Intelligent Alert system for HEP experiments

Rahul Indra

Mentors
Valentin Kuznetsov
Federica Legger
Christian Ariza

Agenda

- About Me
1. Overview
2. Proposed Architecture
3. Alerting Module
4. Alerting Services
5. AlertManager - one place for all alerts
6. Use of Slack & Karma
7. Intelligence Module
8. Future Work
9. Tools Used
10. Important Links
- Appendix

About Me

Rahul Indra
Computer Science Undergraduate Student
Indian Institute of Engineering Science & Technology
Shibpur, India

Google Summer of Code ‘20
Student Developer
@ CERN-CMS

1. Overview

“The growth of distributed services introduces a challenge to properly monitor their status and reduce operational costs.”

Tools in use :-

- ElasticSearch
- Kafka
- Grafana
- Prometheus
- AlertManager
- VictoriaMetrics
- Custom Solutions like GGUS, SSB system etc.

CMS infrastructure can produce significant amount of data on :-

- various anomalies
- intermittent problems
- outages
- scheduled maintenance.

So, in short our operational teams deal with a large amount of alert notifications and tickets !

Solution

An intelligent Alert Management System

Aim

- detect
- analyse
- spot anomalies
- silence false alerts
- automate operation procedures

The system’s abilities include, but are not limited to :-

- Consuming tickets from various ticketing systems. (GGUS & SSB have been implemented). Being modular architecture, there’s always a scope to add more services in future.
- Extracting alerts, relevant to the specific CMS services which gets affected by such interventions
- Intelligently grouping and ranking those alerts.
- Silencing false alerts.
- Making them visible in our monitoring tools (Grafana, Slack, Karma etc.).

2. Proposed Architecture

Components Developed

- Parser
- Alerting Module
- Alerting Service
- Intelligence Module
- Alert CLI Tool

Tools

- Grafana
- Prometheus
- AlertManager
- Slack
- Karma

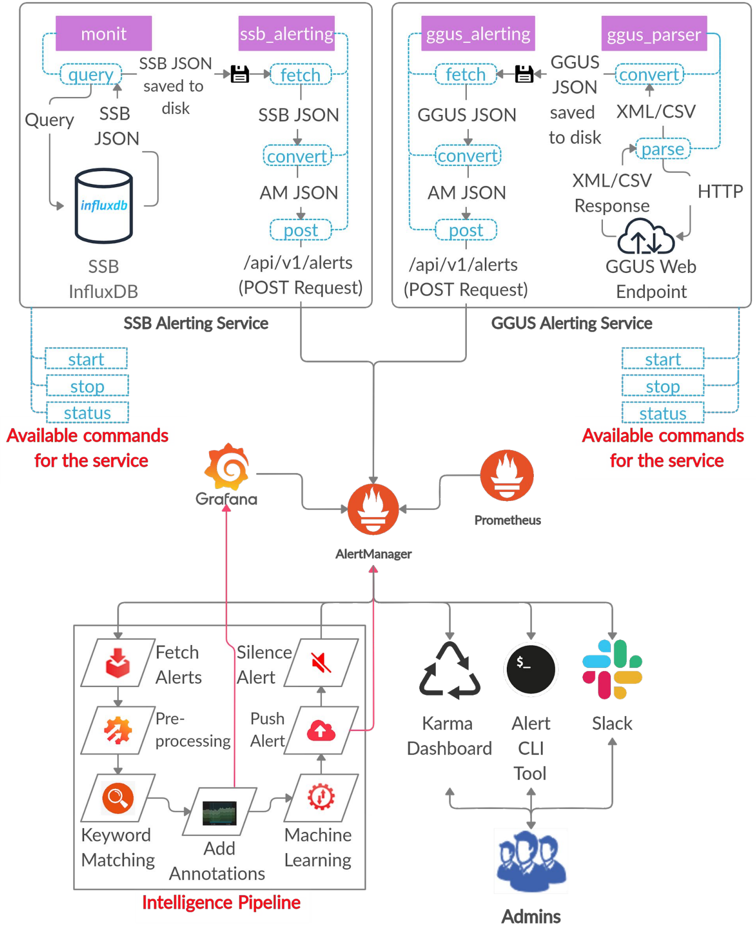

4. Alerting Service

- Parser fetches data and saves to disk
- Alerting module gets fetched data as input, converts it and pushes to AM.
- This whole process is bundled as a Linux Service with three commands :-
  - start
  - stop
  - status

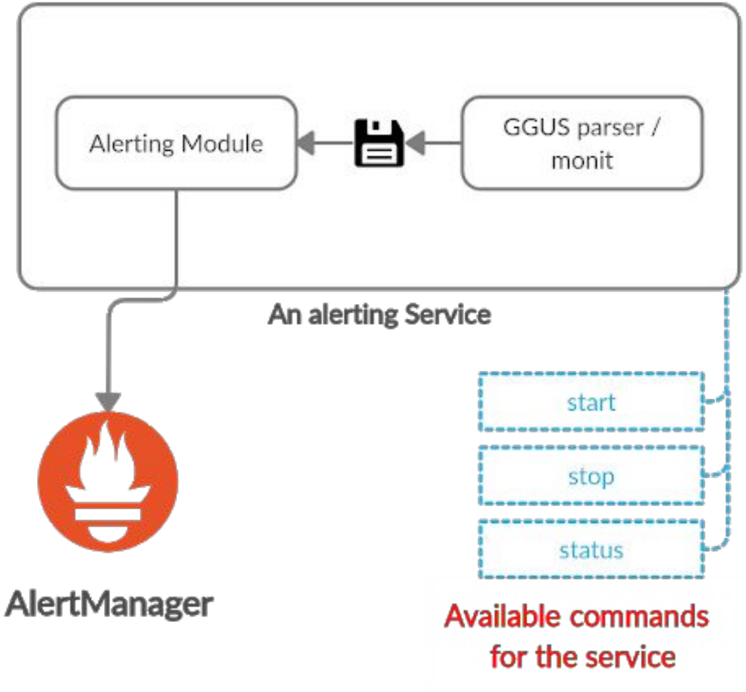

5. AlertManager - one place for all alerts

- Alerting services which has been developed push GGUS & SSB alerts to AM at defined time interval.
- Grafana & Prometheus push their alerts to AM as well.
- Karma Dashboard fetches all alerts from AM, and displays in better format.
- Slack channels are populated when an alert is fired.
- AM, Slack and Karma give all required info for alerts to our Admins.

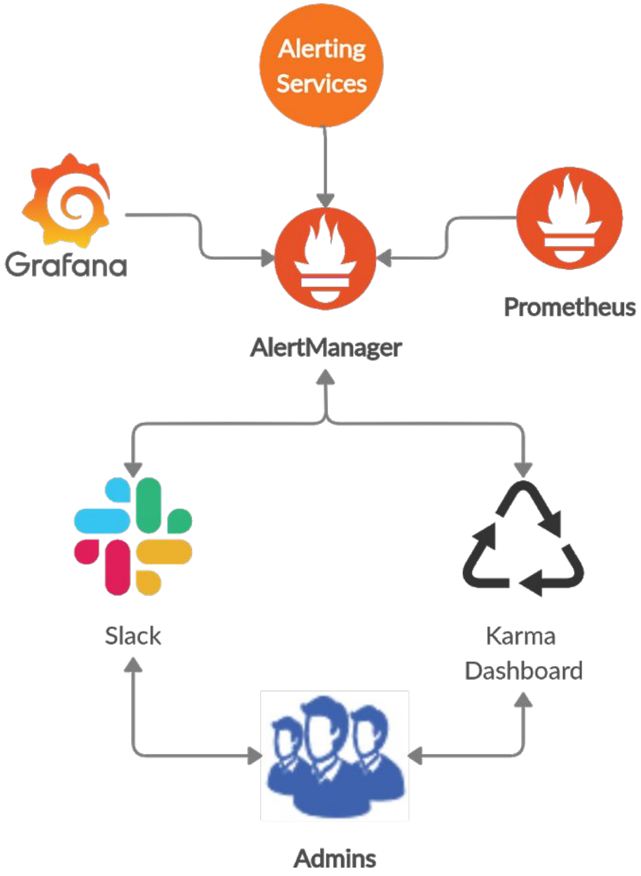

6. Use of Slack & Karma

Slack

- Slack has defined channels for particular service alerts.
- Users are notified about fired alerts.
- AlertManager bots are at work.

Karma

- A dashboard which pulls all alerts from AM.
- Availability of multi grids arrangement based on filters.
- Bundling similar alerts
- Concise and better view than AM.
- Wrote Dockerfile and Kubernetes config files.

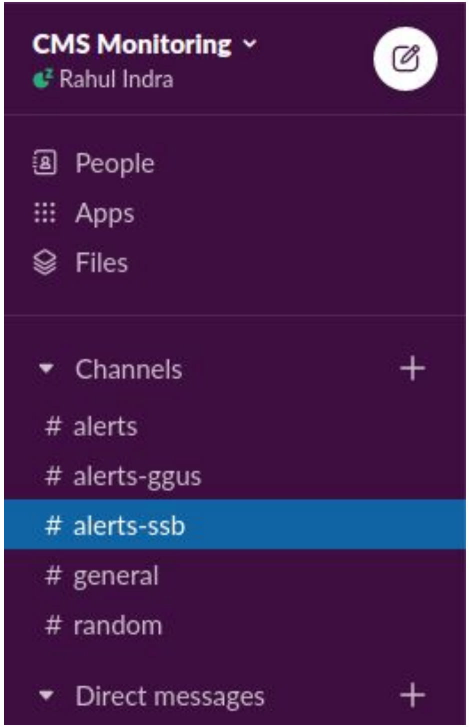

Karma Dashboard showing all alerts under “tag=monitoring” (GGUS)

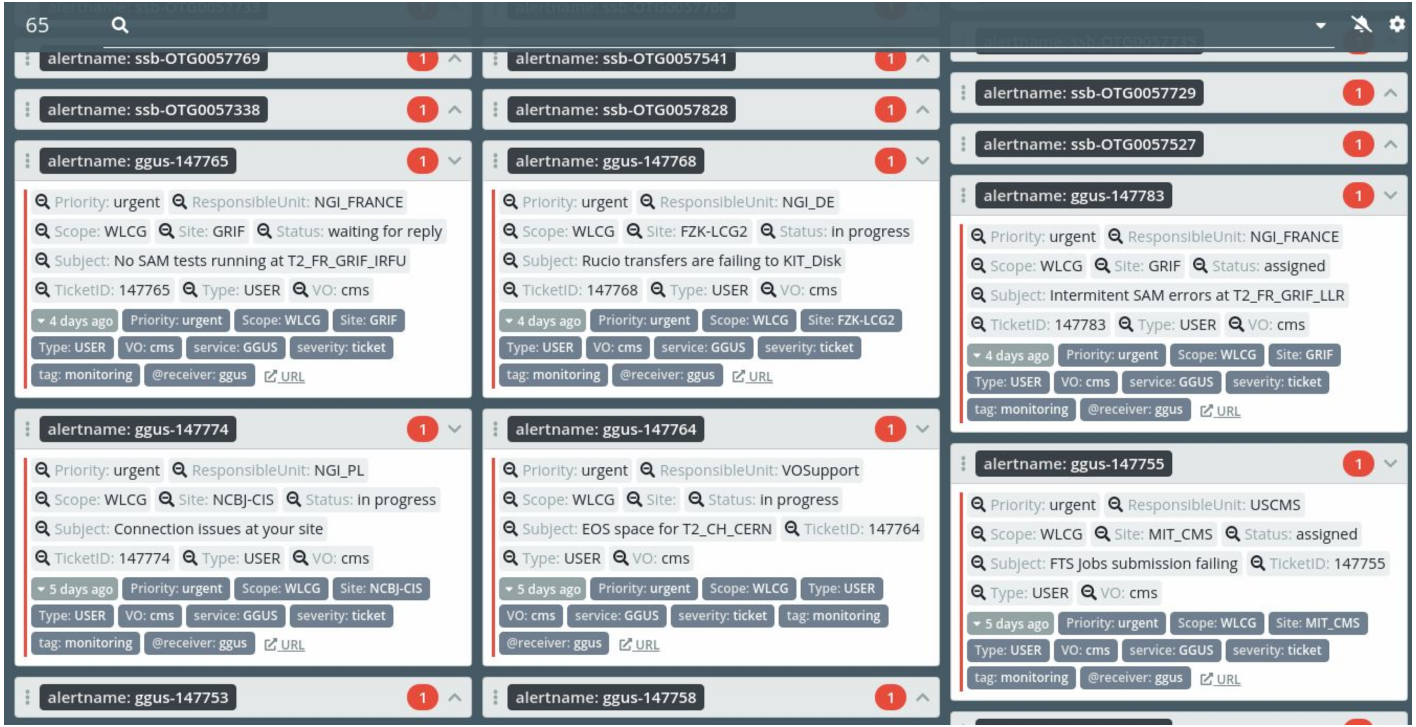

65

alertname: ssb-OTG0057769 1

alertname: ssb-OTG0057338 1

alertname: ggus-147765 1

Priority: urgent · ResponsibleUnit: NGI_FRANCE
Scope: WLCG · Site: GRIF · Status: waiting for reply
Subject: No SAM tests running at T2_FR_GRIF_IRFU
TicketID: 147765 · Type: USER · VO: cms
4 days ago · Priority: urgent · Scope: WLCG · Site: GRIF
Type: USER · VO: cms · service: GGUS · severity: ticket
tag: monitoring · @receiver: ggus · URL

