

CE 528 Cloud Computing



Lecture 16: Performance Profiling Spring 2026

Prof. Yigong Hu

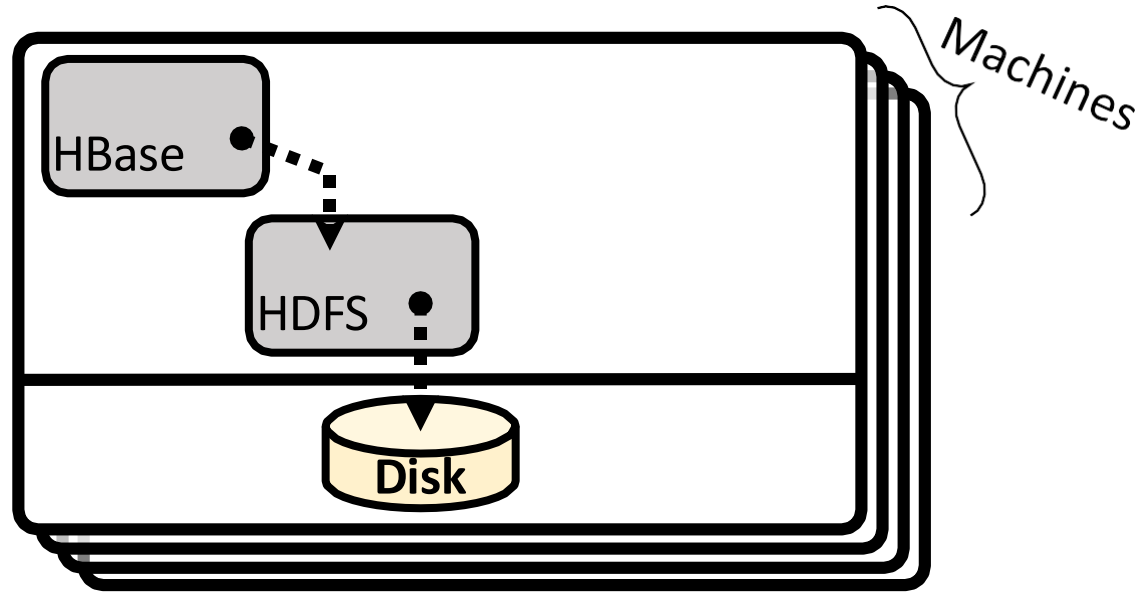



Slides courtesy of Michael Chow and Jonathan Mace



Case Study: Hadoop



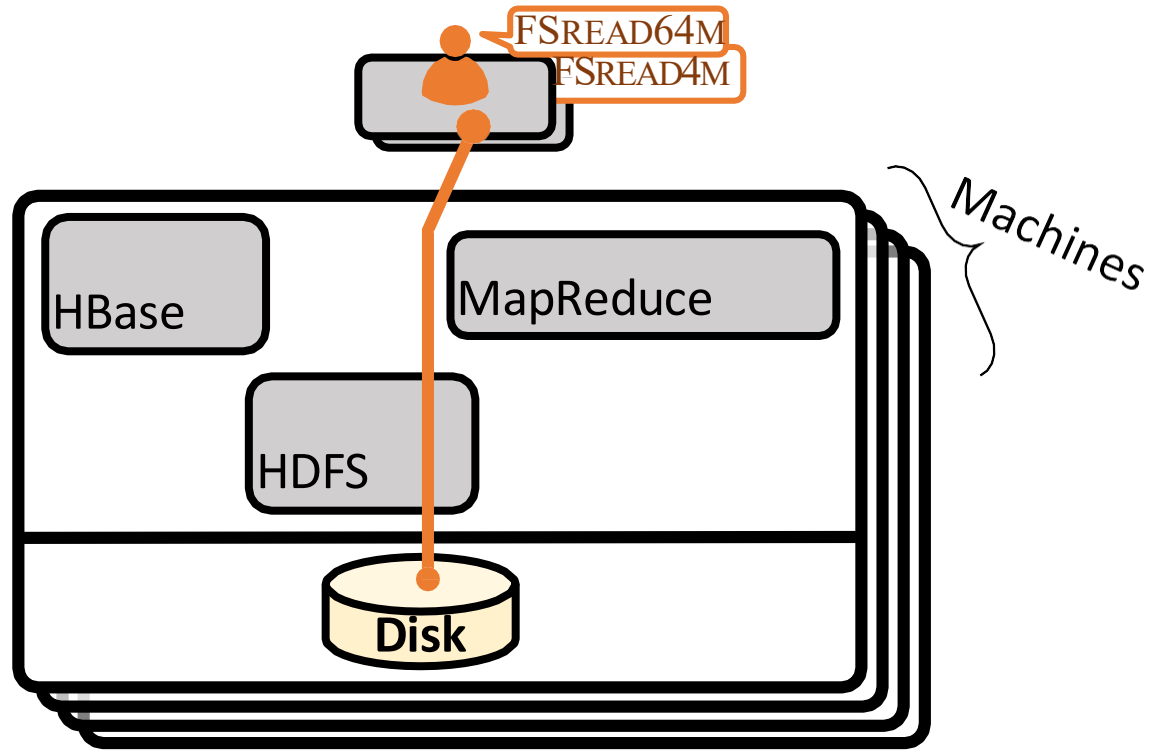

Case Study: Hadoop



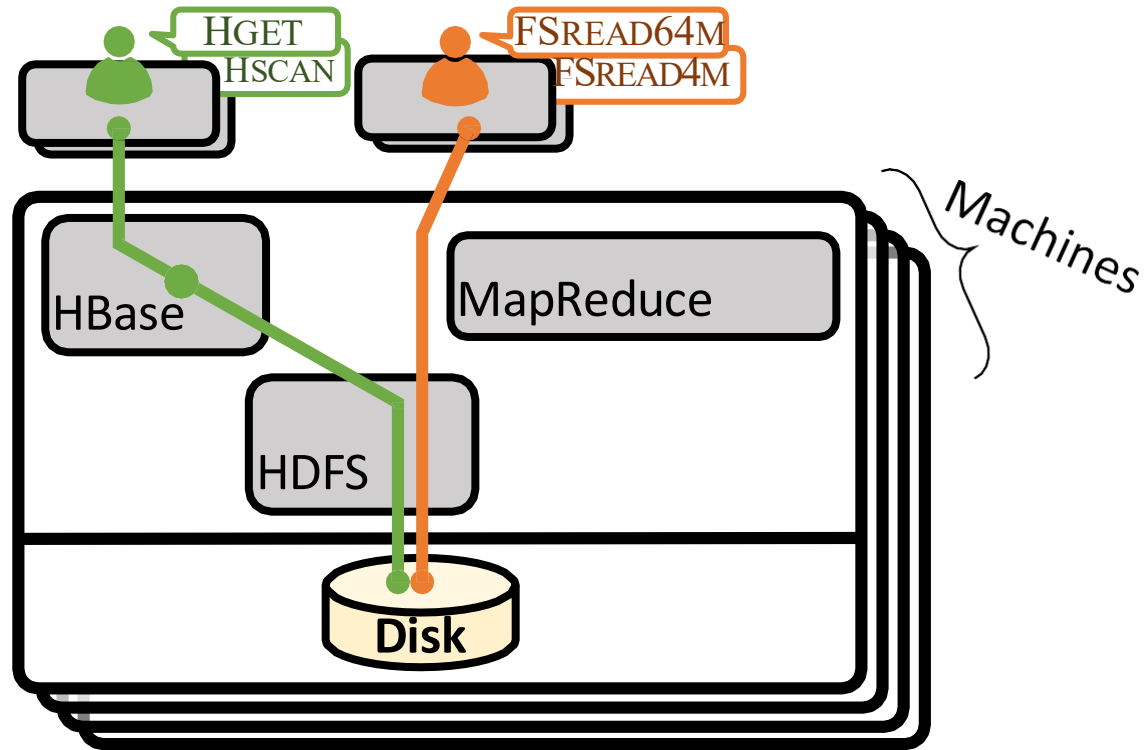

Case Study: Hadoop



