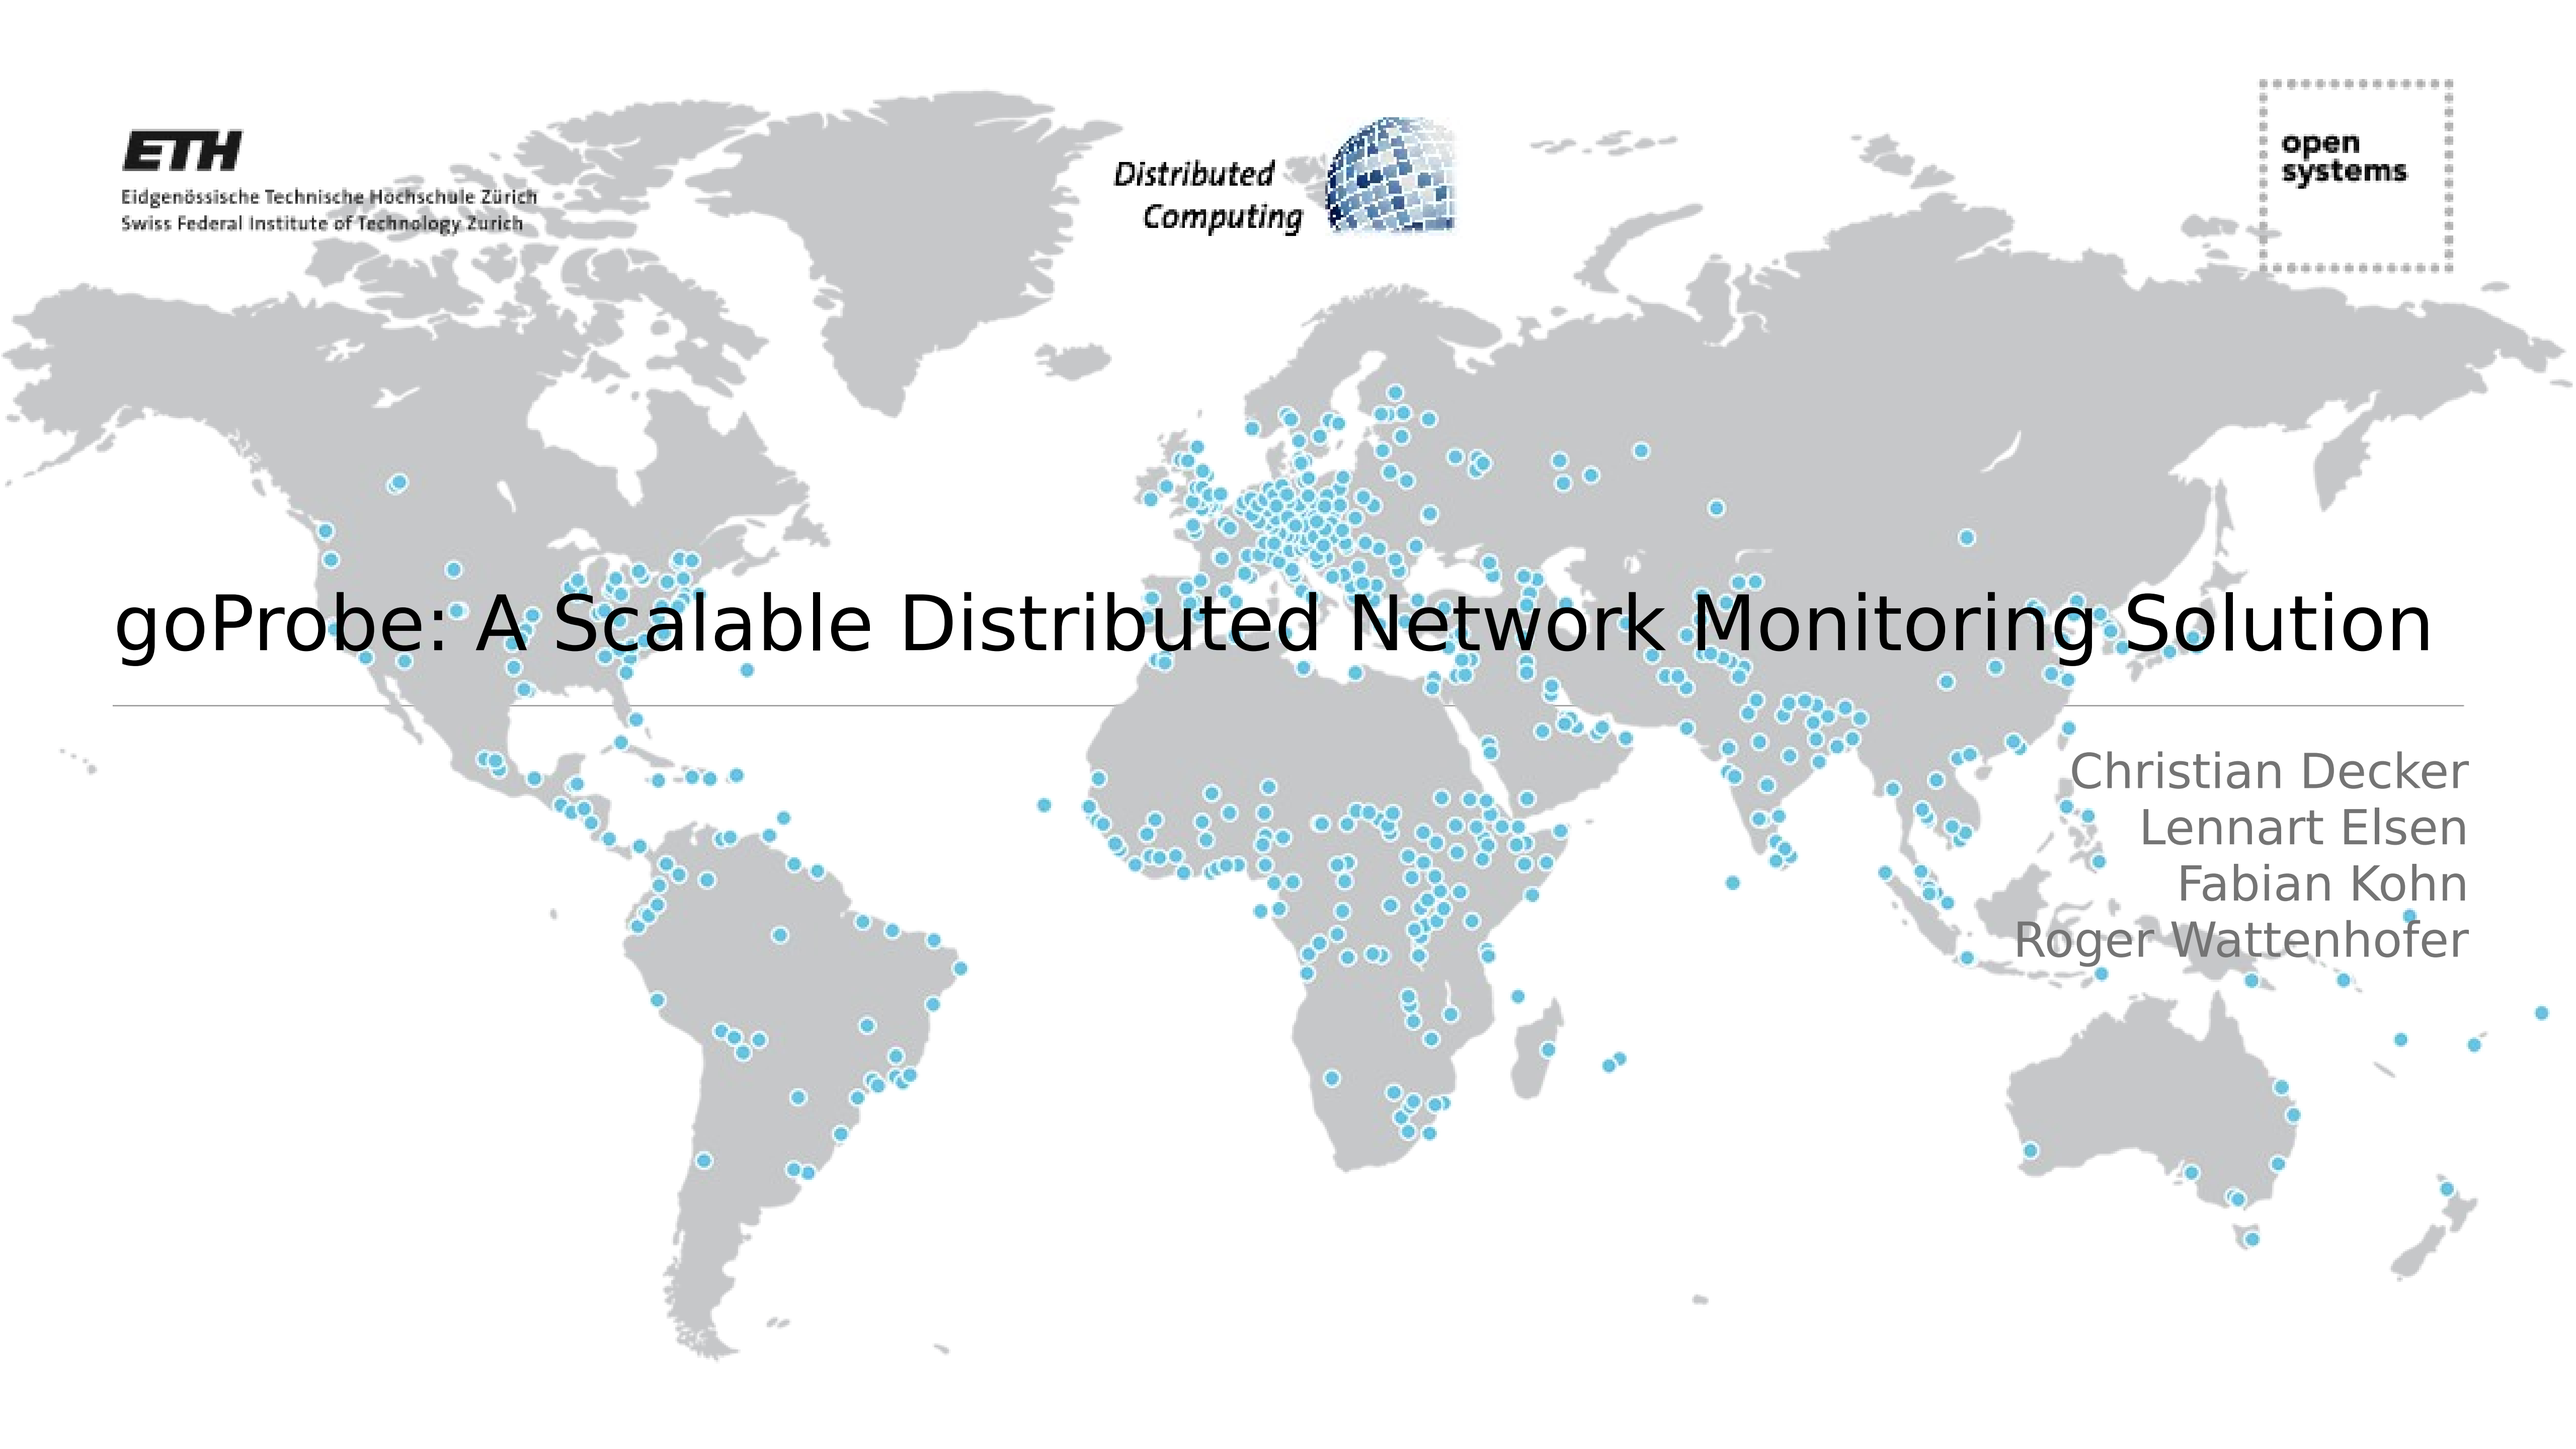

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

open systems

goProbe: A Scalable Distributed Network Monitoring Solution

Christian Decker
Lennart Elsen
Fabian Kohn
Roger Wattenhofer


BEST SCORE:
4,205
PLAY
SHOP
SOCIAL
COSTUMES
GOALS
OPTIONS

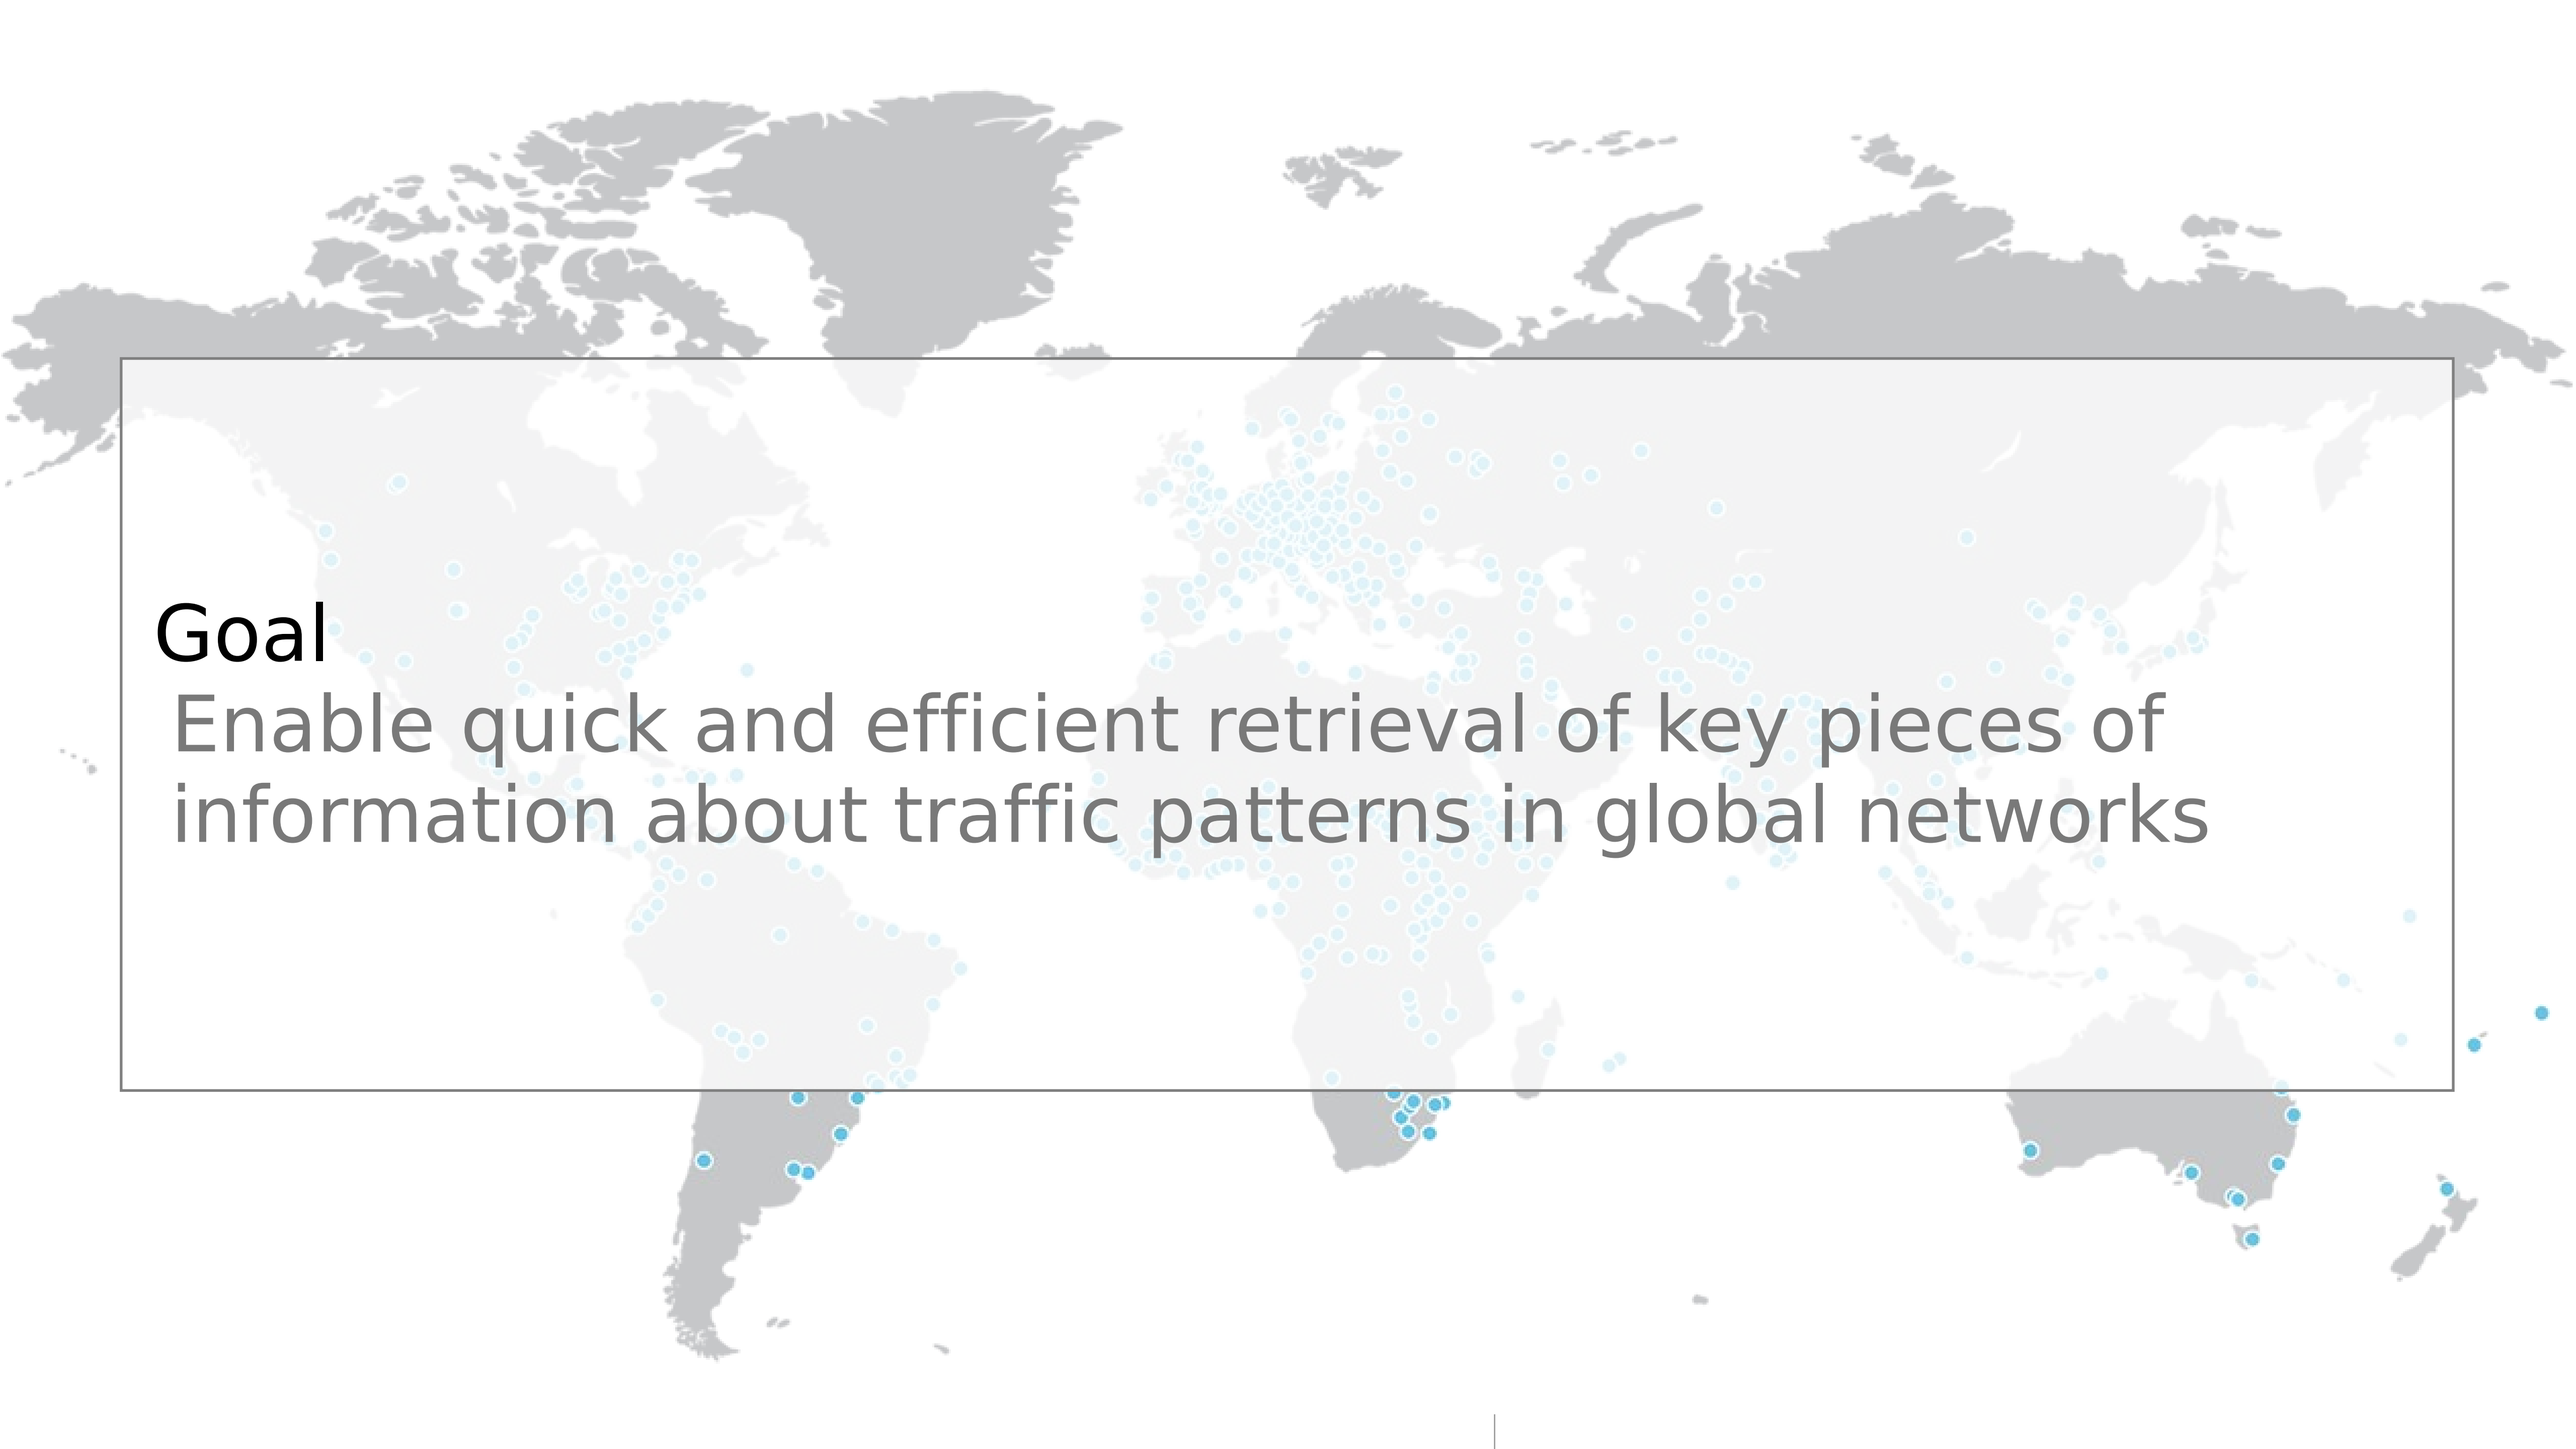

Goal

Enable quick and efficient retrieval of key pieces of information about traffic patterns in global networks

