CLP: Efficient and Scalable Search on Compressed Text Logs

Kirk Rodrigues, Yu Luo, Ding Yuan

UNIVERSITY OF TORONTO

YScope

Compressed Log Processor

Lossless log compression

…better than general-purpose compressors

Can search compressed logs

…without decompression

…with good performance

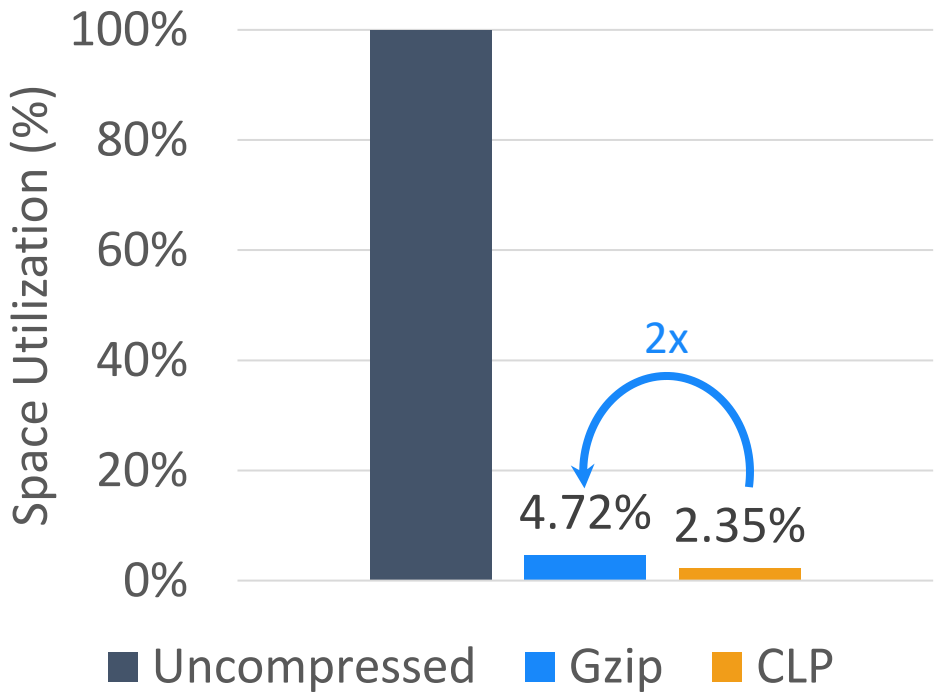

The Log Management Pipeline

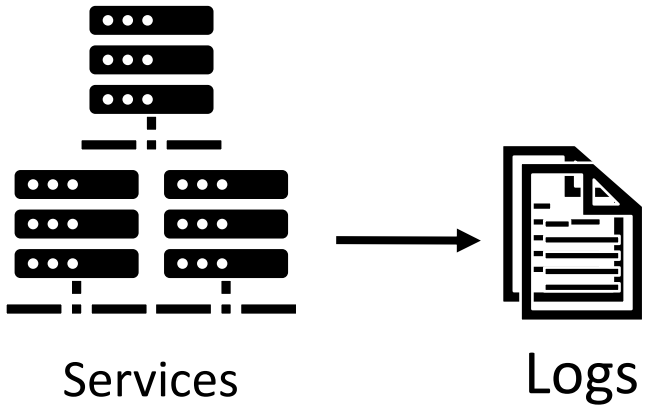

Example log message Timestamp Variables Static text

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container:
[NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

