+33 (0)1 48 10 09 46
formation@ambient-it.net
140 Avenue Jean Lolive 93500 Pantin - France

Updated on 10/26/2026

Sign Up

# VictoriaMetrics Training

3 days (21 hours)

## Overview

VictoriaMetrics is a metrics storage and querying solution for modern environments. Combined with VictoriaLogs and VictoriaTraces, it enables you to build an architecture that covers all three telemetry signals. With OpenTelemetry, you have the components needed to deploy unified pipelines, from collection to visualization.

Our VictoriaMetrics training will enable you to design and deploy telemetry pipelines tailored to the needs of DevOps and SRE teams, combining shared collection with specialized storage.

You'll learn how to configure the OpenTelemetry Collector and vmagent, normalize attributes, filter data, and route each signal to its destination. You'll leverage metrics, logs, and traces to analyze the behavior of your services.

By the end of the training, you'll be able to deploy this architecture on Kubernetes, manage cardinality, secure access, and size queues to better handle service outages.

The workshops will guide you through building a complete pipeline that integrates Grafana, alerting, and collector monitoring. You'll validate its resilience through simulated incidents and formalize the procedures necessary for its operation.

Like all our training courses, this one will introduce you to **the latest stable version** of the technology and its new features.

## Objectives

- Design a unified telemetry architecture with separate storage for metrics, logs, and traces
- Configure OpenTelemetry Collector pipelines and metric collection with vmagent
- Standardize attributes, manage cardinality, and filter sensitive data before exporting it
- Deploy VictoriaMetrics ecosystem components on Kubernetes
- Correlate signals in Grafana and set up operational alerts
- Test the pipeline’s resilience and diagnose telemetry delays or losses

# Target Audience

- DevOps Engineer
- Site Reliability Engineer
- Platform Engineer
- System and Cloud Administrator
- Observability Engineer

# Prerequisites

- Practical experience with Linux administration and common network commands
- Experience with containers using Docker and knowledge of fundamental Kubernetes concepts
- Knowledge of Prometheus metrics, labels, and data collection principles
- Ability to read and edit YAML configuration files

# Technical Prerequisites

- Computer running Linux, macOS, or Windows with WSL 2, equipped with 4 cores, 16 GB of RAM, and 30 GB of available storage for local workshops
- A working installation of Docker Engine or Docker Desktop, with Docker Compose v2 and permissions to run containers
- kubectl, Helm, and Visual Studio Code installed in versions compatible with the educational cluster provided prior to the session
- Access to the educational Kubernetes cluster and a stable internet connection that supports , software repositories, and lab services

# VictoriaMetrics Training Program

## [Day 1 - Morning]

### Architecture of a Unified Telemetry System

- The role of metrics, logs, and traces in service diagnostics
- Division of Responsibilities Between VictoriaMetrics, VictoriaLogs, and VictoriaTraces

- Common collection architecture with OpenTelemetry Collector and specialized storage per signal
- Single-node and cluster deployment models, data persistence, and retention
- Laboratory Setup with Docker Compose, an Instrumented Application, and Data Flow Validation Criteria
- Hands-on Workshop: Deploy the three storage solutions and a demo application, then verify their availability and persistent volumes.

## [Day 1 - Afternoon]

### Deploying OpenTelemetry pipelines

- OpenTelemetry Collector Contrib structure and enabling receivers, processors, and exporters in each pipeline
- Receiving the three signals via OTLP and exporting them to the OTLP HTTP ingestion endpoints specific to
  each storage system
- Standardization of resource attributes and common conventions for identifying services and environments
- Scheduling of processing tasks with `memory_limiter`, attribute enrichment, and batching
- Filtering of unnecessary data and removal of sensitive attributes before export
- Hands-on workshop: Configure metrics, logs, and traces pipelines; harmonize their attributes; and verify that each signal reaches its storage destination.

## [Day 2 - Morning]

### Metrics collection and volume control

- Positioning vmagent as a complement to the Collector and preventing duplicate collections
- Configuring Prometheus scraping, discovering targets, and routing via `remote_write`
- Transforming labels using relabel_configs, metric_relabel_configs, and relabeling rules before export
- Controlling cardinality, removing unnecessary time series, and an introduction to data stream aggregation
- Configuring persistent queues, storage limits, and behavior when the destination is unavailable
- Hands-on workshop: Collecting infrastructure metrics with vmagent, applying
  and monitoring queuing during a storage outage.

## [Day 2 - Afternoon]

### Signal analysis and correlation

- Querying metrics with MetricsQL, rates, aggregations, and latency analysis
- Exploring Logs with LogsQL, Time Filters, Structured Search, and Statistics
- Exploring VictoriaTraces traces in Grafana using the Jaeger data source

- Configuring data sources and correlation links based on service, environment, and trace ID
- Creating rules on metrics with vmalert and routing notifications to Alertmanager
- Hands-on workshop: Build a dashboard, configure a link from logs to traces, and diagnose an application error based on an alert.

### [Day 3 - Morning]

### Scaling on Kubernetes

- Reproducible deployment of components using Helm, value files, and fixed image versions
- Organizing the Collector into agents and gateways, local collection, and centralization of exports
- Persistence management with PVCs, retention policies, and resource allocation
- Securing access with vmauth, TLS encryption, secrets, and network restrictions
- High-availability principles, metric deduplication, and duplicate prevention based on collection paths
- Hands-on Workshop: Deploy the unified pipeline on the training cluster using Helm, configure volumes, and securing an access point with vmauth.

### [Day 3 - Afternoon]

### Pipeline Resilience and Operations

- Monitoring the pipeline using internal metrics from collectors and storage systems
- Diagnosing congested queues, export errors, ingestion failures, and delivery delays
- Configuring retry settings, collector send queues, and persistence based on the selected components
- Load and outage tests to measure throughput, catch-up, and potential data loss
- Formalization of operational procedures, post-update checks, and acceptance criteria of the deployment
- Hands-on workshop: Simulate storage unavailability, restore service, measure data catch-up, and deliver a validation report for the unified pipeline.

## Target Audience

This training is intended for both individuals and companies—large or small—seeking to train their teams in new, advanced IT technologies or to acquire specific domain knowledge or modern methodologies.

## Assessment upon enrollment

The pre-training assessment complies with Qualiopi quality standards. Upon final registration, the learner receives a self-assessment questionnaire that allows us to evaluate their estimated proficiency level across various types of technologies, as well as their expectations and personal goals for the upcoming training, within the limits imposed by the selected format. This

questionnaire also allows us to anticipate certain connection or internal security issues within the company (intra-company or virtual classroom) that could pose problems for the monitoring and smooth running of the training session.

## Teaching Methods

Hands-On Course: 60% Hands-On, 40% Theory. Training materials distributed in digital format to all participants.

## Organization

The course alternates between theoretical presentations by the instructor—supported by examples and reflection sessions—and group work.

## Assessment

At the end of the session, a multiple-choice quiz is administered to verify that participants have successfully acquired the skills.

## Certification

A certificate will be issued to each trainee who has completed the entire training program.


Training organization registered under number 11 75 54743 75. This registration does not constitute state accreditation.
© Ambient IT 2015-2026. All rights reserved. Paris, France - Switzerland - Belgium - Luxembourg