

# AWS vs. GCP vs. Azure vs. OVHcloud: Managed Log Archiving for 7+ Years

2026-05-14

Compare managed log archiving on AWS, GCP, Azure, and OVHcloud: 7-year cost at 100 GB/day, storage tiers, Athena, and BigQuery. Part 1 of 6.

Long-term log archiving at scale is an expensive problem. Cloud providers offer a wide range of managed services for ingesting, storing, and querying logs — and the cost differences between them are dramatic. This series covers **AWS, GCP, Azure, and OVH** along the axes that matter for 7-year archival: services available, storage model, ingest cost, 7-year total cost at 100 GB/day, Kubernetes integration, security, and query interfaces.

- **Part 1 — this post:** Services overview, storage model, and 7-year cost comparison
- [Part 2 — Pre-Flight, Flexibility & Auditor Export](#): Pre-flight checklist, retroactive flexibility, auditor export
- [Part 3 — Operations](#): Onboarding, Kubernetes integration, ingest pipeline, backup & DR
- [Part 4 — Security & Compliance](#): IAM/RBAC, encryption, WORM, GDPR, compliance certificates
- [Part 5 — Query, Dashboards & Recommendations](#): Query interfaces, dashboards, alerting, cold-tier behaviour, when to use which
- [Part 6 — Production Checklist, Guardrails & Runbooks](#): Loss detection, privacy, schema, cost guardrails, runbooks, fire drills

*Self-hosted alternative: [Elasticsearch vs. OpenSearch vs. Loki vs. Quickwit vs. ClickHouse](#) — tiering, compression, resource consumption, SaaS options.*

Retaining logs for seven or more years is a very different problem than storing the last 30 days of operational logs. Ingest cost per GB, the ability to query old data without restoring from tape, and long-term pricing predictability dominate over query latency. Each of the four providers covered here has a fundamentally different answer to this problem — and for at least one of them, the honest answer is “our managed offerings are not designed for this; run it yourself instead.”

## △ Pricing Disclaimer — Read Before Acting

All prices are **public list prices** as of **May 2026** for EU regions (AWS eu-central-1, GCP europe-west3, Azure West Europe, OVH Paris). Every figure is an estimate for one specific scenario (100 GB/day raw ingest, 7× gzip compression, 7-year retention). Negotiated contracts, commitment discounts, and workload specifics will change your actual costs significantly —

always verify on the official pricing pages before any purchasing or architectural decision. The author accepts no responsibility for billing surprises arising from reliance on these numbers.

[AWS CloudWatch](#) · [AWS S3](#) · [GCP Cloud Logging](#) · [GCP BigQuery](#) · [Azure Monitor](#) · [Azure Blob Storage](#) · [OVH Public Cloud](#)

## Provider Comparison at a Glance

	AWS	GCP	Azure	OVH
Primary log service	Amazon CloudWatch Logs	Google Cloud Logging	Azure Monitor Logs	Logs Data Platform (LDP)
Archive / analytics layer	S3 + Athena	BigQuery (Log Analytics)	Archive tier + ADX	Managed ClickHouse
Cold object storage	S3 Standard-IA / Glacier IR	GCS Coldline / Archive	Blob Cool / Cold / Archive	Object Storage IA
Ingest price (EU region)	\$0.63/GB (CloudWatch)	\$0.50/GiB (Cloud Logging)	\$2.53–\$2.99/GB (Monitor)	€1.18/GB (LDP, 2–100 GB/mo)
7-year total (100 GB/day, managed path)	≈\$196,000 (≈\$2,333/mo)	≈\$133,000 (≈\$1,584/mo)	≈\$701,000 (≈\$8,345/mo)	Not viable (LDP; 1y max hot, no 7y archive path)
7-year total (bypass managed log service)	≈\$29,000 (≈\$345/mo)	≈\$50,000 (≈\$595/mo)	≈\$12,000 (≈\$143/mo)	—
7-year total (self-hosted ClickHouse on OVH)	—	—	—	≈€54,000 PAYG (≈€642/mo)
Max retention (managed service)	10 years (CloudWatch)	Unlimited (Log Sink export to GCS/BigQuery)	12 years (Monitor archive)	1 year hot (LDP)
Kubernetes-native agent	CloudWatch Agent / Fluent Bit	GKE managed logging agent / Fluent Bit	Azure Monitor Agent	Vector / Fluent Bit
GDPR EU region	✓ eu-central-1	✓ europe-west3	✓ West Europe	✓ Paris / Milan (3-AZ MKS)
Vendor lock-in	High	High	High	Medium

