



Elasticsearch Diagnose & Index Lifecycle management Service


Pei Wang , Minyi Ding
July 2018

Pronto Overview

Built on an open source Elasticsearch(ES) Stack, Pronto powers eBay engineers/BU reliably and securely take data from any source, in any format, and **search, analyze, and visualize** in real time.

- ✓ **Millions** of buyers & sellers benefit from Elasticsearch powered by Pronto
- ✓ **60+** teams and **100+** Elasticsearch clusters, **3000+** nodes.
- ✓ **36Billion** daily doc ingestion and **7Billion** daily queries

Pronto in eBay tech stack





ElasticSearch Diagnose Service

Pei Wang
July 2018



Diagnose service

Why we build it & What it can do

- Summarize problems we've met in daily DevOps work.
- Run diagnose on target ElasticSearch cluster, send report if find any problem.
- Explain to cluster users what the potential risk is and what suggestion we have.
- Auto run without human involve.

Diagnose service


Key concept

- Pattern, define how to find the problem
- Knowledge base, include description, links and suggestion
- Job
- Report

Diagnose service

Work Flow

- User access UI page to manage job/knowledge base/pattern
- Rest service trigger diagnose process, gather information from data provider
- Analyzer generate report for every knowledge base



Diagnose service

Pattern

- Define how to find the problem
- Access cluster REST response as “property”
- Easy to customize
- Support multiple types

Edit

* Name

* Type

Parameter

Diagnose service

Pattern type

- Number
- String
- Script
- Customize

Edit

* Name

* Type

Parameter

Diagnose service

Knowledge base

- Related pattern
- Description
- Links
- Suggestions
- Help user understand the problem

Knowledge base List 🔔 5 Notifications ⚙️ Settings

[ADD](#)

#	Name	Severity	Description	Operations
1	Cluster status should be green	HIGH		 
<p>Pattern name: Cluster status should be green Pattern type: String</p> <p>Pattern params: {"val_path": "cluster_health.status", "equals": "green"}</p> <p>Links: Elastic Official site, Elastic Official site</p> <p>Support versions: 2.3.2 5.2.2 6.2.2</p>				
2	Minimum master node should be set to a quorum of all master eligible node	CRITICAL	minimum_master_node should be set to quorum of all master-eligible nodes	 
3	Shard size check	LOW		 
4	Replica number should bigger than 0	HIGH	Data loss if the machine host shard last	 
5	Maximum file descriptor check	HIGH	Elasticsearch uses a lot of file descriptors or file handles. Running out of file descriptors can ...	 
6	Master eligible node count check	CRITICAL	Cluster should have more than 1 master eligible nodes to avoid single failure, and the numb...	 
7	Shard allocation setting check	MEDIUM	Shard allocation is the process of allocating shards to nodes. This can happen during initial ...	 
8	Big cluster should have dedicated master node	MEDIUM		 

Diagnose service

Job

- A diagnose process
- Support authentication

Create new job

* Job Name

* Cluster Endpoint

Authentication Header

Label

Diagnose service

Job List

Job List 5 Notifications Settings

[ADD](#)

