Requirements for Time Series Data Platforms:
A Technology Overview with Benchmarks

April 25th, 2018
Albuquerque, NM

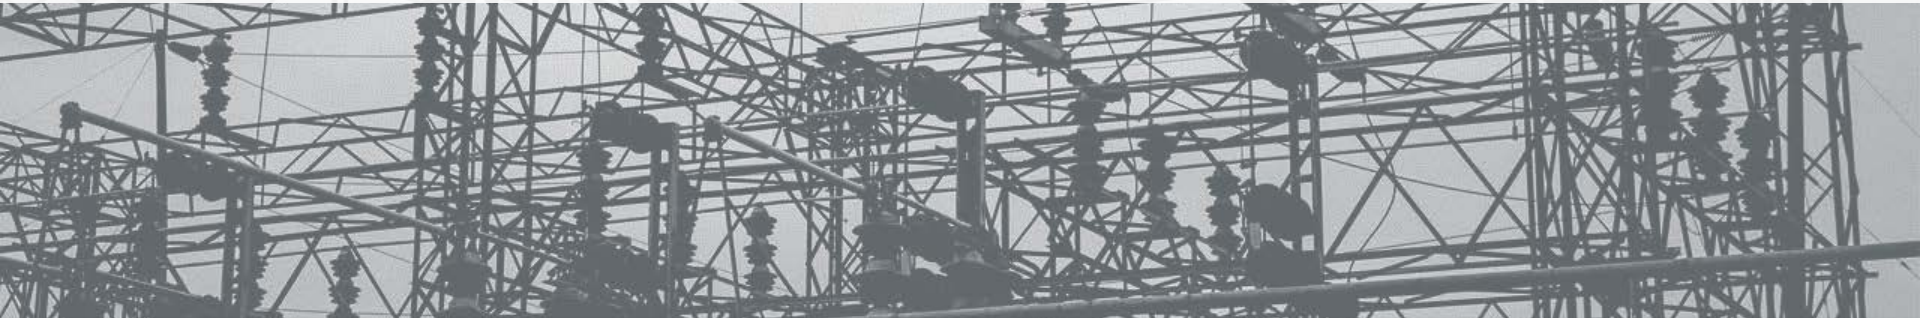

Sean Patrick Murphy
PingThings

Michael Andersen
Berkeley
UNIVERSITY OF CALIFORNIA

Kevin Jones
Dominion Energy®

Outline

- Platform Requirements - 4 Key Drivers
  1. The Nature of Time Series
  2. The Real World is Messy (or Sensors)
  3. Analytics as First Class Citizens
  4. Size Matters
- Benchmarking Implications
- What is a Platform?

The Nature of Time Series

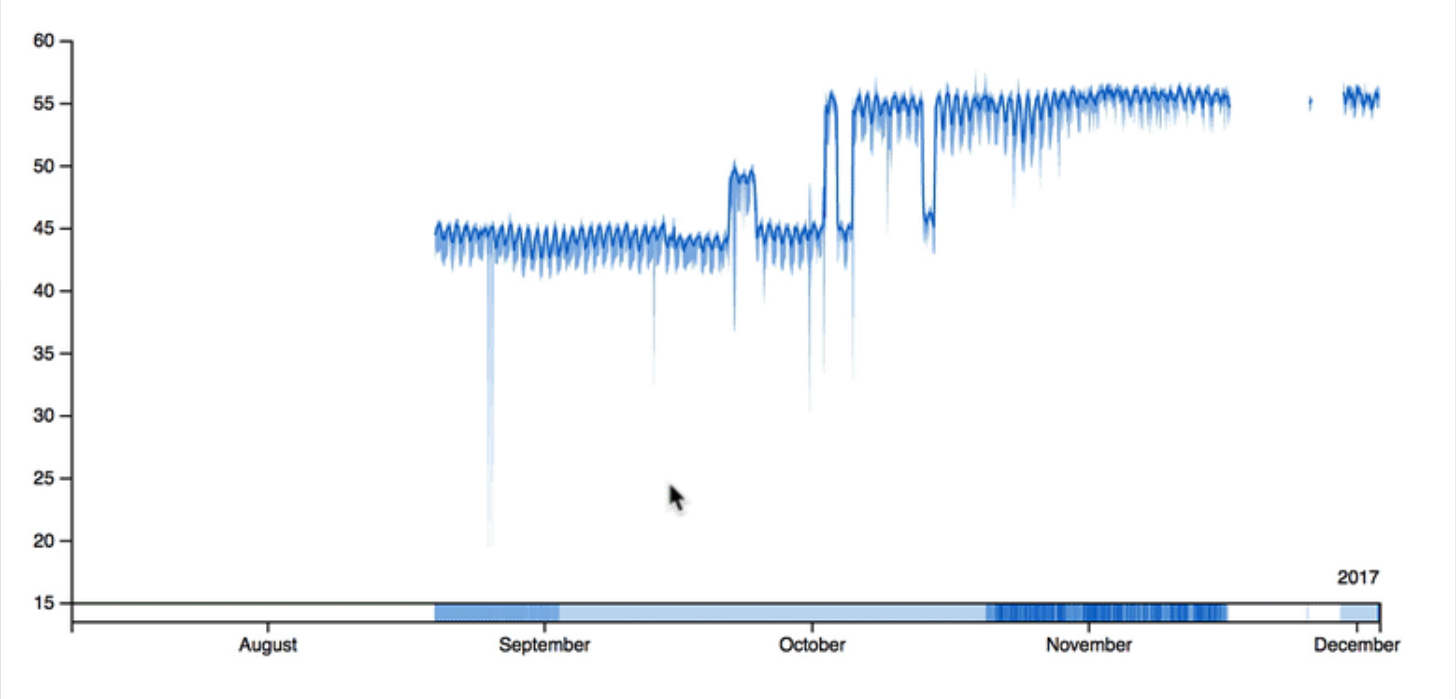

(timestamp, value)

(timestamp, value)

64-bit integer

$[-(2^{63}),\ 2^{63} - 1]$

Or

[-9,223,372,036,854,775,808,
9,223,372,036,854,775,807]

64-bit float

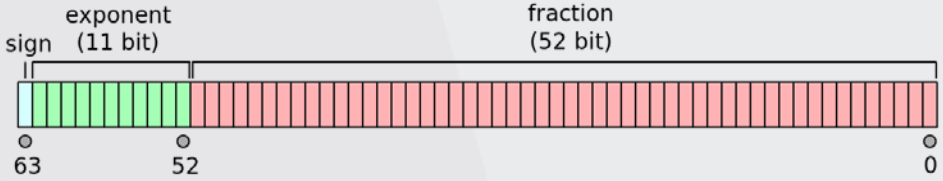

$$(-1)^{\text{sign}}(1.b_{51}b_{50}\ldots b_0)_2 \times 2^{e-1023}$$

(timestamp, value)

```
(151511220000000000, 50.628520965576 1)
(151511220000833333, 50.6273155212402)
(151511220001666666, 50.6269416809082)
(151511220002499999, 50.6258087158203)
(151511220003333332, 50.6216735839843)
(151511220004166665, 50.6205940246582)
(151511220004999998, 50.6227645874023)
(151511220005833331, 50.6207199096679)
(151511220006666664, 50.6192970275878)
(151511220007499997, 50.6227836608886)
(151511220008333330, 50.6249427795410)
```

Intrinsic Redundancy

- Can compress timestamps and values
  - Lossy and Lossless
  - Intra-stream and Inter-stream
- Many approaches available

P. Lindstrom and M. Isenburg. Fast and Efficient Compression of Floating-Point Data. Visualization and Computer Graphics, IEEE Transactions on, 12(5):1245–1250, 2006.
P. Ratanaworabhan, J. Ke, and M. Burtscher. Fast Lossless Compression of Scientific Floating-Point Data. In DCC, pages 133–142. IEEE Computer Society, 2006.

