

Unlock High-Frequency Deployments without Blowing Up Prometheus

Ganesh Vernekar

SRECon Americas 2026

About Me

Staff Software Engineer at [Reddit](#)

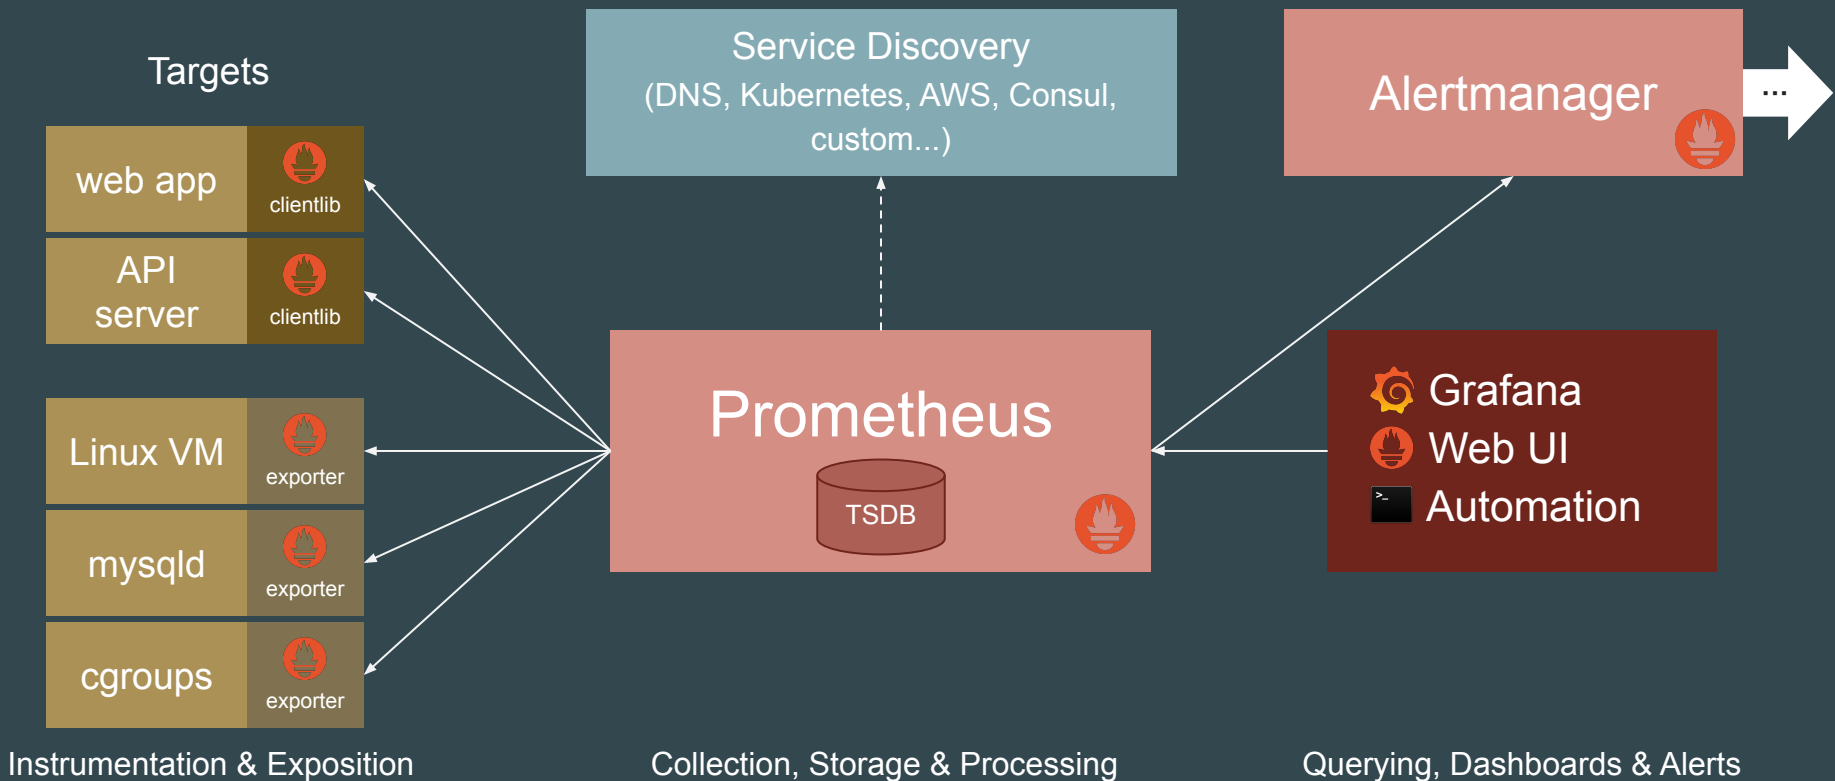
 Prometheus team member & TSDB maintainer

Worked on Mimir and Cortex at Grafana Labs

GitHub: [codesome](#)

The Basics

Prometheus



Time Series

`<identifier>` → `[(t0, v0), (t1, v1), ...]`

string int64 float64

Time Series

What identifies a time series?

```
http_requests_total{job="nginx",instance="1.2.3.4:80",status="200"} 28
```



metric name



labels

- Flexible
- No hierarchy
- Explicit dimensions

Terminologies

Active Series

Time series that is still getting scraped and getting new samples.