The biggest takeaway: **ingest cost dominates** at 100 GB/day. CloudWatch Logs at \$0.63/GB costs ≈\$1,890/month just for ingest — before any storage is paid. The cheapest path for AWS, GCP, and Azure is to bypass the managed log service and write to cheaper storage or analytics backends: S3/Athena (AWS), BigQuery (GCP), or Blob/Synapse (Azure). Against those managed paths, self-hosted ClickHouse on OVH (≈€54,000 over 7 years at PAYG rates) is **2–12× cheaper**. Among the bypass paths, Azure Blob is cheapest (≈\$12,000), followed by AWS S3 (≈\$29,000). OVH self-hosted ClickHouse is more expensive than AWS and Azure at PAYG rates, but wins when frequent SQL queries are needed — no per-query fees vs Athena (\$5/TB) or BigQuery (\$6.25/TiB). OVH Managed ClickHouse is not a viable alternative to self-hosted for 7-year archiving: without S3 tiering, full 7-year retention costs ≈€748,000, and even a 90-day hot-tier hybrid reaches ≈€129,000 — 2.4× more than self-hosted ClickHouse (≈€54,000).

## Services Overview

Each provider offers a layered stack: an operational log service (hot tier, search, alerting), an analytics/archive layer, and object storage for cold data.

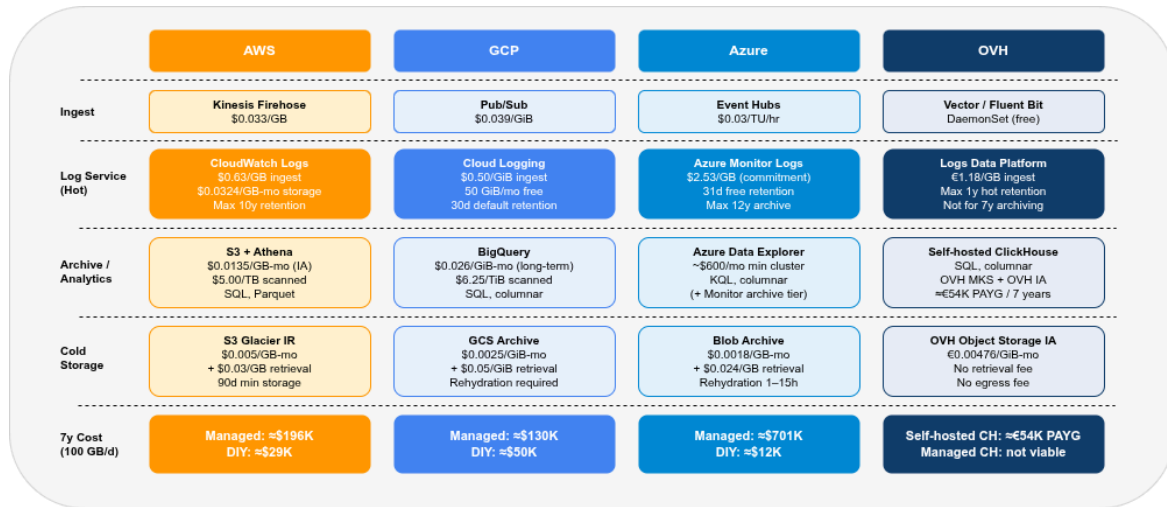


Figure 1: Services overview: AWS, GCP, Azure, OVH managed log archiving layers

## AWS

AWS has the richest set of primitives for log archiving, but they do not compose into a single coherent product. You choose a path and assemble it yourself.

Service	Role	USD (EU region)
<a href="#">CloudWatch Logs</a>	Ingest, hot storage, Log Insights queries	\$0.63/GB ingest, \$0.0324/GB-month storage
<a href="#">Amazon OpenSearch Service</a>	Managed Elasticsearch/OpenSearch, UltraWarm, Cold tier	\$0.024/GB-month (UltraWarm/Cold); compute extra
<a href="#">Amazon S3</a>	Object storage, archival	\$0.0135/GB-month (Standard-IA), \$0.005/GB-month (Glacier IR)
<a href="#">Amazon Athena</a>	SQL queries over S3	\$5.00/TB scanned
<a href="#">Kinesis Data Firehose</a>	Streaming ingest to S3 / OpenSearch	\$0.033/GB (EU, Tier 1)

**CloudWatch Logs** is the default operational log service: it receives logs from EC2 instances, Kubernetes pods (via the CloudWatch agent or Fluent Bit), Lambda functions, and almost every other AWS service. Retention can be set from 1 day to 10 years, or never expire. However, the storage pricing (\$0.0324/GB-month) makes it very expensive for 7-year retention of uncompressed logs. The Log Insights query engine supports a simple filter+aggregate language but is not SQL.