- Classic space/time tradeoff
- We achieve ~3:1 lossless compression

Write Patterns

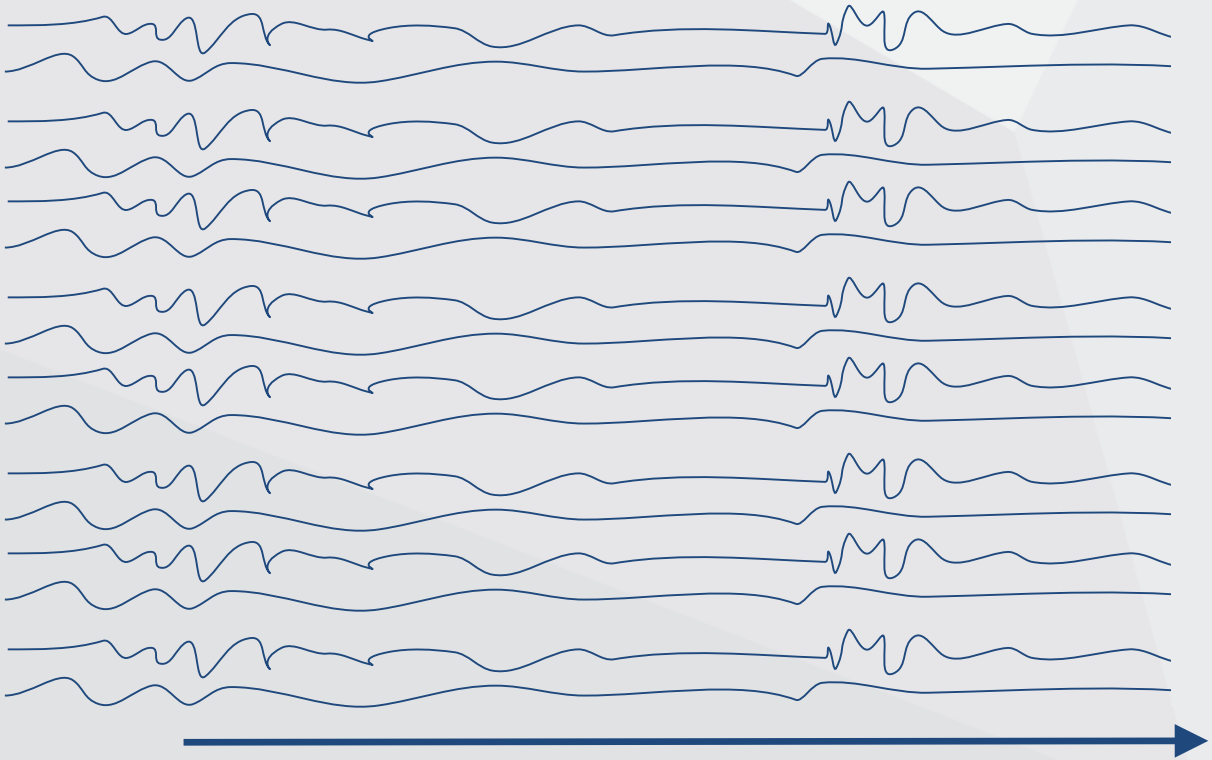

Write Patterns

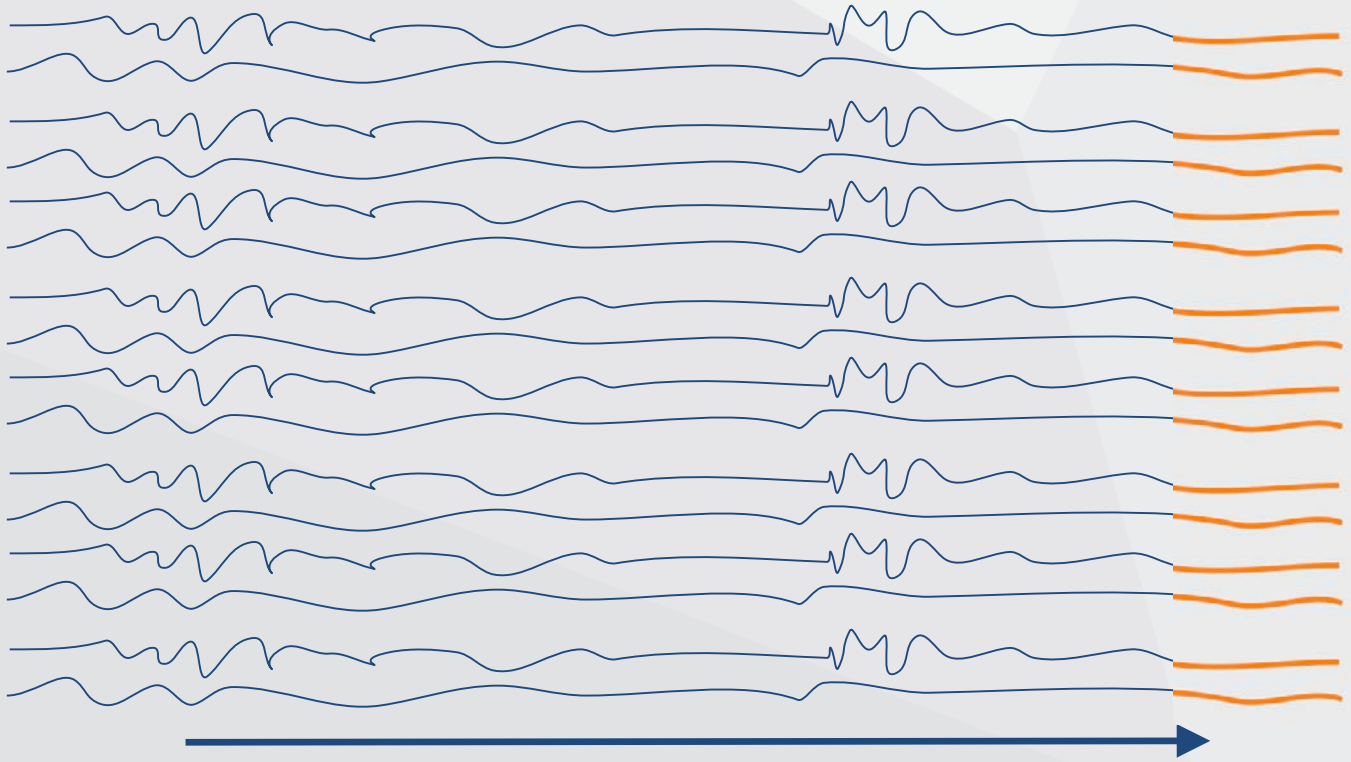

Read Patterns

1. Human interaction and exploration of the data
2. Analytics

Read Patterns

“Overview first, zoom and filter, then details-on-demand.”

The Visual Information-Seeking Mantra [Shneiderman, 1996] summarizes many visual design guidelines and provides an excellent framework for designing information visualization applications.

