

Upskill with StackOps Series

Managing Data Lifecycle

Wen An Tian
Software Engineer, StackOps



StackOps Upskill Series

Webinar (Online)

**Accelerating Deployment
Discovery with SOA**

● Jan 2026

Workshop (In-person)

**Visualisation &
Observability in StackOps**

● Feb 2026

Workshop (In-person)

**Overview & Ingestion in
StackOps**

Webinar (Online)

**Log Analysis (ESQL) &
Dashboarding**

● Mar 2026

Workshop (In-person)

**Applying Machine
Learning in StackOps**

Webinar (Online)

**Managing Data
Lifecycle**

● Apr 2026

Workshop (In-person)

**App Monitoring in
StackOps**

Webinar (Online)

**Smarter Troubleshooting
with StackOps AI Assistant**

● May 2026

Workshop (In-person)

**Visualisation &
Observability in
StackOps**

Upskill Webinars (Online)

This new series will cover short technical sessions focused on specific platform features

Technical Workshops (In-person)

Continuing with our classroom-based sessions with hands-on labs. From foundational ingestion to advanced machine learning, these instructor-led sessions equip you with the skills to transform raw data into actionable insights.

Learning Events Page

Refer to StackOps Learning Events page for latest updates on upcoming and past events:

<https://go.gov.sg/stackops-learning-events>



Agenda

Meaning of... Life(cycle)?

1. Scenarios
2. Where is your data stored?
3. How can I manage them?
4. How does it benefit me?
5. Demo
6. Questions & Answers

| What are we trying to tackle?

Scenarios

| Scenario #1

My team complains that searching of certain logs are slow.

How can I improve the search performance?

| Some Thoughts...

1. Where is my data stored?
2. Is the storage used slow?
3. Is there a storage that is faster?

| Scenario #2

I am already over the budget paying X amount for my data storage.

What can I do to scale back?

Warm

72.22% (1.07 TB of 1.48 TB)



Warm

26.25% (199.52 GB of 760 GB)

| Some Thoughts...

1. How can I retain my data elsewhere?
2. Is there a cheaper option to store my data?

| And the answer is...

Index Lifecycle Management

In short, it is all about
controlling your data movement

Data Storage Tiers

Hot

1. Fast Search
2. Small Storage
3. Expensive (\$\$\$)


Warm

1. Moderate Search
2. Medium Storage
3. Average (\$\$)

Frozen

- Slow Search w/ Cache
- Large Storage (S3)
- Affordable (\$)

Data Storage Nodes



Upskill with StackOps Series

Demo: Viewing Data Nodes

Viewing Data Nodes

 **Management**

Dev Tools

Integrations

Fleet

Osquery

Stack Monitoring

Stack Management

Cluster overview

stackopsdevsyste (fb5bc0)

 **Elasticsearch**

Overview 1 alert

Health Healthy
Version 8.19.8
Uptime 2 months
[Machine learning jobs 40](#)

Nodes: 6
Disk Available 71.67%
1.9 TB / 2.7 TB
JVM Heap 45.22%
16.0 GB / 35.4 GB

Indices: 3,014
Documents 5,474,155,397
Disk Usage 491.8 GB
Primary Shards 3,123
Replica Shards 1,074

Logs
Audit
Server 60 33 27

GET _cat/allocation?v=true&s=node:asc

GET _cat/nodes?v=true&s=name:asc

Optional params: &format=json

| How do we manage them?

Index Lifecycle Policy

In short, this is where you
define your data movement

elastic Find apps, content, and more.

Stack Management Index Lifecycle Man... Edit policy

Management

Ingest ⓘ
Ingest Pipelines
Logstash Pipelines

Data ⓘ
Index Management
Index Lifecycle Policies
Data Set Quality
Snapshot and Restore
Rollup Jobs
Transforms
Cross-Cluster Replication
Remote Clusters

Alerts and Insights ⓘ
Alerts
Rules
Cases
Connectors
Reporting
Machine Learning
Watcher

Edit policy stackops_baseline_policy [Documentation](#)

You are editing an existing policy. Any changes you make will affect 1,266 linked indices and 170 linked index templates that are attached to this policy. Alternatively, you can save these changes in a new policy.

☐ Save as new policy

Policy summary

This policy moves data through the following phases. [Learn about timing](#)

Hot phase Warm phase Frozen phase

Hot phase Required

Store your most recent, most frequently-searched data in the hot tier. The hot tier provides the best indexing and search performance by using the most powerful, expensive hardware.