alertname: ggus-147774 1

Priority: urgent · ResponsibleUnit: NGI_PL
Scope: WLCG · Site: NCBJ-CIS · Status: in progress
Subject: Connection issues at your site
TicketID: 147774 · Type: USER · VO: cms
5 days ago · Priority: urgent · Scope: WLCG · Site: NCBJ-CIS
Type: USER · VO: cms · service: GGUS · severity: ticket
tag: monitoring · @receiver: ggus · URL

alertname: ggus-147753 1

alertname: ssb-OTG0057541 1

alertname: ssb-OTG0057828 1

alertname: ggus-147768 1

Priority: urgent · ResponsibleUnit: NGI_DE
Scope: WLCG · Site: FZK-LCG2 · Status: in progress
Subject: Rucio transfers are failing to KIT_Disk
TicketID: 147768 · Type: USER · VO: cms
4 days ago · Priority: urgent · Scope: WLCG · Site: FZK-LCG2
Type: USER · VO: cms · service: GGUS · severity: ticket
tag: monitoring · @receiver: ggus · URL

alertname: ggus-147764 1

Priority: urgent · ResponsibleUnit: VOSupport
Scope: WLCG · Site: · Status: in progress
Subject: EOS space for T2_CH_CERN · TicketID: 147764
Type: USER · VO: cms
5 days ago · Priority: urgent · Scope: WLCG · Type: USER
VO: cms · service: GGUS · severity: ticket · tag: monitoring
@receiver: ggus · URL

alertname: ggus-147758 1

alertname: ssb-OTG0057729 1

alertname: ssb-OTG0057527 1

alertname: ggus-147783 1

Priority: urgent · ResponsibleUnit: NGI_FRANCE
Scope: WLCG · Site: GRIF · Status: assigned
Subject: Intermitent SAM errors at T2_FR_GRIF_LLR
TicketID: 147783 · Type: USER · VO: cms
4 days ago · Priority: urgent · Scope: WLCG · Site: GRIF
Type: USER · VO: cms · service: GGUS · severity: ticket
tag: monitoring · @receiver: ggus · URL

alertname: ggus-147755 1

Priority: urgent · ResponsibleUnit: USCMS
Scope: WLCG · Site: MIT_CMS · Status: assigned
Subject: FTS Jobs submission failing · TicketID: 147755
Type: USER · VO: cms
5 days ago · Priority: urgent · Scope: WLCG · Site: MIT_CMS
Type: USER · VO: cms · service: GGUS · severity: ticket
tag: monitoring · @receiver: ggus · URL

7. Intelligence Module

- A data pipeline.
- Components independent of each other.
- One component receives the data, adds its logic and forwards the processed data to other component.

Why data pipeline ?

- Low coupling
- Freedom of adding or removing components on demand.
- Power of concurrency

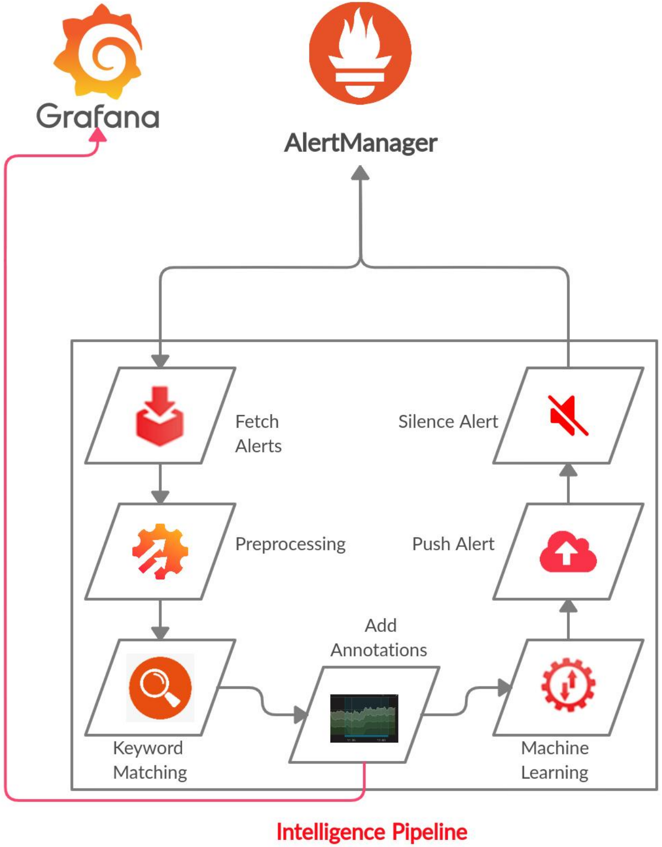

What it does ?

- Assigning proper severity levels to SSB/GGUS alerts which helps operators to understand the criticality of the infrastructure.
  Ex. If Number of Alerts with severity=”urgent” > some threshold, then the infrastructure is in critical situation.
- Annotating Grafana Dashboards when Network or Database interventions.

Scope for additional features include, but are not limited to :-

- Predicting type of alerts and grouping similar alerts with the help of Machine Learning.
- Adds applicable tutorial/instructions doc to alert, on following which an operator can solve the issue.

Let’s watch Intelligence Module live..

https://www.youtube.com/watch?v=vhJ367jaxMo

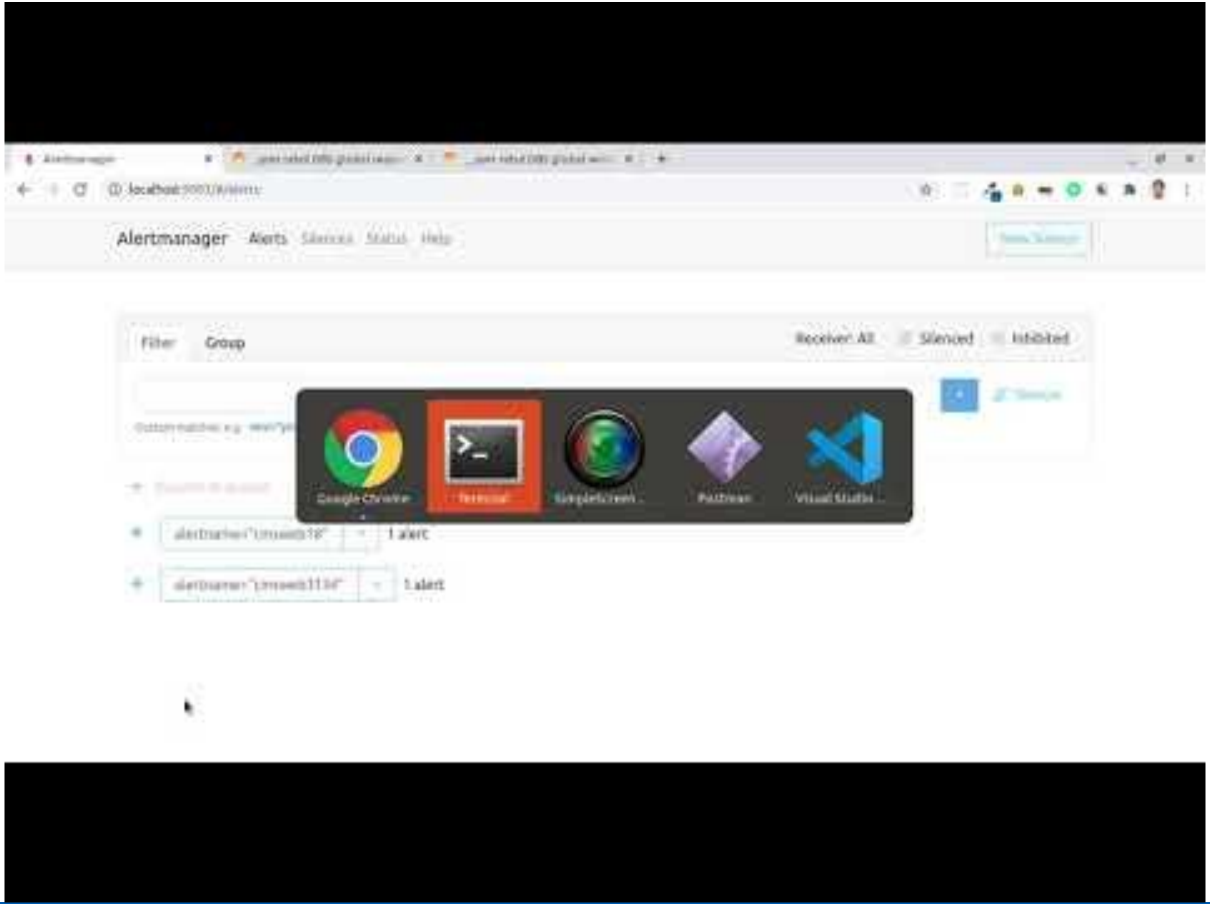
Alertmanager
localhost:9093/#/alerts
Alertmanager
Alerts
Silences
Status
Help
New Silence
Filter
Group
Receiver: All
Silenced
Inhibited
Google Chrome
Terminal
SimpleScreen...
Postman
Visual Studio...
1 alert
1 alert

8. Future Works

- Evaluation of ElastAlert for setting alerts on ElasticSearch and integration of the same in this project.
- Service which takes configuration for operator’s actions and pushes to AM so that it matches alerts with the actions.
- Use of Machine Learning in intelligence module which will predict it’s severity info, priority and type.
- Deployment of finalized project to k8s infrastructure.

9. Tools Used

Programming Language

- GoLang

Editor

- Vim
- Visual Studio Code

Helper Tools

- Github
- git CLI Tool
- golint, goreportcard.com
- Adobe Photoshop
- Google Doc
- Google Slides

10. Important Links

Github repository

https://github.com/dmwm/CMSMonitoring

Contributions in :-

https://github.com/dmwm/CMSMonitoring/tree/master/scripts
https://github.com/dmwm/CMSMonitoring/tree/master/src/go/MONIT
https://github.com/dmwm/CMSMonitoring/tree/master/src/go/intelligence
https://github.com/dmwm/CMSMonitoring/tree/master/doc/AlertManagement

GSoC Progress Report

Thank You !