Stale Series

Stopped getting any samples; not being scraped anymore.

Stale series causes and problems

How do series go stale?

- Targets go away
- Instrumentation changes
- **Rollouts**

Staleness on k8s Deployment rollout

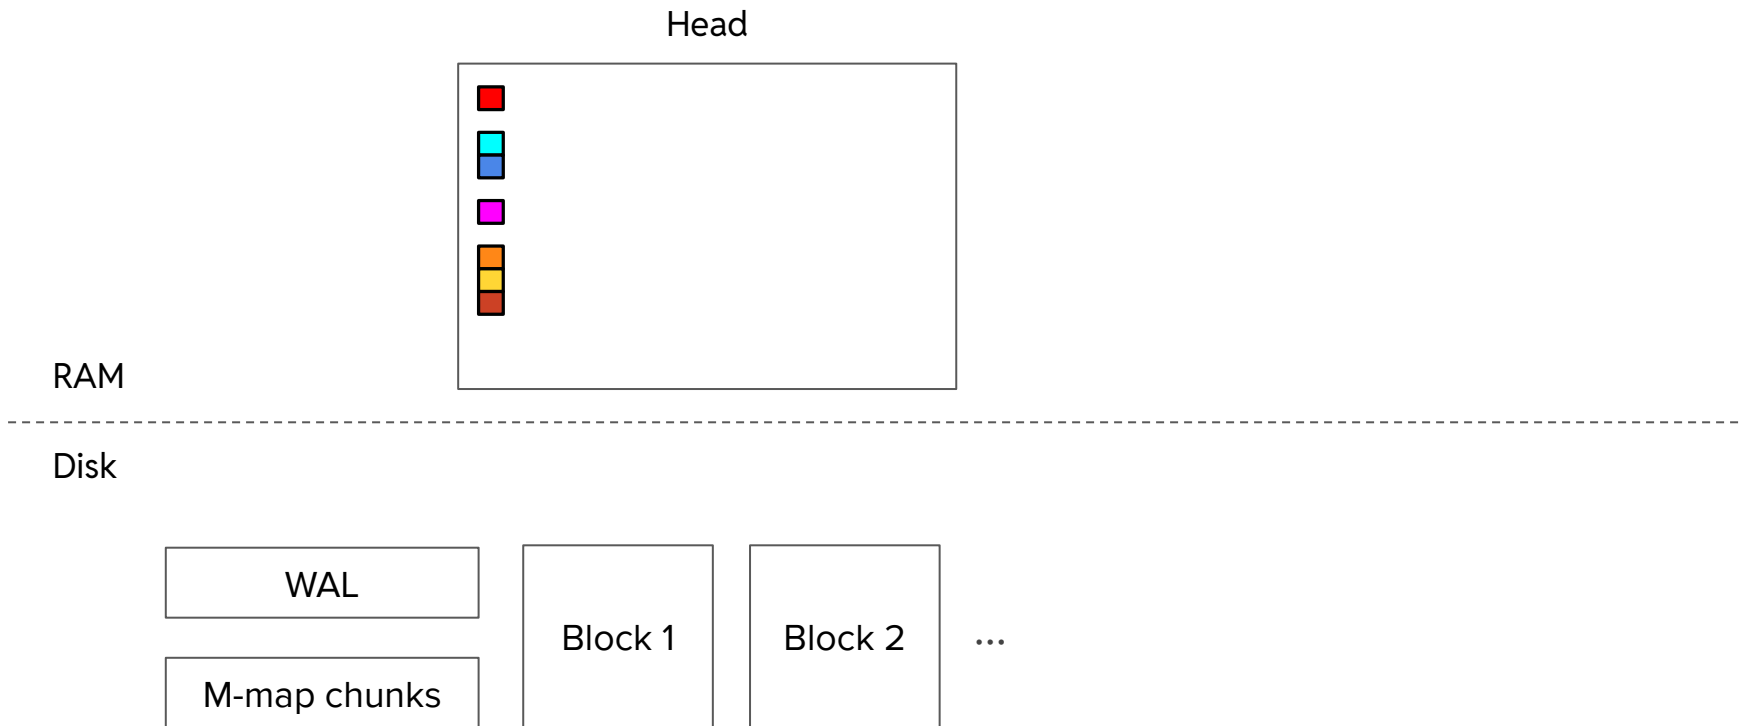
```
my_metric{pod="my-pod-cf97d665c-wbzjs", ...}  
↓  
my_metric{pod="my-pod-cf97d665c-vd5vf", ...}
```

```
my_metric{pod="my-pod-cf97d665c-wbzjs", ...}  
↓  
my_metric{pod="my-pod-bf5897d86-cwgnm", ...}
```