The Log Management Pipeline

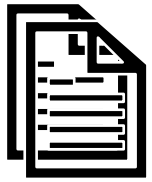

Logs

- Provide crucial runtime information
- Widely used for many purposes

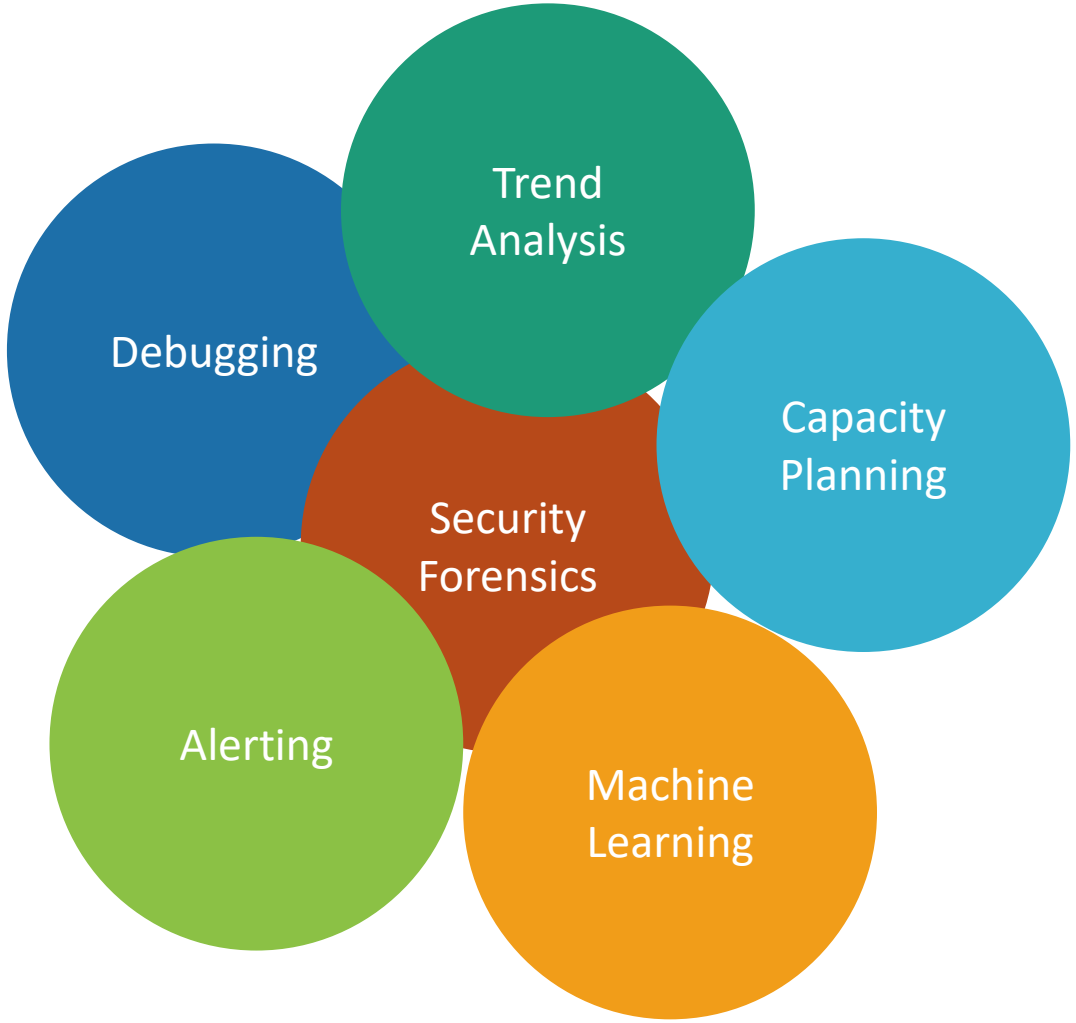

The Log Management Pipeline

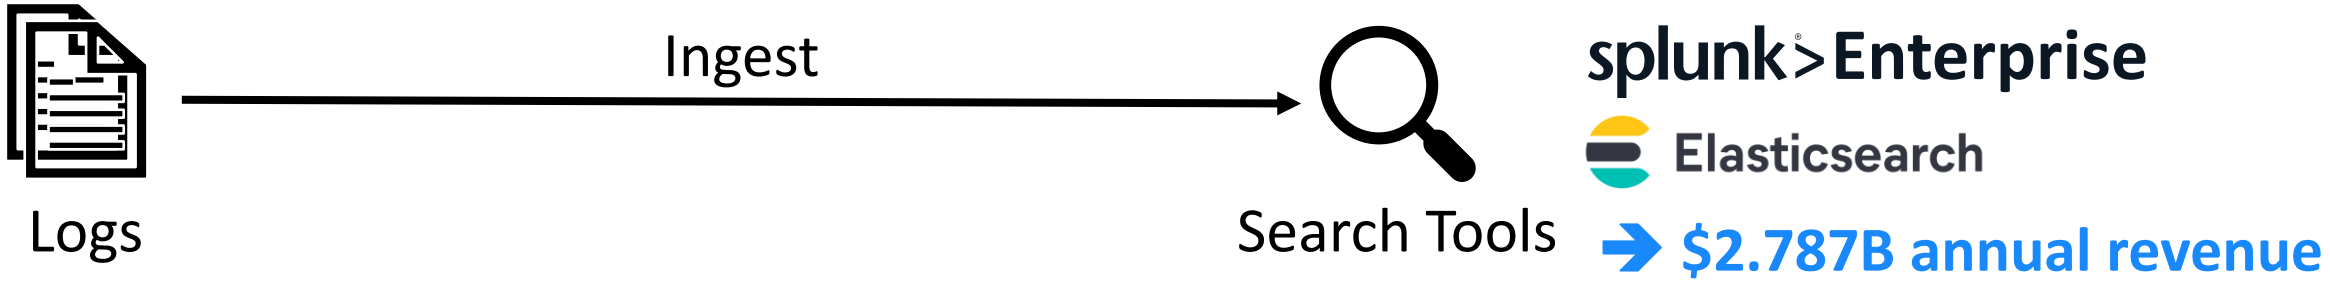

- Provide crucial runtime information
- Widely used for many purposes

- Consume lots of resources

The Log Management Pipeline

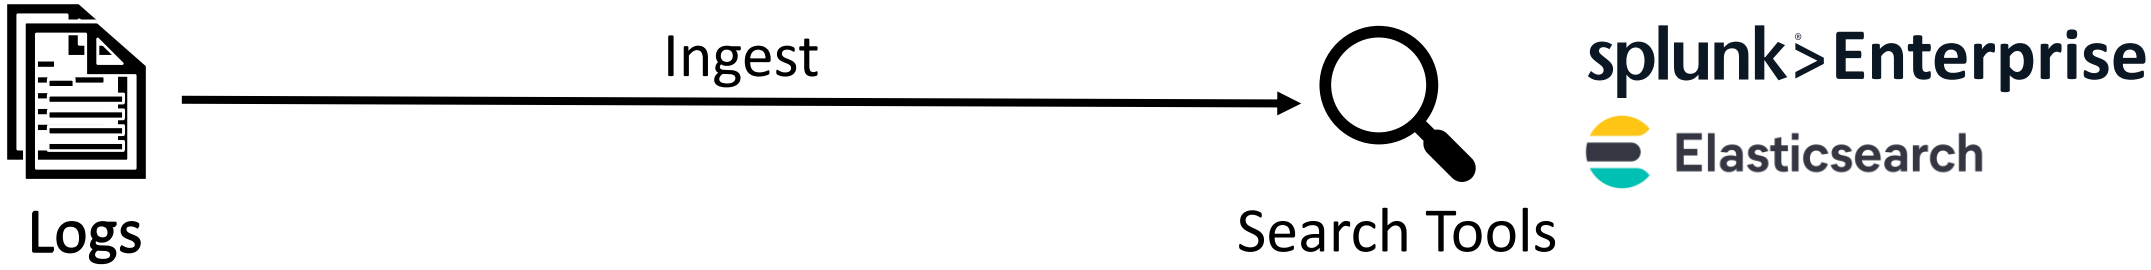

- Provide crucial runtime information
- Widely used for many purposes
- Companies generate petabytes of logs

- Consume lots of resources

eBay generated 1.2 PB of logs per day in 2018

HDD storage cost	$\frac{2c}{\text{GB}}{\text{month}}$
1.2 PB/day annual storage cost	$\$56{,}031{,}707$

The Log Management Pipeline

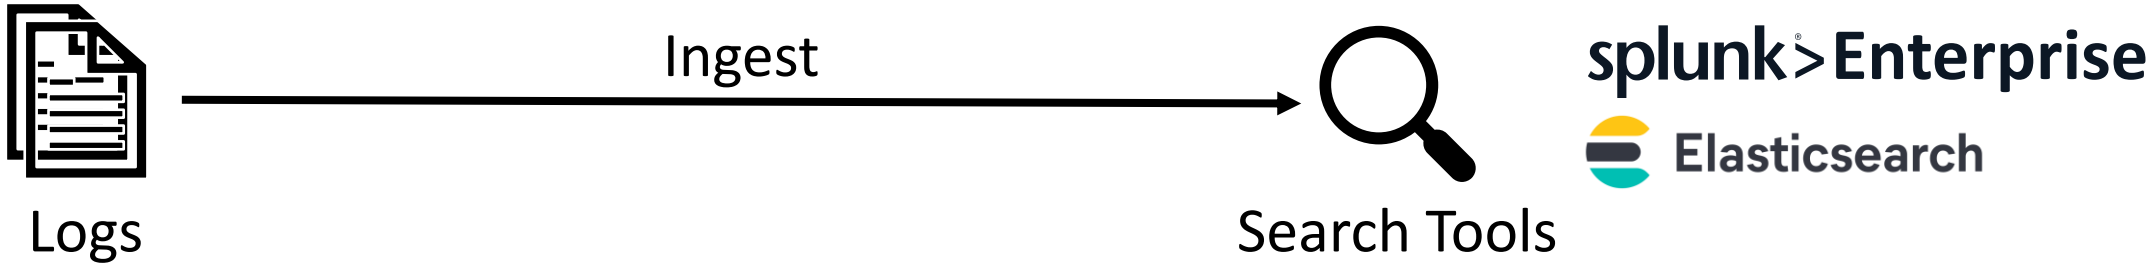

- Provide crucial runtime information
- Widely used for many purposes
- Companies generate petabytes of logs

- Consume lots of resources
- Build indexes ➔ adds storage overhead

The Log Management Pipeline

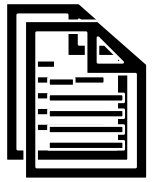

Logs

Ingest

splunk>Enterprise

Elasticsearch

Search Tools

- Provide crucial runtime information
- Widely used for many purposes
- Companies generate petabytes of logs

- Consume lots of resources
- Build indexes ➔ adds storage overhead
- Can only retain indexed logs for weeks

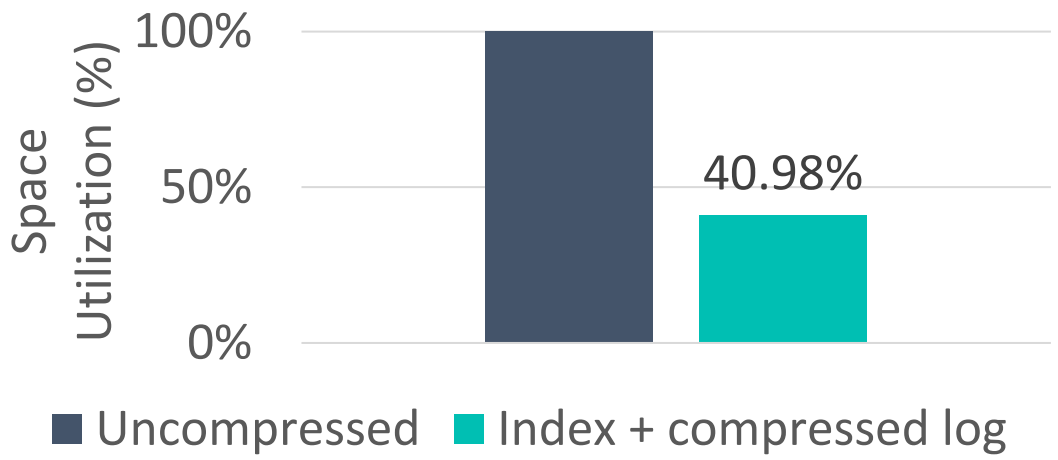

The Log Management Pipeline

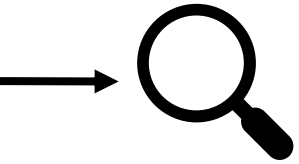

splunk>Enterprise

Elasticsearch

Search Tools

Compress

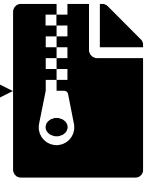

gzip

Archival Tools

- Consume lots of resources
- Build indexes ➔ adds storage overhead
- Can only retain indexed logs for weeks

