DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Fluentd Advanced Training

Design and tune Fluentd for high-volume, highly available logging pipelines with governance, cost control and safe migrations.

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/fluentd/fluentd-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 Fluentd course path, you will be able to:

1

Build unified logging pipelines

Collect, parse, enrich and route logs from servers and applications with Fluentd.

2

Centralise Kubernetes logs

Run Fluentd and Fluent Bit on Kubernetes with metadata, secure forwarding and per-team routing.

3

Operate logging at scale

Tune, monitor and govern highly available Fluentd pipelines for high-volume production use.

WHAT THE ADVANCED LEVEL ADDS

Performance, high availability, scale and governance

Hands-on: Load-test a highly available pipeline. Run multi-worker Fluentd aggregators behind Kafka, generate high log volume, kill an aggregator and prove no logs are lost.

WHO SHOULD ATTEND

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

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Collect Kubernetes logs with Fluentd, build forwarder and aggregator tiers, secure the transport and troubleshoot delivery.
- 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

Fluentd 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: Unified logging concepts, installation and first pipeline

Included from Foundation

1. Logging fundamentals: why centralised logging, structured versus unstructured logs and the unified logging layer
2. What Fluentd is: CNCF graduated log collector written in Ruby and C, plugin ecosystem and common use cases
3. Fluentd versus Fluent Bit: footprint, plugin coverage and the forwarder and aggregator split
4. Event model: tag, time and record, and how events flow from input to output
5. Installation: fluent-package (successor to td-agent) on Linux, the Docker image and gem install
6. Service management: systemctl for fluentd, default config path /etc/fluent/fluentd.conf and log files
7. Configuration directives: source, match, filter, system, label and @include
8. Input plugins: in_tail for files, in_forward, in_http and in_syslog
9. Output plugins: out_stdout, out_file and out_forward
10. Tag routing: match patterns with single and double wildcards, brace alternatives and match order
11. Parsing basics: parse sections with json, regexp, apache2, nginx and none parsers
12. First example: tail an application log, parse it as JSON and write it to stdout and a file

DAY 2 — Essential: Filters, buffering, outputs and log enrichment

Included from Essential

13. Filter plugins: grep for include and exclude, record_transformer for adding and removing fields
14. Parser and formatter plugins: filter_parser, format sections, ltsv, csv and single_value
15. Multiline logs: multiline parser for stack traces, format_firstline and concat plugin
16. Time handling: time_key, time_format, time zones and keeping original timestamps
17. Buffering: memory and file buffers, chunk_limit_size, total_limit_size and flush_interval
18. Retries and errors: retry_wait, retry_max_interval, retry_forever, secondary outputs and @ERROR label
19. Output destinations: Elasticsearch and OpenSearch plugins, S3, Kafka (out_kafka2) and Loki
20. Plugin management: fluent-gem install, plugin versions and checking with fluentd --dry-run
21. Labels and relabel: grouping pipelines with @label and the relabel output
22. Enrichment: adding hostname, environment and service fields, geoip lookups
23. Copy and fan-out: out_copy to send events to several destinations
24. Team workflow: config in Git, @include structure and environment-specific values

DAY 3 — Intermediate: Kubernetes logging, aggregation, security and troubleshooting

Included from Intermediate

25. Kubernetes logging architecture: node-level agents, container log paths and CRI log format
26. Fluentd DaemonSet: fluentd-kubernetes-daemonset images, RBAC, volumes and resource limits
27. Kubernetes metadata: kubernetes_metadata filter for pod, namespace, labels and annotations
28. Forwarder and aggregator topology: Fluent Bit or Fluentd agents forwarding to Fluentd aggregators
29. Forward protocol security: TLS transport, shared_key and user authentication
30. Load balancing and failover: multiple servers in out_forward, weights, standby servers and heartbeat

31. Operators: Fluent Operator and Logging operator (kube-logging) custom resources
32. Monitoring Fluentd: monitor_agent, Prometheus plugin metrics, buffer queue length and retry counts
33. Log routing per team: tag rewriting with rewrite_tag_filter and namespace-based outputs
34. Masking sensitive data: removing or hashing PII, tokens and passwords before shipping
35. Troubleshooting: buffer overflow, chunk errors, unmatched tags, parse failures and plugin errors
36. Debugging techniques: log_level, out_stdout taps, fluent-cat and dry runs

DAYS 4-5 **Advanced: Performance, high availability, scale and governance** **New at this level**

37. Performance tuning: multi-process workers, worker-specific directives, flush_thread_count and chunk sizing
38. Throughput and backpressure: overflow_action options, rate limits and protecting downstream stores
39. High availability: redundant aggregators, persistent file buffers, at-least-once delivery and acknowledgements
40. Scaling patterns: horizontal aggregator scaling, Kafka as a durable buffer and partitioning by tag
41. Writing custom plugins: Ruby plugin structure for input, filter and output, plugin tests with test drivers
42. Observability of the pipeline: SLOs for log delivery latency and loss, alerting on retries and buffer growth
43. Security hardening: least privilege on hosts and pods, secret handling, TLS everywhere and plugin supply chain checks
44. Governance and compliance: retention policies, audit logs, data residency and PII redaction standards
45. Cost control: sampling, dropping noisy logs, routing cold data to object storage
46. OpenTelemetry interoperability: OTLP pipelines, OpenTelemetry Collector alongside Fluentd and Fluent Bit
47. Migration: moving from td-agent to fluent-package, and from Logstash to Fluentd or Fluent Bit
48. Capstone: design and run a highly available, multi-cluster logging platform with tested failover

PRACTICE

Hands-on labs

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

FOUNDATION **Build your first Fluentd pipeline**
Install fluent-package, tail an NGINX access log, parse each line and route the events to stdout and a rotated file.

ESSENTIAL **Ship enriched logs to OpenSearch**
Parse multiline application logs, enrich them with environment fields, buffer to disk and deliver them to OpenSearch and S3 with retries.

INTERMEDIATE **Centralise Kubernetes logs securely**
Deploy Fluent Bit agents and a Fluentd aggregator with TLS forward, enrich logs with Kubernetes metadata and route each namespace to its own index.

ADVANCED **Load-test a highly available pipeline**
Run multi-worker Fluentd aggregators behind Kafka, generate high log volume, kill an aggregator and prove no logs are lost.

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 Fluentd 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 FLUENTD 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/fluentd/

ENROL OR RESERVE A SEAT
Fluentd Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com