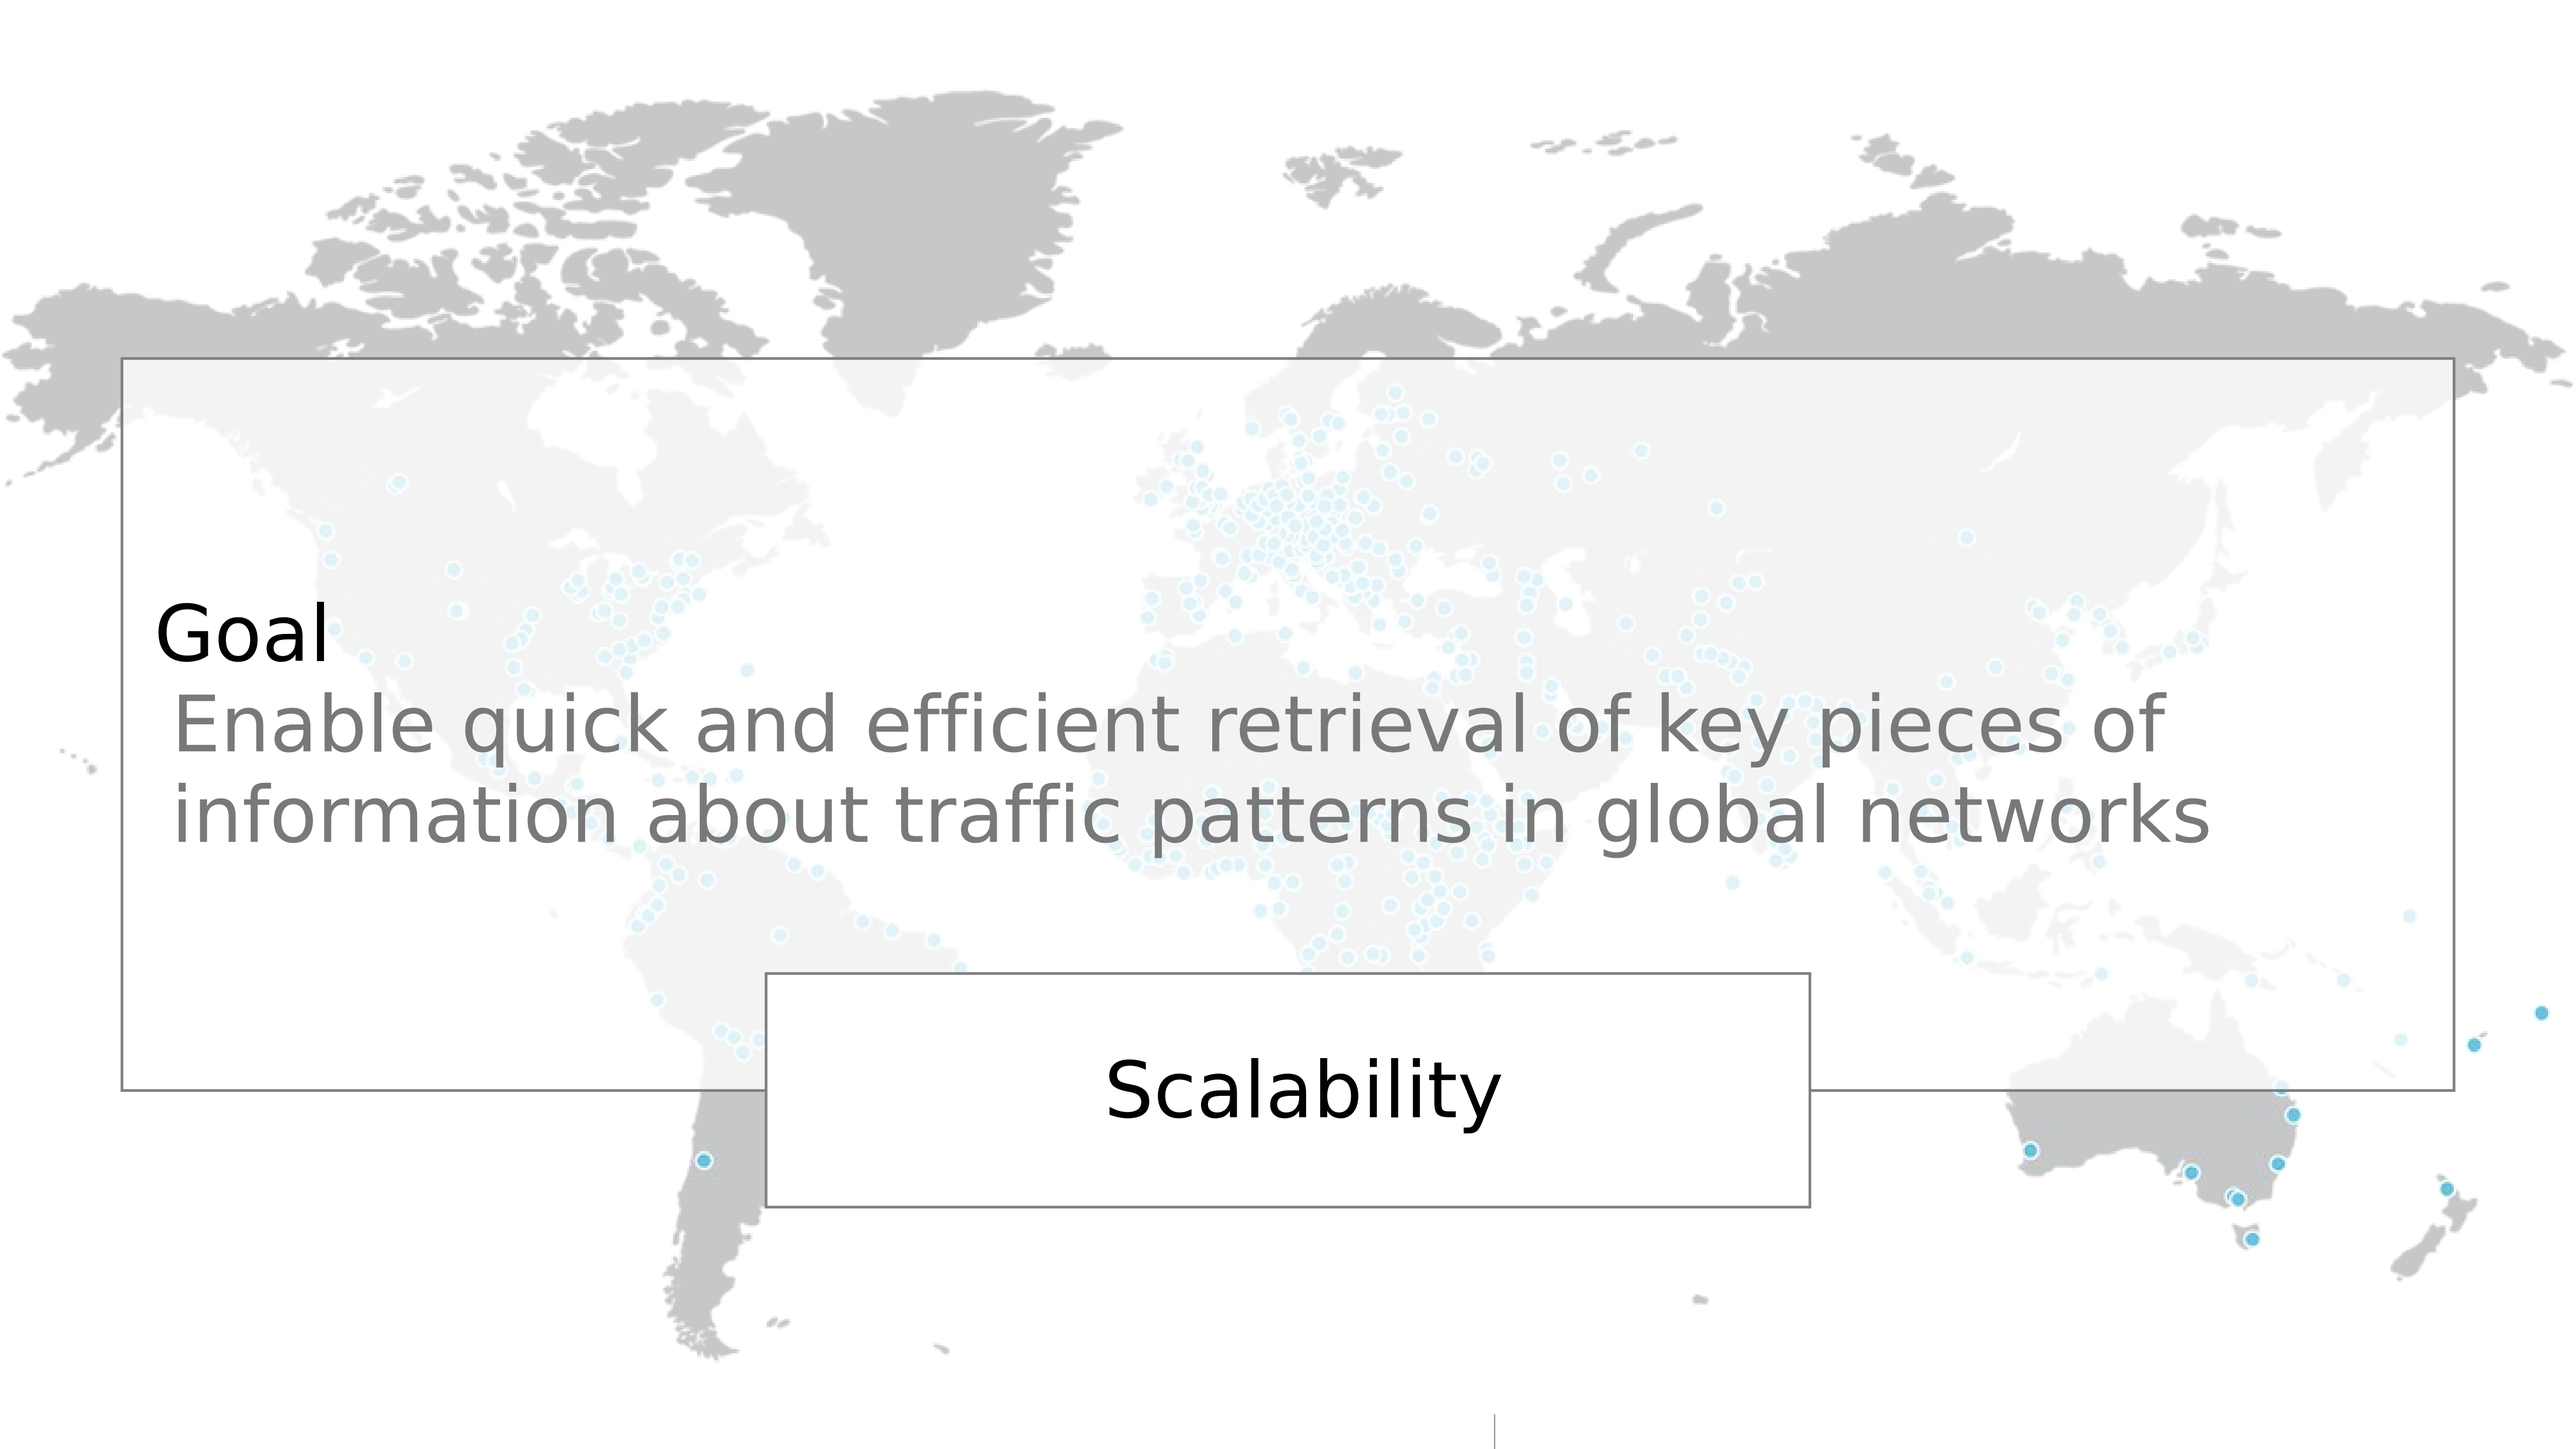

Goal

Enable quick and efficient retrieval of key pieces of information about traffic patterns in global networks

Scalability

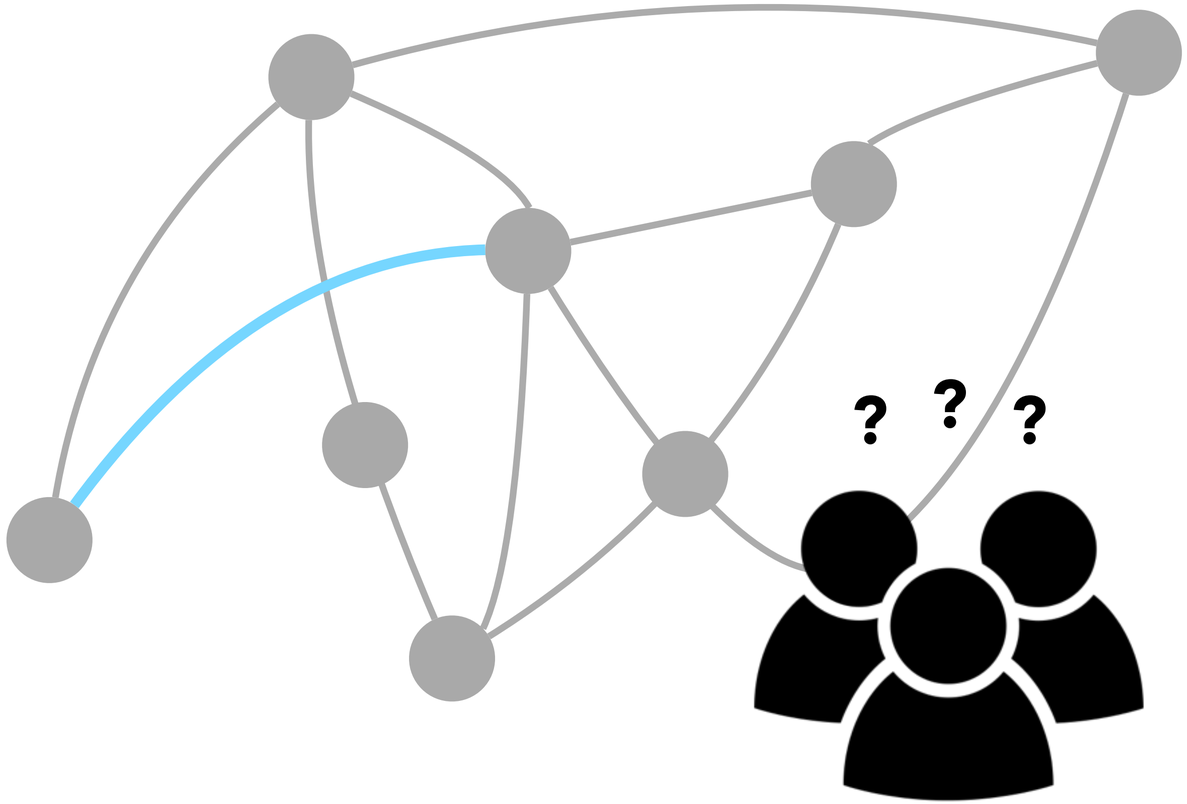

Debugging/Operations

Reporting

Acquisition of Traffic Data

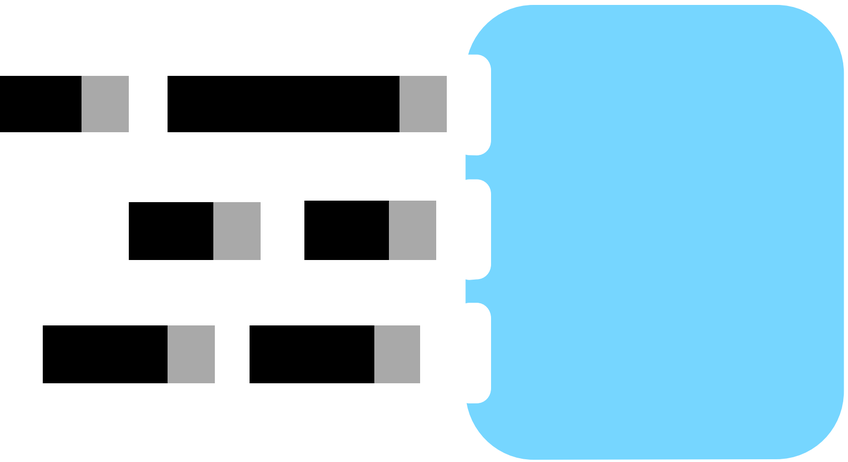

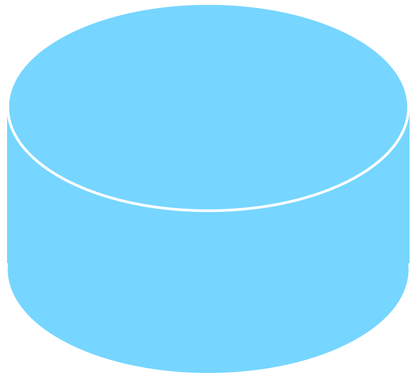

Packet Capture

Storage

Acquisition of Traffic Data

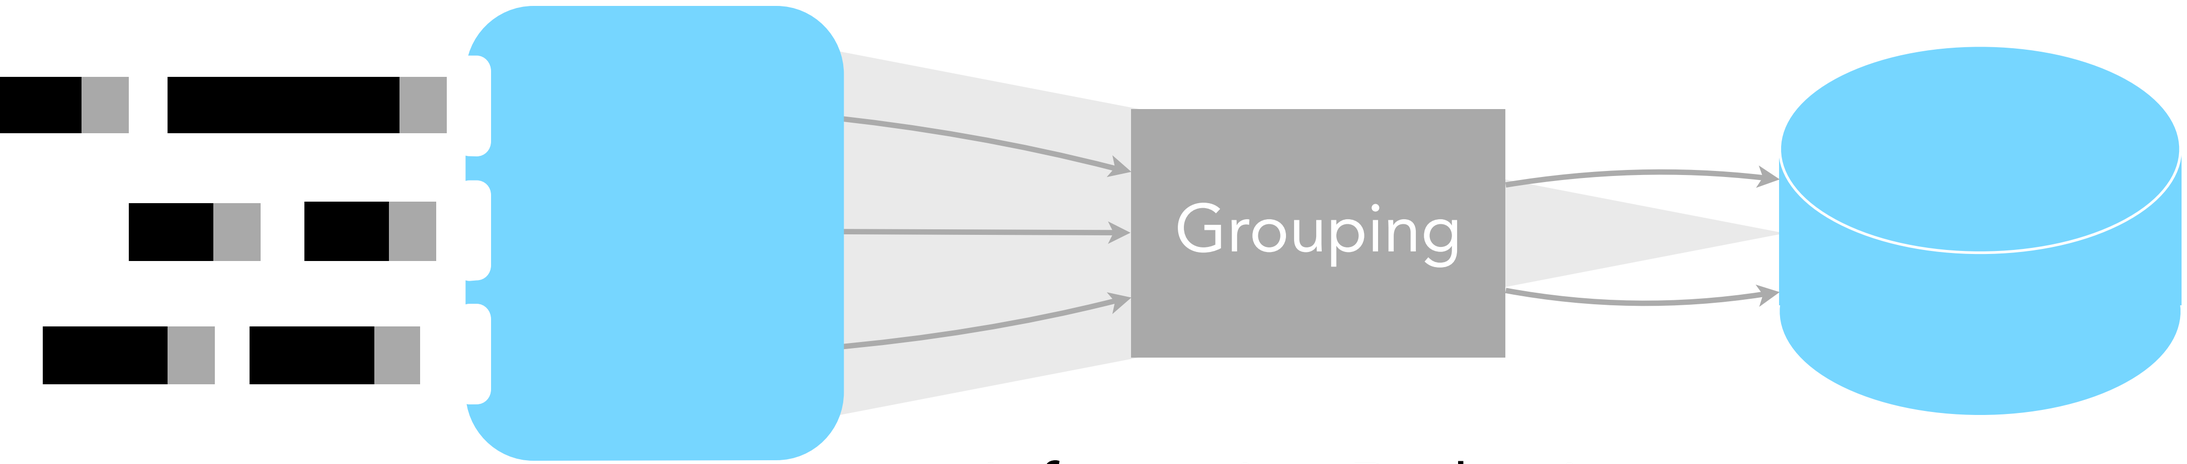

Information Reduction

Packet Capture

Storage

NetFlow

Packet aggregation by set of **shared** attributes

Network packet headers & packet counters

Expiry time

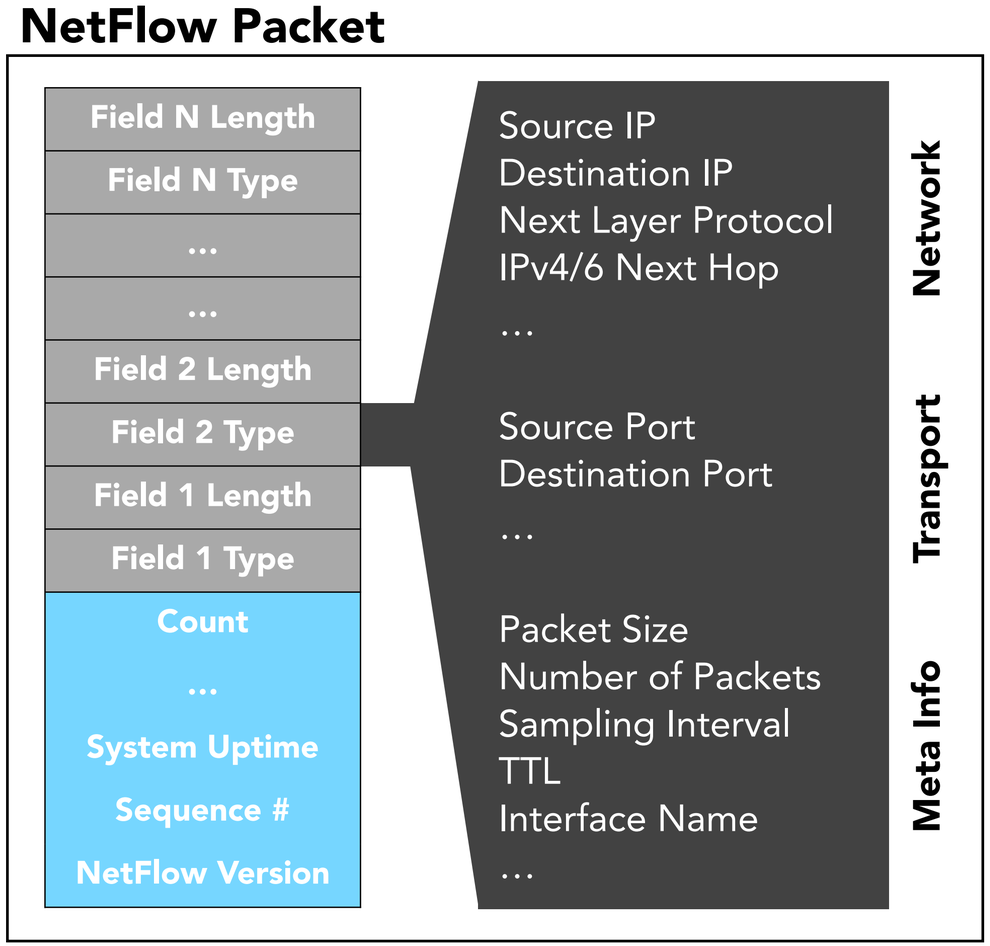

NetFlow

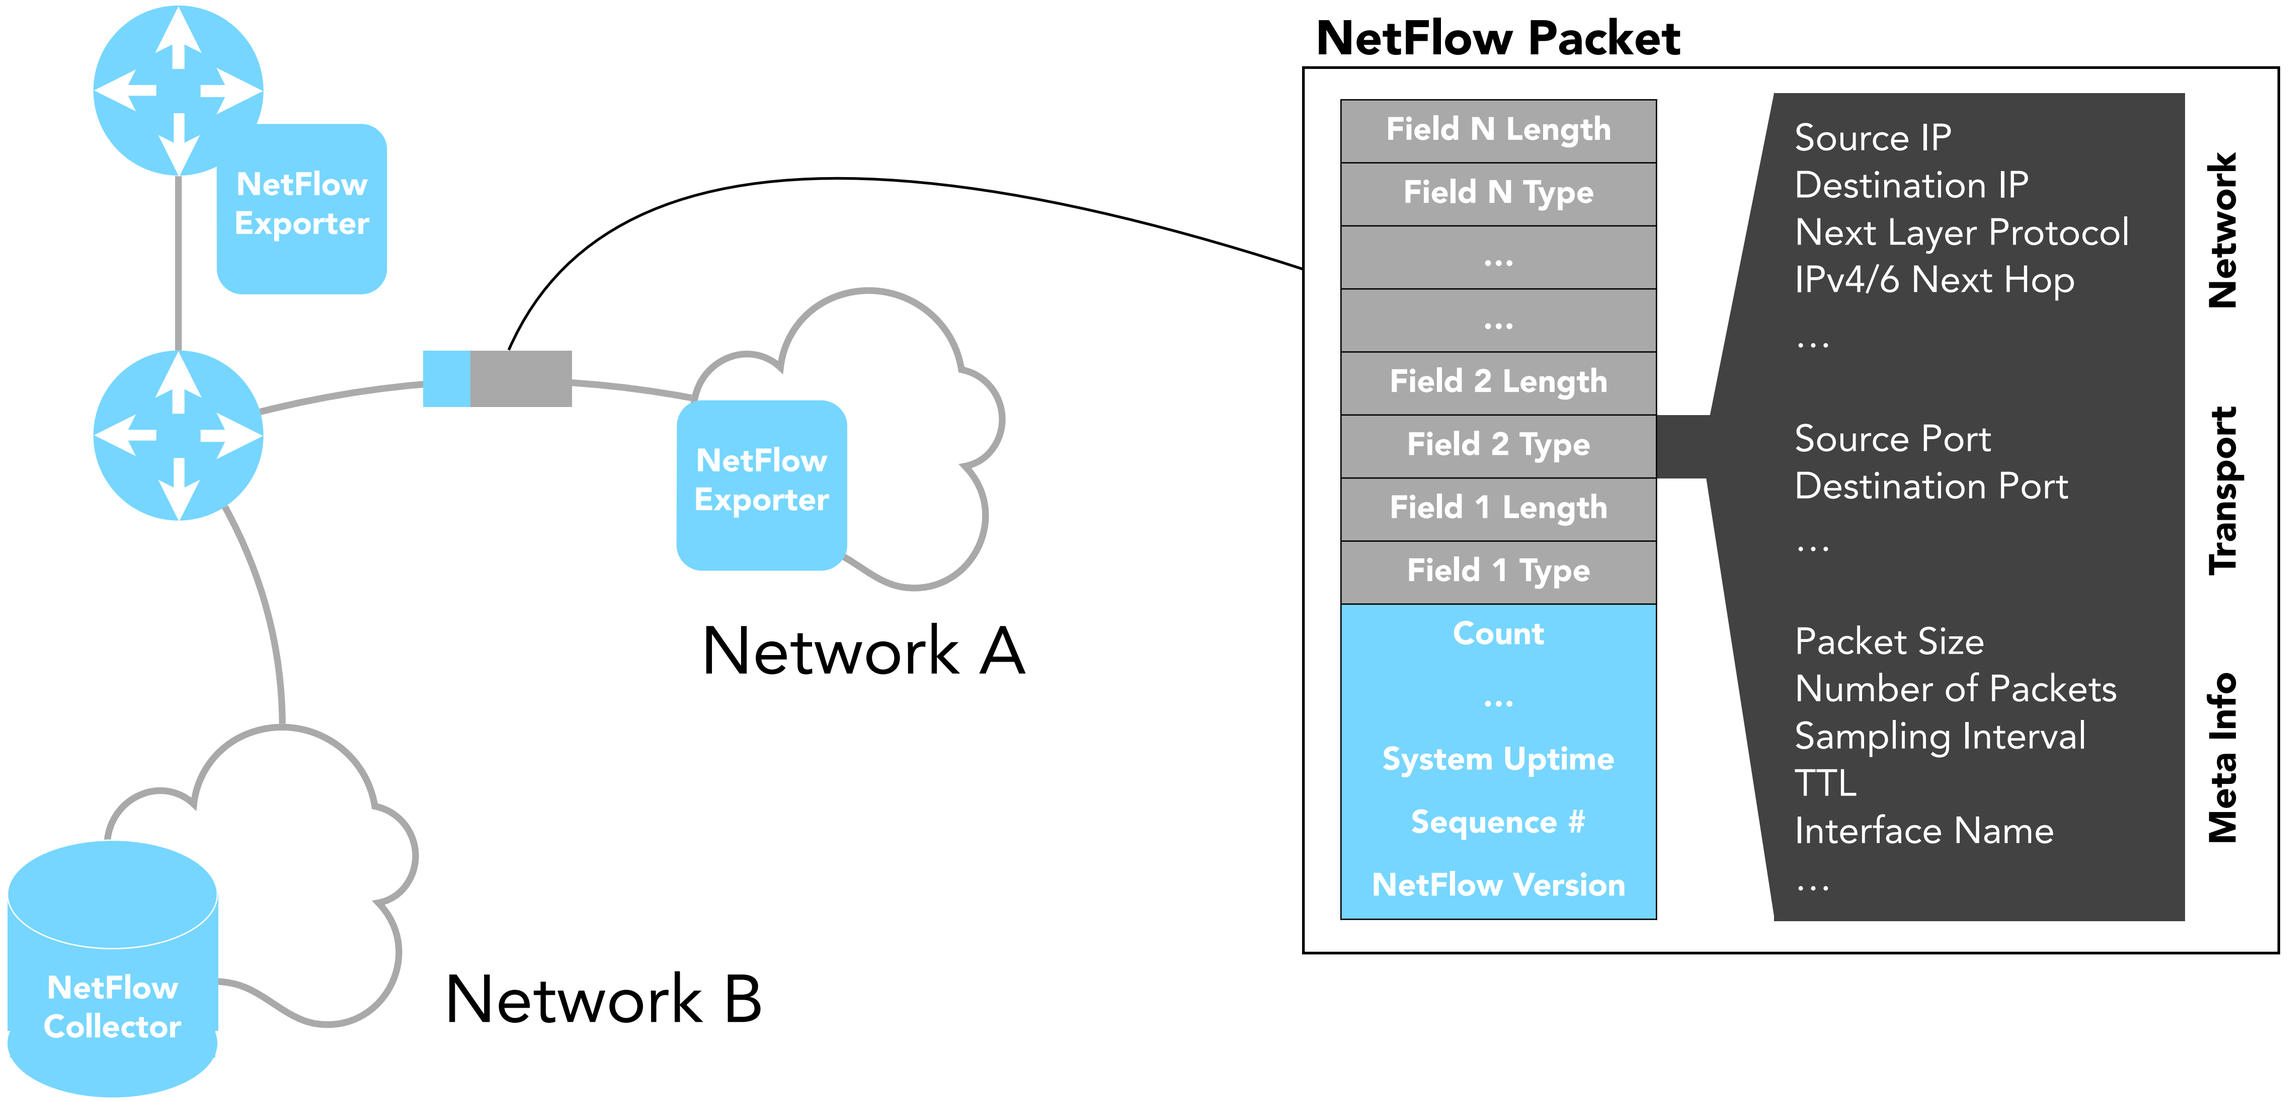
NetFlow Exporter
NetFlow Exporter
Network A
NetFlow Collector
Network B
NetFlow Packet
Field N Length
Field N Type
...
...
Field 2 Length
Field 2 Type
Field 1 Length
Field 1 Type
Count
...
System Uptime
Sequence #
NetFlow Version
Source IP
Destination IP
Next Layer Protocol
IPv4/6 Next Hop
...
Source Port
Destination Port
...
Packet Size
Number of Packets
Sampling Interval
TTL
Interface Name
...
Network
Transport
Meta Info