Random, Multi-Resolution Read Patterns

About 4 billion data points

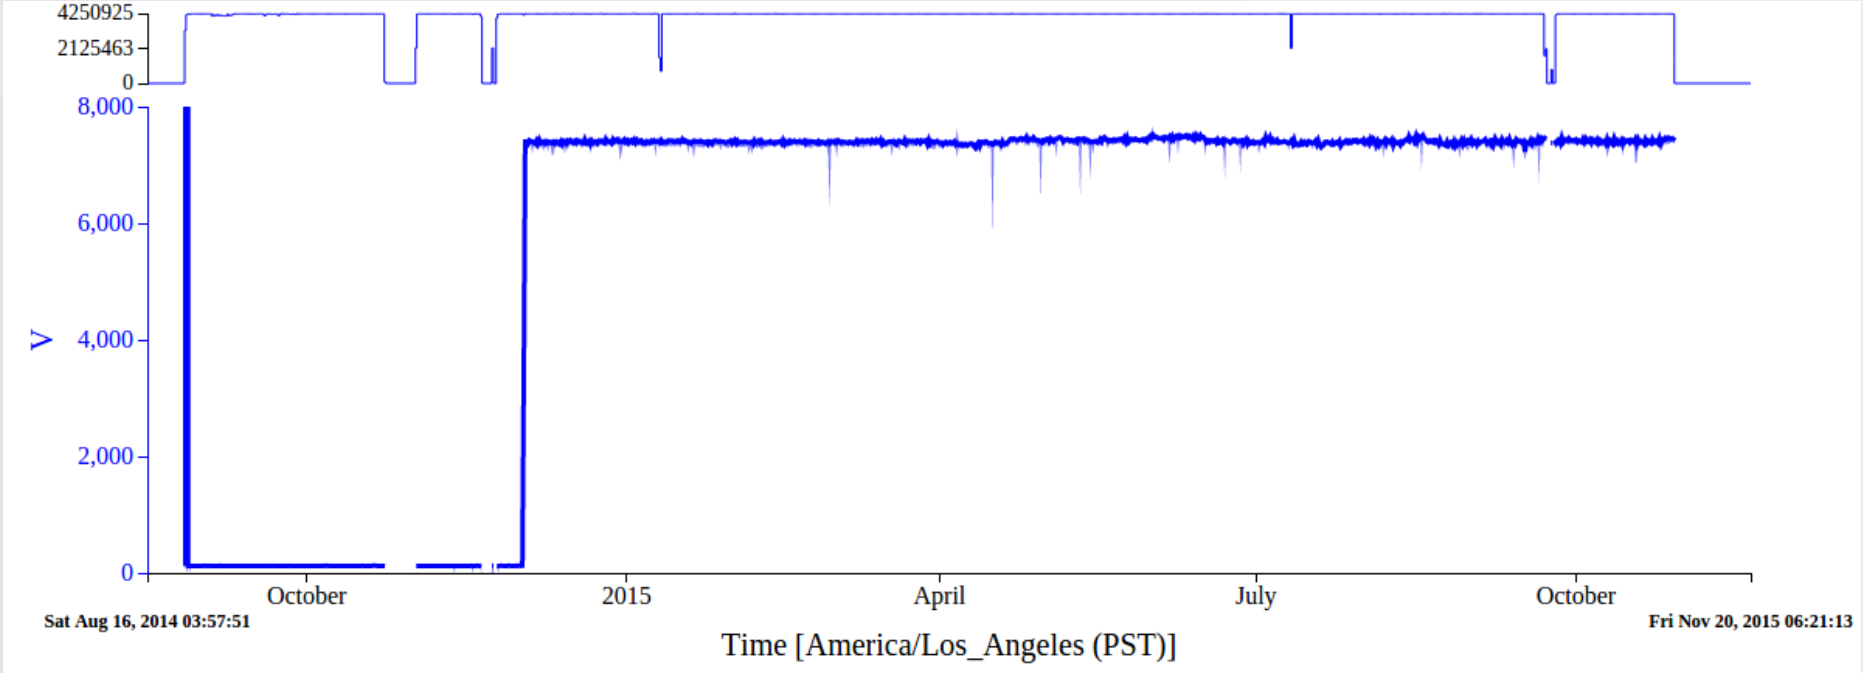

Random, Multi-Resolution Read Patterns

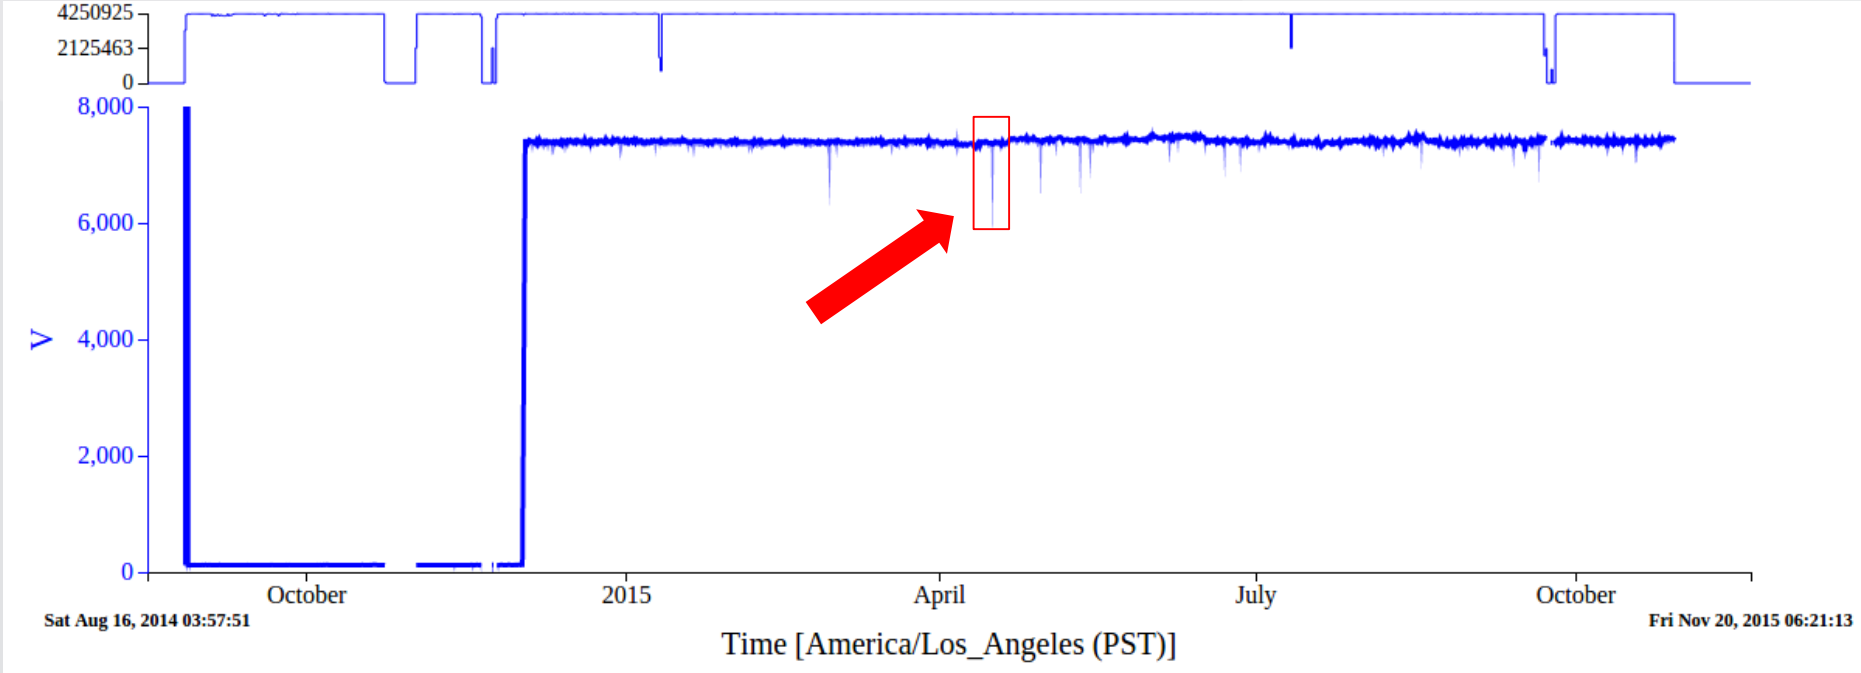

Random, Multi-Resolution Read Patterns

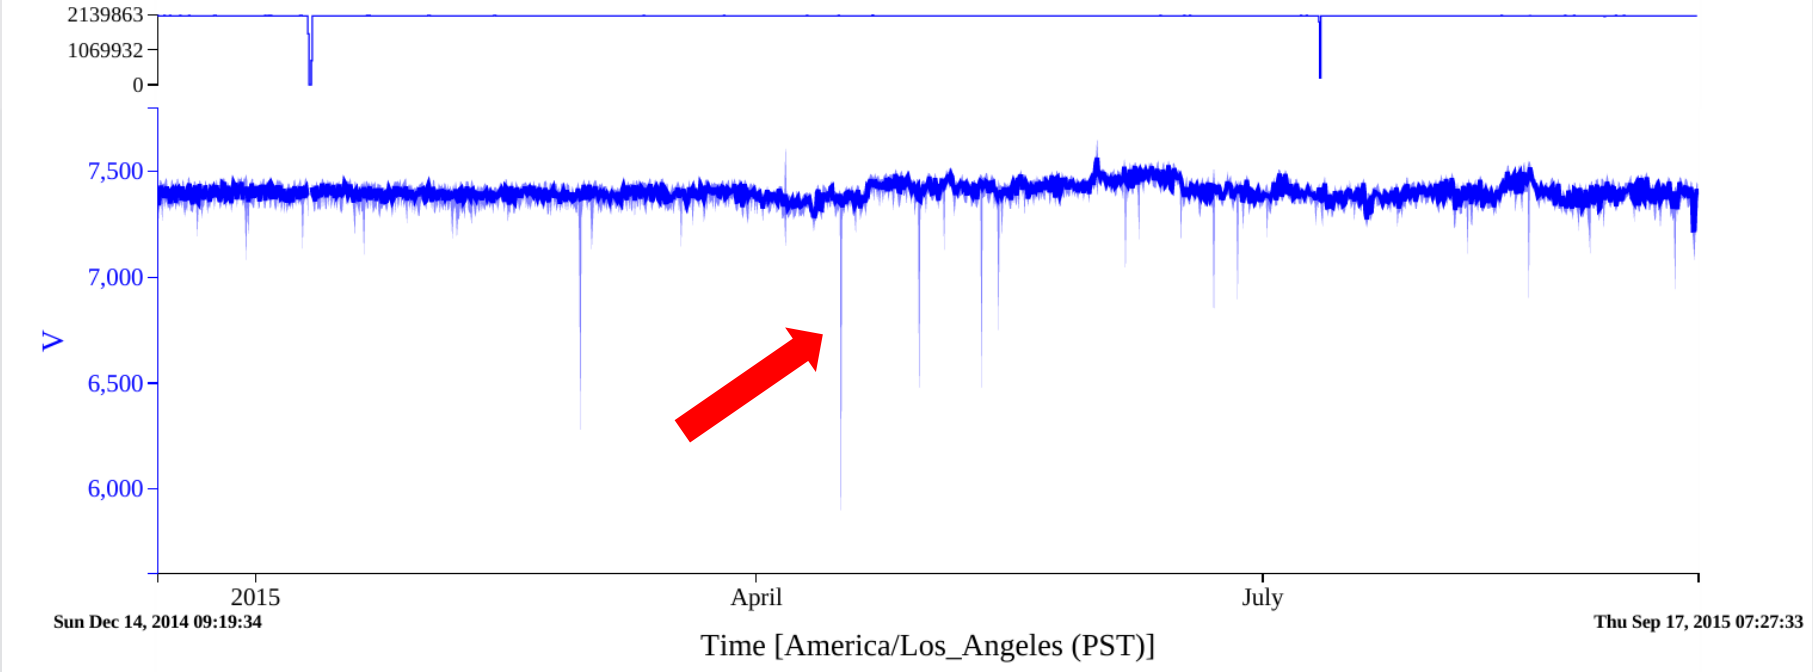

Random, Multi-Resolution Read Patterns

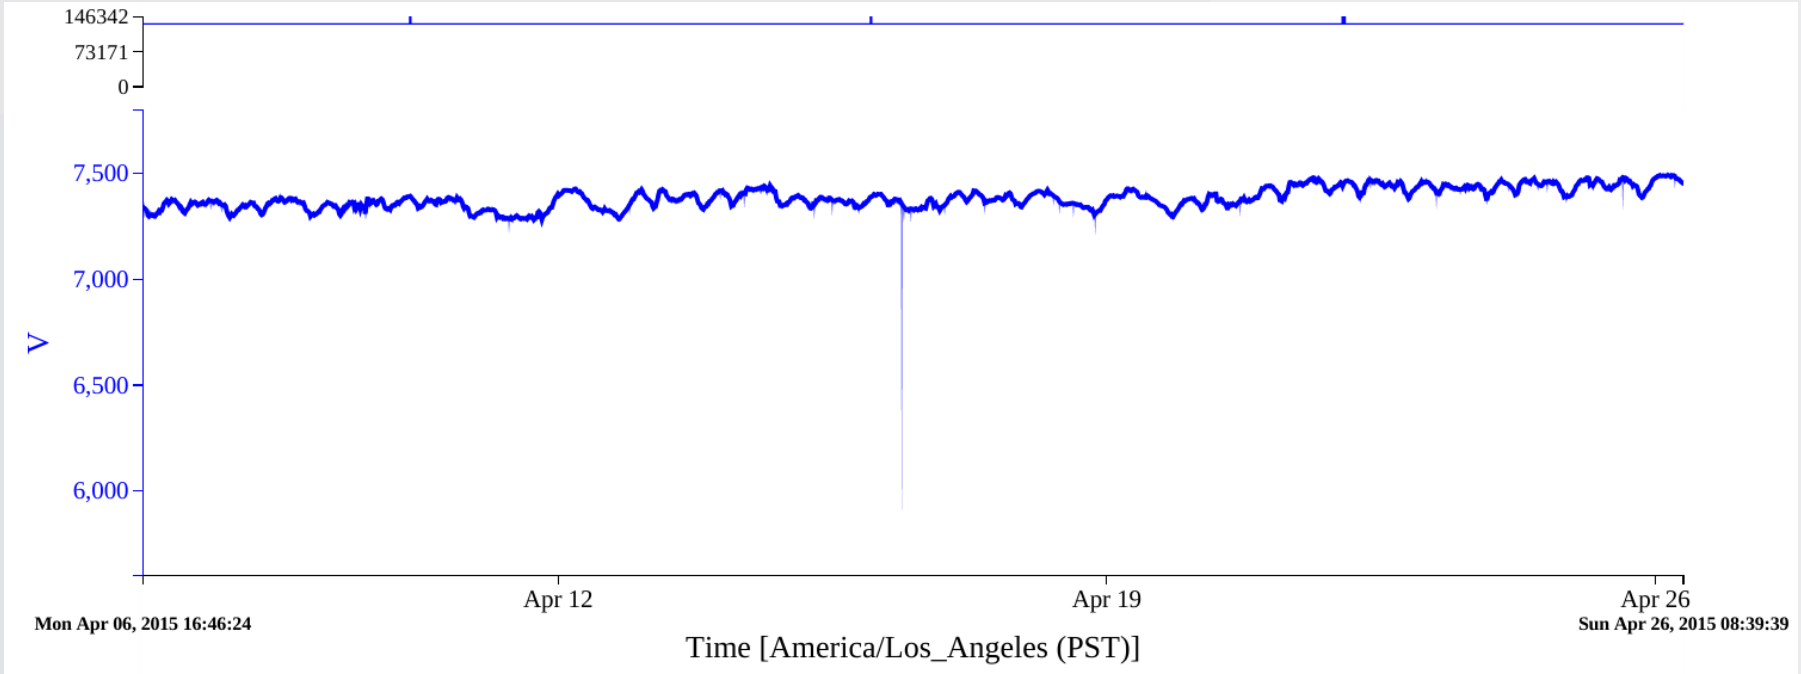

Random, Multi-Resolution Read Patterns

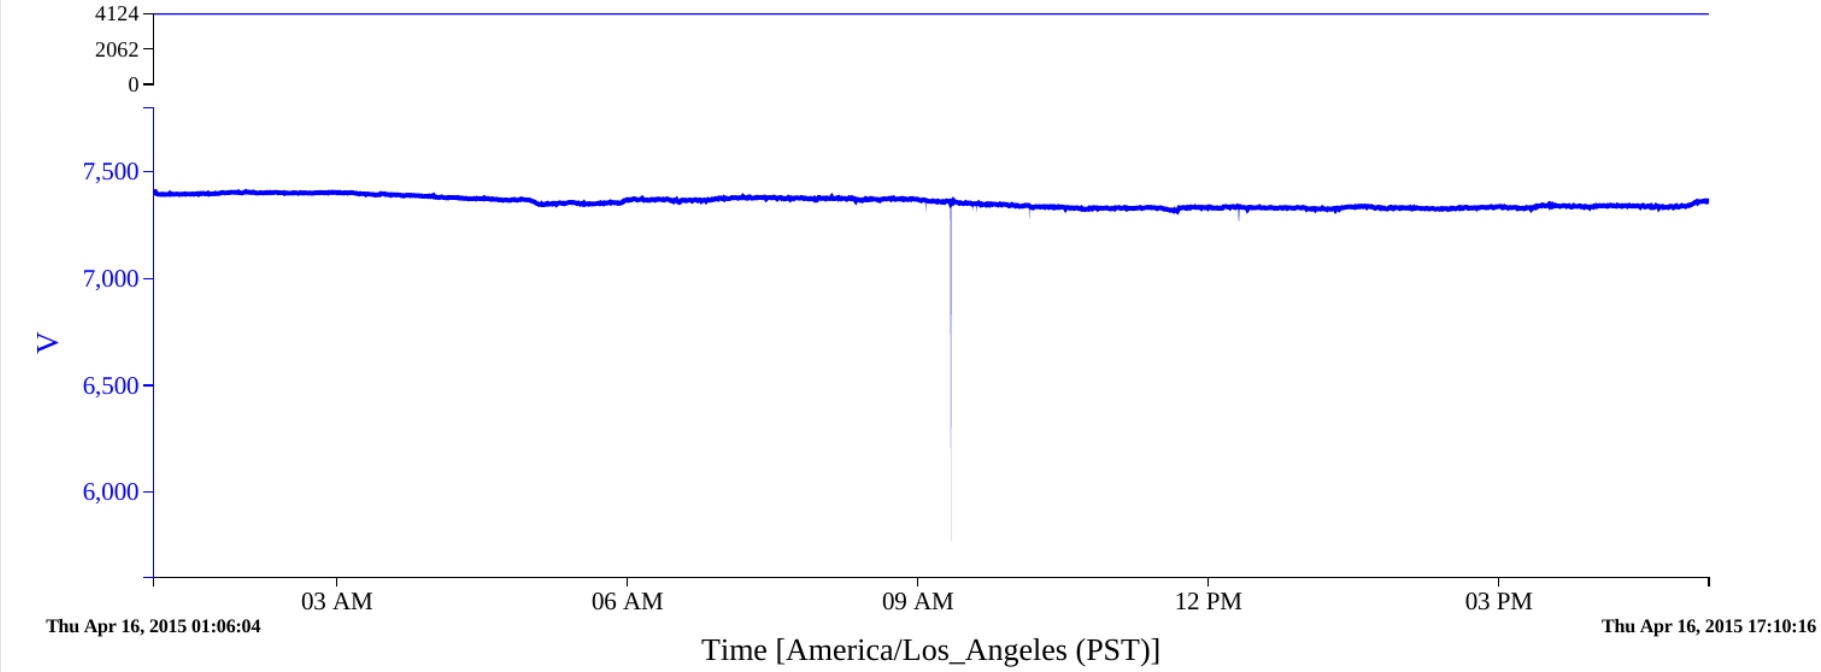

Random, Multi-Resolution Read Patterns

