

Velociraptor for Endpoint Investigation

FOR202

Deploy Velociraptor and use it well: what it reaches and what it never will, the query language, collection that survives a reboot, hunts across a fleet, correlation into a defensible account, and the failure modes the console cannot show you.

Modules	15
Lessons	83
Phases	4
Free preview	Course Orientation (7 lessons)

The method

Velociraptor turns a question into a query that runs on every endpoint at once. This course runs the same loop for every investigation it works.

1. Deploy it so the answers are trustworthy. Client authentication and access control first, because a collection platform anybody can query is a liability rather than a capability.
2. Ask the question in VQL. Artifacts are queries. Learning to write them is what separates running somebody else's collection from running your own investigation.
3. Prove it on one endpoint. A query that works on one host and returns nothing across the fleet is a query with a bug. Establish the shape of a correct answer before scaling it.
4. Hunt, then read the denominator. A fleet hunt reports what answered. How many clients were enrolled, how many responded and how many errored is the part that makes the result mean something.
5. Turn the finding into a standing artifact. The investigation you just ran is worth more as a saved artifact somebody can re-run than as a result in a notebook.

What this course assumes

No minimum experience and no prerequisite course. VQL, the client-server model and the artifact system are built from nothing.

What makes it go faster: two machines to run a server and a client on, virtual is fine, and comfort with SQL-like syntax. Neither is required, and every query in the course is shown with its output.

What this course does not cover: Windows artifact forensics in depth, malware analysis, and incident response process. Velociraptor is how you reach the evidence at scale, and those decide what to do with it.

Full syllabus

Phase 1: The Deployment

Course Orientation FREE PREVIEW

What Velociraptor for Endpoint Investigation teaches: reaching a machine you cannot walk to, asking it a specific question, and turning the answer into evidence.

1. 0.1 What Remote Live Collection Is
2. 0.2 How This Course Is Structured
3. 0.3 The Environment You Will Build
4. 0.4 Where This Shows Up
5. 0.5 What It Cannot Do
6. 0.6 Following Along
7. Module Summary

Deployment and Access Control

What Velociraptor reaches that your other tools do not, the architecture that makes a collection arrive in seconds, choosing a deployment model before you generate a config, standing the server up on a current release, the roles and organizations that decide who can run a query against whose...

1. 1.1 What Velociraptor Reaches, and What It Does Not
2. 1.2 The Architecture Behind a Two-Second Answer
3. 1.3 Choosing a Deployment Model
4. 1.4 Standing Up the Server
5. 1.5 Users, Roles, and Organizations
6. 1.6 Enrolling Clients and Proving the Fleet Is Real
7. Module Summary
8. Check My Knowledge

Driving the Console

Finding a host among hundreds and reading what its record actually claims, why the file browser is a cache rather than a live view and what that means for every conclusion drawn from it, retrieving and inspecting a file without touching the machine, running and reading collections, notebooks as...

1. 2.1 Finding a Host and Reading Its Record
2. 2.2 The Virtual File System, and Why It Is a Cache
3. 2.3 Retrieving and Inspecting a File
4. 2.4 Collections, from Launch to Result
5. 2.5 Notebooks as the Investigation Workspace
6. 2.6 Labels and Metadata
7. Module Summary
8. Check My Knowledge

Phase 2: Asking Questions

VQL for Investigators

The query language underneath every artifact you have collected so far: how a query is shaped, the difference between the things that produce rows and the things that transform them, stored queries and when they materialize, the evaluation order that decides whether a filter protects you, and...

1. 3.1 The Shape of a Query
2. 3.2 Row Sources
3. 3.3 Transforming What Comes Back
4. 3.4 Stored Queries and Composition
5. 3.5 Evaluation Order, and the Trap
6. 3.6 When a Query Fails
7. Module Summary
8. Check My Knowledge

Collecting from One Endpoint

Turning the console and the language into a collection you can defend: the order volatile evidence has to be taken in, choosing artifacts for a question rather than from a checklist, scoping with parameters, building an offline collector for machines with no agent, handling a collection that...