NetFlow

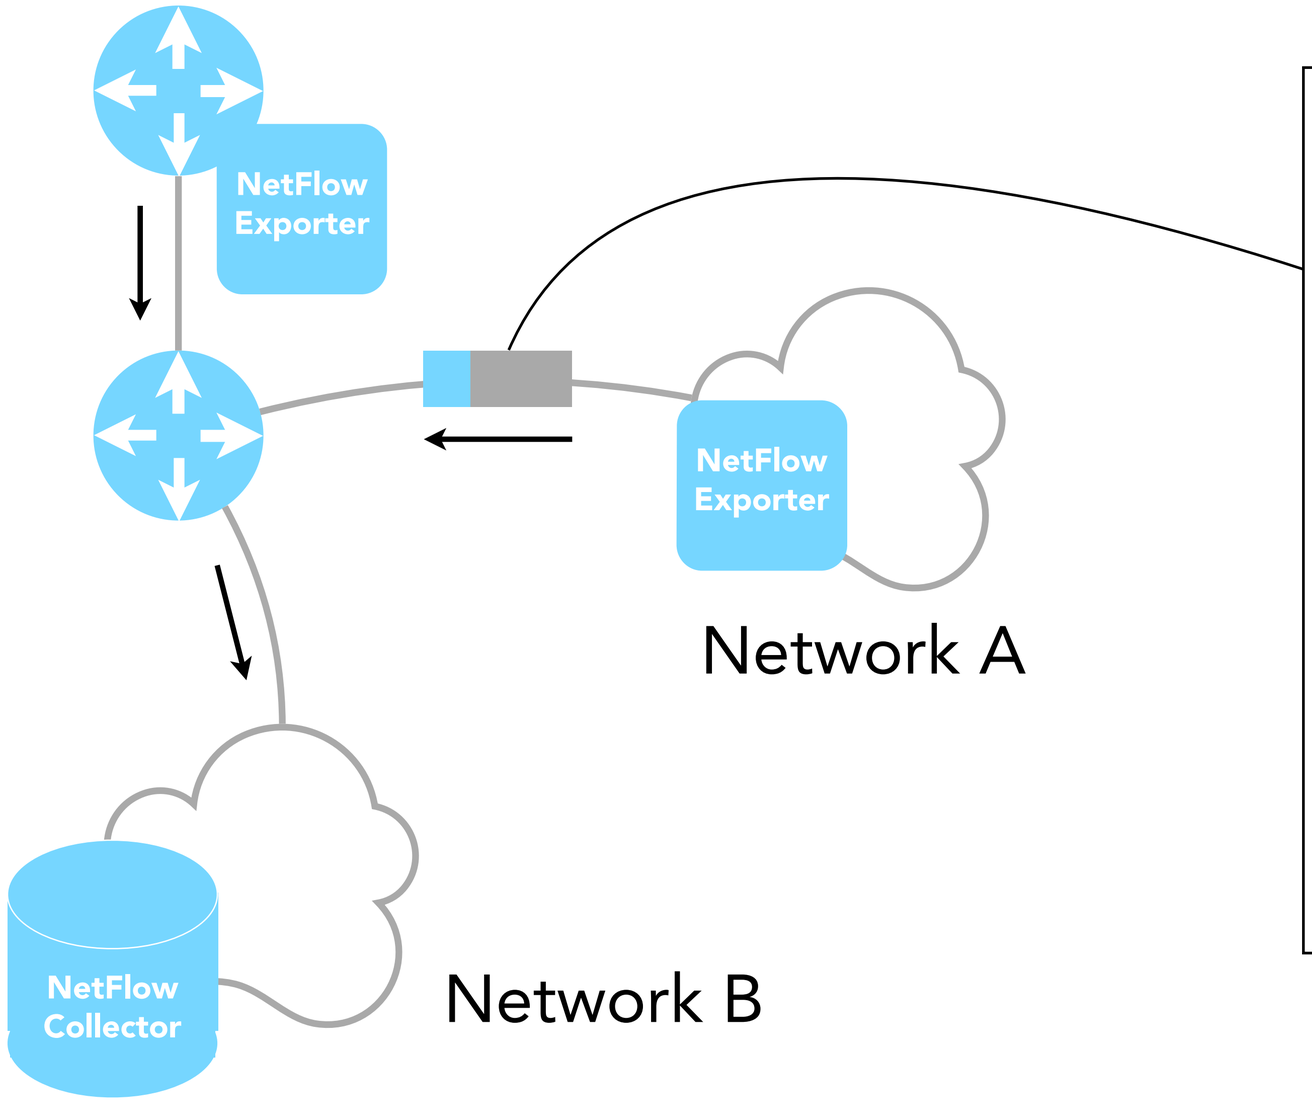

NetFlow Packet

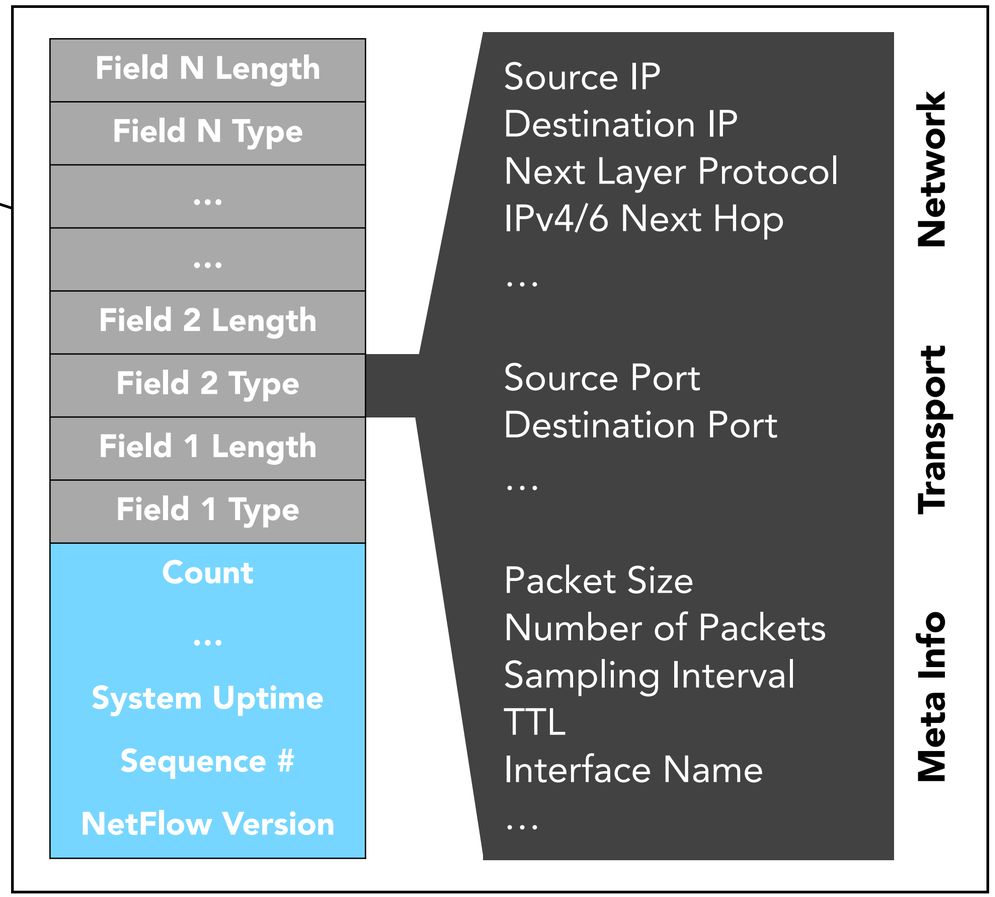

Current Network Monitoring System

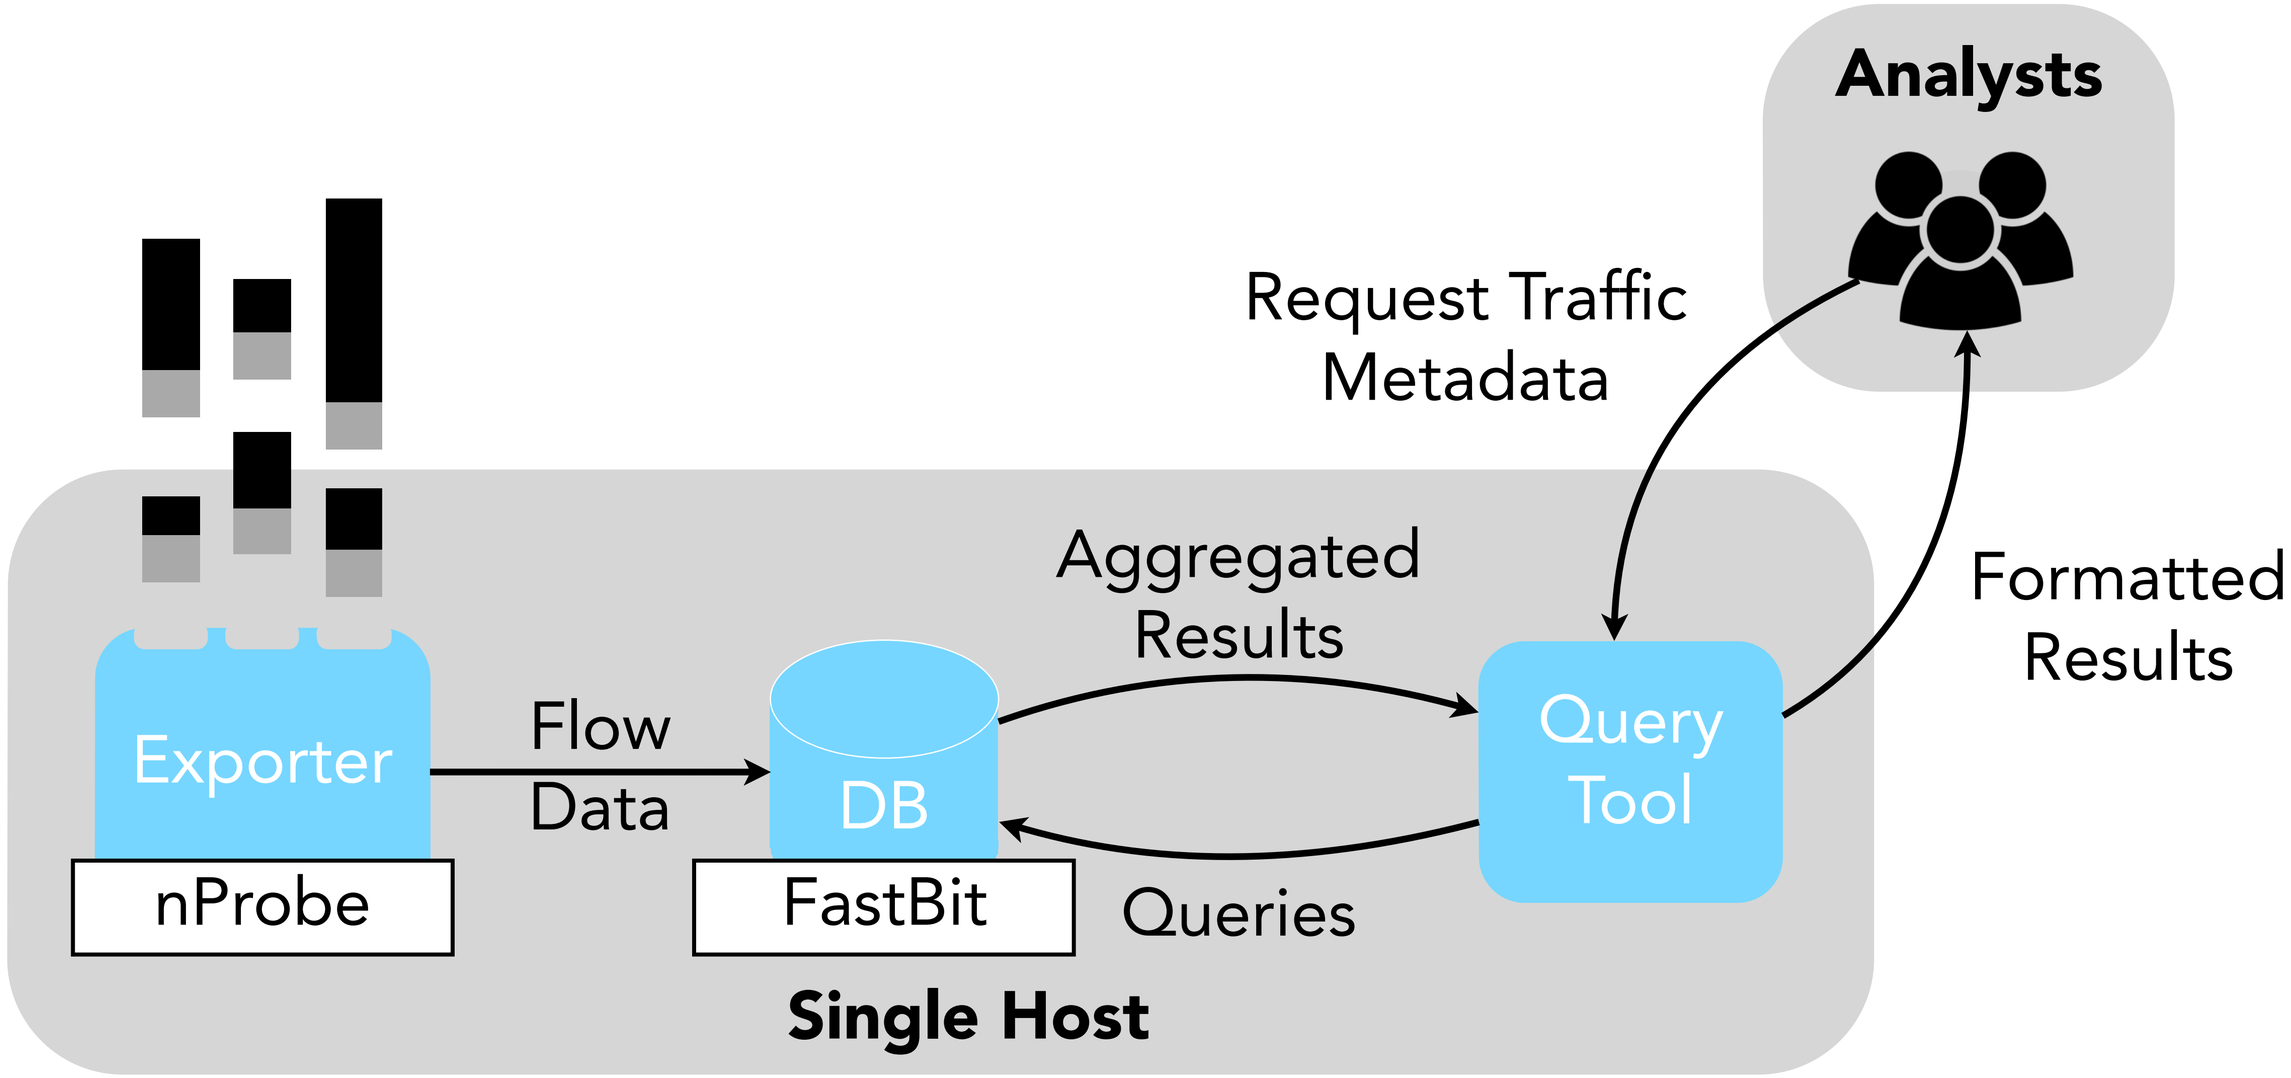

Challenges

Capturing Process

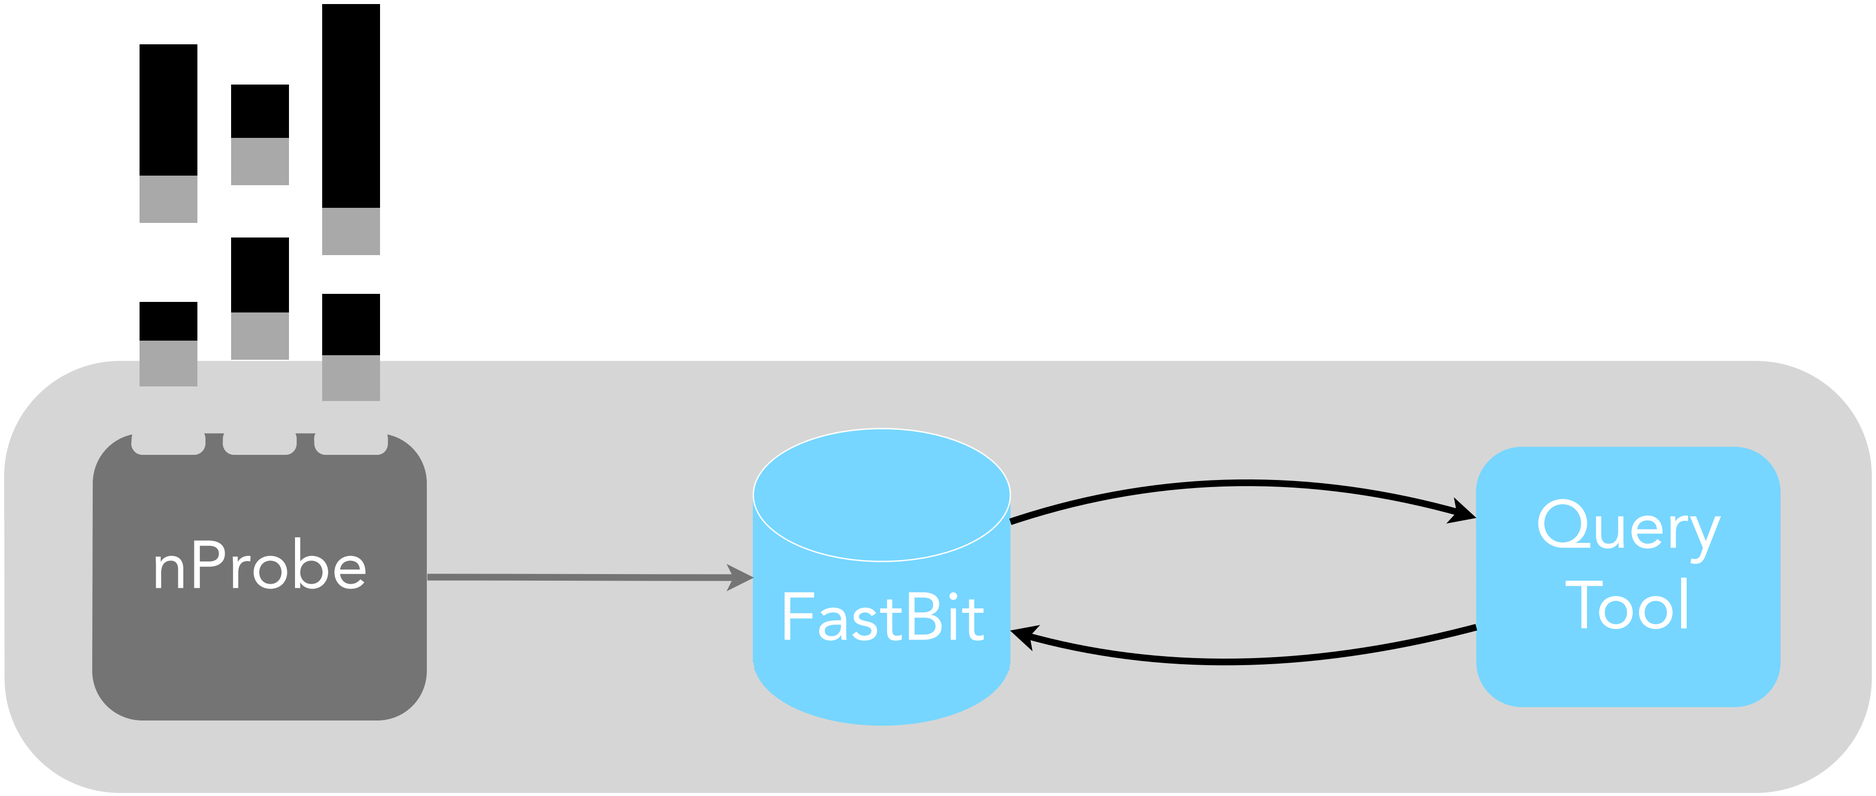

Challenges

Capturing Process

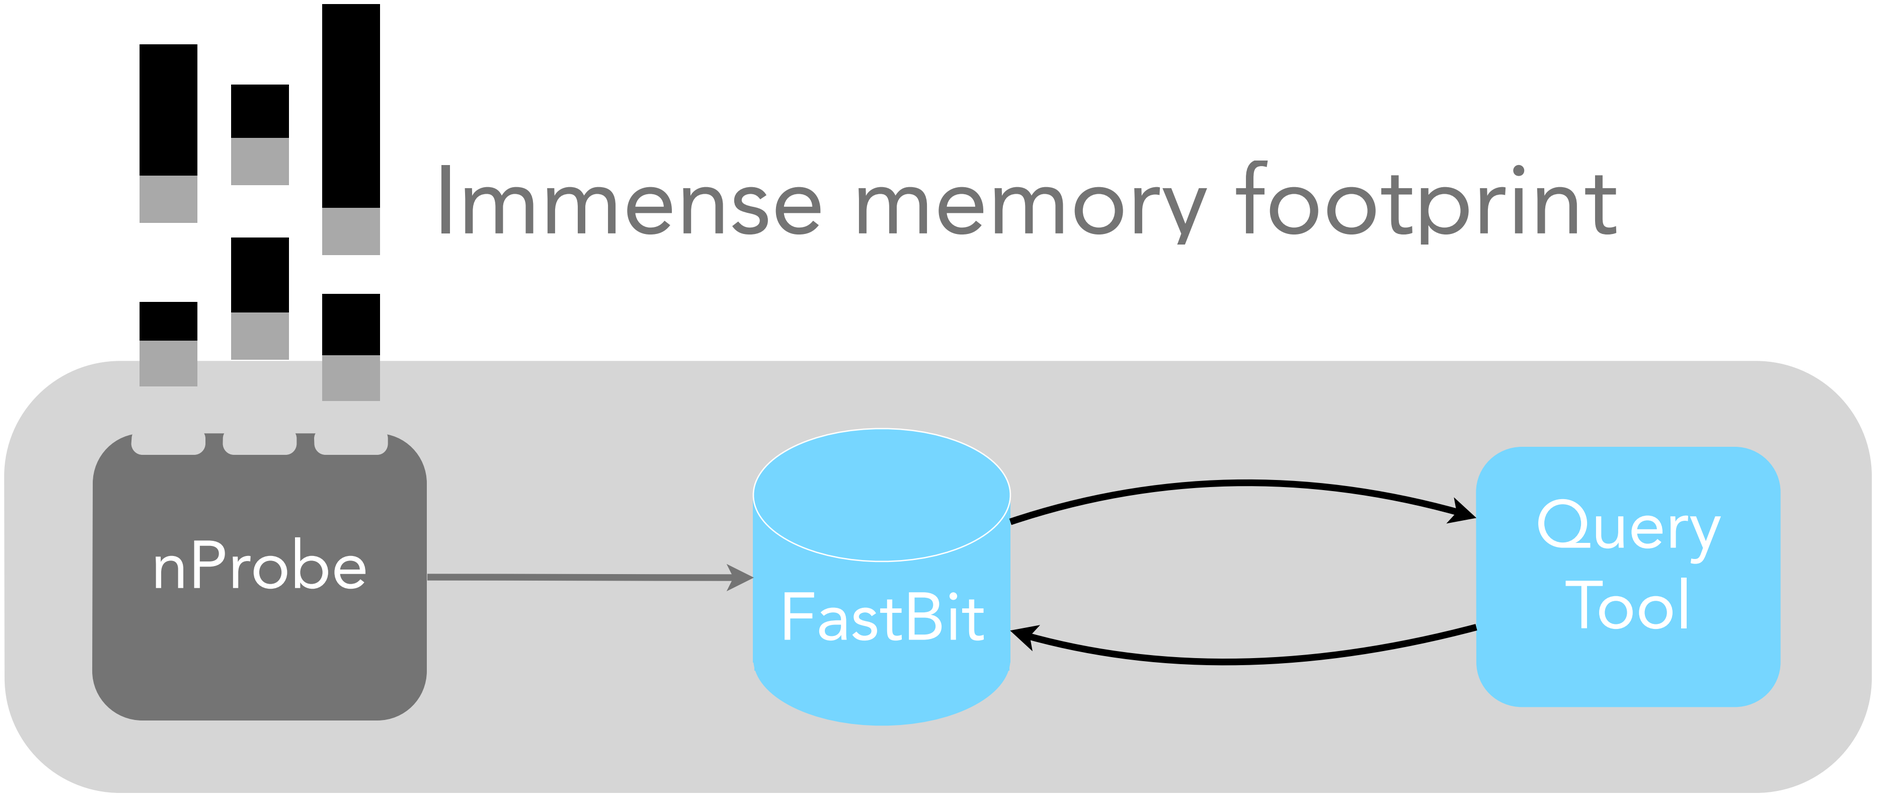

Challenges

Capturing Process

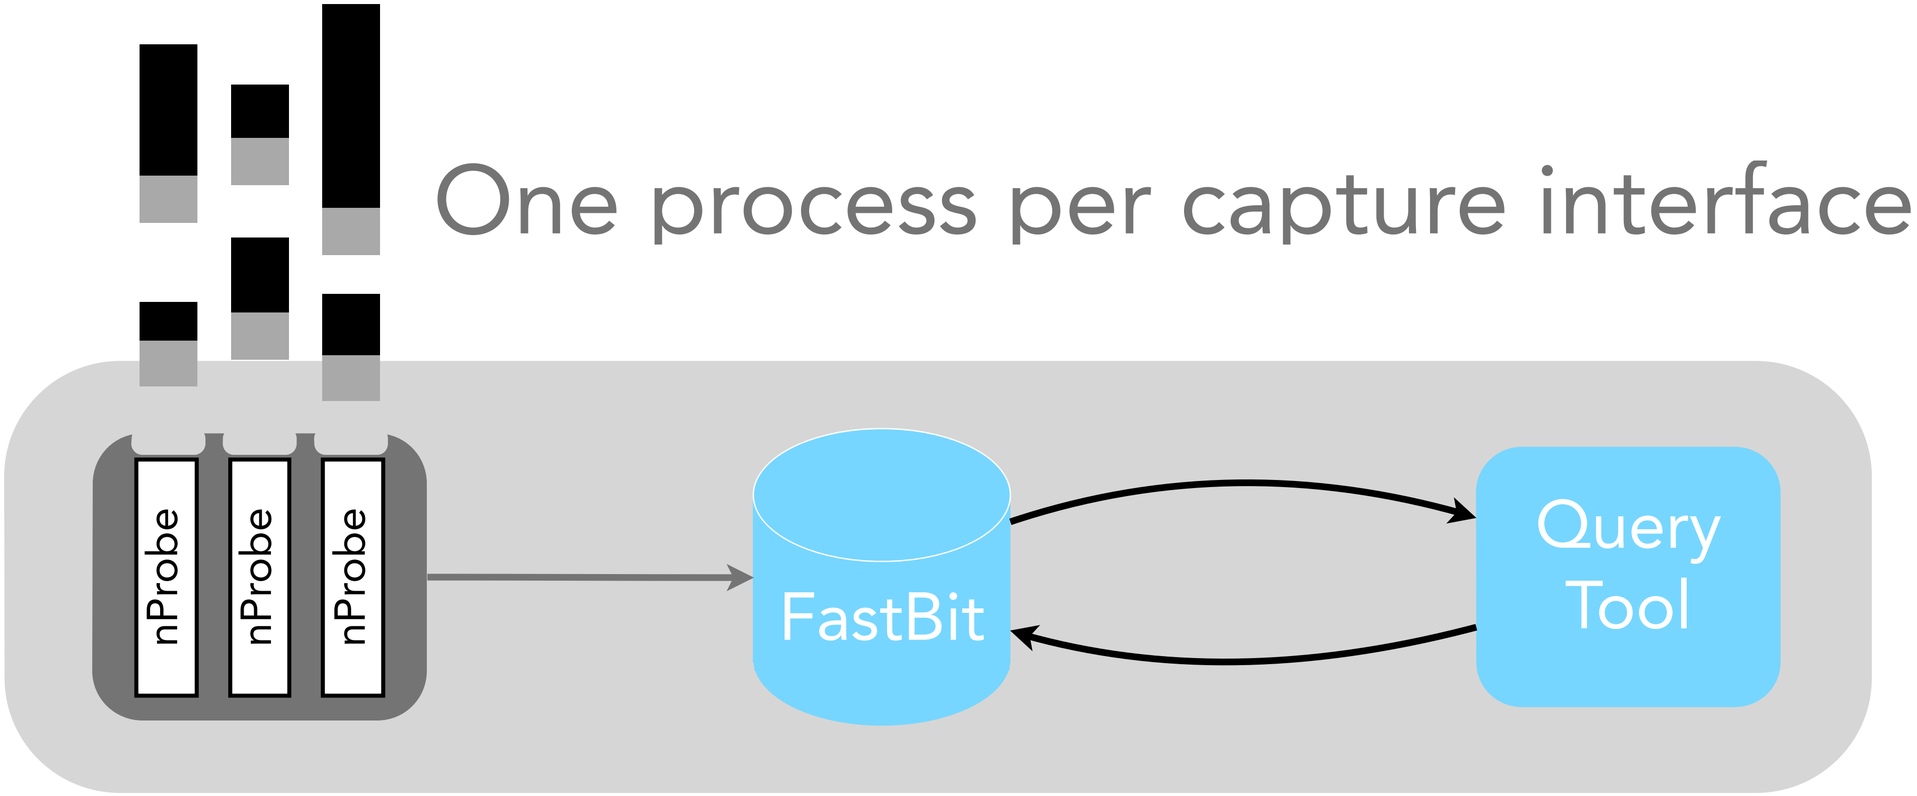

Challenges

Storage Backend

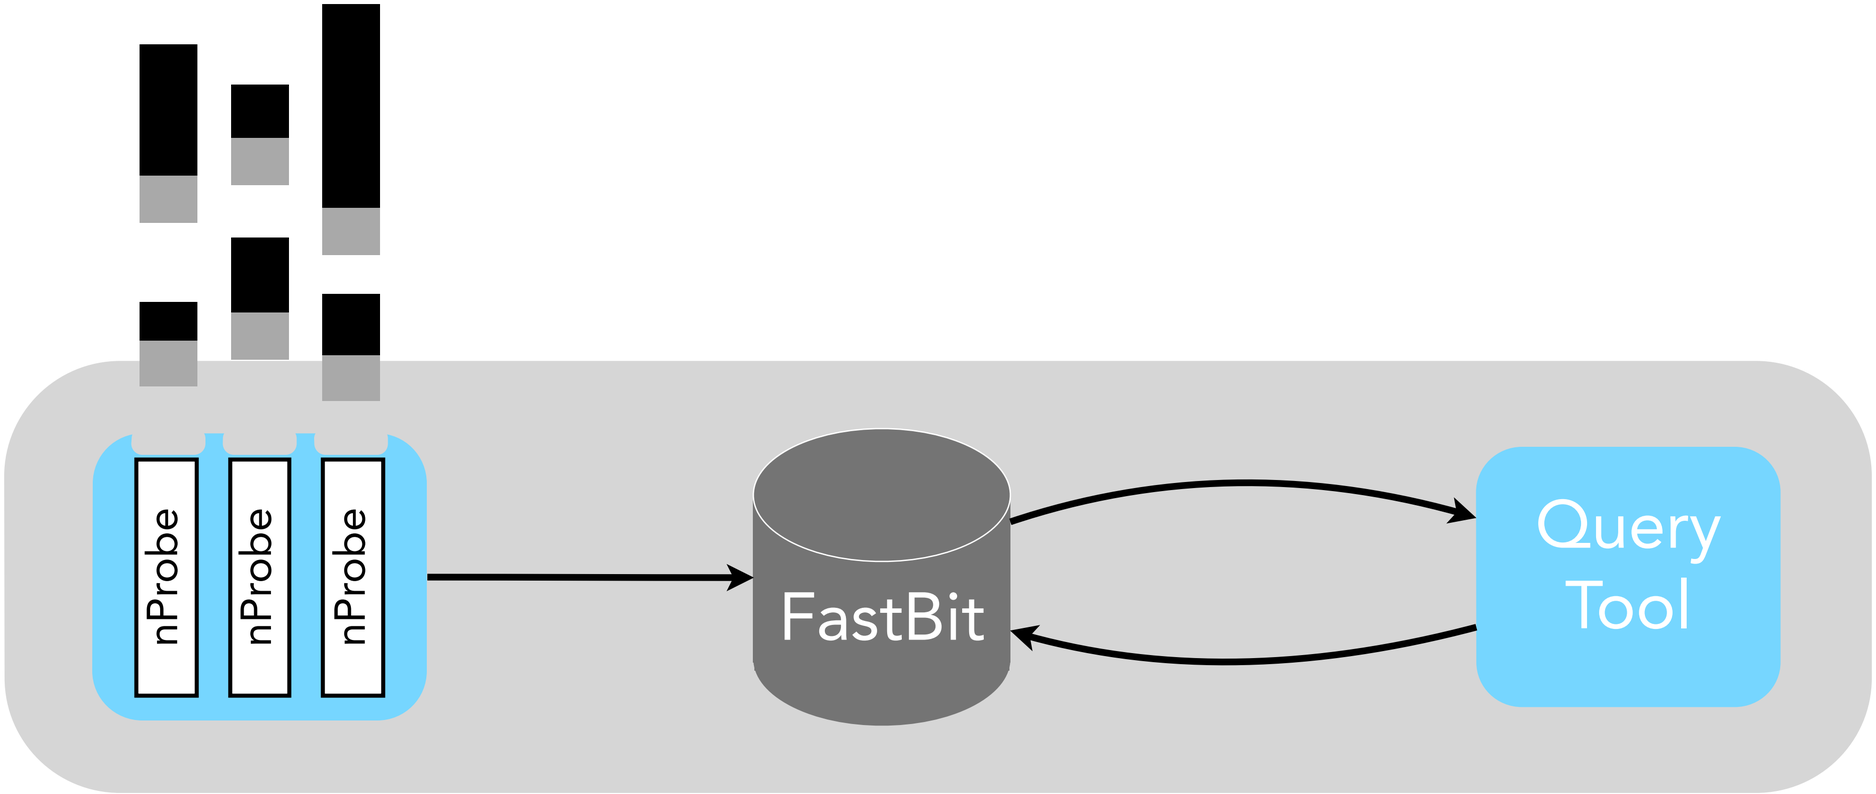

Challenges

Storage Backend

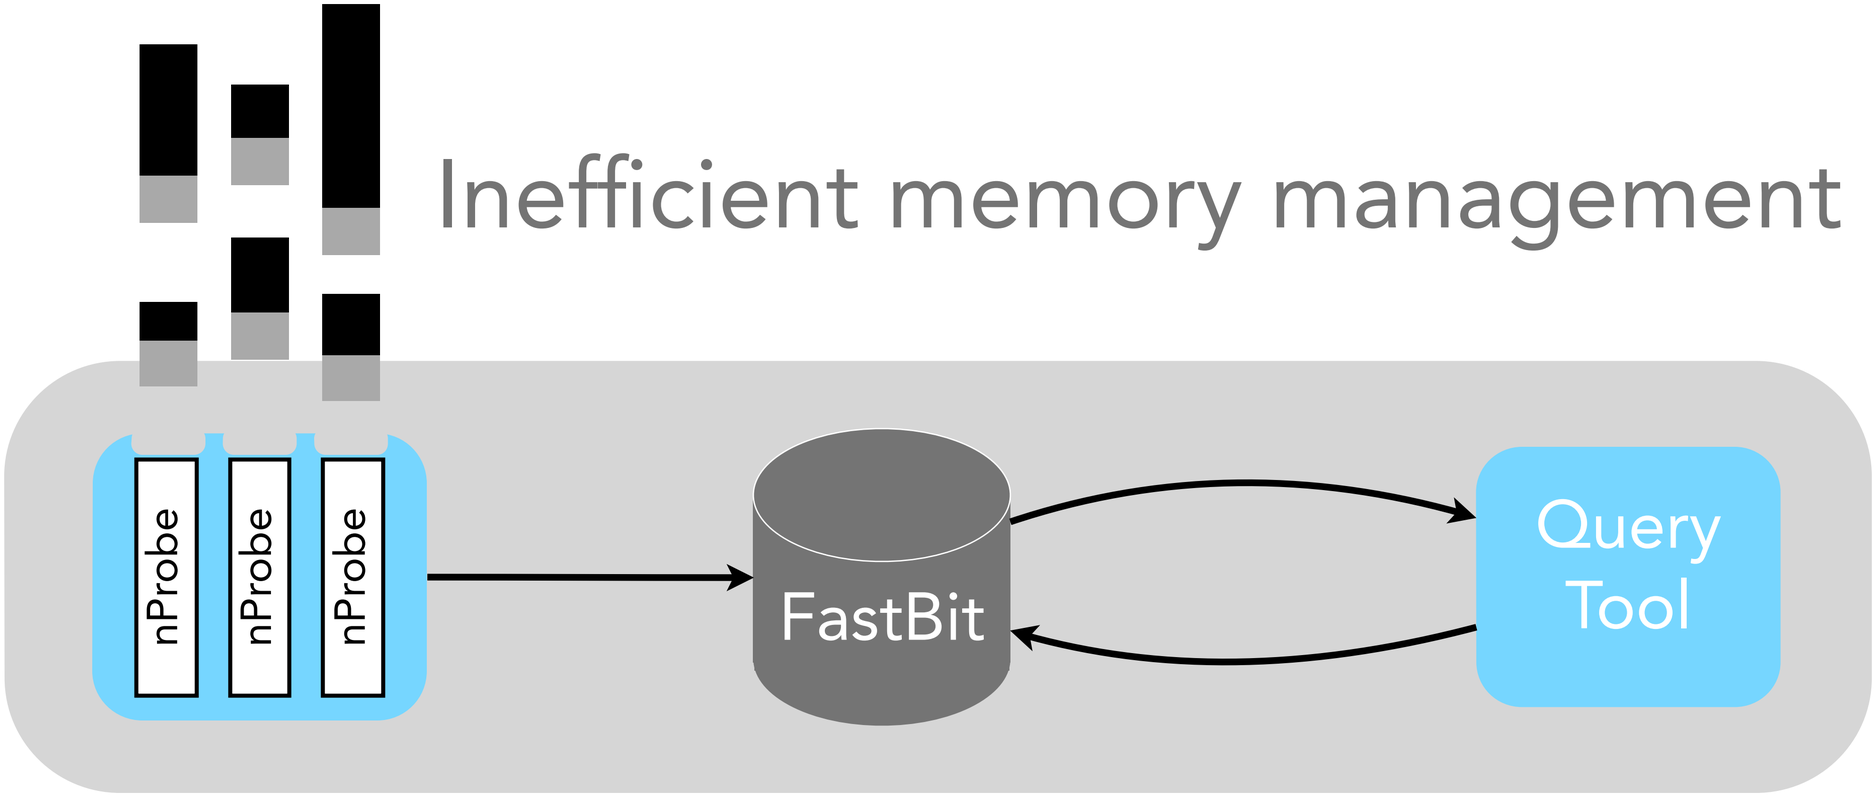

Challenges

Storage Backend

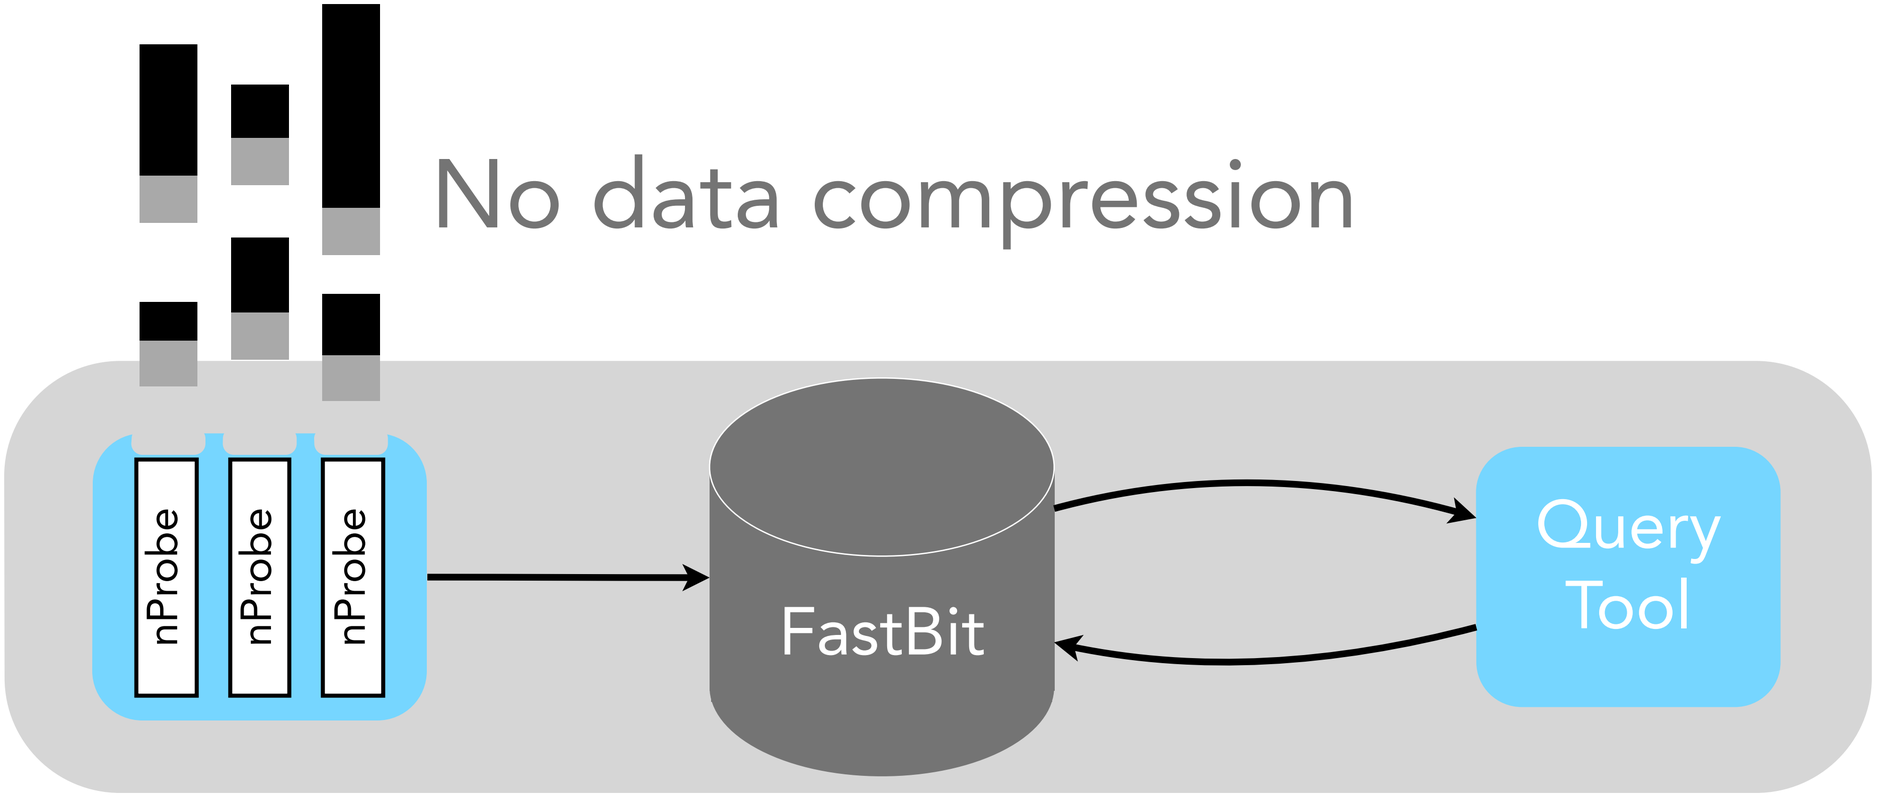

Challenges

Storage Backend

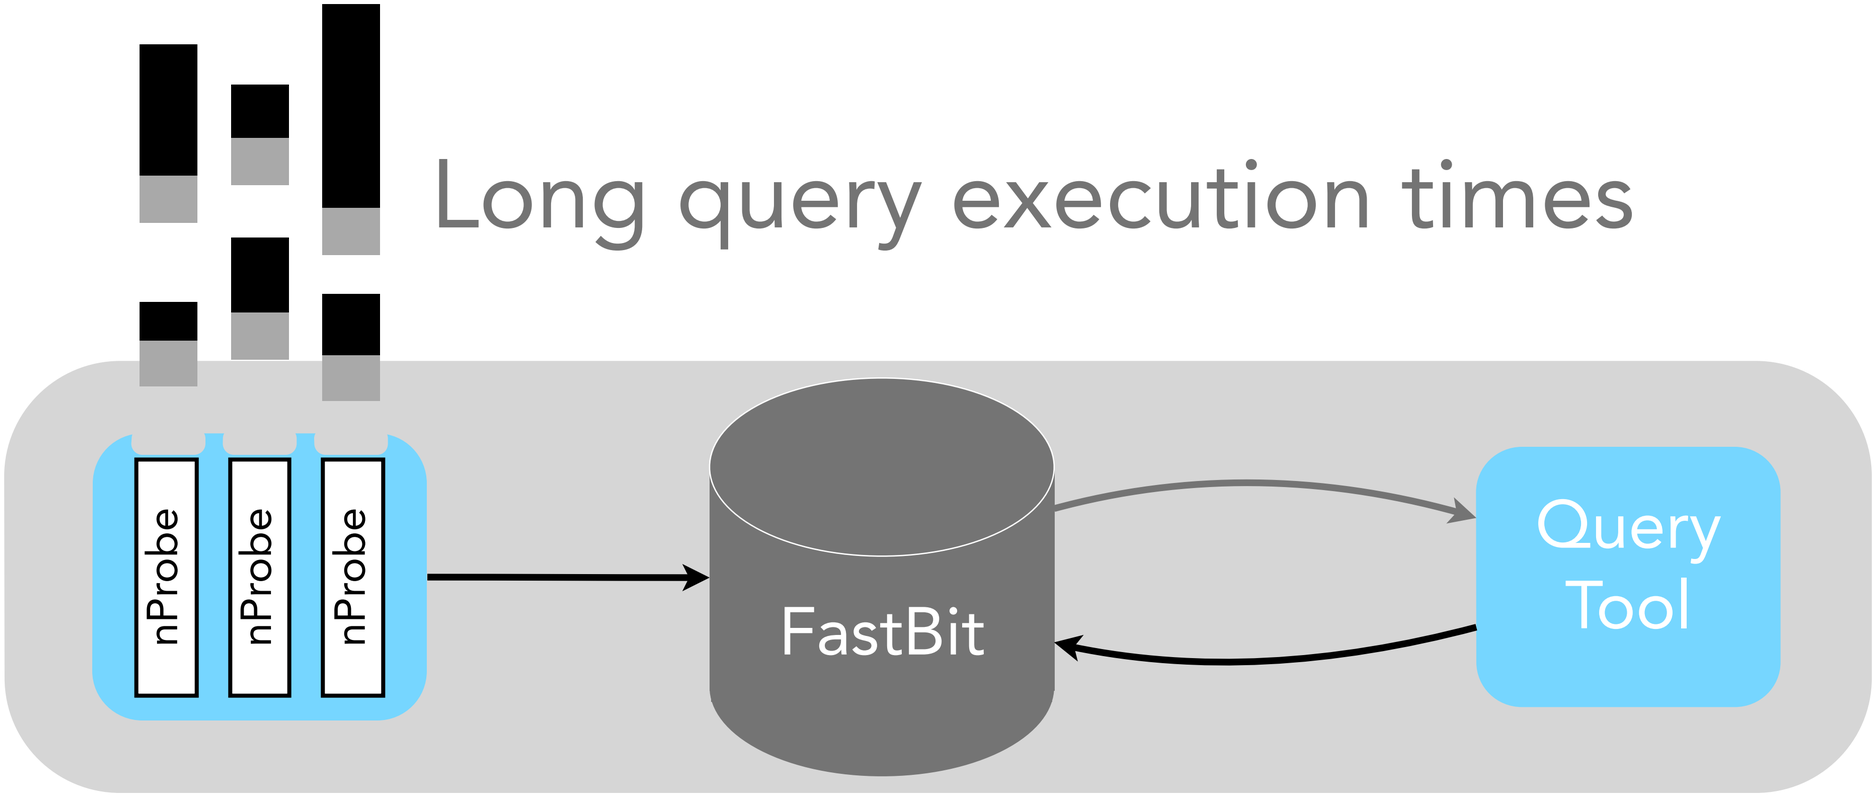