Appendix

A. Parsers
   a. GGUS Parser
   b. monit
B. Alerting Module
C. Alerting Service
D. Slack & Karma
E. Intelligence Module
F. Alert CLI Tool

A. Parsers

- GGUS Ticketing System outputs data either in XML or CSV.
- Developed Parser capable of parsing both formats.
- ggus_parser has two components :-
    - parse - parses the XML or CSV data
    - convert - converts the parsed data
      into JSON format and saves it to disk.
- XML/CSV formats are configurable

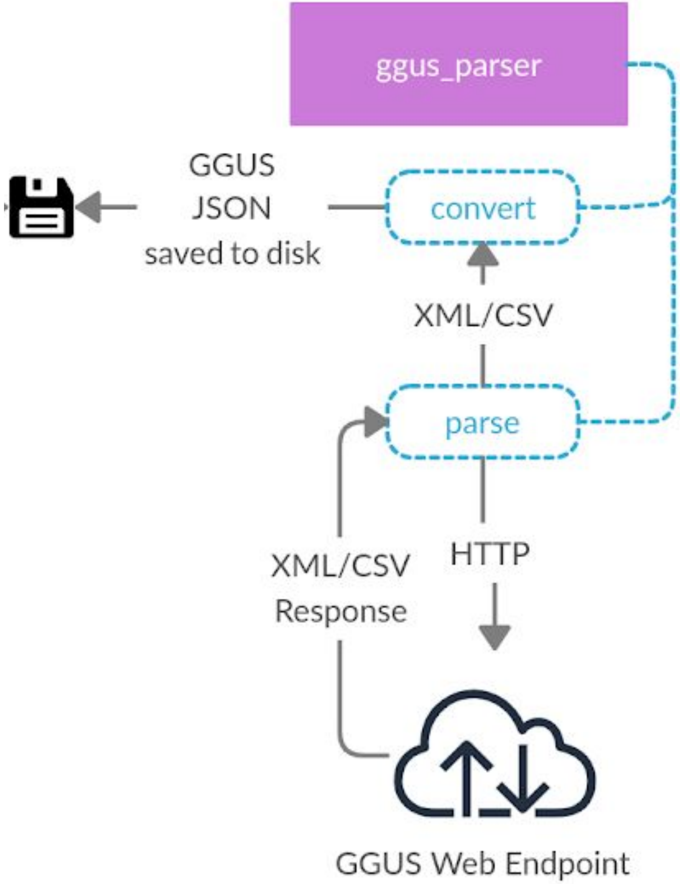

GGUS Ticket (csv)

Ticket-ID,Type,VO,Site,Priority,Resp. Unit,Status,Last Update,Subject,Scope
147196,USER,cms,FZK-LCG2,urgent,NGI_DE,assigned,2020-07-14,FZK-LCG2: issues on data access,WLCG

Which is Parsed and Converted into …..

GGUS Parsed Ticket (JSON)

```json
{
        "TicketID": 147196,
        "Type": "USER",
        "VO": "cms",
        "Site": "FZK-LCG2",
        "Priority": "urgent",
        "ResponsibleUnit": "NGI_DE",
        "Status": "assigned",
        "LastUpdate": "1590670920",
        "Subject": "FZK-LCG2: issues on data access",
        "Scope": "WLCG"
}
```

What about SSB Ticketing System ?

- There was no need of parser for SSB Ticketing System.
- monit tool was developed by CMS.
- Query InfluxDB/ES data sources in MONIT via Grafana proxy
- SSB alerts in JSON format is given on standard output.
- We piped stdout to .json file and saved to disk.

Ref :-
https://github.com/dmwm/CMSMonitoring/blob/master/src/go/MONIT/monit.go

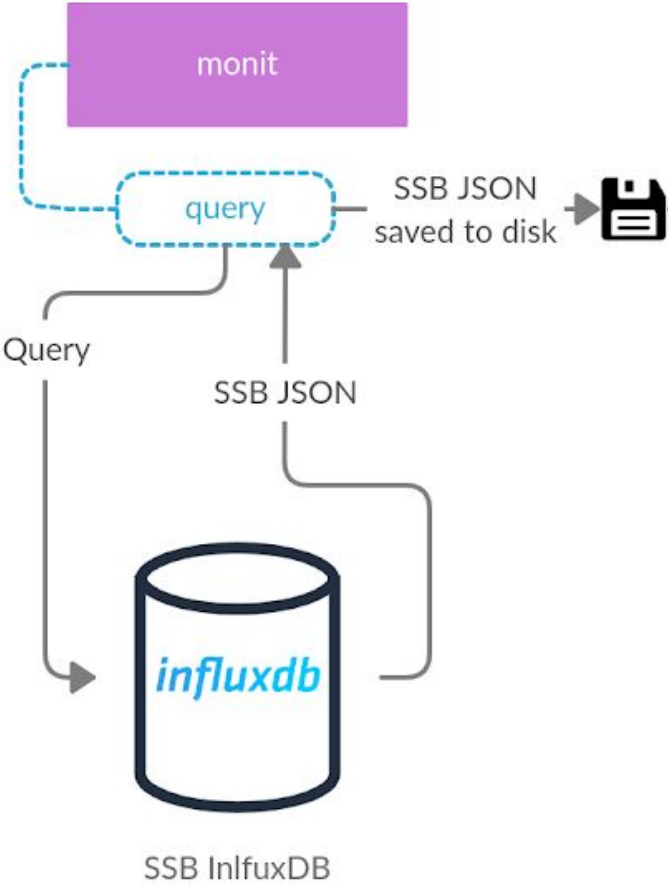

MONIT Query

monit -query=$query -dbname=$dbname -token=$token -dbid=$dbid > ssb_data.json

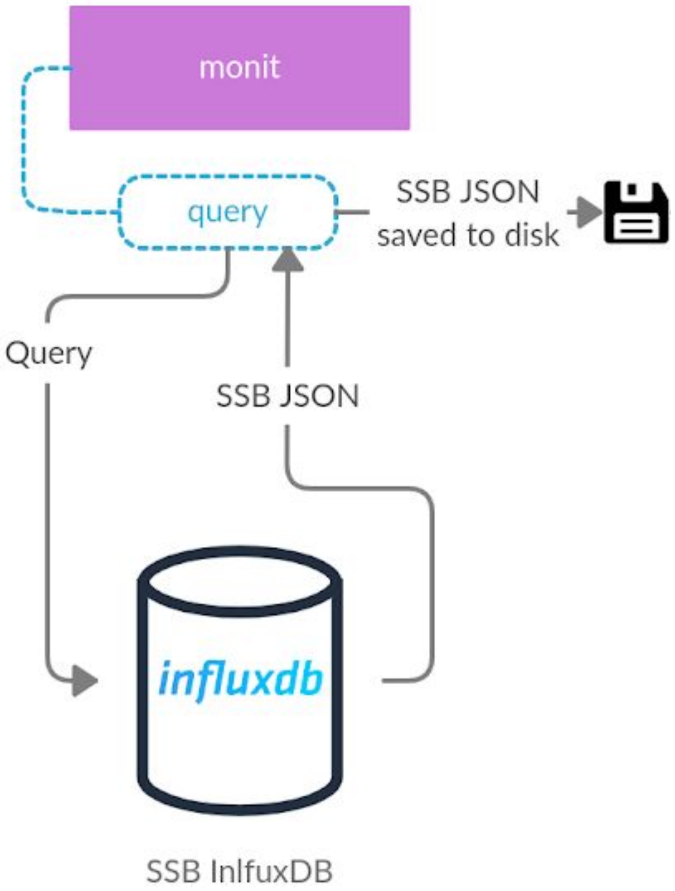

B. Alerting Module

Components Developed

- fetch
- convert
- post
- get
- delete

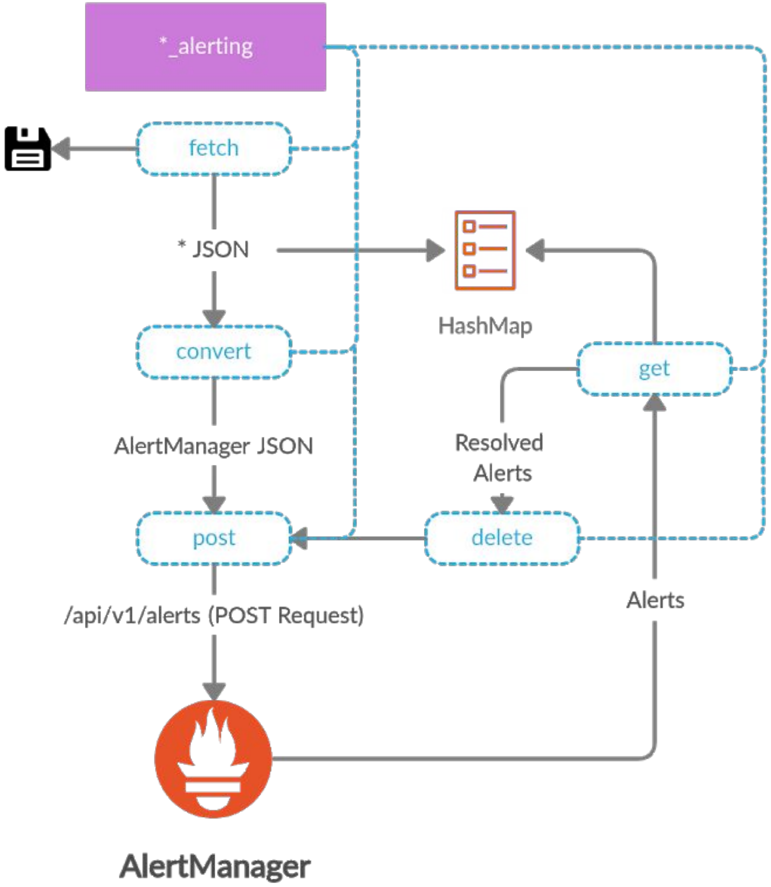

- fetch
  - fetches saved JSON GGUS or SSB data from the disk (ggus_parser or monit)
  - maintains a hashmap for seen alerts
  - map[alert_name] = alert

*now onwards we will call each datapoint from GGUS/SSB as an alert

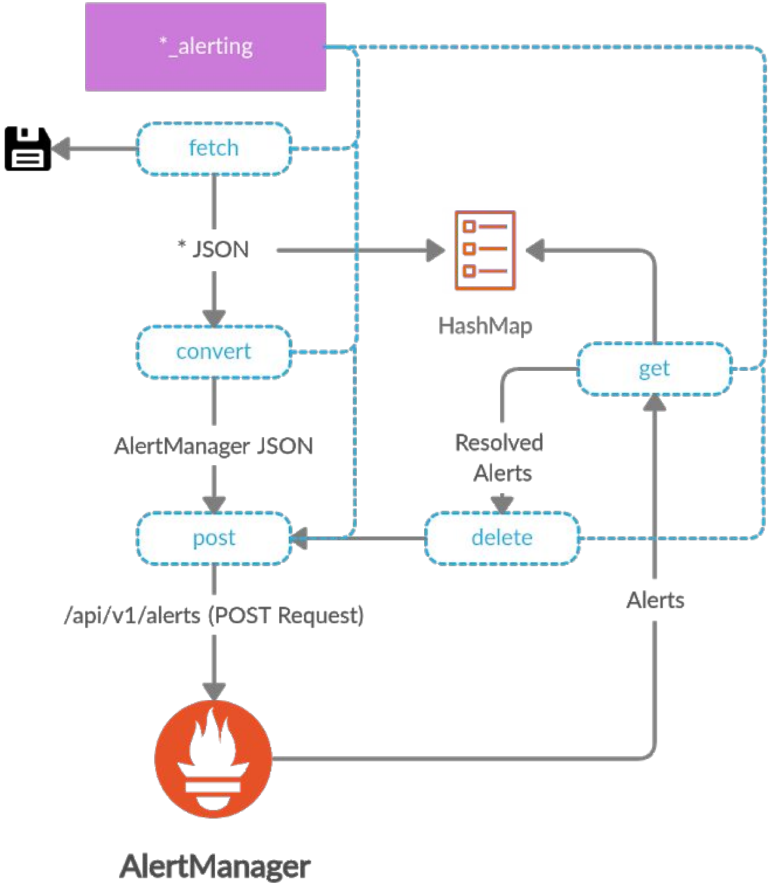

- convert
  - fetched alerts are input here
  - gets converted to JSON data which AlertManager API understands