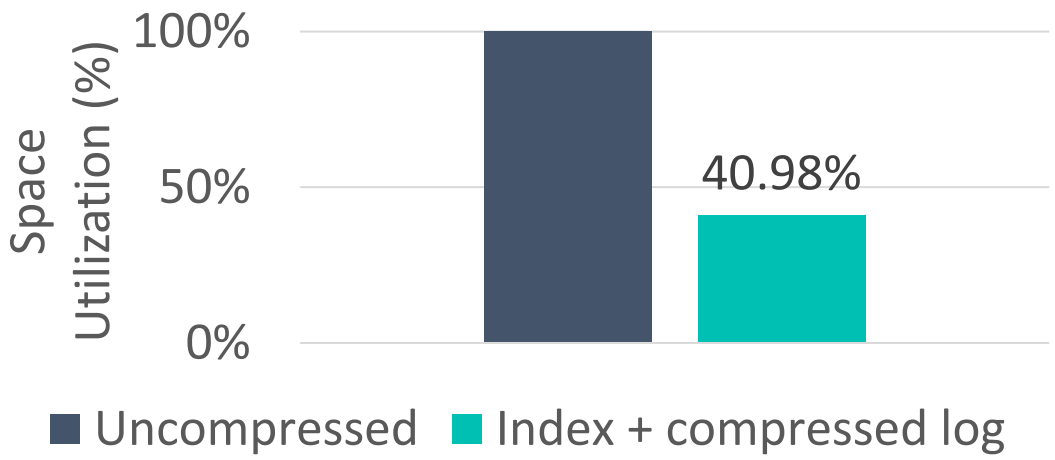

- Unsearchable once compressed

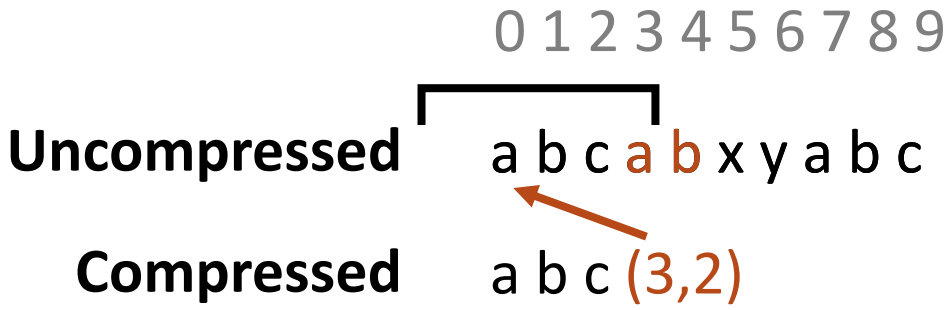

The Log Management Pipeline

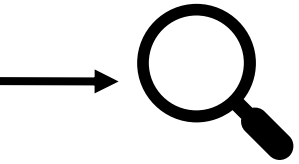

splunk>Enterprise

Elasticsearch

Search Tools

Compress

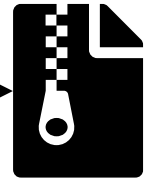

gzip

Archival Tools

- Consume lots of resources
- Build indexes ➔ adds storage overhead
- Can only retain indexed logs for weeks

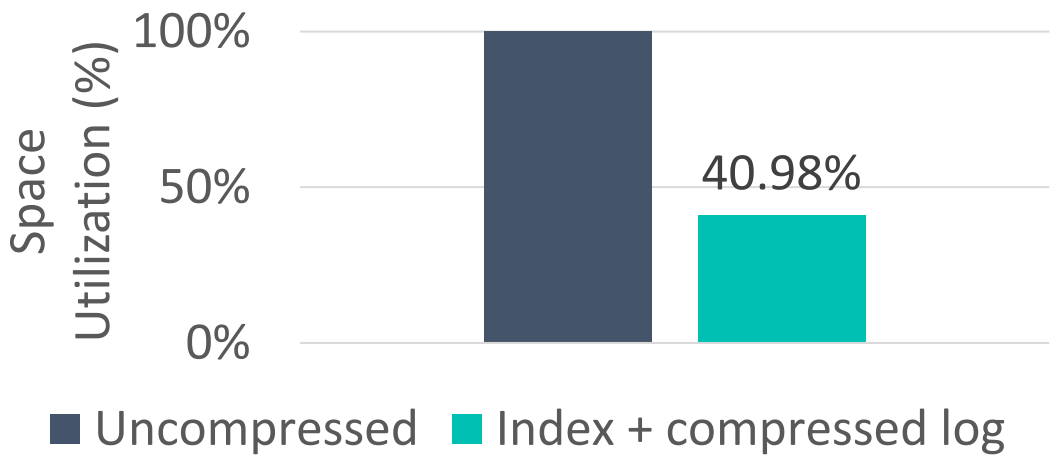

- Unsearchable once compressed

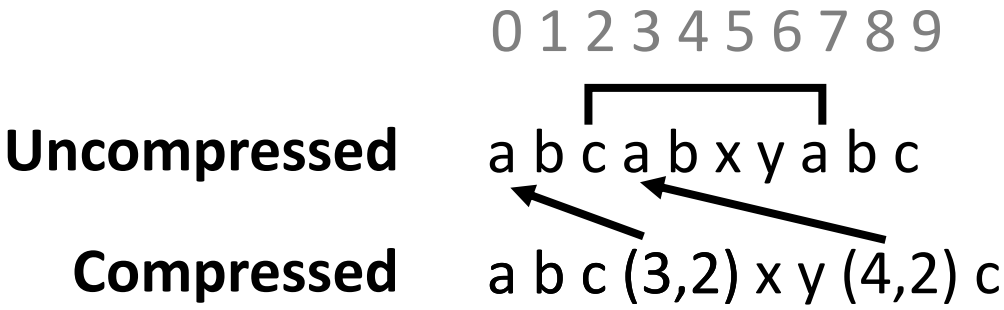

The Log Management Pipeline

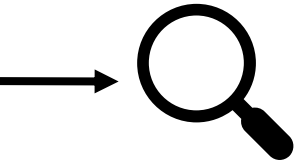

Search Tools

splunk>Enterprise

Elasticsearch

- Consume lots of resources
- Build indexes ➔ adds storage overhead
- Can only retain indexed logs for weeks

