THE NEXT GENERATION OF DATASTORES

HOT QUERIES, COLD STORAGE

PHILIPP KRENN

@XERAA

elastic

DEVELOPER

WHO USES DATASTORES?

WHO ENJOYS MANAGING DATASTORES?

"CLASSIC" APPROACH

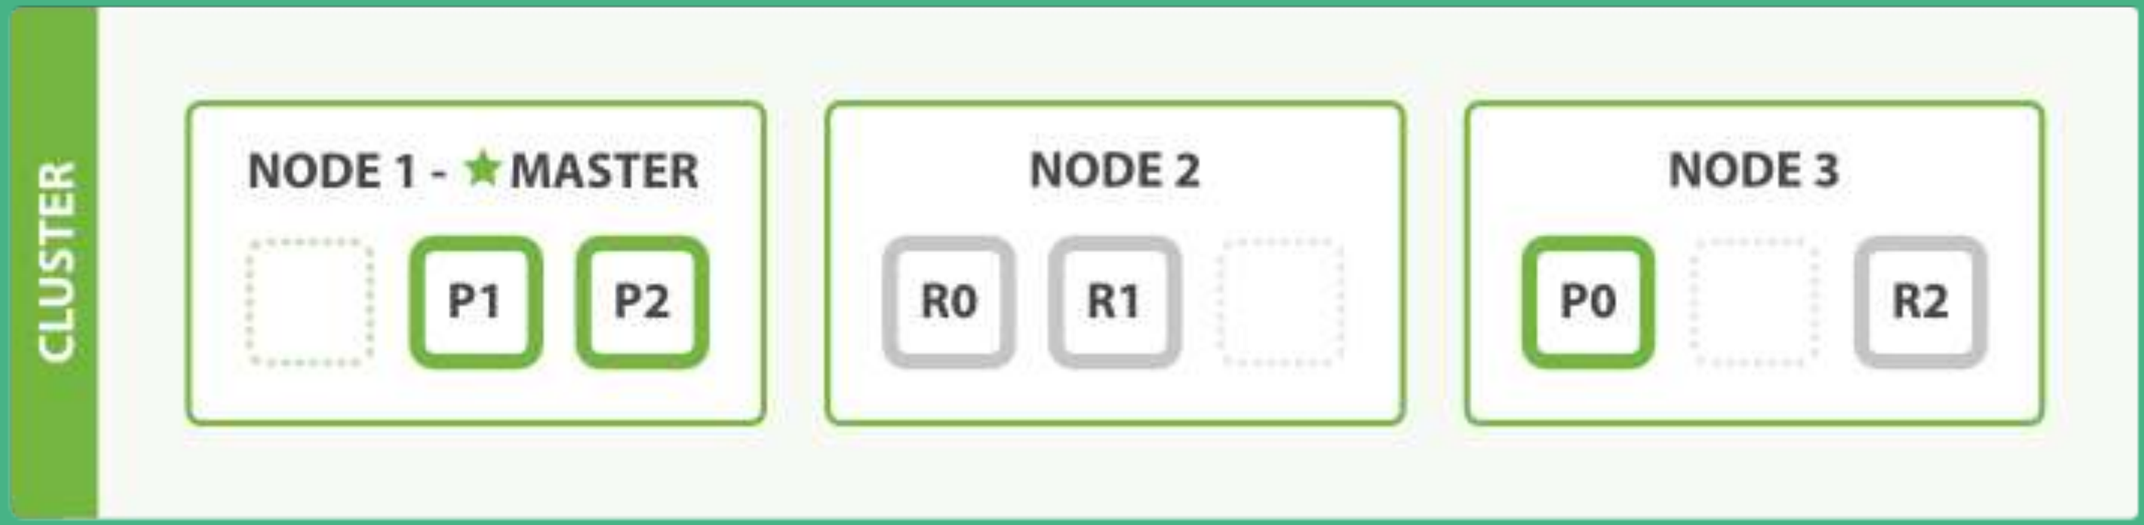

SERVERLESS

Pizza as a Service

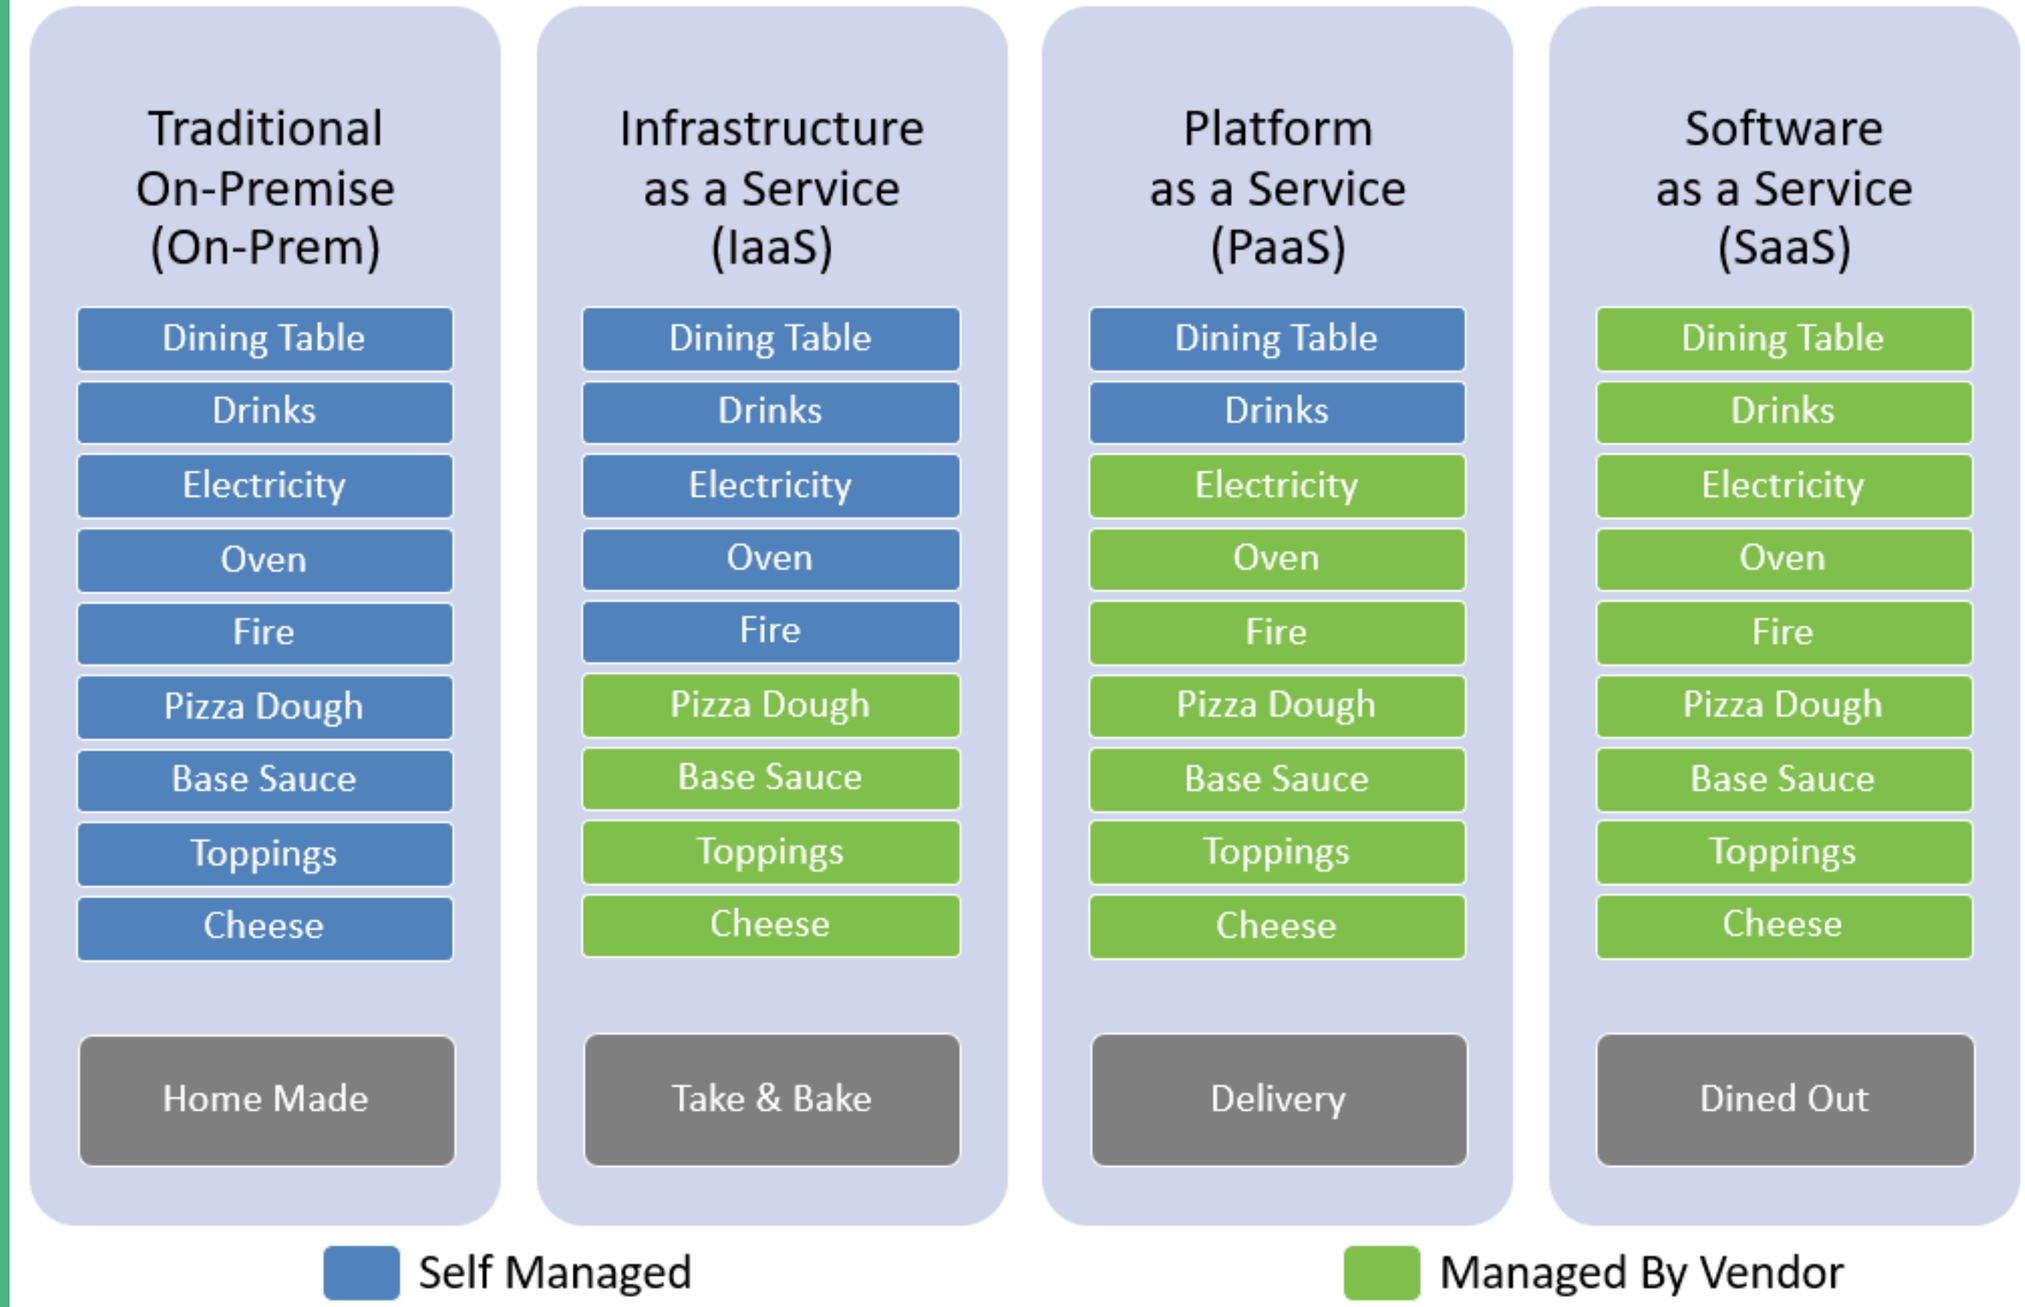

STATELESS

WHAT IS TODAY'S STORAGE STANDARD?

