

Democratizing Time Sync to Level the Playing Field

Balaji Prabhakar

**Professor of Computer Science and Electrical Engineering, Stanford University
Co-Founder, Tick Tock Networks**

Overview

Consider problems in financial trading caused by “random delay jitters”

- Fairness: In-order execution of transactions
- Fair access to data: Market participants should receive order book data simultaneously
- Protection from front running

And ask

- How can nanosecond-level clock sync at scale remedy these problems?
- How to obtain accurate clock sync at scale in *commodity* networks *without* special hardware or infrastructure investment?

Focus on both

- Current on-prem exchange networks
 - Future exchange networks in ad hoc, heterogeneous cloud environments
- Including some demos

A new clock synchronization system

Huygens: Originated at Stanford as an ML-based clock sync solution

- As part of the Self-Programming Networks research program
- Paper: <https://www.usenix.org/node/211256>
- Main ideas and results previously presented at STAC

Developed for commercial application as TTCS by Tick Tock Networks

- Software-based solution scales to 1000s of nodes
- Nanosecond-level clock sync accuracy with NIC hardware timestamps
- Microsecond-level accurate with CPU/VM/container timestamps
- Deploys on-prem or in one or more regions of the cloud

DEMO of Tick Tock: Single Region of the Cloud

Tick Tock in the cloud

No active demo runs

Start new run



Applications to Financial Trading

KREMLIN'S LINKS TO BREXIT PUSH GET A NEW LOOK

COURTING BANKROLLERS

Parallels to U.S. Inquiry
as Investigators Focus
on 2016 Meeting

By DAVID D. KIRKPATRICK
and MATTHEW ROSENBERG

LONDON — Arron Banks, a British financier who bankrolled the campaign for Britain to leave the European Union, has long bragged about his “boozy six-hour lunch” with the Russian ambassador eight months before the vote.

Some also wondered about Mr. Banks’s Russian-born wife and their custom license plate, X M15 SPY, after the British intelligence agency, MI5. But Mr. Banks always laughed off questions about his ties to the Kremlin.

Now, a leaked record of some of Mr. Banks’s emails suggest that he and his closest adviser had a more engaged relationship with Russian diplomats than he has disclosed.

While Mr. Banks was spending more than eight million British pounds to promote a break with the European Union — an outcome the Russians eagerly hoped for — his contacts at the Russian Embassy in London were opening the door to at least three potentially lucrative investment opportunities in Russian-owned gold or diamond mines.

One of Mr. Banks’s business partners, and a fellow backer of Britain’s exit from the European Union, or Brexit, took the Russians up on at least one of the deals.

The extent of these business discussions, which have not been previously reported, raises new questions about whether the Kremlin sought to reward critical figures in the Brexit campaign. Much as in Washington, where investigations are underway into the possibility that Donald J. Trump’s campaign may have cooperated with the Russians, Britain is now grappling with whether Moscow tried to use its close ties with any British citizens to promote Brexit.

In Washington, the investigators for the special prosecutor, Robert S. Mueller III, and Democrats on the House Intelligence Committee have also obtained records of Mr. Banks’s communications, including some with Russian diplomats and about Russian business deals.

And they have taken a special interest in close ties Mr. Banks and other Brexit leaders built to

Continued on Page A9

Firing of Comey Left Rosenstein Feeling ‘Shaken’ and Exploited

By MICHAEL S. SCHMIDT and ADAM GOLDMAN

WASHINGTON — In the days after the F.B.I. director James B. Comey was fired last year, the deputy attorney general, Rod J. Rosenstein, repeatedly expressed anger about how the White House used him to rationalize the firing, saying the experience damaged his reputation, according to four people familiar with his outbursts. In public, Mr. Rosenstein has shown no hint that he had second thoughts about his role in writing a memo about Mr. Comey’s performance that the White House used to justify firing him. “I wrote it. I believe it. I stand by it,” Mr. Rosenstein said to Congress last year.

But in meetings with law enforcement officials in the chaotic days immediately after Mr.



LEITITIA VANCOR FOR THE NEW YORK TIMES

Battling Tanks, Troops and Twisted Truths
American soldiers in Eastern Europe holding training exercises to combat Russia’s military might and manipulation tactics. Page A7.