[Advanced settings](#)

Warm phase

☒ Move data into phase when: 2 days [info](#)

Move data to the warm tier when you are still likely to search it, but infrequently need to update it. The warm tier is optimized for search performance over indexing performance.

[Advanced settings](#)

Cold phase

☐ Move data to the cold tier when you are searching it less often and don't need to update it. The cold tier is optimized for cost savings over search performance.

Frozen phase

☒ Move data into phase when: 28 days [info](#)

Move data to the frozen tier for long term retention. The frozen tier provides the most cost-effective way store your data and still be able to search it.

Searchable snapshot
Convert to a partially-mounted index that caches the index metadata. Data is retrieved from the snapshot as needed to process search requests. This minimizes the index footprint while keeping all of your data fully searchable. [Learn more](#)

Snapshot repository
found-snapshots
Each phase uses the same snapshot repository.

Delete data after this phase: ☐ [info](#)

Delete phase Remove

☐ Move data into phase when: 365 days [info](#)

Delete data you no longer need.

Wait for snapshot policy
Specify a snapshot policy to be executed before the deletion of the index. This ensures that a snapshot of the deleted index is available. [Learn more](#)

Policy name (optional)
Type and then hit "ENTER"

Delete searchable snapshot
Delete the searchable snapshot created in a previous phase. [Learn more](#)

☒ Delete searchable snapshot

[Save policy](#) [Cancel](#) [Show request](#)

1. This is a StackOps Baseline Policy
2. It is created by us to standardize the movement of all ingested data
3. We do not recommend editing it as this is managed by Automation
4. This is a component of StackOps Baseline

| Hmm.. what is that?


StackOps Baseline

It is a **set of configurations**
injected into your Elastic Deployment,
to **comply with organization requirements and best practices**

Upskill with StackOps Series

Demo: Viewing Index Lifecycle Policies


Viewing Index Lifecycle Policies



```
GET _ilm/policy
```

```
GET _ilm/policy/stackops_baseline_policy
```

Index Lifecycle Phases



Hot Phase

Rollover*

What does Rollover do?

1. Ensures indices comply to a certain size based on rollover conditions.
2. So, it can be as memory and storage performant as possible.





Hot Phase

Rollover*

What are the Rollover Conditions?

Hot phase

Rollover

Maximum primary shard size: **20gb**

EITHER OR

Maximum age: **14d**

*Index consist of Shards, Shards contain your Data (Terms Used Interchangeably)

*Books consist of Pages, Pages contain Words

Once one of the conditions are fulfilled, Rollover happens!

☐ .ds-filebeat-wenan-logs-aws.cloudtrail-
2026.04.13-000008

**Rollover to a
new index**

☐ .ds-filebeat-wenan-logs-aws.cloudtrail-
2026.04.15-000009

Hot Phase

Rollover*

Recommendations

1. Avoid Low Rollover Age Range .e.g. 1 – 3 Days
2. Causes Oversharding i.e. Memory Issue
3. Baseline Rollover Age: 14 Days
4. Avoid Rollover Size greater than 50GB
5. Causes Watermark Limit i.e. Storage Issue
6. Baseline Rollover Size: 20GB



Benefits the following scenarios:

1. Rollover happen less often, helps greatly for index with low ingestion
Lesser shard = Lesser Memory Usage
2. Index size is appropriately chunked for ease of data movement
Bite-sized data ease storing, helps greatly for smaller data storage

Warm Phase

Move Data to Warm Nodes

The More You Know

1. Data will only move to Warm Phase ONLY AFTER Rollover
2. Any Other Phases will not happen if there is no rollover

Hot phase

Rollover

Maximum primary shard size: 20gb
Maximum age: 14d **EITHER OR**

Warm phase

ONLY AFTER ROLLOVER

Move data into phase when


2d old

Example #1:

If an index takes 1 Day to ingest 20GB of Data,
it will only enter Warm Phase at the 3rd Day (1 + 2)

Example #2:

If an index takes 14 Days to ingest only 5GB of Data,
it will only enter Warm Phase at the 16th Day (14 + 2)



— Frozen phase

Move data into phase when

28d old

Searchable snapshot

Snapshot repository: `found-snapshots`

— Delete phase

Move data into phase when

366d old

Data are only stored in its respective nodes only after phases are complete!