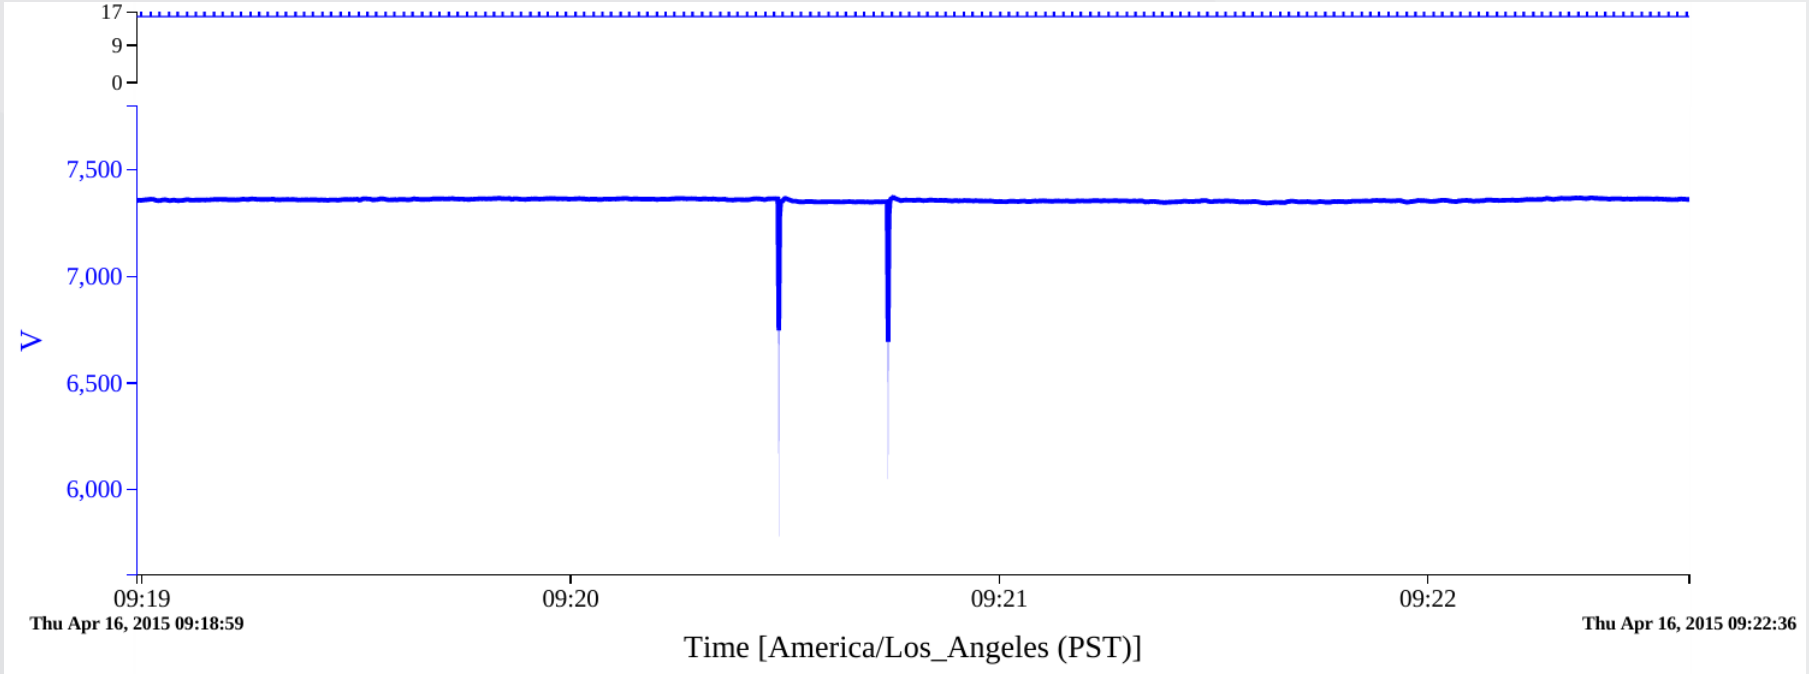

Random, Multi-Resolution Read Patterns

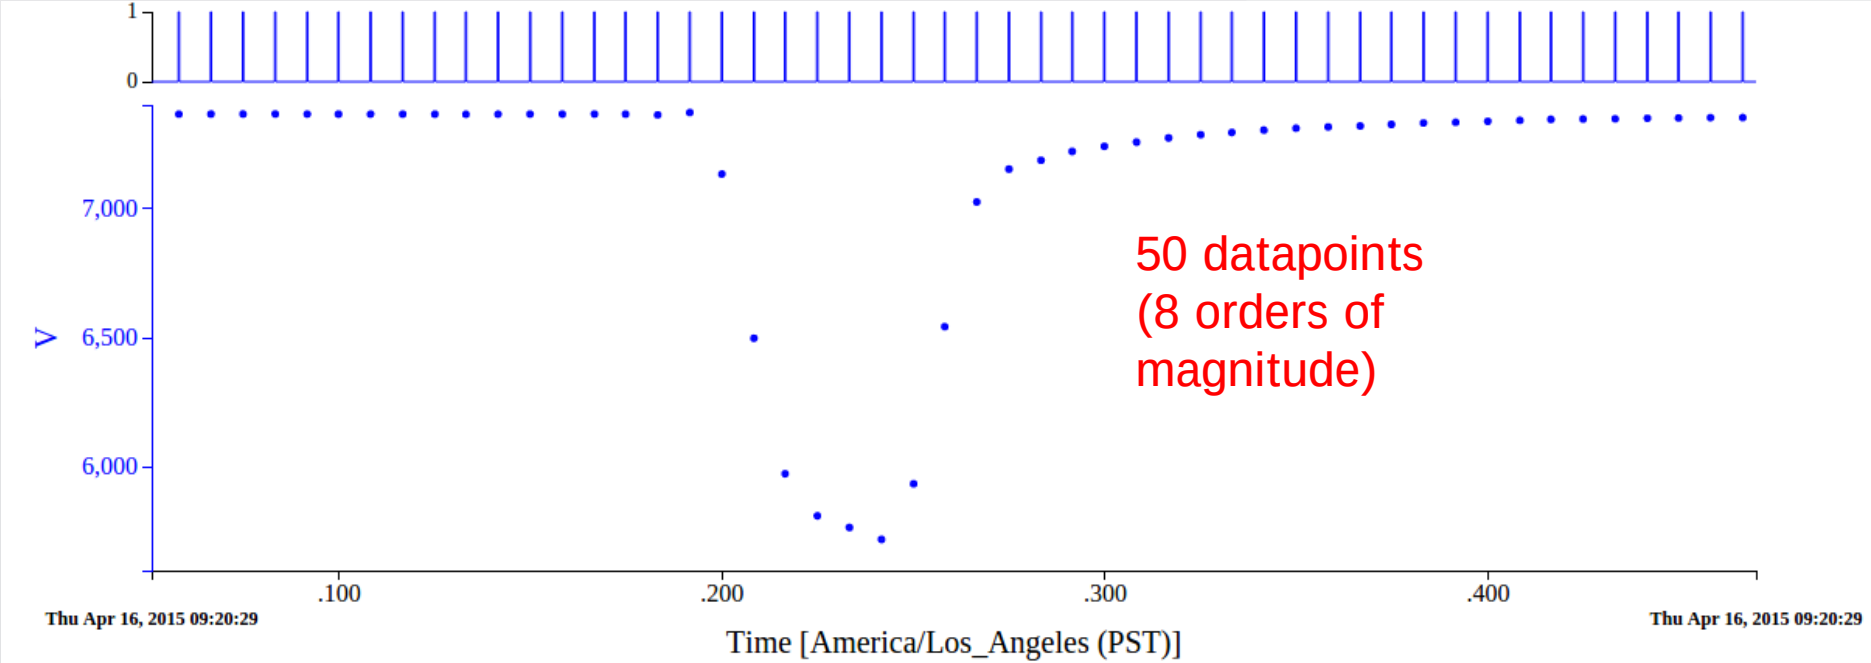

Berkeley tree Data Structure

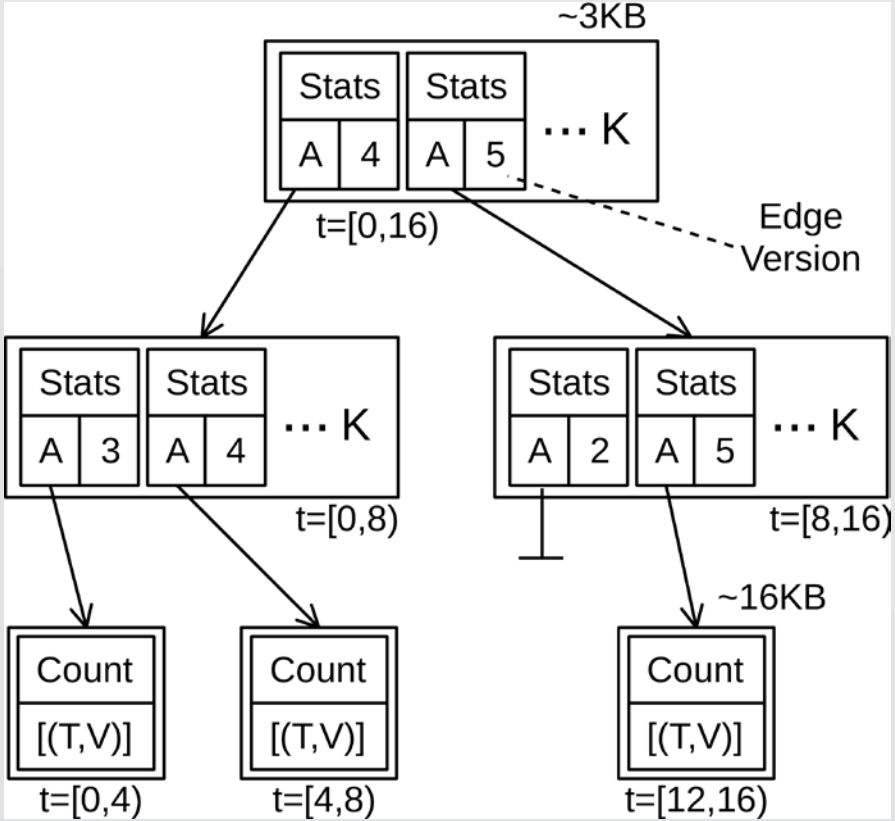

Copy on write K-ary Tree
Partitioning static time (1933 to 2079)

Leaf nodes
- Time, value pairs + length

Internal nodes
- Pointers to children
- **Version annotations** for children
- **Aggregates** for children
    - Min, Mean, Max, Count
    - Any associative operator