Id	Job Name	Status	Has Problem	Label	Report	Create At	End At	Operations
> 19396	es6appspeed-lvs02	COMPLETED	YES	201806301254	Report	07-03 07:07:32	07-03 07:07:32	🔍
> 19395	es6rdlelasticsearch-lvs02	COMPLETED	YES	201806301254	Report	07-03 07:07:30	07-03 07:07:30	🔍
> 19394	es6ufesmlx-lvs02	COMPLETED	NO	201806301254	Report	07-03 07:07:26	07-03 07:07:26	🔍
> 19393	es6maemlx-slc07	COMPLETED	NO	201806301254	Report	07-03 07:07:23	07-03 07:07:24	🔍
> 19392	es6crawler-lvs02	COMPLETED	NO	201806301254	Report	07-03 07:07:21	07-03 07:07:21	🔍
> 19391	es6bnteammonitoring-slc07	COMPLETED	YES	201806301254	Report	07-03 07:07:18	07-03 07:07:18	🔍
> 19390	es6ccssvc-lvs02	COMPLETED	YES	201806301254	Report	07-03 07:06:46	07-03 07:07:16	🔍
> 19389	es6umpevents-slc07	COMPLETED	NO	201806301254	Report	07-03 07:06:44	07-03 07:06:44	🔍
> 19388	es6umpevents-lvs02	COMPLETED	YES	201806301254	Report	07-03 07:06:28	07-03 07:06:43	🔍
> 19387	es6numsgda-slc07	COMPLETED	NO	201806301254	Report	07-03 07:06:26	07-03 07:06:26	🔍

10/page < 1 2 3 4 5 6 ... 1930 >

Diagnose service

Report

Job Report 🔔 5 Notifications - ⚙️ Settings

Job name: es6rcdlelasticsearch-lvs02

Pass count: 9, Not pass count: **1**, Exception count: 0

[Cluster Link](#) [Cerebro](#) Tue Jul 03 2018 07:07:30 GMT+0800

Shard size check NOT_PASS

Severity: LOW
Description:
Suggestion: Create new index with higher primary shard number, make sure shard size is balanced
Items have problem: [logs_write_6-2018.06.18-1](#)

Cluster status should be green PASS


Severity: HIGH
Description:
Links:
[Elastic Official site](#)
[Elastic Official site](#)

Minimum master node should be set to a quorum of all master eligible node PASS

Severity: CRITICAL
Description: minimum_master_node should be set to quorum of all master-eligible nodes
Links:
[Master Election](#)

Diagnose service

Report Overview





Index Lifecycle Management Service


Minyi Ding
July 2018



Index Lifecycle Management

- In broad terms, **Index Management** contains settings, operations and all the things during its lifecycle.

Alias	Index Settings
Allocation	Open
Close	Reindex
Cluster Routing	Replicas
Create Index	Restore
Delete Indices	Rollover
Delete Snapshots	Shrink
Forcemerge	Snapshot



Challenges

- Specialty required
- Daily production maintenance
- Struggling with improper usage, skewed data or unexpected ingestion peaks
- No backchannel from failed index ops
- No high availability

Index Lifecycle Management Service

What it can do

- High availability index lifecycle management with user-friendly web portal
- One-stop management portal for multiple clusters and multiple tenants

Why we build it

- Help customers manage ES data through its full lifecycle
- Let important business applications survive from peak days
- Avoid improper ES usage.

Function	Curator	_rollover	Index Management Service
Index Action	All supported	Create, Alias rotation	Create, Close, Delete, Index settings, Alias rotation
Version Compatibility On ES 2.x 5.x 6.x	Not fully covered	Not fully covered	All covered
Job Scheduling	No	No	Yes
Web Portal	No	No	Yes
High Availability	No	No	Yes
Multi-cluster Support	No	No	Yes

DASHBOARD

CLUSTERS

RULES

JOBS

Cluster Name	Cluster Auth Header	Endpoint	Operations
poc622rc6perf-slc07		http://10.189.211.201:80	
esmon		http://esmonitor-9aqqz.vip.stratus.slc.ebay.com:80	
slrinvsvc-slc07		http://slrinvsvc-slc07-write-vip-945nm.vip.stratus.slc.ebay.com:80	
es5c3metrics-slc07		http://es5c3metrics-slc07-read-vip-l8u2w.vip.stratus.slc.ebay.com:80	
zoom-slc07		http://zoom-slc07-write-vip-9bmfg.vip.stratus.slc.ebay.com:80	
es6ccssvc-lvs02		https://es6ccssvc-lvs02-read-secured-vip-coirx.vip.stratus.lvs.ebay.com:443	
es5comsorderhistory-lvs02		http://es5comsorderhistory-lvs02-write-vip-qd66y.vip.stratus.lvs.ebay.com:80	