Lifecycle policy

[Edit policy in ILM](#)

Policy name `so-analysis-bom-policy`

Current phase `Frozen`

Current action `complete`

Current action time `2025-08-26 16:00:43`

Current step `complete`

Upskill with StackOps Series

Demo: Viewing Lifecycle Phases

Viewing Lifecycle Phases

Indices **Data Streams** Index Templates Component Templates Enrich Policies

Data streams store time-series data across multiple indices and can be created from index templates. [Learn more.](#)

Q Search...	☐ Include stats	View 1	Reload
Name	Health	Indices	Index mode
☐ logs-aws.cloudtrail-default Managed	● green	83	LogsDB

.ds-logs-aws.cloudtrail-default-2026.03.15-000794

Overview Mappings Settings Statistics Index lifecycle

Q data_stream="logs-aws.cloudtrail-default"

Name	Health	Status
☐ .ds-logs-aws.cloudtrail-default-2026.03.15-000794	● green	open

Lifecycle policy

[Edit policy in ILM](#)

Policy name	stackops_baseline_policy
Current phase	Warm
Current action	complete
Current action time	2026-03-22 04:18:30
Current step	complete

GET index_name/_ilm/explain


| What's next..?

**How do we put all these
to good use?!?!?**




| Scenario #1 (Expanded)

My team (developer) complains that searching of certain (application) logs are slow.

Logs searched usually ranges between 7 - 14 days



Scenario #1 Further Investigation



— Warm phase

Move data into phase when
2d old

Logs 7 – 14 days old are probably stored
in the Slower Warm Nodes

| Some Thoughts...


1. Oldest log searched was 14 Days.
2. I can delay the movement into Warm Phase by that amount.
3. So those data can be retained in the Hot Node (Faster)
4. Does it make sense to change stackops_baseline_policy's Warm Phase?
5. But If you do, it will affect CloudTrail Logs too!
6. How can I do that without impacting my CloudTrail Logs?

| And the answer is...

Manage It With Another ILM

Through the usage of Index Template!

Step By Step



Upskill with StackOps Series

Final Demo: Manage It With Another ILM

Step #1: Configure a Different ILM Policy

Data ?

Index Management

Index Lifecycle Policies

Data Set Quality

Snapshot and Restore

Rollup Jobs

Transforms

Cross-Cluster Replication

Remote Clusters

+ Create policy

Name the policy e.g. xxx_policy

Set Warm Phase to 14 Days (Based on Scenario #1)

Can reference from stackops_baseline_policy

Step #2: Create Component Templates with that ILM

Data ?

Index Management

Index Lifecycle Policies

Data Set Quality

Snapshot and Restore

Rollup Jobs

Transforms

Cross-Cluster Replication

Remote Clusters

Indices

Data Streams

Index Templates

Component Templates

Enrich Policies

+ Create component template

Index settings (optional)

Define the behavior of your indices.

```
{  
  "index.lifecycle.name": "my_custom_policy"  
}
```

Step #3: Apply it to the Relevant Index Template



Component templates (optional)

Component templates let you save index settings, mappings and aliases and inherit from them in index templates.

[Component templates docs](#)

Components selected: 7

- = metrics@tsdb-settings M S A ⊖
- = metrics-awsfargate.task_stats@package M S A ⊖
- = metrics-awsfargate.task_stats@custom ⊖
- = ecs@mappings M S A ⊖
- = .fleet_globals-1 M S A ⊖
- = .fleet_agent_id_verification-1 M S A ⊖
- = metrics@custom ⊖

Step #1: Search

Q stackops-ilm@custom Filter 3 v

stackops-ilm@custom M S A Step #2: Add ⊕

stackops-prefer-ilm@custom M S A ⊕

6

Review template

✓ Create template

Step #4: Rollover to Apply Settings

Index Management

Indices **Data Streams** Index Templates Component Templates Enrich Policies

Data streams store time-series data across multiple indices and can be created from index templates

Q Search...			
☐	Name ↕	Health ↕	Indices
☐	logs-aws.cloudtrail-default	Managed	green 83

Management

Dev Tools

Integrations

Fleet


Osquery

Stack Monitoring

Stack Management

POST `datastream_name/_rollover`

What have we learnt?



Thank You!

“ Do your part and care for your Data ~

- *Someone responsible*

Share your feedback on the session

Please share your feedback on the session to help us improve in future!

<https://go.gov.sg/stackops-webinar-feedback>




<https://go.gov.sg/stackops-webinar-feedback>

| Additional Information...