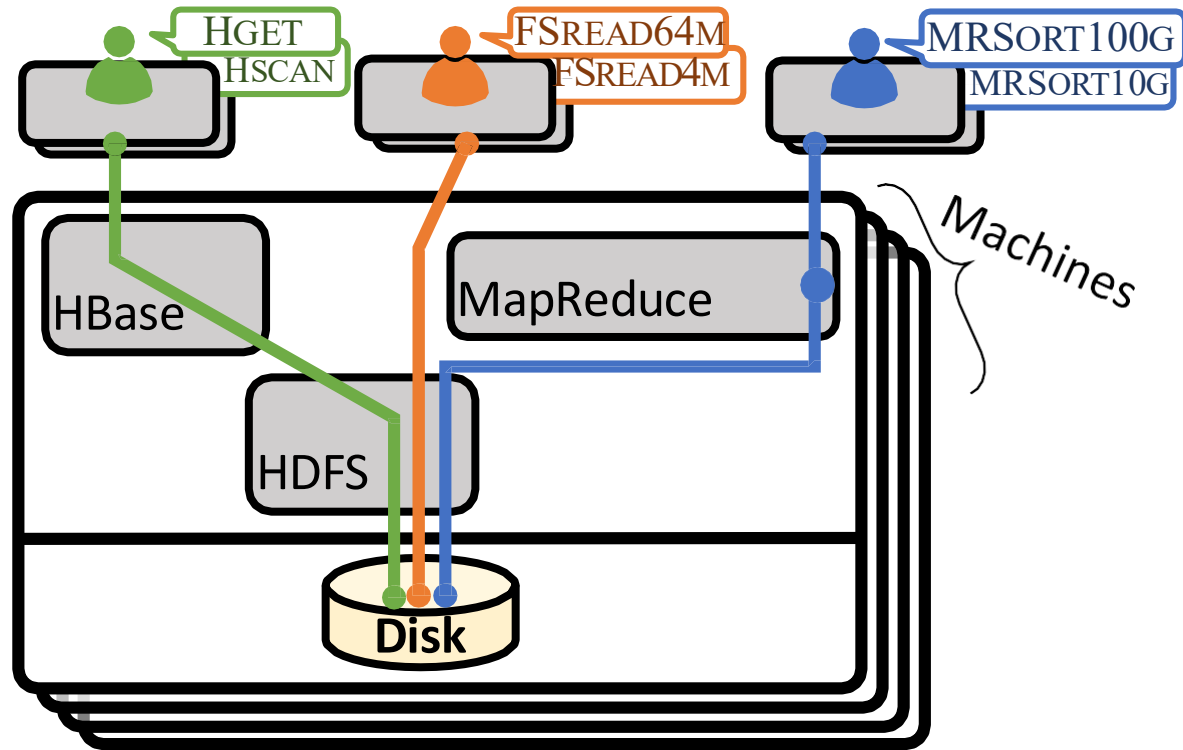

Case Study: Hadoop



Case Study: Hadoop

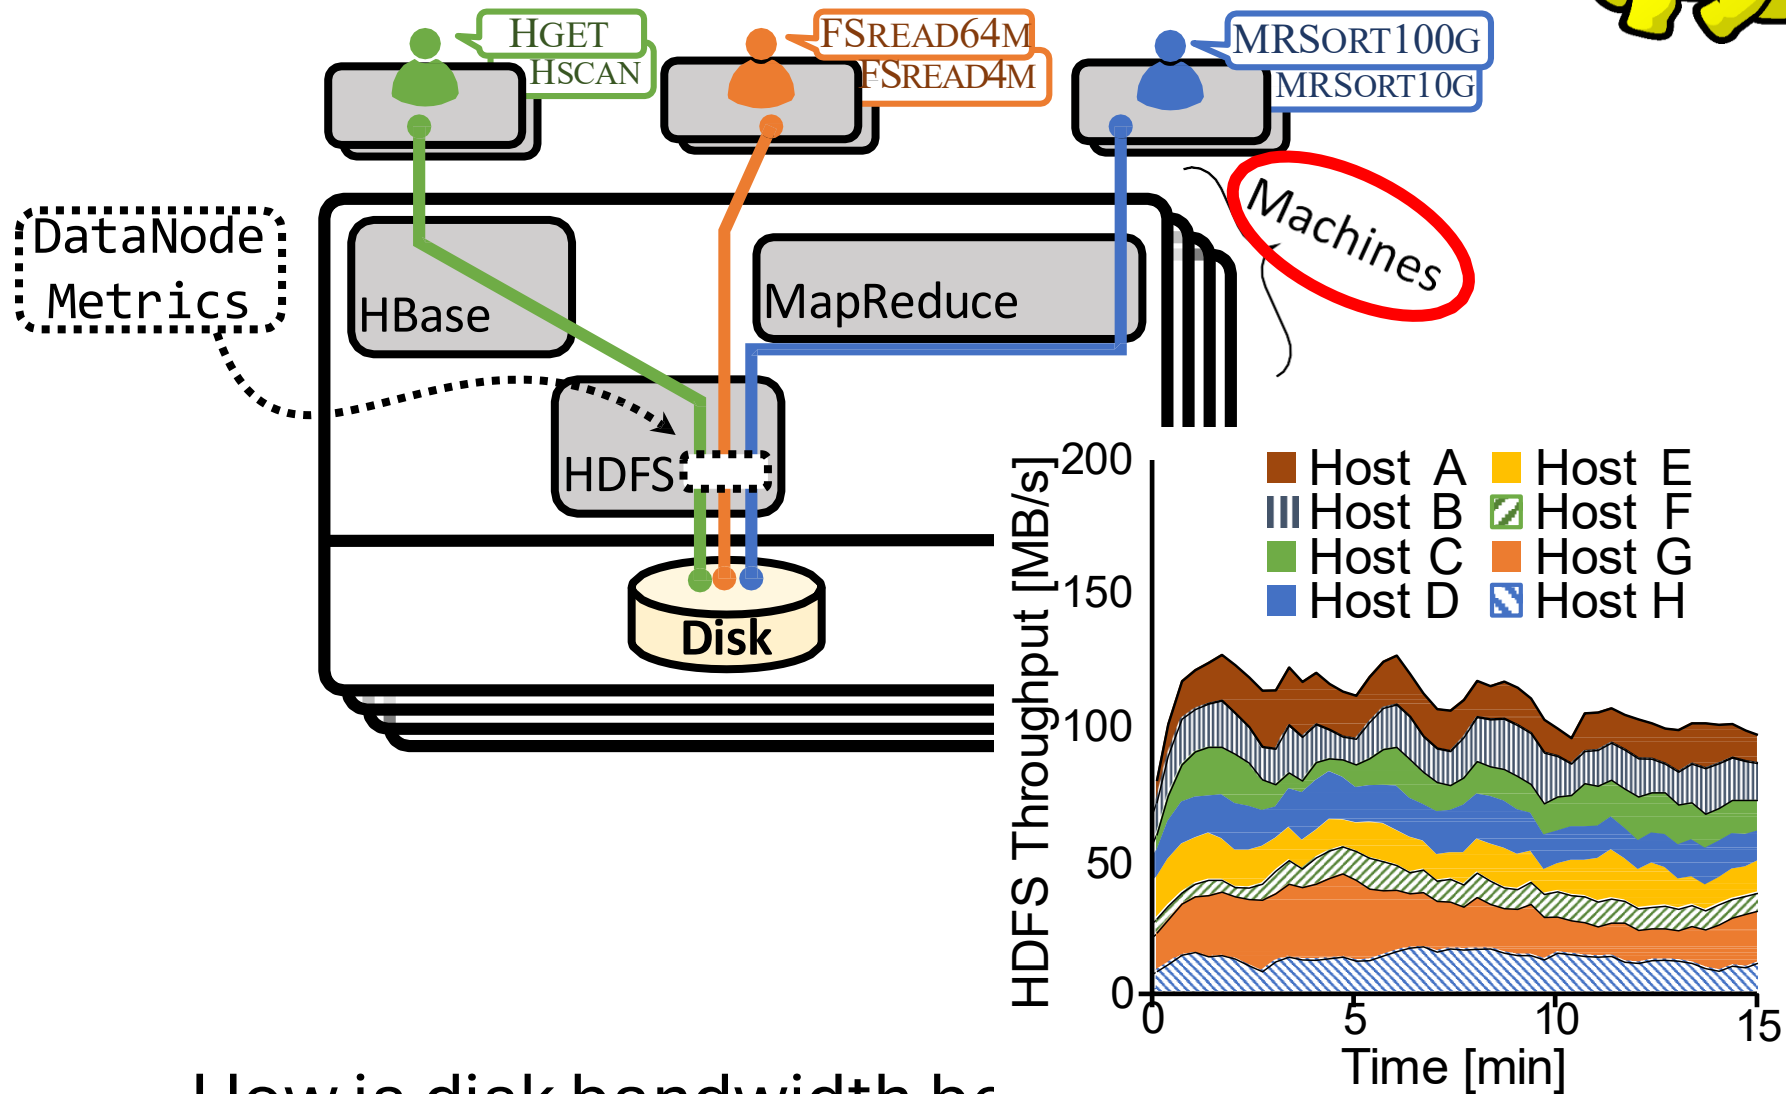



Case Study: Hadoop



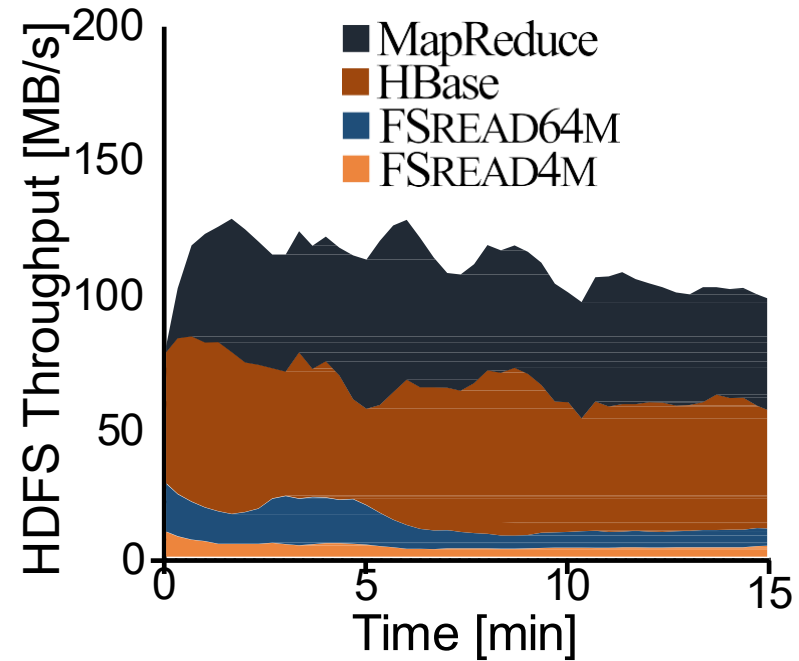
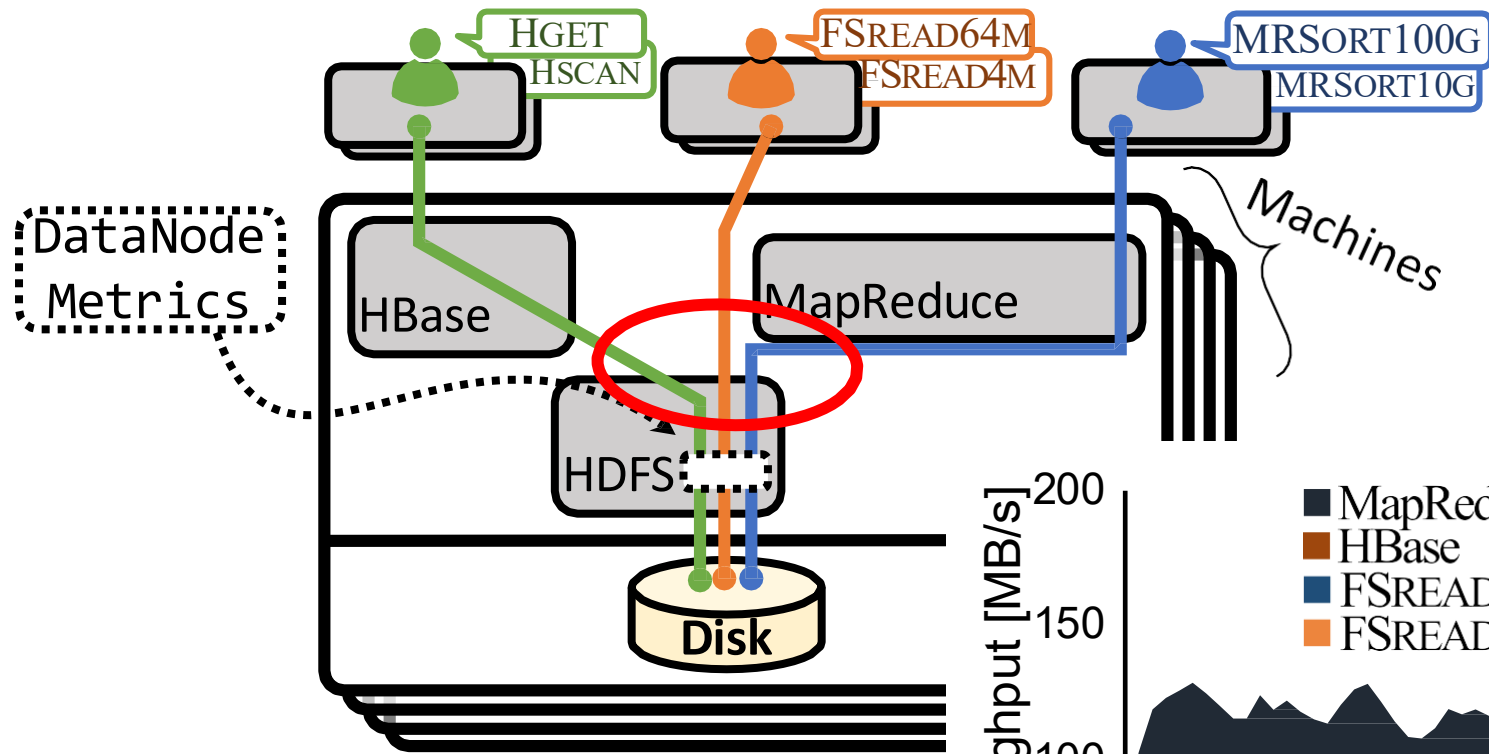

How is disk bandwidth being used?

Case Study: Hadoop


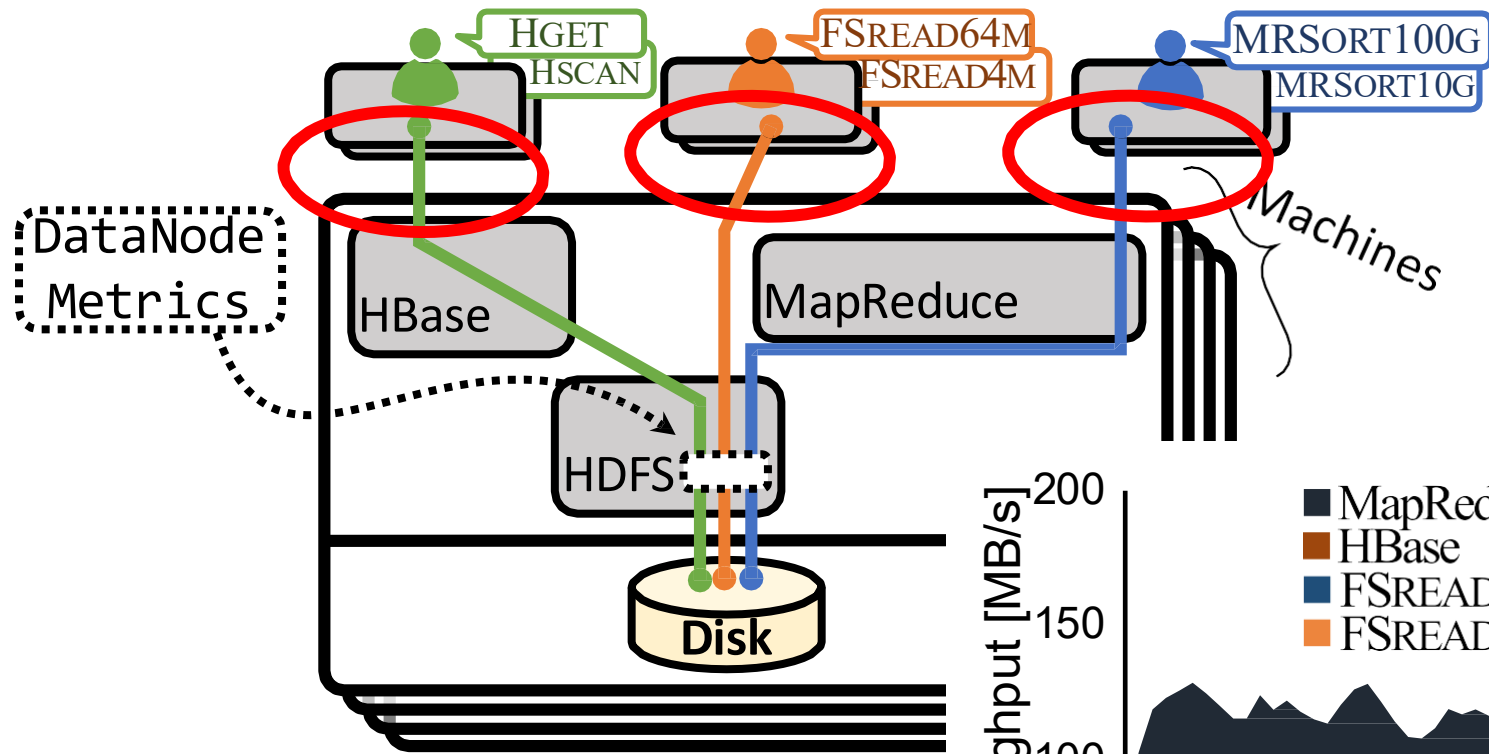



How is disk bandwidth being used?