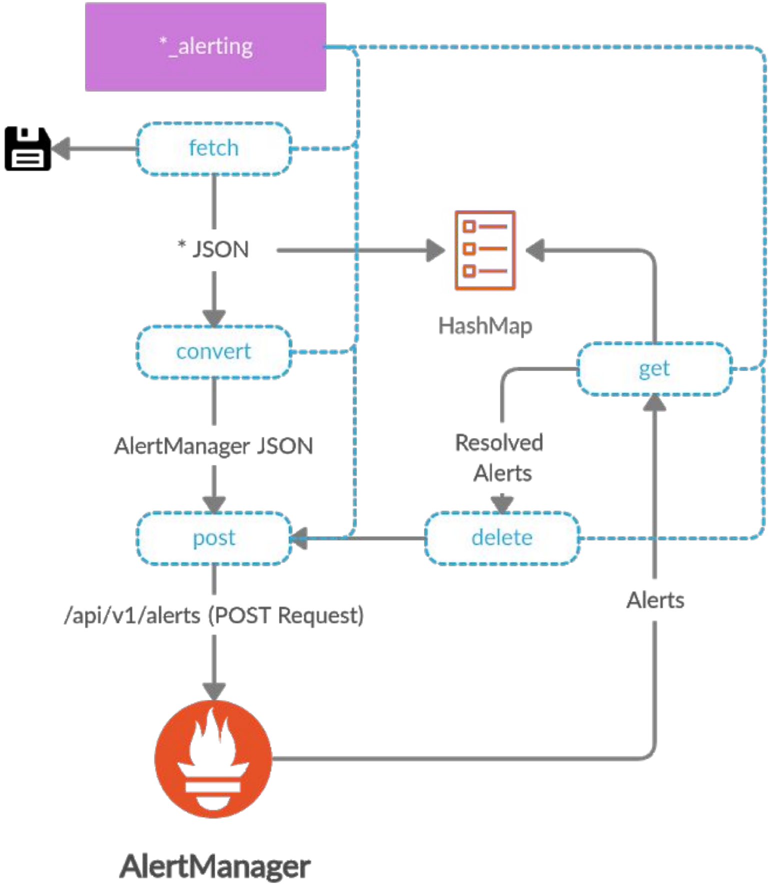

- post
  - converted JSON data which contains GGUS/SSB alerts is pushed to AlertManager.

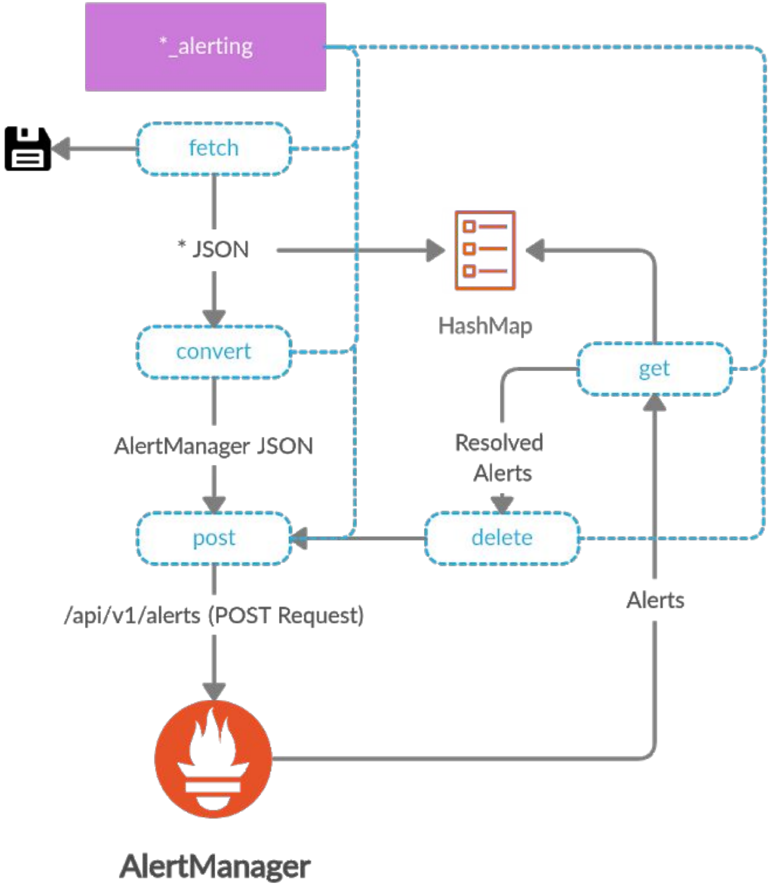

- get
  - Few GGUS/SSB alerts do not have Ending Time, hence open ending.
  - We fetch GGUS/SSB alerts from AlertManager
  - Check with HashMap (which updates), if an alert is resolved or not.
  - Bundle all resolved alerts

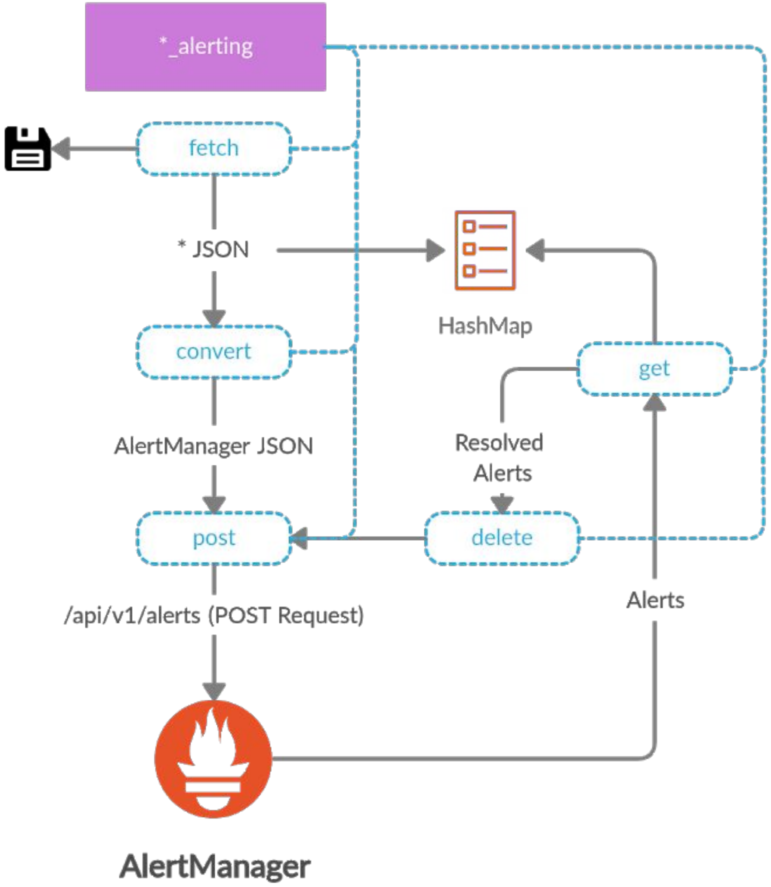

- delete
  - All resolved alerts will now have End Time == time.Now()
  - All open ending alerts in AlertManager get new EndTime,
  - thus get deleted

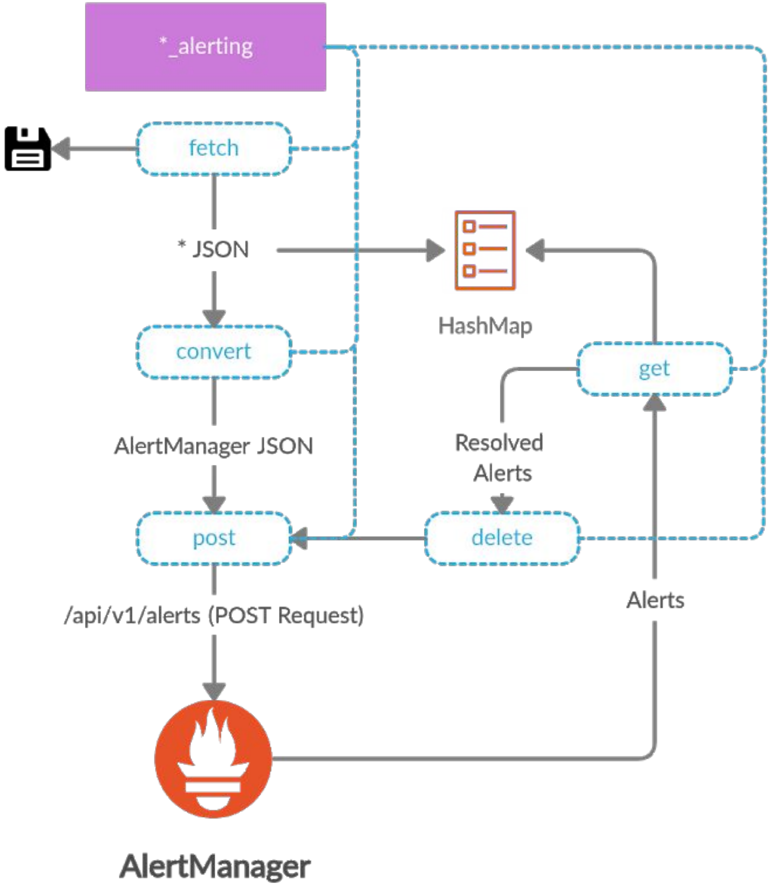

C. Alerting Service

Image beside shows an alerting service architecture

Components

- parser / monit
- *.alerting module

Alerting service -> A linux service running both of these logics at a regular interval in the background.

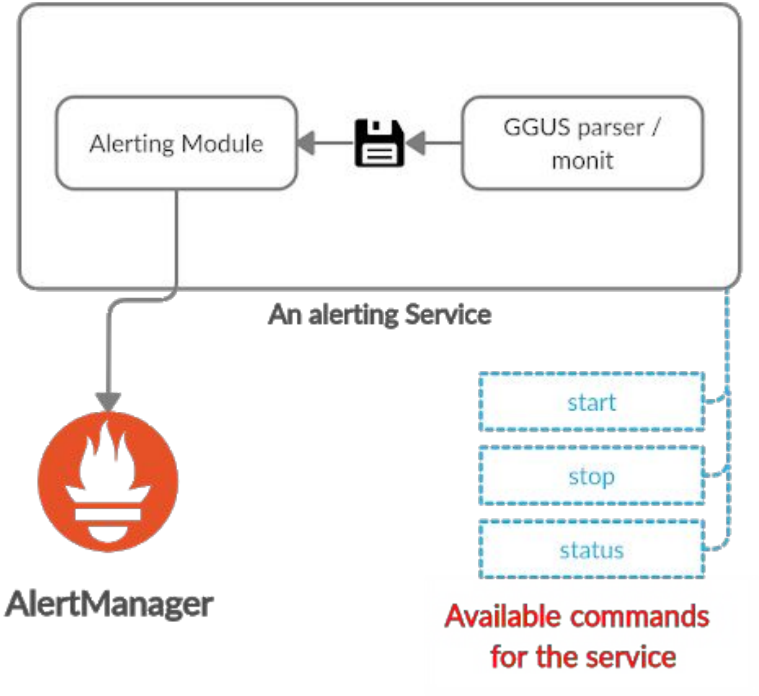

Configuration

- AlertManager URL
- Time Interval for the service
- HTTP Timeout
- Verbosity Level
  - GGUS
- GGUS Format
- VO
  - SSB
- Query
- Token