**Amazon OpenSearch Service** (managed Elasticsearch / OpenSearch) adds a proper inverted index on top of S3 via its UltraWarm and Cold tiers. UltraWarm nodes cache shards from S3 in memory/SSD; Cold storage keeps indices entirely in S3 and pays compute only when rehydrating for queries. Both UltraWarm and Cold storage cost **\$0.024/GB-month** for the actual data stored – plus the UltraWarm node instance hours while the service is running. This makes OpenSearch Service the right choice when you need full-text log search, KQL-like queries, and Kibana dashboards without self-hosting.

**S3 + Athena** is the cheapest AWS archival path. Kinesis Firehose buffers and compresses logs as gzip JSON before writing to S3; Athena provides serverless SQL queries at \$5.00/TB scanned. With 7× compression (typical for gzip-compressed log data), storage costs drop to roughly \$0.0135/GB-month (Standard-IA) or \$0.005/GB-month (Glacier Instant Retrieval). There is no managed ingest charge beyond Firehose — the \$0.033/GB Firehose rate is 19× cheaper than CloudWatch Logs ingest.

## GCP

GCP’s log archiving story is centred on BigQuery. Cloud Logging is the ingest and operational layer; BigQuery is the long-term analytics and archive layer.

Service	Role	USD (EU region)
<a href="#">Cloud Logging</a>	Ingest, hot storage (30d), Log Analytics	\$0.50/GiB ingest (after 50 GiB/month free)
<a href="#">BigQuery</a>	Columnar analytics, Log Analytics backend	\$0.026/GiB-month (long-term physical, europe-west3), \$6.25/TiB queries
<a href="#">Cloud Storage</a>	Cold object archive	\$0.006/GiB-month (Coldline), \$0.0025/GiB-month (Archive)
<a href="#">Pub/Sub</a>	Streaming ingest pipeline (DIY path)	\$0.039/GiB

**Cloud Logging** (formerly Stackdriver) receives logs from GKE, Compute Engine, Cloud Run, and most GCP-managed services automatically, with no agent configuration required for GKE. The default `_Default` log bucket retains logs for 30 days; beyond that, logs must be exported via a Log Sink to GCS or BigQuery. The first 50 GiB/month of ingest per project is free; beyond that, \$0.50/GiB applies.

**Log Analytics** is a Google-specific feature that exposes Cloud Logging data directly as a BigQuery dataset — you write SQL queries against your logs in the Cloud Logging UI without paying extra BigQuery query fees. For longer-term archives exported to BigQuery via a Log Sink, BigQuery’s normal pricing applies with **physical storage billing** (opt-in, billed on actual compressed bytes): active physical storage (\$0.052/GiB-month), long-term physical storage (\$0.026/GiB-month) for data not modified in 90+ days, and \$6.25/TiB for on-demand SQL queries. Logical billing (default) charges on uncompressed size and is significantly more expensive for log data.

**GCS Coldline and Archive** are significantly cheaper than BigQuery for write-once, rarely-read data: \$0.006/GiB-month (Coldline) vs. \$0.026/GiB-month (BigQuery long-term physical), or \$0.0025/GiB-month (Archive) for data read less than once per year. Unlike Azure Blob Archive or AWS Glacier Flexible Retrieval, GCS Archive stays online and does not require offline rehydration before reads. The trade-off is retrieval fees, higher operation costs, lower availability, and a 365-day minimum storage duration. BigQuery remains the better fit when logs must stay directly queryable with SQL.

The cheapest GCP path bypasses Cloud Logging entirely: **Pub/Sub** → **BigQuery** via the Storage Write API (default/batch mode: free). This avoids the \$0.50/GiB Cloud Logging ingest charge, replacing it with \$0.039/GiB Pub/Sub — a 12× reduction. The downside: you lose GCP-native log routing, log-based metrics, and automatic GKE log collection.

## Azure

Azure Monitor Logs / Log Analytics is the native Azure operational log service, but its ingest pricing at **\$2.53–\$2.99/GB** (West Europe) makes it the most expensive option in this comparison for high-volume archiving. Azure's strength is its archive tier (\$0.026/GB-month) and the ability to retain up to **12 years** without exporting to separate storage.