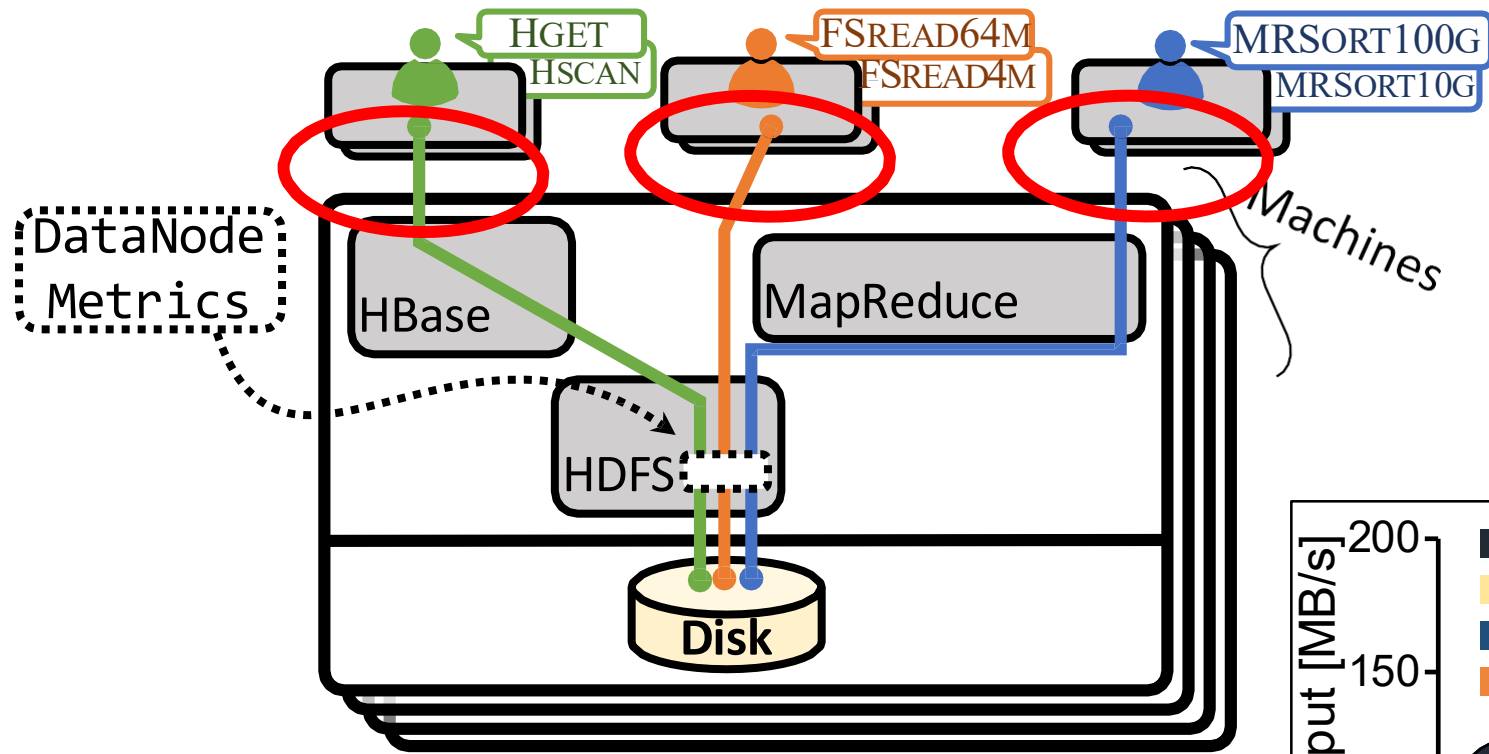

Case Study: Hadoop



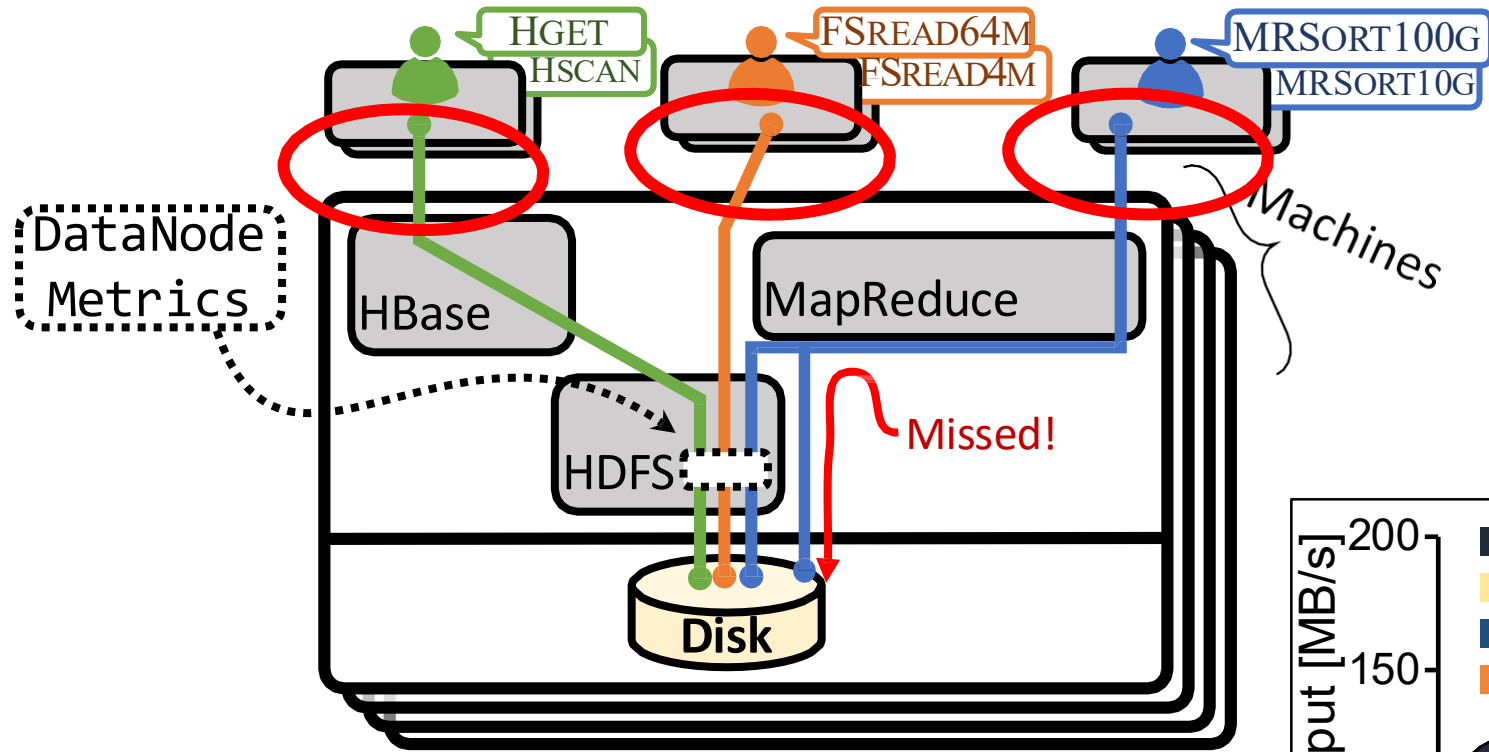

Case Study: Hadoop



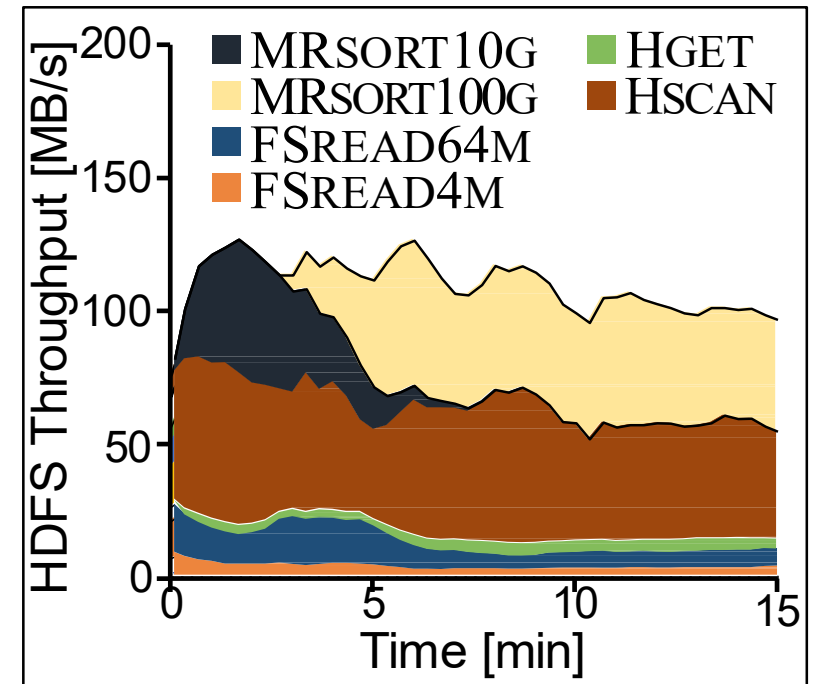
Case Study: Hadoop



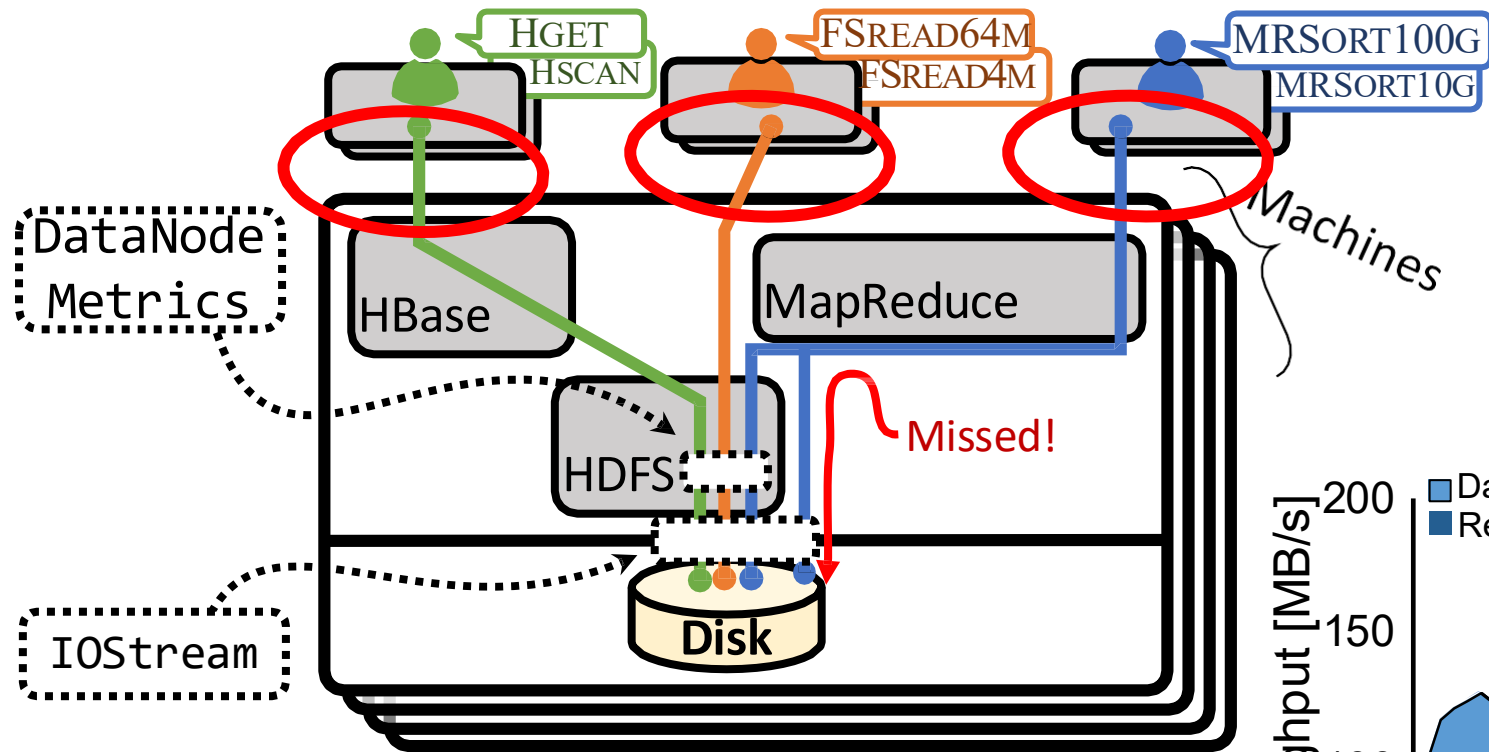

Case Study: Hadoop



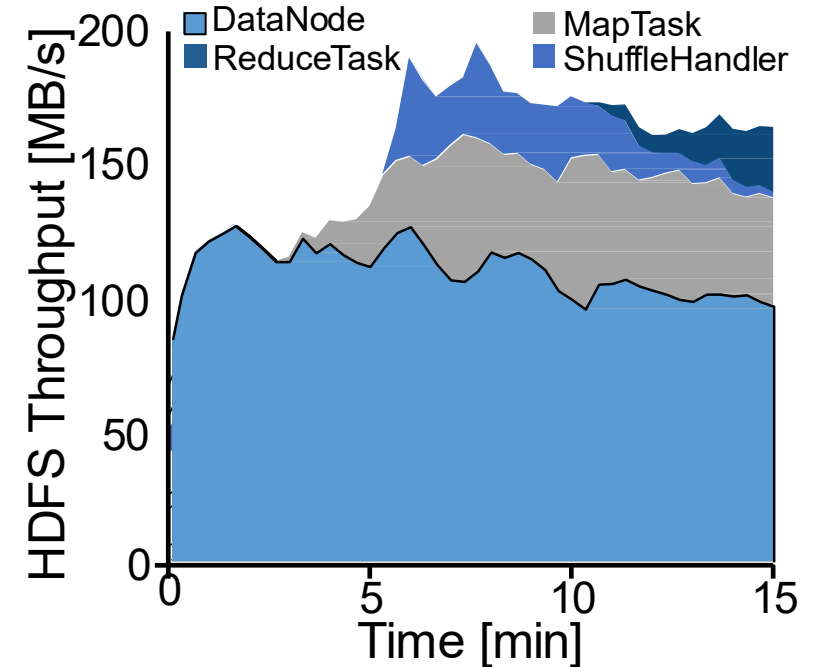
How is disk bandwidth being used?



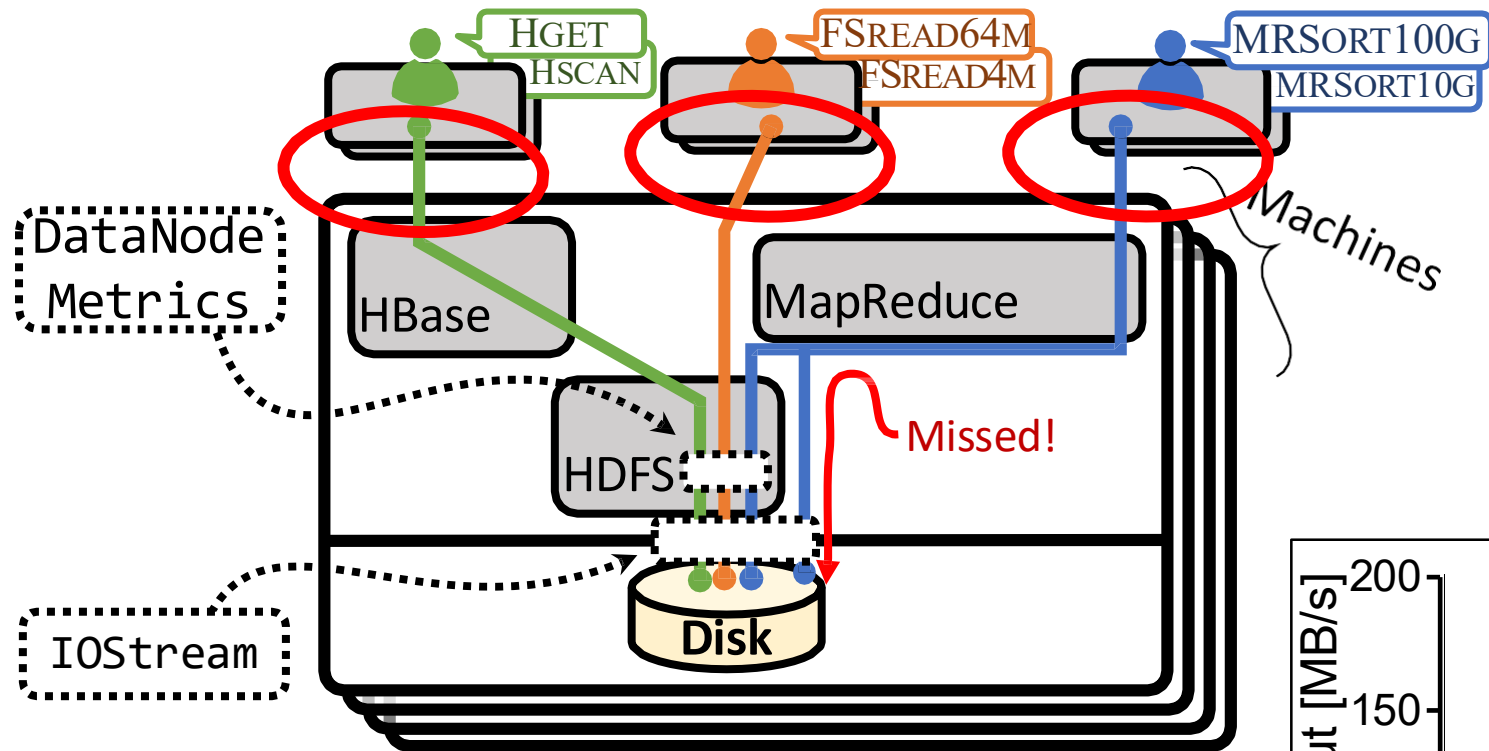

Case Study: Hadoop



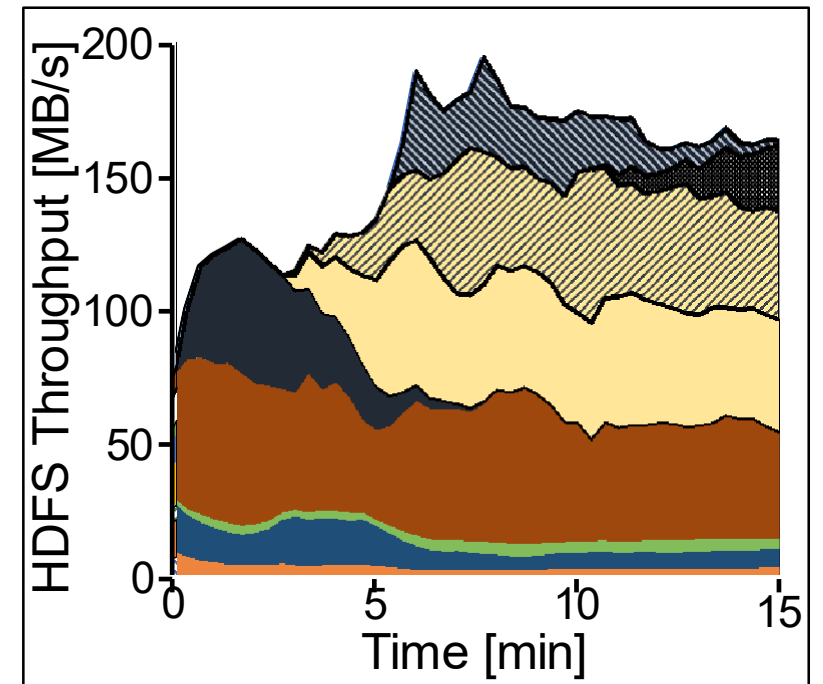
How is disk bandwidth being used?



