DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Graylog Advanced Training

Scale Graylog into a resilient multi-node platform, upgrade it safely and govern retention, security and performance.

LIVE ONLINE **Max 8-10 hours**	CORPORATE (DAY-WISE) **5 days · Days 1-5**	SELF-PACED VIDEO **1-year LMS access**
BATCH SIZE **Max 10 learners**	NEXT BATCHES **1st of every month · weekends: first Saturday**	CERTIFICATE **Verifiable certificate of completion**

Foundation	Essential	Intermediate	Advanced
1 day · Included	2 days · Included	3 days · Included	5 days · This curriculum

INSIDE THIS CURRICULUM

What you will achieve	Outcomes and level focus
Who should attend	Audience and prerequisites
Agenda	48 topics · Foundation to Advanced
Hands-on labs	4 labs
Formats and fees	Live, 1:1, self-paced, corporate
Trainer and reviews	Rajesh Kumar
Enrolment	Batches, payment, contacts

WHY DEVOPSSCHOOL

- Rated 4.6 on Google
- 1,500+ verifiable certificates issued in the public registry
- Trainer Rajesh Kumar: 20 years in industry
- Small batches: max 10 learners, every session recorded

Course page: www.devopsschool.com/courses/graylog/graylog-advance-training.html
Talk to us: WhatsApp +91 99057 40781 · +1 (469) 756-6329 · contact@DevOpsSchool.com · reply within 1 hour

OUTCOMES

What you will achieve

By the end of the Graylog course path, you will be able to:

1

Centralize log collection

Collect logs from servers, network devices and applications into Graylog with inputs and Sidecar.

2

Search, route and alert

Structure, route, enrich and visualize log data and alert on important events.

3

Run Graylog at scale

Design, secure, upgrade and tune multi-node Graylog clusters with clear retention and governance.

WHAT THE ADVANCED LEVEL ADDS

Scaling, resilience, upgrades, tiering and governance

Hands-on: Build a resilient Graylog cluster. Scale to three Graylog nodes with a MongoDB replica set and OpenSearch cluster, tune journals and buffers, and prove recovery from a node failure.

WHO SHOULD ATTEND

- Experienced Graylog practitioners responsible for quality, security and scale.
- Senior engineers, architects and leads who design and troubleshoot Graylog solutions.
- Teams making Graylog an organisation-wide standard.

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Transform and enrich messages with processing pipelines and lookup tables, secure Graylog and integrate it with other tools.
- Basic computer skills and a laptop with a stable internet connection for the live sessions and labs.
- This curriculum includes every earlier level, so it is complete on its own.

AGENDA · FOUNDATION TO ADVANCED

Graylog Advanced Training agenda

48 topics. Each level builds on the one before it, so this curriculum lists every earlier level in full. Days are for corporate (day-wise) delivery; live online batches cover the same agenda in max 8-10 hours.

DAY 1 — Foundation: Log management concepts, installation and first search

Included from Foundation

1. Introduction to Graylog: centralized log management and the Graylog Open, Enterprise and Security editions
2. Architectural considerations: Graylog server, MongoDB and OpenSearch or Graylog Data Node roles
3. Getting started: how messages flow from inputs through streams and pipelines into index sets
4. Choosing your installation method: packages on Ubuntu or RHEL, Docker Compose and cloud images
5. Installing Graylog: MongoDB, Data Node or OpenSearch and the graylog-server package
6. Configuring Graylog: server.conf, password_secret, root_password_sha2 and http_bind_address
7. Sending in log data: Syslog, GELF, Beats and raw inputs and which to choose
8. Creation of inputs: launching global and node inputs, ports and checking received messages
9. Searching: query syntax, field searches, time ranges and saved searches
10. Searching and analyzing your data: field statistics, top values and message details
11. First dashboard: building a dashboard from a search with message count and top values widgets

DAY 2 — Essential: Collectors, extractors, streams, dashboards and alerts

Included from Essential

12. Graylog Collector Sidecar: installing sidecars and managing Filebeat, Winlogbeat and NXLog configuration centrally
13. Graylog Collector (deprecated): moving older collector setups to Sidecar
14. Creation of extractors: regex, grok, JSON, split-and-index and key-value extractors
15. Usage of streams to route messages: stream rules, matching modes and removing messages from the default stream
16. Index sets: index rotation by time or size, retention strategies, shards and replicas
17. Enhancing dashboards and adding widgets: aggregations, charts, tables and dashboard pages
18. Search workflows: search parameters, highlighting and exporting results as CSV
19. Creation of alerts and conditions: event definitions with filter and aggregation conditions
20. Alerts: notifications via email, Slack, HTTP and PagerDuty and alert event details
21. Users and roles: built-in roles, entity sharing and team permissions
22. Field types: field type mappings and custom index field types
23. Content packs: exporting and importing inputs, streams, extractors and dashboards

DAY 3 — Intermediate: Processing pipelines, enrichment, security and integrations

Included from Intermediate

24. Processing pipelines: rules, stages and connecting pipelines to streams
25. Pipeline use cases: flexibility in routing, blacklisting, modifying and enriching messages with processing pipelines
26. Pipeline functions: parse_json, grok, key_value, set_field, route_to_stream and drop_message
27. Blacklisting: dropping noisy messages with pipeline rules to cut storage and noise
28. Message rewriting with Drools: the removed Drools engine and its pipeline-rule replacement
29. Geolocation: GeoIP lookups with MaxMind databases and world map widgets