Challenges

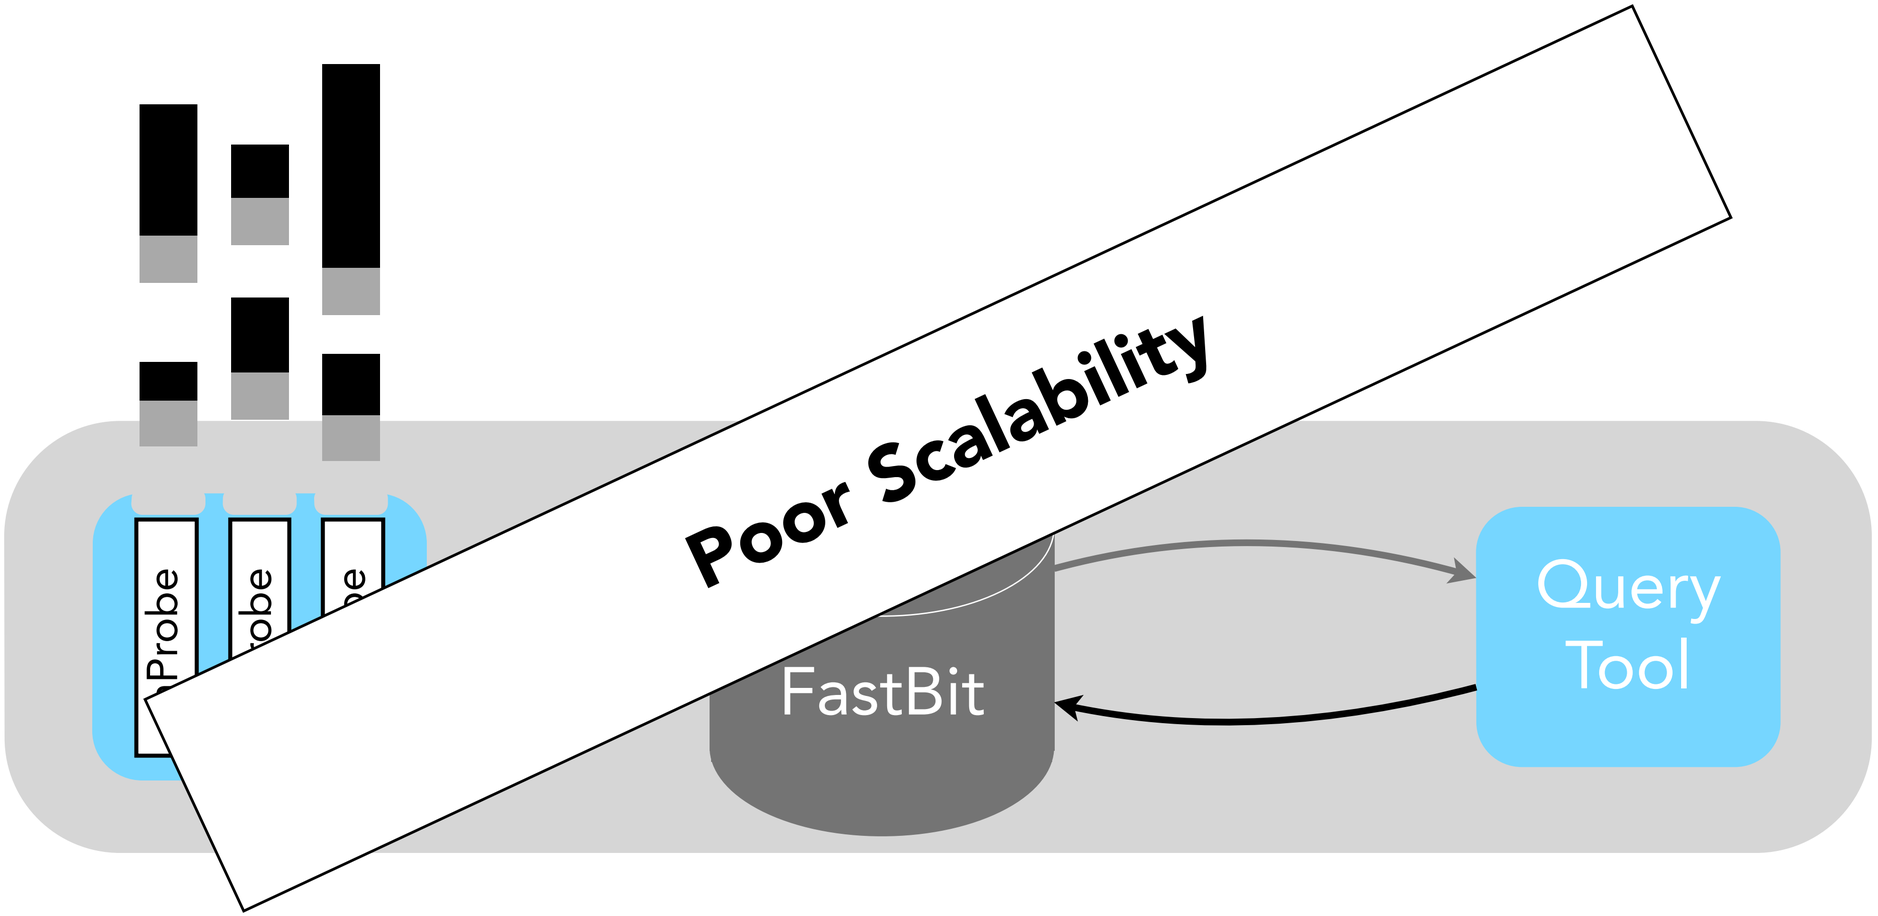

Reduced Flow Format

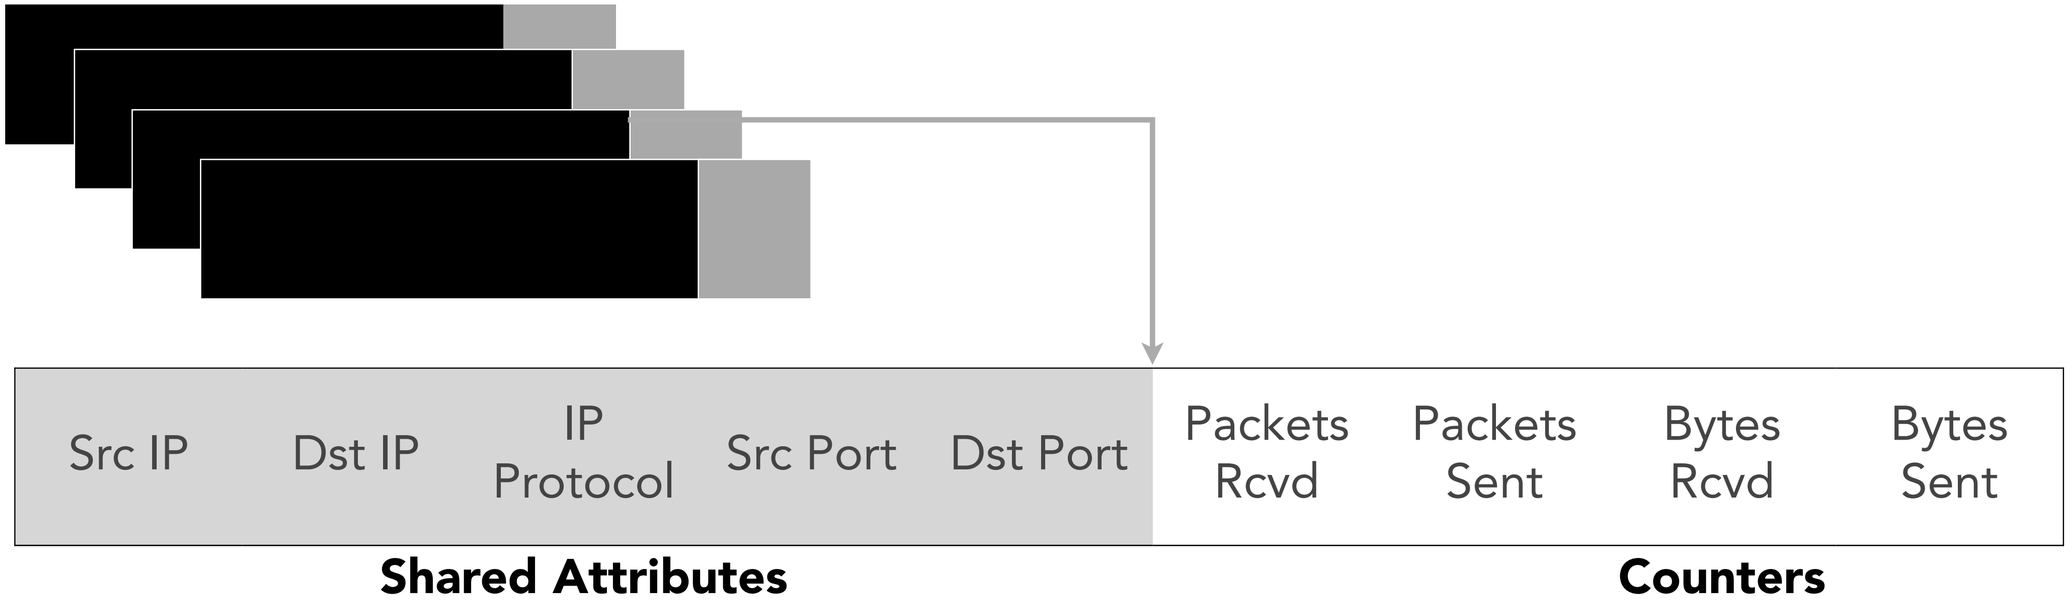

Reduced Flow Format

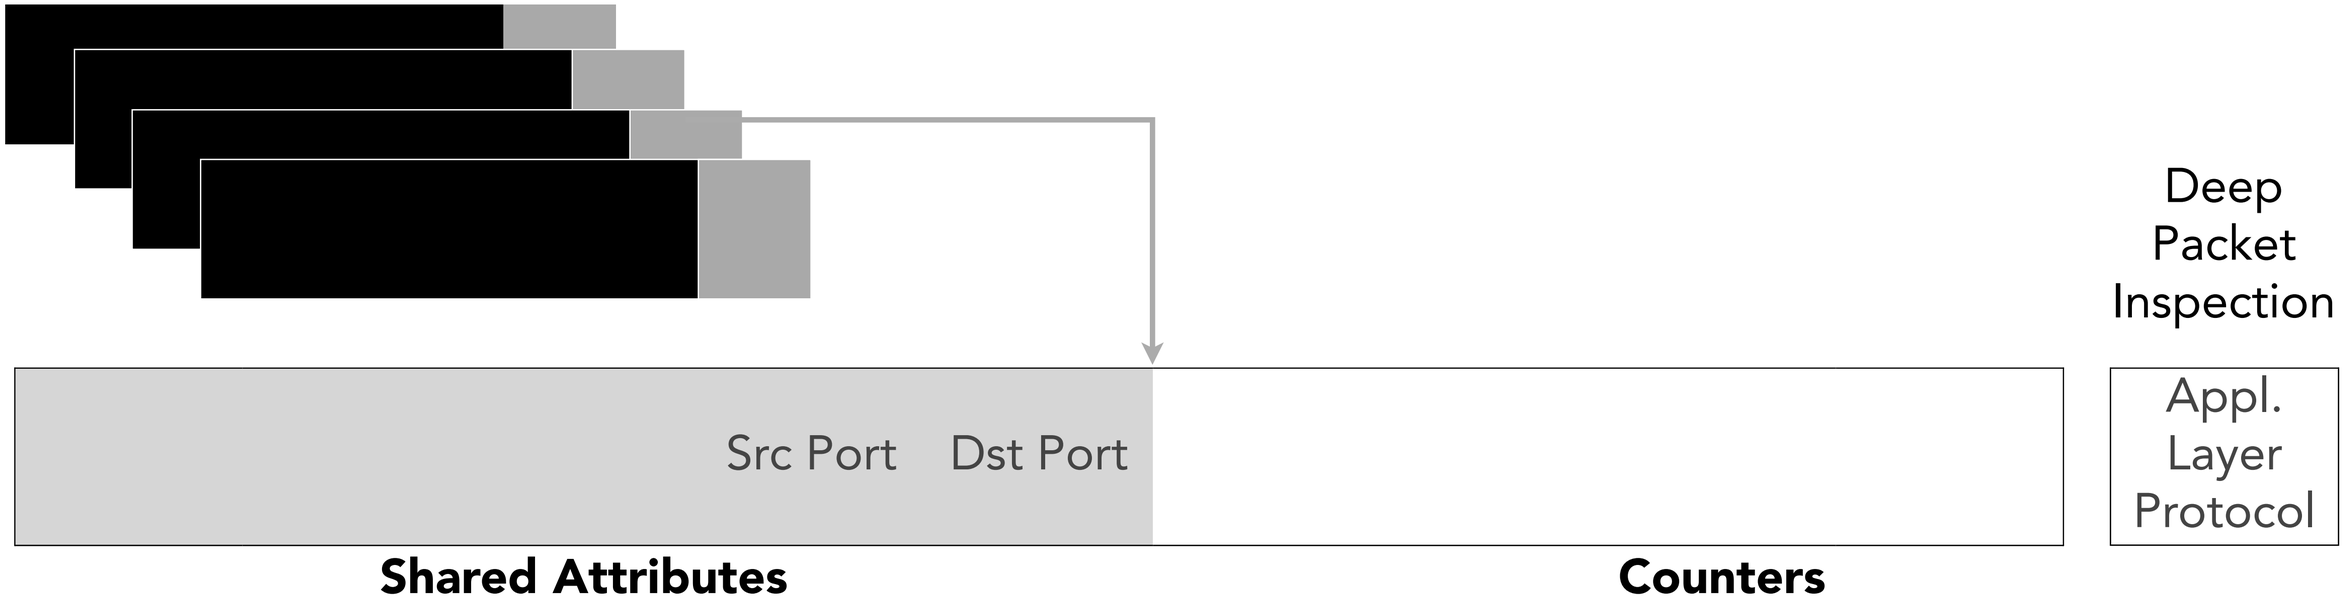

Reduced Flow Format

Deep Packet Inspection

Src Port | Dst Port | Appl. Layer Protocol

Flow in goProbe

Source Port Aggregation

✗

Appl. Layer Protocol | Dst Port

Stored Flow

Collection of Flow Information — goProbe

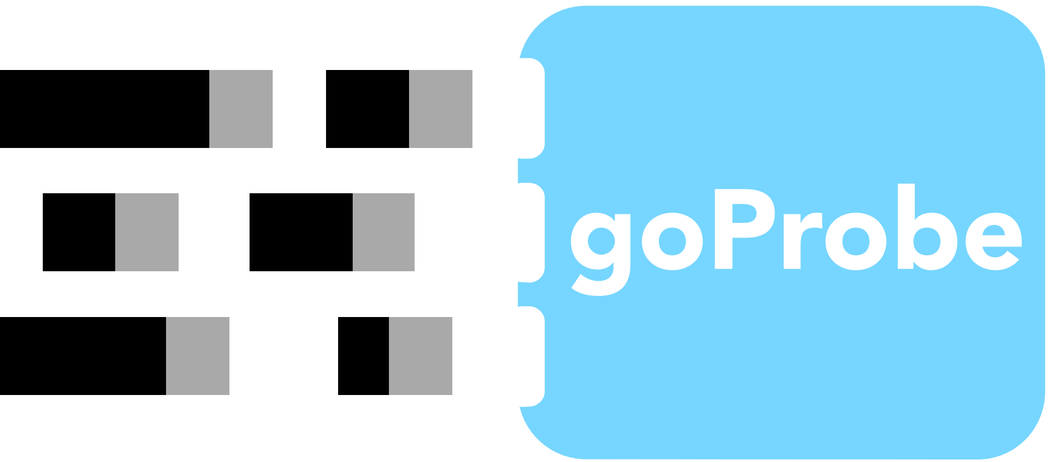

Written in Google **Go**

One capture routine per interface

Packet capture using modified **libpcap**

Database flush in regular intervals

goProbe – Concept (Multiple Interfaces)

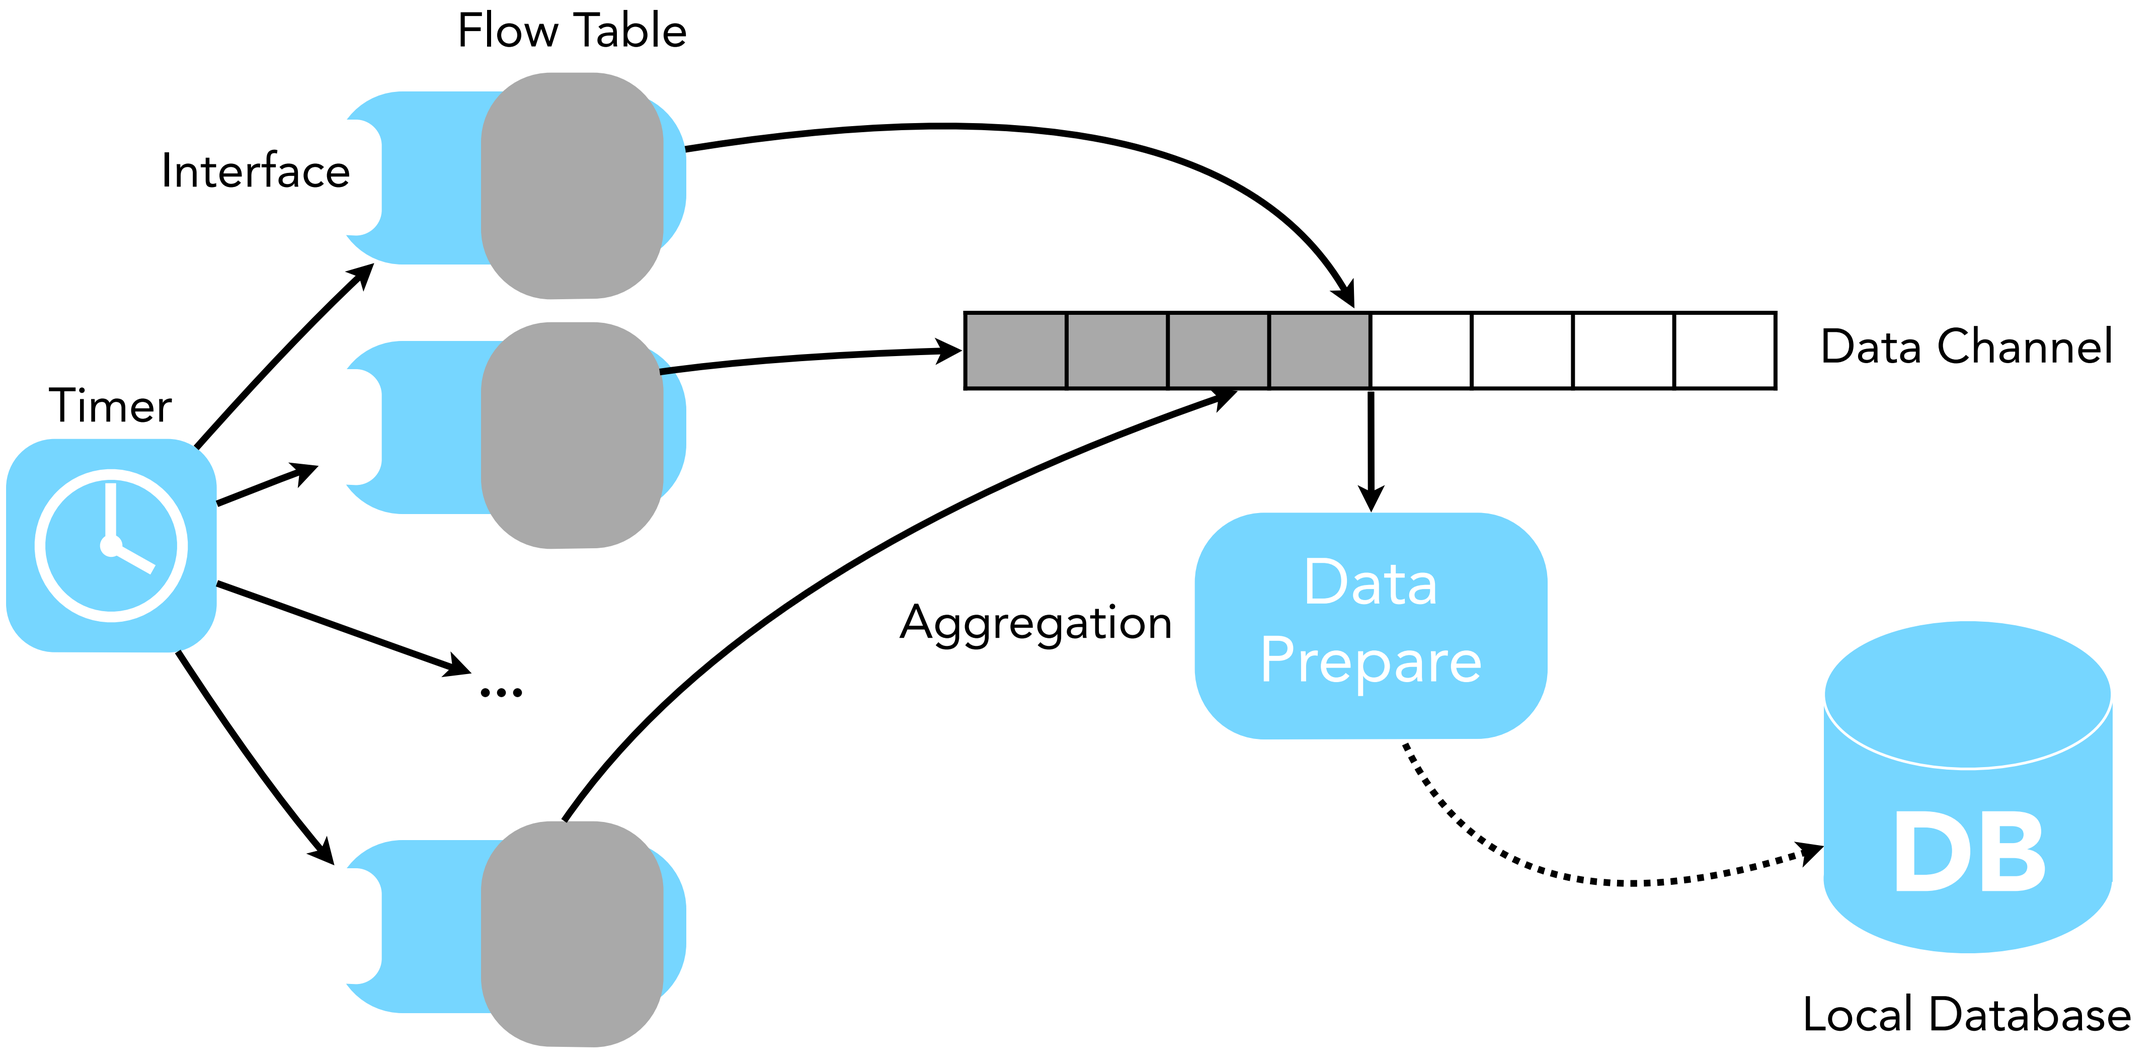

How does it Compare?

