




# cache me if you can

how Grafana Labs scaled up their memcached **42x**  
and **improved reliability**

**danny kopping / sr. software engineer**

# bottom line upfront



- ✓ **65% reduction in object storage reqs**
- ✓ **vastly improved reliability**
- ✓ **~2% overall TCO reduction**
- 🤔 **no change in performance**



but first...

# loki internals



# how does Loki work?

2023-06-16T06:09:05.123456789Z

{app="nginx", env="dev"}

192.168.1.1 [06/16/2023:06:09:05 +0000] "GET / HTTP/1.1" 200

**timestamp**

with nanosecond precision

**labels/selectors**


key-value pairs

**content**

log line

**indexed**

index



**unindexed**

chunks



logs

ingester

### chunks

2023-06-16 11:02.123456789

2023-06-16 11:12.123456789

{app="nginx", pod="nginx-a", az="us-east-1a"}

```

192.168.1.1 - - [06/16/2023:06:09:05 +0000] "GET / HTTP/1.1" 200 612 "-" "curl/7.74.0"
192.168.1.2 - - [06/16/2023:06:09:06 +0000] "GET /index.html HTTP/1.1" 200 612 "-" "wget/1.21.1"
192.168.1.3 - - [06/16/2023:06:09:07 +0000] "GET /about.html HTTP/1.1" 200 396 "-" "Safari/15.4"
192.168.1.4 - - [06/16/2023:06:09:08 +0000] "GET /contact.html HTTP/1.1" 200 299 "-" "Firefox/92.0"
192.168.1.5 - - [06/16/2023:06:09:09 +0000] "GET /blog/index.html HTTP/1.1" 200 824 "-" "Chrome/92.0.4515.131"
192.168.1.6 - - [06/16/2023:06:09:10 +0000] "GET /blog/post-1.html HTTP/1.1" 200 543 "-" "Opera/12.15"
192.168.1.7 - - [06/16/2023:06:09:11 +0000] "GET /blog/post-2.html HTTP/1.1" 200 456 "-" "Edge/182.0.1245.153"
192.168.1.8 - - [06/16/2023:06:09:12 +0000] "GET /blog/post-3.html HTTP/1.1" 200 369 "-" "IE/11.0"

```



2023-06-16 11:02.123456789

2023-06-16 11:12.123456789

{app="nginx", pod="nginx-b", az="us-east-1a"}

```

192.168.1.1 - - [06/16/2023:06:09:13 +0000] "GET / HTTP/1.1" 200 612 "-" "curl/7.74.0"
192.168.1.2 - - [06/16/2023:06:09:14 +0000] "GET /index.html HTTP/1.1" 200 612 "-" "wget/1.21.1"
192.168.1.3 - - [06/16/2023:06:09:15 +0000] "GET /about.html HTTP/1.1" 200 396 "-" "Safari/15.4"
192.168.1.4 - - [06/16/2023:06:09:16 +0000] "GET /contact.html HTTP/1.1" 200 299 "-" "Firefox/92.0"
192.168.1.5 - - [06/16/2023:06:09:17 +0000] "GET /blog/index.html HTTP/1.1" 200 824 "-" "Chrome/92.0.4515.131"
192.168.1.6 - - [06/16/2023:06:09:18 +0000] "GET /blog/post-1.html HTTP/1.1" 200 543 "-" "Opera/12.15"
192.168.1.7 - - [06/16/2023:06:09:19 +0000] "GET /blog/post-2.html HTTP/1.1" 200 456 "-" "Edge/182.0.1245.153"
192.168.1.8 - - [06/16/2023:06:09:20 +0000] "GET /blog/post-3.html HTTP/1.1" 200 369 "-" "IE/11.0"

```



2023-06-16 11:02.123456789

2023-06-16 11:12.123456789

{app="nginx", pod="nginx-c", az="us-east-1b"}

```

192.168.1.1 - - [06/16/2023:06:09:05 +0000] "GET / HTTP/1.1" 200 612 "-" "curl/7.74.0"
192.168.1.2 - - [06/16/2023:06:09:06 +0000] "GET /index.html HTTP/1.1" 200 612 "-" "wget/1.21.1"
192.168.1.3 - - [06/16/2023:06:09:07 +0000] "GET /about.html HTTP/1.1" 200 396 "-" "Safari/15.4"
192.168.1.4 - - [06/16/2023:06:09:08 +0000] "GET /contact.html HTTP/1.1" 200 299 "-" "Firefox/92.0"
192.168.1.5 - - [06/16/2023:06:09:09 +0000] "GET /blog/index.html HTTP/1.1" 200 824 "-" "Chrome/92.0.4515.131"
192.168.1.6 - - [06/16/2023:06:09:10 +0000] "GET /blog/post-1.html HTTP/1.1" 200 543 "-" "Opera/12.15"
192.168.1.7 - - [06/16/2023:06:09:11 +0000] "GET /blog/post-2.html HTTP/1.1" 200 456 "-" "Edge/182.0.1245.153"
192.168.1.8 - - [06/16/2023:06:09:12 +0000] "GET /blog/post-3.html HTTP/1.1" 200 369 "-" "IE/11.0"