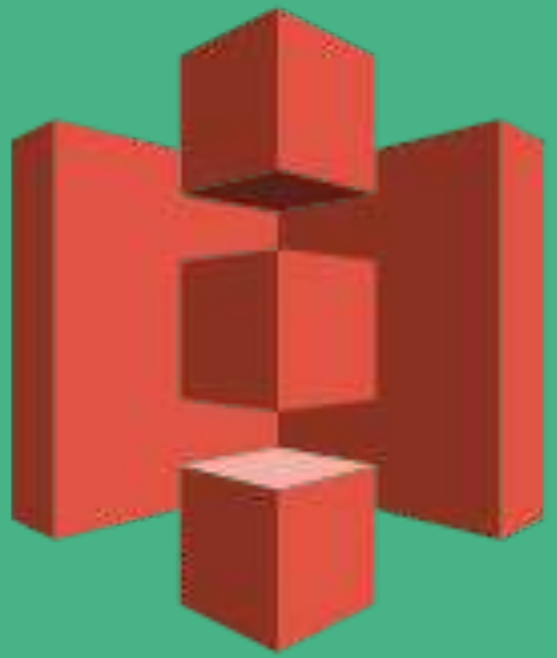

amazon
S3

S3 COMPATIBLE OBJECT STORES

GOOGLE CLOUD STORAGE, AZURE BLOB STORAGE, DIGITALOCEAN SPACES, MINIO,...