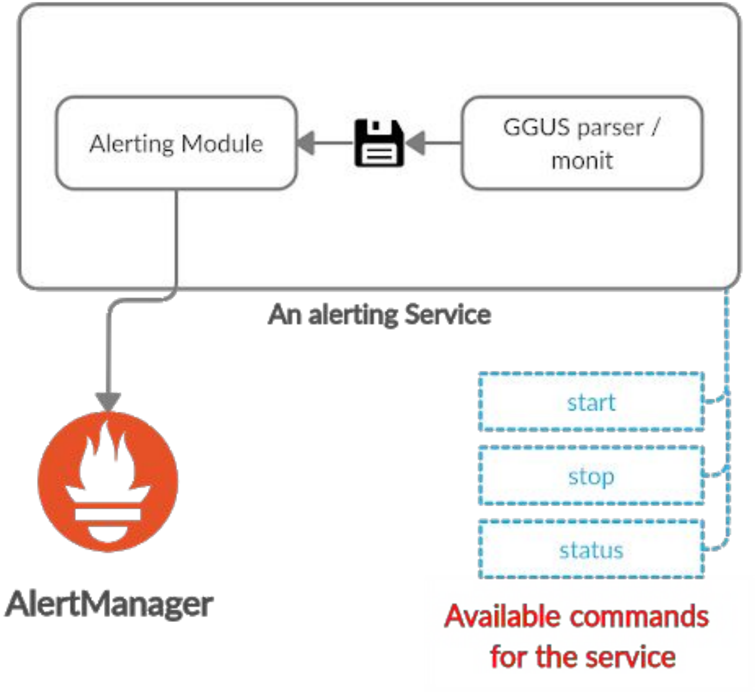

D. Slack & Karma

alerts-ggus channel for GGUS alerts on Slack

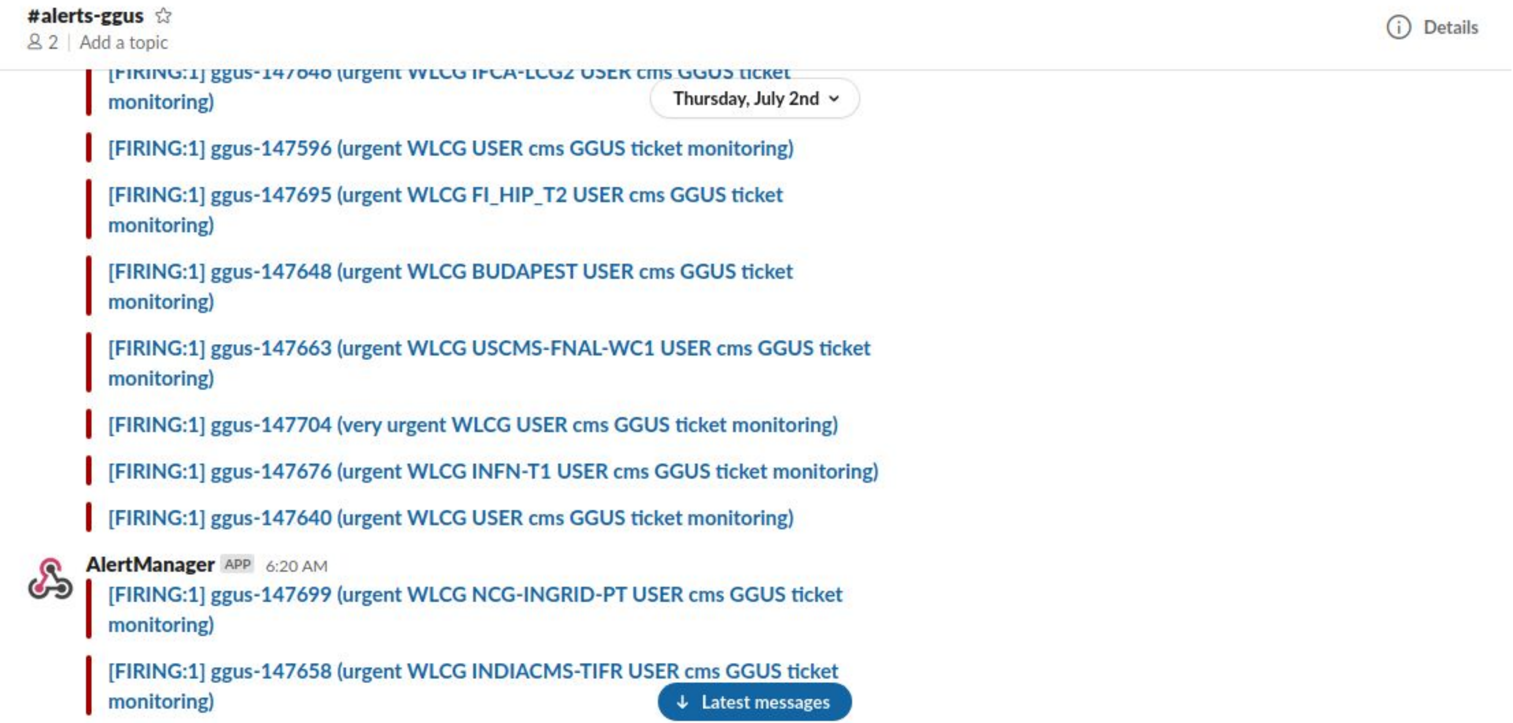

#alerts-ggus ☆
2 | Add a topic

[FIRING:1] ggus-147646 (urgent WLCG IFCA-LCG2 USER cms GGUS ticket monitoring)

[FIRING:1] ggus-147596 (urgent WLCG USER cms GGUS ticket monitoring)

[FIRING:1] ggus-147695 (urgent WLCG FI_HIP_T2 USER cms GGUS ticket monitoring)

[FIRING:1] ggus-147648 (urgent WLCG BUDAPEST USER cms GGUS ticket monitoring)

[FIRING:1] ggus-147663 (urgent WLCG USCMS-FNAL-WC1 USER cms GGUS ticket monitoring)

[FIRING:1] ggus-147704 (very urgent WLCG USER cms GGUS ticket monitoring)

[FIRING:1] ggus-147676 (urgent WLCG INFN-T1 USER cms GGUS ticket monitoring)

[FIRING:1] ggus-147640 (urgent WLCG USER cms GGUS ticket monitoring)

AlertManager APP 6:20 AM

[FIRING:1] ggus-147699 (urgent WLCG NCG-INGRID-PT USER cms GGUS ticket monitoring)

[FIRING:1] ggus-147658 (urgent WLCG INDIACMS-TIFR USER cms GGUS ticket monitoring)

alerts-ssb channel for SSB alerts on Slack

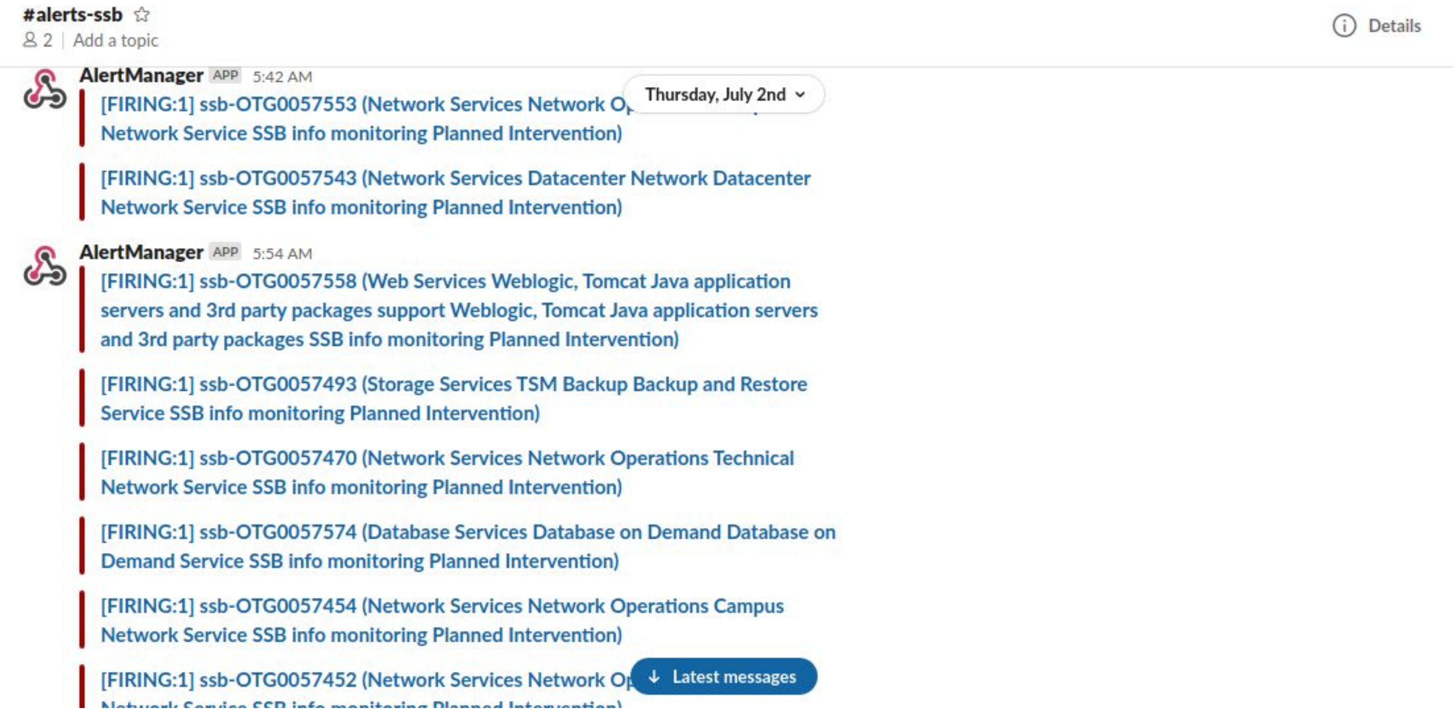

#alerts-ssb ☆

2 | Add a topic

AlertManager APP 5:42 AM

[FIRING:1] ssb-OTG0057553 (Network Services Network O
Network Service SSB info monitoring Planned Intervention)

[FIRING:1] ssb-OTG0057543 (Network Services Datacenter Network Datacenter Network Service SSB info monitoring Planned Intervention)

AlertManager APP 5:54 AM

[FIRING:1] ssb-OTG0057558 (Web Services Weblogic, Tomcat Java application servers and 3rd party packages support Weblogic, Tomcat Java application servers and 3rd party packages SSB info monitoring Planned Intervention)

[FIRING:1] ssb-OTG0057493 (Storage Services TSM Backup Backup and Restore Service SSB info monitoring Planned Intervention)

[FIRING:1] ssb-OTG0057470 (Network Services Network Operations Technical Network Service SSB info monitoring Planned Intervention)

[FIRING:1] ssb-OTG0057574 (Database Services Database on Demand Database on Demand Service SSB info monitoring Planned Intervention)

[FIRING:1] ssb-OTG0057454 (Network Services Network Operations Campus Network Service SSB info monitoring Planned Intervention)

[FIRING:1] ssb-OTG0057452 (Network Services Network O
Network Service SSB info monitoring Planned Intervention)