```



2023-06-16 11:02.123456789

2023-06-16 11:12.123456789

{app="nginx", pod="nginx-d", az="us-east-1b"}

```

192.168.1.1 - - [06/16/2023:06:09:05 +0000] "GET / HTTP/1.1" 200 612 "-" "curl/7.74.0"
192.168.1.2 - - [06/16/2023:06:09:06 +0000] "GET /index.html HTTP/1.1" 200 612 "-" "wget/1.21.1"
192.168.1.3 - - [06/16/2023:06:09:07 +0000] "GET /about.html HTTP/1.1" 200 396 "-" "Safari/15.4"
192.168.1.4 - - [06/16/2023:06:09:08 +0000] "GET /contact.html HTTP/1.1" 200 299 "-" "Firefox/92.0"
192.168.1.5 - - [06/16/2023:06:09:09 +0000] "GET /blog/index.html HTTP/1.1" 200 824 "-" "Chrome/92.0.4515.131"
192.168.1.6 - - [06/16/2023:06:09:10 +0000] "GET /blog/post-1.html HTTP/1.1" 200 543 "-" "Opera/12.15"
192.168.1.7 - - [06/16/2023:06:09:11 +0000] "GET /blog/post-2.html HTTP/1.1" 200 456 "-" "Edge/182.0.1245.153"
192.168.1.8 - - [06/16/2023:06:09:12 +0000] "GET /blog/post-3.html HTTP/1.1" 200 369 "-" "IE/11.0"