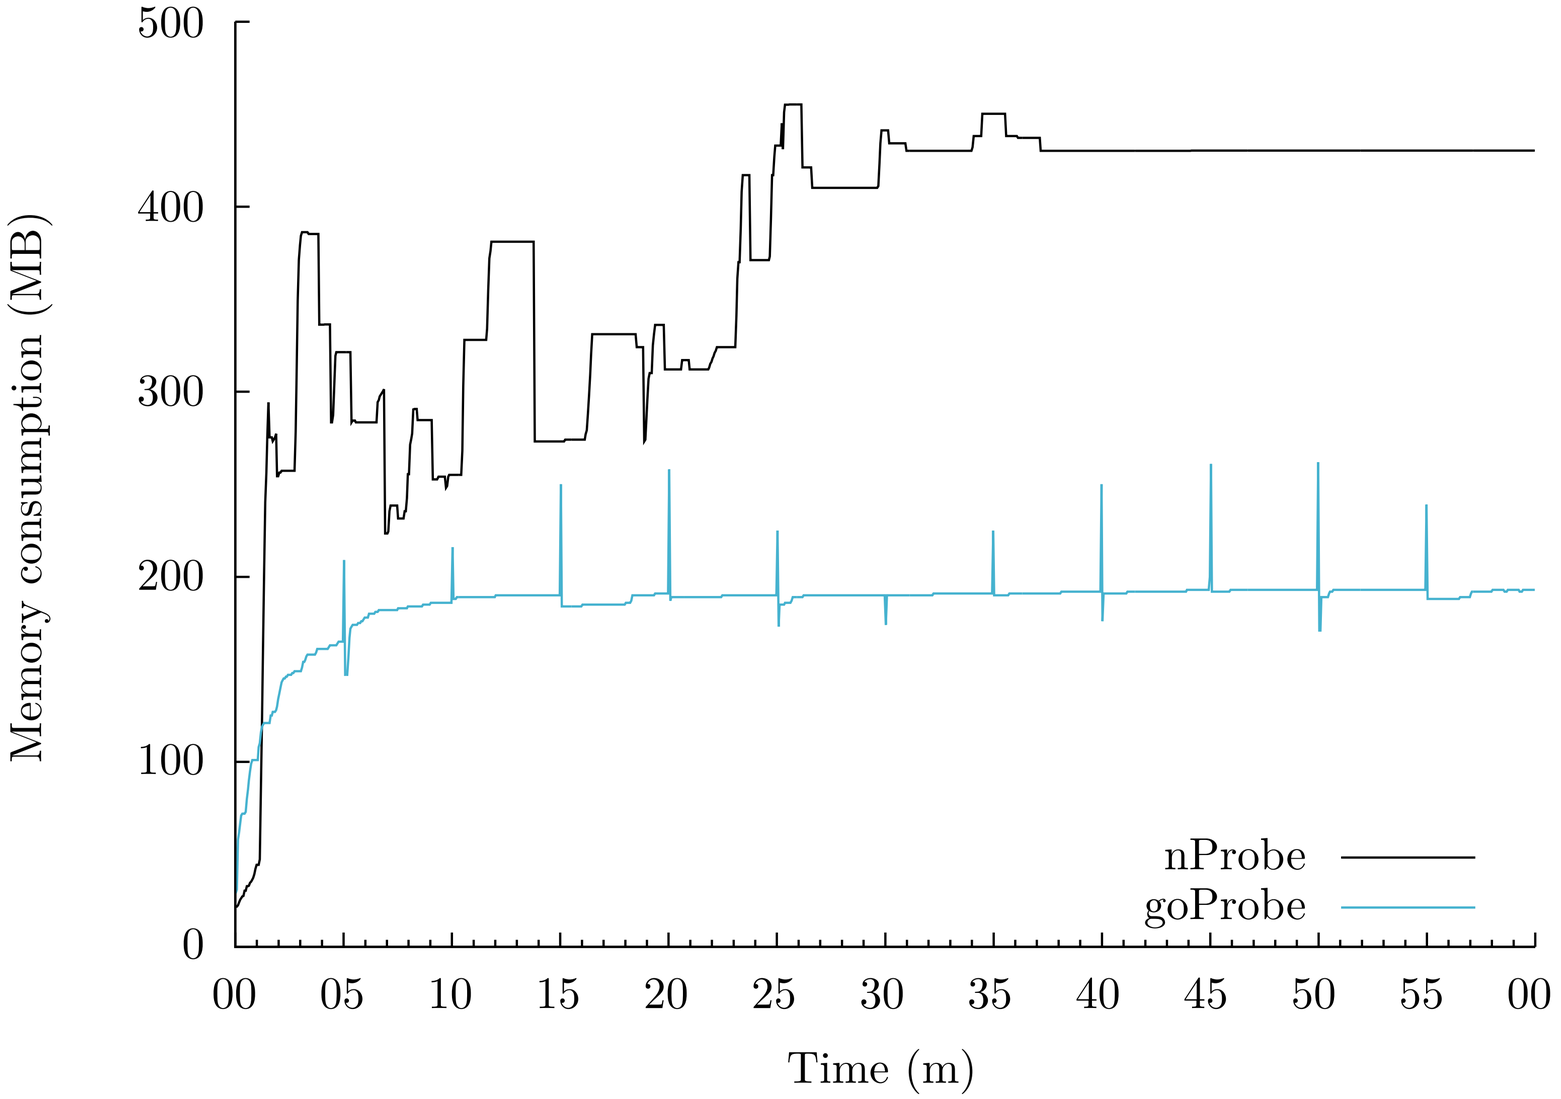

Database Performance Evaluation

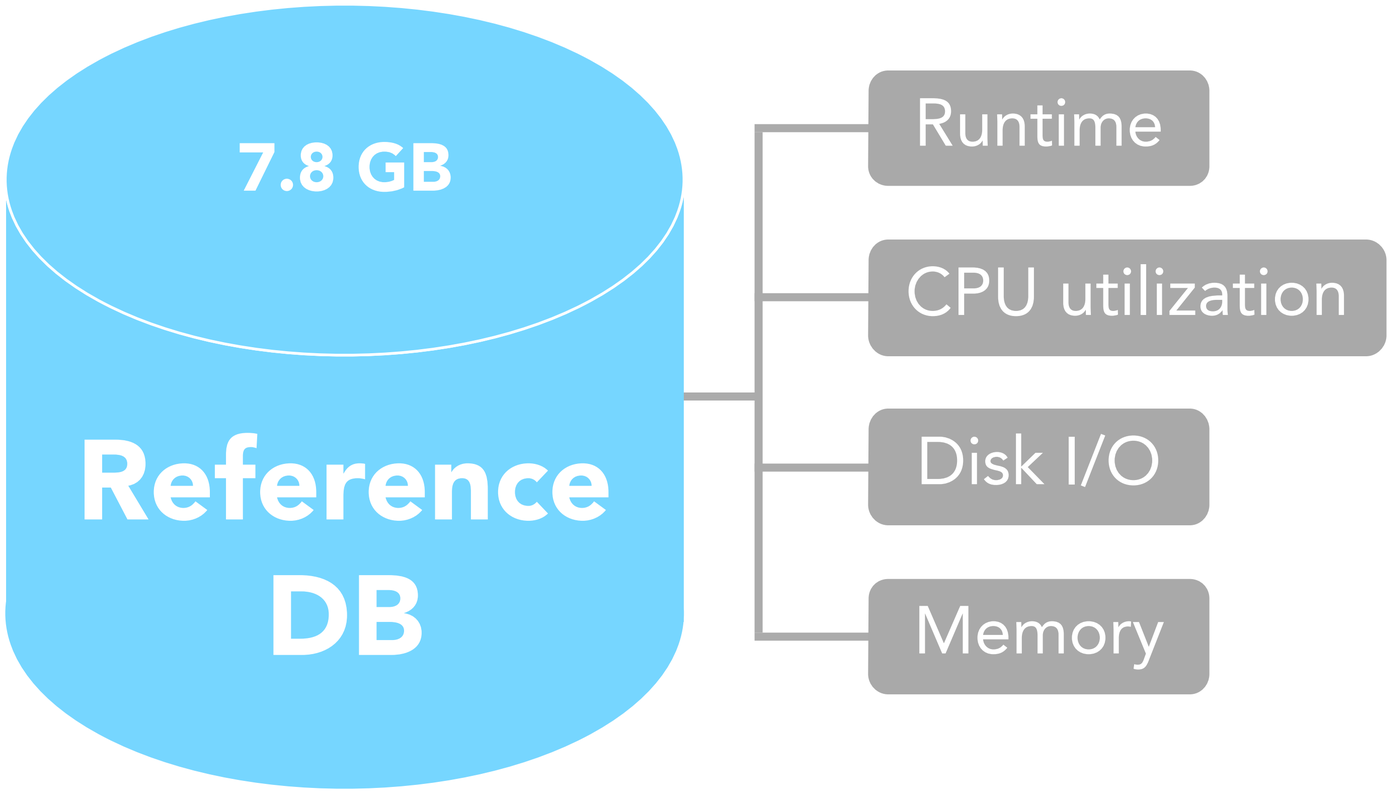

120 Million Entries

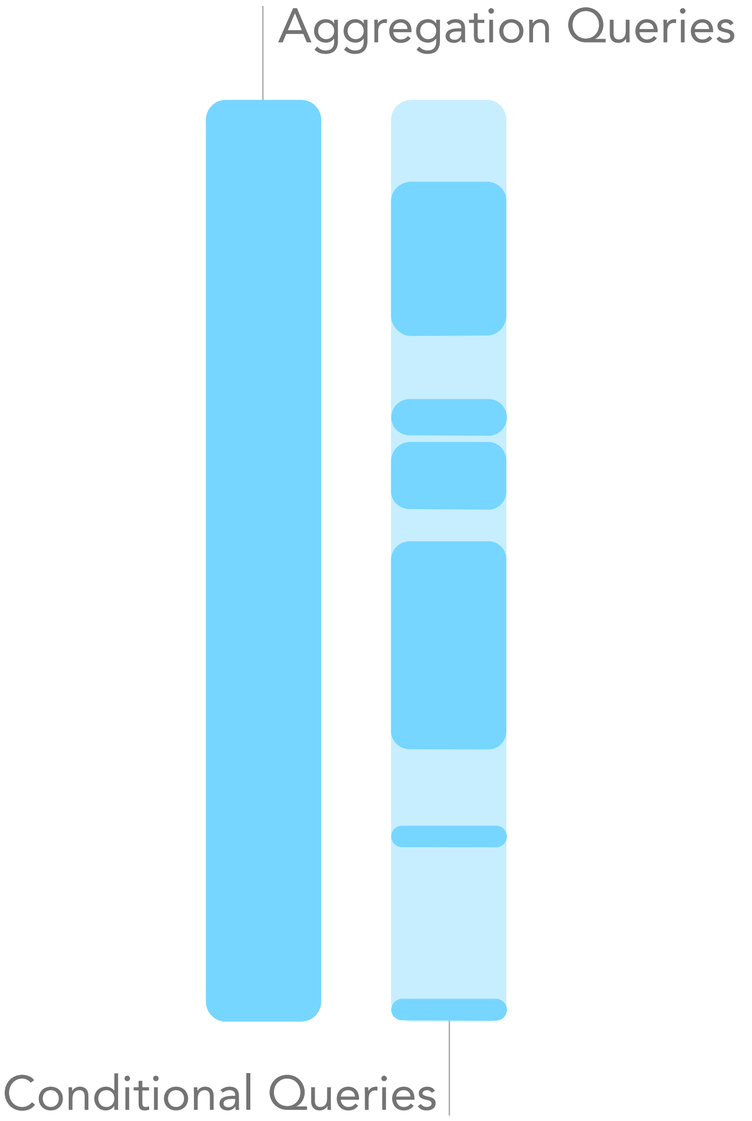

Results

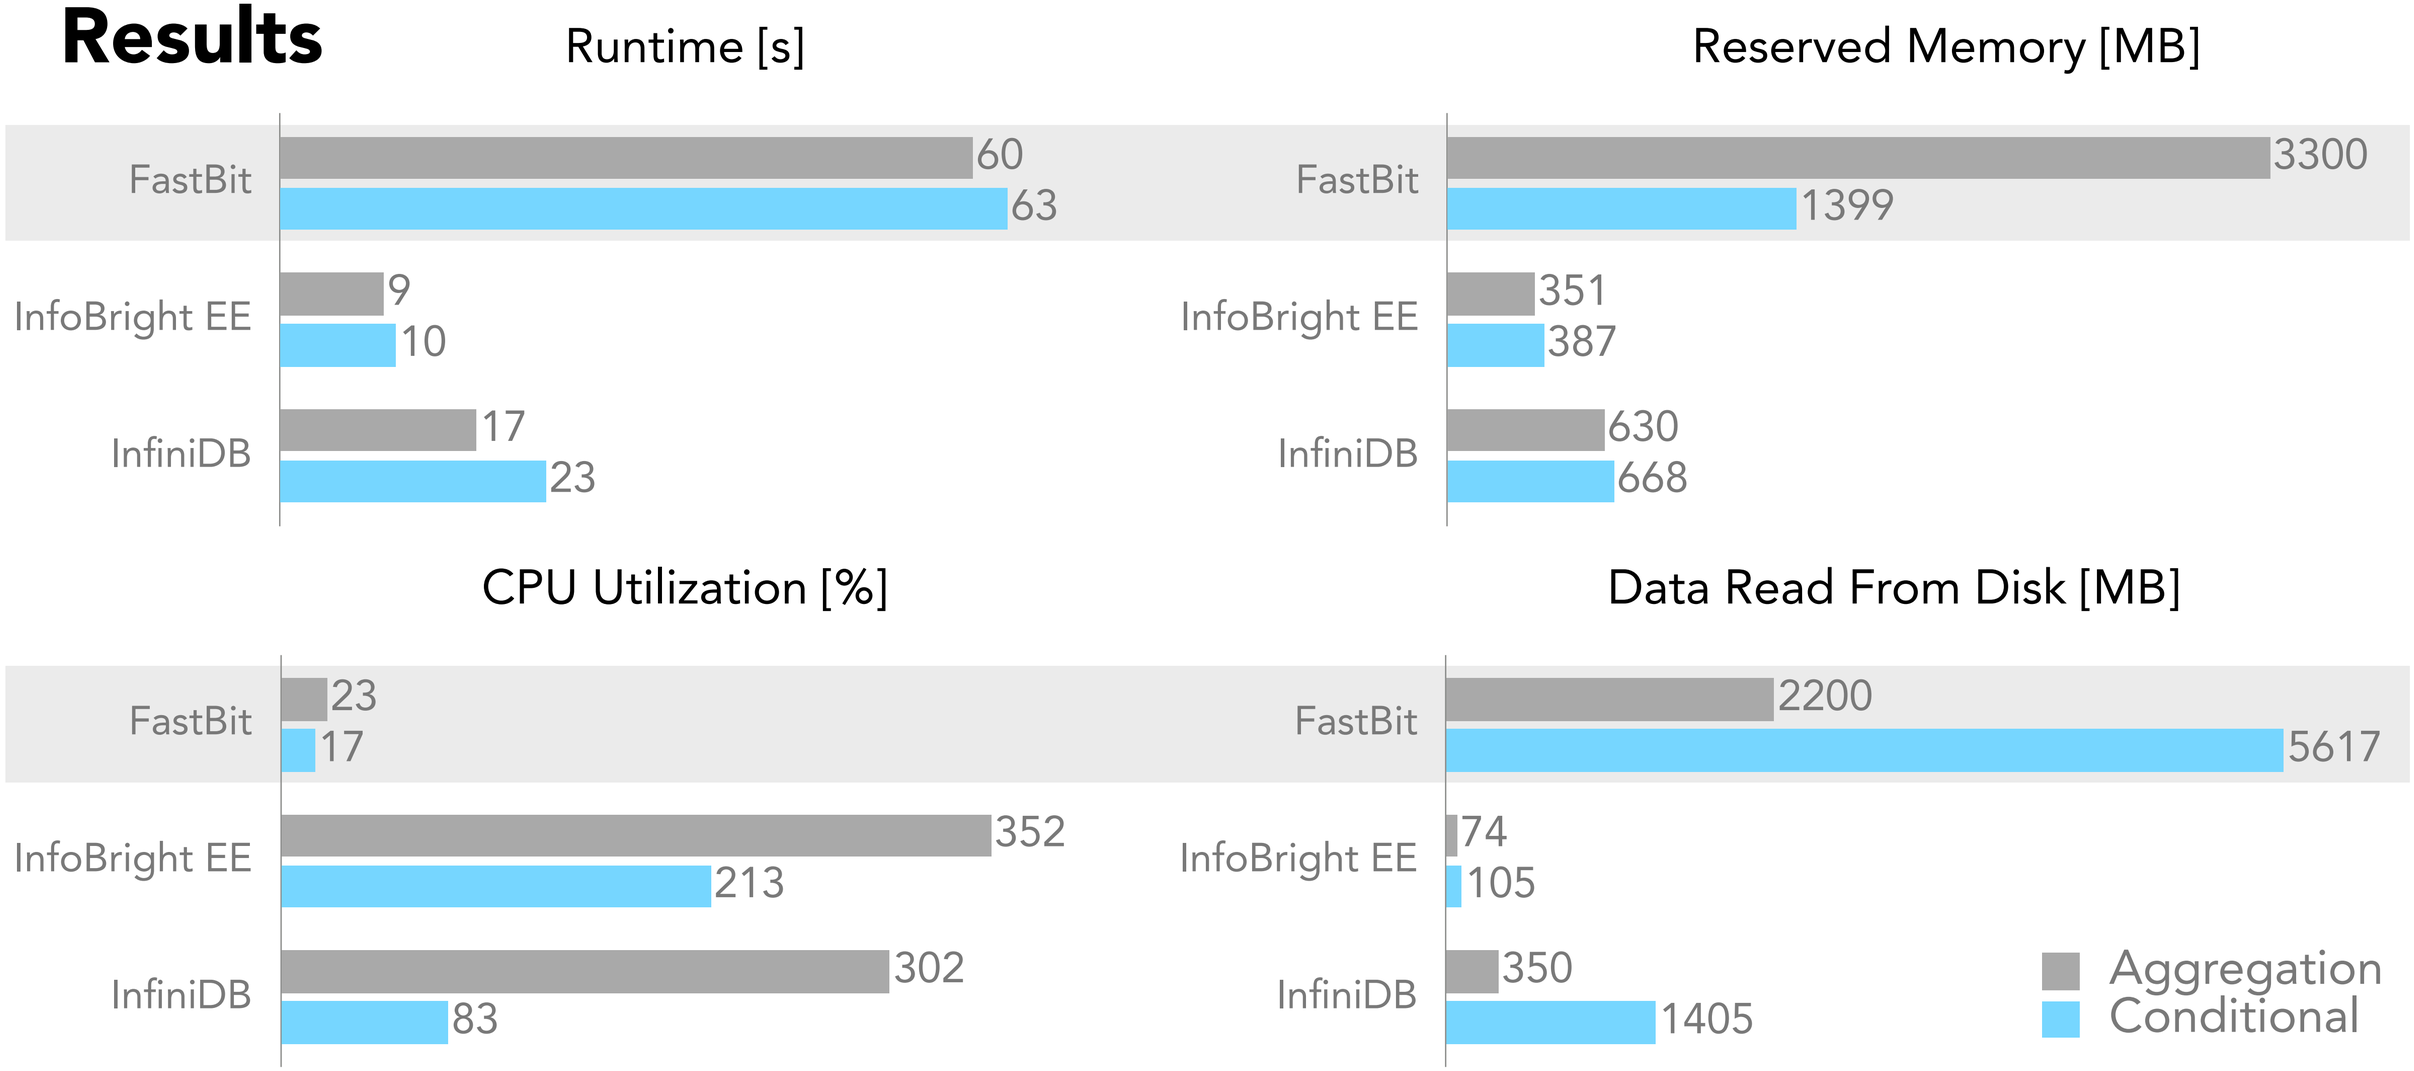
Runtime [s]
FastBit
60
63
InfoBright EE
9
10
InfiniDB
17
23
Reserved Memory [MB]
FastBit
3300
1399
InfoBright EE
351
387
InfiniDB
630
668
CPU Utilization [%]
FastBit
23
17
InfoBright EE
352
213
InfiniDB
302
83
Data Read From Disk [MB]
FastBit
2200
5617
InfoBright EE
74
105
InfiniDB
350
1405
Aggregation
Conditional