Compress

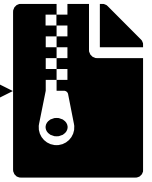

gzip

Archival Tools

- Unsearchable once compressed

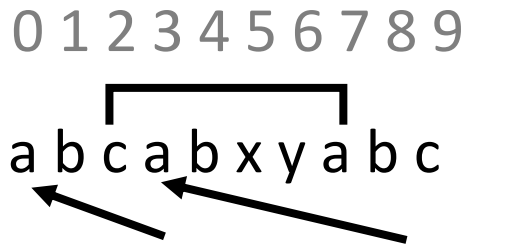

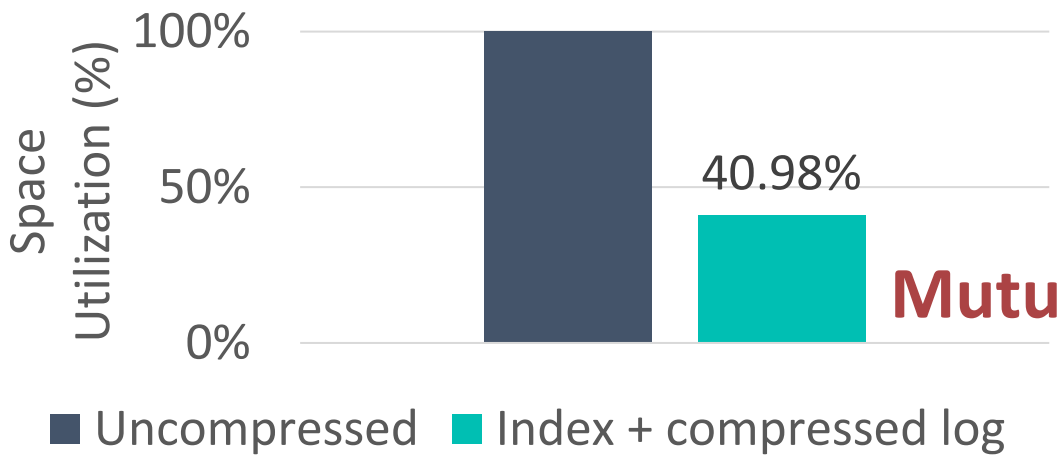

Mutually Exclusive

Demo

Compression

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container: [NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Compression

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container:
[NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

2020-01-02T03:04:06.006 INFO Task task_13 assigned to container:
[NodeAddress:172.128.0.42, ContainerID:container_16], operation took 1.221 seconds

2020-01-02T03:04:09.006 INFO Task task_14 assigned to container:
[NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.115 seconds

Compression

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container: [NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Log Type Dictionary

ID	Log Type

Variable Dictionary

ID	Variable Value

Compression

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container: [NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Log Type Dictionary

ID	Log Type
4	INFO Task assigned to container: [NodeAddress: , ContainerID:], operation took seconds

Variable Dictionary

ID	Variable Value

Compression

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container:
[NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Log Type Dictionary

ID	Log Type
4	INFO Task assigned to container: [NodeAddress: , ContainerID:], operation took seconds

Variable Dictionary

ID	Variable Value
8	task_12
9	172.128.0.41
10	container_15

Compression

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container:
[NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Log Type Dictionary

ID	Log Type
4	INFO Task assigned to container: [NodeAddress: , ContainerID:], operation took seconds

Variable Dictionary

ID	Variable Value
8	task_12
9	172.128.0.41
10	container_15

Encoded Message

Timestamp	Log Type ID	Variable Values
1577934245006	4	8 9 10 0x3FD570A3D70A3D71

Compression

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container:
[NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Log Type Dictionary

ID	Log Type
4	INFO Task assigned to container: [NodeAddress: , ContainerID:], operation took seconds

Variable Dictionary

ID	Variable Value
8	task_12
9	172.128.0.41
10	container_15

Encoded Message

Timestamp	Log Type ID	Variable Values
1577934245006	4	8 9 10 0x3FD570A3D70A3D71

Compression

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container:
[NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Log Type Dictionary

ID	Log Type
4	INFO Task assigned to container: [ess: , ContainerID:], operation took seconds