```




80ba6fba

72447153

94d10ab7

14a8d9b3

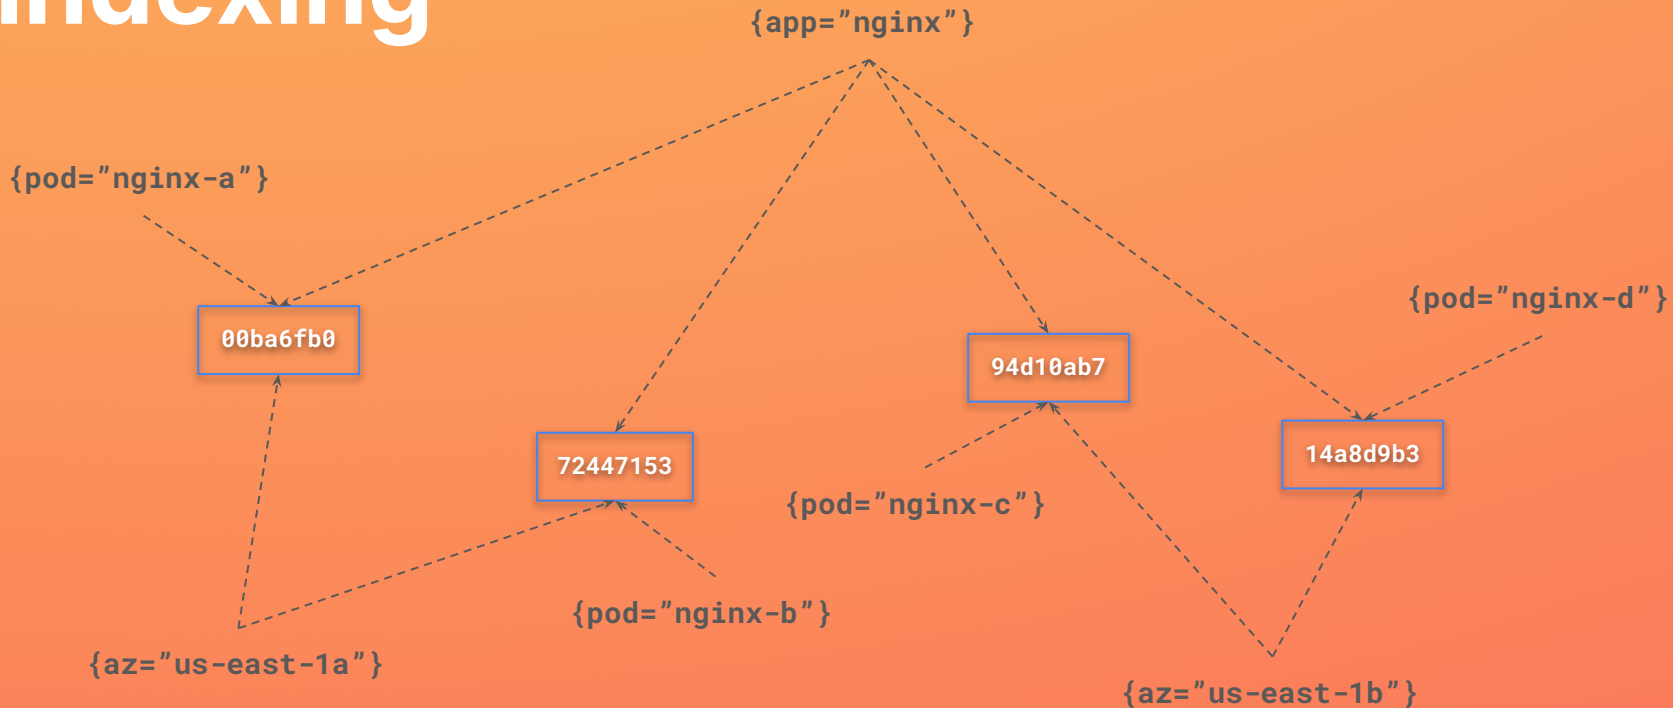


object storage

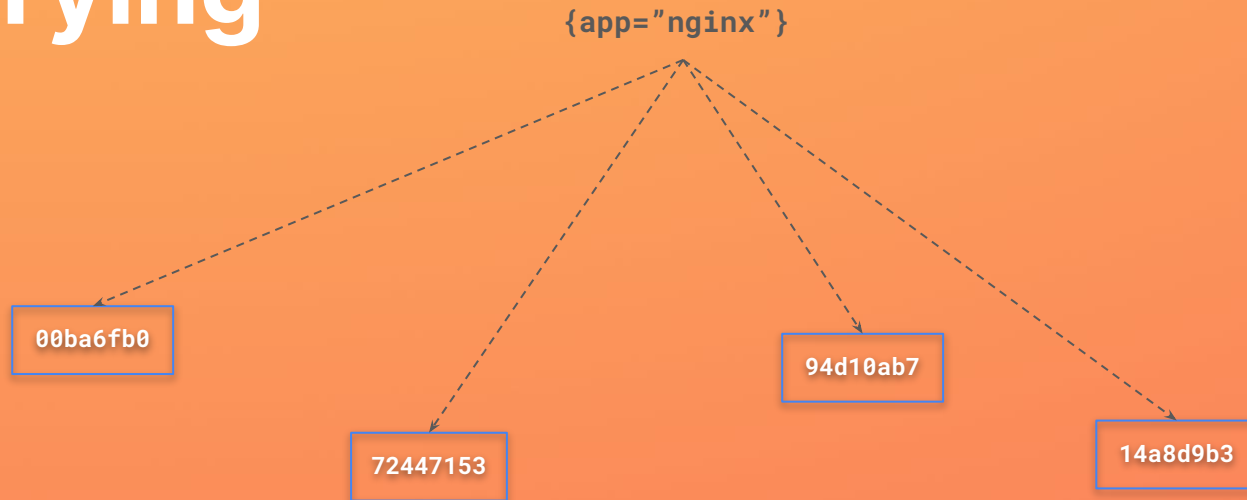
# ingestion



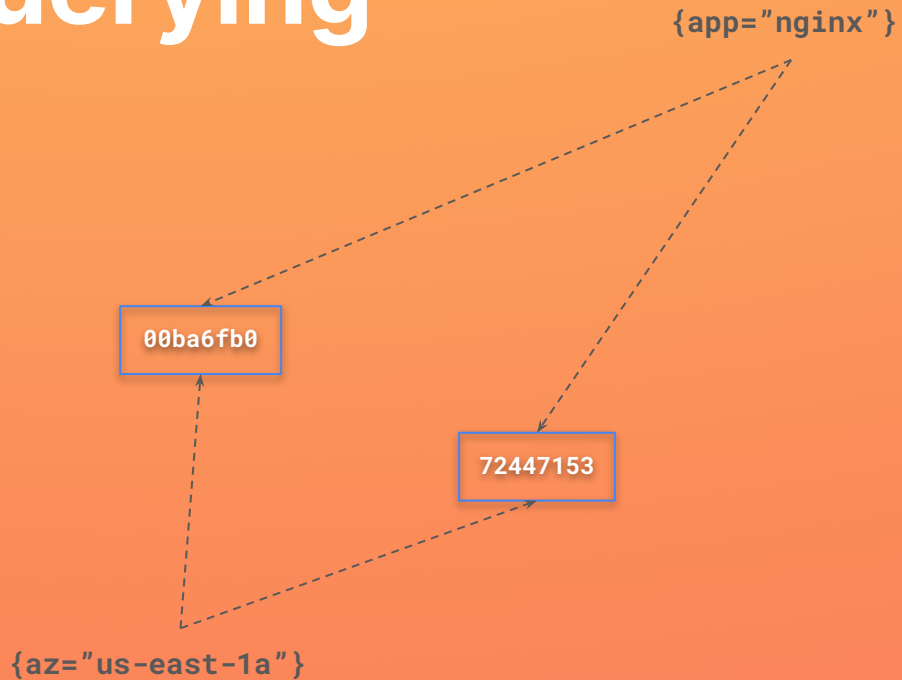
# indexing



# querying

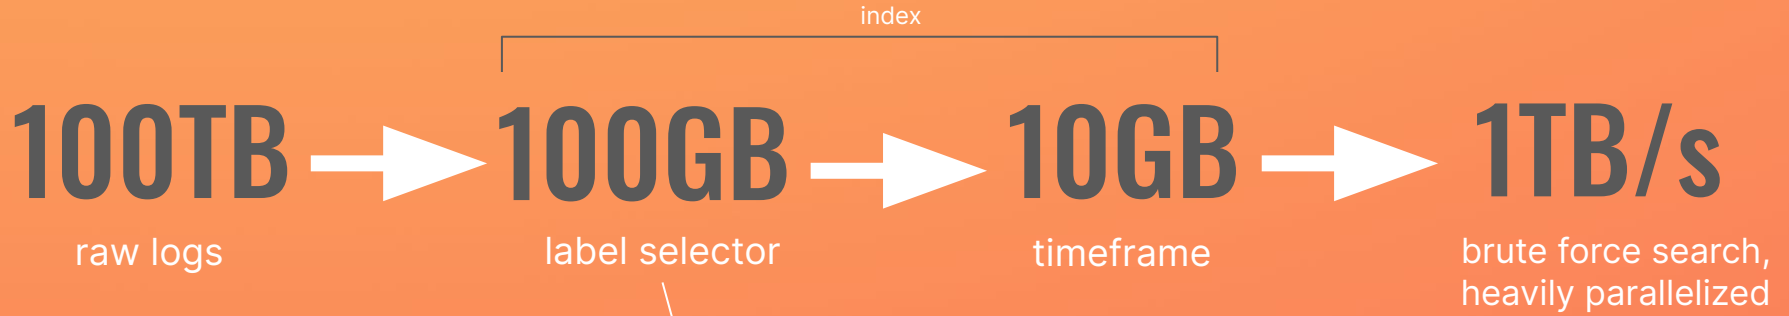


# querying





# fast queries



`{app="nginx", az="us-east-1a"} |= "12.34.56.78"`



# slow queries

**100TB**

raw logs

**100GB**


label selector


**10GB**

timeframe

**1TB/s**

brute force search,  
heavily parallelized





Grafana  
**Cloud**



# rate-limiting



# rate-limiting

- ⇒ decreased throughput
- ⇒ slower queries
- ⇒ frustrated users
- ⇒ SLO budget burn alerts
- ⇒ operator toil
- ⇒ "the incident"



two choices:

**query less or cache more**



choice A:

**query less**

**improve labels  
accelerated "needle-in-haystack"  
smarter query engine  
chunk compaction**



the problem:

**this shit takes time**





choice B:

**cache more**

**buy time to do the *right thing*<sup>TM</sup>  
easier lever to pull  
shorter delivery timeline**



# caching



logs

ingester

### chunks

2023-06-16 11:02.123456789 2023-06-16 11:12.123456789

```
{app="nginx", pod="nginx-a", az="us-east-1a"}
```

```
192.168.1.1 - - [06/16/2023:06:09:05 +0000] "GET / HTTP/1.1" 200 612 "-" "curl/7.74.0"
192.168.1.2 - - [06/16/2023:06:09:06 +0000] "GET /index.html HTTP/1.1" 200 612 "-" "wget/1.21.1"
192.168.1.3 - - [06/16/2023:06:09:07 +0000] "GET /about.html HTTP/1.1" 200 396 "-" "Safari/15.4"
192.168.1.4 - - [06/16/2023:06:09:08 +0000] "GET /contact.html HTTP/1.1" 200 299 "-" "Firefox/92.0"
192.168.1.5 - - [06/16/2023:06:09:09 +0000] "GET /blog/index.html HTTP/1.1" 200 824 "-" "Chrome/92.0.4515.131"
192.168.1.6 - - [06/16/2023:06:09:10 +0000] "GET /blog/post-1.html HTTP/1.1" 200 543 "-" "Opera/12.15"
192.168.1.7 - - [06/16/2023:06:09:11 +0000] "GET /blog/post-2.html HTTP/1.1" 200 456 "-" "Edge/102.0.1245.153"