Infobright EE

InfiniDB

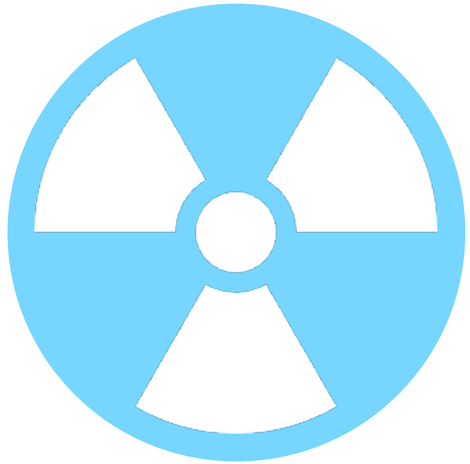

Tailored Column Store

File Based

Compression

Concurrency

Independent Processing

Tailored Column Store — goDB

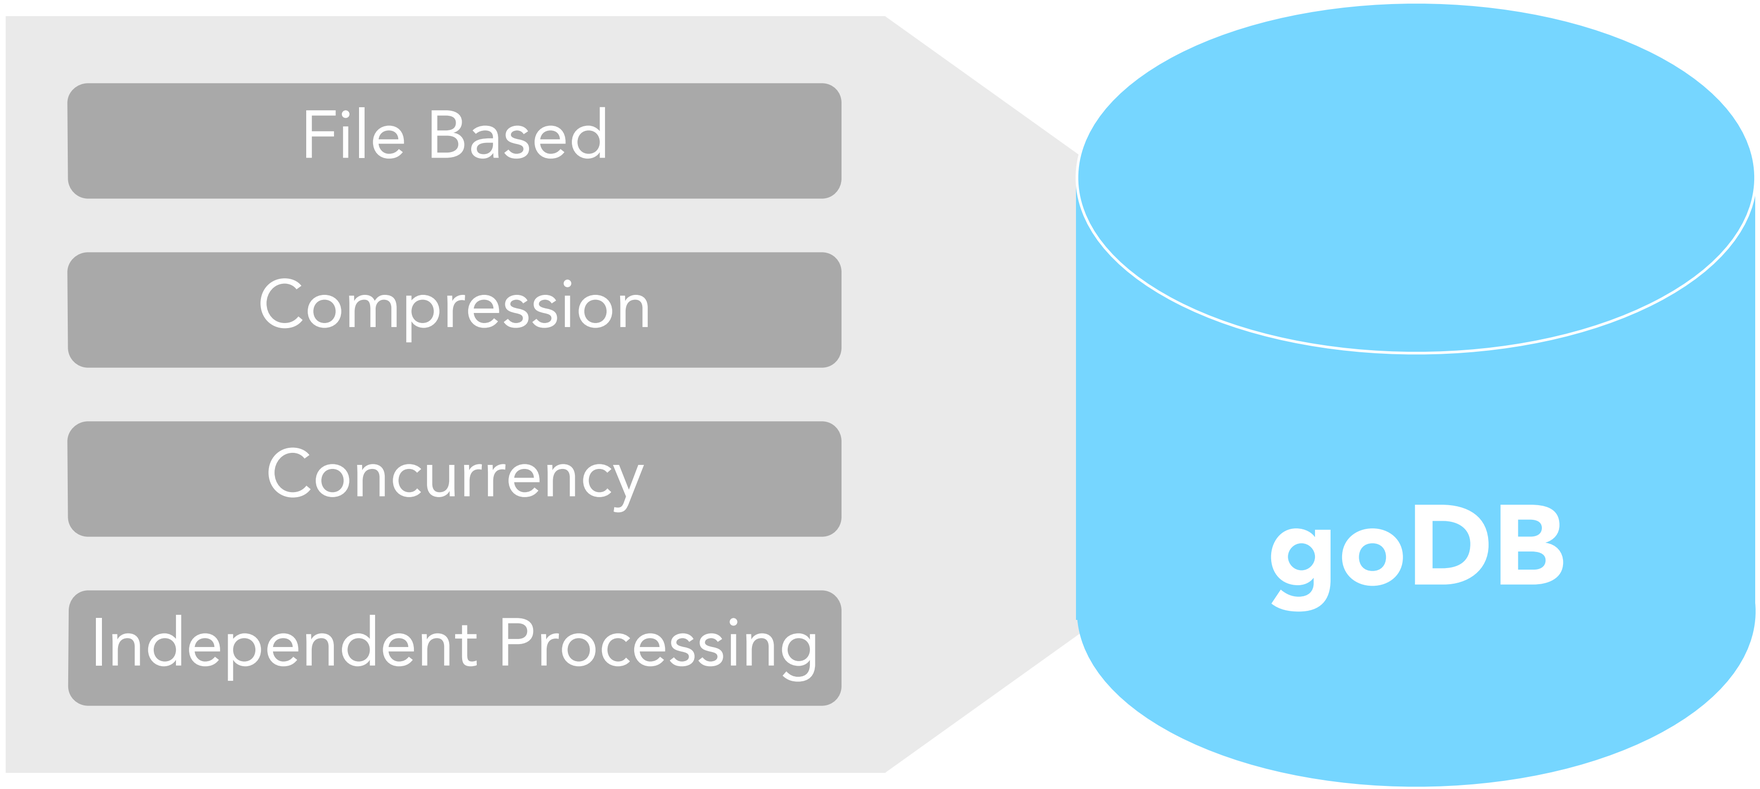

One File per Attribute

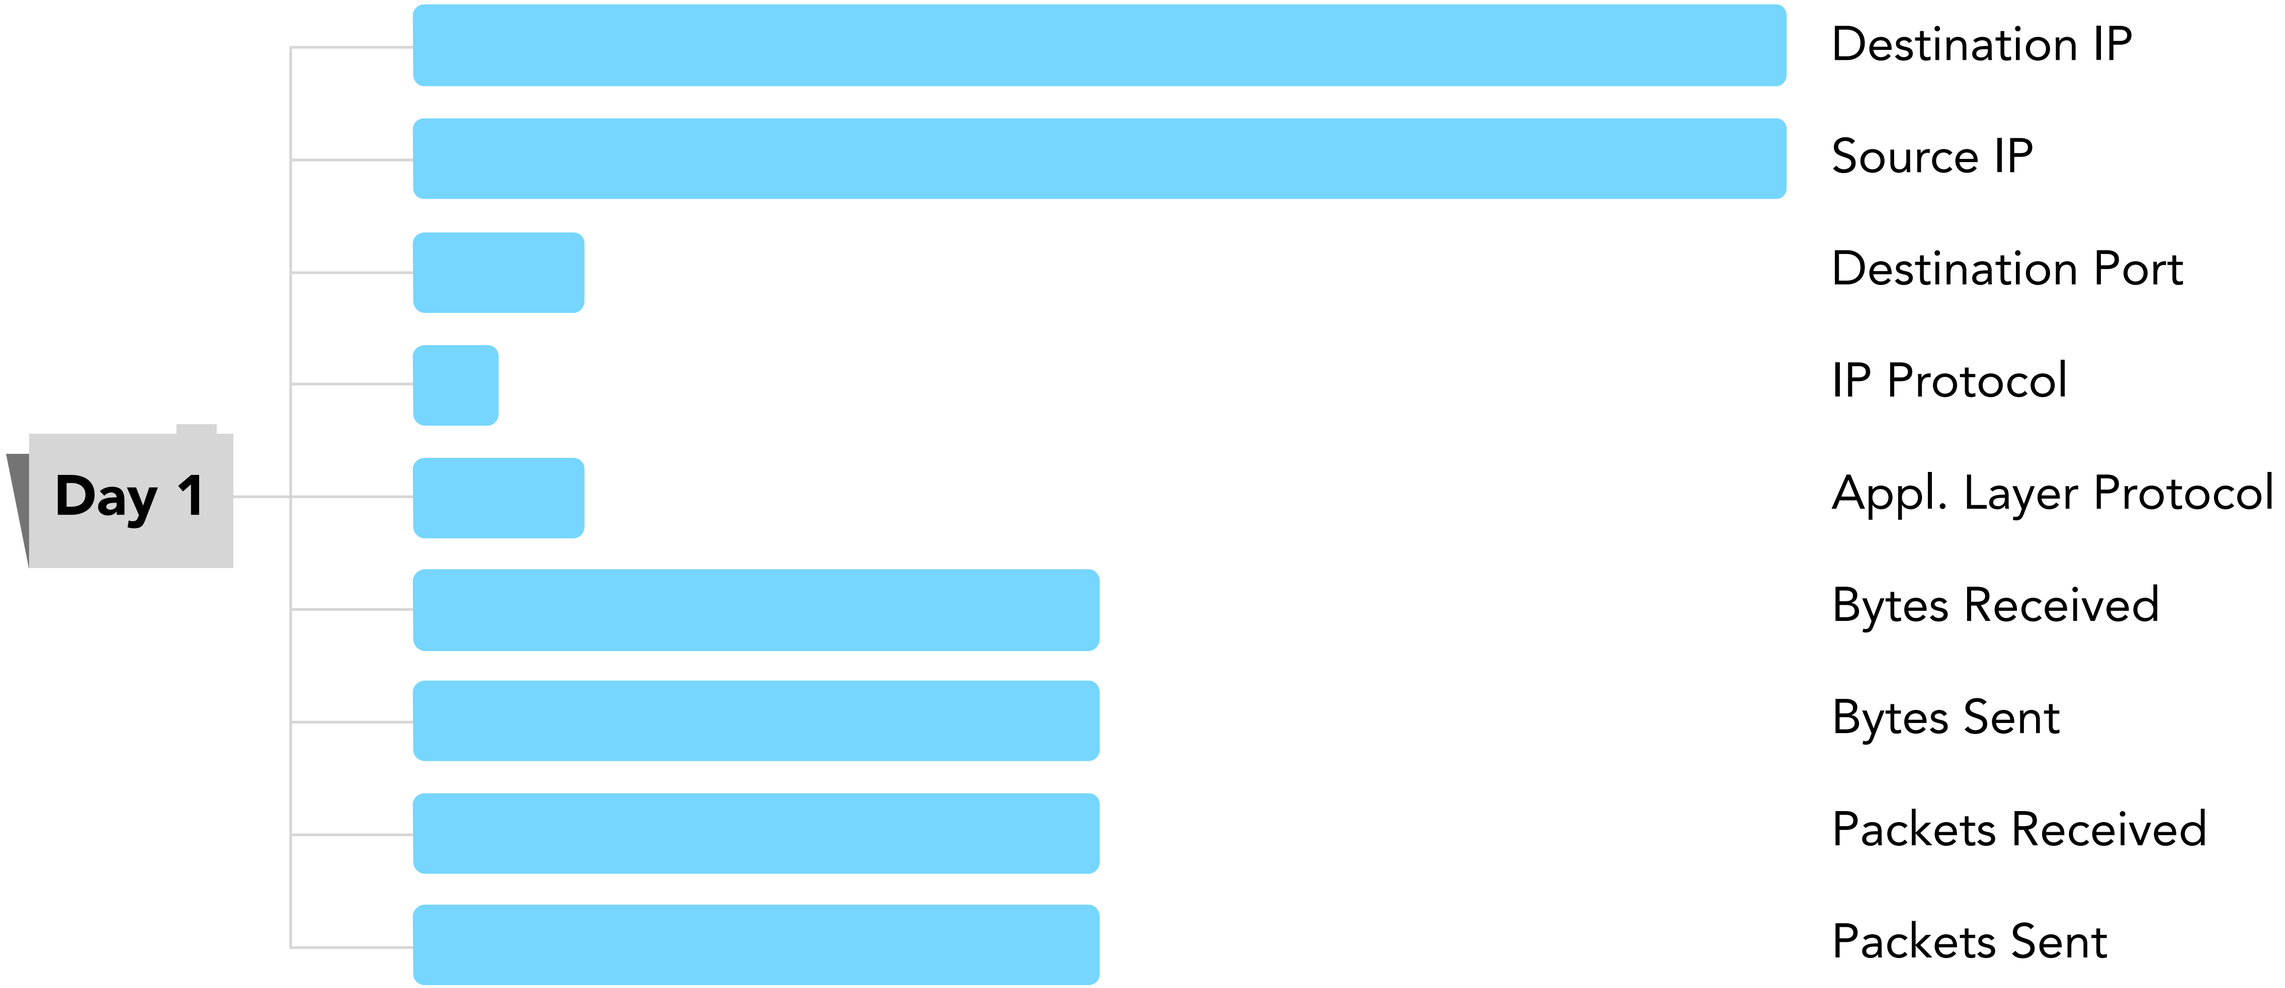

One File per Attribute

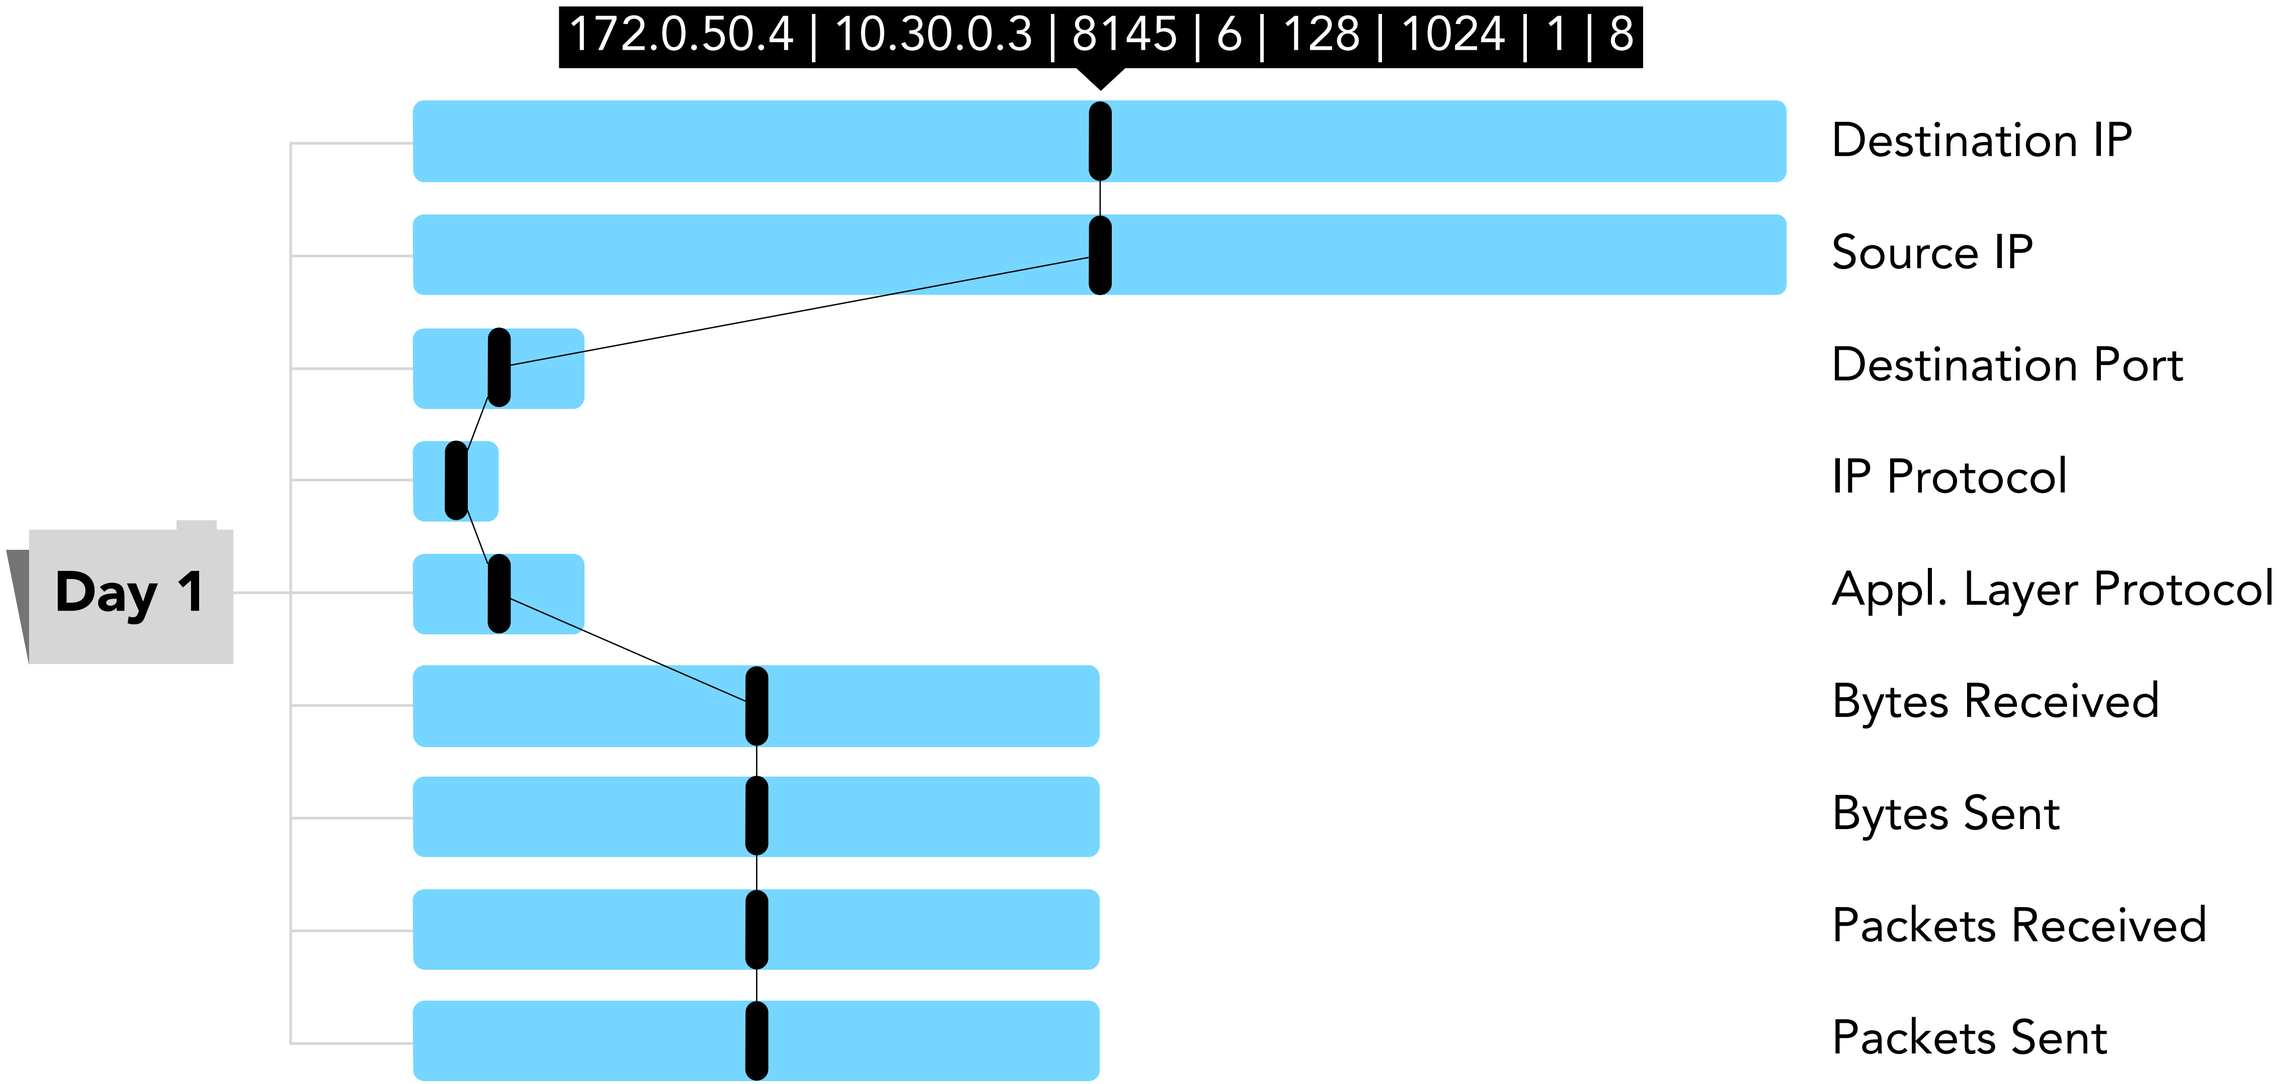

