



Integral
Ad Science

**PATTERNS OF RELIABLE IN-STREAM
PROCESSING @ SCALE**

Alexey Kharlamov



BRAND SAFETY



VIEWABILITY



AD FRAUD

50+ Billion events/day

20+ Terabytes/day

Up to 1.5M event/second

10+ Data Centers

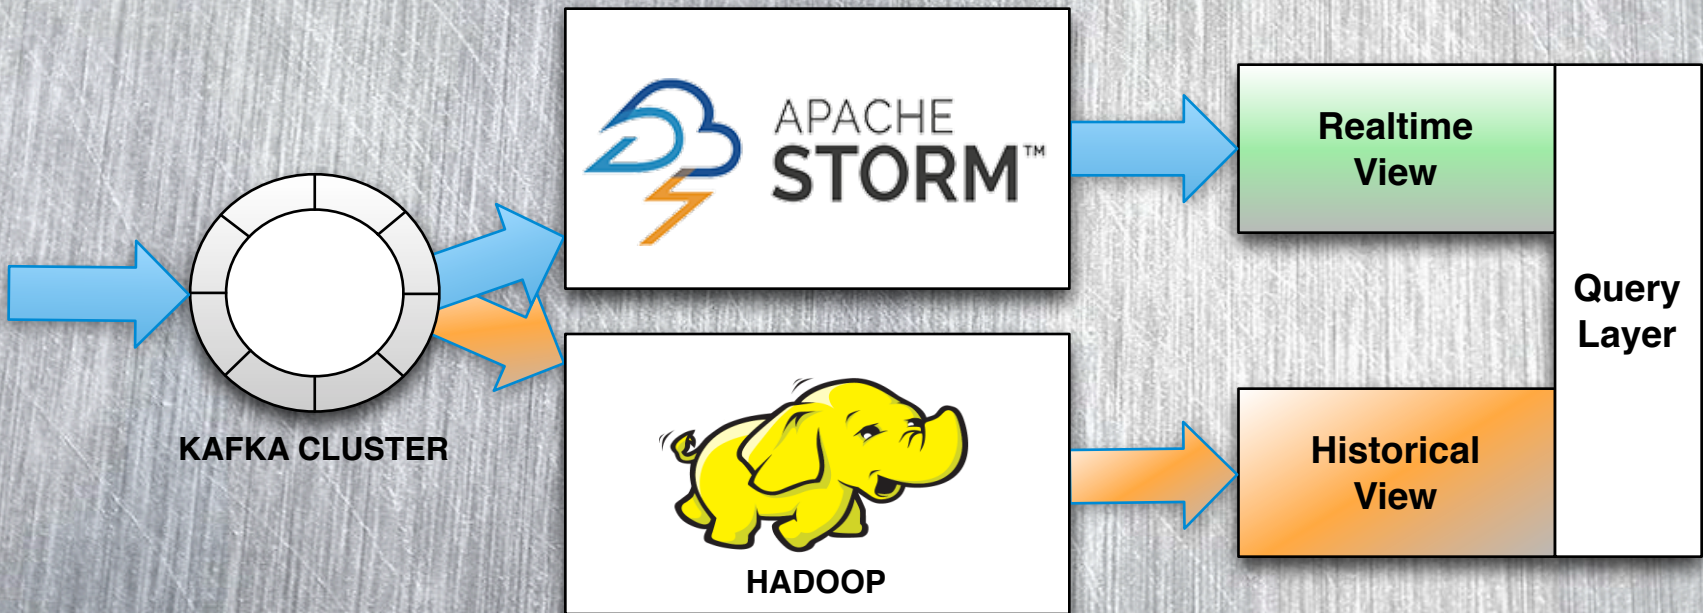
The screenshot shows a monitoring dashboard for a system named 'BATMAN'. The dashboard displays a grid of data points for 10 data centers: NJ01 (New Jersey), AMI (Amsterdam), DAL (Dallas), SJE (San Jose), AU (Australia), AZJ (Japan), AZV (Virginia), AZI (Ireland), and STAG (Stage). The services being monitored are Tailer, Kafka, Join, Fraud, Camus, and Cmprsn. The 'Fraud' row is partially obscured by a Pac-Man icon. The 'Cmprsn' row shows values ranging from 101 to 105 across the data centers.

	NJ01 New Jersey	AMI Amsterdam	DAL Dallas	SJE San Jose	AU Australia	AZJ Japan	AZV Virginia	AZI Ireland	STAG Stage
Tailer									N/A
Kafka									
Join									
Fraud									
Camus	15	16	18	14	15	15	15	15	16
Cmprsn	101	102	101	102	105	101	101	101	101

Why did we decide to go real-time?