192.168.1.8 - - [06/16/2023:06:09:12 +0000] "GET /blog/post-3.html HTTP/1.1" 200 369 "-" "IE/11.0"
```

---

2023-06-16 11:02.123456789 2023-06-16 11:12.123456789

```
{app="nginx", pod="nginx-b", az="us-east-1a"}
```

```
192.168.1.1 - - [06/16/2023:06:09:13 +0000] "GET / HTTP/1.1" 200 612 "-" "curl/7.74.0"
192.168.1.2 - - [06/16/2023:06:09:14 +0000] "GET /index.html HTTP/1.1" 200 612 "-" "wget/1.21.1"
192.168.1.3 - - [06/16/2023:06:09:15 +0000] "GET /about.html HTTP/1.1" 200 396 "-" "Safari/15.4"
192.168.1.4 - - [06/16/2023:06:09:16 +0000] "GET /contact.html HTTP/1.1" 200 299 "-" "Firefox/92.0"
192.168.1.5 - - [06/16/2023:06:09:17 +0000] "GET /blog/index.html HTTP/1.1" 200 824 "-" "Chrome/92.0.4515.131"
192.168.1.6 - - [06/16/2023:06:09:18 +0000] "GET /blog/post-1.html HTTP/1.1" 200 543 "-" "Opera/12.15"
192.168.1.7 - - [06/16/2023:06:09:19 +0000] "GET /blog/post-2.html HTTP/1.1" 200 456 "-" "Edge/102.0.1245.153"
192.168.1.8 - - [06/16/2023:06:09:20 +0000] "GET /blog/post-3.html HTTP/1.1" 200 369 "-" "IE/11.0"
```

---

2023-06-16 11:02.123456789 2023-06-16 11:12.123456789

```
{app="nginx", pod="nginx-c", az="us-east-1b"}
```


```
192.168.1.1 - - [06/16/2023:06:09:05 +0000] "GET / HTTP/1.1" 200 612 "-" "curl/7.74.0"
192.168.1.2 - - [06/16/2023:06:09:06 +0000] "GET /index.html HTTP/1.1" 200 612 "-" "wget/1.21.1"
192.168.1.3 - - [06/16/2023:06:09:07 +0000] "GET /about.html HTTP/1.1" 200 396 "-" "Safari/15.4"
192.168.1.4 - - [06/16/2023:06:09:08 +0000] "GET /contact.html HTTP/1.1" 200 299 "-" "Firefox/92.0"
192.168.1.5 - - [06/16/2023:06:09:09 +0000] "GET /blog/index.html HTTP/1.1" 200 824 "-" "Chrome/92.0.4515.131"
192.168.1.6 - - [06/16/2023:06:09:10 +0000] "GET /blog/post-1.html HTTP/1.1" 200 543 "-" "Opera/12.15"
192.168.1.7 - - [06/16/2023:06:09:11 +0000] "GET /blog/post-2.html HTTP/1.1" 200 456 "-" "Edge/102.0.1245.153"
192.168.1.8 - - [06/16/2023:06:09:12 +0000] "GET /blog/post-3.html HTTP/1.1" 200 369 "-" "IE/11.0"
```