Staleness on k8s StatefulSet rollout

```
my_metric{pod="my-pod-0", instance="10.0.12.1", ...}  
↓  
my_metric{pod="my-pod-0", instance="10.0.10.5", ...}
```


Stale Series Visualization



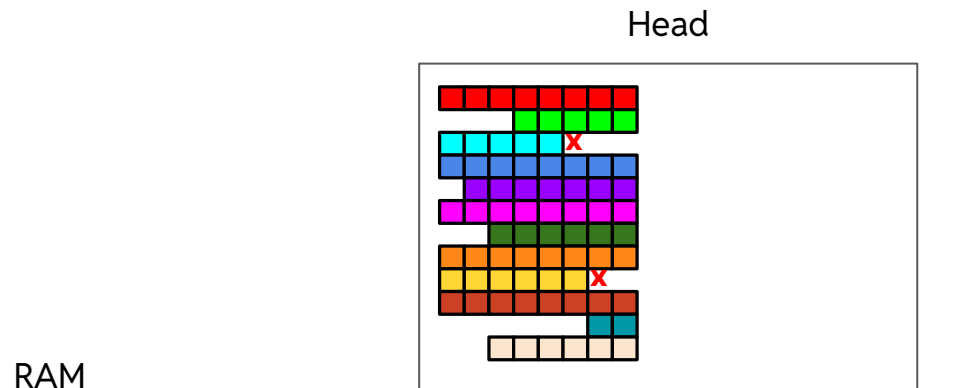
Stale Series Visualization




Disk




Stale Series Visualization




Disk



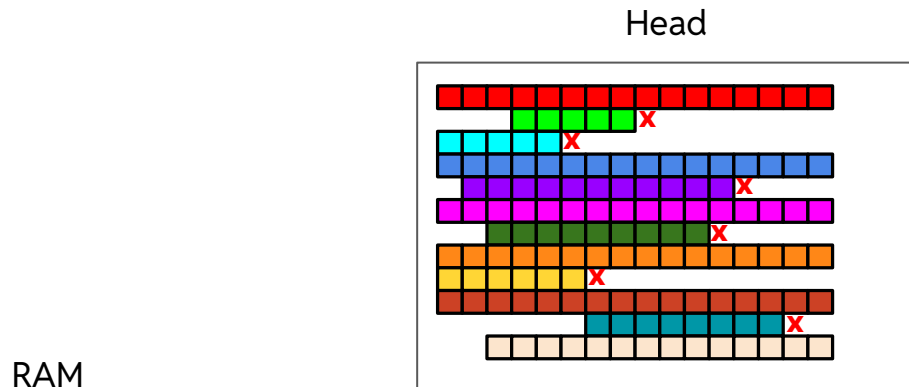
Stale Series Visualization




Disk



Stale Series Visualization




Disk




Real world failure example


Memory Usage (Normalized to Requested)



Stale Series Percentage




Series (Normalized)

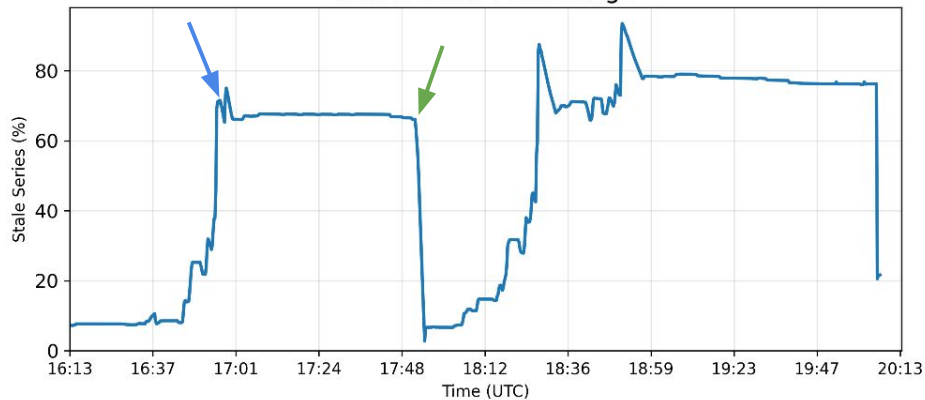


Real world failure example


Memory Usage (Normalized to Requested)



Stale Series Percentage




Series (Normalized)




Real world failure example


Memory Usage (Normalized to Requested)



Stale Series Percentage




Series (Normalized)




Real world failure example


Memory Usage (Normalized to Requested)



Stale Series Percentage




Series (Normalized)




