the-Land Activity

Module 3 of Adversarial Tradecraft for Defenders: what an operator ran on a compromised machine, how it ran and as whom, and how to tell it from an administrator running the same commands.

1. 3.1 Process Creation, Parentage and Execution Context
2. 3.2 PowerShell and Command-Shell Activity
3. 3.3 Script Hosts and Native Execution Utilities
4. 3.4 Remote Execution versus Local Execution
5. 3.5 Services and Scheduled Execution
6. 3.6 Download, Decode and Execute Sequences
7. 3.7 User Context, Integrity Levels and Privileges
8. 3.8 Process and Network Correlation
9. 3.9 Legitimate Administration and Misleading Indicators
10. 3.10 Practical: Reconstruct Execution within Administrative Noise

Discovery and Attack-Path Analysis

Module 4 of Adversarial Tradecraft for Defenders: what an operator learned about a network from one foothold, which routes that knowledge opened, and how to tell a route that exists from a route that was used.

1. 4.1 Host, Network and Domain Discovery
2. 4.2 Accounts, Groups and Privileged Membership
3. 4.3 Domain Relationships and Trust Boundaries
4. 4.4 Shares, Services and Accessible Resources
5. 4.5 Session and Administrative Relationship Discovery
6. 4.6 Reading Identity Relationships in BloodHound
7. 4.7 Validating Graph Edges and Required Conditions
8. 4.8 Potential Paths versus Observed Traversal
9. 4.9 Prioritizing Defensive Choke Points
10. 4.10 Practical: Validate a Path and Explain How to Break It

Credential, Authentication and Session Abuse

Module 5 of Adversarial Tradecraft for Defenders: how operators obtain and reuse access, in the forms a credential takes, and what the records can and cannot prove about each.

1. 5.1 Passwords, Hashes, Tickets, Tokens and Sessions
2. 5.2 Credential Exposure in Files and Configuration
3. 5.3 Controlled Credential-Access Demonstrations
4. 5.4 Kerberos Abuse and Its Prerequisites
5. 5.5 NTLM Authentication and Relay Conditions
6. 5.6 Reuse of Valid Credentials
7. 5.7 Session and Token Abuse
8. 5.8 Service Accounts and Non-Human Identities
9. 5.9 Credential Access versus Confirmed Credential Use
10. 5.10 Practical: Trace a Compromised Identity Across Evidence Sources

Phase 3: Escalation, Movement and Persistence

Privilege Escalation and Permission Abuse

Module 6 of Adversarial Tradecraft for Defenders: how operators gain higher access through groups, services, permissions, delegated identities and Linux sudo rules, and how the records show that access rose and what it was used for.

1. 6.1 Local Privilege and Domain Privilege
2. 6.2 Excessive Group Membership and Delegated Rights
3. 6.3 Service and Task Configuration Weaknesses
4. 6.4 File, Directory and Registry Permissions
5. 6.5 Application and Service Identity Privileges
6. 6.6 Linux Sudo and Permission Misconfiguration
7. 6.7 Certificate-Service Attack Paths: A Bounded Case
8. 6.8 Evidence of Changed Privilege and Subsequent Use
9. 6.9 Remediation That Removes the Enabling Condition
10. 6.10 Practical: Repair an Escalation Path and Retest It

Lateral Movement and Remote Administration

Module 7 of Adversarial Tradecraft for Defenders: how operators move from one machine to the next over SMB, remote desktop, WinRM, WMI, SSH and remote-management tools, and how to reconstruct each hop from both ends while accounting for legitimate administration.

1. 7.1 Preconditions for Moving to Another Host
2. 7.2 SMB and Remote-Service Activity
3. 7.3 RDP Sessions and Their Evidence
4. 7.4 WinRM and PowerShell Remoting
5. 7.5 WMI-Based Remote Activity
6. 7.6 SSH and Linux Remote Access
7. 7.7 Remote Monitoring and Management Tools
8. 7.8 Source, Destination and Identity Correlation
9. 7.9 Attribution Limits, Shared Accounts and Missing Records
10. 7.10 Practical: Reconstruct Movement Across Three Hosts

Persistence and Redundant Access

Module 8 of Adversarial Tradecraft for Defenders: what operators leave behind so they can return after the first foothold is removed, the privilege each mechanism needs, the records each leaves, and how to verify that removal worked.

1. 8.1 Persistence Objectives and Required Privileges
2. 8.2 Registry and User-Level Mechanisms
3. 8.3 Scheduled Tasks and Services
4. 8.4 WMI and Event-Driven Mechanisms
5. 8.5 Linux Scheduled Jobs, Services and SSH Access
6. 8.6 Account, Group and Permission Changes
7. 8.7 Application and Cloud-Identity Persistence
8. 8.8 Multiple Mechanisms and Alternative Access Routes
9. 8.9 Removal, Recurrence and Verification
10. 8.10 Practical: Identify the Access Left After Initial Cleanup

Phase 4: Command and Control, Exfiltration and Cloud

Command and Control, Concealment and Telemetry Interference

Module 9 of Adversarial Tradecraft for Defenders: how an operator talks to the machines it controls, how it hides that conversation and its own traces, and what a defender can still see when the content is encrypted, the log is cleared or a sensor is attacked.