---

2023-06-16 11:02.123456789 2023-06-16 11:12.123456789

```
{app="nginx", pod="nginx-d", az="us-east-1b"}
```

```
192.168.1.1 - - [06/16/2023:06:09:05 +0000] "GET / HTTP/1.1" 200 612 "-" "curl/7.74.0"
192.168.1.2 - - [06/16/2023:06:09:06 +0000] "GET /index.html HTTP/1.1" 200 612 "-" "wget/1.21.1"
192.168.1.3 - - [06/16/2023:06:09:07 +0000] "GET /about.html HTTP/1.1" 200 396 "-" "Safari/15.4"
192.168.1.4 - - [06/16/2023:06:09:08 +0000] "GET /contact.html HTTP/1.1" 200 299 "-" "Firefox/92.0"
192.168.1.5 - - [06/16/2023:06:09:09 +0000] "GET /blog/index.html HTTP/1.1" 200 824 "-" "Chrome/92.0.4515.131"
192.168.1.6 - - [06/16/2023:06:09:10 +0000] "GET /blog/post-1.html HTTP/1.1" 200 543 "-" "Opera/12.15"
192.168.1.7 - - [06/16/2023:06:09:11 +0000] "GET /blog/post-2.html HTTP/1.1" 200 456 "-" "Edge/102.0.1245.153"
192.168.1.8 - - [06/16/2023:06:09:12 +0000] "GET /blog/post-3.html HTTP/1.1" 200 369 "-" "IE/11.0"
```