Real world failure example

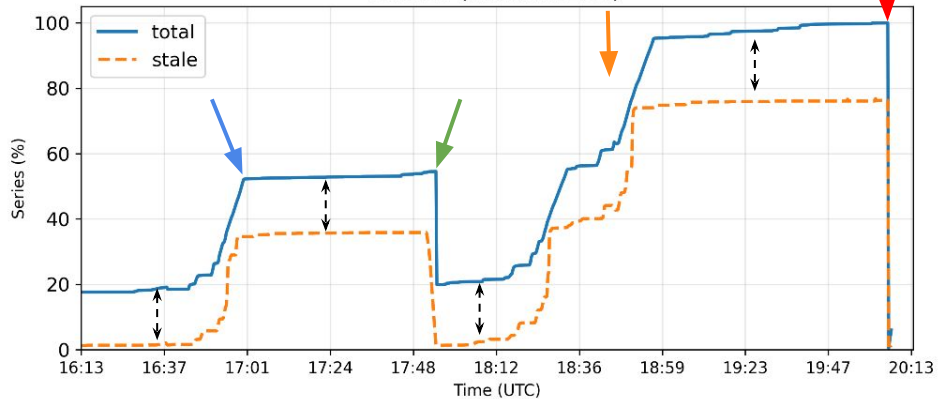
Memory Usage (Normalized to Requested)



Stale Series Percentage




Series (Normalized)