The Real World is Messy

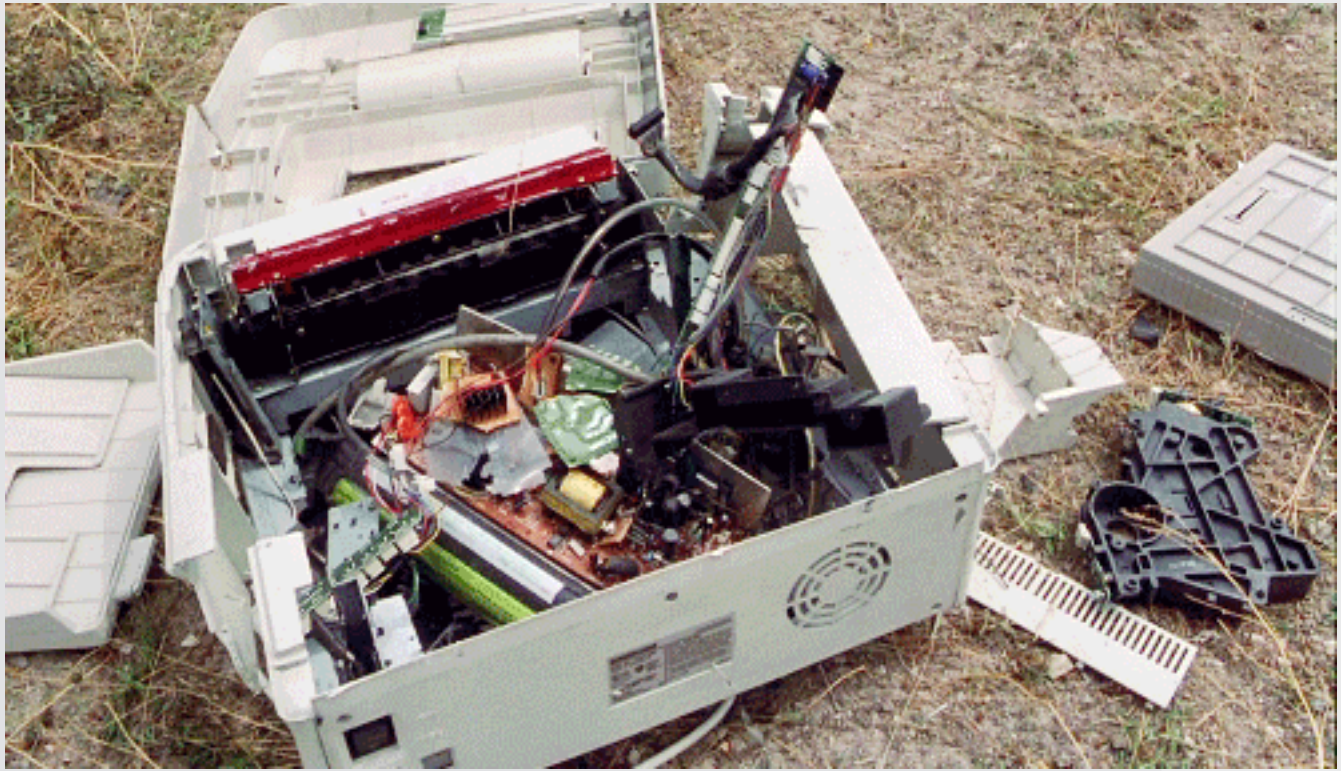


v2

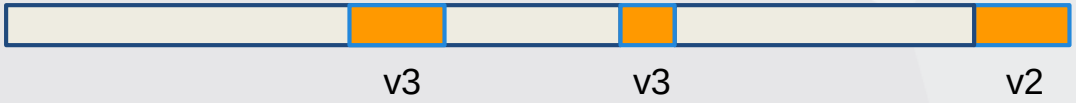
v3
v3
v2

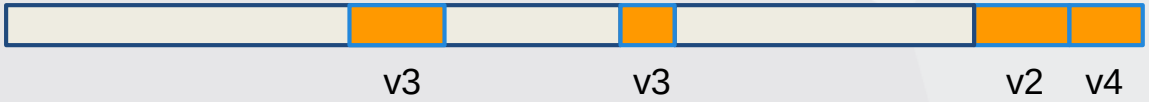
v3
v3
v2
v4

Q: What changed between v2 and v3?

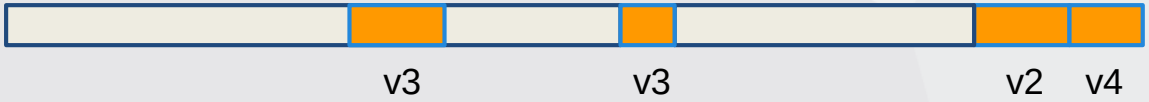

Q: What changed between v2 and v3?

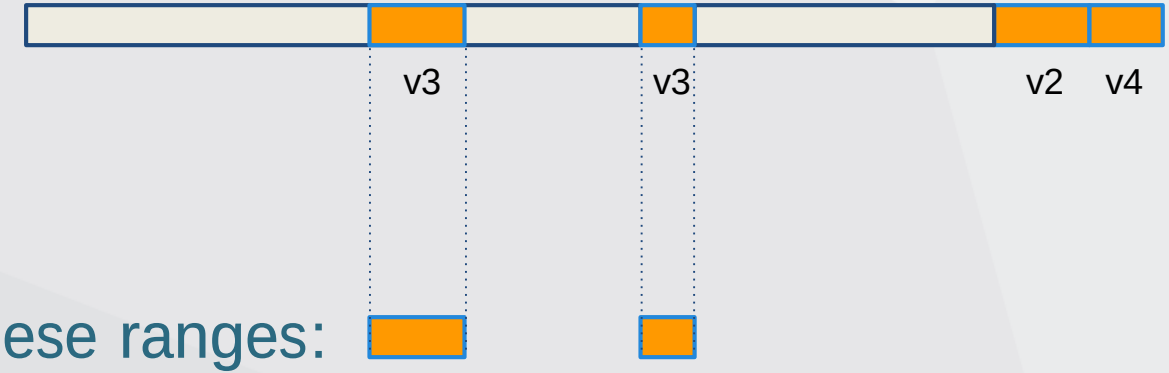

A: These ranges:

Q: What changed between v2 and v3?

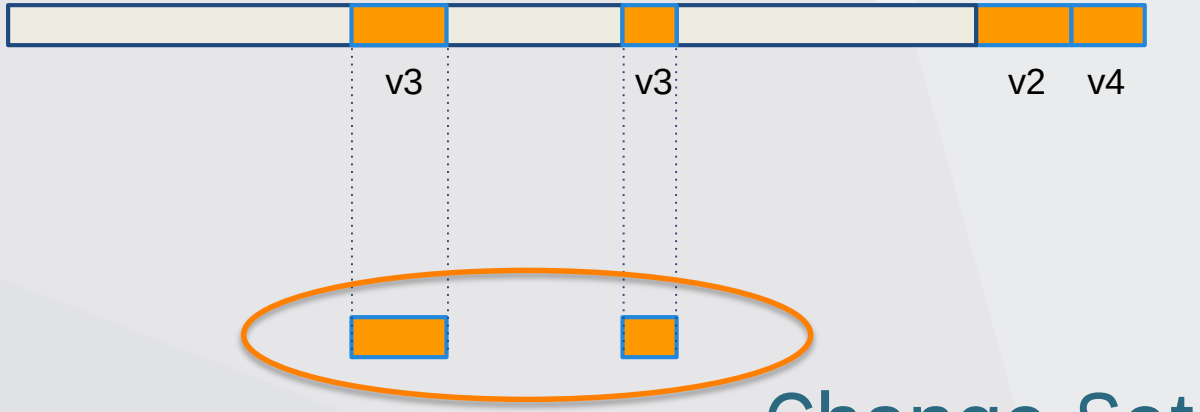

Who Cares?

Who Cares?

- Efficiently update calculations on out of order data

Who Cares?

- Efficiently update calculations on out of order data
- Idempotent calculations????

Who Cares?

- Efficiently update calculations on out of order data
- Idempotent calculations????
- Rewinding data arrival to understand and diagnose problems

Versioning Source Code

5 internal/cephprovider/cephprovider.go

```
@@ -148,12 +148,15 @@ func (seg *CephSegment) Write(uuid []byte, address uint64,
148 148         //start of an object. This is why we do not add the object max size here
149 149         //NEW NOTE:
150 150         //We cannot go past the end of the allocation anymore because it would br
151     -       if ((naddr + MAX_EXPECTED_OBJECT_SIZE) >> 24) != (address >> 24) {
    151 +       if ((naddr + MAX_EXPECTED_OBJECT_SIZE + 2) >> 24) != (address >> 24) {
152 152                 //We are gonna need a new object addr
153 153                 naddr = <-seg.sp.alloc
    154 +               seg.naddr = naddr
154 155                 seg.flushWrite()
    156 +               return naddr, nil
155 157         }
156 158         seg.naddr = naddr
    159 +
157 160         return naddr, nil
158 161 }
```