Shared between messages

Variable Dictionary

ID	Variable Value
8	task_12
9	172.128.0.41
10	container_15

Encoded Message

Timestamp	Log Type ID	Variable Values
1577934245006	4	8 9 10 0x3FD570A3D70A3D71

Search

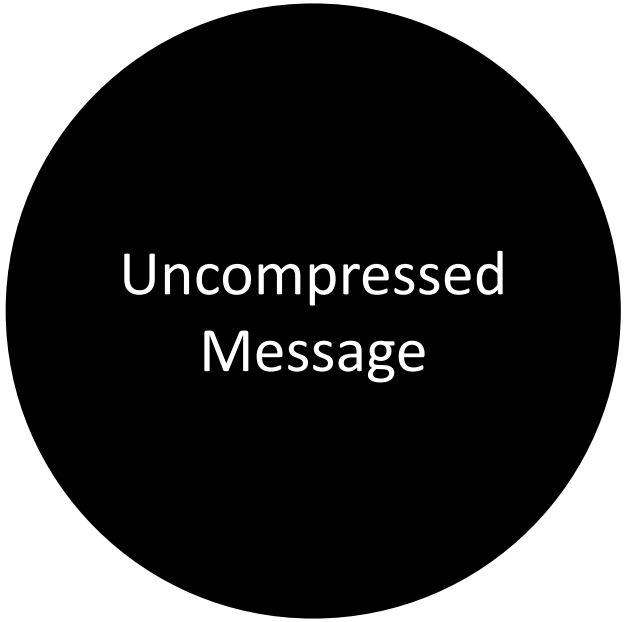

Search

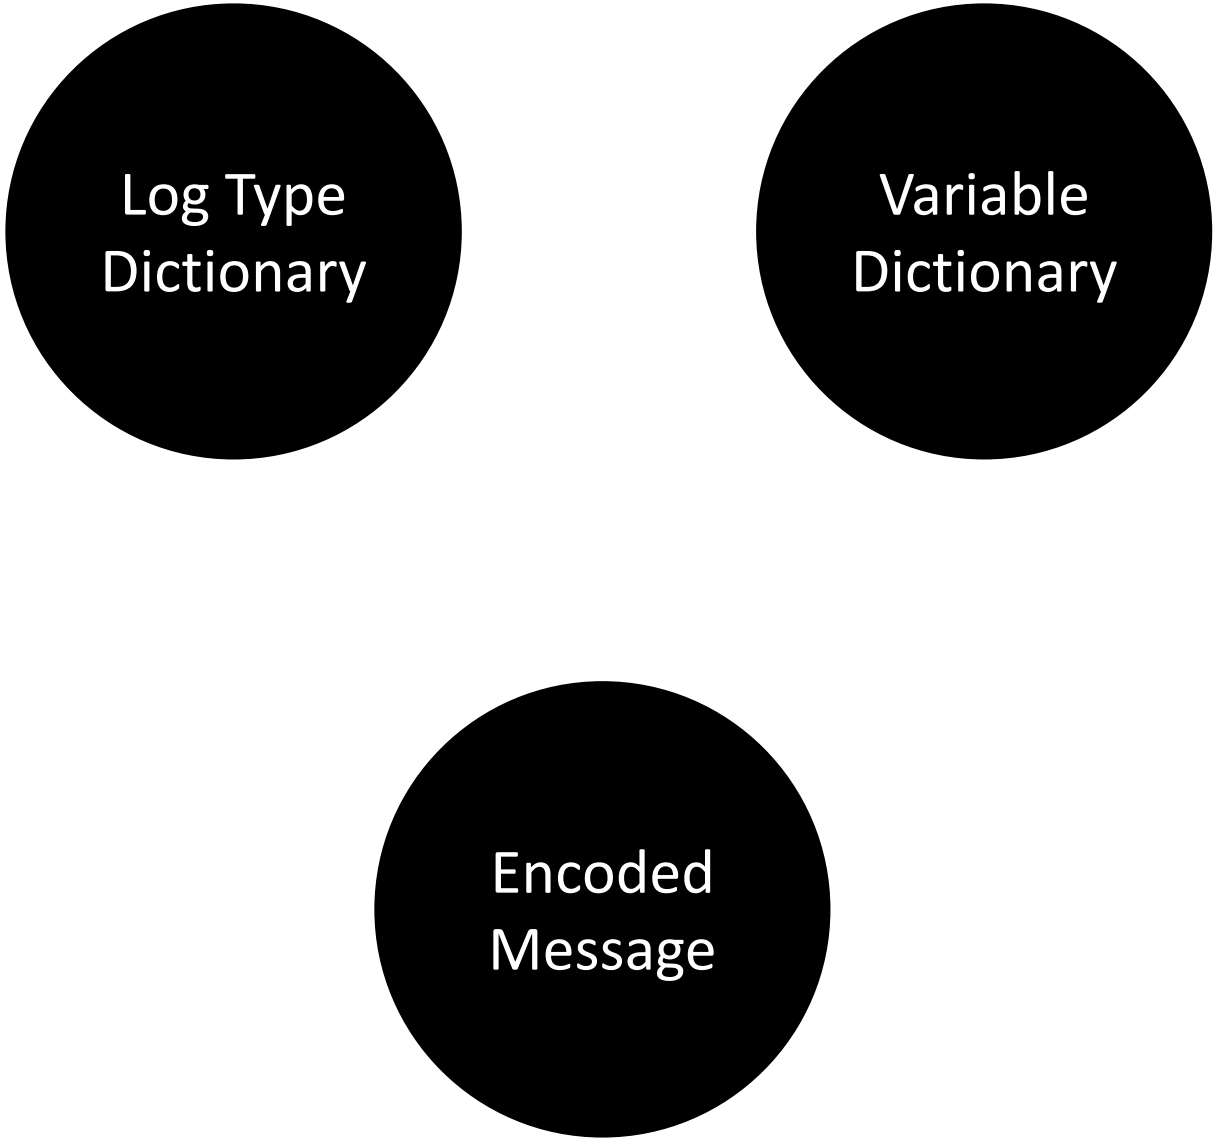

Search

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container: [NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Task * assigned to container*:172.128*

Search

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container: [NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Task * assigned to container*:172.128*

Search

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container:
[NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Task * assigned to container*:172.128*

- Dictionary variable?
- Log type?

Search

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container: [NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Task * assigned to container*:172.128*

Search

2020-01-02T03:04:05.006 INFO Task task_12 assigned to container:
[NodeAddress:172.128.0.41, ContainerID:container_15], operation took 0.335 seconds

Task * assigned to container*:172.128*

- Dictionary variable?
- Encoded variable?
- Log type?

Search

Task * assigned to container*:172.128*

#	Log type	Variables
1	Task * assigned to container*:172.128*	-
2	Task * assigned to container*:	172.128* (IP address)
3	Task * assigned to container*:	172.128* (float)
4	Task * assigned to :172.128*	container*
5	Task * assigned to :	container*, 172.128* (IP address)
6	Task * assigned to :	container*, 172.128* (floating point)

Evaluation

CLP’s compression ratio & speed

CLP’s search performance

Lots more detail in the paper!