LAMBDA Architecture

50 servers



300 servers



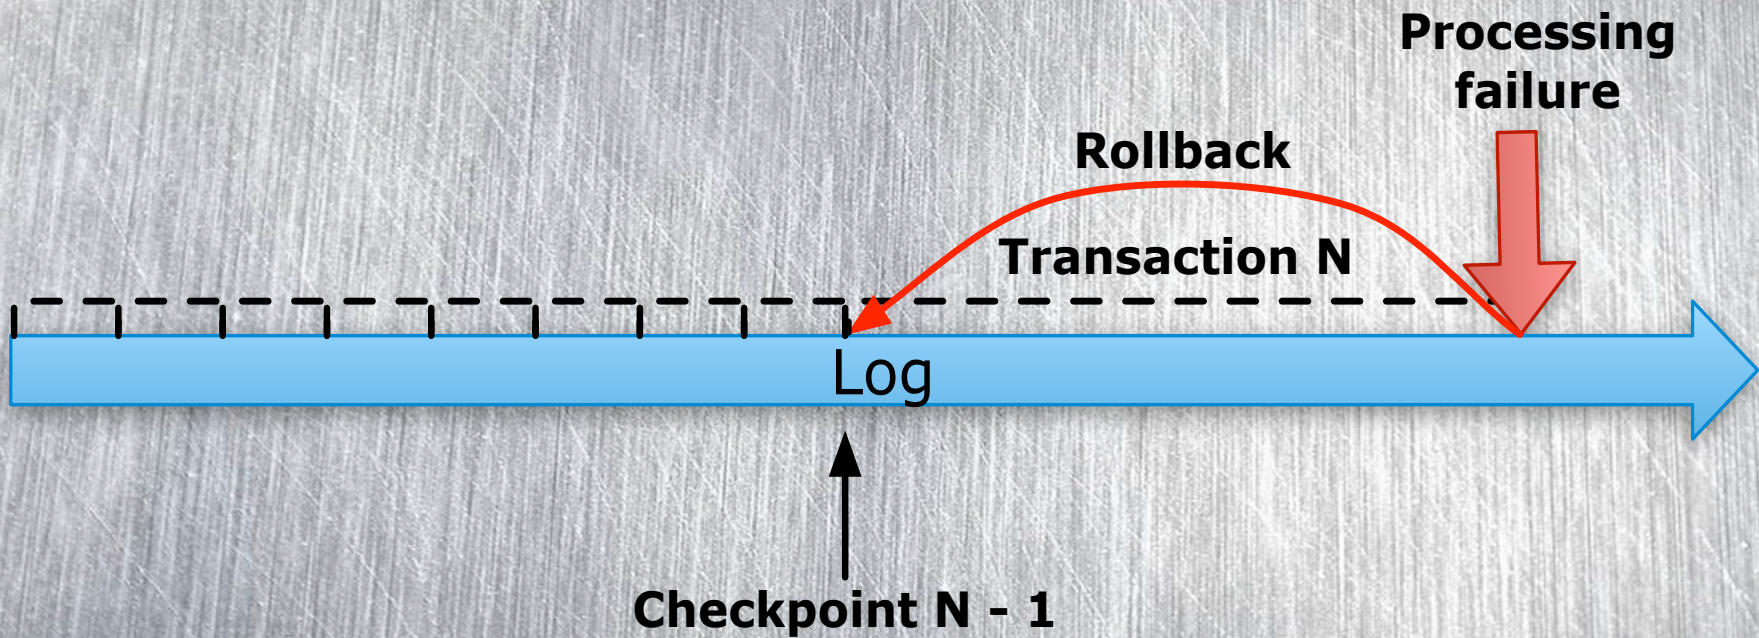


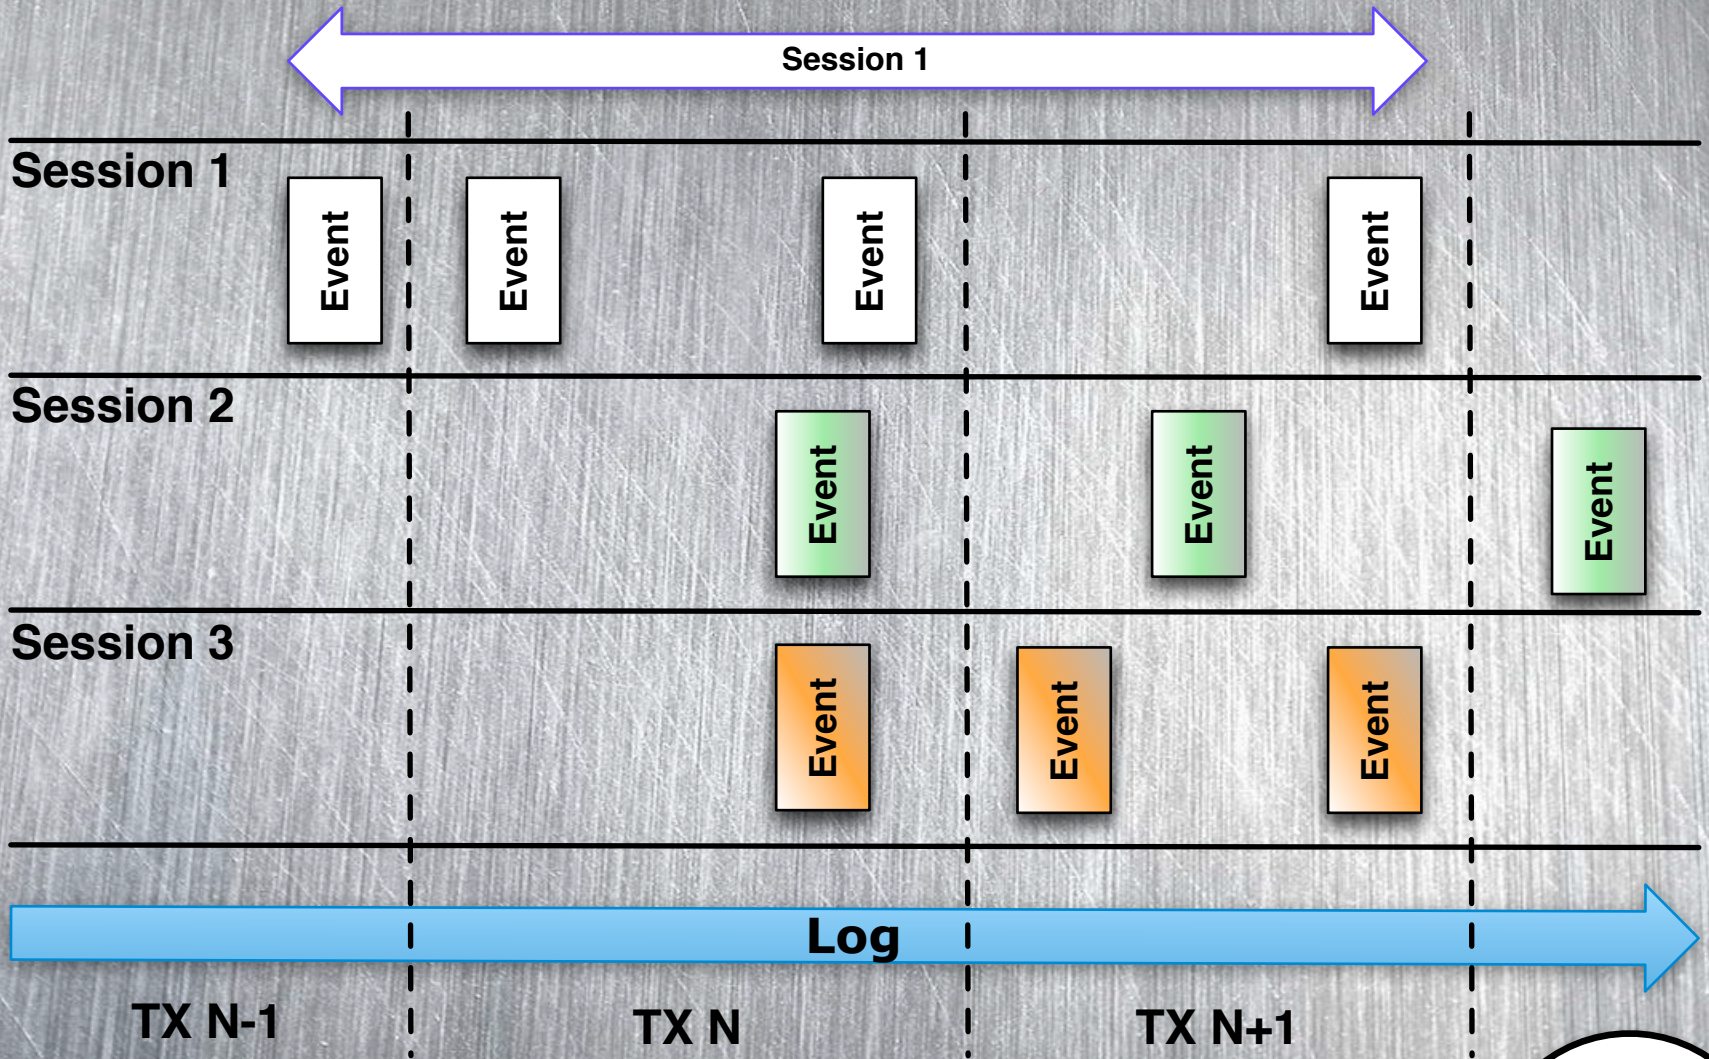
Could we lower costs?