The Solution

It's what you think

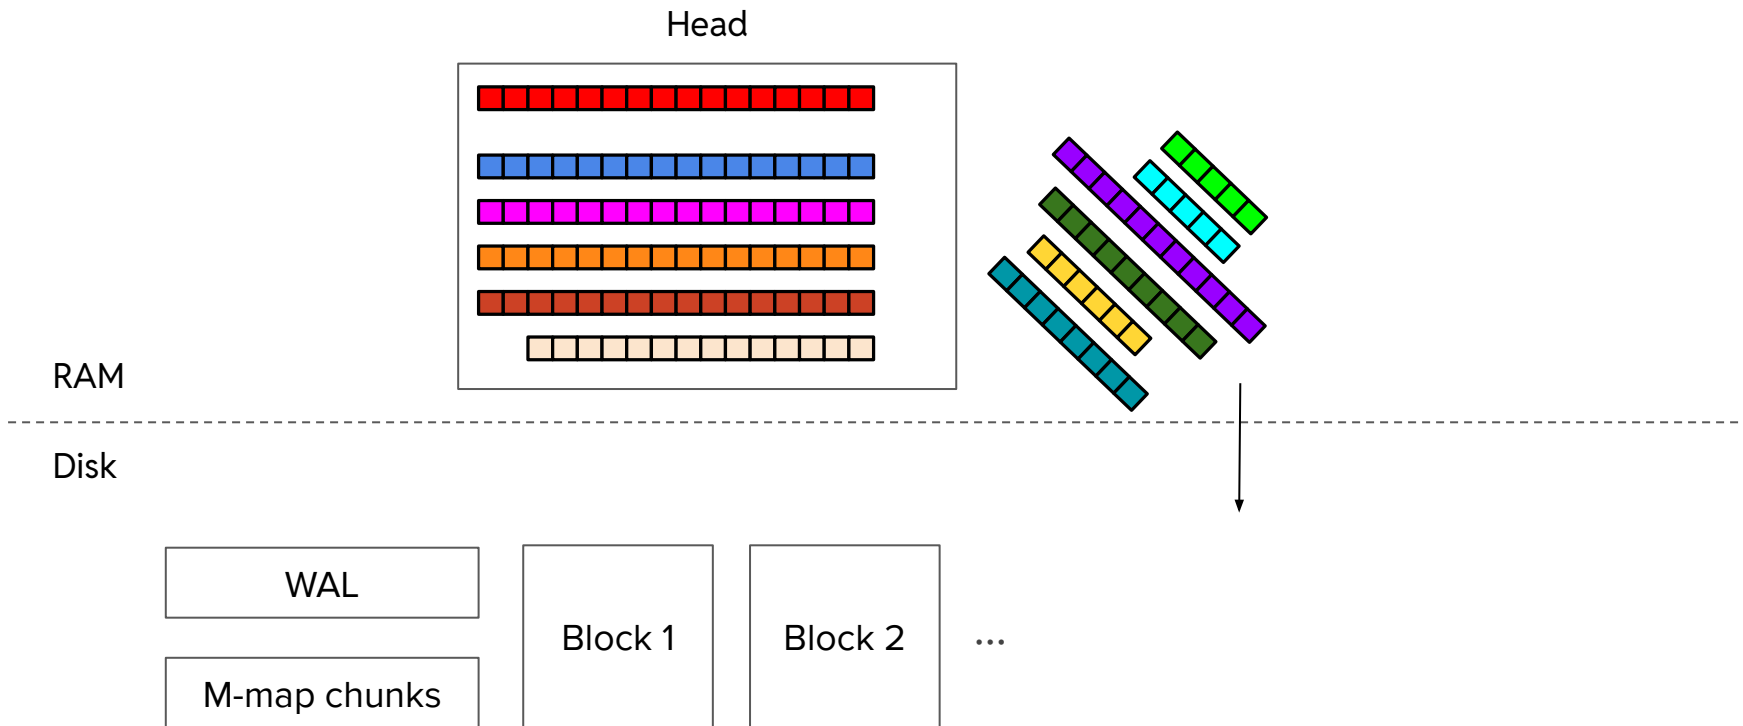


RAM

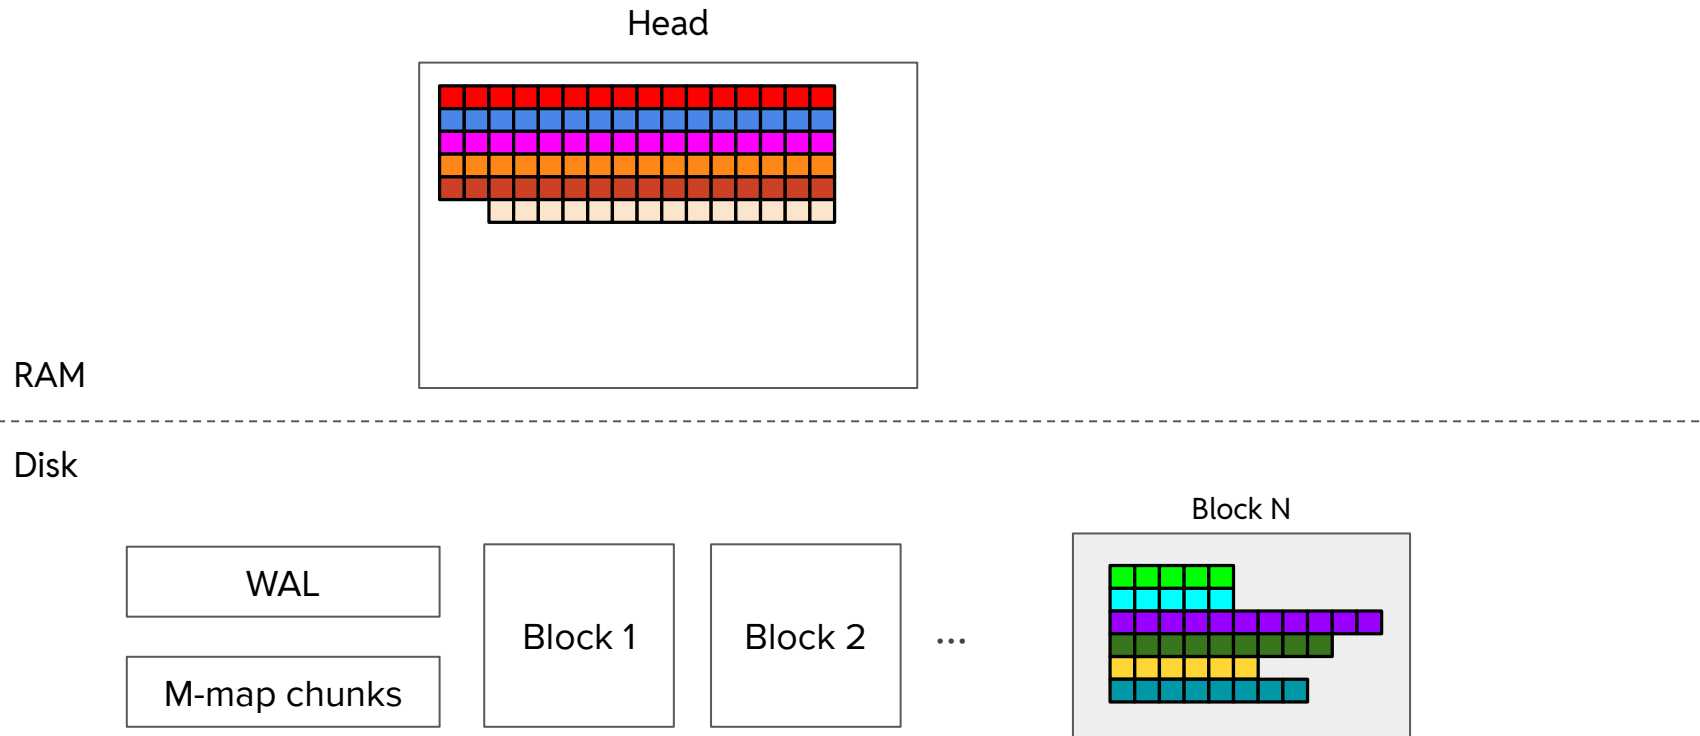
Disk



Stale Series Compaction



Stale Series Compaction



How do configure this?

Trigger stale series compaction when it hits this threshold

Config file

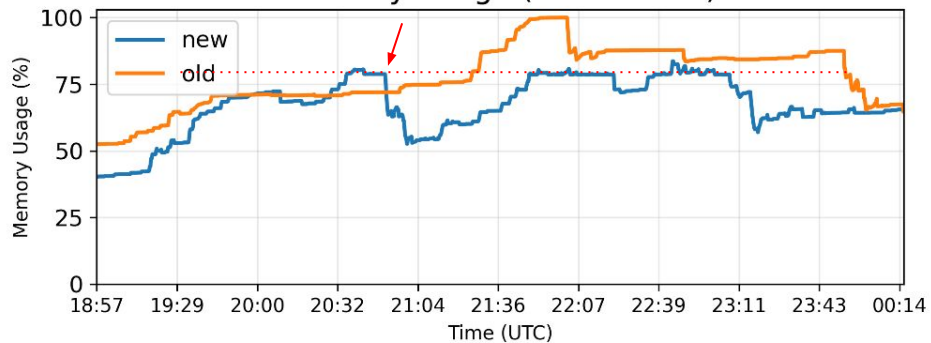
```
storage:  
  tsdb:  
    stale_series_compaction_threshold: <0.0 to 1.0>
```