[Register Cluster](#)

DASHBOARD

CLUSTERS

RULES

JOBS

Index Prefix	Cluster Name	Status	Total # of Alias	Operations
> perf-evaluation	poc622rc6perf-slc07	green	0	
> perftest-0	poc622rc6perf-slc07	green	0	
> clustermetrics	esmon	green	1	
> etl_log	esmon	green	1	
> externalmetrics	esmon	green	1	
> irislog	slrinvsvc-slc07	green	0	
> vm_availability_details	es5c3metrics-slc07	green	1	
> zpoller-downloads	zoom-slc07	green	8	
> buysvc_data_	es6ccssvc-lvs02	green	1	
> buysvc_service_status_	es6ccssvc-lvs02	green	1	
> ccs_data_	es6ccssvc-lvs02	green	1	
> ccs_service_status_	es6ccssvc-lvs02	yellow	1	

[Create Rule](#)

M Index Management

DASHBOARD

CLUSTERS

RULES

JOBS

All cluster status

Name: ccs_service_status_2018_06_24 Data Size: 240435.2MB	Name: ccs_service_status_2018_05_21 Data Size: 6243.2MB	Name: ccs_service_status_2018_07_01 Data Size: 6324.8MB	Name: ccs_service_status_2018_06_18 Data Size: 60998.4MB	Name: ccs_service_status_2018_06_08_daily Data Size: 66968.8MB	
	Name: ccs_service_status_2018_06_11 Data Size: 60179.2MB	Name: ccs_service_status_2018_06_28 Data Size: 76595.2MB	Name: ccs_service_status_2018_05_16 Data Size: 76185.6MB	Name: ccs_service_status_2018_05_08 Data Size: 76083.2MB	
	Name: ccs_service_status_2018_07_03 Data Size: 60076.8MB	Name: ccs_service_status_2018_05_20 Data Size: 75679.4MB	Name: ccs_service_status_2018_05_17 Data Size: 72908.8MB	Name: ccs_service_status_2018_05_23 Data Size: 72601.6MB	
Name: ccs_service_status_2018_06_20 Data Size: 194150.4MB	Name: ccs_service_status_2018_06_19 Data Size: 78448MB	Name: ccs_service_status_2018_06_03 Data Size: 75776MB	Name: ccs_service_status_2018_05_27 Data Size: 72294.4MB	Name: ccs_service_status_2018_06_30 Data Size: 67481.6MB	
	Name: ccs_service_status_2018_06_17 Data Size: 78336MB	Name: ccs_service_status_2018_05_30 Data Size: 75673.6MB	Name: ccs_service_status_2018_06_01 Data Size: 71782.4MB	Name: ccs_service_status_2018_05_09 Data Size: 67276.8MB	Name: ccs_service_status_2018_05_21 Data Size: 60825.6MB
Name: ccs_service_status_2018_06_26 Data Size: 151347.2MB	Name: ccs_service_status_2018_06_21 Data Size: 78396MB	Name: ccs_service_status_2018_05_06 Data Size: 75571.2MB	Name: ccs_service_status_2018_06_04 Data Size: 71168MB	Name: ccs_service_status_2018_05_25 Data Size: 66867.2MB	Name: ccs_service_status_2018_05_11 Data Size: 66217.6MB
	Name: ccs_service_status_2018_06_05 Data Size: 77926.4MB	Name: ccs_service_status_2018_07_05 Data Size: 74956.8MB	Name: ccs_service_status_2018_05_24 Data Size: 71095.6MB	Name: ccs_service_status_2018_06_15 Data Size: 66867.2MB	Name: ccs_service_status_2018_05_22 Data Size: 66217.6MB
Name: ccs_service_status_2018_06_27 Data Size: 127283.2MB	Name: ccs_service_status_2018_06_07 Data Size: 77926.4MB	Name: ccs_service_status_2018_06_07 Data Size: 74854.4MB	Name: ccs_service_status_2018_05_10 Data Size: 70246.4MB	Name: ccs_service_status_2018_06_23 Data Size: 65423.6MB	Name: ccs_service_status_2018_05_10 Data Size: 6038.4MB
	Name: ccs_service_status_2018_06_10 Data Size: 7731.2MB	Name: ccs_service_status_2018_06_14 Data Size: 74547.2MB	Name: ccs_service_status_2018_05_15 Data Size: 69636.8MB	Name: ccs_service_status_2018_05_18 Data Size: 65311.2MB	Name: ccs_service_status_2018_05_09 Data Size: 3996.8MB
Name: ccs_service_status_2018_06_06 Data Size: 88473.6MB	Name: ccs_service_status_2018_06_29 Data Size: 77209.6MB	Name: ccs_service_status_2018_06_31 Data Size: 73830.4MB	Name: ccs_service_status_2018_06_08 Data Size: 69734.4MB	Name: ccs_service_status_2018_06_16 Data Size: 65024MB	Name: ccs_service_status_2018_05_08 Data Size: 25800MB
	Name: ccs_service_status_2018_06_29 Data Size: 84787.2MB	Name: ccs_service_status_2018_06_13 Data Size: 76800MB	Name: ccs_service_status_2018_07_04 Data Size: 73523.2MB	Name: ccs_service_status_2018_06_02 Data Size: 69120MB	Name: ccs_service_status_2018_05_05 Data Size: 62771.2MB
Name: ccs_service_status_2018_07_02 Data Size: 84684.8MB	Name: ccs_service_status_2018_05_22 Data Size: 76697.8MB	Name: ccs_service_status_2018_05_28 Data Size: 73011.2MB	Name: ccs_service_status_2018_06_22 Data Size: 68300.8MB	Name: ccs_service_status_2018_05_26 Data Size: 61952MB	Name: ccs_service_status_2018_05_04 Data Size: 6077.6MB