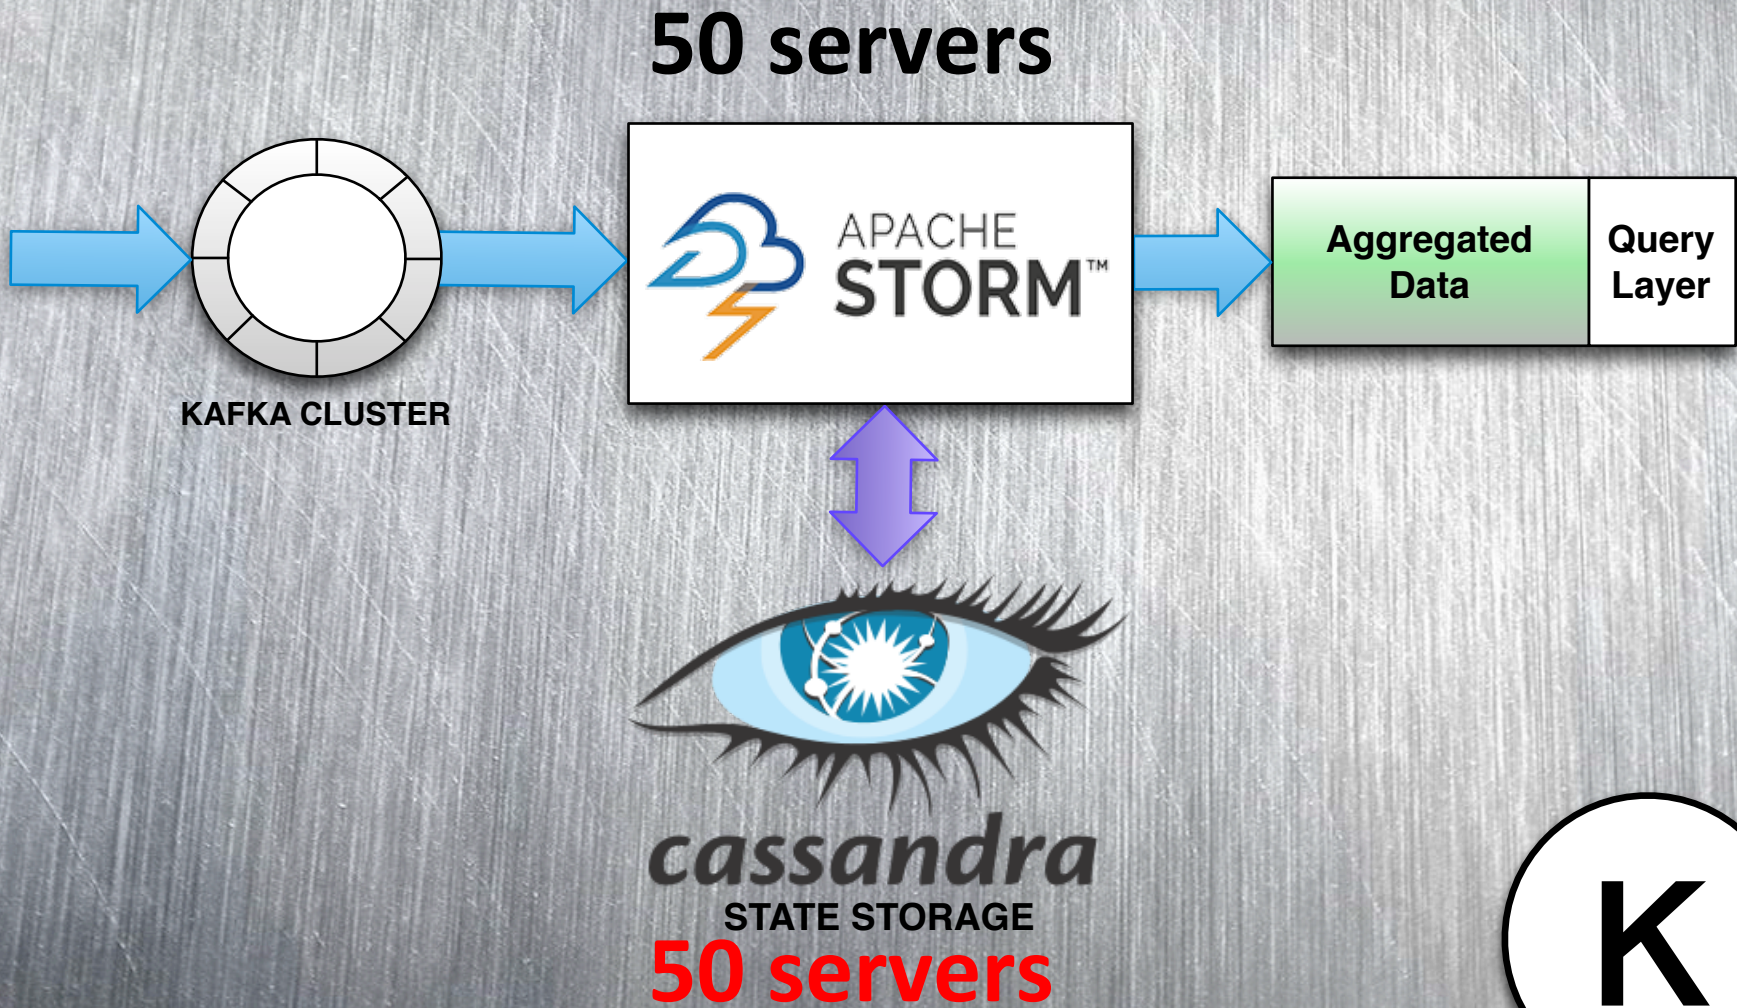
KAPPA Architecture








K

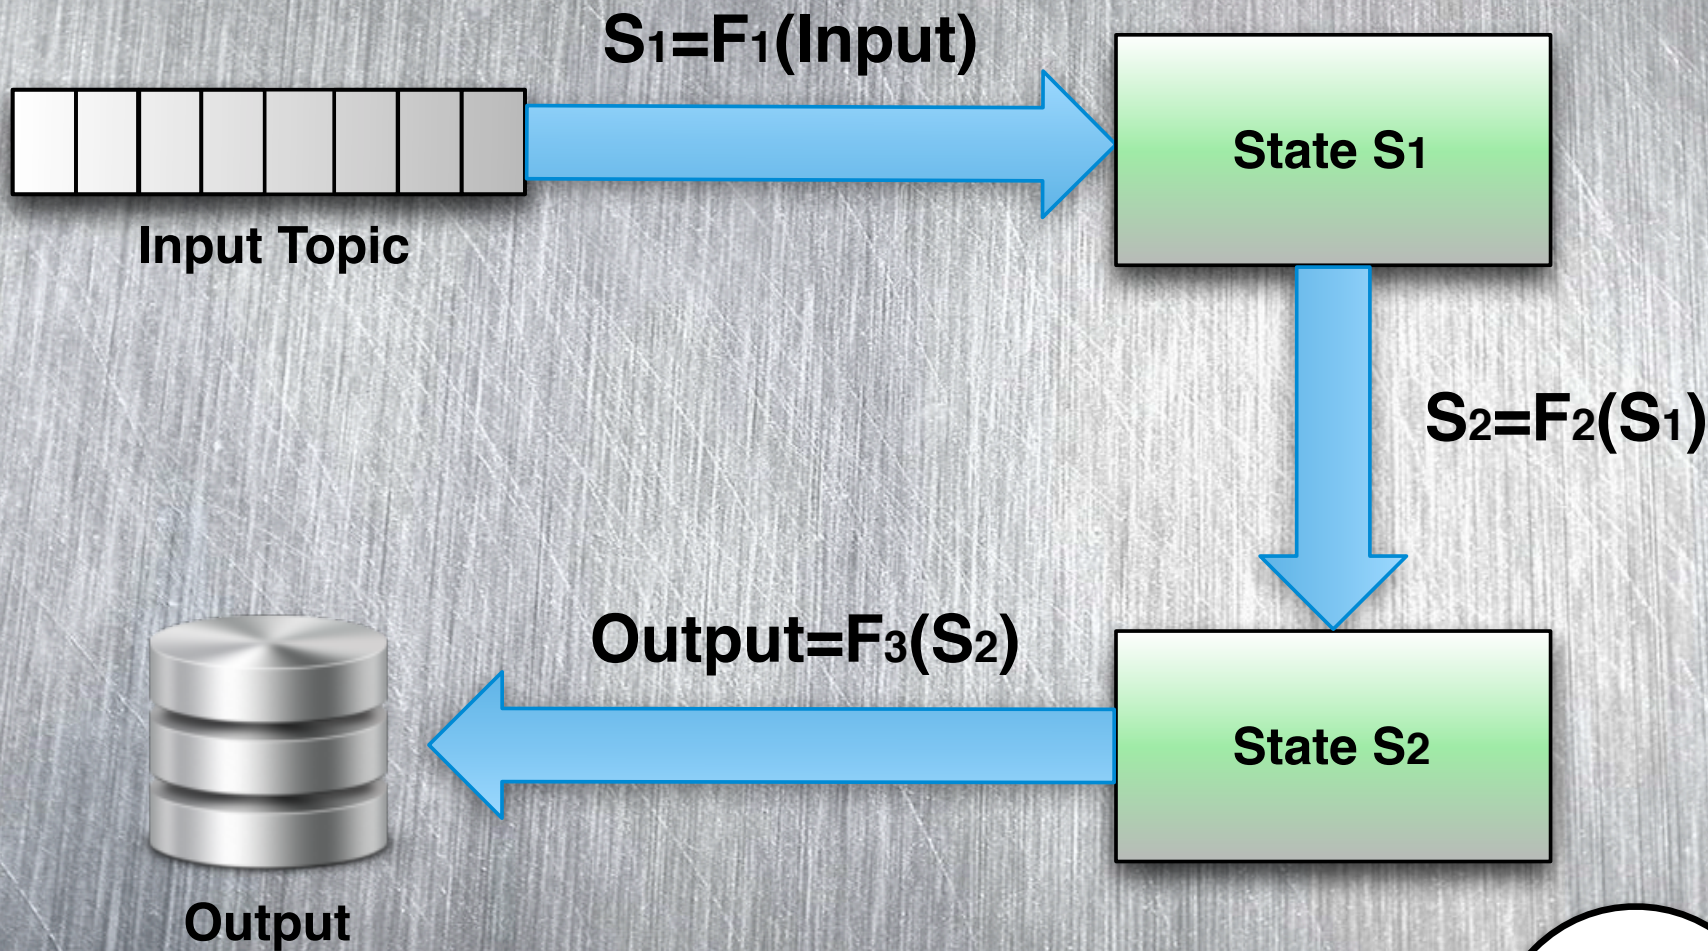


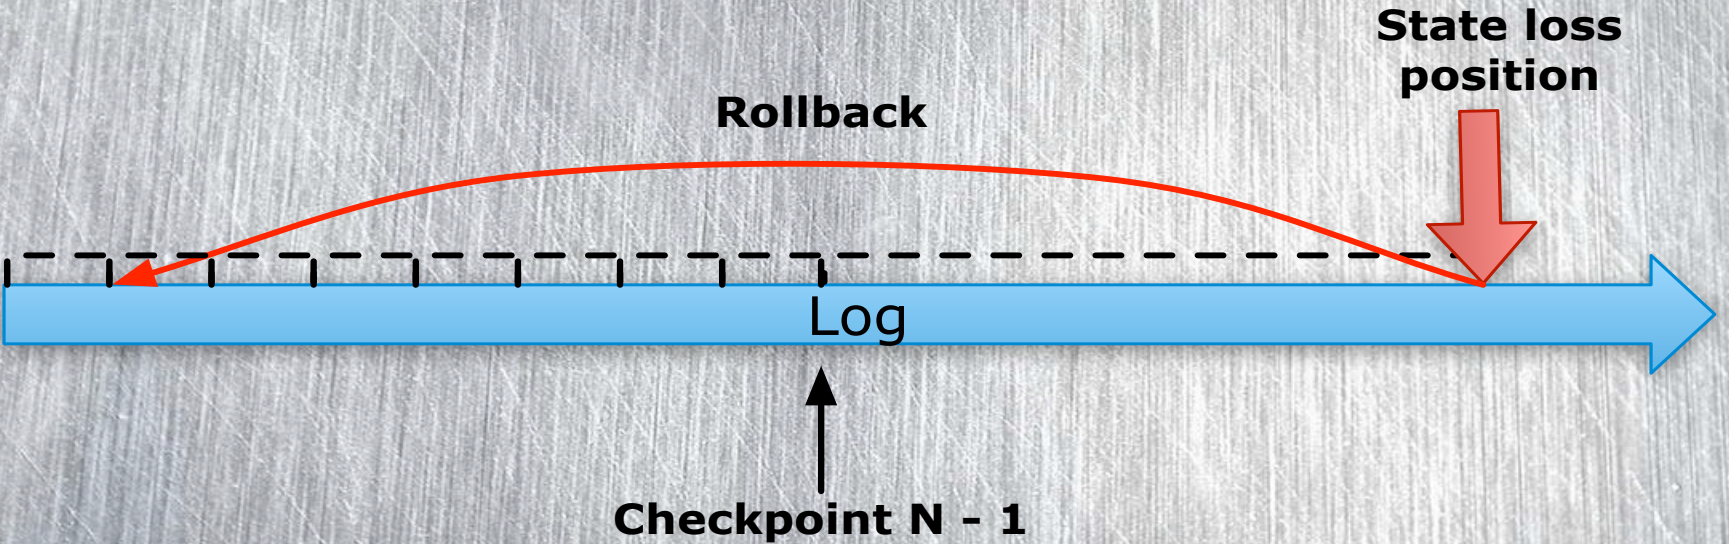




All data transformations are
idempotent and state is volatile.