Once Deemed Too Radical for Mexico, He Now May Be President

By AZAM AHMED

GUADALAJARA, Mexico — Andrés Manuel López Obrador nodded at the sea of red T-shirts and flag-waving devotees jammed into a plaza in Guadalajara.

Never before had such a crowd welcomed him here. In his previous campaigns for the president’s office, residents of Guadalajara, the wealthy capital of the state of Jalisco, shunned him, considering

his leftist platform too radical.

But this time, only days before one of Mexico’s most important elections in decades, the cheers reflected a nationwide shift — and the ability of Mr. López Obrador to ride it.


Brandishing a deep connection with the poor, built over more than a decade of visits to every corner of this country of 120 million, he has managed a staggering lead ahead of Sunday’s vote.

If the poll numbers bear out on Election Day, Mr. López Obrador — who has promised to sell the

presidential plane and convert the opulent presidential palace into a public park — could win by a landslide, putting a leftist leader in charge of Latin America’s second-largest country for the first time in decades.

He is currently 20 to 30 percentage points ahead of his closest rival, a stunning reversal for a politician whose future was far from clear just a few years ago. But a broad disgust with Mexico’s political establishment has brought

Continued on Page A8



ALFREDO ESTRADA/AGENCE FRANCE PRESSE — GETTY IMAGES

The presidential candidate Andrés Manuel López Obrador during the closing rally of his campaign on Wednesday in Mexico City.

Whittling a New York Minute To 100 Billionths of a Second

By JOHN MARKOFF

SAN FRANCISCO — Computer scientists at Stanford University and Google have created technology that can track time down to 100 billionths of a second. It could be just what Wall Street is looking for.

System engineers at Nasdaq, the New York-based stock exchange, recently began testing an algorithm and software that they hope can synchronize a giant network of computers with that nanosecond precision. They say they have built a prototype, and are in the process of deploying a bigger version.

For an exchange like Nasdaq, such refinement is essential to accurately order the millions of stock trades that are placed on their computer systems every

second. Ultimately, this is about money. With stock trading now dominated by computers that make buying and selling decisions and execute them with blazing speed, keeping that order also means protecting profits. So-called high-frequency trading firms place trades in a fraction of a second, sometimes in a bet that they can move faster than bigger competitors.

The pressure to manage these high-speed trades grows when the stock market becomes more volatile, as it has been in recent months, in part to prevent the fastest traders from taking unfair advantage of slower firms. High-frequency traders typically ac-

Continued on Page A17

In Gunman’s Rage, Something Eerily Familiar for Newsrooms

By TIM ARANGO and JOHN HERRMAN

Years ago, Joe Kieta was out to dinner at a nice restaurant with his wife, celebrating their second anniversary.

Mr. Kieta, then the editor of The Merced Sun-Star in Merced, Calif., had just published a series of articles that led to the ouster of the local district attorney. As he was eating, a friend of the district attorney showed up and asked Mr. Kieta to go outside.

The man “challenged me to a fight right in the middle of this fancy restaurant,” Mr. Kieta, now the editor of The Fresno Bee in California, said on Friday.

The fatal shooting a day earlier at The Capital Gazette in Annapolis, Md., reverberated throughout newsrooms across the country, not only for its tragedy but also for

the familiarity of conflicts like the suspected gunman’s long-running feud with the paper. Many reporters and editors, especially at the local level, have stories of being confronted or harassed by a resident upset by something in the newspaper. Unlike Thursday’s shooting, in which five people were killed, few of these situations end in violence.

The suspect, Jarrod W. Ramos, 38, had hounded the paper for years, after the publication of an article about his conviction in a criminal harassment case involving online threats against a former classmate. He filed lawsuits, posted highly personal comments about reporters online and routinely intimidated violence on social

Continued on Page A14

Whittling a New York Minute To 100 Billionths of a Second

By JOHN MARKOFF

SAN FRANCISCO — Computer scientists at Stanford University and Google have created technology that can track time down to 100 billionths of a second. It could be just what Wall Street is looking for.

System engineers at Nasdaq, the New York-based stock exchange, recently began testing an algorithm and software that they hope can synchronize a giant network of computers with that nanosecond precision. They say they have built a prototype, and are in the process of deploying a bigger version.