Designed by Pronto Team.

M Index Management

- DASHBOARD
- CLUSTERS
- RULES**
- JOBS

Rule Add

test

Cluster Name:

Index Prefix:

Timezone Setting:

Alias Patterns: read_all write +

Last: ⊗

Index Rotation Strategies: create close purge setting_change

Creation Type:

Index Creation Interval:

Max size per shard:

Default ES settings:

Close indices older than:

Purge indices older than:

Setting change rules:

Change settings to below after ⊕ ⊗

Index Management Rule

Cluster Name:	<input type="text" value="Please select cluster"/>
Index Prefix:	<input type="text" value="require not to be empty"/>
Timezone Setting:	<input type="text" value="-0700"/>

Alias Rotation

Alias Patterns: read_all write +

Last	<input type="text" value="5"/>	<input type="text" value="Please select time unit ^"/>	<input type="button" value="x"/>
Last	<input type="text"/>	<input type="text" value="day"/>	<input type="button" value="x"/>
Last	<input type="text"/>	<input type="text" value="week"/>	<input type="button" value="x"/>
Last	<input type="text"/>	<input type="text" value="month"/>	<input type="button" value="x"/>

Action: Index Creation

Index Rotation Strategies: create close purge setting_change

Creation Type: hybrid

Index Creation Interval: 1 Please select time unit


Max size per shard: 30 GB

Default ES settings: require json format


Submit Cancel

Hybrid vs Time Based Creation

Time Based Rotation



Hybrid Rotation



- Built on monitoring on shard size of target indices
- No worries to configure primary shard number for daily indices

Action: Index Creation

Index Rotation Strategies: create close purge setting_change

Creation Type: hybrid

Index Creation Interval: 1 Please select time unit

Max size per shard: 30 GB


Default ES settings: require json format

Submit Cancel

Features

Query a bunch of indices by:

- Index Pattern
- Index Alias
- Date Math Expressions



Action: Index Close and Delete

Index Rotation Strategies: create close purge setting_change

Close indices older than: Please select time unit ^

Submit Cancel

- day
- week
- month

Index Rotation Strategies: create close purge setting_change

Purge indices older than: Please select time unit ^

Submit Cancel

- day
- week
- month

Action: Index Settings

Index Rotation Strategies: create close purge setting_change

Setting change rules:


Change settings to below after 3 week

```
{
  "index": {
    "routing": {
      "allocation": {
        "require": {
          "box_type": "warm"
        }
      }
    }
  }
}
```

Feature Summary

Business Specifications

- Index Creation
 - Condition :Time, Size, Shard-Size
 - Supported time unit : hour, day, week, month
 - Supported size unit : GB
- Index Close
- Index Deletion
- Index Setting Change
 - Condition :Time
 - Supported time unit : hour, day, week, month
 - Use case: hot-warm, replica number, etc.
- Alias Rotation



The screenshot shows the 'Rule Add' configuration window. It includes the following fields and options:

- Cluster Name:** A dropdown menu with the text 'Please select cluster'.
- Index Prefix:** A text input field containing 'require not to be empty'.
- Timezone Setting:** A text input field containing '-0700'.
- Alias Patterns:** A list of checkboxes. 'read_all' is checked, and 'write' is also checked. There is a blue plus icon next to 'read_all'.
- Index Rotation:** A set of radio buttons. 'create' is selected (checked), while 'close', 'purge', and 'setting_change' are unselected.
- Strategies:**
 - Creation Type:** A dropdown menu with 'hybrid' selected.
 - Index Creation Interval:** A numeric input field with '1' and a 'Please select time unit' dropdown.
 - Max size per shard:** A numeric input field with '30' and a 'GB' unit dropdown.
 - Default ES settings:** A text input field containing 'require json format'.

At the bottom of the form are two buttons: 'Submit' (in blue) and 'Cancel' (in white).

Feature Summary


Engineering Specifications

- Multiple cluster support
- Production level scheduler and HA deployment
- Index health dashboard
- Time zone support for different use cases
- Alerting with back-channel*
- Rotation job viewer*

Future Enhancement

- Backup and restore automation

<input type="checkbox"/>	Service Instance	Health ?	Traffic ?	Monitored ?
<input type="checkbox"/>	phxidmgmtwrapper-2219417.stratus.phx.ebay.com	Healthy	ON	ON
<input type="checkbox"/>	phxidmgmtwrapper-2219419.stratus.phx.ebay.com	Healthy	ON	ON
<input type="checkbox"/>	phxidmgmtwrapper-2238569.stratus.phx.ebay.com	Healthy	ON	ON
<input type="checkbox"/>	lvsidmgmtwrapper-2389878.stratus.lvs.ebay.com	Healthy	ON	ON
<input type="checkbox"/>	lvsidmgmtwrapper-2392313.stratus.lvs.ebay.com	Healthy	ON	ON
<input type="checkbox"/>	slcidmgmtwrapper-1916296.stratus.slc.ebay.com	Healthy	ON	ON
<input type="checkbox"/>	slcidmgmtwrapper-1946971.stratus.slc.ebay.com	Healthy	ON	ON
<input type="checkbox"/>	lvsidmgmtwrapper-2521194.stratus.lvs.ebay.com	Healthy	ON	ON
<input type="checkbox"/>	slcidmgmtwrapper-1959003.stratus.slc.ebay.com	Healthy	ON	ON



Q&A

Open source soon



eBay ELK开源爱好者群



该二维码7天内(7月30日前)有效, 重新进入将更新



elastic 中文社区

专业、垂直、纯粹的 Elastic 开源技术交流社区
<https://elasticsearch.cn/>