


GovTech Serverless Migration Workshop

Moving StackOps tenants from Elastic Cloud
Hosted to Elastic Cloud Serverless

PRESENTED BY
John Carlo Avila
Principal Consulting Architect, Elastic



Elastic Cloud Hosted v8.19



- Elastic: End of Support in 2nd half of 2027
- Decision to move to Elastic Cloud Serverless
- Trial will be available from Nov'26 (look out for announcement)
- New onboarding to be available by end of this quarter
- Migration of existing tenants planned from 2027

Serverless Cost & Pricing



Continue to enjoy 7% discount on your Elastic usage.



Potential Cost Reduction (Indicative only) – directionally positive, but savings vary by tenant usage profile and requires further modeling to confirm.

Cost on Serverless will be usage-based.

Serverless Cost & Pricing (Reference links on cost estimation)

- [Elastic Cloud Serverless Pricing](#)

Elastic Cloud pricing table

Rates for Elastic Cloud, broken out by the components that make up your bill.

Elastic Cloud Serverless

Serverless projects

Charges that follow from the Elastic Cloud Serverless solution you choose. Pricing applies per project.

Cloud provider

Region

Amazon Web Services

Singapore (ap-southeast-1)

Elasticsearch

Vector Database

Observability

Security

Collapse all

Agent Builder

Product	Volume tier	Price	Unit
Agent Builder Executions (Observability)	0 to 10,000	\$0.0000	Execution
	10,000 to 100,000	\$0.0500	Execution
	100,000 to +	\$0.0250	Execution

Cross Project Search

Product	Price	Unit
CPS - Observability - Retained data in linked Elasticsearch projects	\$0.0375	Mounted GB-month
CPS - Observability - Retained data in linked Security Projects	\$0.0113	Mounted GB-month
CPS - Observability - Retained data in linked Vector Database projects	\$0.1250	Mounted GB-month
CPS - Observability - Retained logs, traces, and non-TSDS data in linked Observability projects	\$0.0113	Mounted GB-month
CPS - Observability - Retained TSDS metrics data in linked Observability projects	\$0.0029	Mounted GB-month

Ingest

Product	Volume tier	Essentials	Complete	Unit
Data Ingestion	0 to 1,500	-	\$0.6250	GB
	1,500 to 3,000	-	\$0.4063	GB
	3,000 to 6,000	-	\$0.2813	GB
	6,000 to 15,000	-	\$0.1875	GB
	15,000 to 30,000	-	\$0.1313	GB
	30,000 to 60,000	-	\$0.1250	GB
	60,000 to 150,000	-	\$0.1188	GB
	150,000 to +	-	\$0.1156	GB

- [Cost Estimator](#)

Serverless plan

Pay only for the data you use. Elastic Serverless scales automatically — no infrastructure to manage.

Select Your Product

Observability

Monitor and resolve issues.

Security

Detect and respond to threats.

Vector Database

Store and search vectors.

Elasticsearch

Build and query search-powered apps.

Coming Soon

Choose plan

Logs Essentials

Core logging and analytics for smaller workloads.

Complete

Recommended

Full stack observability for cloud-native and hybrid environments.

Pick Support Tier

Info

Limited

Community and email support, standard response.

Base

Business-hours support with prioritized response and phone access.

Enhanced

24/7 global support with 1-hour response SLA.

Premium

Priority 24/7 support with 30-minute critical-issue response.

Set Environment

Cloud and Region

Choose where to deploy — Elastic manages scale and uptime.

Cloud Provider

Amazon Web Services

Region

ap-northeast-1

Add Data Sources

Logs

Centralized log collection and analytics for apps and infrastructure.

Raw log data volume

GB/day

25

Raw daily log volume before enrichment. Billing meters the enriched, indexed size — about 1.68x this value.

Data retention period

months

12

\$1,650.29

per month

Cost breakdown

Data ingest

25 GB/day raw = 41.50 GB/day materialized

\$757.38 / mo

Data retention

15,148 GB/month retained

\$677.66 / mo

Support

Premium (15% of \$1,435.04)

\$215.26 / mo

Total

\$1,650.29 / mo

Yearly estimate: \$19,803.54

Contact us

Start free trial

Estimate limits

The calculator supports estimates up to \$10K/month. For higher usage, our experts can help you build a tailored plan.