From Prometheus v3.10.0 onwards

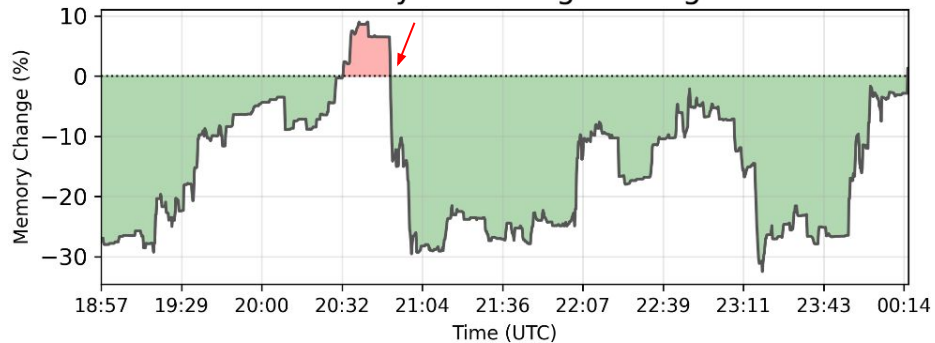
Experimental feature with a known bug [#18379](#)

Real World Example

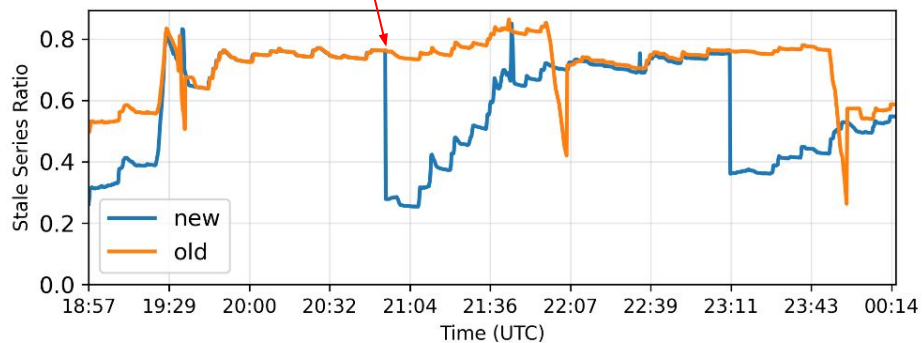
Memory Usage (Normalized)



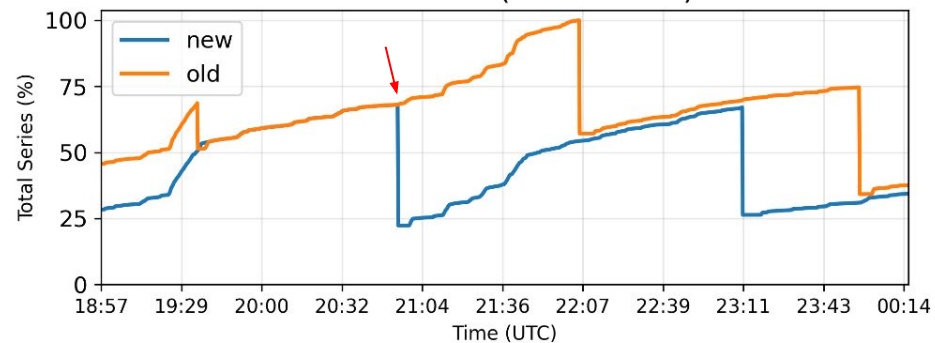
Memory Percentage Change



Stale Series Ratio




Total Series (Normalized)