Service	Role	USD (EU region)
<a href="#">Azure Monitor Logs</a>	Ingest, hot storage (31d), archive tier	\$2.53/GB (100 GB/day commitment), \$0.026/GB-month archive
<a href="#">Azure Data Explorer (ADX)</a>	Columnar analytics, KQL	Cluster-based; min ≈\$600/month (2-node D11 v2)
<a href="#">Azure Blob Storage</a>	Object archive	\$0.0045/GB-month (Cold LRS), \$0.0018/GB-month (Archive LRS)
<a href="#">Azure Event Hubs</a>	Streaming ingest (DIY path)	\$0.03/hour per TU (Standard)

**Azure Monitor Logs** uses a tiered model: the first 31 days of interactive retention are included in the ingest price. Beyond 31 days, data moves automatically to the **archive tier** at \$0.026/GB-month (up to 12 years), or stays in the more expensive interactive tier (\$0.13/GB-month beyond 31 days) for continued KQL query access. Restoring data from archive back to interactive for querying costs an additional \$0.13/GB/day with a 2 TB minimum and 12-hour minimum window — not suitable for ad-hoc queries on individual log lines.

At 100 GB/day, the **commitment tier** applies (\$252.84/day in West Europe, effective \$2.53/GB — 15% off pay-as-you-go). Even with the discount, ingest is more expensive than the equivalent AWS CloudWatch (\$0.63/GB) or GCP Cloud Logging (\$0.50/GiB) by a factor of 4× and 5× respectively.

**Azure Data Explorer** is a columnar analytics engine with KQL (Kusto Query Language), purpose-built for telemetry and log analytics at scale. It supports hot/cold cache policies and can use Azure Blob Storage as its cold tier. However, ADX is cluster-based with no consumption/serverless tier — a minimum 2-node Standard\_D11\_v2 cluster costs ≈\$0.82+/hour (≈\$600/month) regardless of query volume. This makes ADX economical only for organizations that run near-continuous analytics queries against historical logs.

**Azure Blob Archive** is an excellent pure cold storage tier at \$0.0018/GB-month — about 2.8× cheaper than OVH Infrequent Access (≈\$0.0051/GB-month at EUR/USD parity). The catch: Archive tier requires **rehydration before reading**, which takes 1–15 hours and costs \$0.024/GB retrieved. For compliance archiving where data is rarely queried, Blob Archive is competitive; for regular investigative log queries, Azure Cold or Cool tier is more practical.

**Three Azure archival paths — explicitly:**

	Log Analytics archive (Path A)	Log Analytics + ADX	Event Hubs → Blob (Path B)
Ingest cost	\$2.53/GB (Monitor)	\$2.53/GB (Monitor)	≈\$0.003/GB (Event Hubs)
Long-term storage	\$0.026/GB-month (in workspace)	ADX hot/cold cache	\$0.0018–\$0.0045/GB-month
Query mechanism	Search jobs or KQL after restore	Direct KQL on hot/warm cache	Synapse Serverless or ADX
Extra services required	None	ADX cluster (≈\$600+/month min)	Blob Storage + query engine
Best for	Compliance archiving, infrequent queries	Near-continuous historical analytics	Cost-first, no Monitor dependency

Path A is the simplest Azure archive strategy: data auto-moves from the Log Analytics Workspace to its own archive tier after 31 days, no export pipeline required, and it can be accessed later with search jobs or restored for interactive KQL. The cost driver in Path A is **ingest at \$2.53/GB** — not the archive tier itself (\$0.026/GB-month is competitive with AWS Standard-IA). If Monitor’s native features (KQL alerting, Container Insights, metric correlation) are required, Path A avoids building a separate pipeline entirely.

ADX only makes economic sense when you run **near-continuous** historical queries — its cluster cost (\$600+/month) amortises only at high query volumes. For occasional investigative or audit queries, Path A’s search/restore-on-demand model is sufficient and significantly cheaper.

## OVH

OVH’s managed log offerings are designed for operational workloads (days to months), not 7-year archival. For the 7-year use case on OVH infrastructure, the right answer is self-hosted ClickHouse — covered in depth in [Part 1 of the self-hosted series](#).

Service	Role	Price
<a href="#">Logs Data Platform (LDP)</a>	OpenSearch-based SaaS log management	€1.18/GB (2–100 GB/month tier)
LDP Cold Storage add-on	Long-term S3-backed extension	€0.12/GB-month
<a href="#">Managed ClickHouse</a>	Columnar analytics (no S3 tiering)	€917/month ex. VAT (B3-16, 4 vCores/16 GB × 3 nodes; 100 GB base; +€0.438/GB-month additional)
<a href="#">Object Storage IA</a>	Infrequent Access cold tier	€0.00476/GiB-month (≈\$0.0052/GiB-month)

**Logs Data Platform** is an OpenSearch-based managed log service with zero-ops onboarding. The pricing model is straightforward: pay per GB ingested, with volume discounts at higher tiers. The critical limitation for 7-year archiving: the maximum hot retention is **1 year**, and the Cold Storage add-on costs €0.12/GB-month — **25× more expensive than OVH’s own Infrequent Access object storage tier** (€0.00476/GiB-month). At 100 GB/day with 7x compression and 7-year retention, LDP Cold Storage alone would cost ≈€459,000 in storage fees (plus ≈€36,000 in ingest). LDP is an operational product, not an archival one.