SIDENOTE ON "COMPATIBLE"

https://www.elastic.co/guide/en/elasticsearch/reference/current/repo-analysis-api.html

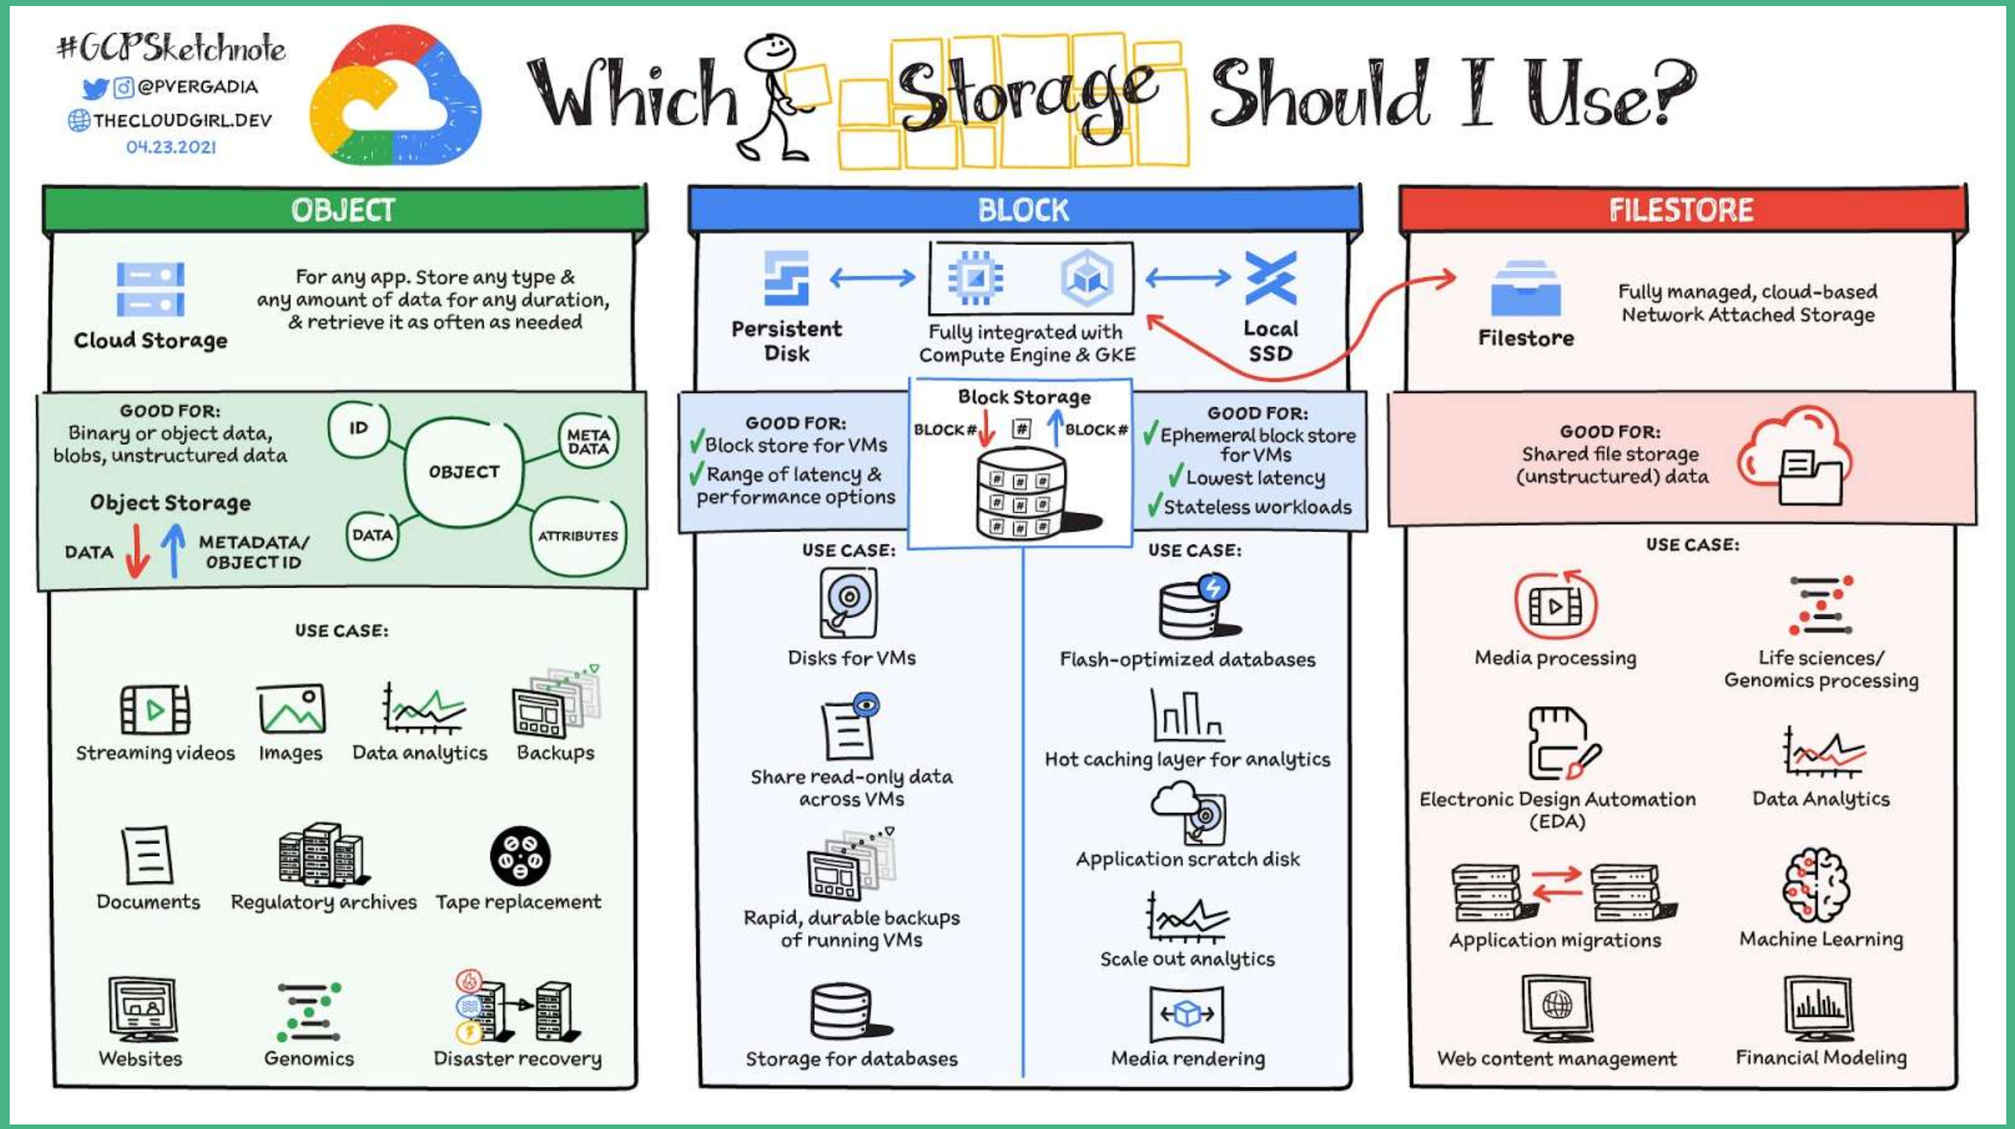

https://cloud.google.com/blog/topics/developers-practitioners/map-storage-options-

S3 FOR DATASTORES?

+ DURABILITY
? COST
- LATENCY

TRADEOFFS

“We ultimately decided that a few hundred milliseconds increase in the median latency was the right decision for our customers to save seconds, tens of seconds, and hundreds of seconds of p95/p99/max latency, respectively.”
https://www.datadoghq.com/blog/engineering/introducing-husky/

SPLIT STORAGE & COMPUTE

SCALE WRITES & READS INDEPENDENTLY

SCALE TO 0

LOCAL DISK FOR CACHING

PAY PER EXECUTION

OBJECT STORE PERFORMANCE CHARACTERISTICS

ABSTRACTION != MAGIC

DIFFERENT FOR EACH PROVIDER

GCP: RAMP UP

“If your request rate is less than 1000 write requests per second or 5000 read requests per second, then no ramp-up is needed. If your request rate is expected to go over these thresholds, you should start with a request rate below or near the thresholds and then double the request rate no faster than every 20 minutes.”
https://cloud.google.com/storage/docs/request-rate