Case Study: Hadoop




- MRSORT10G HGET
- MRSORT100G HSCAN
- FSREAD64M
- FSREAD4M
- SORT10G (ShuffleHandler)
- SORT10G (ReduceTask)
- SORT100G (MapTask)




Happens-Before

- Consider three processes: P1, P2, and P3
- Notation: Event a happens before event b ($a \rightarrow b$)




Happens-Before

1. If same process and a occurs before b, then $a \rightarrow b$




Happens-Before

1. If same process and a occurs before b, then $a \rightarrow b$
2. If c is a message receipt of b, then $b \rightarrow c$
3. If $a \rightarrow b$ and $b \rightarrow c$, then $a \rightarrow c$



Happens-Before

- Not all events are related by a
- a, d not related by a so concurrent, written as $a \parallel d$



Pivot Tracing: Dynamic Causal Monitoring for Distributed Systems

Jonathan Mace, Ryan Roelke, Rodrigo Fonseca


Brown University

Key insight of Pivot Tracing

- **Model system events as tuples in a streaming, distributed dataset**
- **Dynamically evaluate relational queries over this dataset**
- **Happened-before Join**

Pivot Tracing Design

Dynamic instrumentation



PT Agent



Causal tracing




Baggage

Pivot Tracing Workflow


Instrumented System (+  Baggage, +  PT Agent)



Pivot Tracing Workflow



Pivot Tracing Workflow



Places where PT can add instrumentation

Export identifiers accessible to queries

Defaults: host, timestamp, pid, proc name

Only references – not materialized until query is installed

Tracepoint A

Class: A

Method: A1()

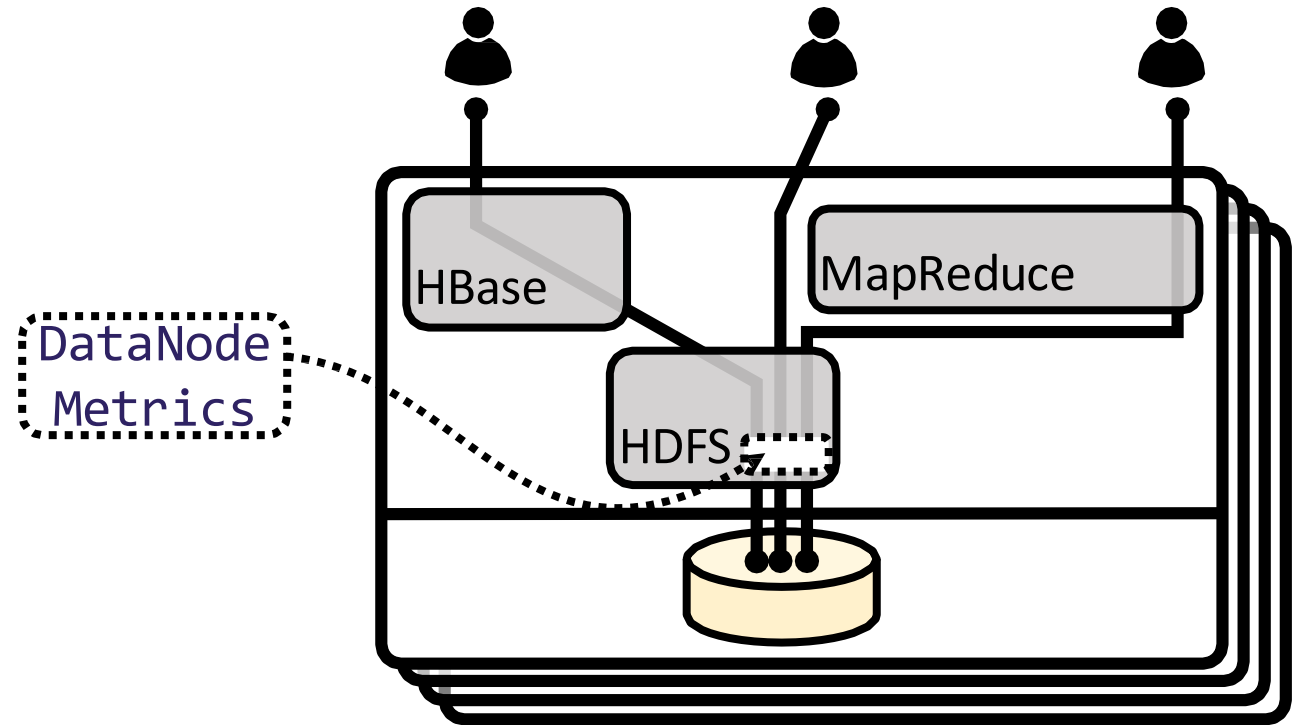
Tracepoint B

Class: B

Method: B1()

Exports: "delta"=delta

Tracing Point in Pivot Tracing



Tracing Point in Pivot Tracing

DataNodeMetrics.java

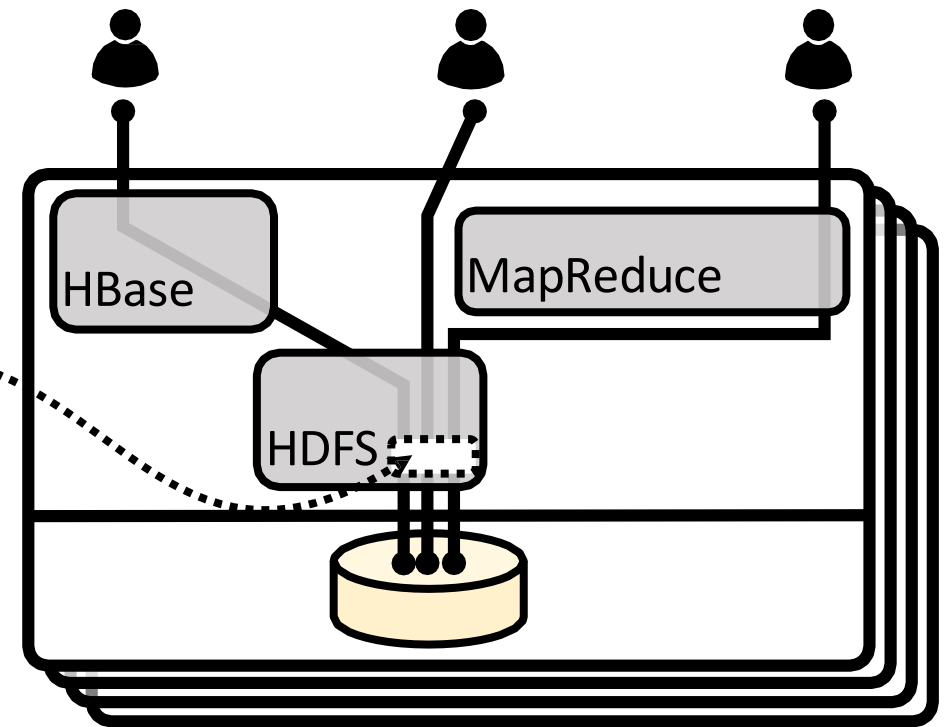
```
50 public class DataNodeMetrics {  
    ...  
266 public void incrBytesRead(int delta)  
    {  
267     ...  
268 }  
    ...  
407 }
```

Tracepoint


Method: incrBytesRead
Class: DataNodeMetrics
Exports: "delta"=delta

(“DataNodeMetrics”, delta=10,
host=“hop01”, ...)

DataNode
Metrics




Pivot Tracing Workflow



Relational query language, similar to SQL, LINQ


- Selection
- Projection
- Filter
- GroupBy
- Aggregation
- Happened-Before Join

Refers to trace point-exported identifiers



Pivot Tracing Workflow

Instrumented System (+  Baggage, +  PT Agent)



Query Language

Relational query language, similar to SQL, LINQ

- Selection
- Projection
- Filter
- GroupBy
- Aggregation
- Happened-Before Join

Refers to trace point-exported identifiers

Output: stream of tuples



```
From a In A
Join b In B On a -> b
GroupBy a.procName
Select a.procName, SUM(b.delta)
```

Tracepoint A
Class: A
Method: A1()

Tracepoint B
Class: B
Method: B1()
Exports: "delta"=delta

Query in Pivot Tracing


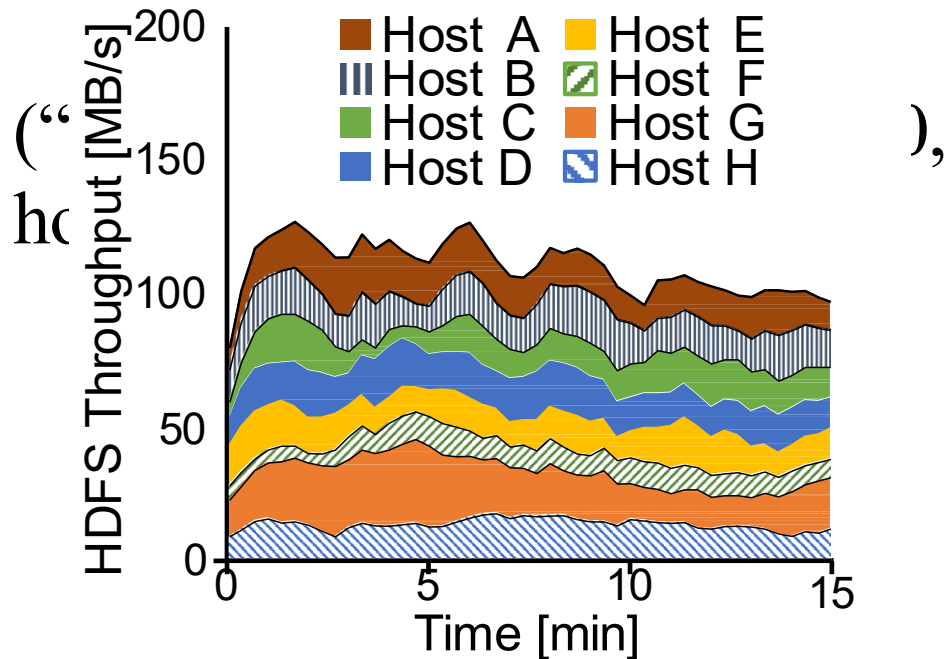
DataNodeMetrics.java

```
50 public class DataNodeMetrics {  
    ...  
266 public void incrBytesRead(int delta)  
    {  
267     ...  
268 }  
    ...  
407 }
```

Tracepoint



Method: incrBytesRead
Class: DataNodeMetrics
Exports: "delta"=delta

DataNode
Metrics


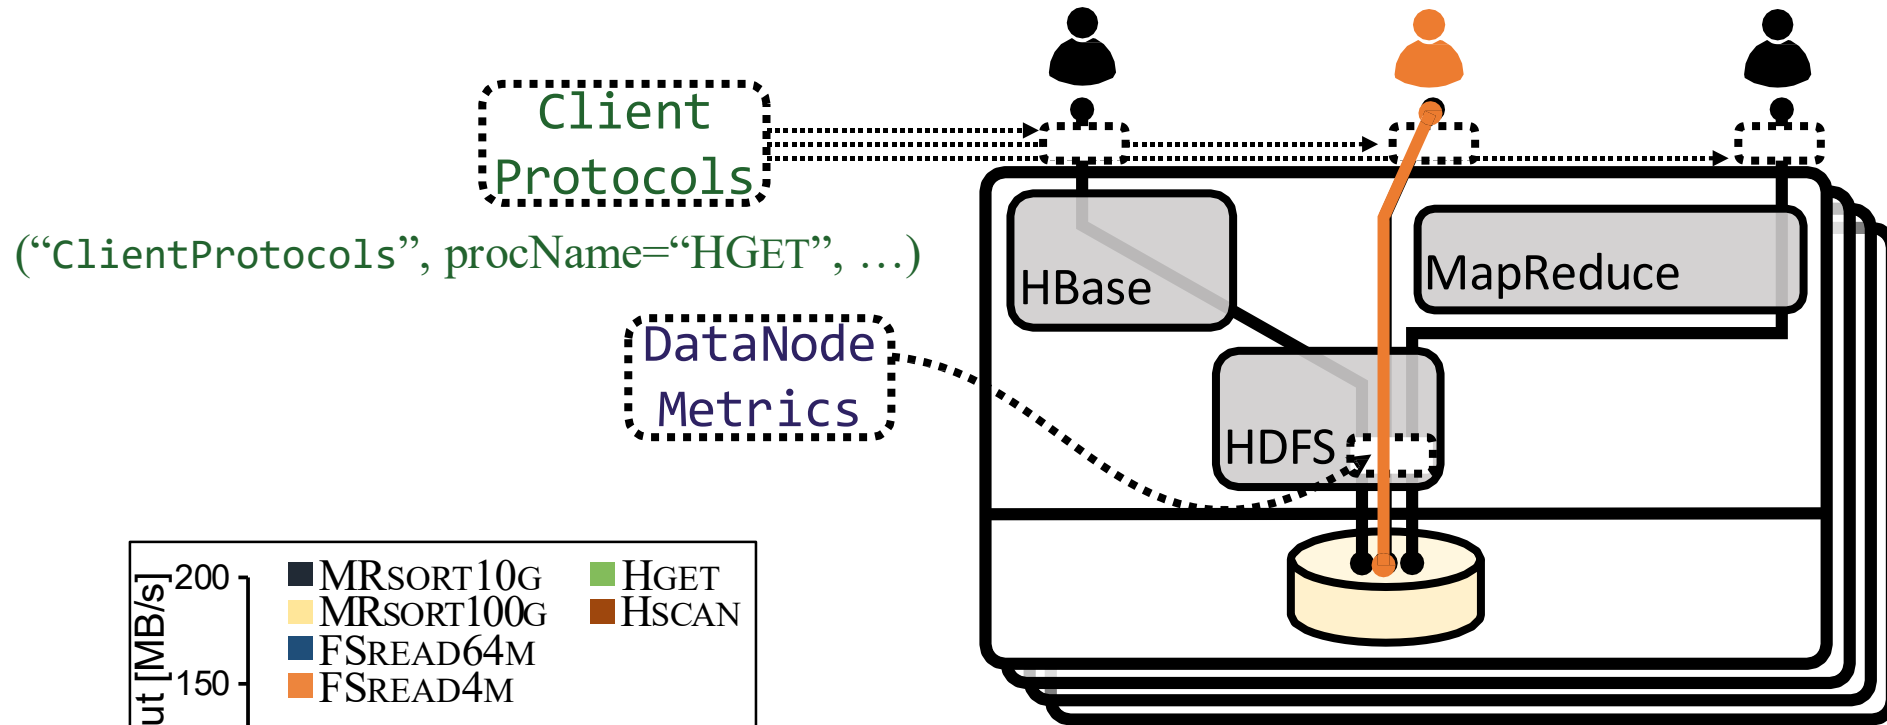


From incr In
DataNodeMetrics.incrBytesRead
GroupBy incr.host
Select incr.host, SUM(incr.delta)