**Managed ClickHouse** is OVH's managed columnar database. ClickHouse requires a minimum of **3 nodes** for high availability and replication — a single-node deployment is not offered. The minimum production plan, **B3-16**, provides 3 nodes (16 GB RAM, 4 vCores per node) at **€917/month ex. VAT** with 100 GB of base storage. Additional storage is billed at **€0.438/GB-month** (€21.90/month per 50 GB block), up to a maximum of ~40 TB total.

OVH's own capabilities documentation confirms that **Managed ClickHouse has no S3 or object storage tiering**: all data lives on the cluster's local block storage. This has severe cost implications for long-term retention:

- **Full 7-year retention in Managed CH**: instance €917/mo × 84 = €77,028 + cumulative 1,531,000 GB-months × €0.438/GB-month = **≈€748,000 total** — not viable. This is 3.8× more expensive than the AWS managed path (≈\$196K).
- **90-day hot tier only** (the only practical use): constant ~1.3 TB compressed → instance + storage ≈€1,443/mo → 7y compute+storage ≈€121,000. Older data requires a custom export pipeline to OVH IA (no native tiering). Total including OVH IA: **≈€129,000**.

The self-hosted ClickHouse path (Vector → ClickHouse with 3-tier OVH IA tiering) costs **≈€54,000 over 7 years (≈€642/month avg.)** — **1.4× cheaper** than the 90-day Managed CH hybrid, and **2–12× cheaper** than all managed log service paths. Self-hosted CH handles multi-year tiering natively; Managed ClickHouse does not.

## 7-Year Cost Model

The following calculations use a consistent workload: **100 GB/day raw logs ingested, 7× compression** (typical for gzip-compressed log data), **7 years (84 months)**.

### i Cost model assumptions

The model uses 30-day months, 84 months total, USD/EUR parity for rough comparison, and public EU-region list prices as of May 2026. GB/GiB differences are kept as providers publish them; this creates small rounding differences but does not change the ranking. Query cost assumes 1 TiB scanned per month where the query engine bills by scan volume. Retrieval, egress, support plans, taxes, and negotiated discounts are excluded unless called out explicitly.

Key numbers: - Monthly raw ingest: 3,000 GB - Monthly compressed: ~429 GB - Total storage at year 7 (compressed): ~36,000 GB - Cumulative GB-months of compressed storage over 7 years: ~1,531,000

All prices are EU-region list prices as of May 2026 (AWS eu-central-1, GCP europe-west3, Azure West Europe, OVH Paris).

### Path A: Managed log service (fully managed, hot + archive)

This is the native path: ingest directly into the provider's log service, retain hot data for 30 days, archive to the provider's built-in cold tier.

Component	AWS (CloudWatch + S3 IA)	GCP (Cloud Logging + GCS Coldline)	Azure (Monitor + Archive tier)
Ingest	$\$0.63/\text{GB} \times 3,000 \times 84 =$ <b>\$158,760</b>	$\$0.50/\text{GiB} \times 2,950 \times 84 =$ <b>\$123,900</b>	$\$2.53/\text{GB} \times 3,000 \times 84 =$ <b>\$637,560</b>
Hot storage (30d rolling)	$3,000 \times \$0.0324 \times 84 =$ <b>\$8,165</b>	Included (30d free)	Included (31d free)
Export to cold (Firehose/Sink)	$3,000 \times \$0.033 \times 84 =$ <b>\$8,316</b>	Included (Log Sink free)	Included (auto-archive)
Cold storage (compressed, 7y)	$1,531\text{K} \times \$0.0135 =$ <b>\$20,669</b>	$1,531\text{K} \times \$0.006 =$ <b>\$9,186</b>	$1,531\text{K} \times \$0.026 =$ <b>\$39,806</b>
Query	$1 \text{ TiB}/\text{mo} \times \$5 \times 84 =$ <b>\$420</b>	Log Analytics for hot-tier data; GCS archive requires external query tooling – not included	Search jobs / KQL restore costs not included
<b>7-year total</b>	<b>≈\$196,000</b>	<b>≈\$133,000</b>	<b>≈\$701,000</b>
<b>avg. monthly</b>	<b>≈\$2,333/mo</b>	<b>≈\$1,584/mo</b>	<b>≈\$8,345/mo</b>