What we'll cover today

Part 1 · Refresher (brief intro)

- Elastic Serverless: core fundamentals
- Business value: why migrate?
- Choosing the right project type
- Security & access management
- Data management & storage changes
- Feature changes

Part 2 · Migration deep dive (full coverage)


- Migration strategies & best practices
- Content migration: dashboards, alerts, saved objects
- Data ingestion migration: agents, Fleet, APM
- Machine learning workloads: jobs & collaterals
- Feature & use case migration
- Operational readiness: roles, responsibilities, logistics
- Migration timeline & project scheduling

BRIEF INTRO

PART 1

Refresher

A quick recap of the Serverless basics from the intro webinar



Why StackOps is moving to Serverless

- 1 Platform modernization**
Move the multi-tenant platform from ECH to fully managed Serverless for cloud-native scale and faster access to new capabilities, especially Agentic AI.
- 2 Less operational toil**
No ES version upgrades, SLM, ILM tuning or cluster sizing. Elastic owns infrastructure, scaling and snapshots.
- 3 End of Life: Elasticsearch v8.x**
Current Elasticsearch version StackOps is using is reaching its end-of-maintenance / end-of-life. Migration target to complete is September 2027
- 4 New capabilities**
Workflows, Agent Builder, Elastic Managed Integrations (agentless), Streams, ES|QL-based alerting, Columnar Data Store (upcoming).

A technology refresh of StackOps, not a lift-and-shift.

For tenants: consistent onboarding, capacity that scales with your ingest, and access to AI and automation features without waiting for platform upgrades.

Elastic Cloud Serverless

Fully managed

- No clusters, nodes, data tiers or manual scaling to manage
- You deploy projects; Elastic scales and maintains them
- No stack version: Elastic keeps it current, always the latest version

New architecture

- Data lives in cloud object storage (S3, GCS, Azure Blob)
- Separate indexing and search tiers; independently scales

Built for resilience

- Nodes spread across AZs
- Auto-scales based on need

For tenants: no upgrade windows to plan around, and new features arrive continuously. You work in projects (not deployments), each with its own Kibana and Elasticsearch endpoints.

StackOps Suites

Observability Project

- Full Observability solution: Logs, Metrics, Traces (APM), Synthetics
- Includes ML/AIOps, SLOs, Agent Builder and Workflows

Security Project


- Includes Agent Builder, Workflows, UEBA and Threat Intel
- Endpoint Protection Complete add-on (no additional cost per endpoint)
- [Optional](#) Cloud Protection add-on (CSPM)

StackOps default: every tenant gets a paired Observability + Security projects, linked by Cross-Project Search.

For tenants: logs, metrics, APM and SLOs live in your Observability project; detections, Defend and SIEM content live in your Security project. You can search across both.

Common features: Agent Builder, Workflows, Machine Learning

From data tiers to indexing and search tiers



Existing (ECH) vs new (Serverless) architecture

Existing (ECH)

- Hot primary + replica, cold and frozen nodes
- Data lives on node disks; tiers are sized and managed

New (Serverless)

- Indexing tier and search tier scale independently
- Both tiers share an object store (index store + translog)

How data storage and retention change

Area	Elastic Cloud Hosted	Elastic Cloud Serverless
Storage	Local disks, hot/warm/cold/frozen tiers	Object storage with no tiers to manage
Lifecycle	ILM policies (hot → warm → cold → delete)	Data stream lifecycle (data_retention)
Retention	Per-policy	Per-data-stream setting
Snapshots	SLM policies owned by StackOps	Elastic-managed, kept about 7 days; restore via Support ticket;
Durability	Replica shards	Object store: AZ-replicated
Indexing mode	logsdb for logs-*-*	logsdb for logs and traces (OpenTelemetry); columnar storage (tech preview)

For tenants: the StackOps default keeps data for 1 year (365 days). No more hot/warm/cold choices. Data restores go through StackOps to Elastic Support, from snapshots kept about 7 days.