Karma Dashboard
https://cms-monitoring.cern.ch

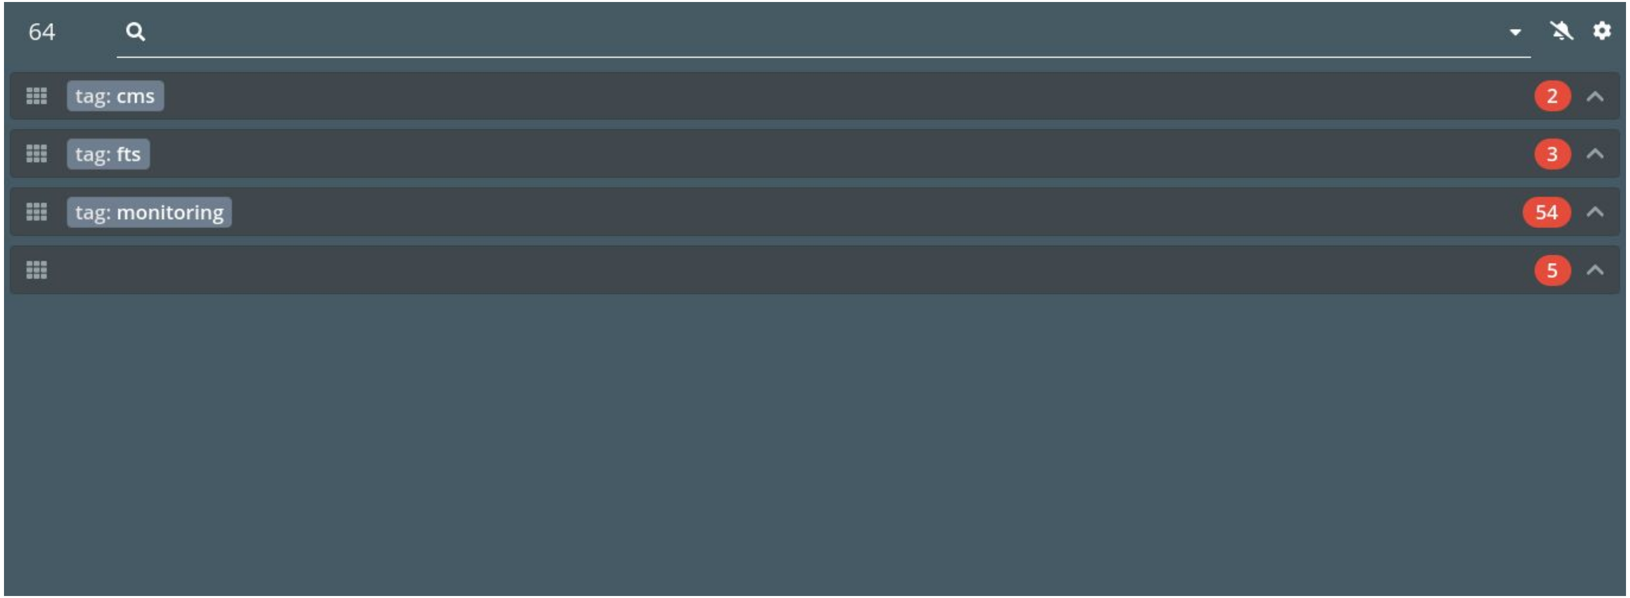

Karma Dashboard showing all alerts under “tag=monitoring” (SSB)

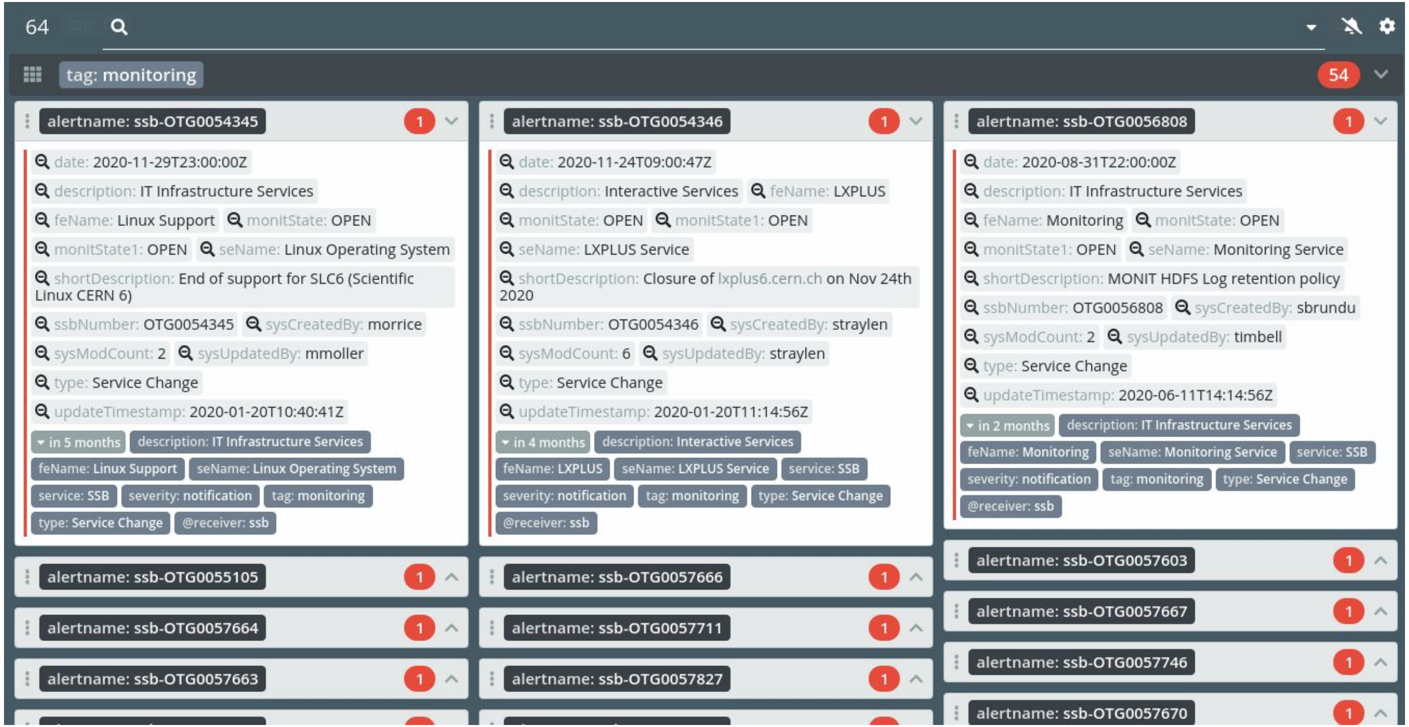

64

tag: monitoring 54

alertname: ssb-OTG0054345 1

date: 2020-11-29T23:00:00Z
description: IT Infrastructure Services
feName: Linux Support
monitState: OPEN
monitState1: OPEN
seName: Linux Operating System
shortDescription: End of support for SLC6 (Scientific Linux CERN 6)
ssbNumber: OTG0054345
sysCreatedBy: morrice
sysModCount: 2
sysUpdatedBy: mmoller
type: Service Change
updateTimestamp: 2020-01-20T10:40:41Z

in 5 months
description: IT Infrastructure Services
feName: Linux Support
seName: Linux Operating System
service: SSB
severity: notification
tag: monitoring
type: Service Change
@receiver: ssb

alertname: ssb-OTG0054346 1

date: 2020-11-24T09:00:47Z
description: Interactive Services
feName: LXPLUS
monitState: OPEN
monitState1: OPEN
seName: LXPLUS Service
shortDescription: Closure of lxplus6.cern.ch on Nov 24th 2020
ssbNumber: OTG0054346
sysCreatedBy: straylen
sysModCount: 6
sysUpdatedBy: straylen
type: Service Change
updateTimestamp: 2020-01-20T11:14:56Z

in 4 months
description: Interactive Services
feName: LXPLUS
seName: LXPLUS Service
service: SSB
severity: notification
tag: monitoring
type: Service Change
@receiver: ssb

alertname: ssb-OTG0056808 1

date: 2020-08-31T22:00:00Z
description: IT Infrastructure Services
feName: Monitoring
monitState: OPEN
monitState1: OPEN
seName: Monitoring Service
shortDescription: MONIT HDFS Log retention policy
ssbNumber: OTG0056808
sysCreatedBy: sbrundu
sysModCount: 2
sysUpdatedBy: timbell
type: Service Change
updateTimestamp: 2020-06-11T14:14:56Z

in 2 months
description: IT Infrastructure Services
feName: Monitoring
seName: Monitoring Service
service: SSB
severity: notification
tag: monitoring
type: Service Change
@receiver: ssb

alertname: ssb-OTG0055105 1

alertname: ssb-OTG0057664 1

alertname: ssb-OTG0057663 1

alertname: ssb-OTG0057666 1

alertname: ssb-OTG0057711 1

alertname: ssb-OTG0057827 1

alertname: ssb-OTG0057603 1

alertname: ssb-OTG0057667 1

alertname: ssb-OTG0057746 1

alertname: ssb-OTG0057670 1

E. Intelligence Module

Components

- Fetch Alerts
- Preprocessing
- Keyword Matching
- Add Annotations
- Machine Learning
- Push Alert
- Silence Alert

Tools

- AlertManager
- Grafana

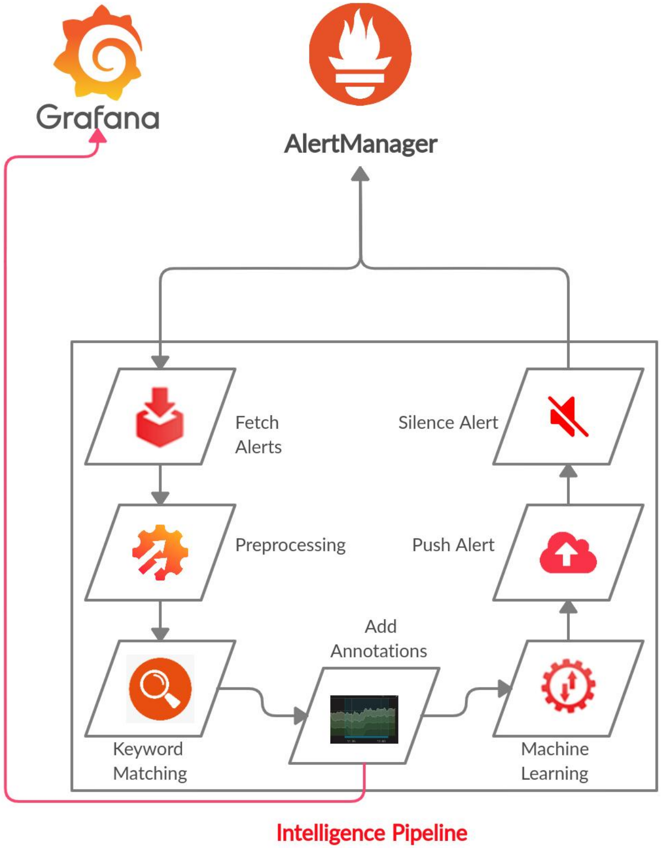

Fetch Alerts

- Fetches all alerts from AlertManager
- Bundles them and put them on a channel.
- Channel (Analogy) - baggage belt at Airports. You put data into it, data will be picked up when required by other party.

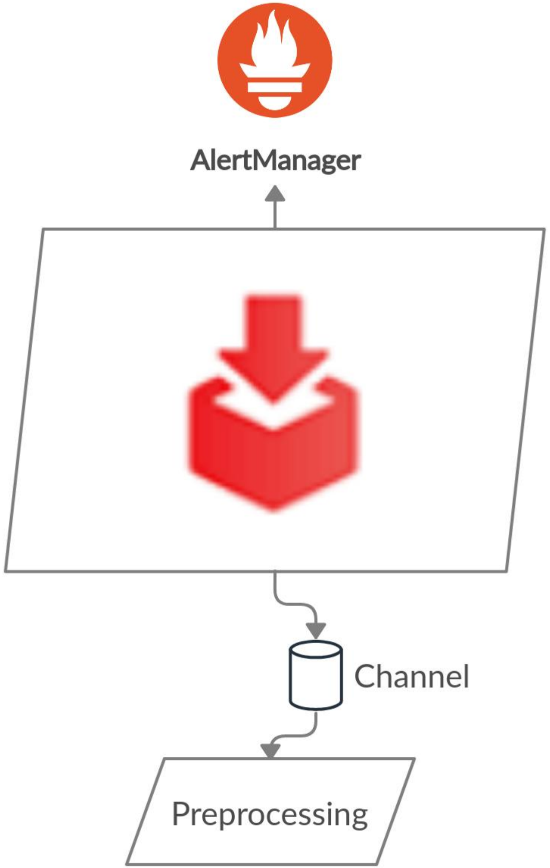

Preprocessing

- Filtering based on configuration.
- Only filtered alerts are forwarded.
- Here we also manage one map for keeping track of active silenced alerts to avoid redundant silences.
- If an alert is already silenced that means it has been processed by the intelligence module before.