Notes: - **AWS**: 50 GiB/month CloudWatch free tier excluded (negligible at this scale). Firehose export is an architecture choice; CloudWatch Log Subscription Filter to S3 could be used instead at lower cost. - **GCP**: First 50 GiB/month of Cloud Logging ingest is free (~50 GiB/month deducted). Log Sink to GCS is free; GCS Coldline storage is billed on compressed data. GCS retrieval fees (\$0.02/GiB) not included – negligible for infrequent queries. - **Azure**: Commitment tier at \$252.84/day (\$2.53/GB). Azure Monitor’s billable ingest size is typically smaller than raw incoming JSON (Microsoft applies internal column trimming); \$2.53/GB is worst-case. Archive tier at \$0.026/GB-month applied to data after Azure’s internal compression (~5×). Search jobs can retrieve targeted long-term data; full interactive KQL analysis still uses restore jobs. Search / restore costs are not included. - **OVH LDP**: Not viable at 100 GB/day for 7-year searchable archive (1-year hot max, €0.12/GB-month cold = ≈€495,000 over 7 years: €459,000 storage + ≈€36,000 ingest).

### **Path B: Bypass managed log service (skip managed log ingest for archiving)**

Skip the managed log service for long-term data. Use a streaming ingest agent (Firehose, Pub/Sub, Event Hubs) to bypass the expensive managed log ingest layer – writing directly to cold object storage (AWS, Azure) or a queryable analytics store (GCP BigQuery). Sacrifice some operational log features (Log Insights UI, Cloud Logging routing, Azure Monitor metrics) in exchange for dramatically lower cost.

Component	AWS (Firehose → S3 IA + Athena)	GCP (Pub/Sub → BigQuery)	Azure (Event Hubs → Blob Cold)	OVH (self-hosted CH)
Ingest	$3,000 \times \$0.033 \times 84 =$ <b>\$8,316</b>	$3,000 \times \$0.039 \times 84 =$ <b>\$9,828</b>	<b>≈\$4,600</b> (1 TU + events)	Vector DaemonSet (free)
Storage (compressed, 7y)	$1,531K \times \$0.0135 =$ <b>\$20,669</b>	$1,531K \times \$0.026 =$ <b>\$39,806</b>	$1,531K \times \$0.0045 =$ <b>\$6,890</b>	OVH IA: <b>≈€7,400</b>
Query	$1 \text{ TiB/mo} \times \$5 \times 84 =$ <b>\$420</b>	BQ: $\approx 1 \text{ TiB/mo}$ free tier	Synapse serverless: <b>\$420</b>	ClickHouse SQL (free)
ClickHouse compute	Not needed (serverless)	Not needed (serverless)	Not needed (serverless)	MKS + $3 \times \text{c3-16}$ : <b>≈€46,534</b> (≈€554/mo)
<b>7-year total</b>	<b>≈\$29,000</b>	<b>≈\$50,000</b>	<b>≈\$12,000</b>	<b>≈€54,000</b>
<b>avg. monthly</b>	<b>≈\$345/mo</b>	<b>≈\$595/mo</b>	<b>≈\$143/mo</b>	<b>≈€642/mo</b>

Notes: - **AWS Firehose → S3**: No hot-tier log search. Add OpenSearch Service for search capability: +≈\$60,000–\$100,000 depending on instance sizes. S3 Glacier Instant Retrieval (\$0.005/GB-month) can reduce storage to ≈\$7,700 at the cost of \$0.03/GB retrieval fees. - **GCP Pub/Sub → BigQuery**: BigQuery physical storage billing (opt-in) used – long-term physical (\$0.026/GiB-month) kicks in after 90 days of no modifications, billed on compressed bytes. First 1 TiB/month of BQ queries is free; on-demand pricing (\$6.25/TiB) applies beyond that. Capacity-based slot pricing (Standard/Enterprise editions) is an alternative for high query volumes but adds fixed cost regardless of usage – not suitable for infrequent archive queries. Bypasses Cloud Logging entirely – no GCP log-based alerting or routing without custom Pub/Sub subscriptions. - **Azure Event Hubs → Blob Cold**: Blob Cold tier (\$0.0045/GB-month) allows direct read access without rehydration (unlike Archive tier). Write lifecycle policy from Hot → Cold after 30 days. Blob Archive (\$0.0018/GB-month) is 2.5× cheaper but requires rehydration (hours). - **OVH self-hosted ClickHouse**: MKS Control Plane (€78/month × 84 = €6,567) + 3 × c3-16 nodes required for HA (PAYG: 3 × €158.54/mo × 84 = €39,967) + OVH IA storage (≈€7,400) = ≈€54,000 total. With a 36-month Savings Plan (3 × €71.34/mo): ≈€33,000. Vector DaemonSet for ingest (runs on existing nodes). Hot tier (SSD 0–7d), cold tier (OVH IA, 7d–7y). Full SQL queries on all tiers via ClickHouse, no per-query fees. Details in [Part 1 of the self-hosted series](#).