K

Are we all set?



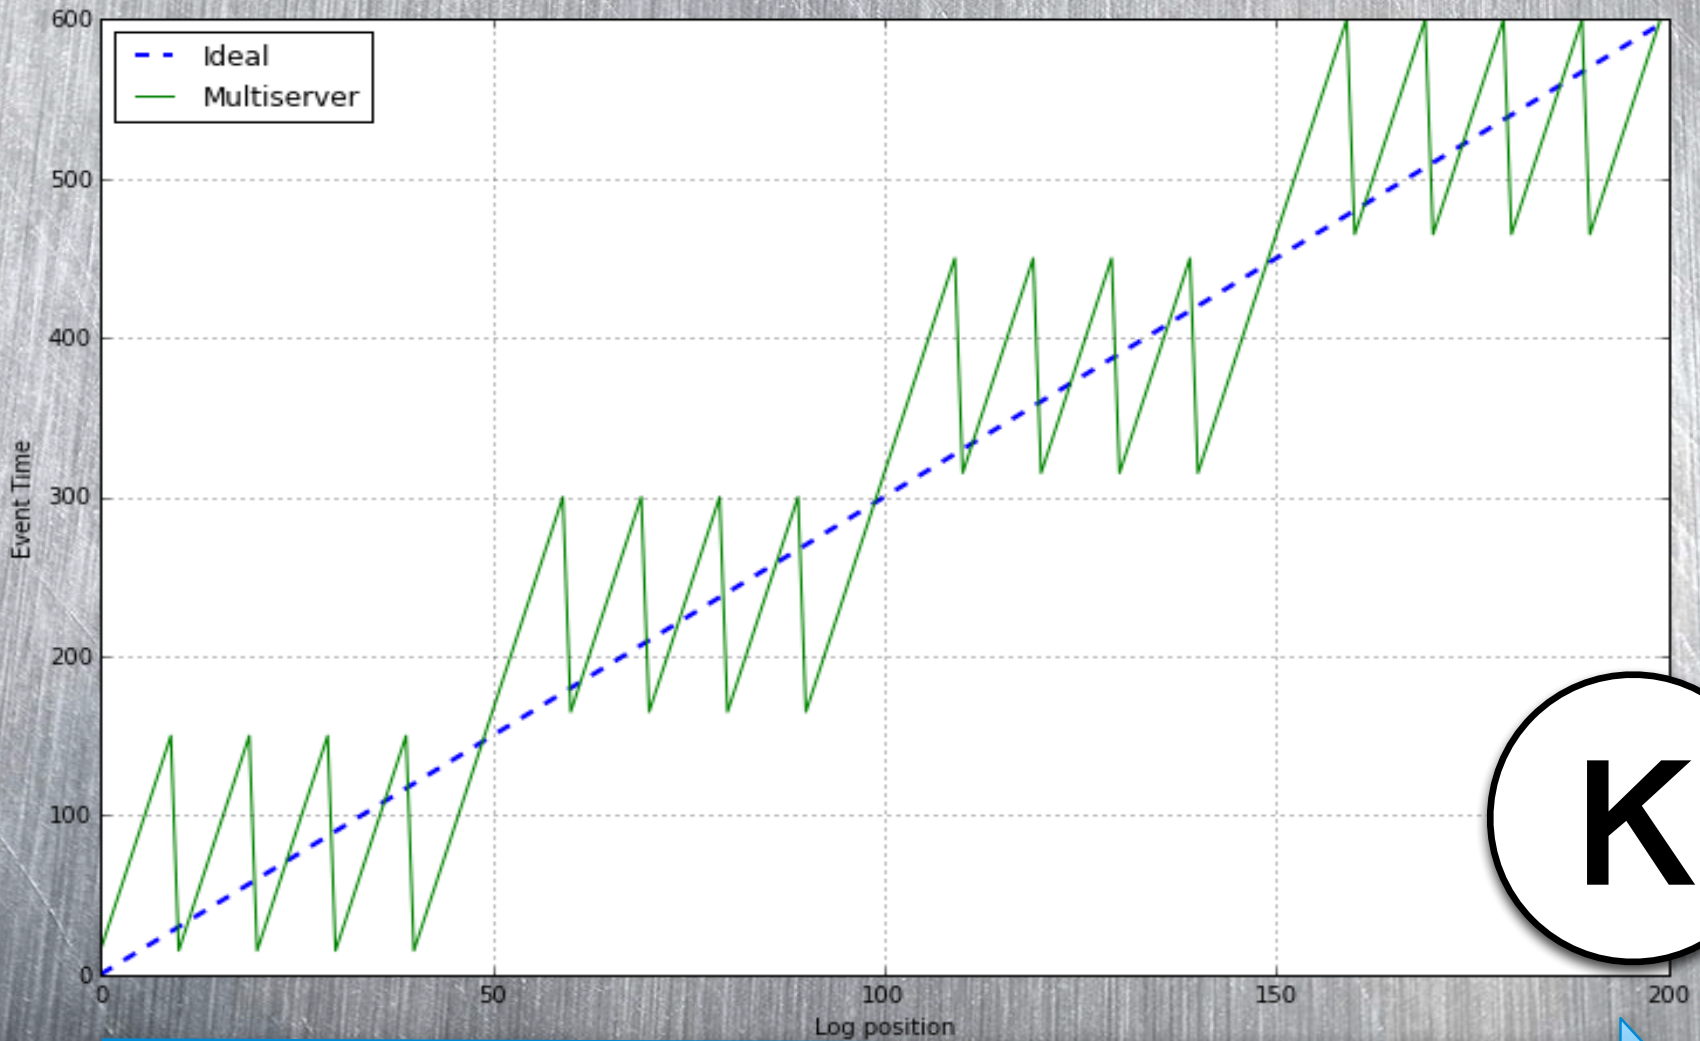
Oops! Something wrong with time!