For an exchange like Nasdaq, such refinement is essential to accurately order the millions of stock trades that are placed on their computer systems every

second.

Ultimately, this is about money. With stock trading now dominated by computers that make buying and selling decisions and execute them with blazing speed, keeping that order also means protecting profits. So-called high-frequency trading firms place trades in a fraction of a second, sometimes in a bet that they can move faster than bigger competitors.

The pressure to manage these high-speed trades grows when the stock market becomes more volatile, as it has been in recent months, in part to prevent the fastest traders from taking unfair advantage of slower firms. High-frequency traders typically ac-

Continued on Page A17



9/11 Safety Chief Laid to Rest
Chief Ronald R. Spadafora, who kept rescue workers safe at ground zero, died of blood cancer at 63. PAGE A19

Tough Math in School Plan
The mayor’s proposal to drop the test for the city’s elite high schools could upend its elite middle schools. PAGE A18

INTERNATIONAL A4-9

E.U. Reaches Migration Deal
Despite making no commitments, the agreement seemed to be enough to save Chancellor Angela Merkel of Germany while appeasing Italy. PAGE A7

NATIONAL A10-17

Medicaid Work Rule Blocked
A ruling in a Kentucky case may affect other states’ efforts to require the poor to work to get health care. PAGE A16

Riding an Outsider Wave
Midterm votes brought a wave of progressive female candidates powered by strong personal narratives. PAGE A11



Mission: Crispy for an Hour
As delivery services ramp up, one giant company’s goal is to keep your fries succulent for up to 60 minutes. PAGE B1

China Curtails Lending Spree
A bold plan to gain greater global influence by funding big projects is lagging as China grows cautious. PAGE B1

SPORTSSATURDAY D1-6

The N.B.A.’s Decider
LeBron James might join a new team, instantly boosting the fortunes of a city and its N.B.A. franchise. Or he might stay in Cleveland. PAGE D1

ARTS C1-7

In a Pixar Short, a Dark Twist
“Bao,” running with “Incredibles 2,” is about a Chinese dumpling and a mother’s love. The director says, “Part of me wanted to shock audiences.” PAGE C1