1. 4.1 Order of Volatility
2. 4.2 Choosing Artifacts for a Question
3. 4.3 Scoping with Parameters
4. 4.4 The Offline Collector
5. 4.5 When a Collection Fails Partway
6. 4.6 The Custody Record
7. Module Summary
8. Check My Knowledge

Phase 3: The Fleet

Hunts Across the Fleet

The collection you already know how to build, addressed to every machine at once: why a hunt expires rather than completes, scoping by label and condition, launching from the paused state with an expiry and limits that mean something, watching one run against a fleet that is never all online...

1. 5.1 What a Hunt Is, and Why It Never Completes
2. 5.2 Scoping a Hunt
3. 5.3 Launching Safely
4. 5.4 Watching a Hunt Run
5. 5.5 Incremental Hunting
6. 5.6 Reading Fleet Results
7. Module Summary
8. Check My Knowledge

Fleet Analysis and Correlation

Turning separate results into one account: correlating findings across machines, baselining an estate so abnormal becomes visible, building a timeline that spans hosts, handling clocks that disagree, pivoting from a finding to the next question, and writing the sequence in a form somebody else...

1. 6.1 From Samples to a Sequence
2. 6.2 Baselining and Differential Analysis
3. 6.3 Building a Timeline Across Machines
4. 6.4 When Clocks Disagree
5. 6.5 Pivoting
6. 6.6 Writing the Account
7. Module Summary
8. Check My Knowledge

Phase 4: Extending and Operating

Extending and Operating It

Writing your own artifact when nothing published answers the question, reading and trusting community definitions, monitoring instead of asking, containing a host, keeping the deployment healthy, and taking one case from alert to written account using everything in the six modules behind it.

1. 7.1 Writing Your First Artifact
2. 7.2 Artifacts Worth Borrowing
3. 7.3 Monitoring Instead of Asking
4. 7.4 Containment
5. 7.5 Operating the Deployment
6. 7.6 One Case, End to End
7. Module Summary
8. Check My Knowledge

Phase 0: Course Resources

Cheatsheets

Queries and commands by the question you arrived with, each with the fields that carry the decision and what the result does not establish.

1. Query Patterns
2. First Response
3. Hunting the Fleet
4. Correlation and Timeline
5. Operating It
6. States, Limits and Lookups

Cookbooks

Six runbooks for the procedures this work repeats: triaging an alerted host, hunting an indicator across the estate, collecting from a machine with no agent, containing one, promoting a query to an artifact, and handing evidence to somebody else.

1. Triaging an Alerted Host
2. Hunting an Indicator Across the Estate
3. Collecting from a Machine with No Agent
4. Containing a Host
5. Promoting a Query to an Artifact
6. Handing Evidence to Somebody Else

Walkthroughs

Endpoint investigations reasoned end to end, including the cases where the tool was the wrong answer and the ones that finish with a weaker finding than the evidence first suggested.

1. The Alert That Was the Last Host
2. The Hunt That Found Nothing
3. The Collection That Said Finished
4. The Five Minutes That Were Not There
5. The Task on Six Hundred Machines
6. The Laptop That Should Have Been Imaged

Playground

Every practice surface available for this course, what each one gives you, where the gaps are, and the exercises worth building yourself.

Operational Reference

Every command, query, artifact invocation, state and check from the Velociraptor for Endpoint Investigation course in one place, organized by task and linked back to the section that explains when not to use it.

1. Operational Reference

References & Further Reading

Vendor documentation, the artifact reference and community exchange, endpoint forensics sources, incident response standards and frameworks, and the operational research used throughout the Velociraptor for Endpoint Investigation course.

Course Completion

Course Exam

Velociraptor for Endpoint Investigation end-of-course exam: fleet scoping inside the hour during a live intrusion, testing whether you can match a hunt design to the question, cost it before launching, and know what complete means when a fifth of the fleet is offline.

1. Course Completion. Velociraptor for Endpoint Investigation

Generated 2026-09-01 from the published course. The current syllabus is always at ridgelinecyber.com/training/courses/velociraptor-endpoint-investigation/