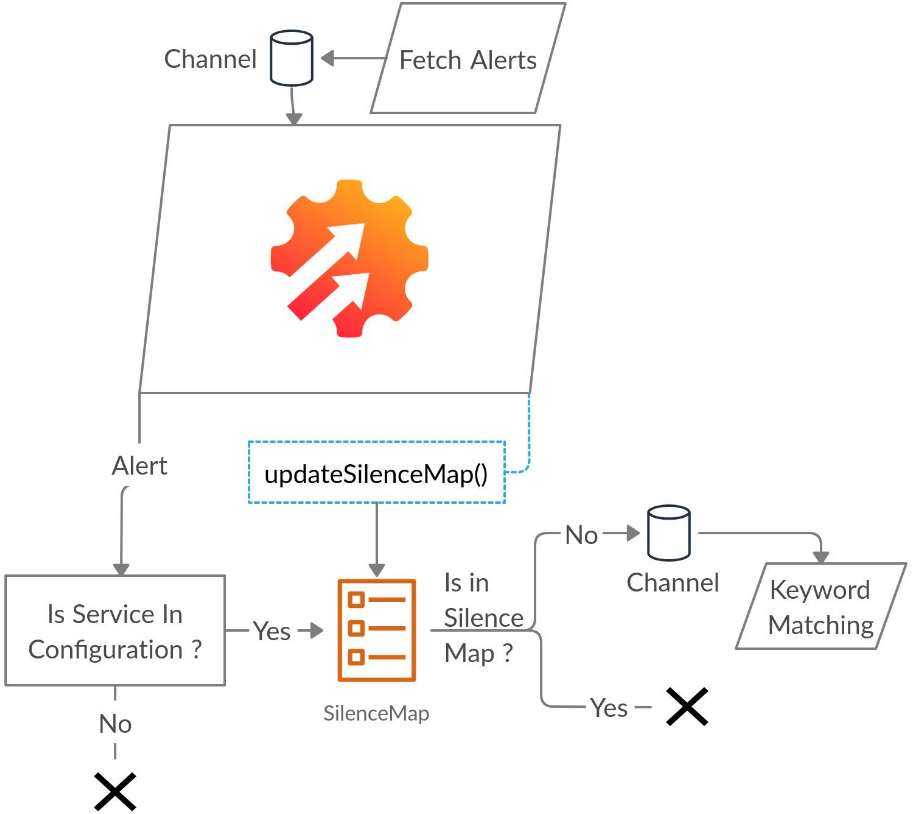

Keyword Matching

- Analysis of Alerts showed us repetitive use of a few important keywords.
- These keywords help in assigning severity levels.
- We search for these keywords in alerts, if found we assign severity level mapped to that keyword.

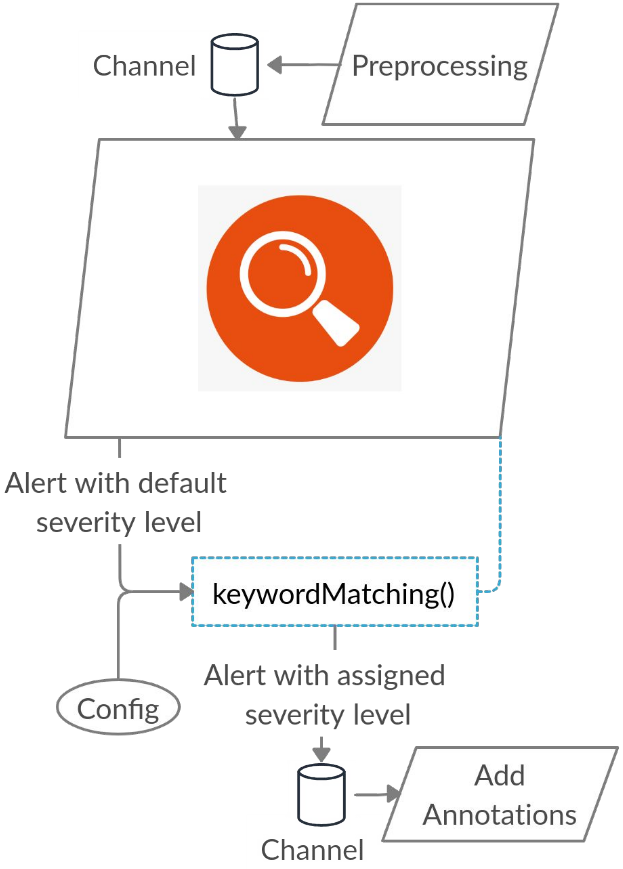

Channel
Preprocessing
Alert with default severity level
keywordMatching()
Config
Alert with assigned severity level
Channel
Add Annotations

Add Annotations

- Grafana has dashboards which shows running services' metrics in the form of graphs.
- Grafana has add Annotation feature.
- SSB alert mentioning intervention in network / DB affects these services.
- We push such interventions info in the form of annotations into Grafana dashboards.

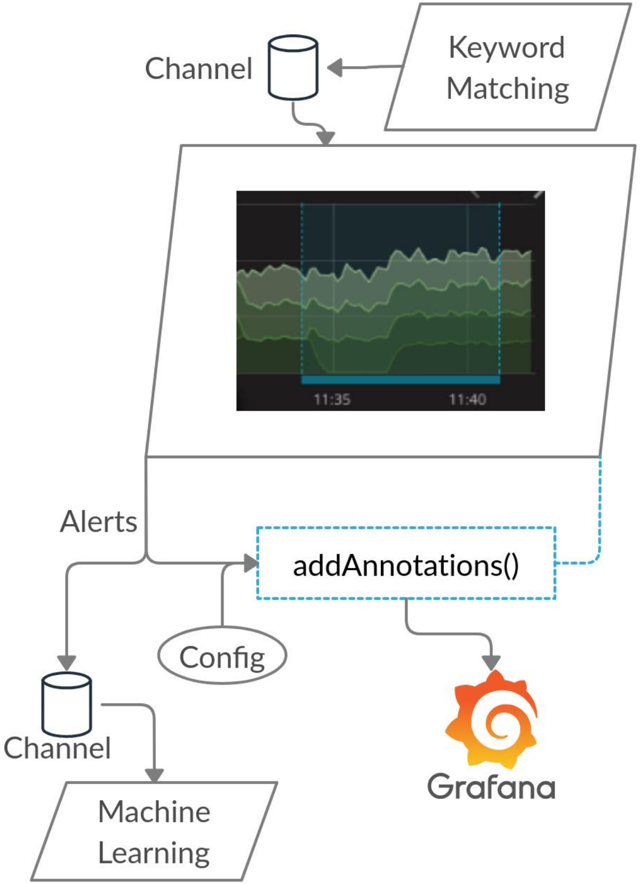

Machine Learning

FUTURE WORK

As of now forwards the same data that it gets

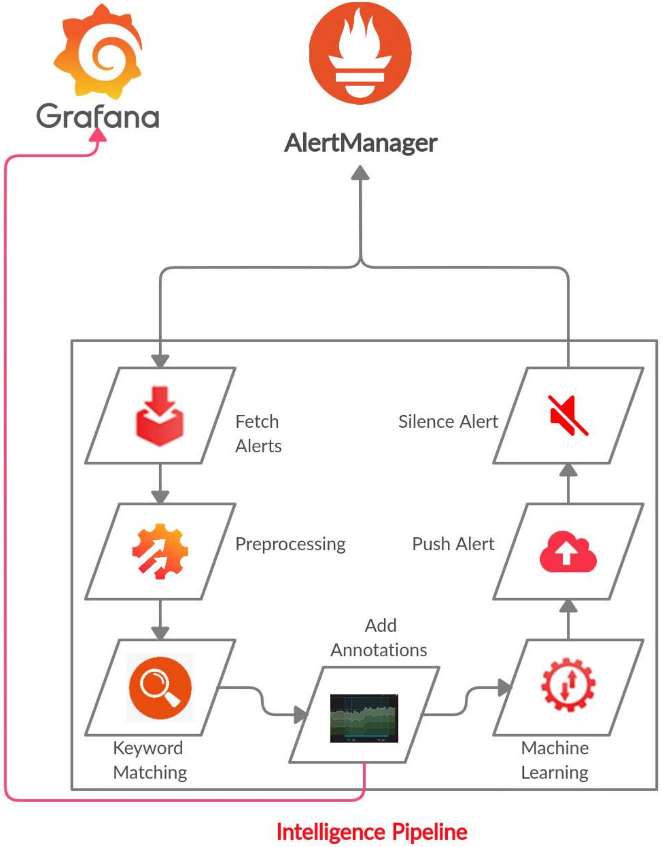

Push Alert

- Alerts with modified information are pushed to AlertManager
- Incoming alerts are then forwarded to Silence Alert.

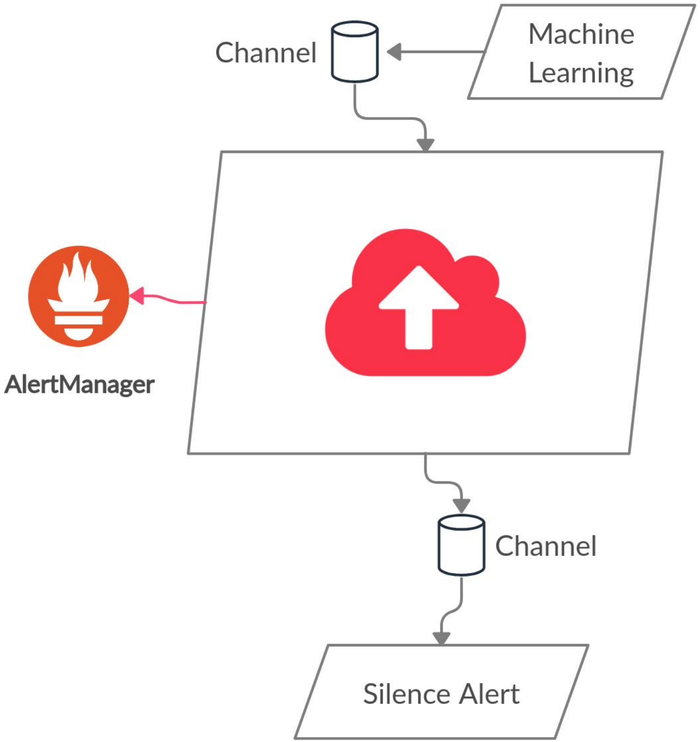

Silence Alert

- Alerts which get modified and pushed to AlertManager get copied.
- Older alert is redundant
- We silence the older one for the duration of its lifetime.