03ba6f89

724a7153

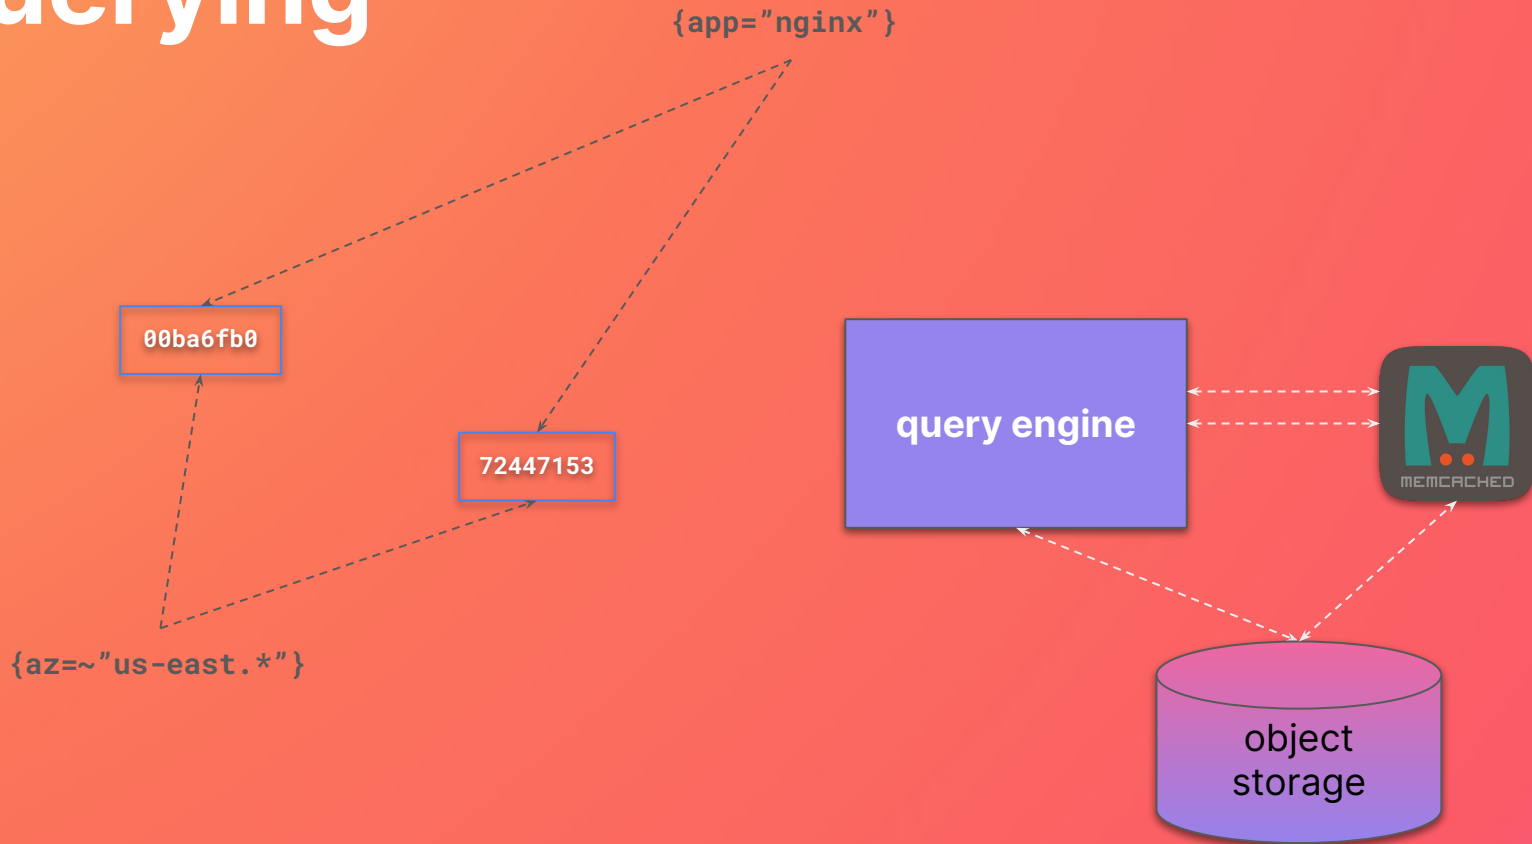
9ad10ab7

14abd9b3


# ingestion



# querying




# recency bias



global



# recency bias



loki-prod3



success metric:


**hit rate**



looks kinda... *good*, right?



problem:  
**churn**




items being evicted before  
being fetched **even once!**






problem:  
**churn**



**full cache replacements per hour**



# how much cache do we need?





**200 instances @ 1 vCPU, 6GB RAM**


**total capacity: ~1.2TB RAM**

running on shared


**n2-standard-32** nodes



# how much cache do we need?



# how much cache do we need?



**~50TB**  
cache

**~500Gbps**  
throughput



the challenge:

can we do it **cost-effectively**,  
whilst **maintaining performance**,  
in an **operationally-familiar way?**





**200 instances @ 1 vCPU, 6GB RAM**

**total capacity: ~1.2TB RAM**

running on shared

**n2-standard-32 nodes**

CPU: `<list> / 2 / <cpus>` = ~\$17.72/vCPU/month

RAM: `<list> / 2 / <gb>` = ~\$4.43/GB/month

200 vCPUs  $\Rightarrow$  \$3544

1200 GB RAM  $\Rightarrow$  \$5316

$\Rightarrow$  **\$8860 per month**







200 replicas @ 1 vCPU, 6GB RAM

Total capacity: ~1.2TB RAM  
running on shared  
n2-standard-32 nodes

$\text{vCPU} \times \text{list} / 2 / \text{cpus} = \sim \$17.72 / \text{vCPU} / \text{month}$

RAM: **~\$4.43 / GB / month**

200 vCPUs  $\Rightarrow$  \$3544

1200 GB RAM  $\Rightarrow$  \$5316

$\Rightarrow$  **\$8860 per month**



the goal:

**drive down cost per GB of cache**



**use memcached  
...with SSDs!**



# memcached “extstore”

the industry’s best-kept secret



[Home](#)

[About](#)

[Downloads](#)

[Blog](#)

[Mailing List](#)

[Wiki](#)

[Bugs](#)

## Caching beyond RAM: the case for NVMe - [Dormando](#) (June 12th, 2018)

Caching architectures at every layer of the stack embody an implicit tradeoff between performance and cost. These tradeoffs however are constantly shifting: new inflection points can emerge alongside advances in storage technology, changes in workload patterns, or fluctuations in hardware supply and demand.


**In this post we explore the design ramifications of the increasing cost of RAM on caching systems.** While RAM has always been expensive, DRAM prices have risen by over 50% in 2017, and high densities of RAM involve multi-socket NUMA machines, bloating power and overall costs. Concurrently, alternative storage technologies such as Flash and Optane continue to improve. They have specialized hardware interfaces, consistent performance, high density, and relatively low costs. While there is increasing economic incentive to explore offloading caching from RAM onto NVMe or NVM devices, the implications for performance are still not widely understood.

We will explore these design implications in the context of [Memcached](#), a distributed, simple, cache-focused key/value store. For a quick overview, see the [about page](#) or the [story tutorial](#).