Block-wise Writing and Reading

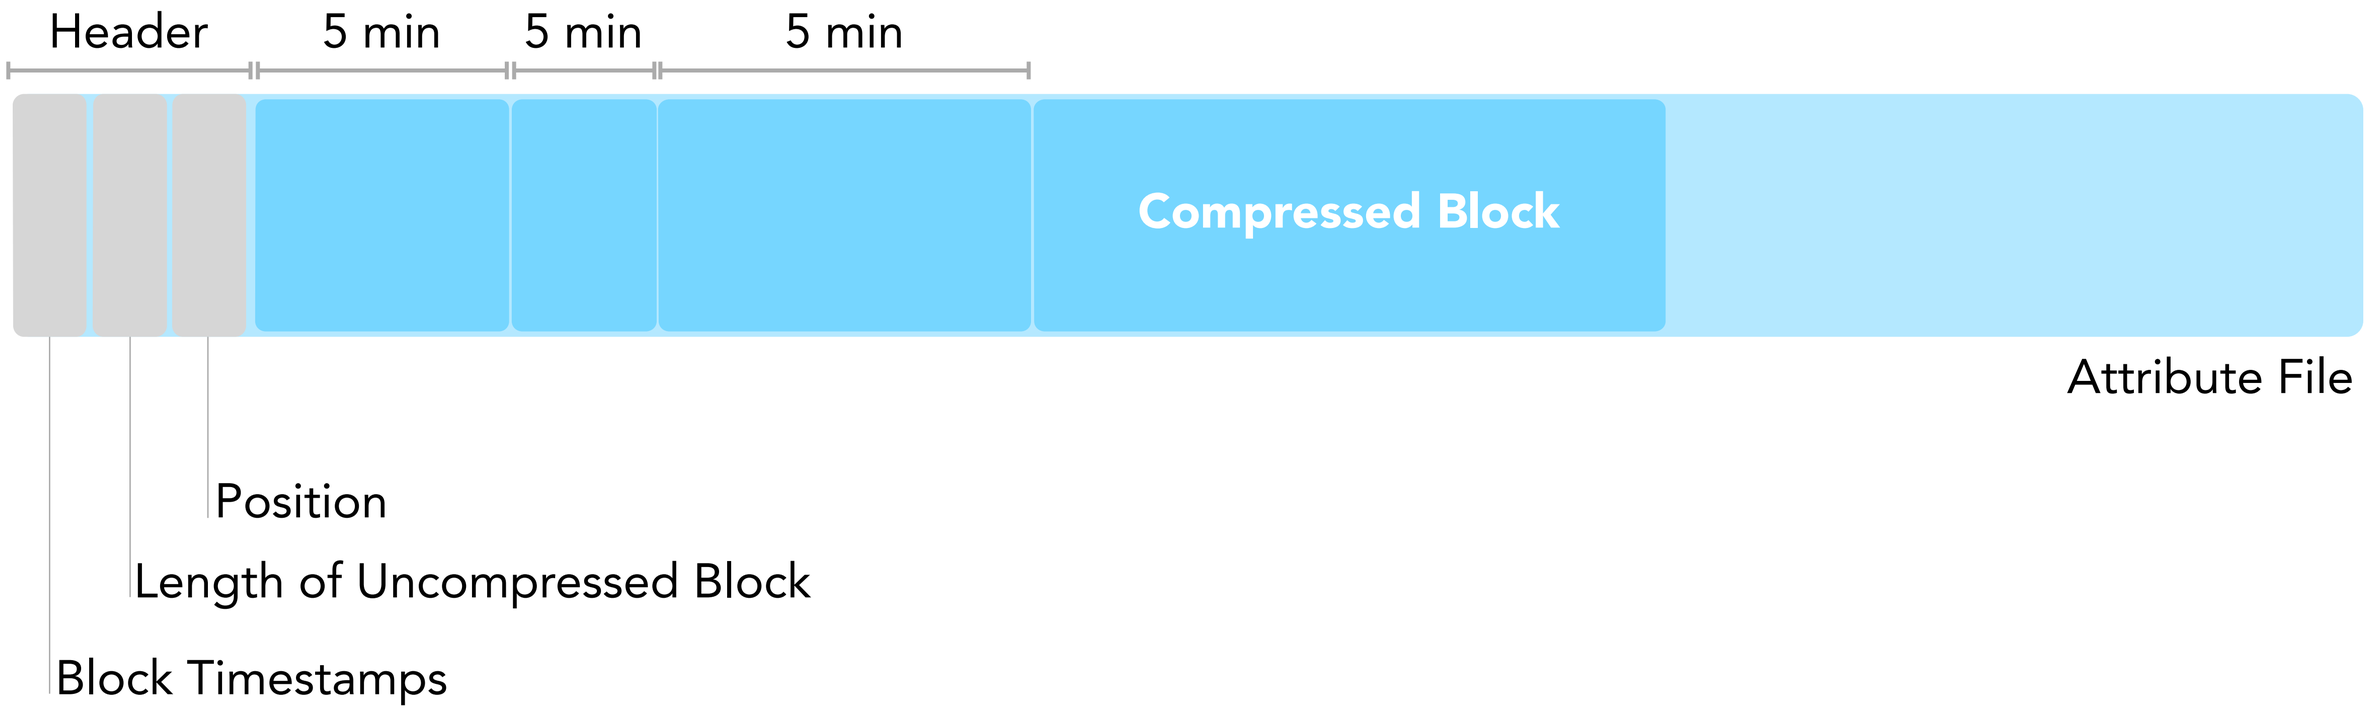

Concurrent Processing

Concurrent Processing

Concurrent Processing

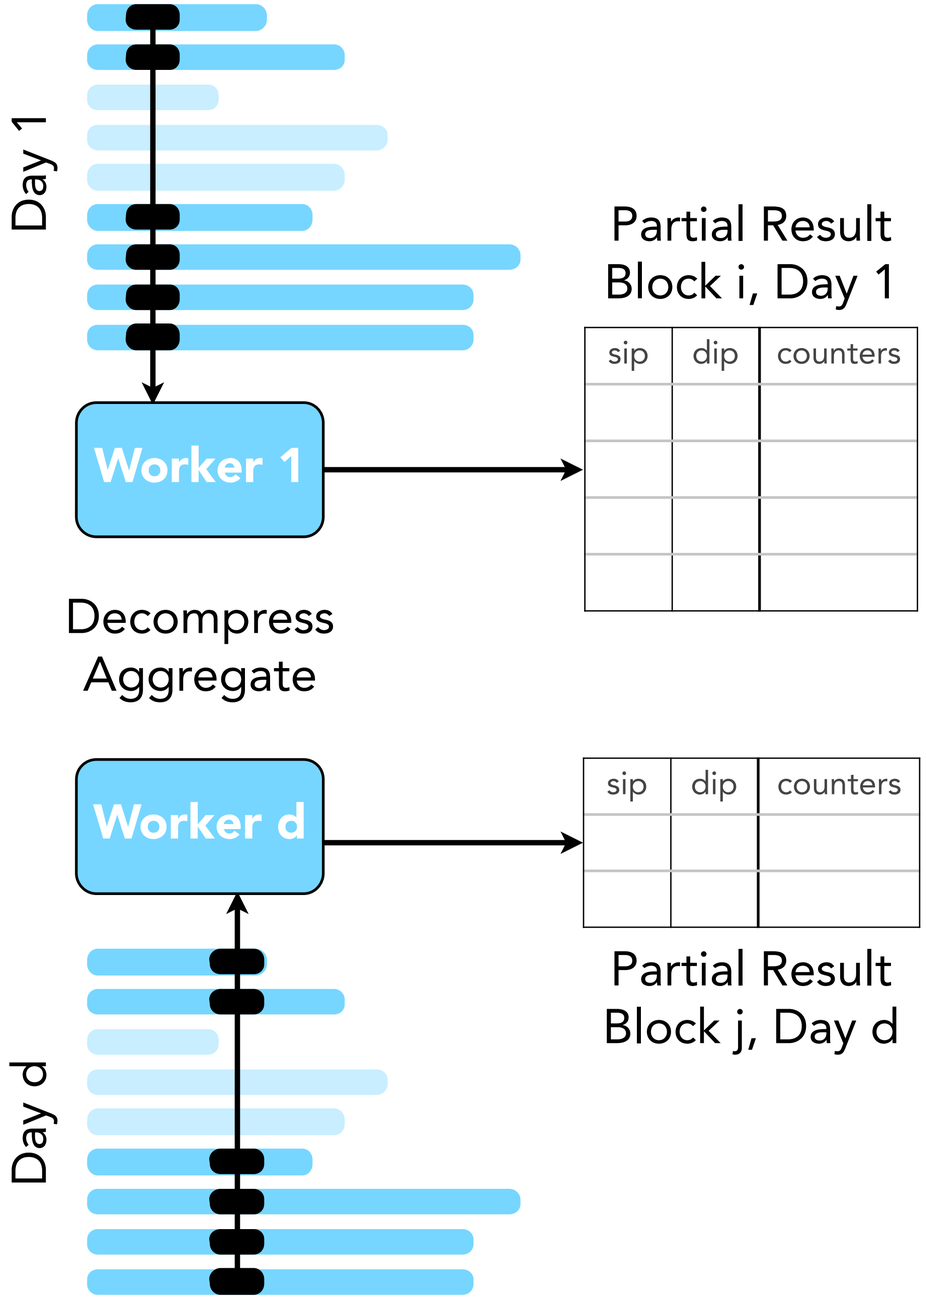

Concurrent Processing

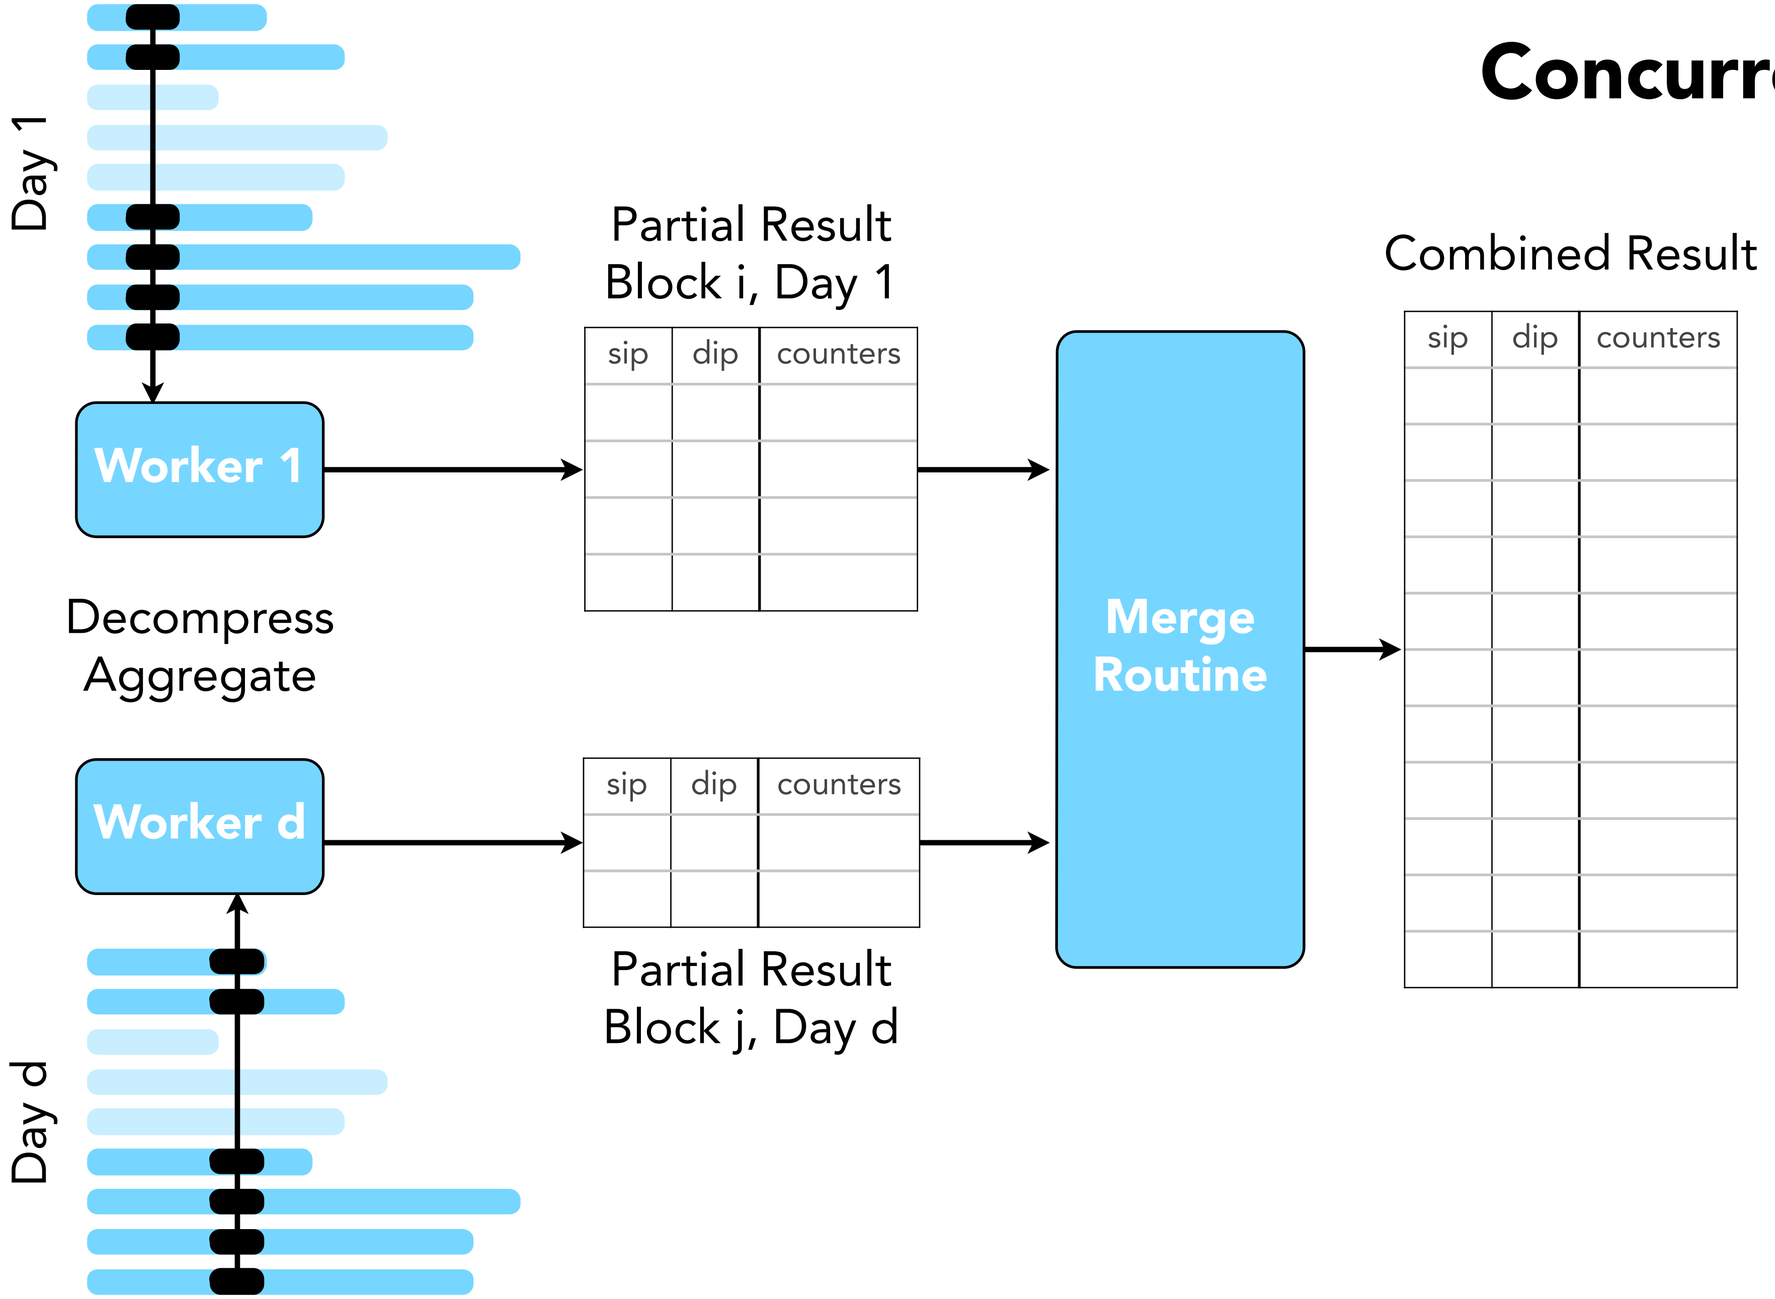

Concurrent Processing

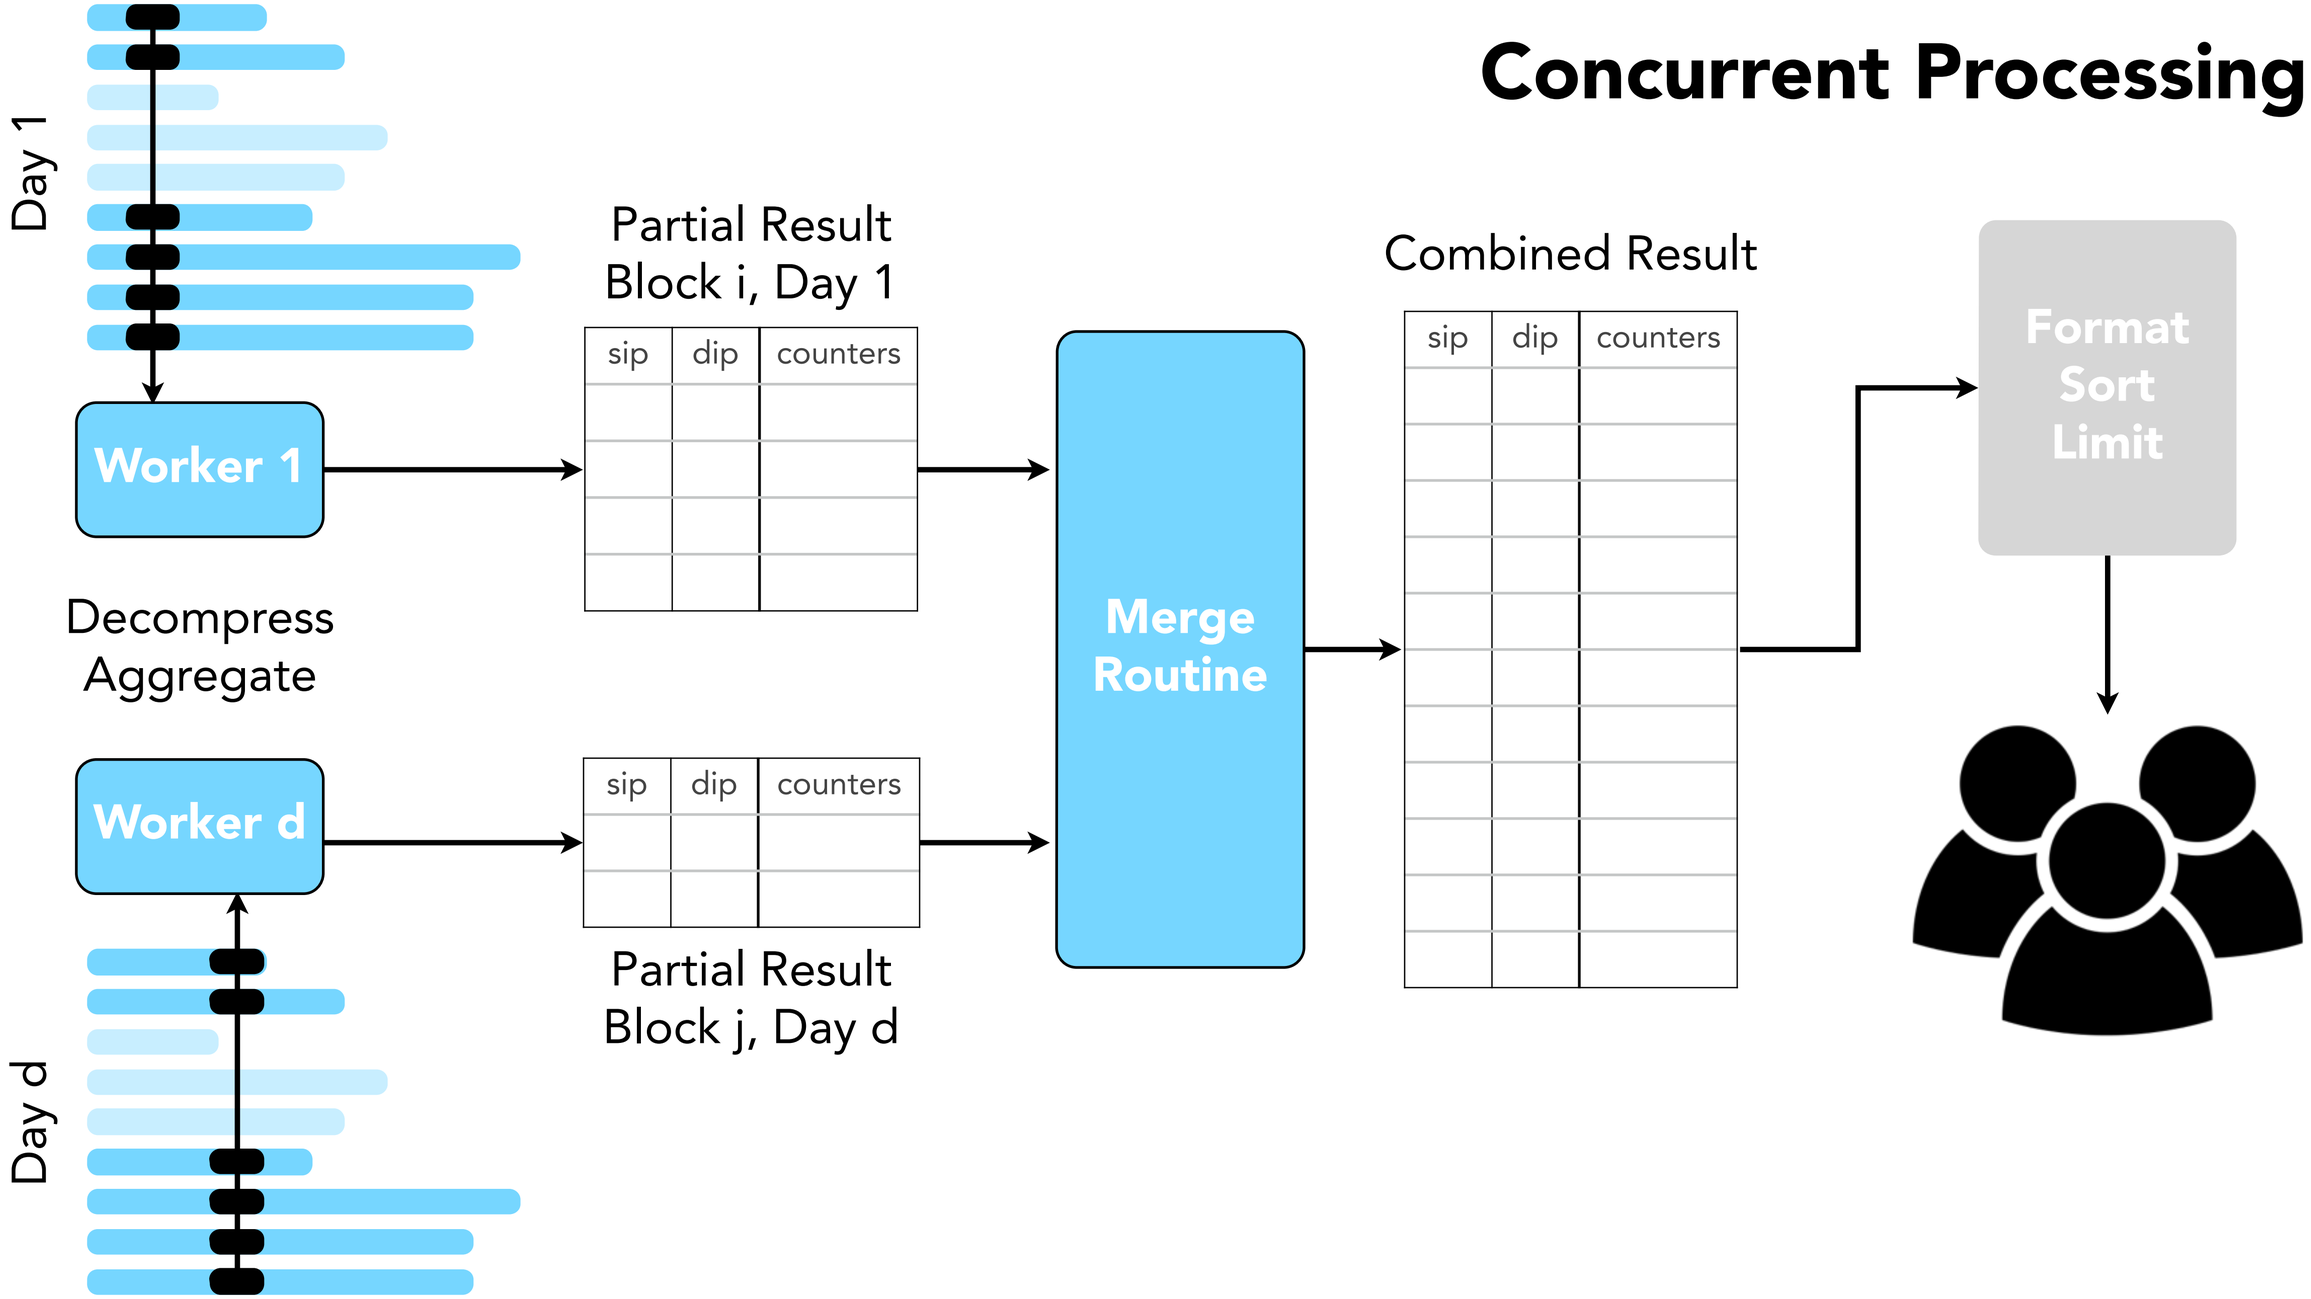

How does it Compare?