Identity, roles and network access

Users & roles

- No change: Login is still via Techpass
- Privilege Changes: manage_ilm, manage_ccr, manage_slm and manage_watcher are **unavailable**
- No replacement roles as these features are no longer relevant in Serverless

Automation identity

- No local users / service accounts (Est. 14% of deployments impacted)

Network & encryption

- Static egress/ingress IPs provided for AWS ap-southeast-1 and Azure southeastasia.
 - Will be published later in StackOps documentation
- Private endpoints for ES and Fleet are auto-configured

Gaps, workarounds and tenant impact

Feature & Impact	Recommended approach	Tenant impact
Stack Monitoring / .monitoring-* Est. 100% affected	N/A	Platform monitoring is covered by Elastic
Universal Profiling Est. 65% affected	OTel profiling integration (tech preview)	Migrate to OTel-based profiling (tech preview)
Dot-prefixed index templates Est. 25% affected	Don't use dot-prefixed indices for user data	Behaviour change
Watcher Est. 16% affected	Kibana Alerts + Workflows; ES QL alerting	Watchers have to be migrated to Kibana alerting rules
PNG / PDF export Est. 13% affected	PNG / PDF report generation coming soon	Known gap; CSV export remains
Real User Monitoring (RUM) Est. 10% affected	No workaround today; excluded from Phase 1	Not currently available; Elastic to share roadmap / availability later
Cross-Cluster Replication / remote clusters Est. 2% affected	Workflows push / remote reindexing	Explore Cross-project search; workflows / remote reindexing as replacement
On-prem Fleet Server Est. 1% affected	Cloud-only Fleet in Phase 1	No on-prem Fleet Server currently
Custom plugins & analyzer dictionaries 0% affected	Synonyms API	Currently not in use; No Tenant impact

APIs your automation can no longer call


API group	Examples	Status on Serverless	What to do
Cluster & node	_cluster/health, _cluster/settings, _nodes/stats	Blocked (HTTP 410)	Remove health checks that rely on them
Snapshot & ILM	_snapshot/*, _ilm/*	Blocked	Retention via data stream lifecycle; restores via StackOps support ticket
Index maintenance	_close, _open, _forcemerge, _shrink, _split, _stats	Blocked	Not needed on Serverless; remove from scripts
Legacy templates, Watcher, rollups	_template, _watcher/*, _rollup/*	Blocked	Composable templates; Kibana alerting rules
Cat APIs	_cat/health, _cat/nodes, _cat/shards	Blocked	_cat/indices, _cat/aliases, _cat/count still work
Search & write	_search, _bulk, _msearch, _count, _reindex, PIT, scroll	Works	Only update endpoints and API keys
Wildcard delete	DELETE /pattern*	Rejected (400)	Delete indices by exact name

For tenants: review any scripts, CI jobs or third-party tools that call Elasticsearch directly before your migration wave.

PART 2

Migration deep dive

How tenant content, ingestion and operations move to Serverless.



01

Migration strategies & best practices



Two-phase Migration Approach

Phase 1 · No Migration (Configuration / Data)

- Tenants with no legacy data on ECH are onboarded straight onto Serverless
- Focus on cloud-native, internet-connected workloads
- Elastic-provided Fleet Server only

Phase 2 · With Existing ECH Data / Deployment

- Migrate current ECH tenants with a fresh suite or a remote reindex to backfill history
- Complex networking and on-prem data sources

Phase 1 Onboarding: Direct to Serverless

Direct Onboarding Requirements

- Tenants without any data or configuration that needs to be migrated
- No dependency on legacy configurations or existing ECH deployments (cross ECH / Serverless search)
- Doesn't require RUM in the short-term

Provisioned Environment Suite

INCLUDED PROJECT

Observability (Obs) Project

Full visibility across metrics, logs, and traces out-of-the-box

INCLUDED PROJECT


Security (Sec) Project

Pre-configured unified threat detection and security analytics suite

Streamlined Onboarding Experience

Eligible Phase 1 tenants receive access to their dedicated Observability and Security project suite without migration overhead.