Event Time



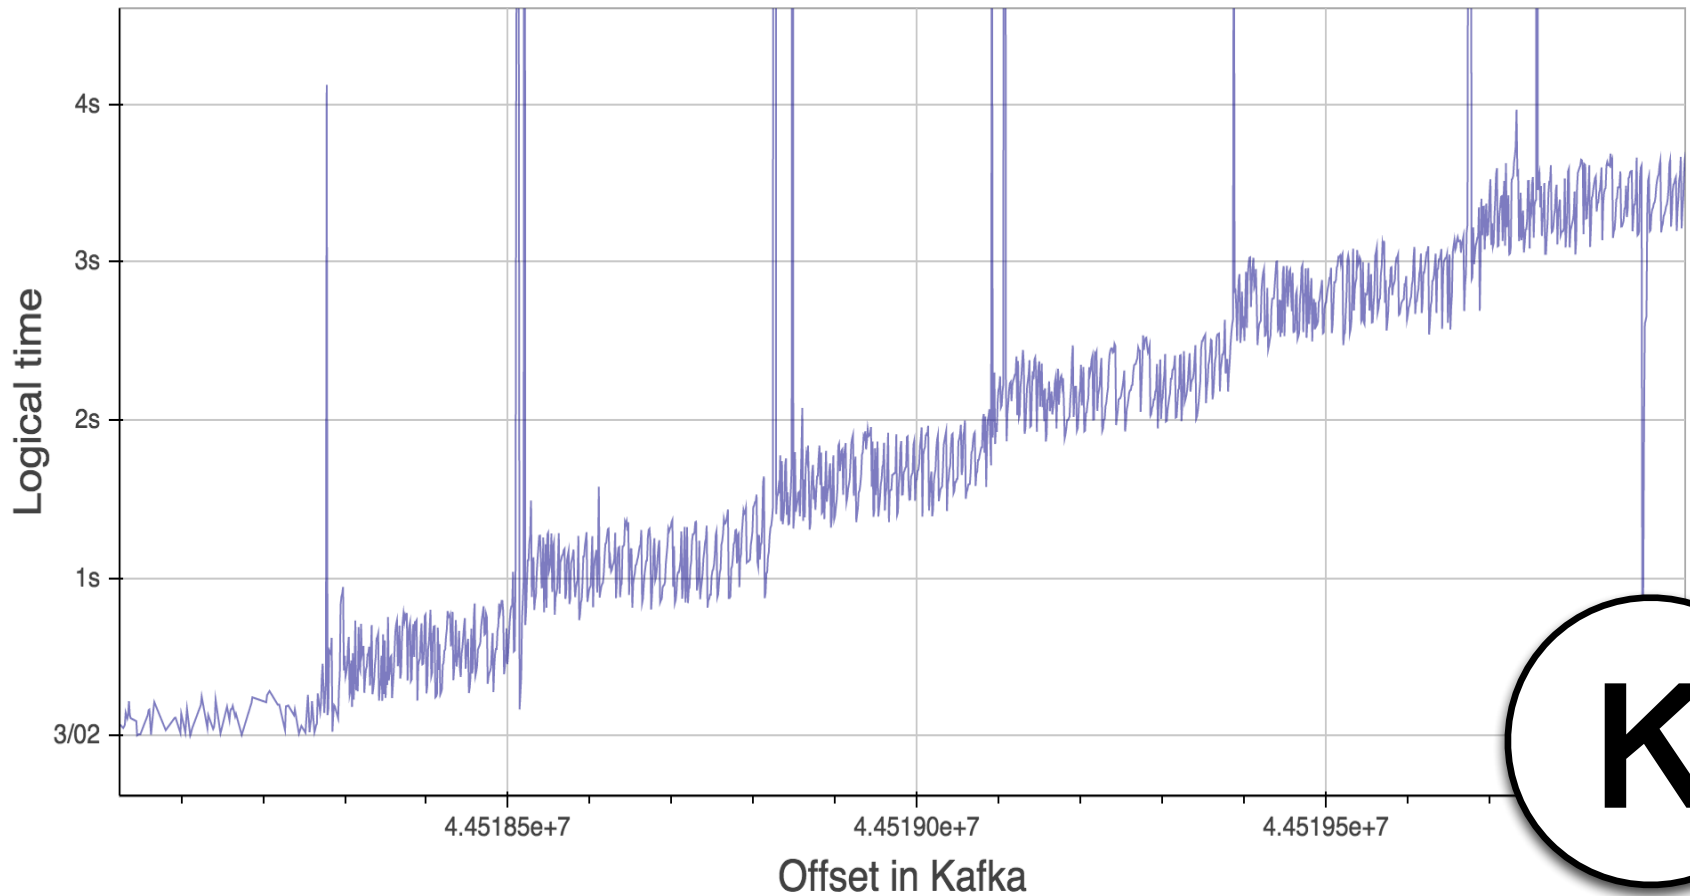
Actually... More like that



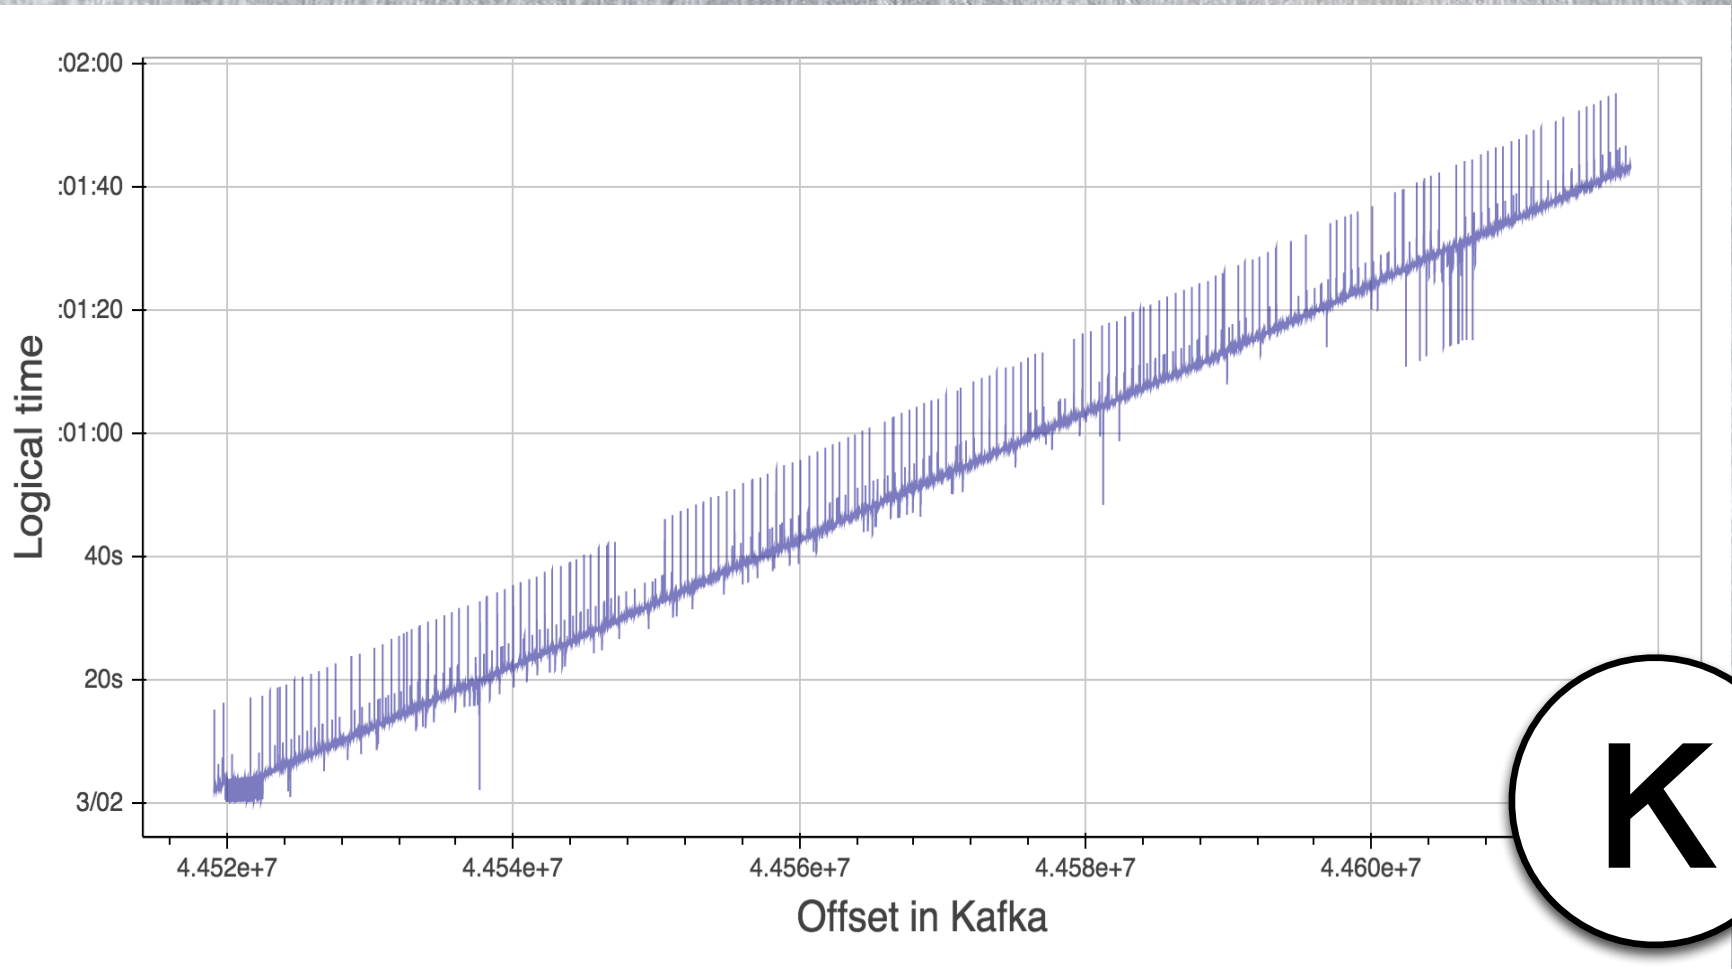
K

Log

In reality...



In reality...