Join Query



Join Query



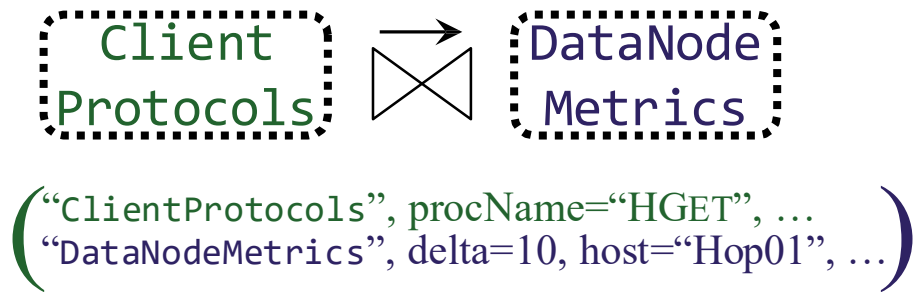


Join Query



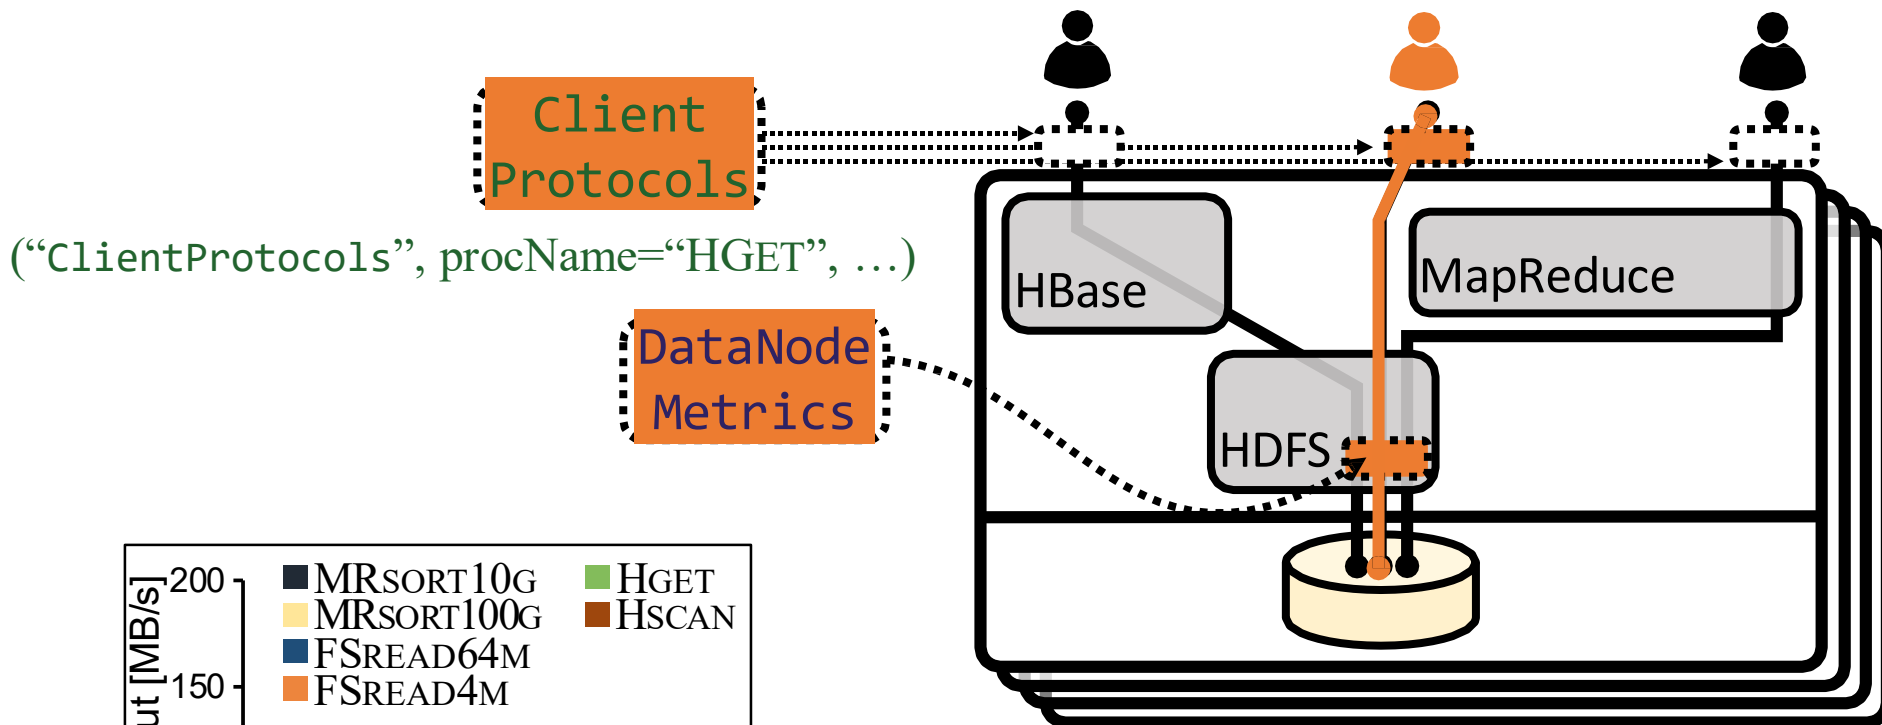
Client Protocols → DataNode Metrics

(`\"ClientProtocols\", procName=\"HGET\", ...`
`\"DataNodeMetrics\", delta=10, host=\"Hop01\", ...`)


Join Query



Join Query



(“ClientProtocols”, procName=“HGET”, ...)



From incr In


DataNodeMetrics.incrBytesRead

Join client In First(ClientProtocols) On
client -> incr

GroupBy client.procName

Select client.procName, SUM(incr.delta)

Instrumented System (+  Baggage, +  PT Agent)



Advice

Pivot Tracing
Front End

Query



```
From a In A  
Join b In B On a -> b  
GroupBy a.procName  
Select a.procName, SUM(b.delta)
```

Query is compiled to advice
(intermediate representation for instrumentation)

Advice will be installed at tracepoints


Limited instruction set

- OBSERVE
- PACK
- FILTER
- UNPACK
- EMIT

Advice A1 A
OBSERVE procName
PACK procName


Advice B1 B
OBSERVE delta
UNPACK procName
EMIT procName, SUM(delta)

Instrumented System (+  Baggage, +  PT Agent)



Weaving

PT Agent dynamically enables advice at tracepoints

 **Query**

```
From a In A
Join b In B On a -> b
GroupBy a.procName
Select a.procName, SUM(b.delta)
```

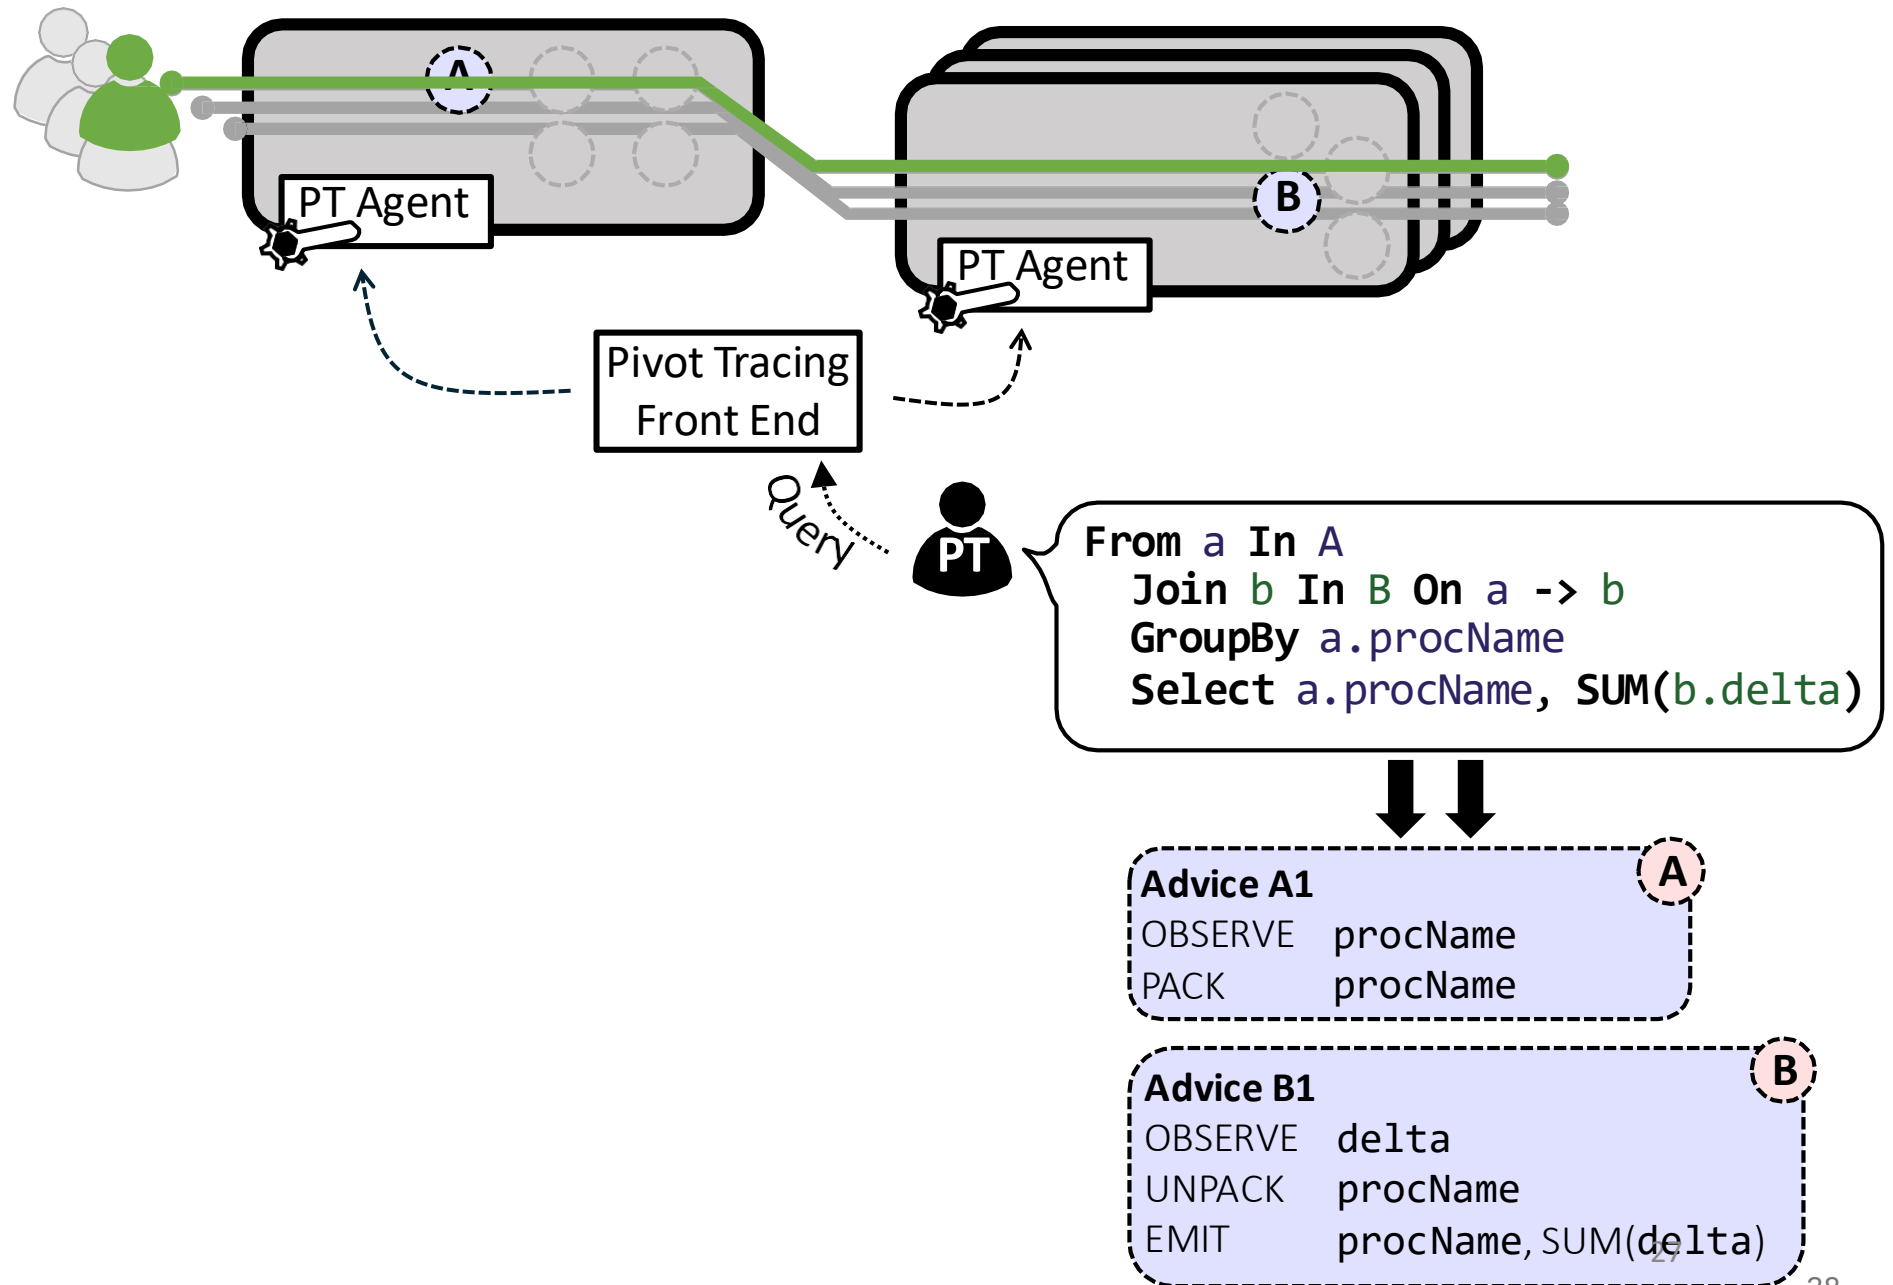
Advice A1 A

```
OBSERVE  procName
PACK     procName
```