How a tenant moves (Phase 2, proposed)



Proposed flow, to be finalized in the Phase 2 runbook. Pilot sub-waves (Internal StackOps) will validate it before standard waves.

For tenants: your active involvement is in steps 2, 5 and 6. **Plan for the cutover window** and the validation sign-off.

Phase Qualification and Options

Phase 1 Qualification

- No 3rd-party tools / libraries that need older Elasticsearch versions (before v9.5)
- No legacy templates
- No on-prem Fleet Server
- Cloud-native / internet-connected / internet-accessible data sources

Phase 2 Options

- **Age-out:** new ingest goes to Serverless and old data expires on ECH
- **Backfill:** reindex-from-remote
- Both options are supported equally
- Cost implication to be discussed later

Special topologies

- Hub-and-spoke tenants
- Central projects
- StackOps deployments with inter-dependencies / data-sharing requirements

Age-out or backfill: what each means for you

	Age-out (no backfill)	Backfill
Where new data goes	Serverless - from cutover date	Serverless - from cutover date
Where history lives	Stays on ECH until it expires under its retention	Copied into Serverless
Dashboards & queries over history	Use the old ECH deployment for pre-cutover periods	Work across the full period in Serverless
ML anomaly detection	Jobs retrain from new data (~3 months to train)	Jobs can train on backfilled history
Effort & timing	Simplest cutover; ECH retired after data ages out	Extra reindex step; source must be on Elastic Cloud
Good fit when	Short retention, or history rarely queried	Long look-backs, compliance searches, ML reliance

For tenants: tell StackOps your preference and why during assessment; cost implications will be shared separately.

Accelerating Migration with Elastic Tooling

01 Assessment & Discovery



Automated Analysis

Automated inspection and scan of active features, settings, and workloads using the Migration Tool.



Gap Identification

Pinpoints functional gaps, legacy configurations, and unsupported dependencies prior to migration.



Architectural Review

Comprehensive evaluation of current deployment landscape and target Serverless topology.

02 Guided Migration



Tailored Strategy

Expert recommendations on Age-out vs. Backfill migration paths per tenant requirement.

- **Age-out:** New ingest to Serverless; history expires on ECH
- **Backfill:** Reindex-from-remote for historical continuity



Automated Provisioning & Baseline

Automated setup and baseline configuration executed directly by StackOps.

- Observability & Security suite environment setup
- Role definitions, permissions & baseline policies

What tenants should prepare before their wave



Content & alerting

Collect **connector secrets** (email, Slack, webhooks) for manual re-entry
100% of Tenants

List **Watchers** to rebuild as Kibana rules
Estimated 16% of Tenants

List **scheduled PNG/PDF reports**
Estimated 13% of Tenants

Flag **TSVB / Timelion visualizations** to modernize
Estimated 11% of Tenants



Data & ingestion

Agent inventory

Versions (≥ 8.11 ; at least 8.19.x recommended) and on-prem locations

Logstash, Beats, APM agents & API clients

Clients that need new endpoints and keys

Custom templates & runtime fields

Includes audit of any dot-prefixed indices

Migration preference

Selection between Age-out vs Backfill data migration strategy



Features & access

Solutions inventory

Inventory use of Defend, OSQuery, APM, Synthetics, RUM, Profiling

Machine Learning **Estimated 20% of Tenants**

Identify active ML jobs and trained models

API compatibility

Audit custom scripts calling blocked APIs

Users & roles **Estimated 14% of Tenants**

Users and roles needed per project - local / service account replacement

Retention with data stream lifecycle

How retention is applied

Security projects

Default configuration baseline applied by StackOps (1 year retention)

Observability projects

Default configuration baseline applied by StackOps (1 year retention)

Non-data-stream indices (Not configured by baseline)

Not time series, no standard time field. Examples include lookup indices, caches, references, knowledge bases, etc.

Transform Jobs

Destination indices have no retention; configure retention on the transform job itself

What this means for tenants

Automatic Stream Deletion

Data in data streams is deleted automatically once it passes the retention threshold.

Plain Indices (Manual Cleanup)

Plain (non-data-stream) indices don't expire; requires manual cleanup (e.g. delete-by-query).