Analytics as First Class Citizens

Common Analytics Patterns

Common Analytics Patterns

Common Analytics Patterns

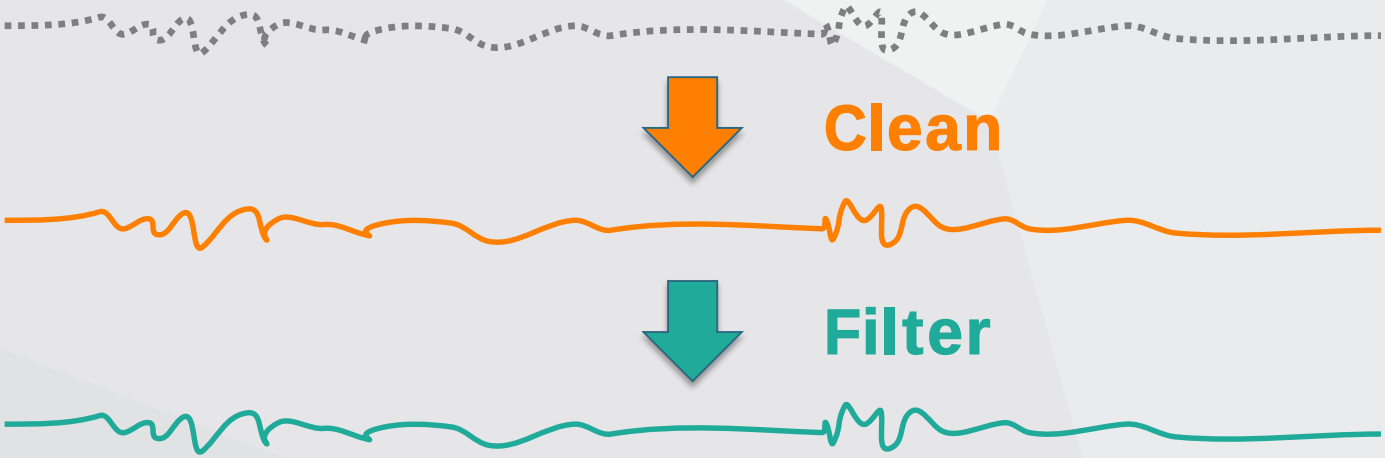

Common Analytics Patterns

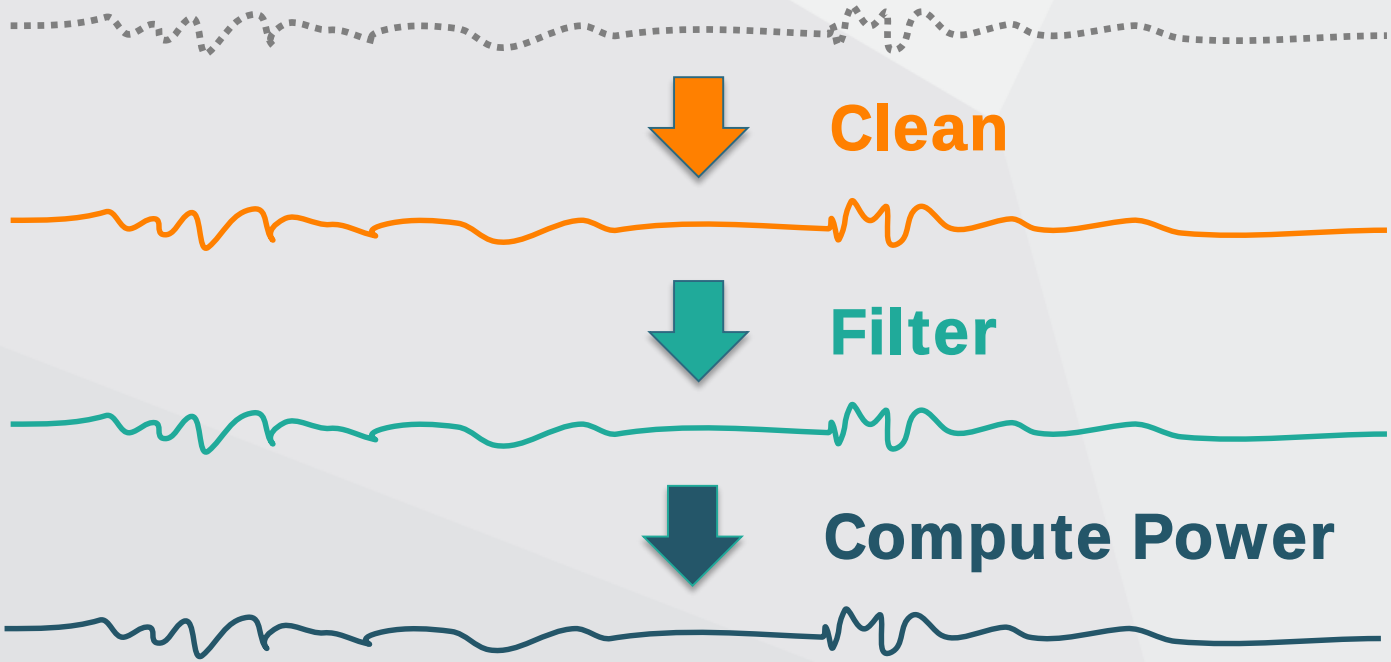

Directed Acyclic Graph of Calculations

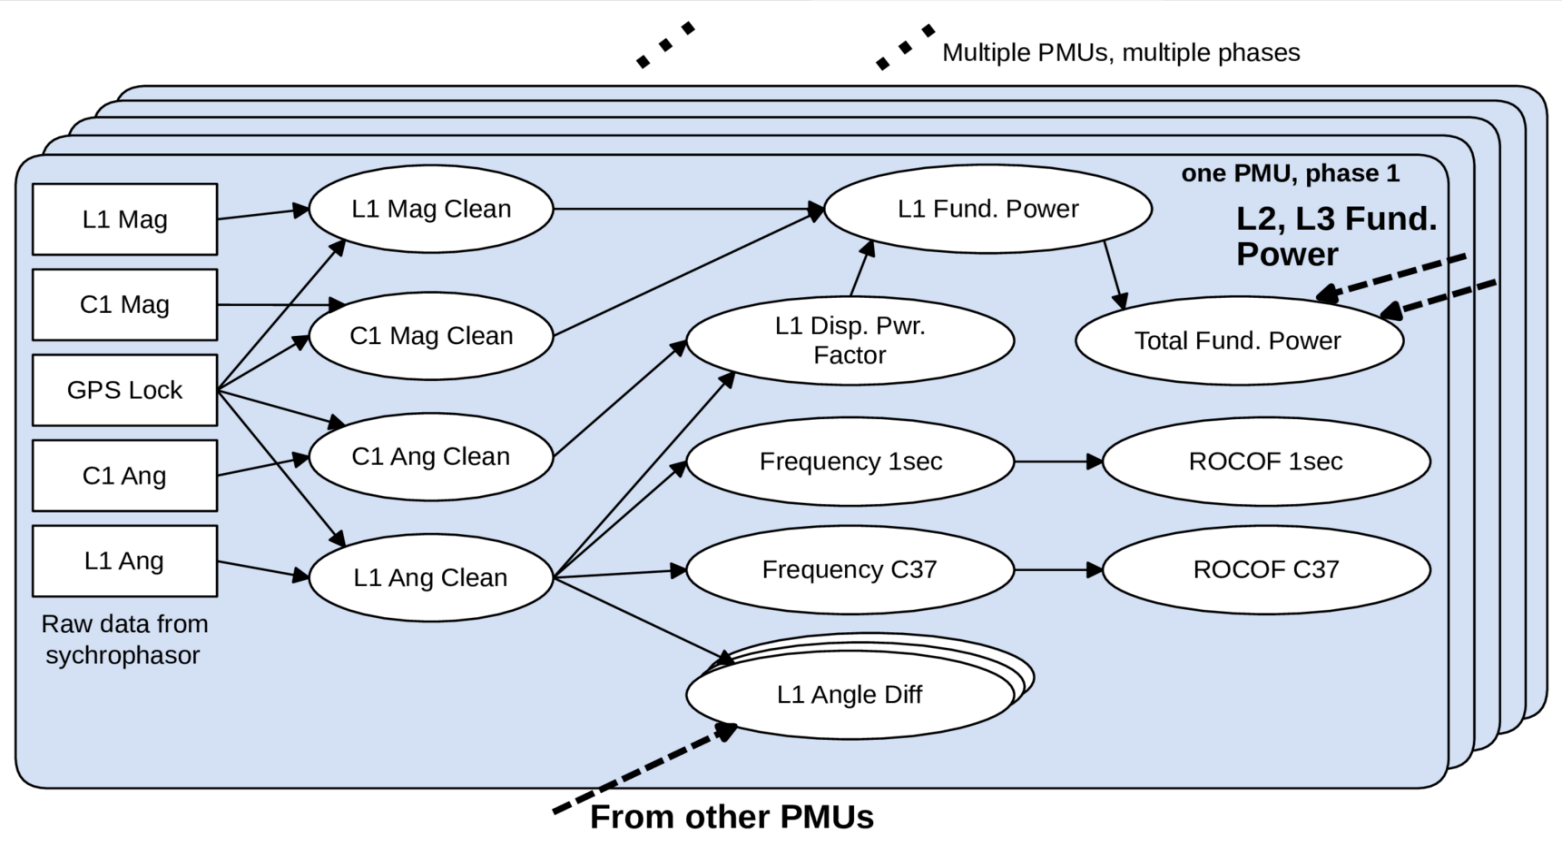

But Wait, There's More

- Windowing operations
- Spatial/Frequency Transforms
  - Wavelet
  - Fourier
  - Stockwell
- Indexing
- Clustering
- Classification
- Categorization
- Anomaly/Event/Novelty Detection
- Motif Discovery

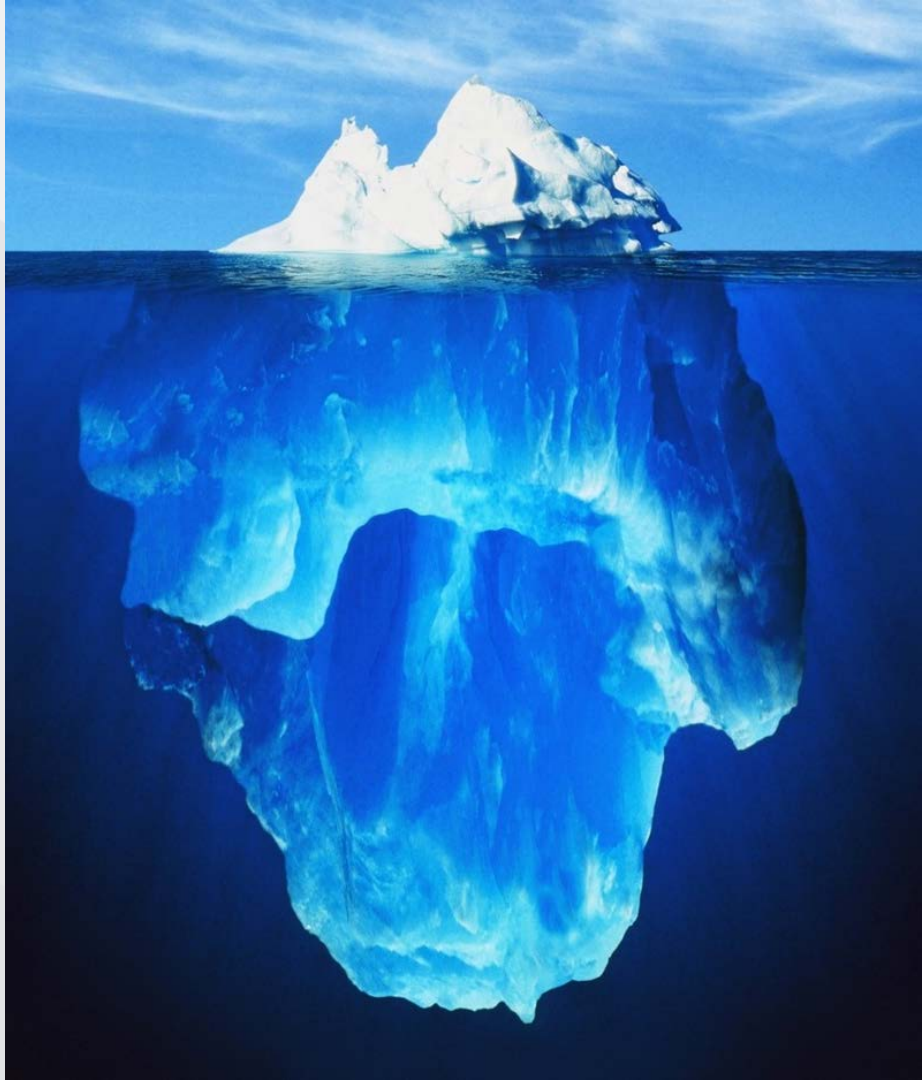

But Wait, There's Even More!

Old Paradigm - Software Engineering

- Humans write the code
- Limited by ability to describe exactly what must be done without error

New Paradigm - Machine Learning

- Data teaches algorithms to perform function or task
- Limited by the amount of data and algorithms
- Algorithms need ***ALL*** available data
- Capable of tackling high dimensional problems

Deep Blue beats Gary Kasparov 1997

Watson beats champions 2011

AlphaGo beats Lee Sedol 2016

Who Cares?