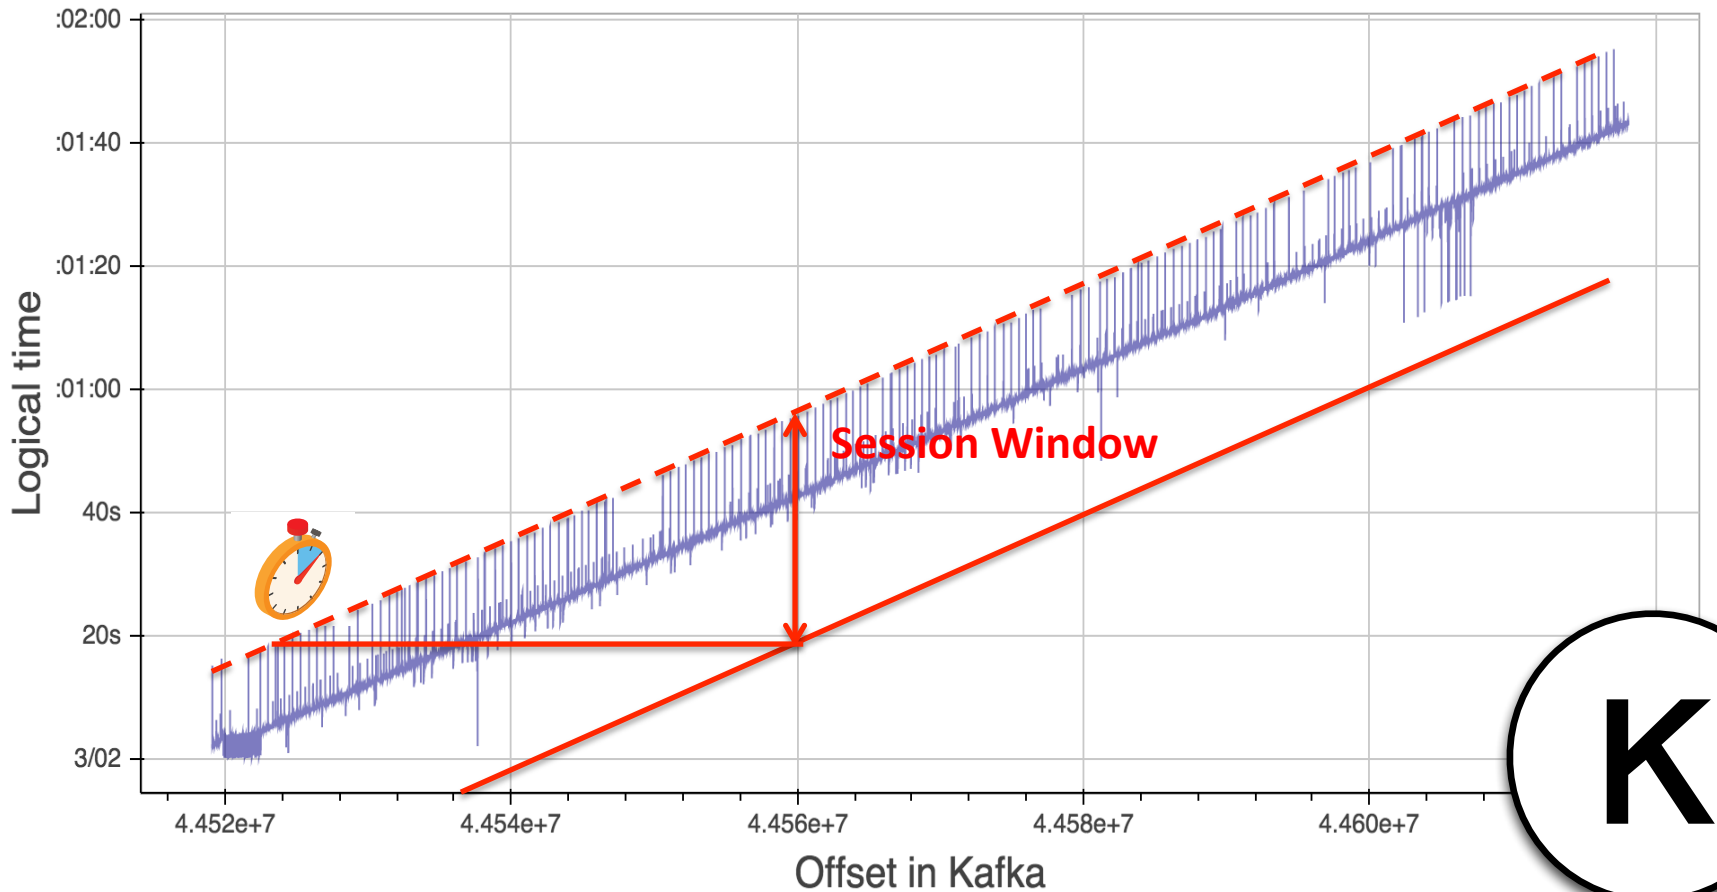




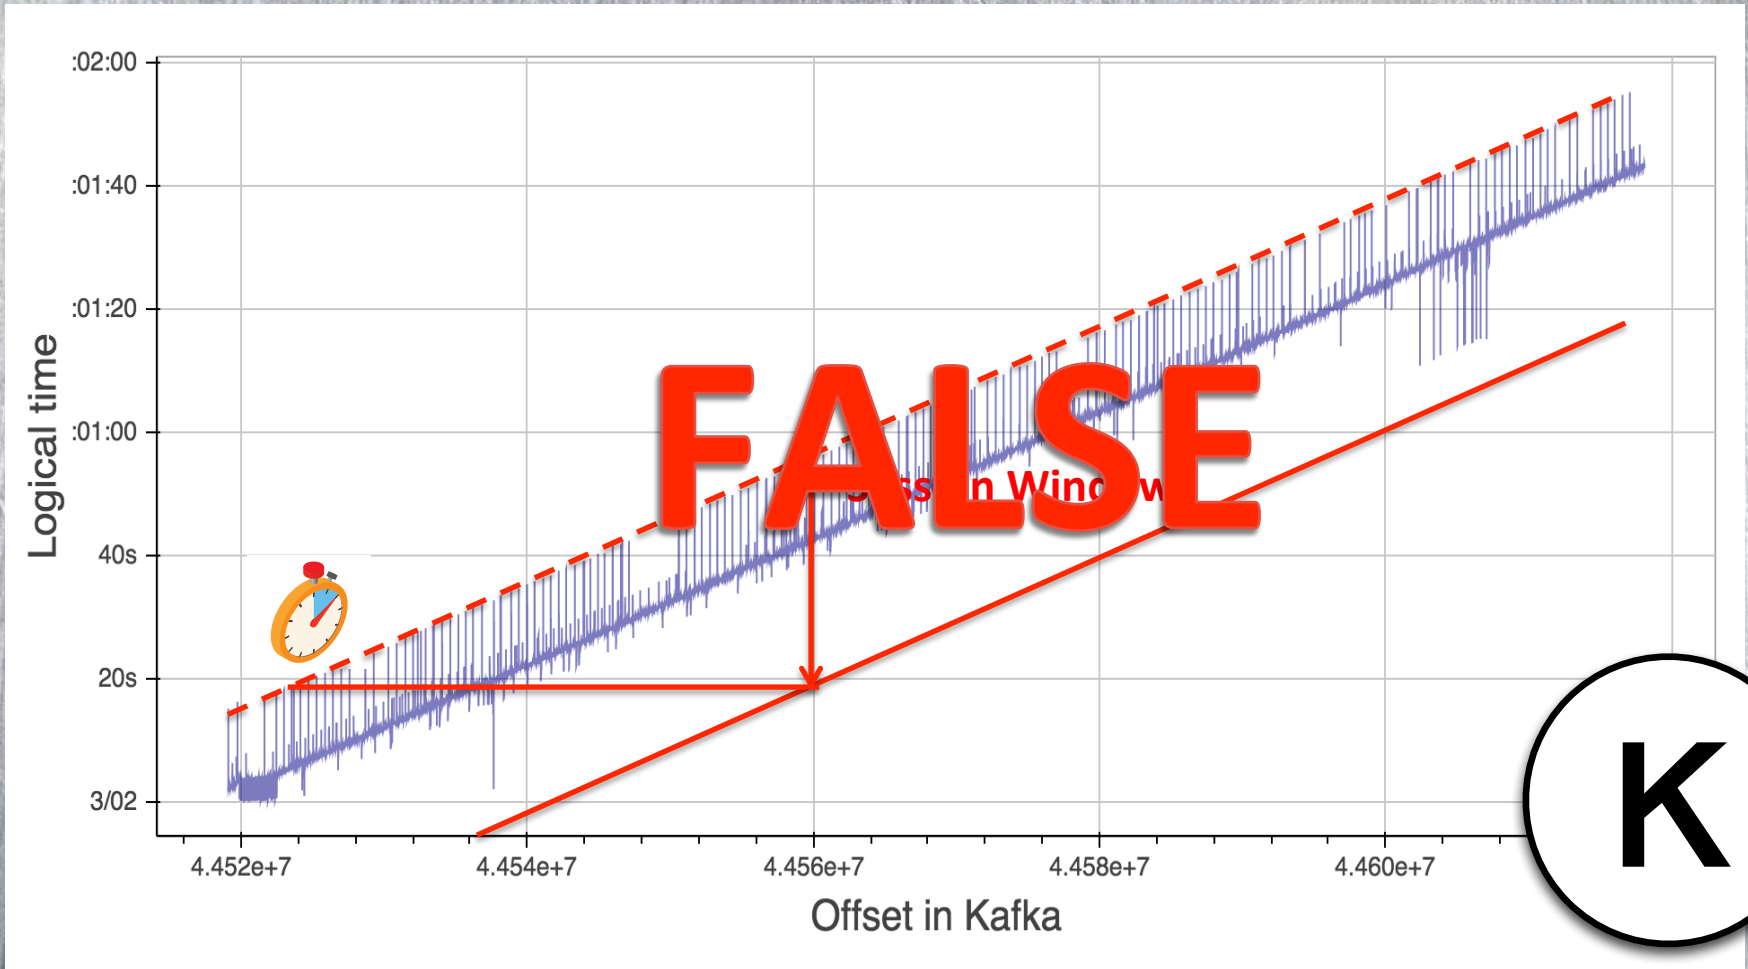
Logical time is cumulative maximum of
observed event timestamps




In reality...