30. Lookup tables: data adapters such as CSV, DSV over HTTP and threat intel, caches and enrichment
31. Securing Graylog: TLS for web, API and inputs, certificates and LDAP or Active Directory authentication
32. Indexer failures: diagnosing mapping conflicts and failed indexing on the indexer failures page
33. REST API: the API browser, access tokens and automating configuration
34. Plugins: installing plugins, the Graylog Marketplace and community integrations
35. External dashboards: visualizing Graylog data in Grafana and sharing outside Graylog
36. Troubleshooting: server.log, the message journal and process and output buffer metrics

DAYS 4-5 **Advanced: Scaling, resilience, upgrades, tiering and governance** — New at this level

37. Scaling and building resilient environments: multi-node Graylog clusters behind a load balancer
38. Backend clusters: MongoDB replica sets and OpenSearch or Data Node clusters with sizing, shards and heap
39. Message journal and buffers: disk journal sizing, processor counts and throughput tuning
40. Upgrading Graylog: upgrade paths across major versions, MongoDB and OpenSearch compatibility and backups
41. Data tiering and archiving: hot and warm tiers, archiving and restoring indices
42. Graylog Security: Sigma rules, investigations and anomaly detection in the Enterprise editions
43. Ingest at scale: Graylog Forwarder, load-balanced syslog and Kafka inputs
44. Monitoring Graylog: internal metrics, the Prometheus exporter, health checks and backlog alerts
45. Backup and disaster recovery: MongoDB backups, OpenSearch snapshots and configuration export
46. Governance: audit logging, retention policies, PII handling and access reviews
47. Performance tuning: index sizing, field limits, search performance and JVM settings
48. Capstone: a resilient multi-node Graylog platform ingesting Linux, Windows and application logs with alerting

PRACTICE

Hands-on labs

You build these yourself during the course, guided by the trainer.

FOUNDATION **Install Graylog and search logs**
Deploy Graylog with MongoDB and a Data Node using Docker Compose, send syslog and GELF messages and investigate them with searches.

ESSENTIAL **Collect, route and alert on logs**
Ship Linux and Windows logs with Sidecar, extract fields, route them into streams with their own index sets and alert on failed logins.

INTERMEDIATE **Enrich logs with processing pipelines**
Build pipelines that parse web logs, add GeoIP and threat-intel lookups, drop noise and route security events to a dedicated stream over TLS inputs.

ADVANCED **Build a resilient Graylog cluster**
Scale to three Graylog nodes with a MongoDB replica set and OpenSearch cluster, tune journals and buffers, and prove recovery from a node failure.

HOW YOU CAN LEARN

Formats and fees

The same agenda in four formats. Fees are in Indian rupees as of November 2026; the course page shows them in your currency.

FORMAT	DURATION	WHAT YOU GET	FEE
Live & interactive (online)	Max 8-10 hours	Instructor-led public batch of max 10 learners with demos, labs, Q&A and assignment review. Weekday batches start on the 1st of every month, weekend batches on the first Saturday of every month.	**INR 34,999**
1:1 live training (online)	Max 8-10 hours, one-to-one	The same agenda, taught one-to-one by the trainer. Starts immediately, without waiting for the next batch; sessions planned around your working hours, at your pace.	**INR 99,999**
Self-paced video	1-year LMS access	All four Graylog levels, plus 20+ more courses and 50+ tools, at your own pace.	**INR 833/month** billed yearly at INR 9,996
Corporate (day-wise)	5 days for this level	A private batch for your team, on demand. Foundation 1 day, Essential 2, Intermediate 3, Advanced 5 (cumulative).	**Quote on request**

ASSESSMENT AND CERTIFICATE

Your progress is checked through hands-on labs and assignments at every level. On completion you receive a DevOpsSchool certificate of completion with a unique credential ID that anyone can verify at devopsschool.com/certificates.

INCLUDED WITH EVERY LIVE BATCH

- Every live session is recorded and stays in your LMS for a year.
- Can't attend your batch? Move to next month's batch.
- Free demo class before you pay.
- Free 10-15 minute call with an advisor or the trainer to check the course fits you.

YOUR TRAINER

Learn from Rajesh Kumar

Principal Architect, trainer and mentor in DevOps, SRE, DevSecOps and platform engineering, with 20 years across 8+ software MNCs, including Staff DevOps/SRE Engineer for a leading Japanese mobile-payment platform and Principal DevOps Architect and Manager in Europe.

Profile: www.rajeshkumar.xyz · Videos: youtube.com/@TheDevOpsSchool

20 years in industry	**200+** corporate trainings
10,000+ engineers mentored	**70+** organisations coached

REVIEWS

What learners say

Rated 4.6 on Google · more reviews at devopsschool.com/testimonials and in our video testimonials.

> “Rajesh is very good trainer. Rajesh was able to resolve our queries and question effectively. We really liked the hands-on examples covered during this training program.”
>
> Indrayani, India

> “The training was very useful and interactive. Rajesh helped develop the confidence of all.”
>
> Abhinav Gupta, Pune

START YOUR GRAYLOG TRAINING

Join the next batch

Weekday batches start on the 1st of every month and weekend batches on the first Saturday of every month, with max 10 learners each. 1:1 training can start immediately; corporate batches run on demand. We reply within 1 hour.

Pay by UPI, credit and debit cards, net banking or PayPal. Before you pay, read our refund policy. All levels of this course: www.devopsschool.com/courses/graylog/

ENROL OR RESERVE A SEAT
Graylog Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com