Advice B1 B


```
OBSERVE  delta
UNPACK   procName
EMIT     procName, SUM(delta)
```

Instrumented System (+  Baggage, +  PT Agent)





Baggage



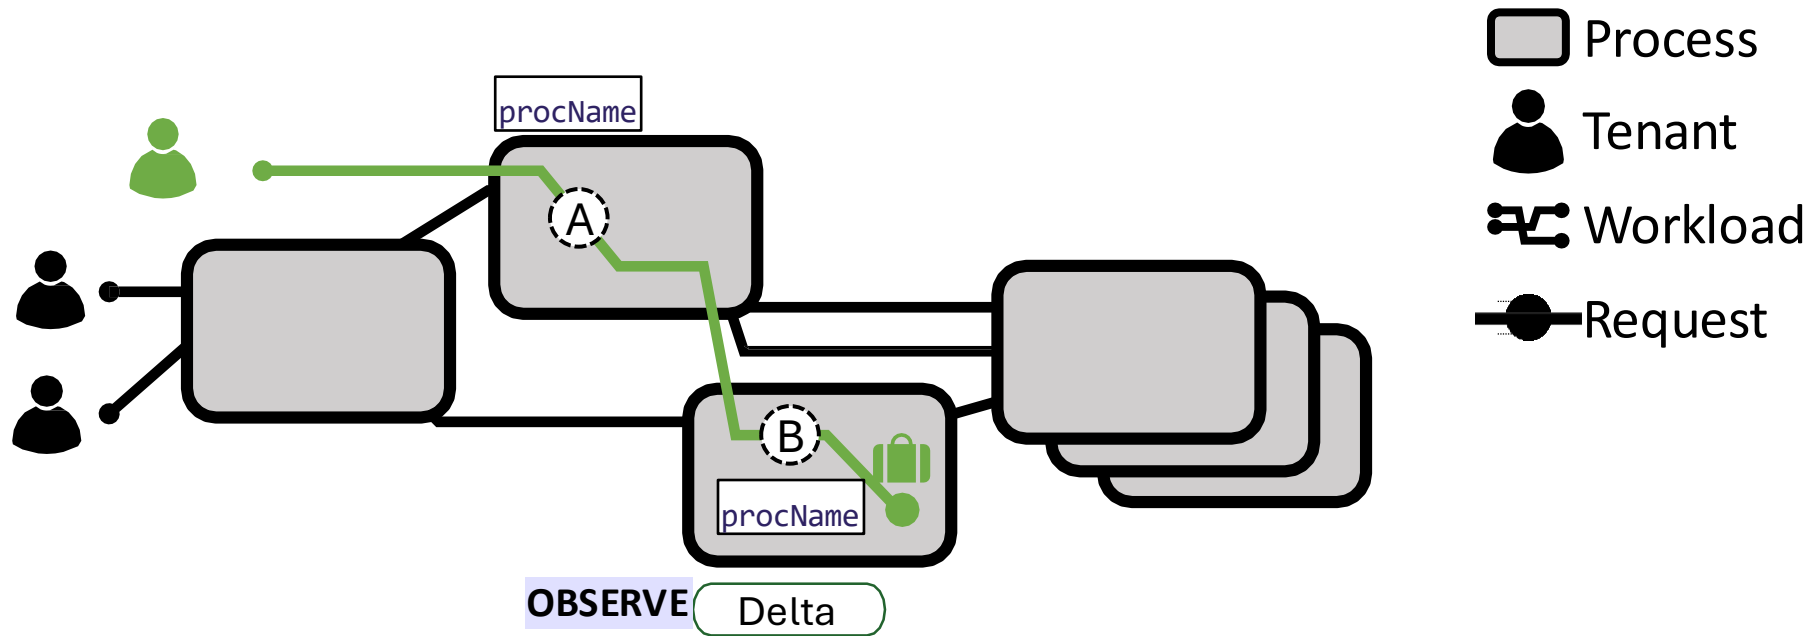
-  Baggage is a Key: Value container propagated alongside a request
- Generalization of metadata in end-to-end tracing
- One instance per request

Baggage



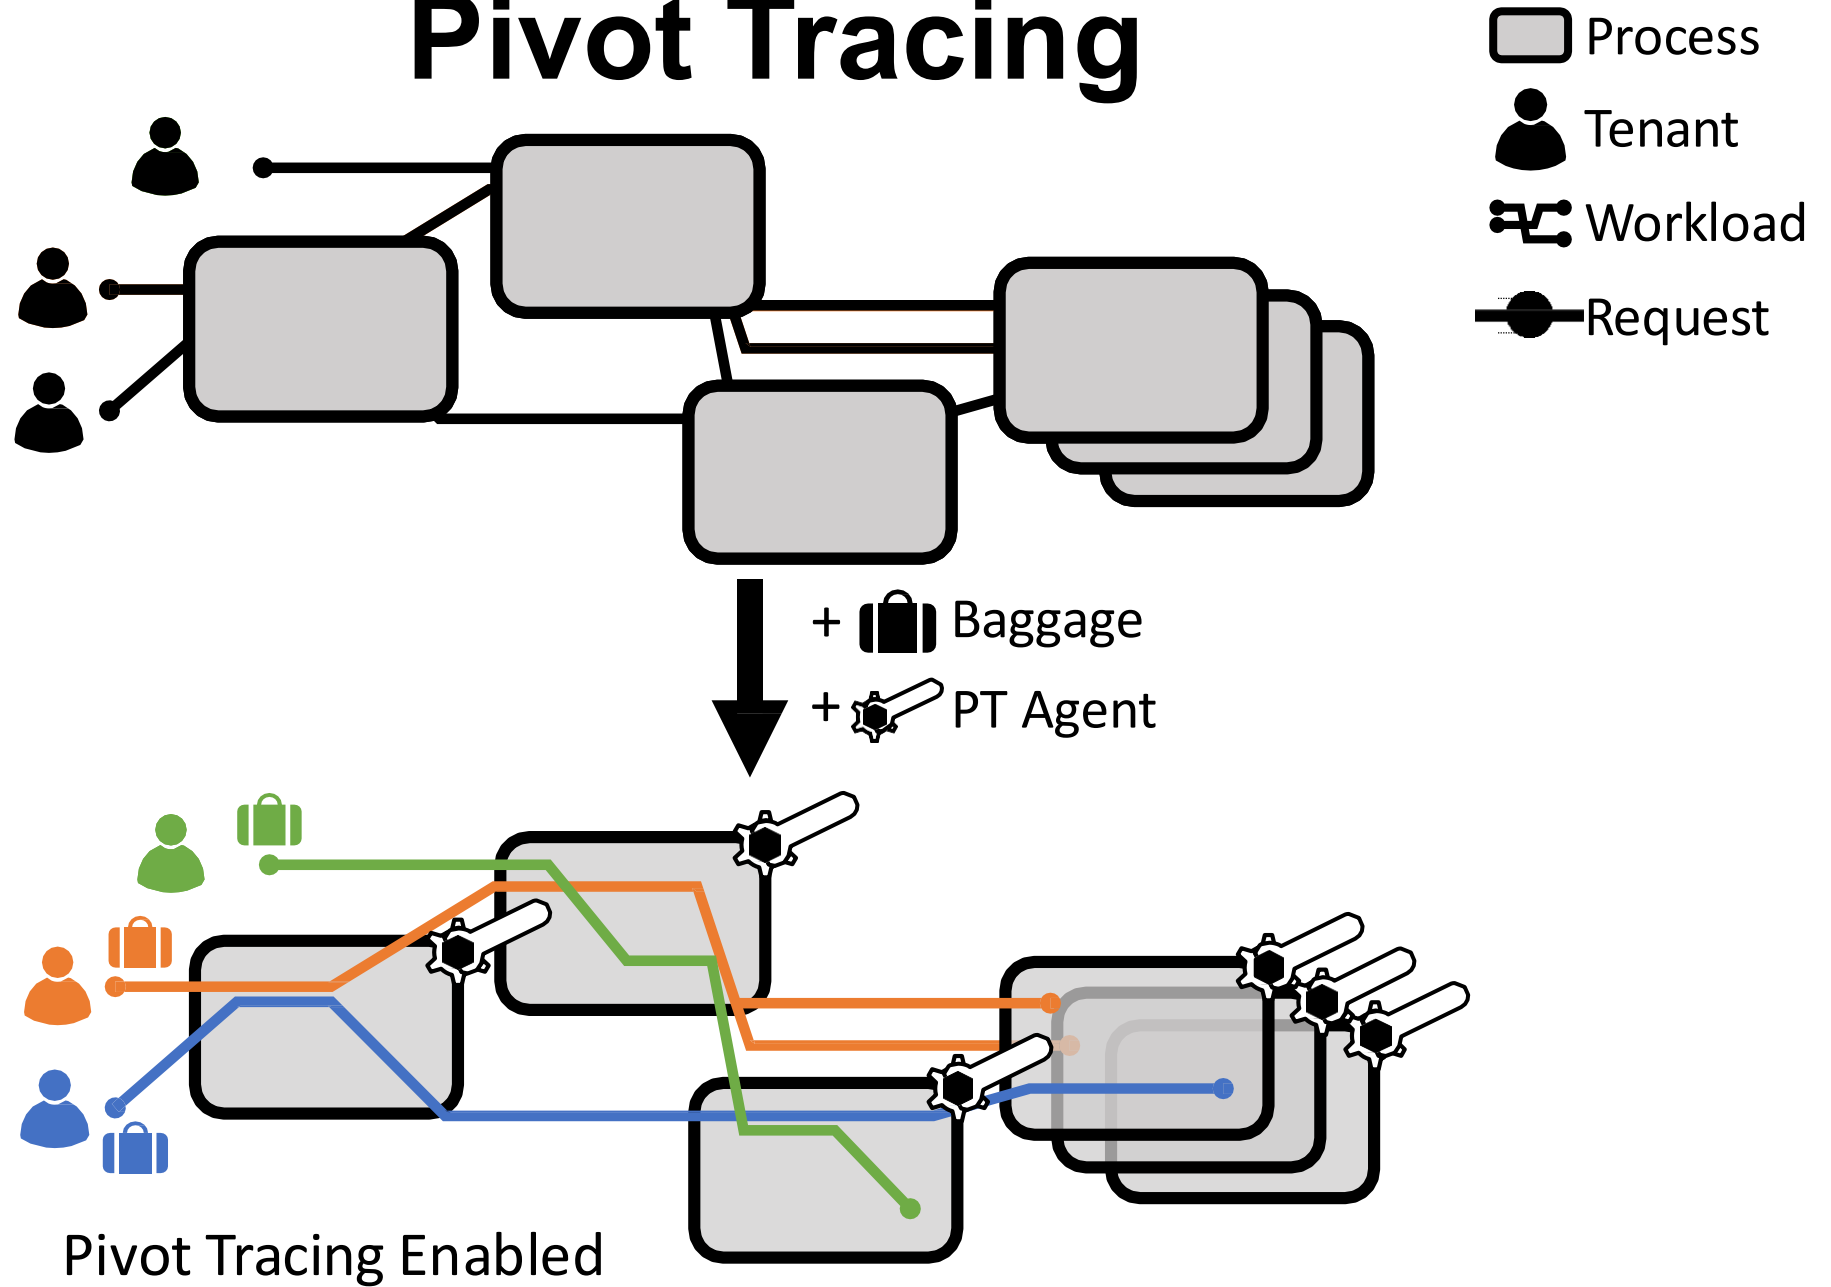
-  Baggage is a Key: Value container propagated alongside a request
- Generalization of metadata in end-to-end tracing
- One instance per request

Baggage



- Baggage is a Key: Value container propagated alongside a request
 - Generalization of metadata in end-to-end tracing
 - One instance per request

Pivot Tracing




The Mystery Machine: End-to-end performance analysis of large-scale Internet services