Compliance Retention Needs

Need longer retention for compliance? Raise it with StackOps before your wave.


No ILM Phases

There's no ILM: hot/warm/cold phases no longer apply under this stream lifecycle.


03

Data ingestion migration

Elastic Agent, Fleet and APM



Where do you manage Elastic Agents?




Integrations	Count
Observability project	443
Security project	462
Obs exclusive (apm, synthetics...)	3
Sec exclusive (endpoint, osquery, ML...)	19

Exceptions

- **Defend + Synthetics private monitor** or APM: secondary Fleet on the Obs project
- **Obs-focused tenant needing OSQuery:** Fleet on the Sec project


For tenants: Need decision on which project agents will be managed.

One agent, two projects: remote ES output



Integration assets are preloaded

The project receiving remote output doesn't get dashboards or templates automatically, so the integration assets (dashboards, ingest pipeline, templates) need to be pre-loaded



The same pattern works in reverse: Fleet on Obs with a remote output to Security ES.

For tenants: one Elastic Agent per host can still feed both projects. Security data lands in Security and observability data in Observability, each with its integration dashboards.

On-prem agents & Fleet: open questions for tenants

On-prem agents



On-prem agents that cannot go through a proxy?



Is an internet proxy a viable alternative?



Reinstallation of existing Fleet-managed Elastic Agents (can use [SORA](#))

Fleet limits & changes



Agents $\geq$ 8.11 can connect; 8.19.x recommended



Up to 30,000 agents and 500 policies per project



No automatic agent upgrades

Ingest architecture, indexing and APM

Ingest architecture

Supported Frameworks

All ECH ingest architectures are supported except air-gapped on-prem Fleet Server

Pipelines & Proxies

Logstash, Kafka and proxies remain valid

Agentless Support

Elastic Managed Integrations (agentless): optional, new feature

Indexing modes

Logs & OpenTelemetryTraces Mode

Logs and traces: logsdb

APM Standards [Estimated 25% of Tenants](#)

logs-apm* and traces-apm*: keep standard
[\[Highly recommended to use OpenTelemetry\]](#)

Metrics Configuration

Metrics: keep each integration's mode

APM & RUM


Fleet & Observability

APM Server and Synthetics private locations need Fleet on the Observability project

Serverless Status

RUM isn't currently supported in Serverless [\[on the roadmap\]](#)

Elastic Managed Integrations (agentless)



How it works

Kibana Configuration
Configure the integration in Kibana; Elastic runs the collectors

Zero Agent Overhead
No Elastic Agent to install or upgrade for supported cloud sources

Optional Feature
Optional, new feature

04

Machine learning workloads

Est. 20% affected



ML migration runbook

Estimated 20% of Tenants

STEP 01 Inventory

Per tenant: anomaly detection jobs, datafeeds, trained models and ML SIEM rules (Phase 2 assessment)

Tenant action

Confirm jobs still in use

STEP 02 Place

Security ML packages go to Security project: ded, dga, lmd, pad, problemchild, beaconing, hta; AIOps jobs go to Observability

Tenant action

None required

STEP 03 Recreate

Export job & datafeed configs; recreate per project, pointing datafeeds at local indices / cross-project search (no rc_* remote patterns)

Tenant Role

Export ML job and recreate in Serverless

STEP 04 Link

After CPS linking, validate the datafeed, and start ML job

Tenant Role

Validate datafeed and start the ML job

STEP 05 Validate

Compare job state and results with ECH before cutover or age-out

Tenant Role

Sign off on results

STEP 06 Time to train

Generally, 3 months of past data is preferred as input for anomaly detection jobs to establish a pattern aligned with production behaviour

Tenant Role

Monitor

06

Operational readiness

Roles, responsibilities and logistics



Who owns what during and after migration

PLATFORM

Elastic (platform)

- Infrastructure, scaling, upgrades and HA
- Snapshots and data restore (via Support)
- Platform monitoring

PLATFORM

StackOps

- Baseline automation and tenant onboarding
- Enrichment workflow
- Audit Logs, initial SIEM content and alert rules
- Tenant comms and documentation


APPLICATION TEAMS / AGENCIES

Tenants


- Readiness inventory and age-out/backfill choice
- Re-enroll agents; update endpoints and API keys
- Rebuild Watchers into Kibana Alerts / Workflows
- Re-enter connector secrets
- Validate data, dashboards and alerts; sign off

07

Migration timeline & project scheduling



Migration roadmap (Plan)



Serverless migration: FAQ

TIMELINE

Do we have to migrate?