**DRAM**



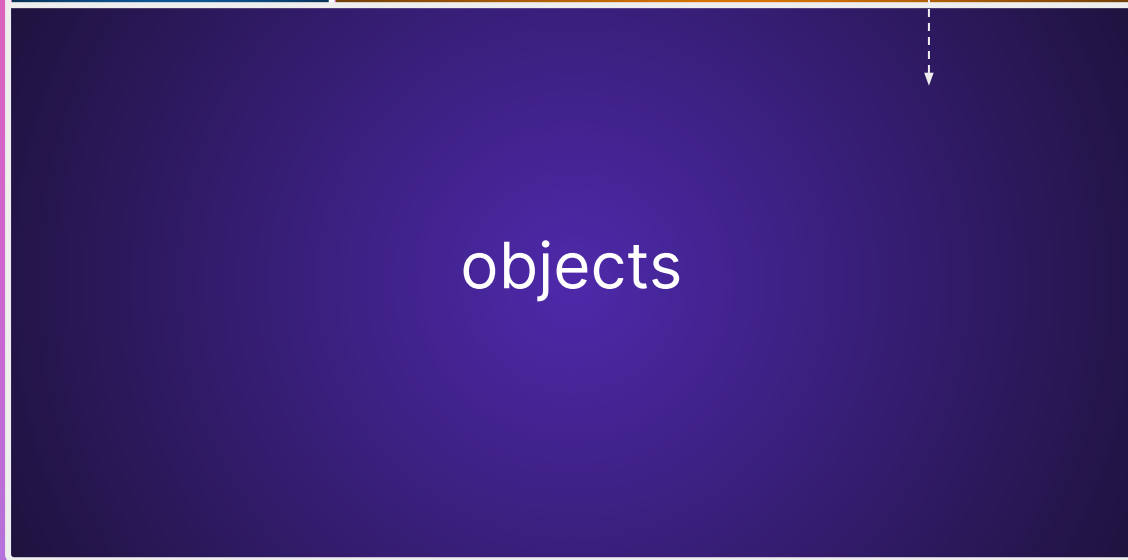


**mextstore**

**DRAM**



**SSD**



# using extstore



```
memcached <other flags>  
--extended  
ext_path=/mnt/disks/ssd0/datafile:345G
```



# using extstore



```
memcached <other flags>  
--extended  
ext_path=/mnt/disks/ssd0/datafile:345G,  
        /mnt/disks/ssd1/datafile:345G,  
        /mnt/disks/ssd2/datafile:345G...
```





# SSDs in cloud?



no "disk over network"

no multi-tenancy

**physically attached to hypervisor**





Google Cloud

## “Local SSDs”

- **375GB each @ \$30/month**
- **add up to 24 SSDs to most machine types**
- **660MB/s reads, 350MB/s writes**





# “Instance Storage”

- varies by instance type
- included in cost



**## ttv: time-to-value**  
2 weeks!





**33 instances @ 6 vCPU, 5GB RAM, 4 SSDs**

**total capacity: ~50TB SSD, 528Gbps**  
running on *dedicated*  
**n2-highcpu-8** nodes (16Gbps each)

CPU: `<list> / 2 / <cpus>` = ~\$13.08/vCPU/month  
RAM: `<list> / 2 / <gb>` = ~\$13.08/GB/month  
SSD: \$30 per disk (375GB) = \$0.08/GB/month

198 vCPUs ⇒ \$2590  
165 GB RAM ⇒ \$2158  
132 SSDs ⇒ \$3960

⇒ **\$8708 per month**





33 instances @ 6 vCPU, 5GB RAM, 4 SSDs

Total capacity: ~50TB SSD, 528Gbps  
running on dedicated  
n2-highcpu-8 nodes (16Gbps each)

CPU:  $\langle 112 \rangle \times 2 \times 60\%$

RAM:  $\langle 112 \rangle \times 2 \times 5\text{GB}$

SSD:  $\langle 112 \rangle \times 2 \times 4 \times 100\text{GB}$  = **\$0.08/GB/month**





**200 instances @ 1 vCPU, 6GB RAM**

**total capacity: ~1.2TB RAM**

running on shared

**n2-standard-32 nodes**

CPU: `<list> / 2 / <cpus>` = ~\$17.72/vCPU/month

RAM: `<list> / 2 / <gb>` = ~\$4.43/GB/month


200 vCPUs  $\Rightarrow$  \$3544

1200 GB RAM  $\Rightarrow$  \$5316

$\Rightarrow$  **\$8860 per month**







- ✓ **65% reduction in object storage reqs**
- ✓ **vastly improved reliability**
- ✓ **~2% overall TCO reduction**
- 🤔 **no change in performance**






Google Cloud Storage

# object storage

**\$5 per million writes**

\$0.005 per 1000



Amazon S3

**\$0.4 per million reads**

\$0.0004 per 1000


*no charge for bandwidth  
within the same region*



# rate-limits



# SLO (query throughput)




# hit rate




rollout




# churn rate



# cache effectiveness



# query performance



rollout





and then everything  
was perfect... the end







# trade-offs



# latency

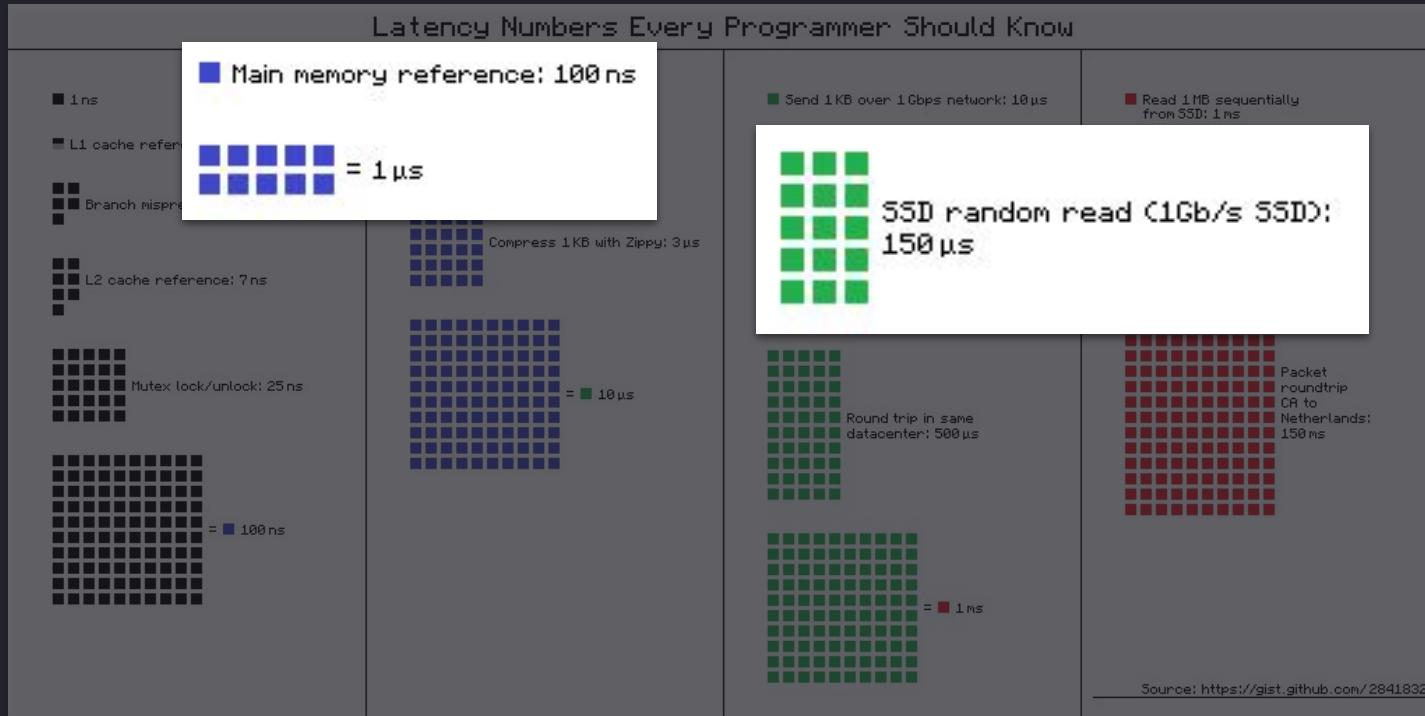
## Latency Numbers Every Programmer Should Know