GCP: AVOID HOTSPOTS

“Auto-scaling of an index range can be slowed when using sequential names, such as object keys based on a sequence of numbers or timestamp. This occurs because requests are constantly shifting to a new index range, making redistributing the load harder and less effective.”
https://cloud.google.com/storage/docs/request-rate#naming-convention

S3 CONSIDERATIONS

SHORT ROUND-TRIPS
FAST ACCESS CLASS

EXAMPLES

NEON (SERVERLESS POSTGRESQL)
YUGABYTEDB
COCKROACHDB CLOUD
CLICKHOUSE CLOUD
FIREBOLT CLOUD
...MANY MORE

ELASTICSEARCH

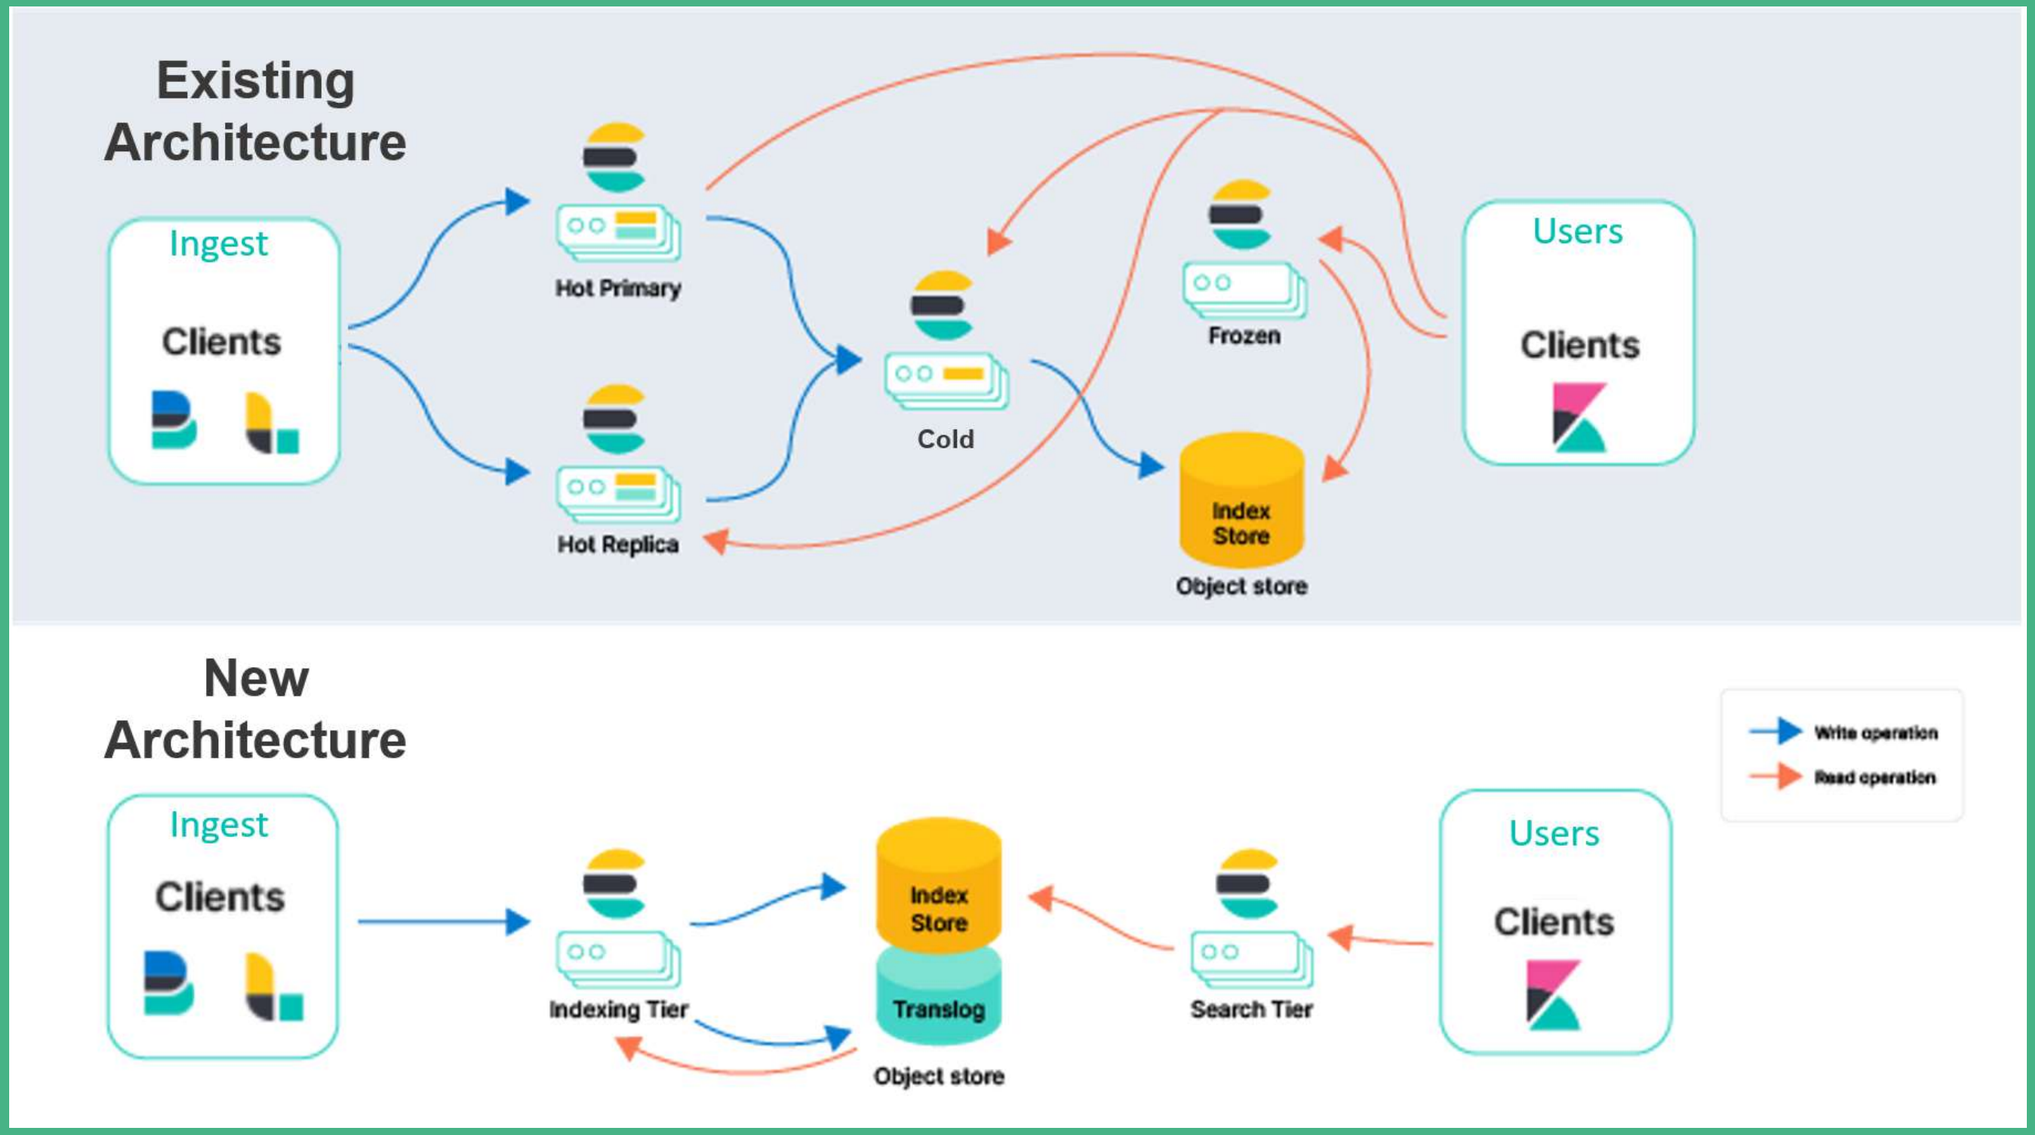

CHALLENGES

STATELESS MASTER
TRANSLOG
REAL-TIME GET

Indexing Throughput (Higher the better)

