DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Site Reliability Engineering SRE Advanced Training

Operate multi-cluster platforms with long-term telemetry, self-healing automation, cost control and reliability governance.

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	45 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/sre/site-reliability-engineering-sre-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 Site Reliability Engineering SRE course path, you will be able to:

1

Build an observability stack

Instrument services with Prometheus, Grafana, OpenTelemetry and central logging tied to SLOs.

2

Engineer resilient systems

Use Kubernetes controls, progressive delivery, load tests and chaos experiments to prove resilience.

3

Automate reliability at scale

Run multi-region platforms with self-healing automation, governance and cost awareness.

WHAT THE ADVANCED LEVEL ADDS

Multi-region scale, automation and governance

Hands-on: Regional failover drill. Run a service across two clusters with Thanos for global metrics, fail one region over with automated remediation and report SLO impact and recovery time.

WHO SHOULD ATTEND

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

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Harden Kubernetes workloads, ship with progressive delivery and prove resilience with load and chaos experiments.
- 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

Site Reliability Engineering SRE Advanced Training agenda

45 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: SRE foundations and the observability toolchain

Included from Foundation

1. SRE fundamentals: principles, SRE roles and how SRE puts DevOps into practice
2. Linux and networking for SRE: processes, systemd, file systems, DNS, TCP and HTTP
3. Cloud and containers recap: virtual machines, Docker images and the Kubernetes objects SREs operate
4. Golden signals in practice: latency, traffic, errors and saturation
5. SLIs, SLOs and error budgets: the vocabulary and simple calculations
6. Metrics with Prometheus: exporters, scraping, time series and labels
7. Dashboards with Grafana: panels, variables and a first service dashboard
8. Logs: structured logging, log levels and central collection with Loki or Elasticsearch
9. Health checks: liveness and readiness probes and synthetic checks with Blackbox Exporter
10. Lab environment: a Kubernetes cluster with Prometheus and Grafana from kube-prometheus-stack

DAY 2 — Essential: PromQL, SLO alerting, tracing and incidents

Included from Essential

11. PromQL: rate, increase, histogram_quantile, aggregations and recording rules
12. SLO as code: defining SLOs with Sloth or Pyrra and generating Prometheus rules
13. Alertmanager: routing, grouping, inhibition, silences and receivers
14. Burn rate alerts: fast and slow burn alerts tied to error budgets
15. Distributed tracing: OpenTelemetry SDKs, the Collector, spans and context propagation
16. Trace backends: Jaeger or Grafana Tempo and finding latency hotspots
17. On-call tooling: PagerDuty, incident.io or Jira Service Management schedules and escalations
18. Runbooks and automation: linking alerts to runbooks and scripted remediation
19. Incident workflow: severity, roles, chat incident channels and status pages
20. Postmortems: timelines from telemetry, root cause analysis and tracked actions
21. CI/CD for reliability: deployment pipelines, smoke tests and automatic rollback

DAY 3 — Intermediate: Kubernetes resilience, delivery, chaos and IaC

Included from Intermediate

22. Kubernetes resilience: requests and limits, HPA, PodDisruptionBudgets and topology spread
23. Progressive delivery: Argo Rollouts canaries with Prometheus analysis
24. Resilience patterns: timeouts, retries with backoff, circuit breakers and bulkheads
25. Load testing: k6 or Locust scenarios, capacity models and headroom
26. Chaos experiments: Chaos Mesh or LitmusChaos faults and structured game days
27. Infrastructure as code: Terraform modules, GitOps with Argo CD and drift detection
28. Fleet configuration: Ansible for configuration changes and patching
29. Backup and restore: Velero for Kubernetes, database backups and restore drills
30. Observability pipelines: OpenTelemetry Collector processors, sampling and cardinality control

31. Log analytics: LogQL queries and correlating logs, metrics and traces
32. Security for SREs: Vault secrets, Kubernetes RBAC and image scanning in pipelines
33. Troubleshooting methodology: hypothesis-driven debugging with kubectl, strace, tcpdump and profilers

DAYS 4-5 **Advanced: Multi-region scale, automation and governance** — **New at this level**

34. Multi-cluster and multi-region: failover design, global load balancing and data replication
35. Long-term metrics: Thanos or Grafana Mimir, retention and global views
36. eBPF observability: continuous profiling with Pyroscope or Parca and Cilium Hubble network visibility
37. Self-healing automation: Kubernetes operators, event-driven remediation and safety limits
38. Platform reliability: internal developer platforms, golden paths and reliability guardrails
39. AIOps: anomaly detection, alert correlation and AI-assisted incident analysis
40. Cost and performance: FinOps, right-sizing and node autoscaling with Karpenter or Cluster Autoscaler
41. Disaster recovery at scale: DR runbooks, regional evacuation tests and recovery objectives
42. Reliability governance: production readiness and SLO reviews and error budget policies across teams
43. Policy as code: OPA Gatekeeper or Kyverno guardrails for reliability settings
44. Engineering metrics: DORA metrics, MTTR trends and toil tracking
45. Capstone: run a reliable microservices platform with SLOs, alerts, chaos tests and a DR drill

PRACTICE

Hands-on labs

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

FOUNDATION

Golden signals dashboard on Kubernetes

Install kube-prometheus-stack on a lab cluster, deploy a demo service and build a Grafana dashboard for latency, traffic, errors and saturation.

ESSENTIAL

SLO alerts and traces for a service

Define an SLO with Sloth, route burn rate alerts through Alertmanager and use OpenTelemetry traces in Tempo to find the cause of a latency regression.

INTERMEDIATE

Canary release under chaos

Ship a new version with an Argo Rollouts canary gated by Prometheus analysis while Chaos Mesh injects pod and network faults, and confirm the SLO holds.

ADVANCED

Regional failover drill

Run a service across two clusters with Thanos for global metrics, fail one region over with automated remediation and report SLO impact and recovery time.

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 Site Reliability Engineering SRE 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 SITE RELIABILITY ENGINEERING SRE 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/sre/

ENROL OR RESERVE A SEAT
Site Reliability Engineering SRE Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com