- Analytics must be first class citizens of the platform
- The platform must be built from the ground up to support relevant analytic use cases

Size Matters

Measurements Per Second

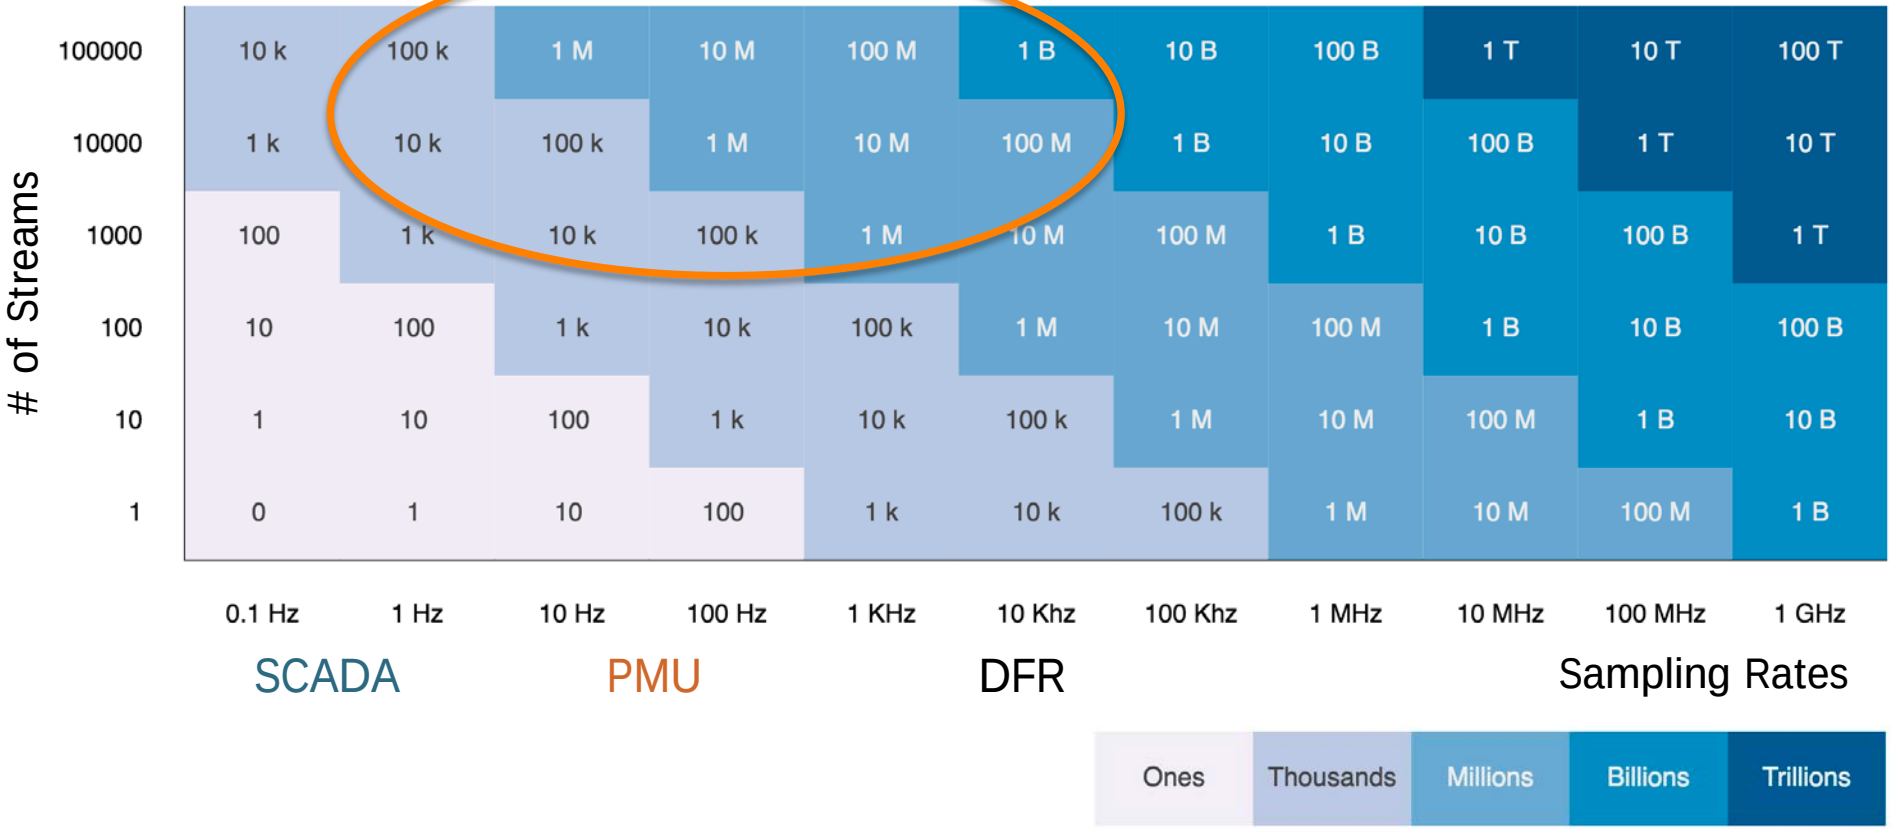

Required Network Bandwidth (Lower Bound)

Annual Data Volumes (Lower Bound)

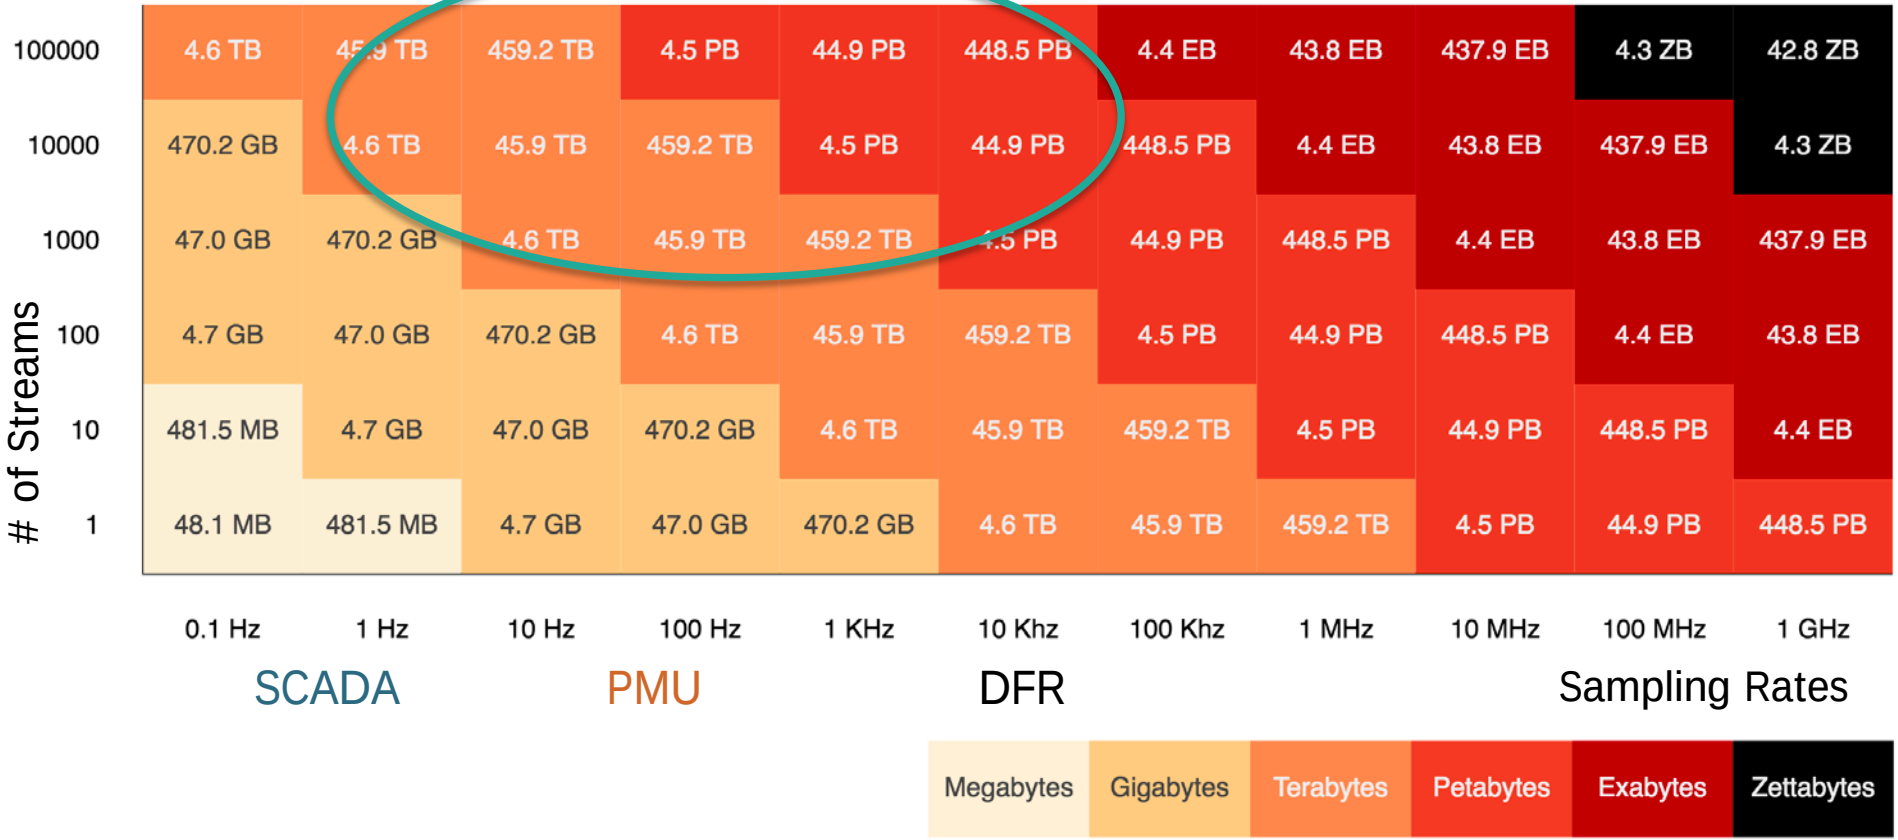

Big Data Conclusions

- Don’t move the data, move the calculations.
  - (Why analytics are first class citizens)
- We are going to need a bigger [machine(s)].