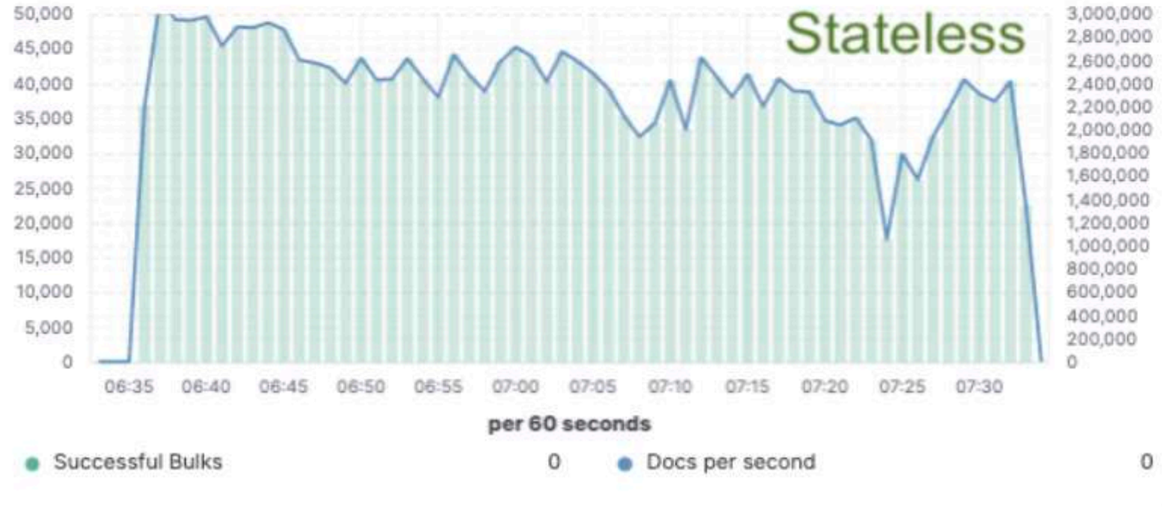

CPU Usage (Lower the better)

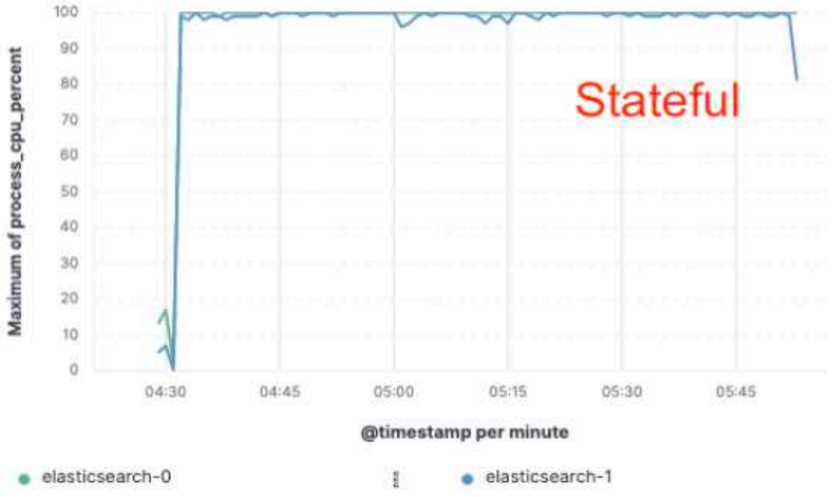

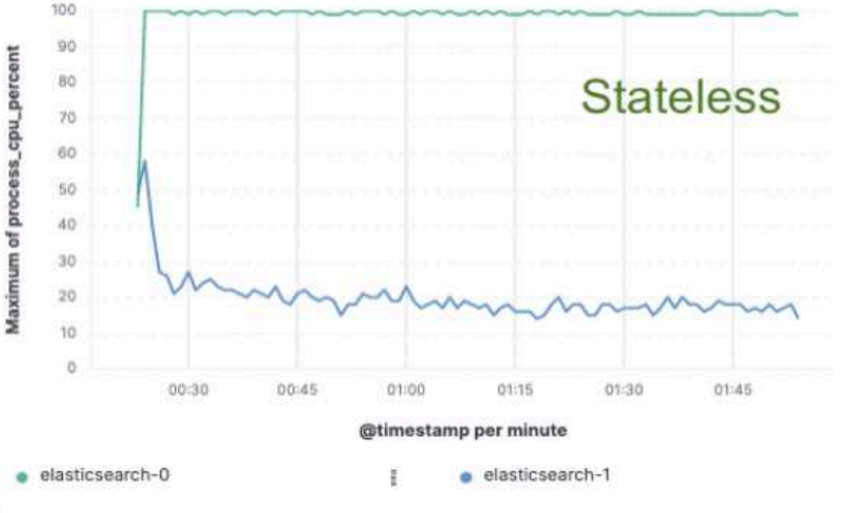

CONCLUSION

SERVERLESS & STATELESS

CLOUD NATIVE

PS: MORE THAN AUTOSCALING

THE NEXT GENERATION OF DATASTORES

HOT QUERIES, COLD STORAGE

PHILIPP KRENN

@XERAA