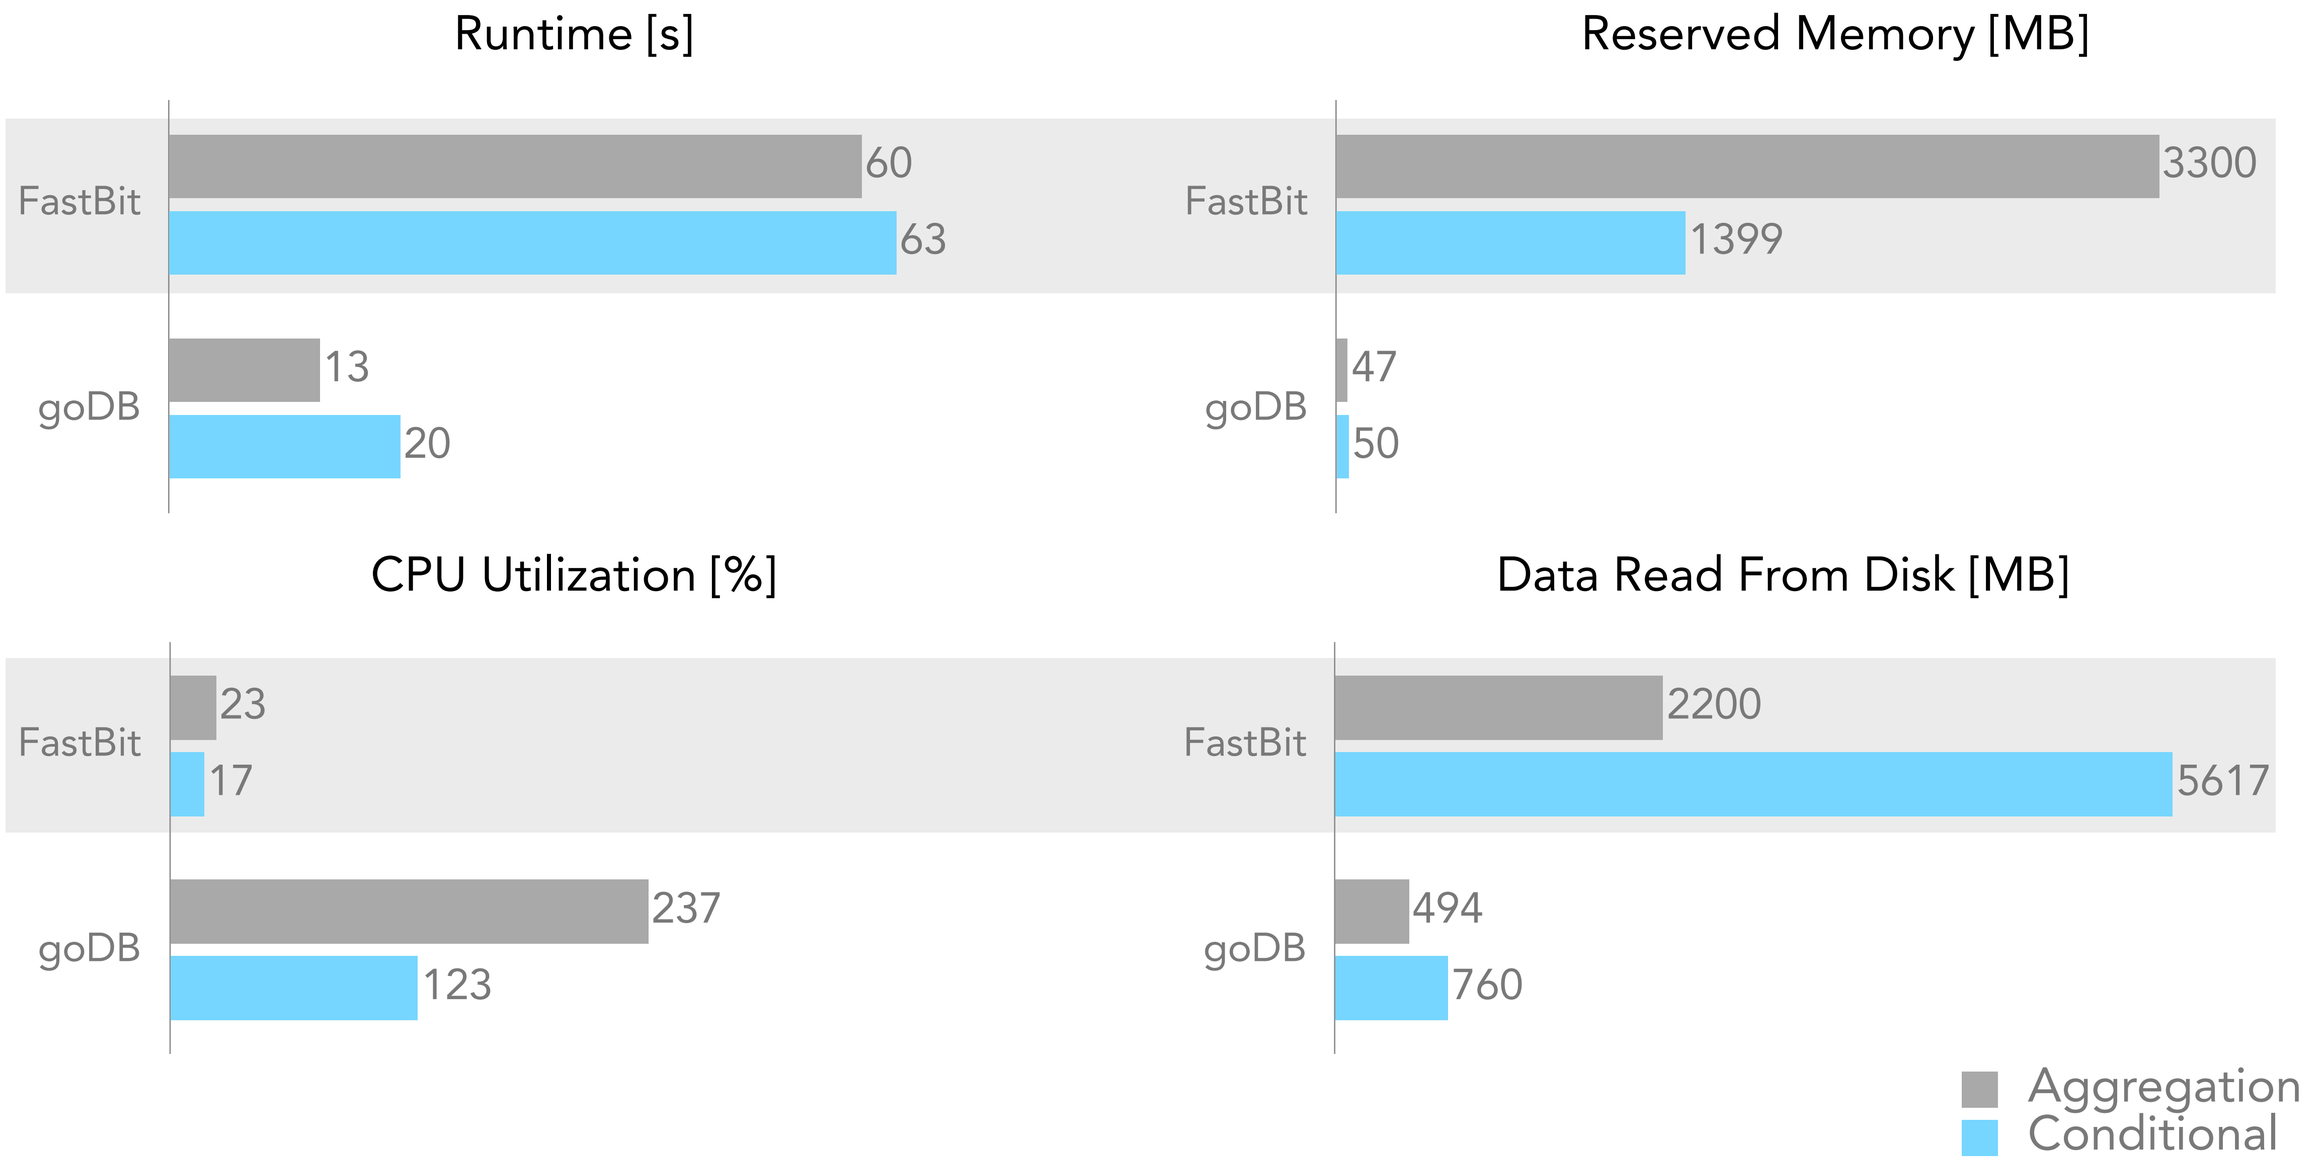

Traffic Portfolio of an NGO Customer

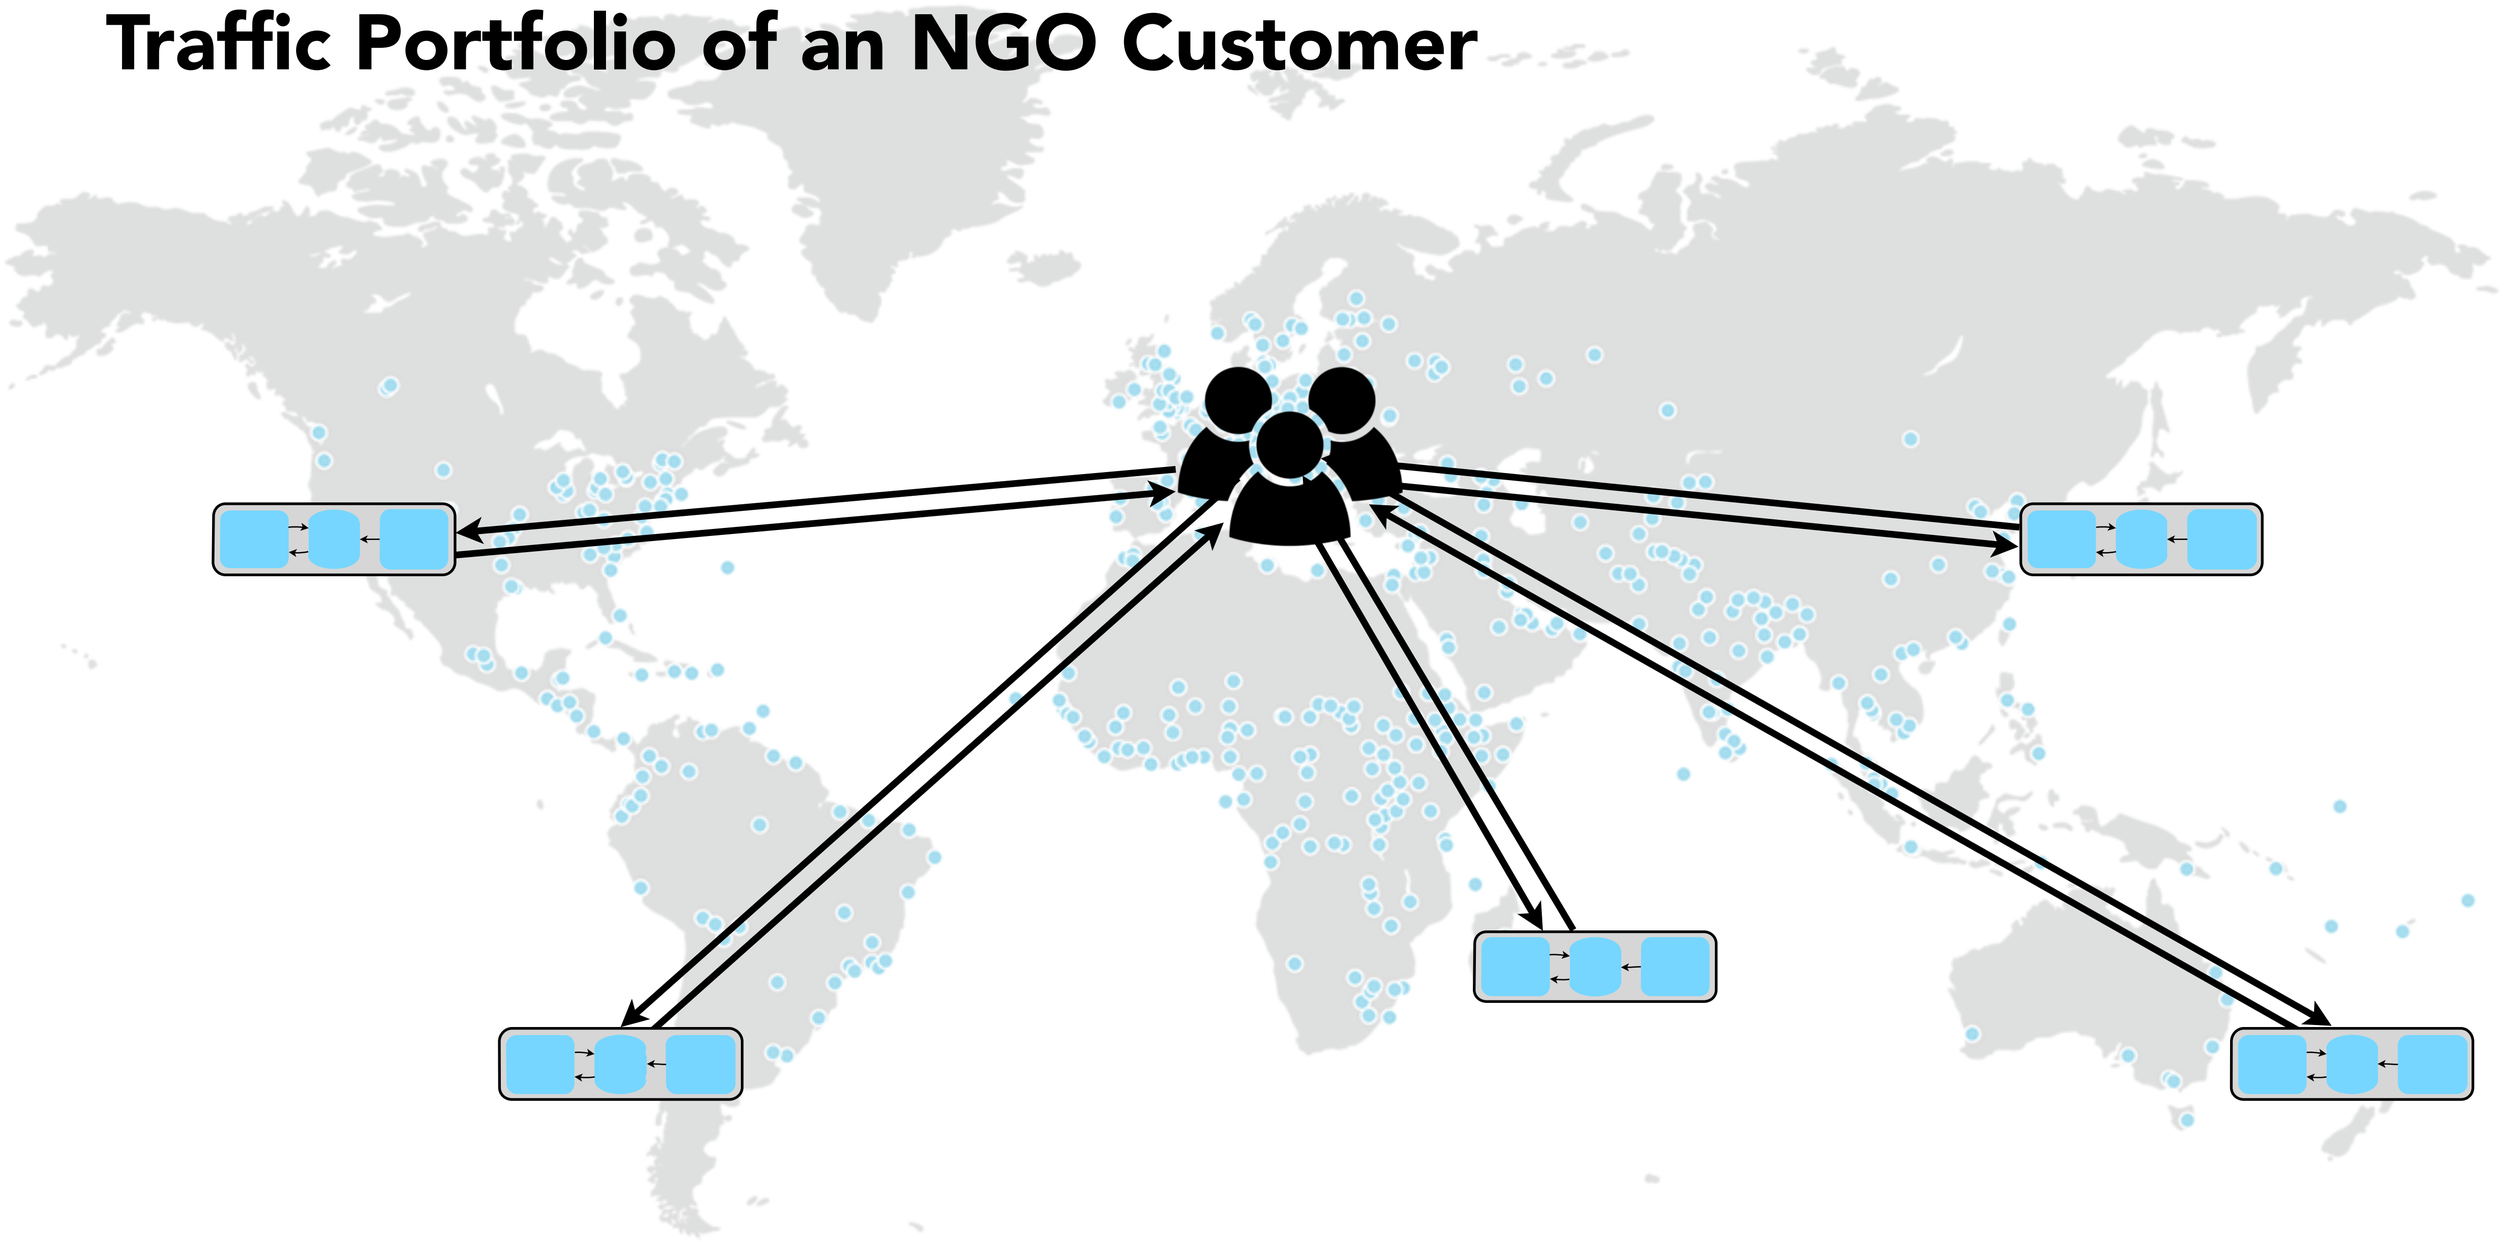

Global Breakdown of Ports

External Traffic

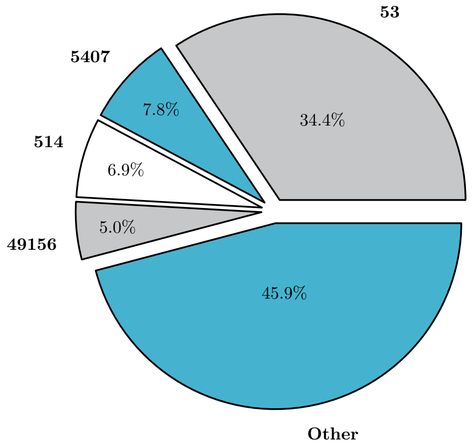

(a) **UDP:** 32.3 GB

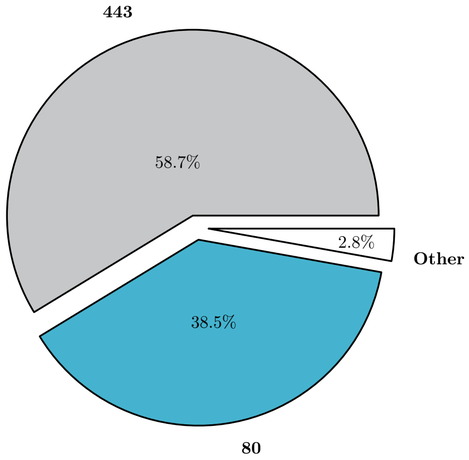

(b) **TCP:** 4.7 TB

Internal Traffic

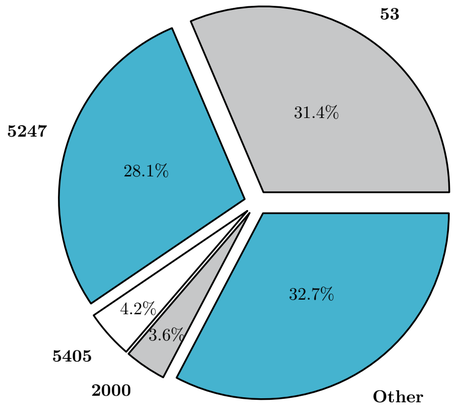

(c) **UDP:** 674.1 GB

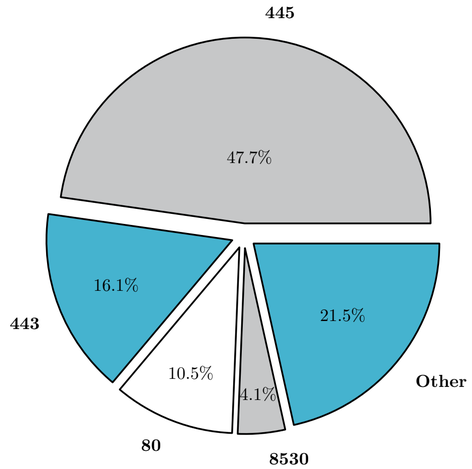

(d) **TCP:** 20.7 TB

Global Breakdown of Ports

External Traffic

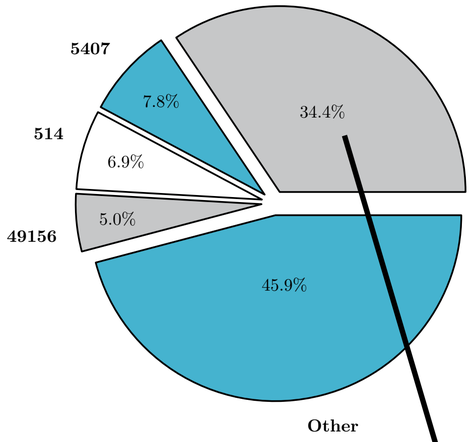

(a) **UDP:** 32.3 GB

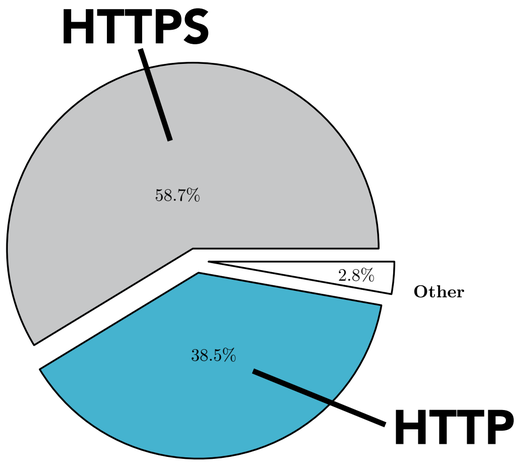

(b) **TCP:** 4.7 TB

Internal Traffic

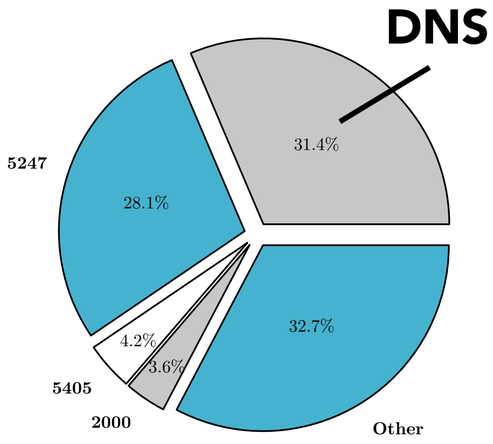

(c) **UDP:** 674.1 GB

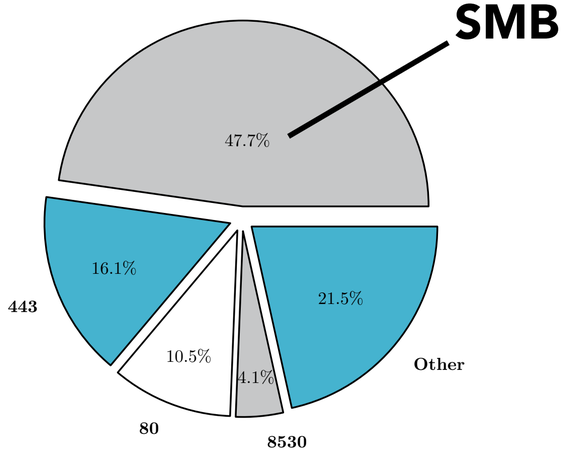

(d) **TCP:** 20.7 TB

Global Breakdown of Ports

External Traffic

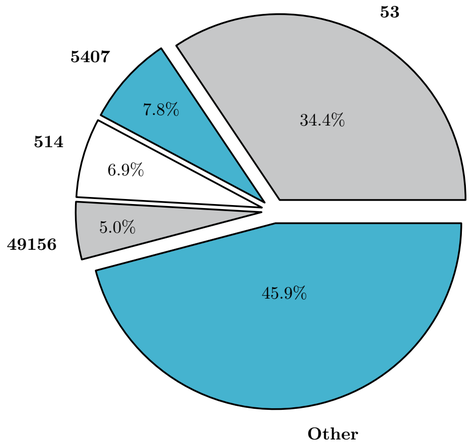

(a) **UDP:** 32.3 GB

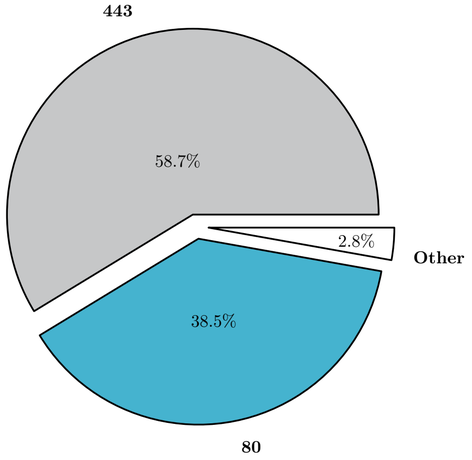

(b) **TCP:** 4.7 TB

Internal Traffic

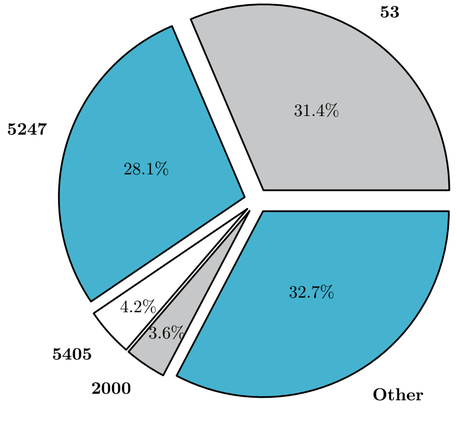

(c) **UDP:** 674.1 GB

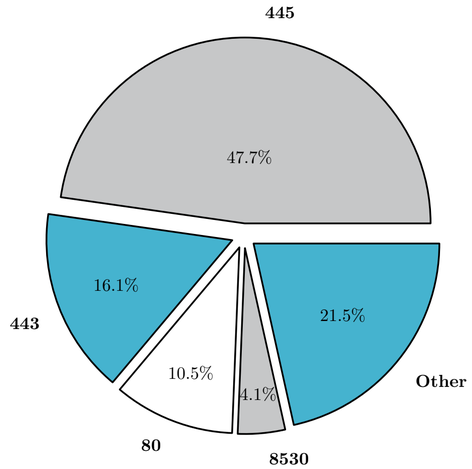

(d) **TCP:** 20.7 TB

European Hub Traffic Usage

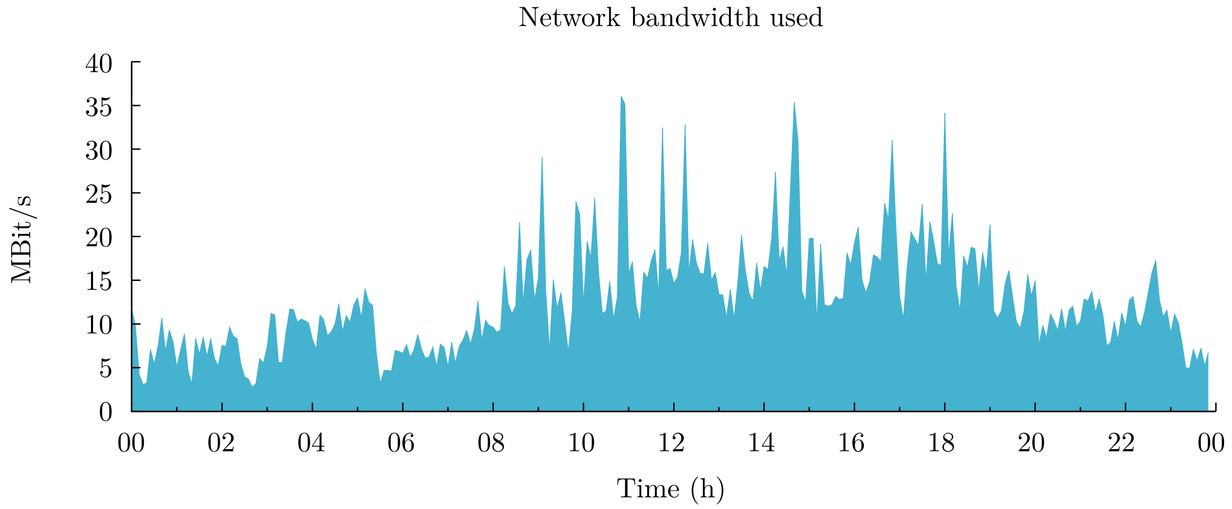

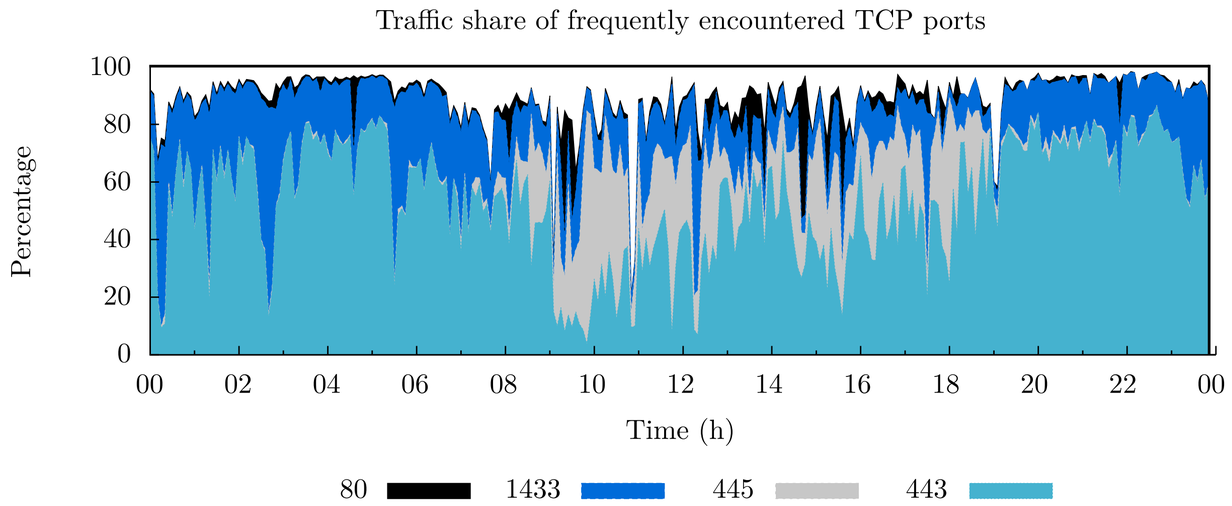

Conclusion

Improved capturing and flow logic

High performance DB written from scratch

Global deployment

Open source:

https://github.com/open-ch/