It's not free

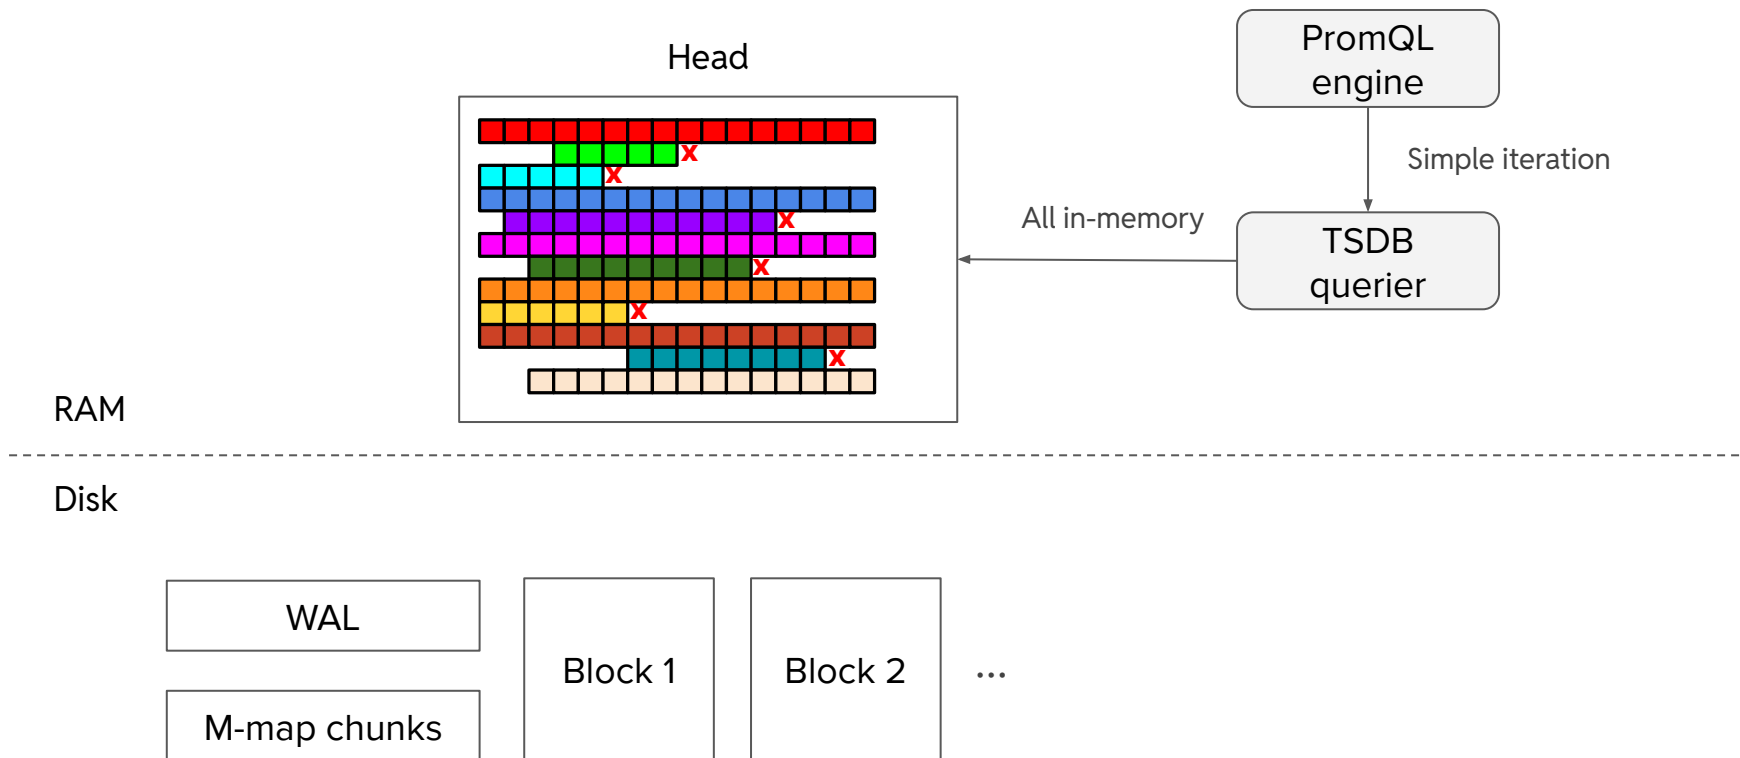
CPU Usage (Normalized)