## Cost at a glance

### ☒ 7-year cost summary at 100 GB/day

Path	AWS	GCP	Azure	OVH
Managed log service	≈\$196,000 (≈\$2,333/mo)	≈\$133,000 (≈\$1,584/mo)	≈\$701,000 (≈\$8,345/mo)	Not viable (LDP; 1y max hot)
Bypass managed log service	≈\$29,000 (≈\$345/mo)	≈\$50,000 (≈\$595/mo)	≈\$12,000 (≈\$143/mo)	—
Self-hosted ClickHouse (OVH)	—	—	—	≈€54,000 PAYG (≈€642/mo)

All EU-region list prices, May 2026. Monthly figures are averages over 84 months (costs rise as storage accumulates). Self-hosted OVH includes MKS control plane + 3 × c3-16 nodes + OVH IA.

Bypassing managed log ingest cuts costs dramatically: **AWS 7×** (\$196K vs. \$29K), **GCP 3×** (\$133K vs. \$50K), **Azure 58×** (\$701K vs. \$12K). The dominant factor is always **ingest cost**: CloudWatch at \$0.63/GB, Cloud Logging at \$0.50/GiB, and Azure Monitor at \$2.53/GB make 7-year retention of 100 GB/day inherently expensive. Bypass ingest agents (Firehose, Pub/Sub, Event Hubs) cost 10–80× less.

## Storage Model Comparison

Understanding *how* each provider stores logs determines query capability, compression efficiency, and cold-tier behaviour.

### CloudWatch Logs

CloudWatch stores logs in its own proprietary format (not user-accessible S3). No user-controlled compression: billing is on raw ingested bytes. Log Insights scans raw log events (not columnar compressed data), making queries expensive at scale. The managed export path (Log Subscription Filter → Firehose → S3 as gzip) enables compression, but requires additional infrastructure.

**Compression at query time:** Not applicable (proprietary, billing on raw bytes)

### Google Cloud Logging / BigQuery

Cloud Logging stores logs in a Google-managed format. When exported to BigQuery, Google uses columnar storage internally; the long-term storage rate (\$0.026/GiB-month, physical billing) applies to the stored (compressed) size. In practice, JSON application logs compress to roughly 5–10× in BigQuery columnar format. BigQuery queries are billed on bytes scanned before columnar pruning — use SELECT only the columns you need, and partition by day, to reduce scan costs.

**Compression:** ~5–10× (BigQuery internal columnar)

### Azure Monitor Logs

Azure Monitor uses its own columnar format (similar to a Kusto database). Ingest pricing is based on the raw data size; storage pricing is based on the stored (compressed) size. Azure reports typical compression of 4–8× for application logs. Archive tier stores the same compressed data but restricts live KQL queries: use search jobs for targeted retrieval, or restore jobs for full interactive KQL over a time range.

**Compression:** ~4–8× (Azure internal columnar, billed on stored size for archive)

### S3 / GCS / Blob + gzip JSON (bypass path)

When logs are written to object storage as **gzip JSON** via Firehose/Dataflow/Event Hubs, you control compression. gzip typically gives 7–10× compression on structured log data. Athena, BigQuery (external tables), and Synapse Analytics charge by bytes scanned — tight partitioning (by day, namespace, service) reduces query costs by 10–30× regardless of format. Parquet conversion is available as an optional upgrade (\$0.02/GB extra for Firehose) and cuts Athena scan costs further for high-query workloads.

**Compression:** 7–10× (gzip JSON, no format conversion fee)

## When to Use Managed vs. Self-Hosted

Situation	Recommendation
Small team, no Kubernetes expertise	Managed log service (CloudWatch / Cloud Logging) for simplicity
100+ GB/day, 7+ year retention, cost-sensitive	Bypass managed log service path; or self-hosted ClickHouse on OVH
Full-text log search required	AWS OpenSearch Service (UltraWarm) or GCP BigQuery Log Analytics
SQL analytics on archived logs	GCP BigQuery or AWS Athena
Compliance archiving, rarely queried	Azure Blob Archive or S3 Glacier IR
Kubernetes-native, existing GKE fleet	GCP (automatic log collection from GKE)
BSI C5 required (German public sector / KRITIS)	AWS (eu-central-1), GCP (europe-west3), or Azure (West Europe) – all hold BSI C5 attestation; OVH does not
EU-sovereign, no US-jurisdiction risk (Cloud Act / FISA)	OVH (Paris or Milan, EU 3-AZ MKS) – not subject to US extraterritorial access law
Lowest possible cost, engineering capacity	Self-hosted ClickHouse on OVH ( <a href="#">Part 1 of the self-hosted series</a> )