Tested Tools

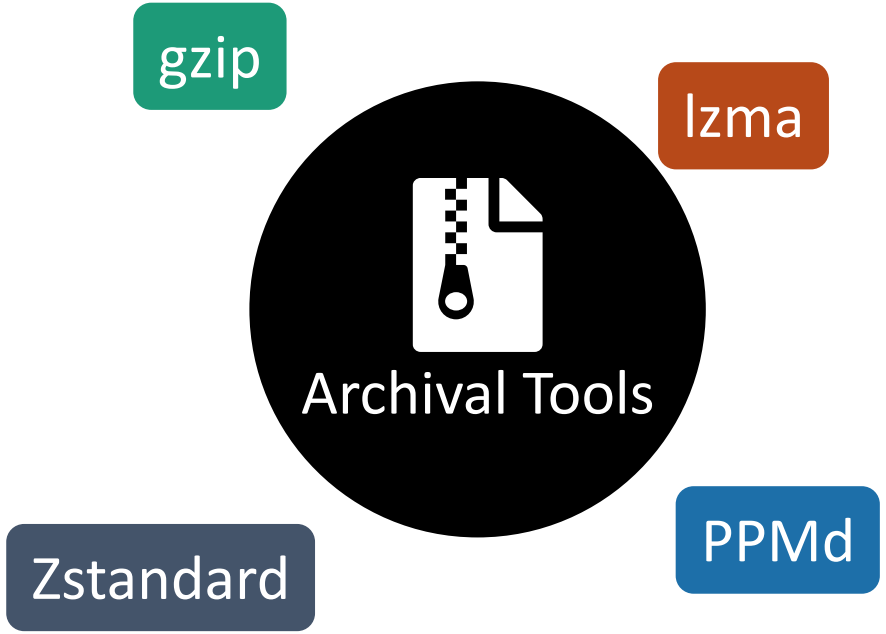

Compression Ratio

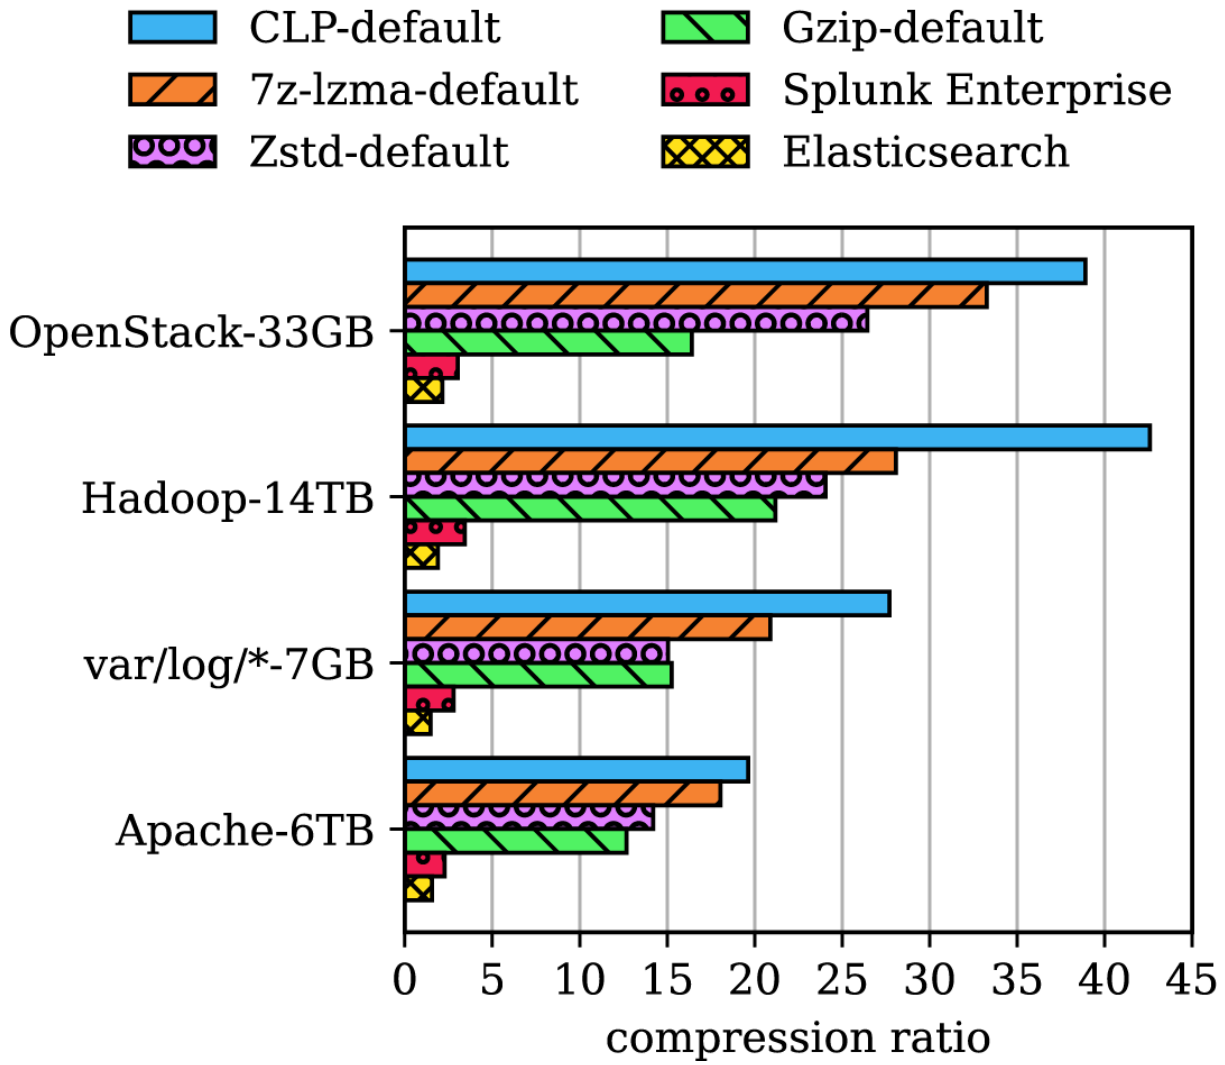

Tool	Average Compression Ratio
CLP	32.20
Gzip	16.38
Splunk Enterprise	2.86
Elasticsearch	1.75

Compression vs Ingestion Speed

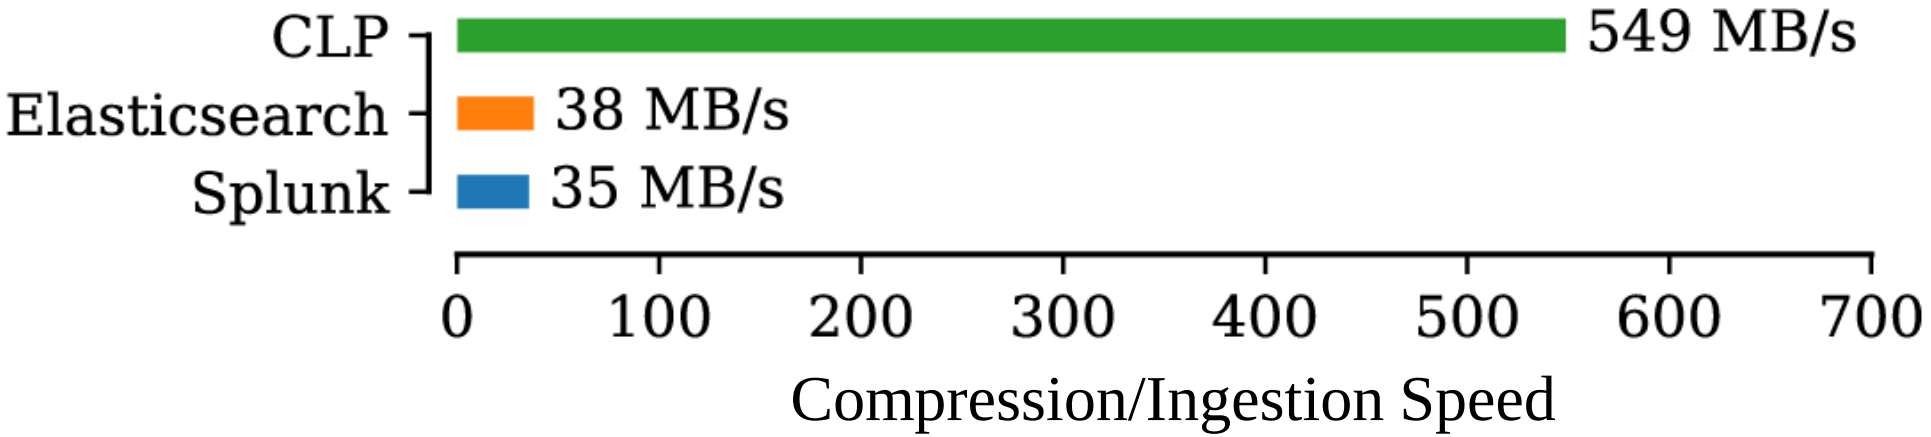

Search Performance

Query Benchmark

- Designed to exercise all of CLP's execution paths
  - Log type queries, variable queries, etc.
- Queries which return few and many results