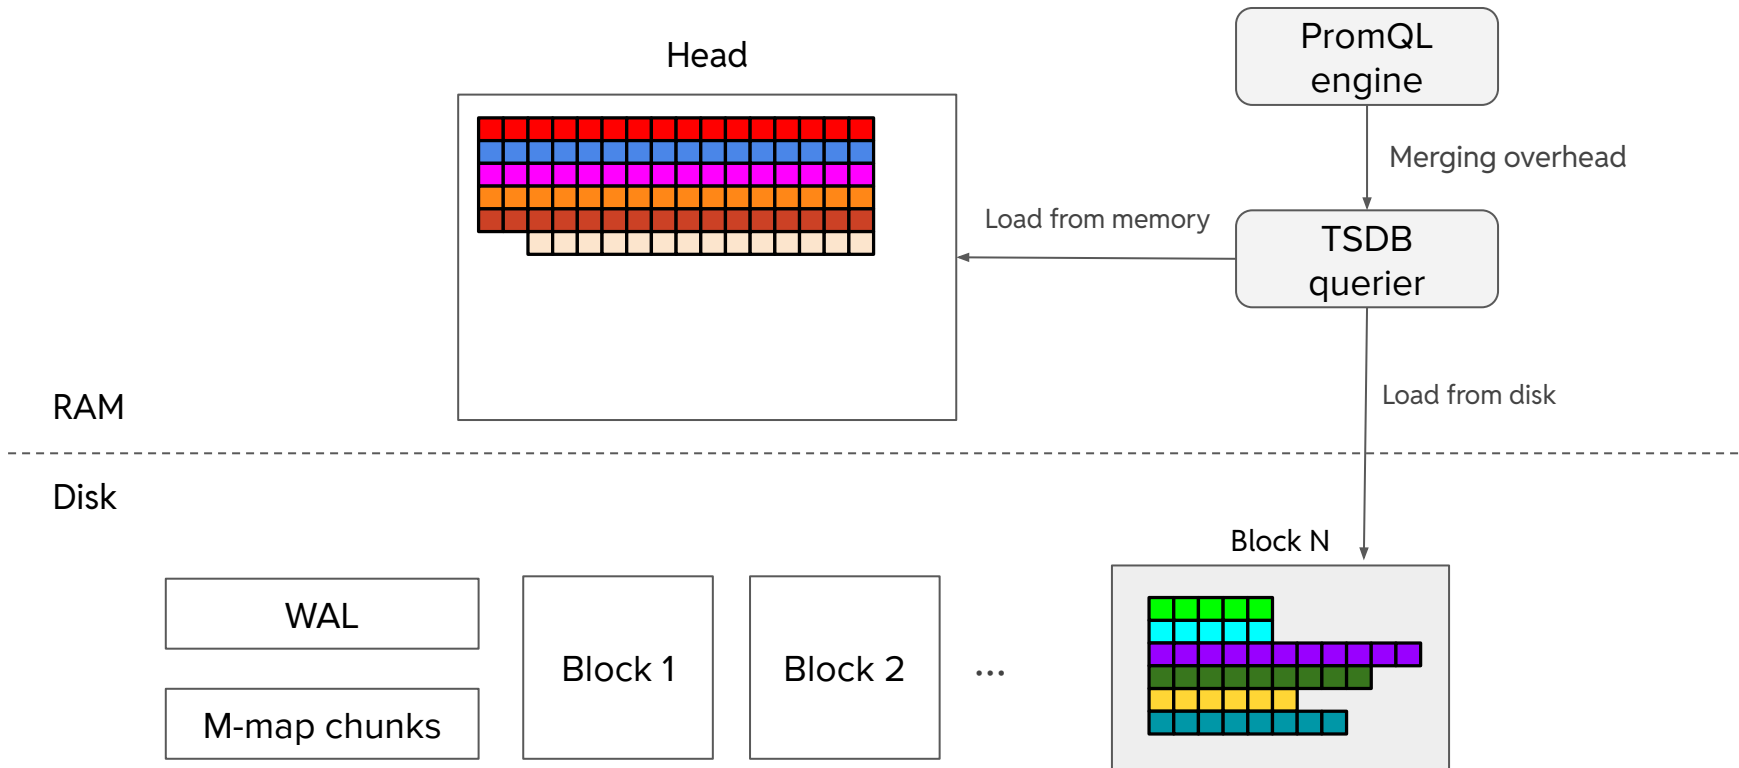
CPU Percentage Change



Recent data query - Before stale series compaction



Recent data query - After stale series compaction



How to choose stale series compaction trigger?

Disclaimer: Below guidelines are based on my experiments. You may need some trial-and-error.

Consider this feature only if high churn is blowing up a *large* Prometheus

Depends on


- Recording rule load
- Other query loads
- Size of Prometheus
(to balance the query overhead)

```
Stale series ratio PromQL (v3.6.0+):  
    prometheus_tsdb_head_stale_series{  
        / prometheus_tsdb_head_series{}
```

Peak stale series ratio	Action
< 0.3	Don't bother
0.3 - 0.5	Keep a close eye for potential issues
> 0.5	Choose a number that is not touched often but also not touched rarely and protects your Prometheus (needs trial-and-error)


From The Real World Example

Stale Series Ratio

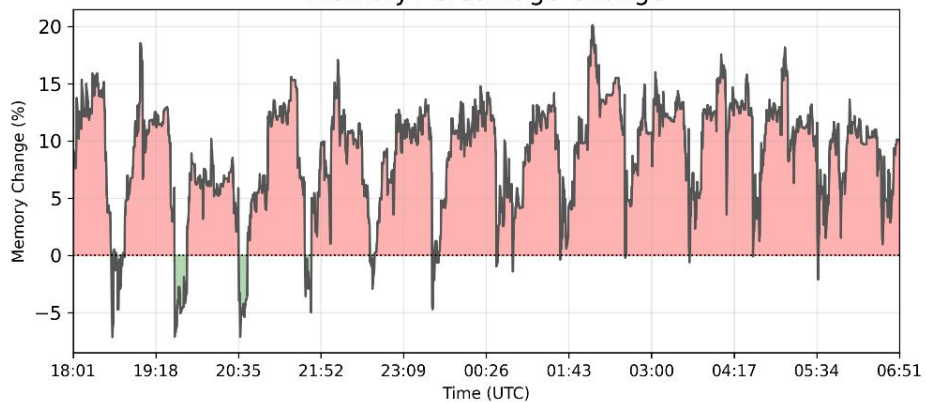


Consequences of wrong config


Stale Series Ratio



Memory Percentage Change



CPU Percentage Change



Try it out

And help us mature this feature.

Give feedback in `#prometheus` or `#prometheus-dev` channels on CNCF slack.

ETA for the bugfix: Prometheus v3.12.0 - v3.13.0

Thank You