Michael Chow

David Meisner, Jason Flinn, Daniel Peek,
Thomas F. Wenisch

OSDI'14

Internet Services Are Complex



Scale and heterogeneity make Internet services complex

Challenges

Previous methods derive a causal model

- Instrument scheduler and communication
- Build model through human knowledge

Need method that works at scale with
heterogeneous components

Opportunities

Component-level logging is ubiquitous

- Tremendous detail about a request's execution

Handle a large number of requests

- Coverage of a large range of behaviors

The Mystery Machine

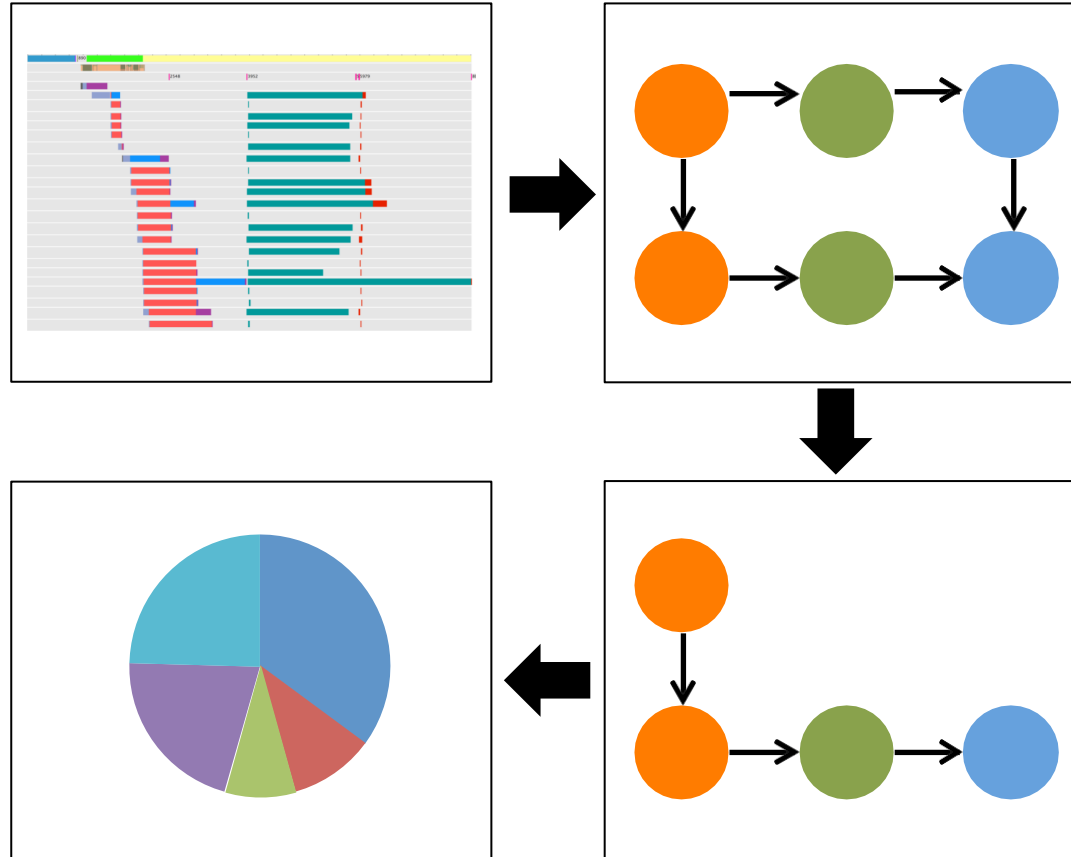
1) Infer causal model from large corpus of traces

- Identify segments
- Hypothesize all possible causal relationships
- Reject hypotheses with observed counterexamples

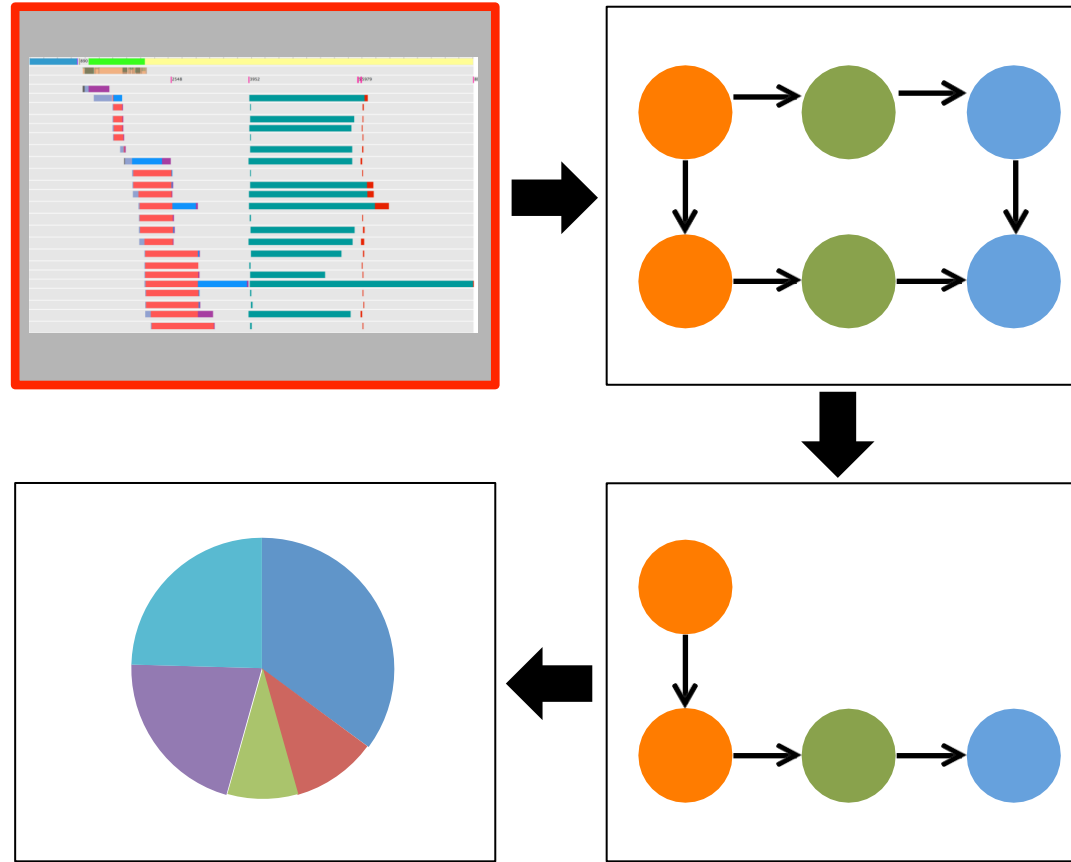
2) Analysis

- Critical path, slack, anomaly detection, what-if

Analysis Pipeline




Step 1: Identify Segments



Define a Minimal Schema

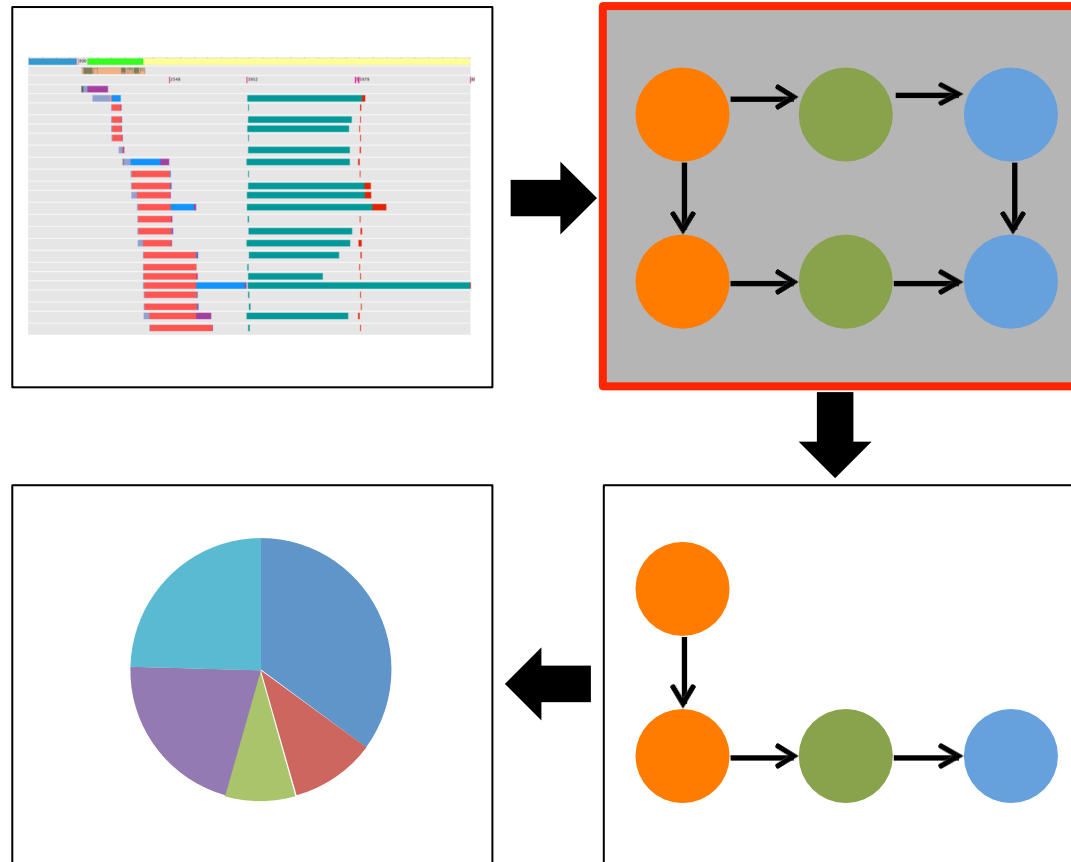
Task



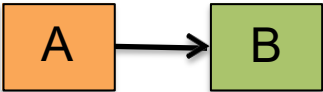
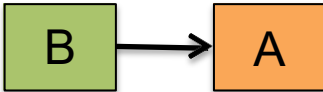
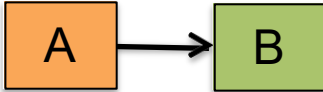
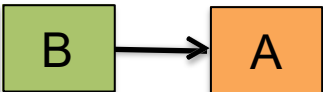
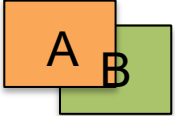
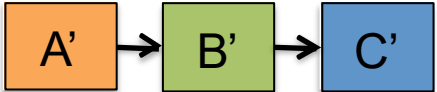
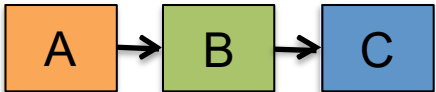
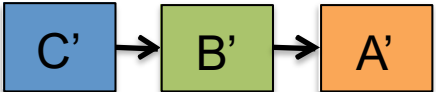
Request identifier
Machine identifier
Timestamp
Task
Event