Yes. The current hosted version reaches end of support, so every tenant must be on Serverless by Sep 2027.

HOSTING

Where does Serverless run?

Entirely on Elastic Cloud, not in GCC or CEP.

COST

Will we save money?

On average, per latest data, there may be savings of around 5% or more. Model your own workload with the Elastic Serverless pricing calculator.

COST

How is ingest billed?

On uncompressed ingest volume, plus how long you retain it. You pay for what you ingest and keep.

COST

Will we pay twice during migration?

Some parallel-run cost is possible and is assessed separately. Migrate promptly once you start.

DATA

What happens to historical data?

Let it age out on Hosted, or copy it to Serverless at extra cost. Need over 1 year of retention? Contact StackOps.

TIMELINE

Onboarding a new project soon?

Go straight to Serverless if it is within the next couple of months.

SUPPORT

Who supports us afterwards?


StackOps, with Elastic. Upgrades become automatic, no more scaling, shard optimization, cluster upscaling / downscaling

Feature gaps or edge cases? Bring them to a migration clinic: 26 Oct, 28 Oct or 2 Nov (sign up via the QR code).

Workshops & Migration Clinic session Registration


Registration is open

(Do register early to secure your slot!)



<https://go.gov.sg/stackops-serverless-clinic>

Reach out to [StackOps Support](#) for any further clarification



Thank you

Questions & discussion

⚡ Hands-on with Elastic Serverless Observability

Join us on **3 Nov '26** at **Stack Conference Workshop**

Explore, experience, and unlock insights with **Elastic Serverless!**

02:00 PM - 05:00 PM

[Workshop 11] Less Plumbing, More Insight: Hands-On with Elastic Serverless Observability

Mr Sean MATTHEWS, Elasticsearch Pte Ltd

Mr YAP Yi Sheng, GovTech Singapore

Workshops


This hands-on workshop combines modern observability practices with AI agent development in Elastic. Participants begin by deploying an EDOT-instrumented OpenTelemetry demo, validating end-to-end telemetry, and injecting controlled failures to troubleshoot a feature-flagged product catalog issue. Using logs, metrics, and distributed traces, they analyze service dependencies and apply AI-assisted investigation techniques during a payment stress scenario to accelerate root-cause analysis.

Building on this foundation, learners shift to a UI-first agent workflow in Kibana. They inspect the Kafeju agent, understand how tools are defined, and create custom ES|QL-backed tools that integrate directly into the agent. Through exercises like exploring ML anomalies, reverse-engineering existing tools, and designing new ones, participants enable multi-tool reasoning that combines built-in and custom capabilities. The workshop culminates in agents answering realistic, data-backed FinOps-style questions, demonstrating how observability data and AI agents work together to reduce time-to-resolution and enhance operational insight.

Elastic Serverless Observability removes five operational concerns you currently manage — shards, storage tiering, version upgrades, monitoring pipelines, and fixed infrastructure nodes. They stop being your problem entirely. That frees you up to focus on what the platform now enables: AI-driven investigations that surface root cause automatically, 100+ built-in ML jobs for anomaly detection and forecasting, ES|QL for sub-second queries across millions of metrics, and 450+ integrations out of the box. You get unified logs, metrics, traces, and synthetic monitoring in a single project; SLOs and alerting without infrastructure sizing; and continuous access to new capabilities like Knowledge Indicators and agentic remediation workflows — no upgrade cycle required. This workshop walks through the platform hands-on so you can see what's available from day one.

Who should attend:

Software developers, Project Managers



Mr Sean MATTHEWS

Senior Customer Architect, Elasticsearch Pte Ltd

[View Profile](#) →



Mr YAP Yi Sheng

Engineering Lead, GovTech Singapore

[View Profile](#) →