1. 9.1 Command-and-Control Objectives and Components
2. 9.2 Controlled Callbacks and Beaconing
3. 9.3 Encrypted Connections and Metadata Limitations
4. 9.4 Proxies, Relays and Trusted-Service Abuse
5. 9.5 Native Tools and Legitimate Traffic Lookalikes
6. 9.6 Obfuscation and Encoded Content Interpretation
7. 9.7 Logging Changes and Event-Log Clearing
8. 9.8 Security-Control Interference and Health Signals
9. 9.9 Distinguishing Concealment from Ordinary Maintenance
10. 9.10 Practical: Explain a Communications Chain and Its Visibility Gaps

Collection, Exfiltration and Impact

Module 10 of Adversarial Tradecraft for Defenders: what the operator collected, how it staged and sent it, what the records support about what left, and how to scope impact without equating a note or encryption with the whole story.

1. 10.1 Identifying Data and Services of Value
2. 10.2 File Discovery and Collection
3. 10.3 Staging, Compression and Temporary Storage
4. 10.4 Synthetic-Data Transfer to a Controlled Destination
5. 10.5 Cloud Storage and Approved-Service Abuse
6. 10.6 Transfer Attempts versus Completed Transfers
7. 10.7 Network Volume versus Identified Data Exposure
8. 10.8 Non-Destructive Ransomware and Disruption Simulation
9. 10.9 Scoping Impact without Equating Encryption with Compromise
10. 10.10 Practical: Produce an Exposure and Impact Assessment

Cloud and Hybrid Tradecraft

Module 11 of Adversarial Tradecraft for Defenders: following the operation beyond the endpoint into a cloud identity, an application it was granted, and the files it read, from the tenant's own logs.

1. 11.1 Cloud Control Planes and Shared Responsibility
2. 11.2 User and Workload Authentication
3. 11.3 Exposed Credentials and Excessive Permissions
4. 11.4 Application Consent and Service Principals
5. 11.5 Cloud Resource and Storage Discovery
6. 11.6 Hybrid Identity Dependencies and Sensitive Infrastructure
7. 11.7 Correlating Endpoint Activity with Cloud Audit Records
8. 11.8 Session Revocation and Remaining Access
9. 11.9 Cloud Containment and Recovery Verification
10. 11.10 Practical: Follow an Endpoint Foothold into a Cloud Application

Phase 5: Defensive Intervention

Defensive Intervention and Detection Improvement

Module 12 of Adversarial Tradecraft for Defenders: choosing where to interrupt the operation, building and testing the detections that would have seen it, sequencing containment, and recording what remains.

1. 12.1 Prevention, Observation, Detection and Response
2. 12.2 Selecting High-Value Intervention Points
3. 12.3 Telemetry Requirements and Coverage Checks
4. 12.4 Behavioral Detections and Tool-Specific Indicators
5. 12.5 Positive Tests and Legitimate Comparison Cases
6. 12.6 Narrow Exclusions and Lost-Coverage Risks
7. 12.7 Containment Sequencing and Business Dependencies
8. 12.8 Attacker Adaptation and Remaining Paths
9. 12.9 Verifying Repairs and Recording Residual Exposure
10. 12.10 Practical: Interrupt a Campaign and Test an Alternative Route

Phase 0: Course Resources

Cheatsheets

One sheet per module: the question, the lesson's query on the campaign, what to read and what it does not prove.

1. Threat Reports and Evidence Plans
2. Reconnaissance and Initial Access
3. Execution and Living off the Land
4. Discovery and Attack Paths
5. Credentials, Authentication and Sessions
6. Privilege Escalation
7. Lateral Movement
8. Persistence
9. Command and Control and Concealment
10. Collection, Exfiltration and Impact
11. Cloud and Hybrid Identity
12. Defensive Intervention
13. Lookup Values

Cookbooks

Runbooks for the course's procedures: each step with its reason, the query, what it produces and where it branches.

1. From an Advisory to an Evidence Plan
2. Reconstructing an Entry Route
3. Rebuilding Execution Inside Administrative Noise
4. Validating an Attack Path
5. Tracing One Identity Across Sources
6. Joining a Hop from Both Ends
7. A Persistence Inventory and Its Removal
8. An Exposure and Impact Assessment

Walkthroughs

Four cases worked end to end on public recordings, wrong turns included.

1. A Screensaver That Opened a Shell
2. Two Ways onto a Second Machine
3. A Block That Was Only Sometimes a Block
4. From a Decode to the Domain Controller

Playbooks

One per thing that fires: pre-flight, the clearing check first, the escalation signals, containment and the record.

1. A Preventive Control Blocked Something
2. Encoded PowerShell from a Task or Office
3. A Service Installed from a Workstation
4. An RC4 Service Ticket Request
5. A Single-Factor Sign-in After a Spray
6. A Broad Delegated Consent
7. A Machine-Paced Bulk Read
8. An Upload to a New Destination
9. A Log Cleared or a Sensor Stopped
10. sudo Started a Shell

Practice Playground

The two workspaces, a first query on each, and every practice task by module.

1. The Two Workspaces and Every Practice

References and Reading

Every source the lessons were written from, by module.

1. References by Module

Course Completion

Course Exam

Adversarial Tradecraft for Defenders end-of-course exam: 40 questions in four parts, from threat reports to the cloud, ending on a case study.

1. Course Completion. Adversarial Tradecraft for Defenders

Generated 2026-11-26 from the published course. The current syllabus is always at ridgelinecyber.com/training/courses/adversarial-tradecraft-for-defenders/