Source: <https://gist.github.com/2841832>



# latency



# latency



rollout



# latency



rollout



# durability vs volatility



**Warning:** The performance gains from local SSDs require certain trade-offs in availability, durability, and flexibility. Because of these trade-offs, Local SSD storage is **not** automatically replicated and **all data on the local SSD may be lost** if the VM stops for any reason. See [Local SSD data persistence](#) for details.

	What happens to your data?
<b>System-initiated instance lifecycle events</b>	
The instance is rebooted	The data persists
The instance is stopped	The data does not persist
The instance is hibernated	The data does not persist
The instance is terminated	The data does not persist
The instance type is changed	The data does not persist *
A Windows AMI is created from the instance	The data does not persist in the created AMI **
An EBS-backed AMI is created from the instance	The data does not persist in the created AMI **
An instance store-backed AMI is created from the instance	The data persists in the AMI bundle uploaded to Amazon S3 ***
<b>User-initiated OS events</b>	
A shutdown is initiated	The data does not persist †
A restart is initiated	The data persists
<b>AWS scheduled events</b>	
Instance stop	The data does not persist
Instance reboot	The data persists
System reboot	The data persists
Instance retirement	The data does not persist
<b>Unplanned events</b>	



# disk-related issues

- disk fills up
- disk has bad sectors
- disk ages, performs poorly





solution:





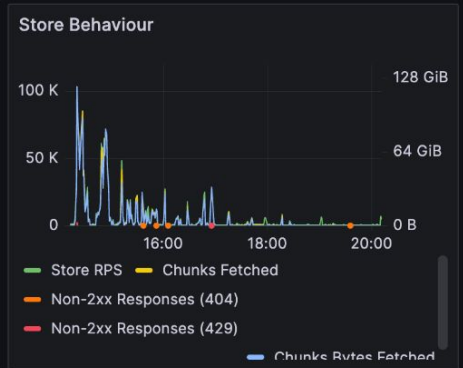



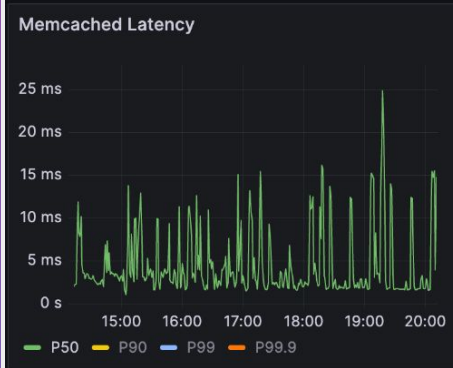
**kill the disk!**



# observability

- **prometheus**
  - memcached-exporter
  - node-exporter
- **alertmanager**
  - disk full alert
  - disk latency alert
  - memcached latency alert
- **grafana & loki (shocker!)**





# implications for the future of cloud databases?

- cost is becoming compelling
- capacious disks at near-DRAM speeds
- higher network throughput
  - getting higher as network is offloaded from CPU





shoutout to  
Ed & Alan!

thanks!  
questions?



@dannykopping

check out the  
blog post for  
more details



# why not Memomystore / Elasticache?

Memomystore for Memcached

37 x Cache  


vCPU usage: 135,050 vCPU hours

RAM usage: 40,515,000 GiB hours

Location: us-central1

**USD 367,336.00**

**Total Estimated Cost: USD 367,336.00 per 1 month**

Estimate Currency  
USD - US Dollar 



## Configure Amazon ElasticCache [Info](#)

### Cluster settings [Info](#)

#### Nodes

500

 Nodes can't be more than 250.

cache.m6g.8xlarge

Selected Instance:

**cache.m6g.8xlarge**

**vCPU: 32**

**Memory: 103.68 GiB**

**Network Performance: 12 Gigabit**

#### Utilization (On-Demand only)

Elasticache Instances will be always at 100% utilization, to stop incurring charges

100

**Total Upfront cost: 0.00 USD**

**Total Monthly cost: 865,780.00 USD**

Show Details 