Annex

Slides placed below are not presented
due to time constraint

Analogies



Hierarchy of Data Streams → Indices → Shards



*Data Stream are groups of similar Indices consisting of Shards which contain your Data (Terms Used Interchangeably)

*Series of Harry Potter Books consisting of Pages, which contain Words

| Hmm.. what is that?

StackOps Baseline

It is a set of configurations injected
into your Elastic Deployment,
to comply with organization requirements and best practices

StackOps Baseline

How do we standardize usage of stackops_baseline_policy?

1. Inject through *best-effort* cloning of Index Templates

Data ?

Index Management

Index Lifecycle Policies

Data Set Quality

Snapshot and Restore

Rollup Jobs

Transforms

Cross-Cluster Replication

Remote Clusters

Index Management

[Index Management docs](#)

Indices Data Streams Index Templates Component Templates Enrich Policies

Use composable index templates to automatically apply settings, mappings, and aliases to indices. [Learn more.](#)

View 1

Reload

logs-aws.cloudtrail

Create template

☐	Name ↑		Index patterns ↕	Component te... ↕	Data stream	Content	Actions
☐	logs-aws.cloudtrail	Managed	logs-aws.cloudtrail-*	8	✓	M S A	...
☐	logs-aws.cloudtrail-stackops	Managed	logs-aws.cloudtrail-*	11	✓	M S A	...

Rows per page: 50

< 1 >

StackOps Baseline (Cont.)

logs-aws.cloudtrail Managed

[Summary](#) [Settings](#) [Mappings](#) [Aliases](#) [Preview](#)

Index pattern	Data stream
logs-aws.cloudtrail-*	Yes
Priority	Version
200	None
Component templates	
logs@mappings	
logs@settings	
logs-aws.cloudtrail@package	
logs@custom	
logs-aws.cloudtrail@custom	
ecs@mappings	
.fleet_globals-1	
.fleet_agent_id_verification-1	

logs-aws.cloudtrail-stackops Managed

[Summary](#) [Settings](#) [Mappings](#) [Aliases](#) [Preview](#)

Index pattern	Data stream
logs-aws.cloudtrail-*	Yes
Priority	Version
210	None
Component templates	
logs@mappings	
logs@settings	
logs-aws.cloudtrail@package	
logs@custom	
logs-aws.cloudtrail@custom	
ecs@mappings	
.fleet_globals-1	
.fleet_agent_id_verification-1	
stackops-schema@custom	
stackops-ilm@custom	
stackops-es-security-user@custom	

stackops-ilm@custom

[Summary](#) [Settings](#) [Mappings](#) [Aliases](#)

```
{
  "index": {
    "lifecycle": {
      "name": "stackops_baseline_policy"
    }
  }
}
```

StackOps Baseline Gaps

Best Effort Cloning is an ad-hoc triggered only

1. Deployment Creation
2. Deployment 8.15 – 8.19 Upgrade

Index Template created in between will be **missed out**

Some Indices will not be attached with ***stackops_baseline_policy***

Additional Mitigation

Selected ILM are aligned with ***stackops_baseline_policy***

StackOps Baseline Improvements

Identify Un-cloned Templates

Scheduling of Best Effort Cloning

Challenges

Lacks Metadata to Identify Tenant-Owned Index Template

Impossible to Prevent Unintentional Cloning

Identify Index Templates w/o stackops_baseline_policy

Name ↕	Linked index templates ↕	Linked indices ↓	Modified date ↕	Actions
stackops_baseline_policy	224	53	Dec 23, 2025	...
metrics	88	47	Dec 23, 2025	...
logs	69	42	Dec 23, 2025	...

☐ Name ↕	Health ↕	Status ↕	Primaries ↕	Replicas ↕	Documents c... ↕	Storage size ↕	Data stream ↕
☐ .ds-metrics-azure.compute_vm-default-2025.12.29-000001	● green	open	1	1	354	2.06mb	metrics-azure.compute_vm-default

Identify Index Templates w/o stackops_baseline_policy (Cont.)

metrics-azure.compute_vm-default Managed

This data stream and its associated indices are managed by ILM

To edit data retention for this data stream, you must edit its associated ILM policy.

— GOOD —
Health ? **green**
Template has been cloned with -stackops

Last updated ?
April 12th, 2026 3:32:00

Index template ?
metrics-azure.compute_vm-stackops

metrics-awsfargate.task_stats-default Managed

This data stream and its associated indices are managed by ILM

To edit data retention for this data stream, you must edit its associated ILM policy.