Two Options

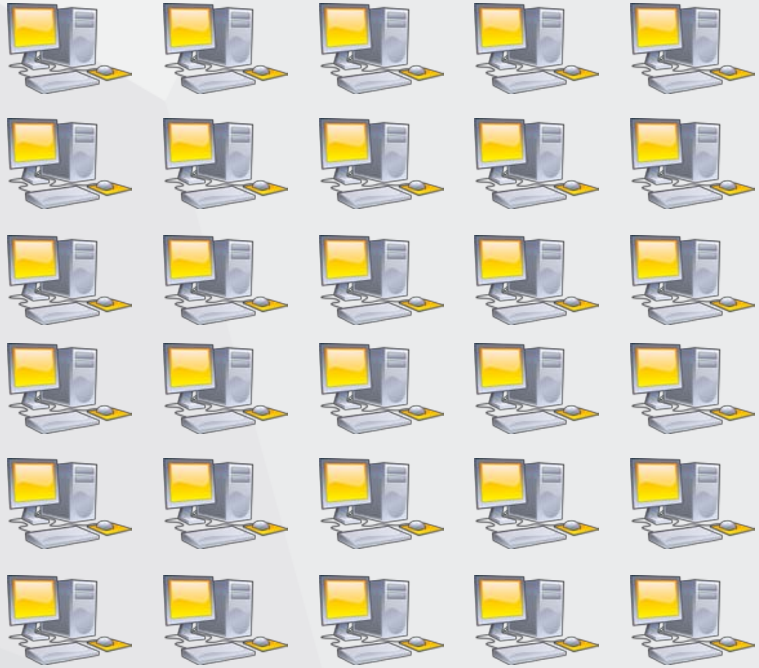

How this Played Out

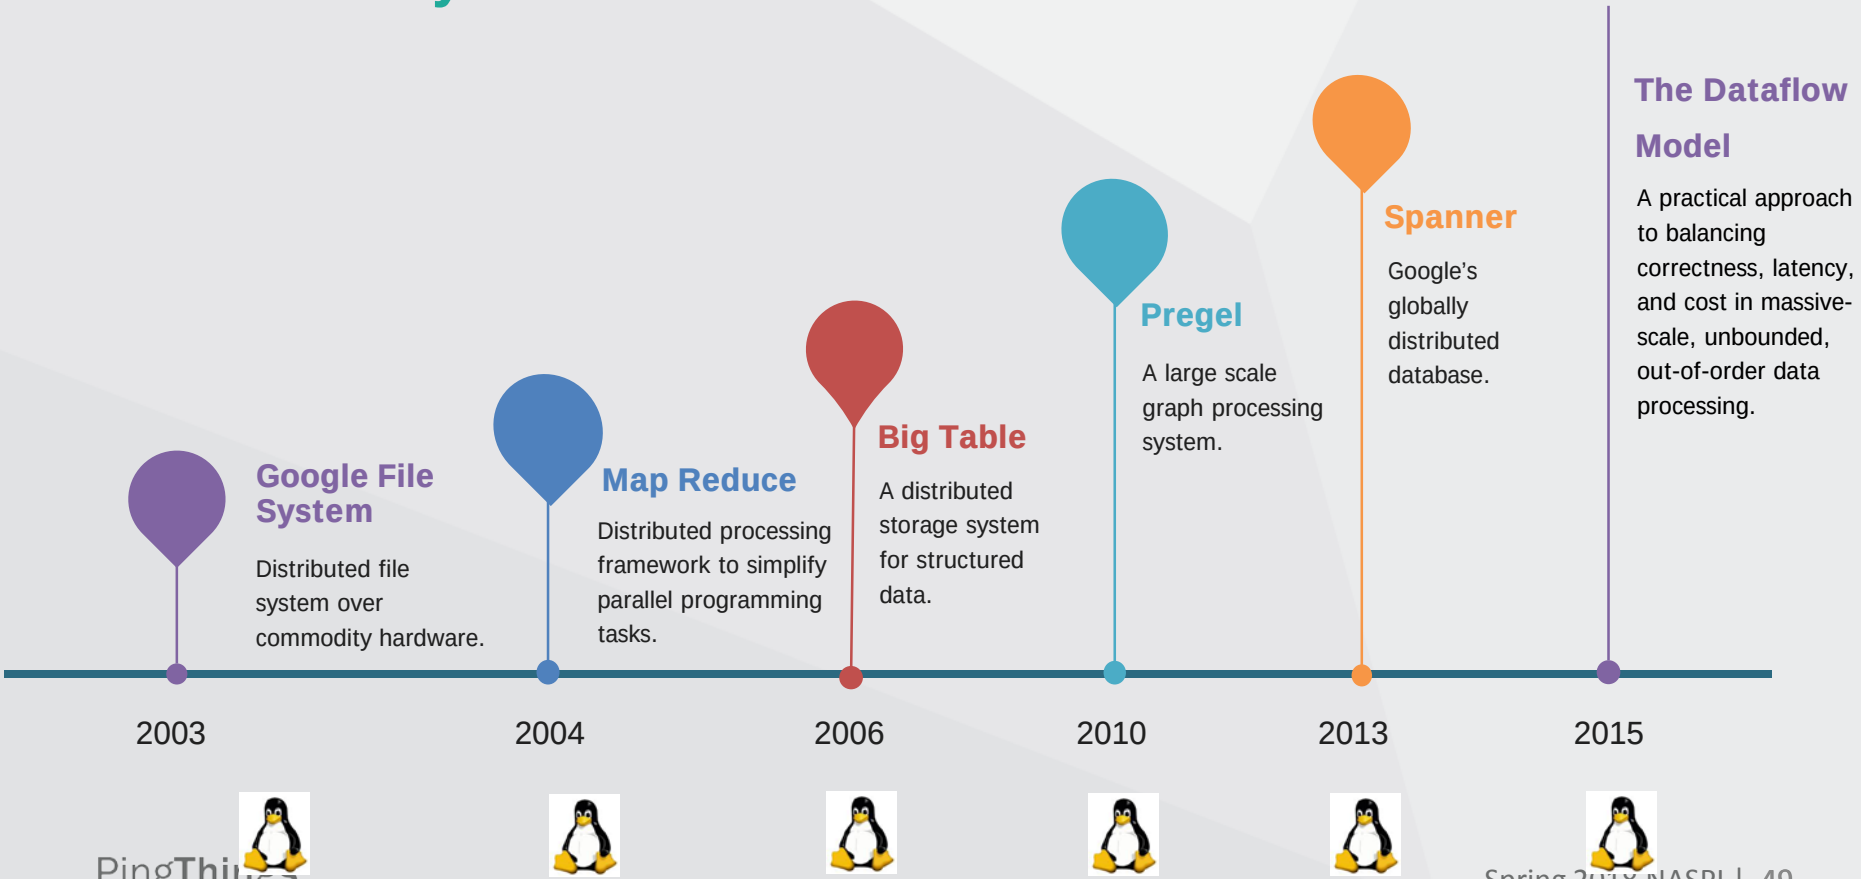

Who Cares?

- No single point of failure/resilient
- Much more cost effective
- Buy more capacity when you need it
- Flexible and adaptable

Benchmark Considerations

Core Benchmarks – Reading and Writing Data

#BTrDB	Streams	Total points	#Conn	Insert [mil/s]	Cold Query [mil/s]	Warm Query [mil/s]
1	50	500 mil	30	16.77	9.79	33.54
2	100	1000 mil	60	28.13	17.23	61.44
3	150	1500 mil	90	36.68	22.05	78.47
4	200	2000 mil	120	53.35	33.67	119.87

Analytics Benchmarks

	Distributed		
	Identity	Phase Difference	Reactive/Fundamental Pwr
Input/Output streams	1/1	2/1	4/2
Compute changeset	972 μs	1659μs	1180μs
Query data [s]	69.8	104.4	196.9
Kernel calculation [s]	10.8	22.7	245.5
Delete old data[s]	6.7	6.9	15.8
Insert new data[s]	40.7	39.8	66.5
Changeset / compute time	1064 x	773 x	259 x

What is a Platform?

What isn't a Time Series Platform?

Time Series

Not Time Series

"Small" Data

"Big" Data

What Does a Time Series Data Platform Do?

Ingest / Egest

Clean / Condition / Fix

Store

Visualize

Access / Use / Build With

Analyze

Learn From

The Goal of the Platform

Allow utility subject matter experts to create value from (sensor) data.

Platform?

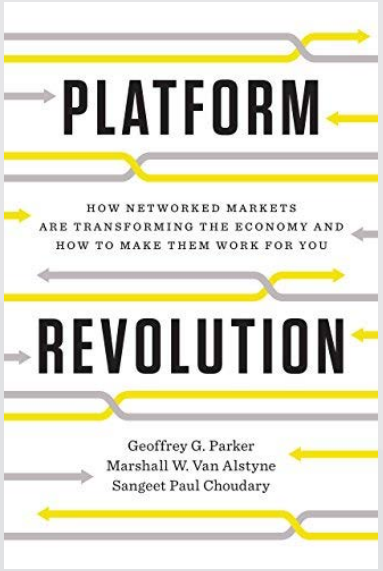

"A platform is a business based on enabling value-creating interactions between external producers and consumers. The platform provides an open, participative infrastructure for these interactions and sets governance conditions for them. **The platform's overarching purpose: to consummate matches among users and facilitate the exchange of goods, services, or social currency, thereby enabling value creation for all participants.**

Strategy has moved from controlling unique internal resources and erecting competitive barriers to orchestrating external resources and engaging vibrant communities. And innovation is no longer the province of in-house experts and research and development labs, but is produced through crowdsourcing and the contribution of ideas by independent participants in the platform. External resources don't completely replace internal resources—more often they serve as a complement. But platform firms emphasize ecosystem governance more than product optimization, and persuasion of outside partners more than control of internal employees."

PingThings PredictiveGrid™

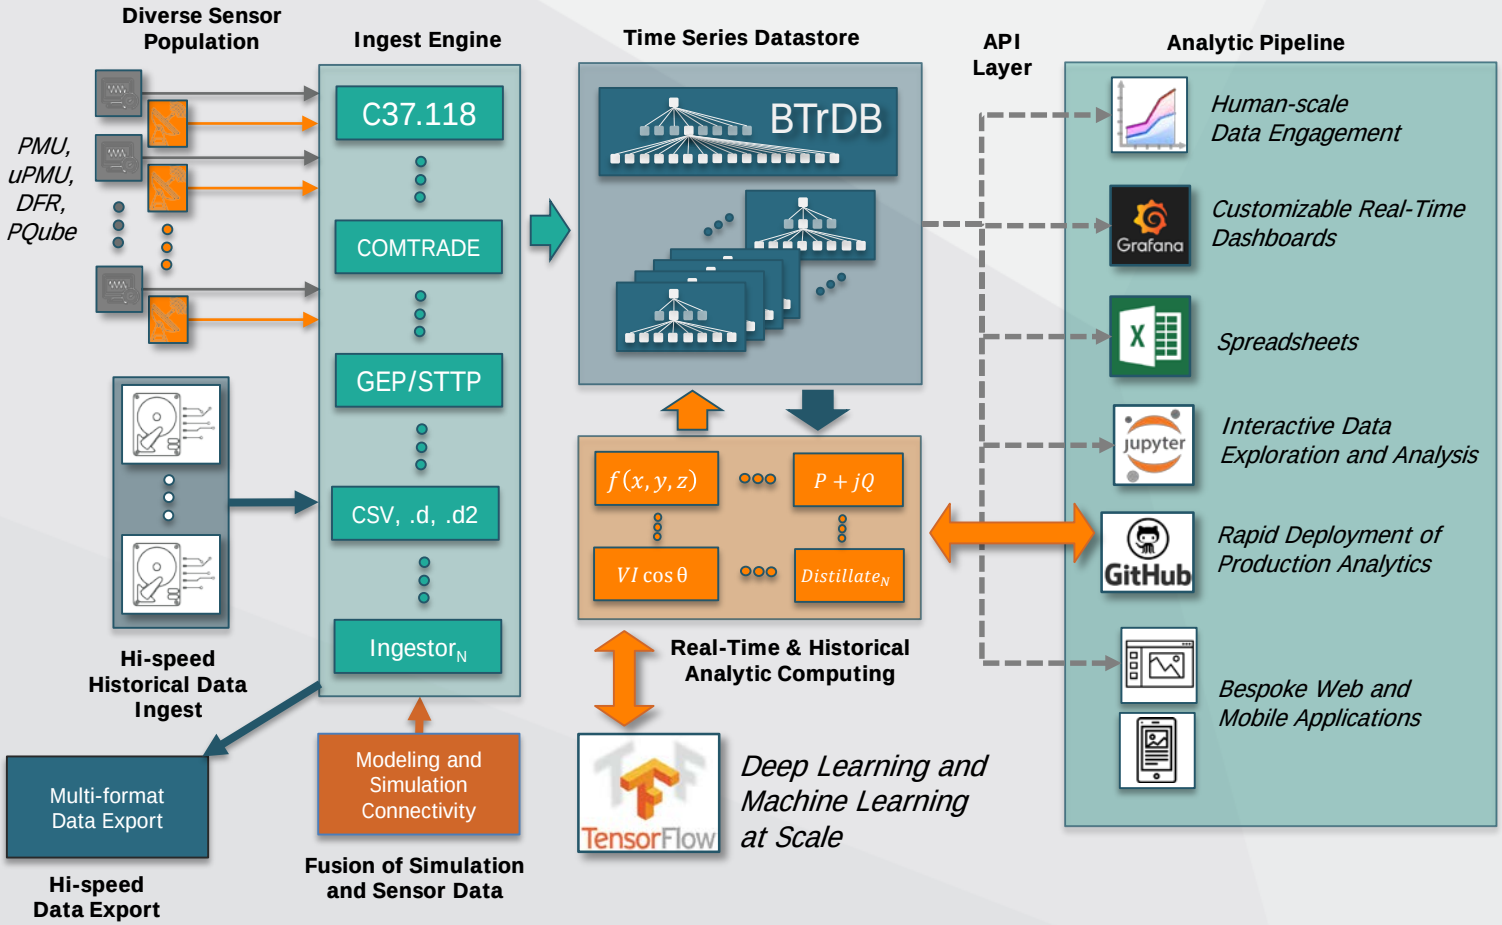

sean@pingthings.ai