Aggregate existing logs using minimal schema

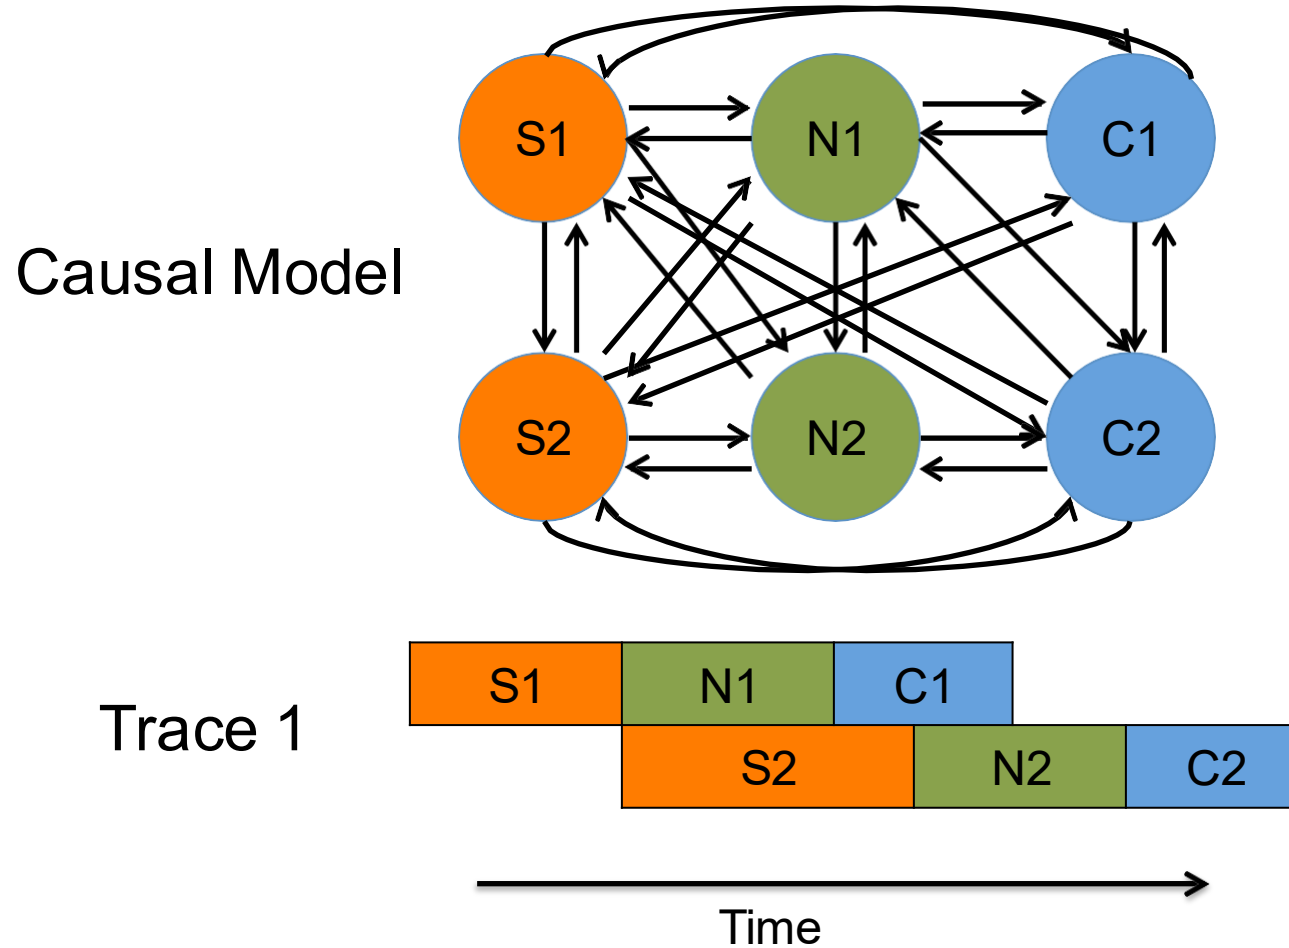
Step 2: Infer Causal Model



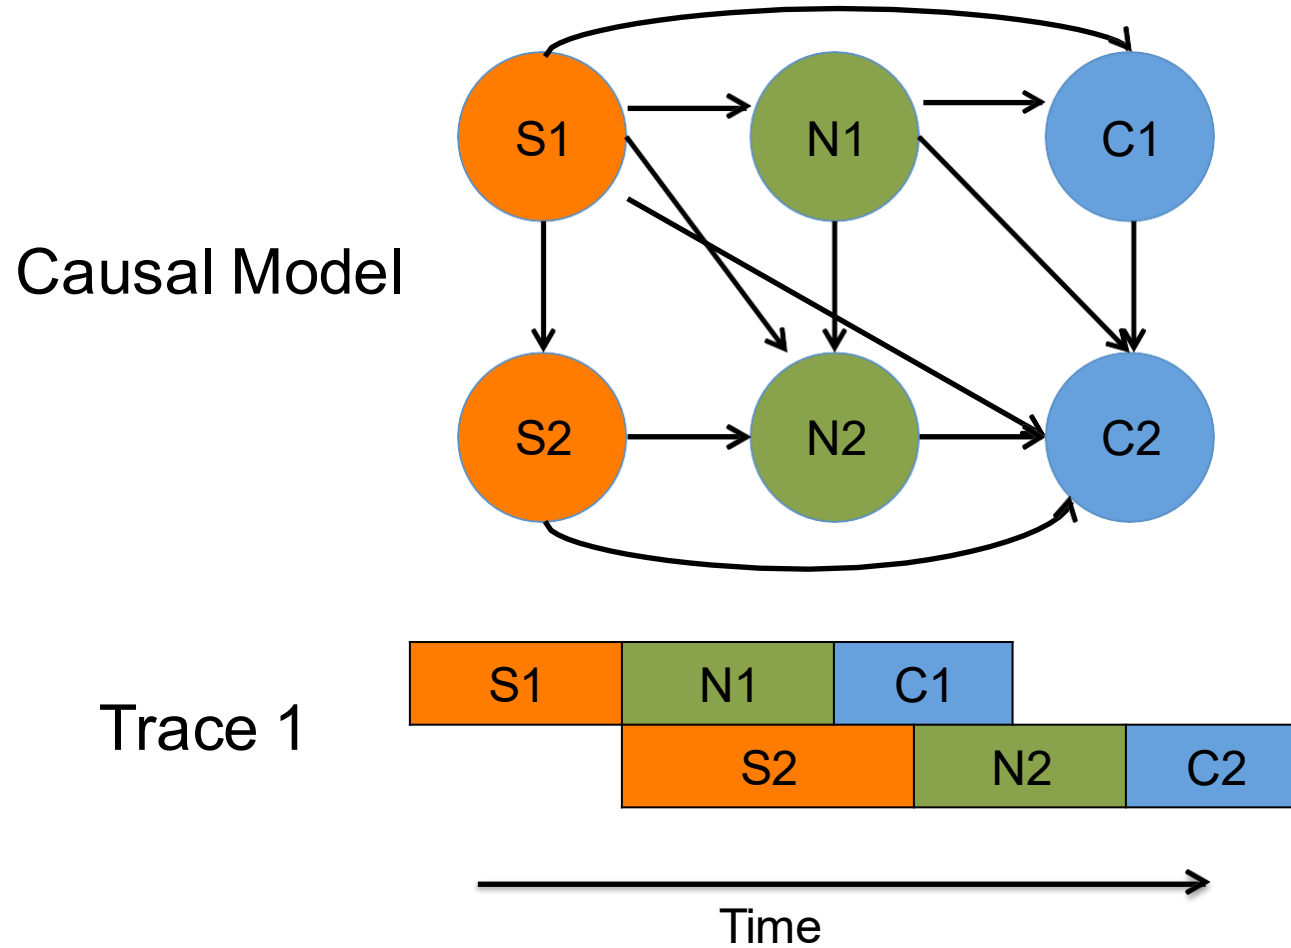
Types of Causal Relationships

Relationship	Counterexample
Happens-Before  A → B	 B → A
Mutual Exclusion  A → B OR  B → A	 A B
Pipeline t ₁  A → B → C t ₂  A' → B' → C'	t ₁  A → B → C t ₂  C' → B' → A'


Producing Causal Model



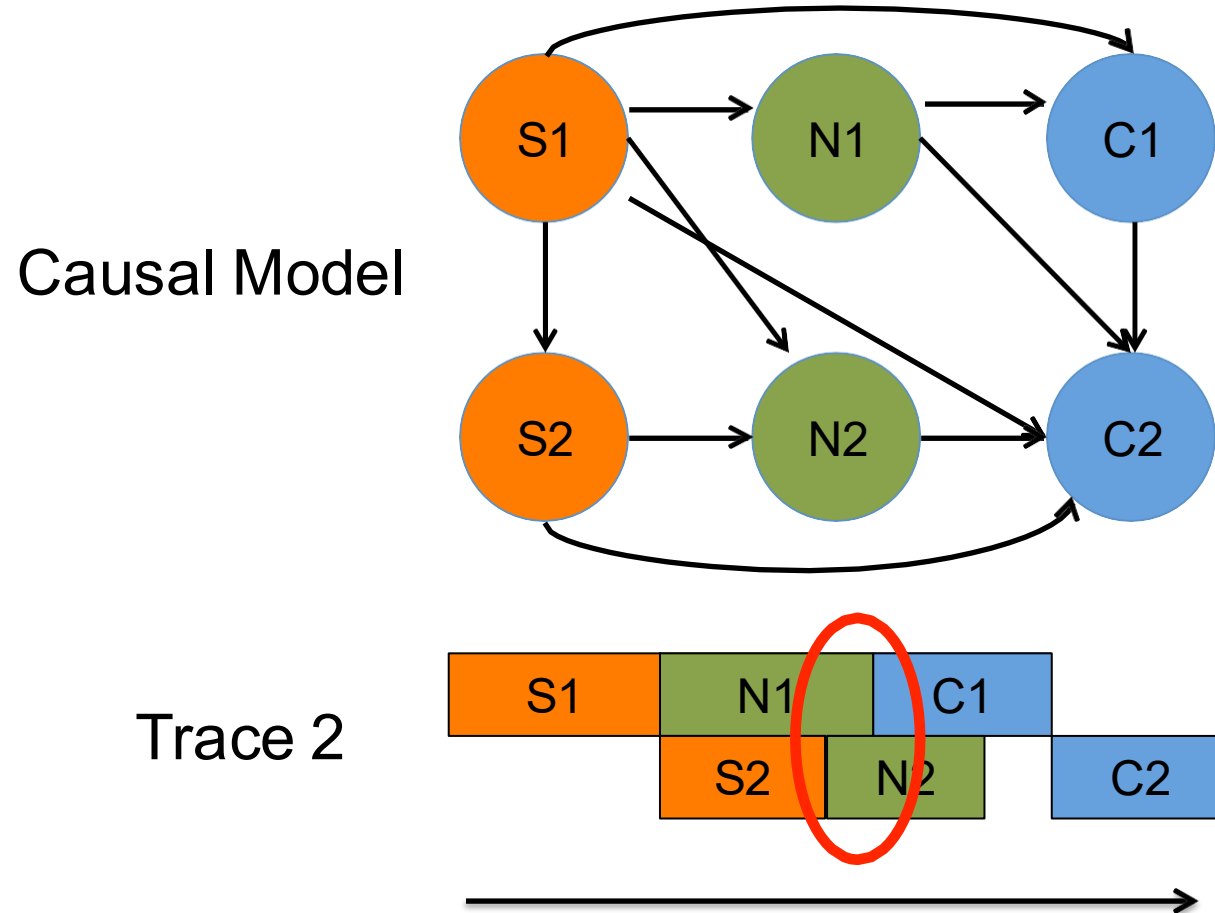
Producing Causal Model



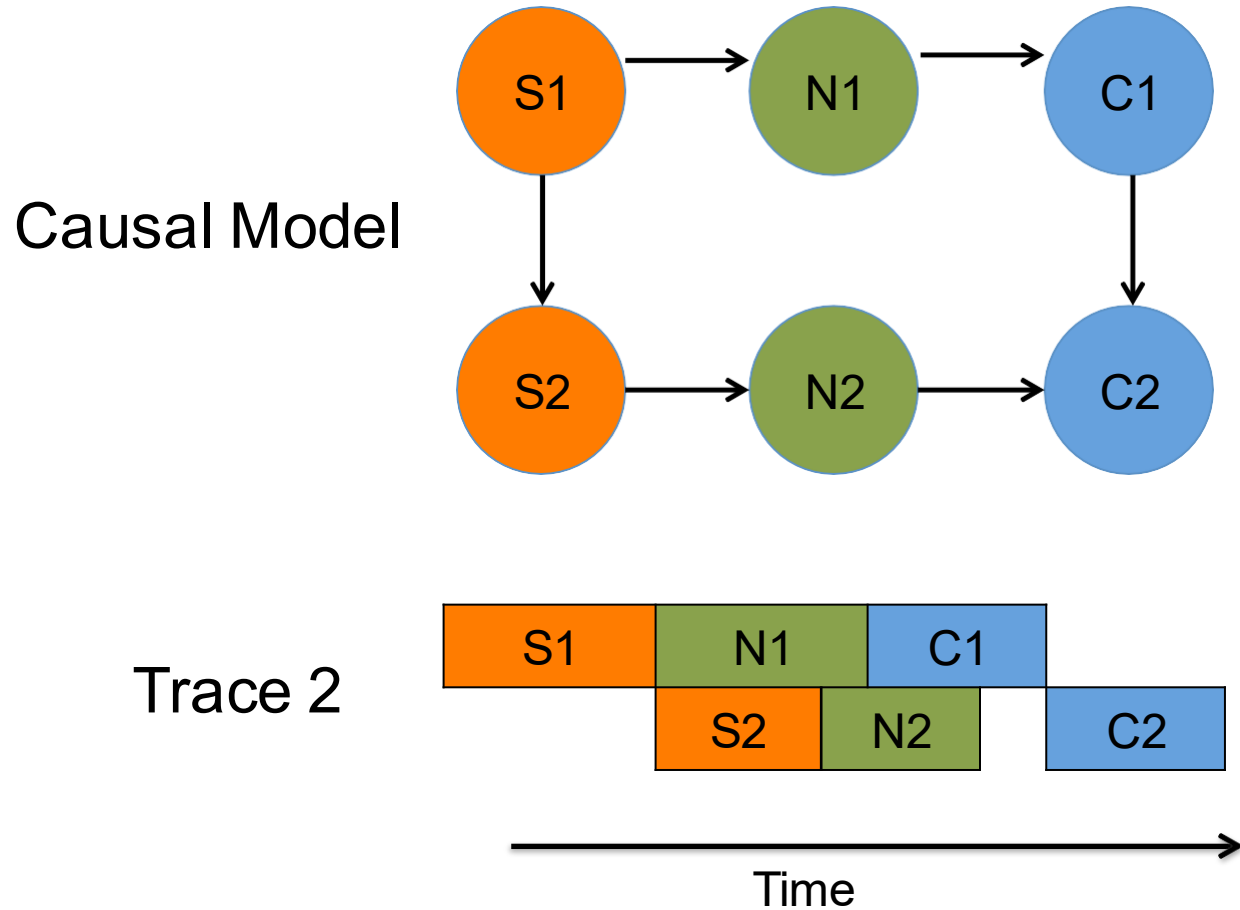
Producing Causal Model



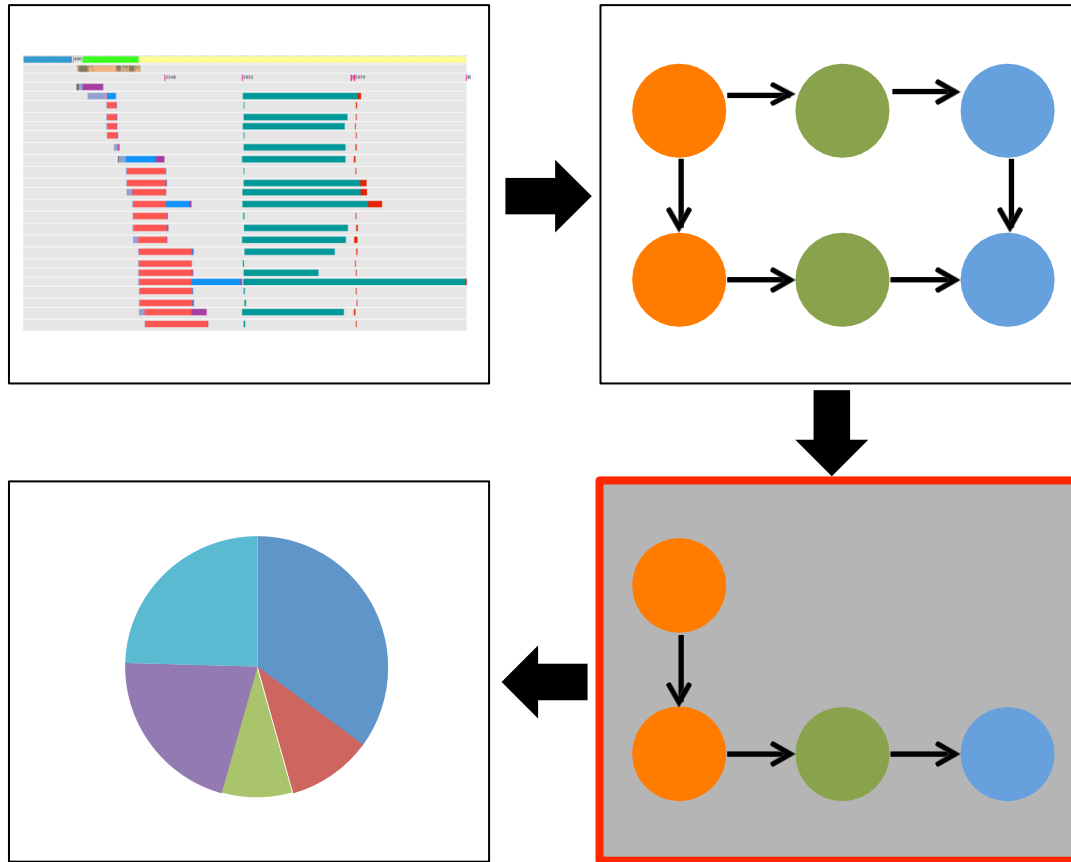
Producing Causal Model



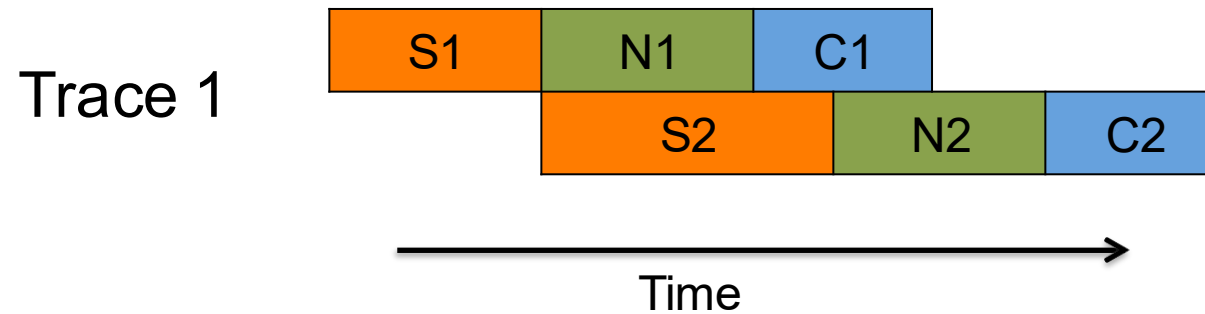

Producing Causal Model




Step 3: Analyze Individual Requests




Critical Path Using Causal Model



Critical Path Using Causal Model



Trace 1



Critical Path Using Causal Model