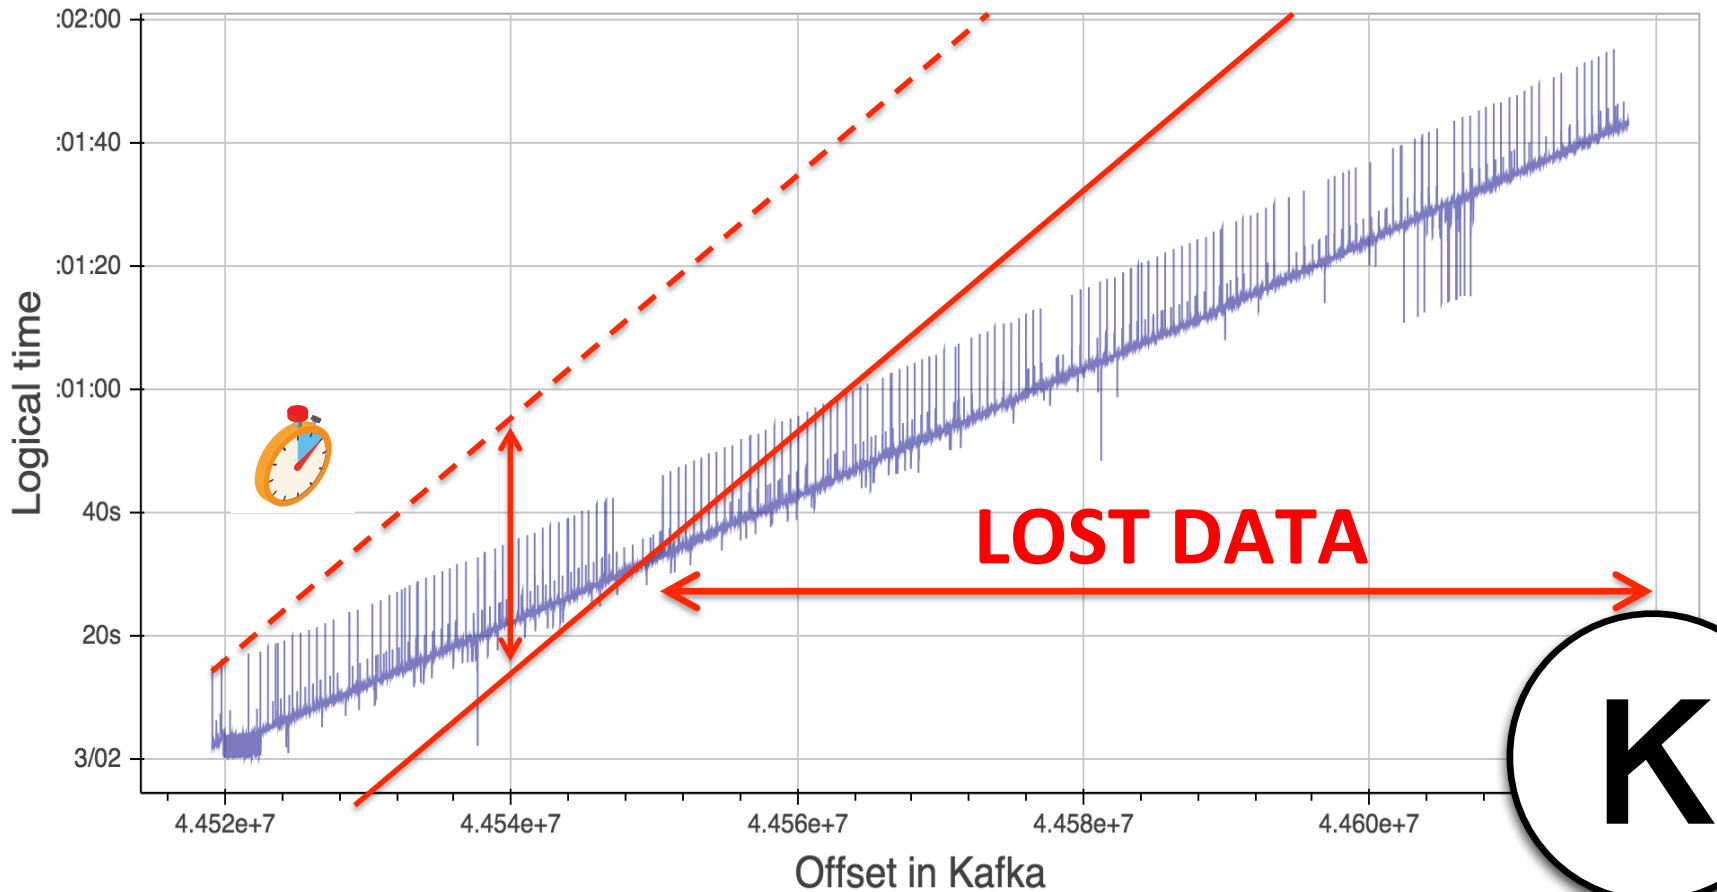
In reality...