Search Performance

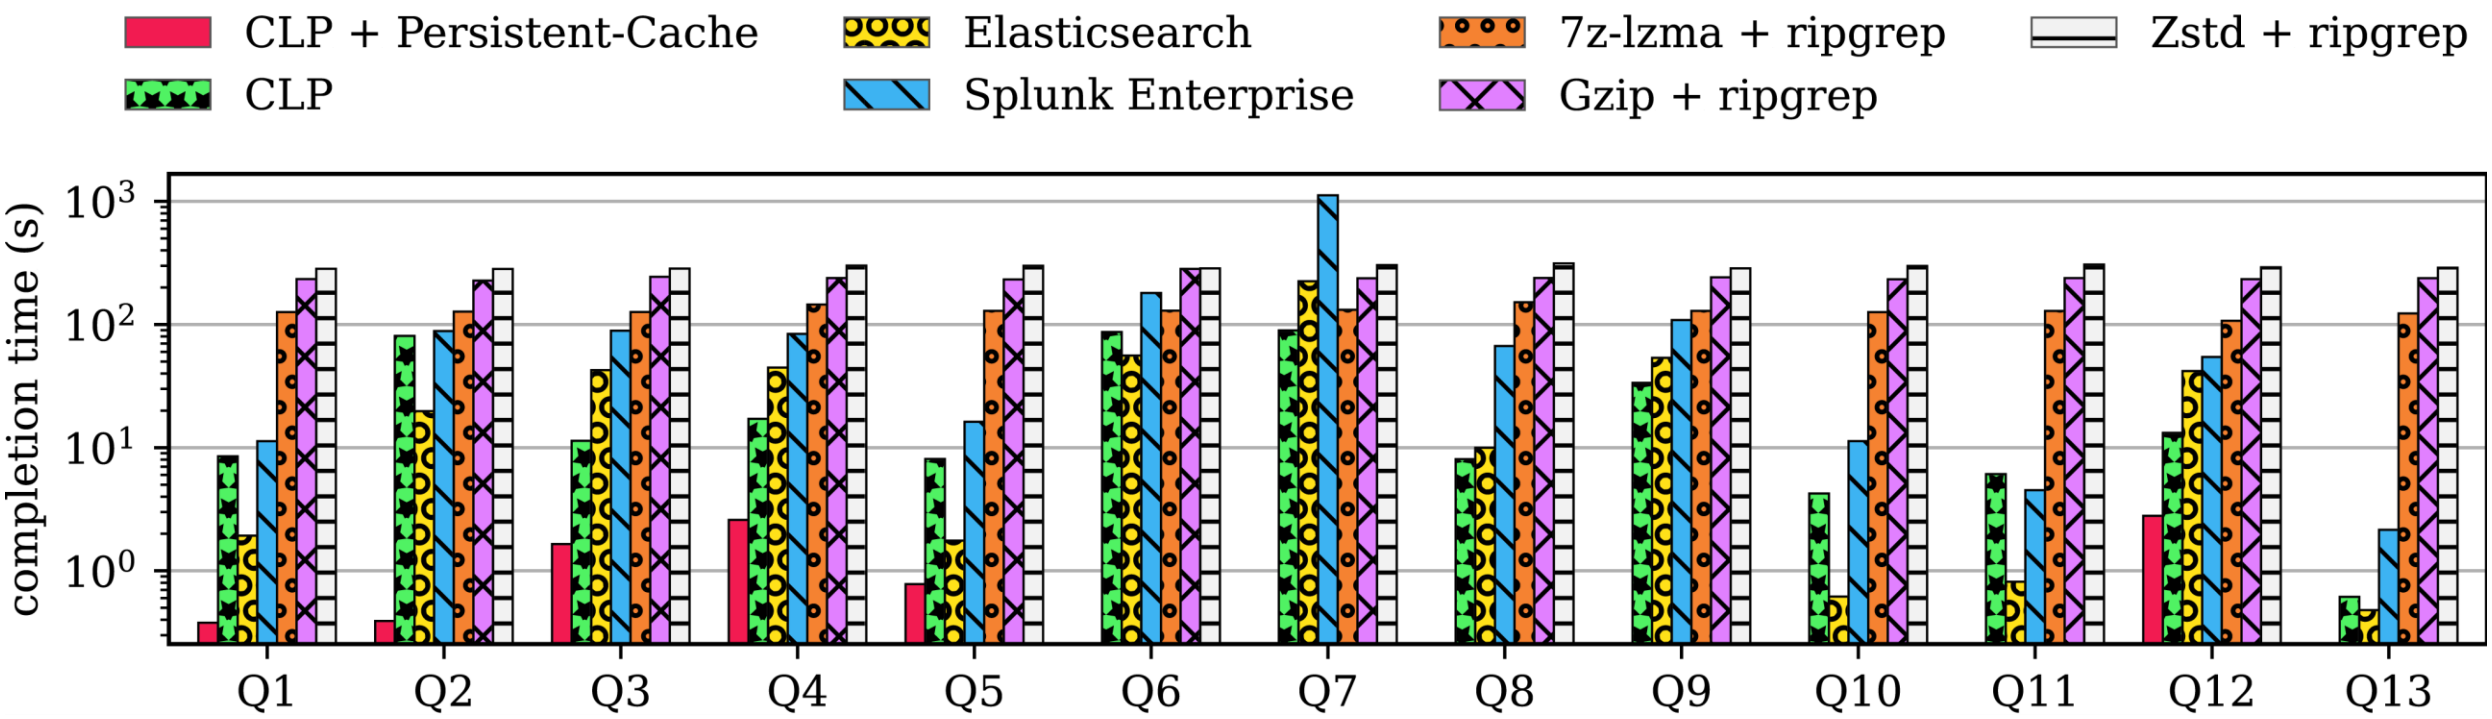

CLP:

- 4.2x faster than Splunk Enterprise
- 1.3x faster than Elasticsearch
- 7.8x faster than ripgrep

Search Performance

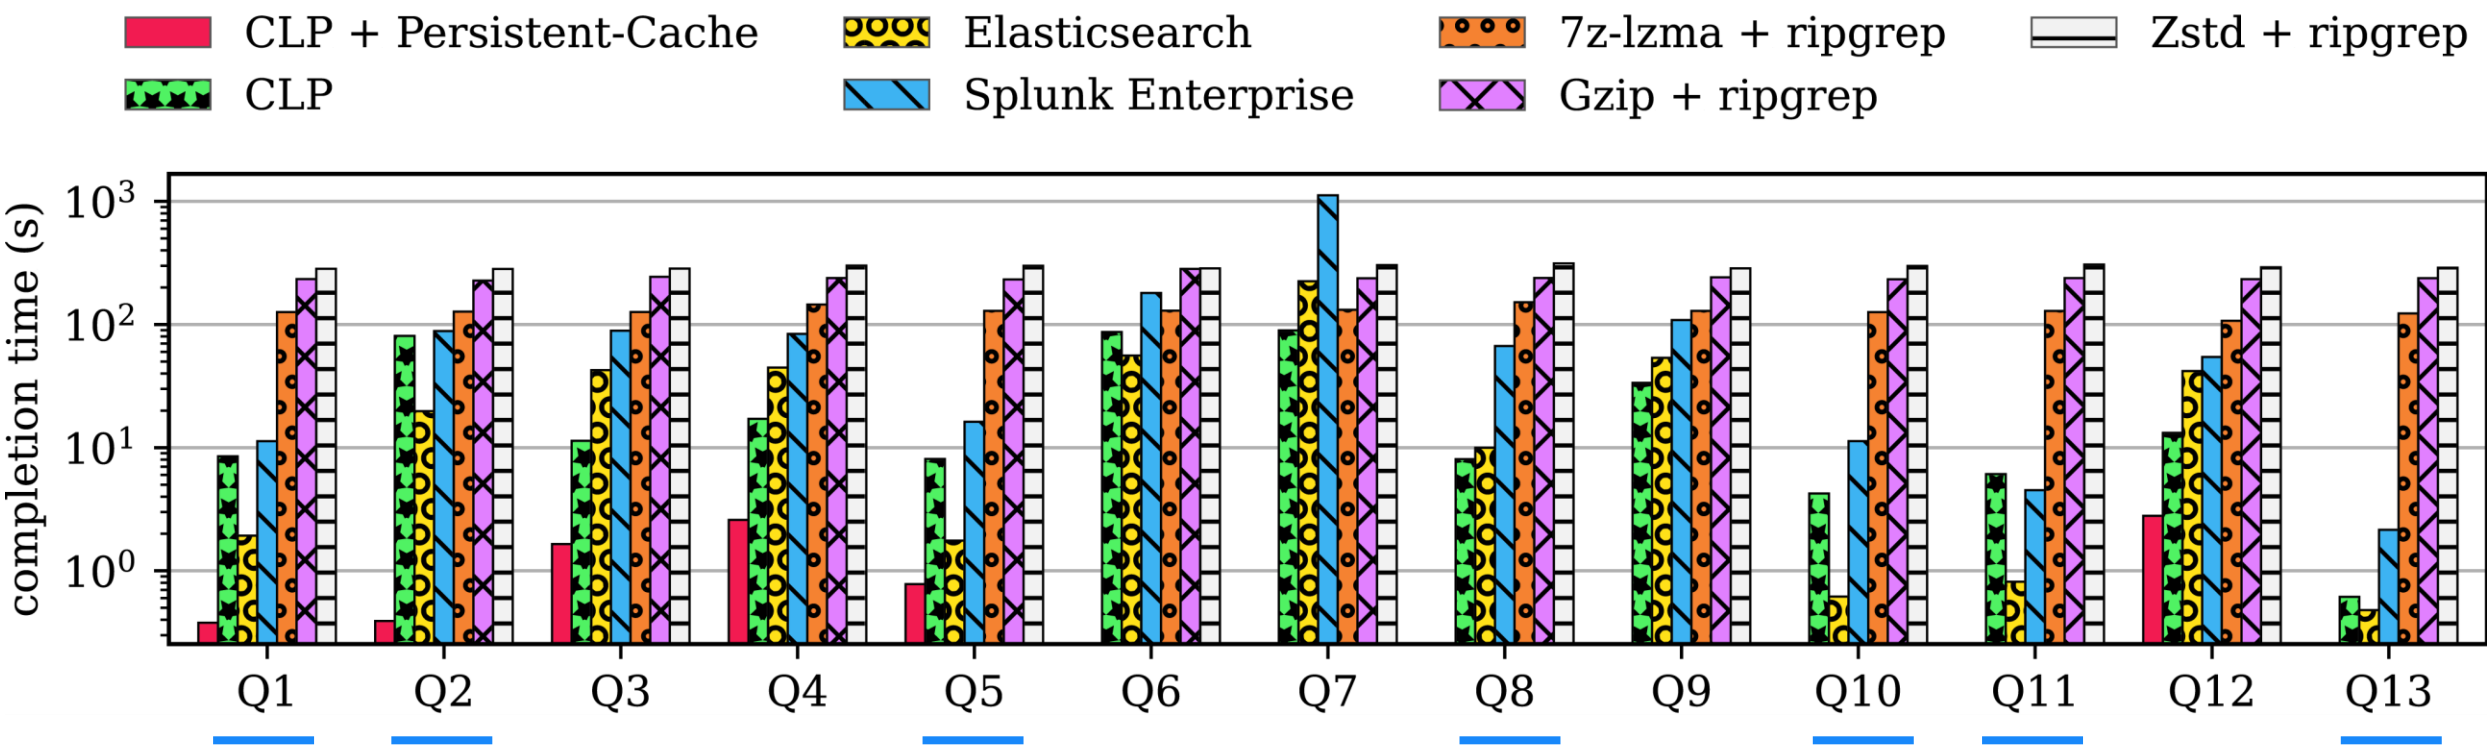

Queries that return few results

CLP:

- 4.2x faster than Splunk Enterprise
- 1.3x faster than Elasticsearch
- 7.8x faster than ripgrep

Search Performance

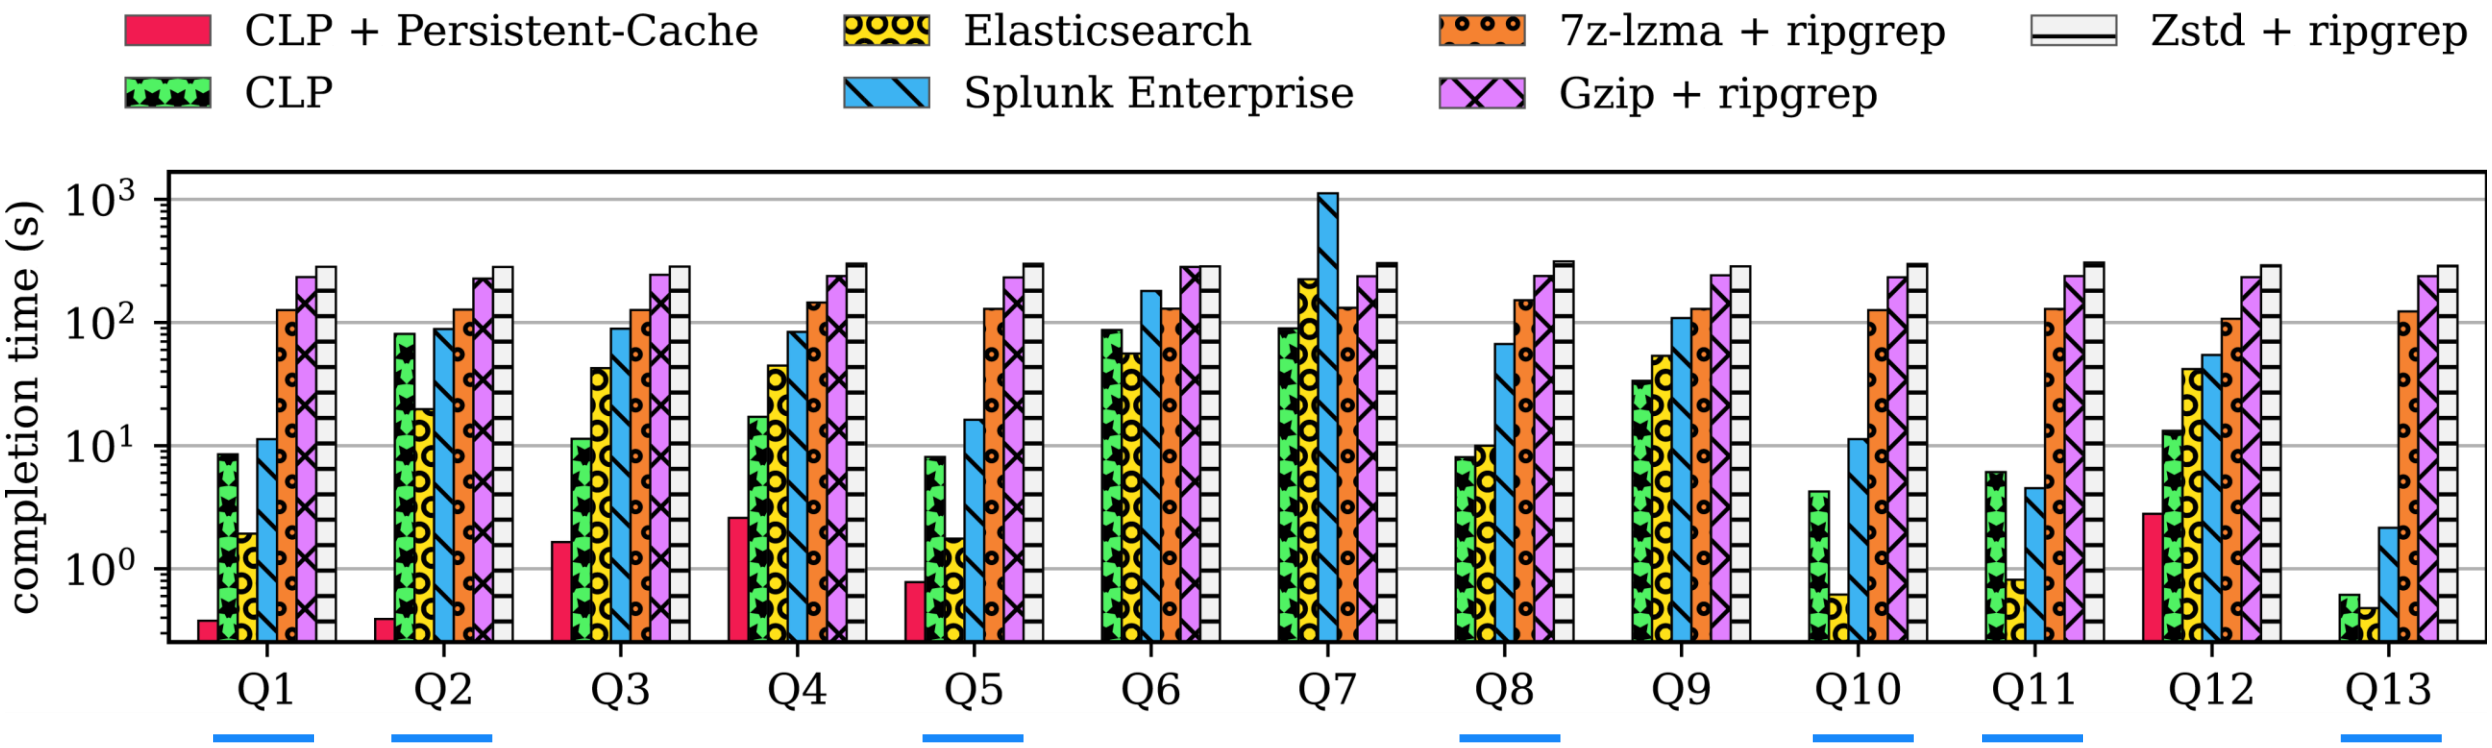

Queries that return few results

CLP + cache:

- 40x faster than Splunk Enterprise
- 17.8x faster than Elasticsearch

Related Work

- Singh and Shivanna [US patent 9,619,478] also aims to deduplicate static text from variable values
  - Does not propose a search algorithm
  - Relies on application source
  - Not entirely lossless
- Tools like Splunk Enterprise and Elasticsearch build text indexes to search logs
- Succinct [Agarwal *et al*. NSDI '15] proposed a method for compressing indexes
  - But any index still carries overhead whereas CLP deduplicates the original data
- Scalyr uses optimizations to search uncompressed logs at 1.25 GB/s
  - CLP works on compressed data with up to 420 GB/s throughput
- Grafana Loki only indexes labels
  - Index still adds overhead
  - Reduced index size but search limited to labels

Conclusion

- Achieves unparalleled log compression
- Allows search without decompression
- Combines archiving & log search

- Open-sourced!
- Try it out at yscope.com!
- CLP is just the beginning…
  - e.g., Stitch [Zhao et al. OSDI '16], Log20 [Zhao et al. SOSP '17]

YScope

We Automate Debugging

Want to get in touch?

kirk.rodrigues@yscope.com

info@yscope.com