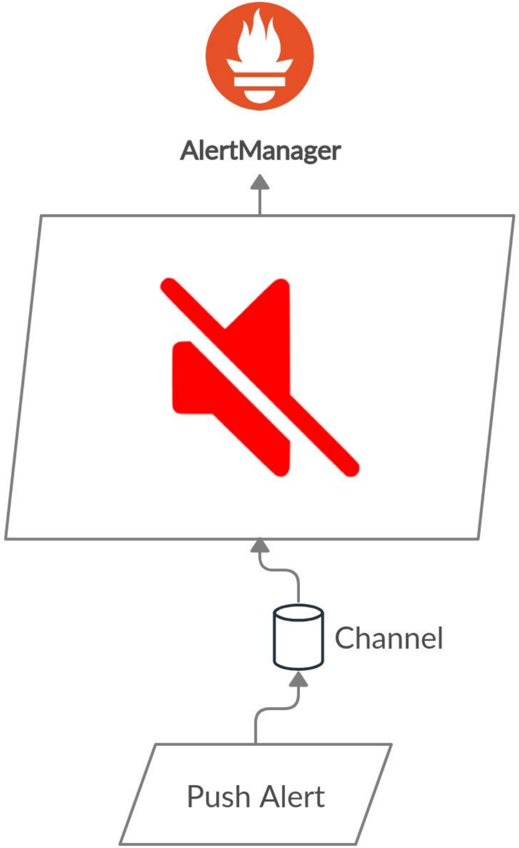

F. Alert CLI Tool

- Gives a nice and clean CLI interface for getting alerts, their details printed on the terminal itself either in tabular form or JSON format.
- Convenient option for operators who prefer command line
- Comes with several options such as :-
  - service, severity, tag - Filters
  - sort - Sorting
  - details - For detailed information of an alert
  - json - information in JSON format

$alert -service=SSB -sort=duration

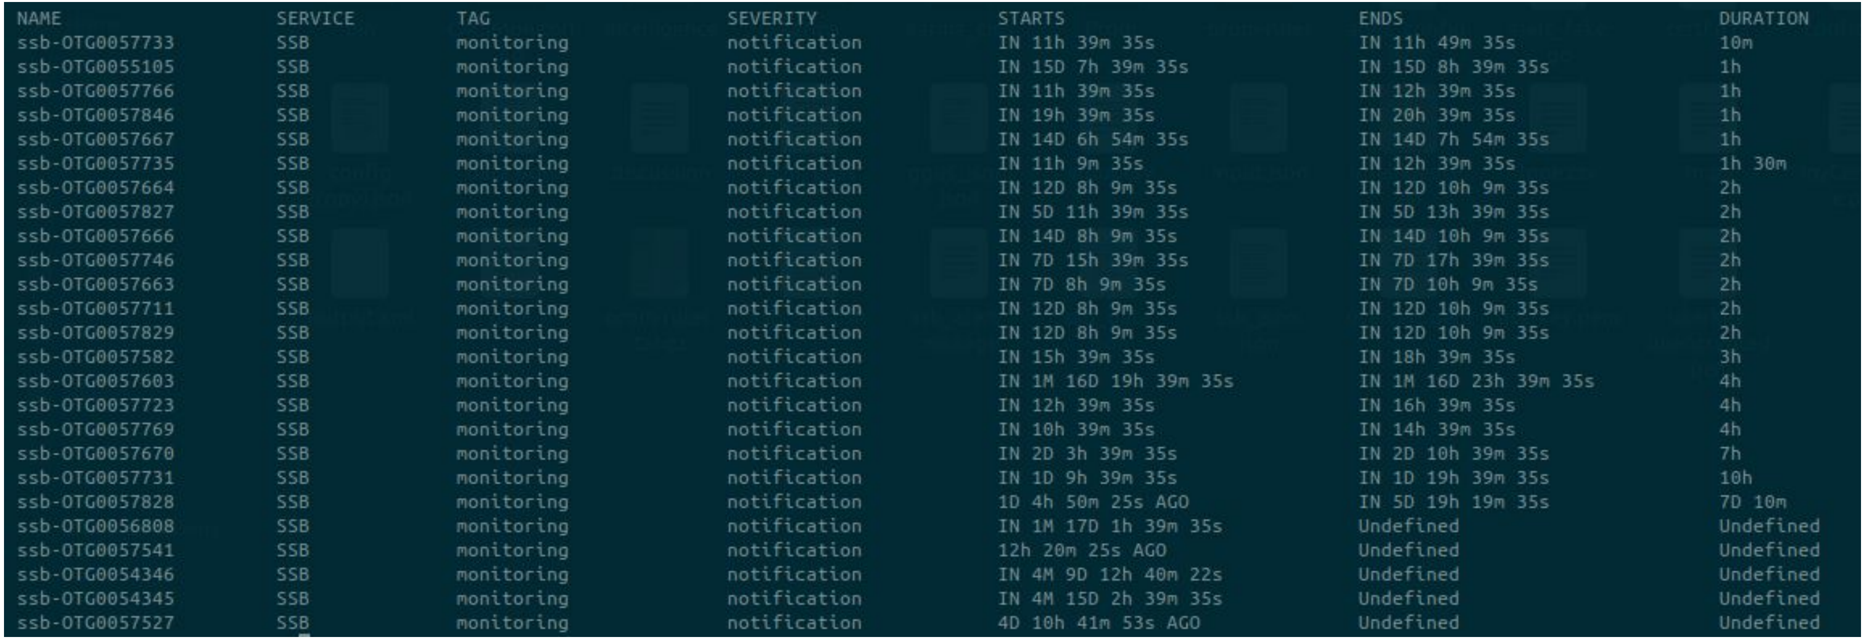

```
NAME                   SERVICE         TAG                     SEVERITY                STARTS                          ENDS                            DURATION
ssb-OTG0057733         SSB             monitoring              notification            IN 11h 39m 35s                  IN 11h 49m 35s                  10m
ssb-OTG0055105         SSB             monitoring              notification            IN 15D 7h 39m 35s               IN 15D 8h 39m 35s               1h
ssb-OTG0057766         SSB             monitoring              notification            IN 11h 39m 35s                  IN 12h 39m 35s                  1h
ssb-OTG0057846         SSB             monitoring              notification            IN 19h 39m 35s                  IN 20h 39m 35s                  1h
ssb-OTG0057667         SSB             monitoring              notification            IN 14D 6h 54m 35s               IN 14D 7h 54m 35s               1h
ssb-OTG0057735         SSB             monitoring              notification            IN 11h 9m 35s                   IN 12h 39m 35s                  1h 30m
ssb-OTG0057664         SSB             monitoring              notification            IN 12D 8h 9m 35s                IN 12D 10h 9m 35s               2h
ssb-OTG0057827         SSB             monitoring              notification            IN 5D 11h 39m 35s               IN 5D 13h 39m 35s               2h
ssb-OTG0057666         SSB             monitoring              notification            IN 14D 8h 9m 35s                IN 14D 10h 9m 35s               2h
ssb-OTG0057746         SSB             monitoring              notification            IN 7D 15h 39m 35s               IN 7D 17h 39m 35s               2h
ssb-OTG0057663         SSB             monitoring              notification            IN 7D 8h 9m 35s                 IN 7D 10h 9m 35s                2h
ssb-OTG0057711         SSB             monitoring              notification            IN 12D 8h 9m 35s                IN 12D 10h 9m 35s               2h
ssb-OTG0057829         SSB             monitoring              notification            IN 12D 8h 9m 35s                IN 12D 10h 9m 35s               2h
ssb-OTG0057582         SSB             monitoring              notification            IN 15h 39m 35s                  IN 18h 39m 35s                  3h
ssb-OTG0057603         SSB             monitoring              notification            IN 1M 16D 19h 39m 35s           IN 1M 16D 23h 39m 35s           4h
ssb-OTG0057723         SSB             monitoring              notification            IN 12h 39m 35s                  IN 16h 39m 35s                  4h
ssb-OTG0057769         SSB             monitoring              notification            IN 10h 39m 35s                  IN 14h 39m 35s                  4h
ssb-OTG0057670         SSB             monitoring              notification            IN 2D 3h 39m 35s                IN 2D 10h 39m 35s               7h
ssb-OTG0057731         SSB             monitoring              notification            IN 1D 9h 39m 35s                IN 1D 19h 39m 35s               10h
ssb-OTG0057828         SSB             monitoring              notification            1D 4h 50m 25s AGO               IN 5D 19h 19m 35s               7D 10m
ssb-OTG0056808         SSB             monitoring              notification            IN 1M 17D 1h 39m 35s            Undefined                       Undefined
ssb-OTG0057541         SSB             monitoring              notification            12h 20m 25s AGO                 Undefined                       Undefined
ssb-OTG0054346         SSB             monitoring              notification            IN 4M 9D 12h 40m 22s            Undefined                       Undefined
ssb-OTG0054345         SSB             monitoring              notification            IN 4M 15D 2h 39m 35s            Undefined                       Undefined
ssb-OTG0057527         SSB             monitoring              notification            4D 10h 41m 53s AGO              Undefined                       Undefined
```

$alert -severity=high

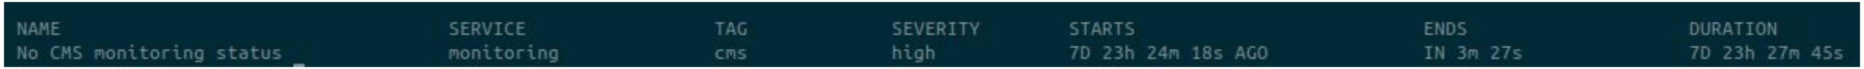

```
NAME                                         SERVICE                   TAG                SEVERITY          STARTS                              ENDS                       DURATION
No CMS monitoring status                     monitoring                cms                high              7D 23h 24m 18s AGO                  IN 3m 27s                  7D 23h 27m 45s
```

$alert -name=ssb-OTG0054345 -details

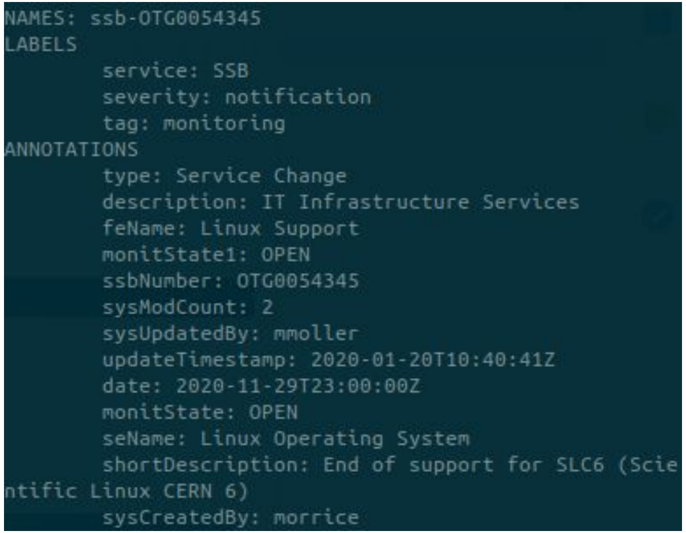

```
NAMES: ssb-OTG0054345
LABELS
        service: SSB
        severity: notification
        tag: monitoring
ANNOTATIONS
        type: Service Change
        description: IT Infrastructure Services
        feName: Linux Support
        monitState1: OPEN
        ssbNumber: OTG0054345
        sysModCount: 2
        sysUpdatedBy: mmoller
        updateTimestamp: 2020-01-20T10:40:41Z
        date: 2020-11-29T23:00:00Z
        monitState: OPEN
        seName: Linux Operating System
        shortDescription: End of support for SLC6 (Scie
ntific Linux CERN 6)
        sysCreatedBy: morrice
```

$alert -name=ssb-OTG0054345 -details -json

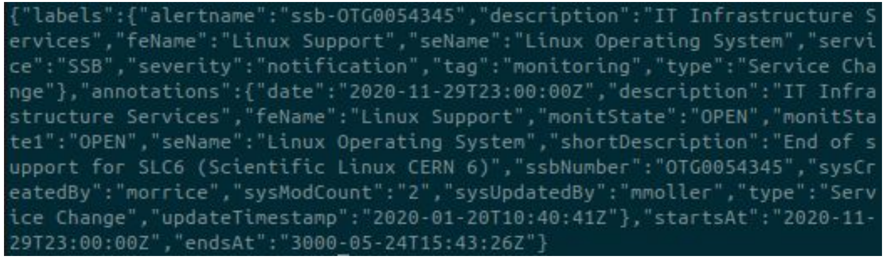

```
{"labels":{"alertname":"ssb-OTG0054345","description":"IT Infrastructure S
ervices","feName":"Linux Support","seName":"Linux Operating System","servi
ce":"SSB","severity":"notification","tag":"monitoring","type":"Service Cha
nge"},"annotations":{"date":"2020-11-29T23:00:00Z","description":"IT Infra
structure Services","feName":"Linux Support","monitState":"OPEN","monitSta
te1":"OPEN","seName":"Linux Operating System","shortDescription":"End of s
upport for SLC6 (Scientific Linux CERN 6)","ssbNumber":"OTG0054345","sysCr
eatedBy":"morrice","sysModCount":"2","sysUpdatedBy":"mmoller","type":"Serv
ice Change","updateTimestamp":"2020-01-20T10:40:41Z"},"startsAt":"2020-11-
29T23:00:00Z","endsAt":"3000-05-24T15:43:26Z"}
```