Rogue Broker



Rogue Broker



K

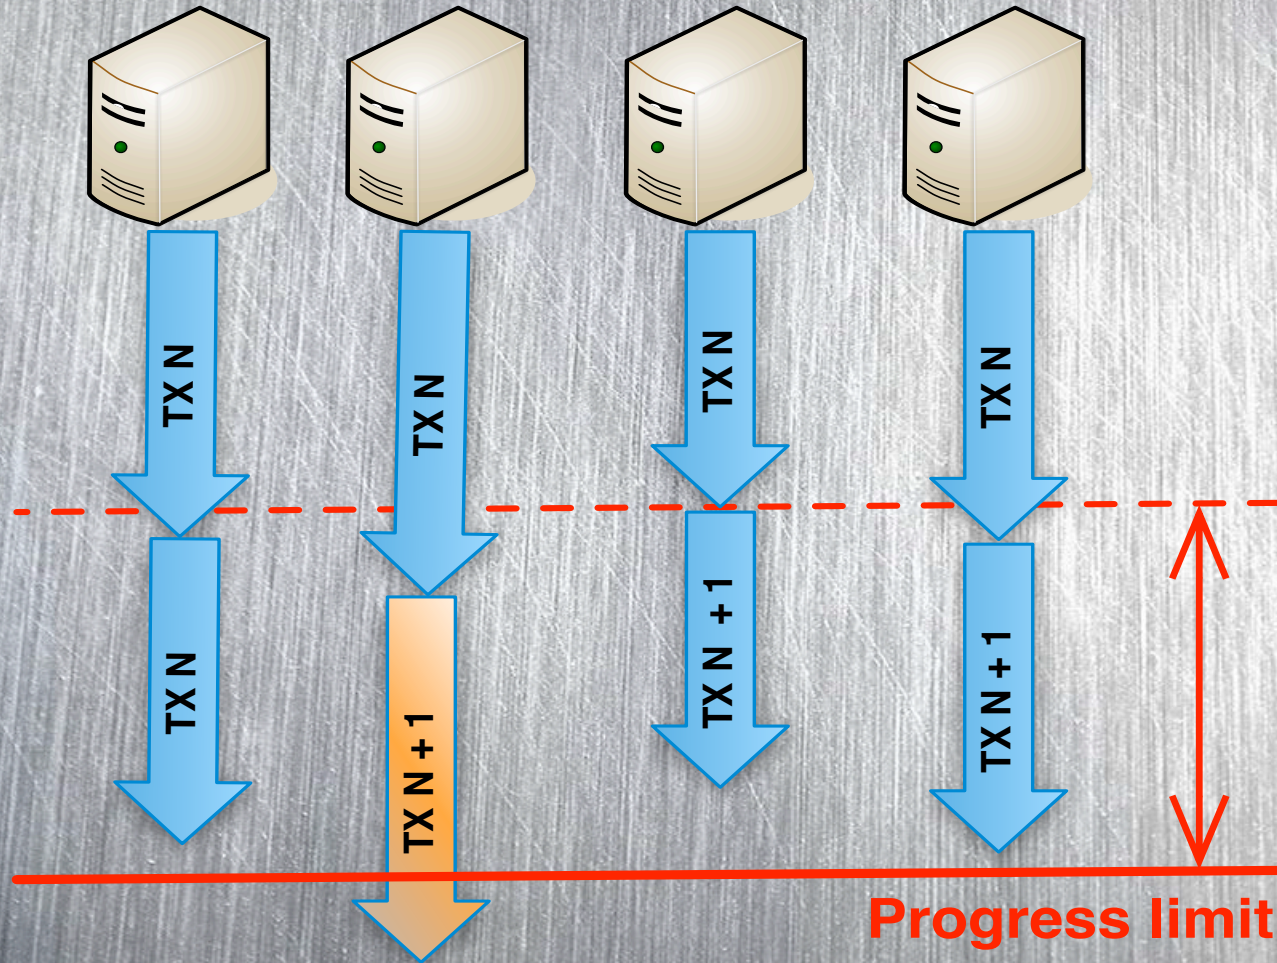
Log



Limit logical time progress across
consumers

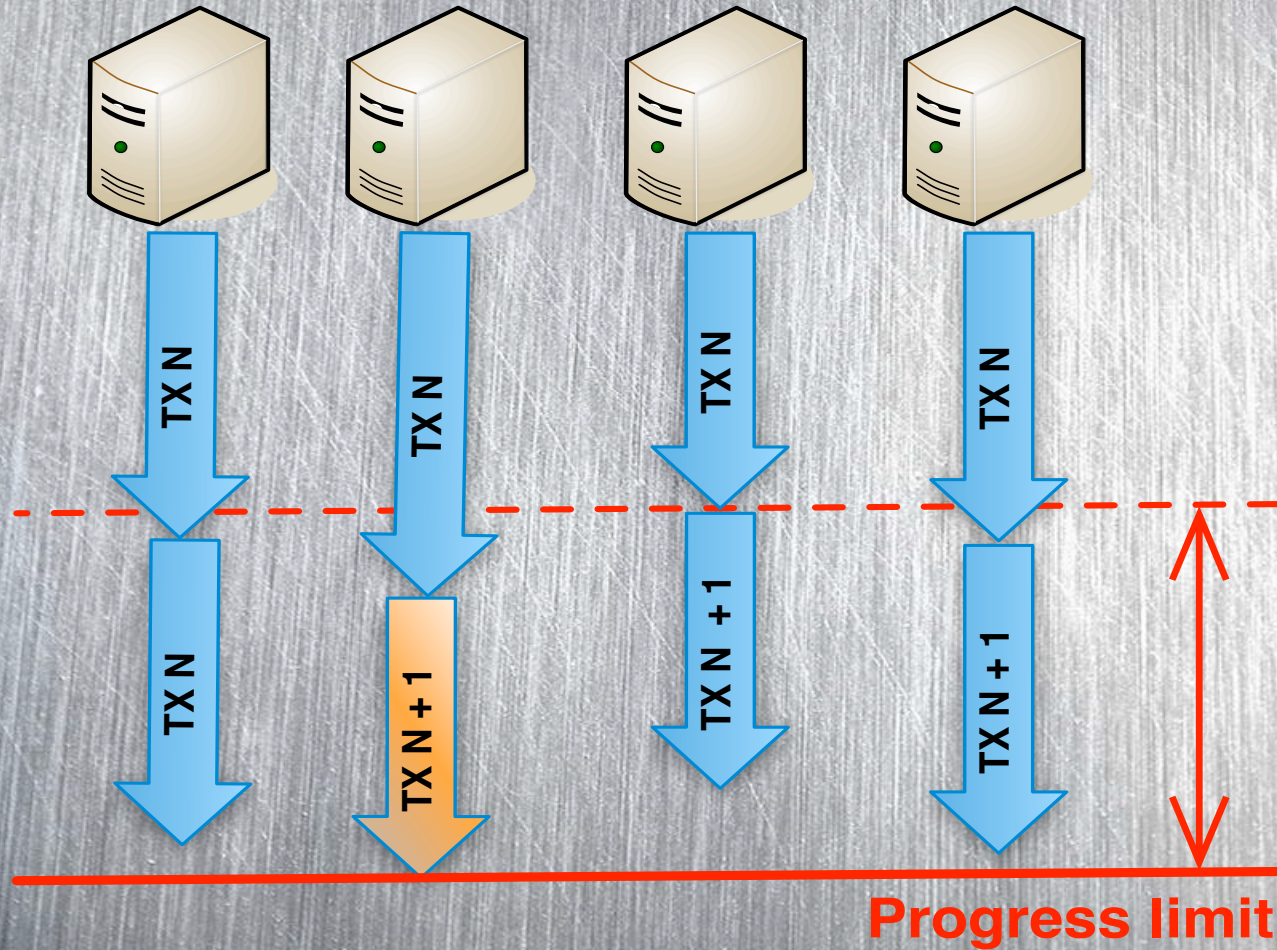


Logical time progress control



K

Logical time progress control





Global synchronization







Assume data distributed evenly

Zipf distribution





Hardware matters!

Questions?



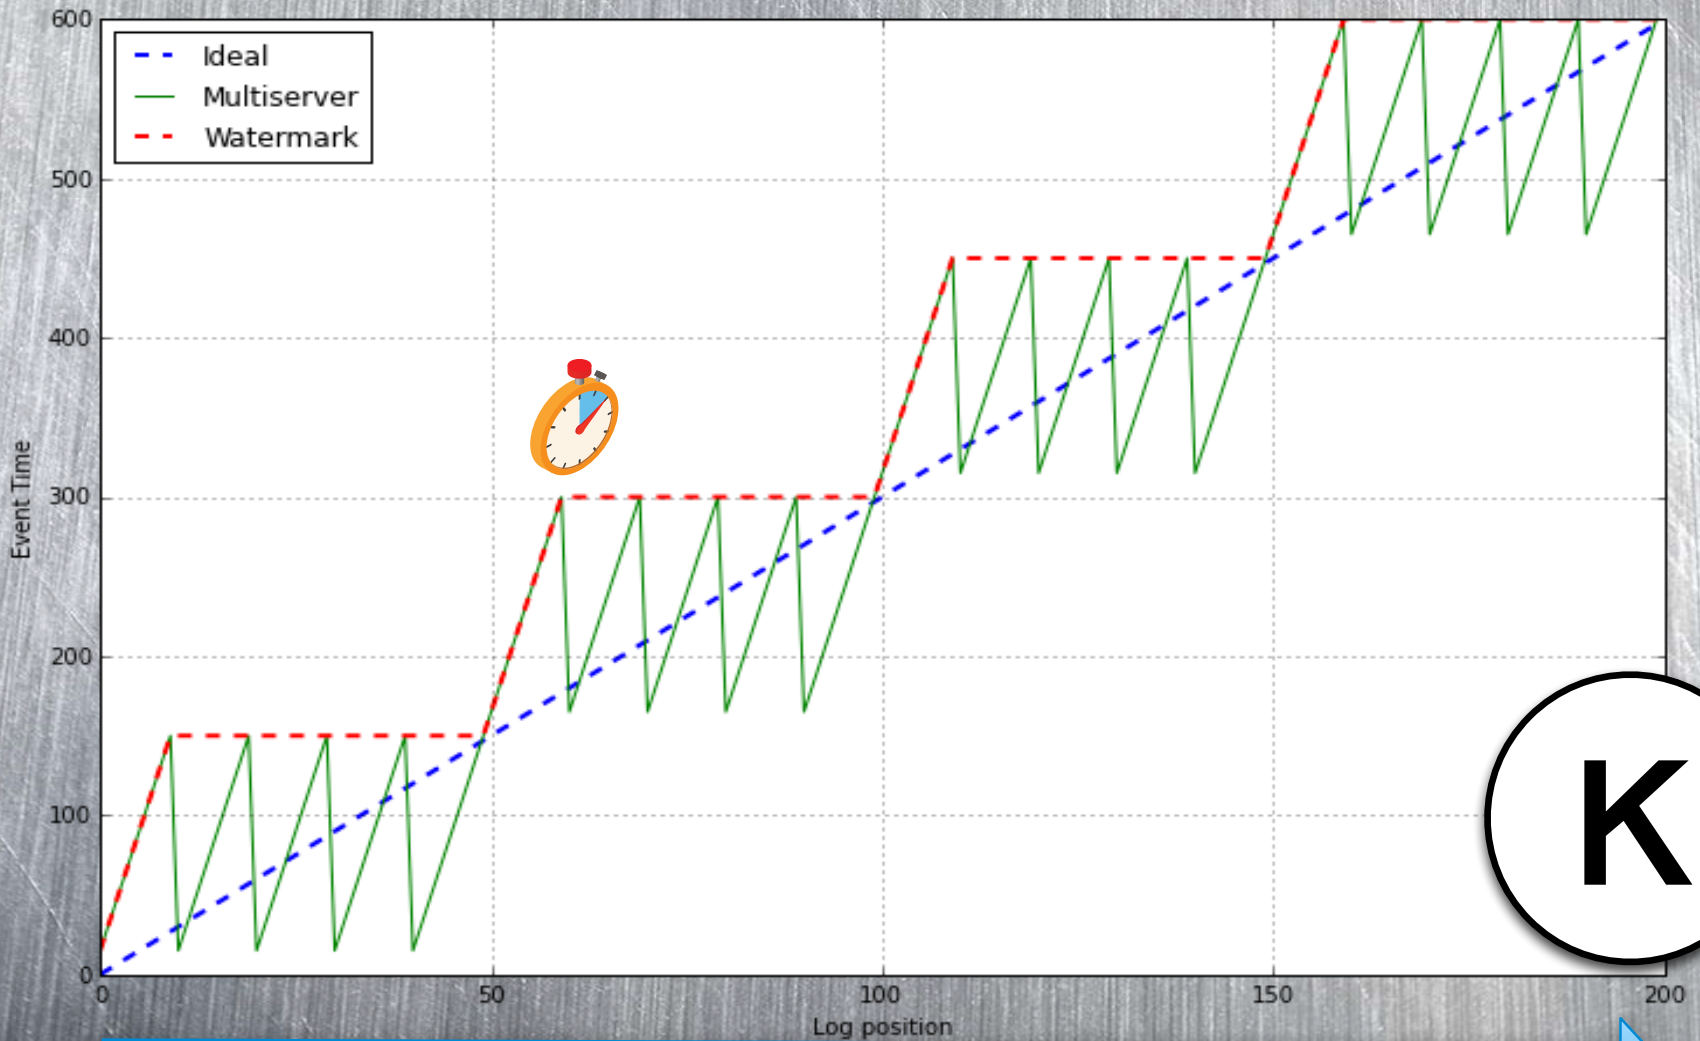
Integral
Ad Science

THANK YOU!

@aih1013

alexey@kharlamov.biz

Logical Time (watermark)



Log

K