— MISSED —
Index template may be added after cloning, hence it was missed

ILM isn't stackops_baseline_policy

Storage size ?
3.4gb

Indices ?
27

Timestamp field ?
@timestamp

Generation ?
150

Index template ?
metrics-awsfargate.task_stats

Index mode ?

Time series

Effective data retention ?


Disabled

Failure store ?

Disabled

Index lifecycle policy ?
metrics

Updating ILM by Cloning Index Templates




Component templates (optional)



[Component templates docs](#)

Component templates let you save index settings, mappings and aliases and inherit from them in index templates.

Step #1: Search



Filter 3



Replacing ILM in Existing Index

Data ?

Index Management

Index Lifecycle Policies

Data Set Quality

Snapshot and Restore

Rollup Jobs

Transforms

Cross-Cluster Replication

Remote Clusters

Management

Dev Tools

Integrations

Fleet

Osquery

Stack Monitoring

Stack Management

Name ↕	Linked index templates ↕	Linked indices ↓	Modified date ↕	Actions
stackops_baseline_policy	224	53	Dec 23, 2025	...
metrics	88	47	Dec 23, 2025	...
logs	69	42	Dec 23, 2025	...

review the response's in_use_by indices

```
GET _ilm/policy/metrics
```

identify the indices you want to update the ILM

detach the old ILM

```
POST index-name1/_ilm/remove
```

```
POST index-name2/_ilm/remove
```

#attach the new ILM

```
PUT index-name1,index-name2/_settings
```

```
{
```

```
  "index.lifecycle.name": "stackops_baseline_policy"
```

```
}
```

verify index is attached to the new ILM

```
GET index-name1/_ilm/explain
```

```
GET index-name2/_ilm/explain
```

Upskill with StackOps Series

Demo: Viewing Data Stream, Indices & Shards

Viewing Data Stream, Indices & Shards

Index Management

Index Lifecycle Policies

Data Set Quality

Snapshot and Restore

Rollup Jobs

Indices Data Streams Index Templates Component Templates Enrich Policies

Data streams store time-series data across multiple indices and can be created from index templates. [Learn more.](#)

Q

Search...

×

Include stats

View

1

Reload

Name

Health

Indices

Index mode

Data retention

Actions

logs-aws.cloudtrail-default

Managed

green

83

LogsDB

Disabled

Index Management

[Index Management docs](#)

Indices Data Streams Index Templates Component Templates Enrich Policies

Update your Elasticsearch indices individually or in bulk. [Learn more.](#)

☒ Include hidden indices
 ☐ Include rollout indices

6 indices have lifecycle errors

Q data_stream="logs-aws.cloudtrail-default"			Lifecycle status		Lifecycle phase		Reload indices	Create index
☐	Name ↑	Health ↕	Status ↕	Primarys ↕	Replicas ↕	Documents ... ↕	Storage size ↕	Data stream ↕
☐	.ds-logs-aws.cloudtrail-default-2026.03.15-000794	● green	open	1	1	19,692,035	40.17gb	logs-aws.cloudtrail-default
☐	.ds-logs-aws.cloudtrail-default-2026.03.19-000797	● green	open	1	1	19,804,053	40.23gb	logs-aws.cloudtrail-default

GET _data_stream

GET _cat/shards?v=true&s=node,index,store

Additional APIs

```
# Prepare Index Template
POST _index_template/wenan-data-stream
{
  "index_patterns": [
    "wenan-data-stream*"
  ],
  "template": {
    "settings": {
      "index": {
        "mode": "standard"
      }
    }
  },
  "composed_of": [],
  "data_stream": {
    "hidden": false,
    "allow_custom_routing": false
  }
}
GET _index_template/wenan-data-stream

# Simulate Data Ingestion by Inserting Data
POST wenan-data-stream/_doc
{"@timestamp":"2026-04-14T20:32:00+0800",
 "event.dataset":"simulate.ingestion", "author":"wenan"}
GET wenan-data-stream/_search

# Create ILM Policy
PUT _ilm/policy/wenan_baseline_policy
{
  "policy": {
    "phases": {
      "frozen": {
        "min_age": "28d",
        "actions": {
          "searchable_snapshot": {
            "snapshot_repository": "found-snapshots",
            "force_merge_index": true
          },
          "unfollow": {}
        }
      },
      "delete": {
        "min_age": "366d",
```

Refer to PDF for details:

ManageDataLifecycleUsingILM-PolicyInStackOps-APIs-17-April-2026