The self-hosted ClickHouse path costs ≈€54,000 over 7 years at PAYG rates (≈€33,000 with a 36-month Savings Plan) and is 2–12× cheaper than all managed log service paths. OVH Managed ClickHouse is not a comparable alternative: without S3 tiering, even a practical 90-day hot-tier hybrid costs ≈€129,000 – 2.4× more than self-hosted. Self-hosted CH beats the AWS/GCP cold-storage path when query volume is high: ClickHouse has no per-query fees, while Athena charges \$5/TB and BigQuery \$6.25/TiB. For pure compliance archiving with rare queries, Azure Blob Archive (≈\$12,000) is cheapest. The OVH self-hosted path requires Kubernetes operational experience but eliminates vendor lock-in entirely.

## FAQ

### Is CloudWatch Logs free for small workloads?

The AWS Free Tier includes **5 GB/month each** (separate allocations) for CloudWatch Logs ingest, storage, and Log Insights queries scanned. For workloads under 5 GB/month (small teams, low-traffic services), CloudWatch is effectively free. Beyond 5 GB/month, the \$0.63/GB ingest charge applies without volume discounts for custom application logs. Vended logs (from AWS services like VPC Flow Logs, RDS, Lambda) have volume tiers starting at \$0.315/GB after 10 TB/month.

### Can I use Google BigQuery Log Analytics without paying Cloud Logging ingest fees?

Yes, by bypassing Cloud Logging: stream logs directly to BigQuery via the Storage Write API (default/batch mode: free) or via Pub/Sub → Cloud Dataflow → BigQuery. You lose Cloud Logging features: automatic GKE log collection, log-based metrics, Log Router routing rules, and log-based alerting. For high-volume archiving where these features are not needed, the Pub/Sub path reduces ingest cost from \$0.50/GiB to \$0.039/GiB.

## What is the difference between Azure Monitor Archive tier and Azure Blob Archive?

Azure Monitor **archive tier** (\$0.026/GB-month): data stays inside Log Analytics Workspace and supports up to 12 years. Use search jobs for targeted retrieval, or restore jobs when you need full interactive KQL over a time range. Best for: compliance archiving where occasional KQL-based investigations are needed.

Azure **Blob Archive** (\$0.0018/GB-month): data is in Blob Storage, not queryable without rehydration, requires Synapse Analytics or ADX for SQL queries. Rehydration to Cool/Hot tier takes 1–15 hours. 14× cheaper storage but requires custom query tooling. Best for: write-once, almost-never-read compliance archives.

## Why is OVH Managed ClickHouse not recommended for 7-year archiving?

Two reasons. First: **no S3/object storage tiering** — confirmed by OVH’s own capabilities documentation. All data must live on local block storage. The minimum production plan, B3-16 (3 nodes, 4 vCores/16 GB RAM), costs **€917/month ex. VAT** with 100 GB base storage; additional storage is billed at **€0.438/GB-month**. Keeping all 7 years of logs in Managed CH costs ≈€748,000: €77,028 in instance fees plus ≈€671,000 in cumulative storage. Second: even a practical 90-day hot tier with a custom export job to OVH IA reaches ≈€129,000, because Managed CH has no native tiering pipeline — you build and operate the cold-tier export yourself. The self-hosted ClickHouse path with OVH IA tiering (≈€54,000) is **1.4× cheaper** than the 90-day hybrid and **13.9× cheaper** than full 7-year retention in Managed CH.

## Are prices shown inclusive of data egress?

No. The cost models above exclude egress (data transfer out of the region). AWS charges ≈\$0.085/GB for egress from EU regions; GCP charges \$0.085/GiB for egress from europe-west3; Azure charges ≈\$0.08/GB from West Europe. OVH charges no egress fees **within the OVH network** (egress to external networks is billed normally). For 7-year log archiving, egress costs only apply when investigative queries retrieve data to an external system — typically a small fraction of total cost. OVH’s zero-egress policy is a meaningful advantage for multi-step investigative workflows that pull large data volumes.

---

*Part 2: [Pre-Flight, Flexibility & Auditor Export](#)*

*Part 3: [Operations — Kubernetes integration, ingest pipeline, backup & DR](#)*

*Part 4: [Security & Compliance — IAM/RBAC, encryption, WORM, GDPR](#)*

*Part 5: [Query, Dashboards & Recommendations](#)*

*Part 6: [Production Checklist, Guardrails & Runbooks](#)*

*Self-hosted alternative: [Elasticsearch vs. OpenSearch vs. Loki vs. Quickwit vs. ClickHouse — Part 1](#)*