EDITORIAL, OP-ED A22-23
Dave Eggers PAGE A23

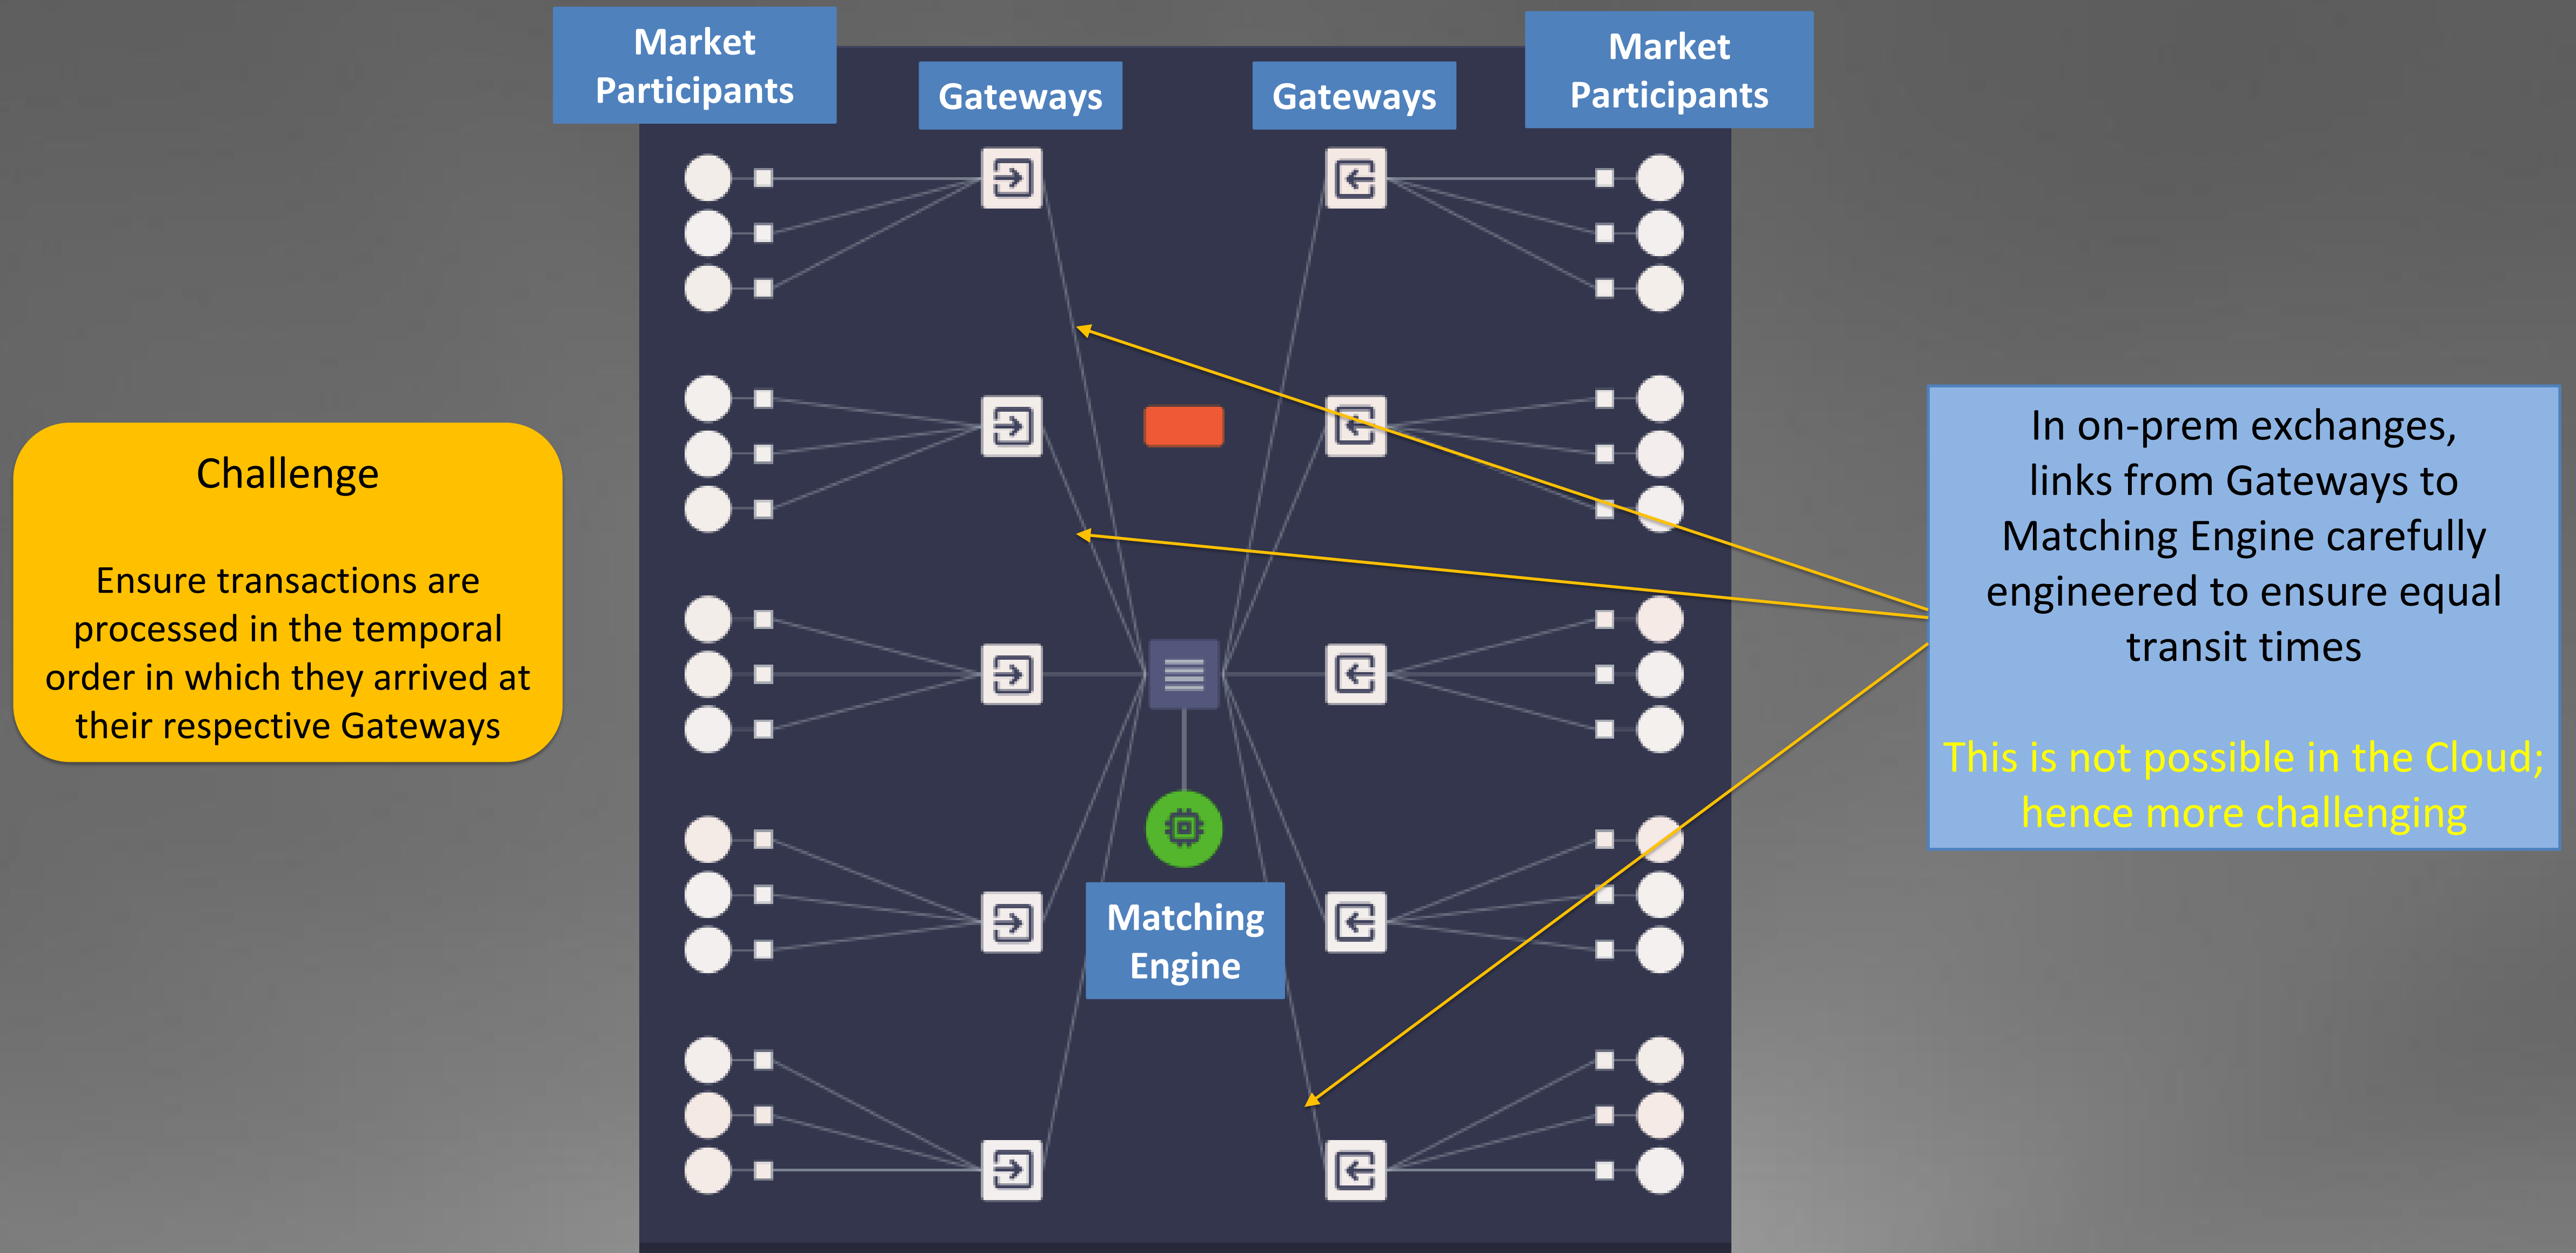


Special Section

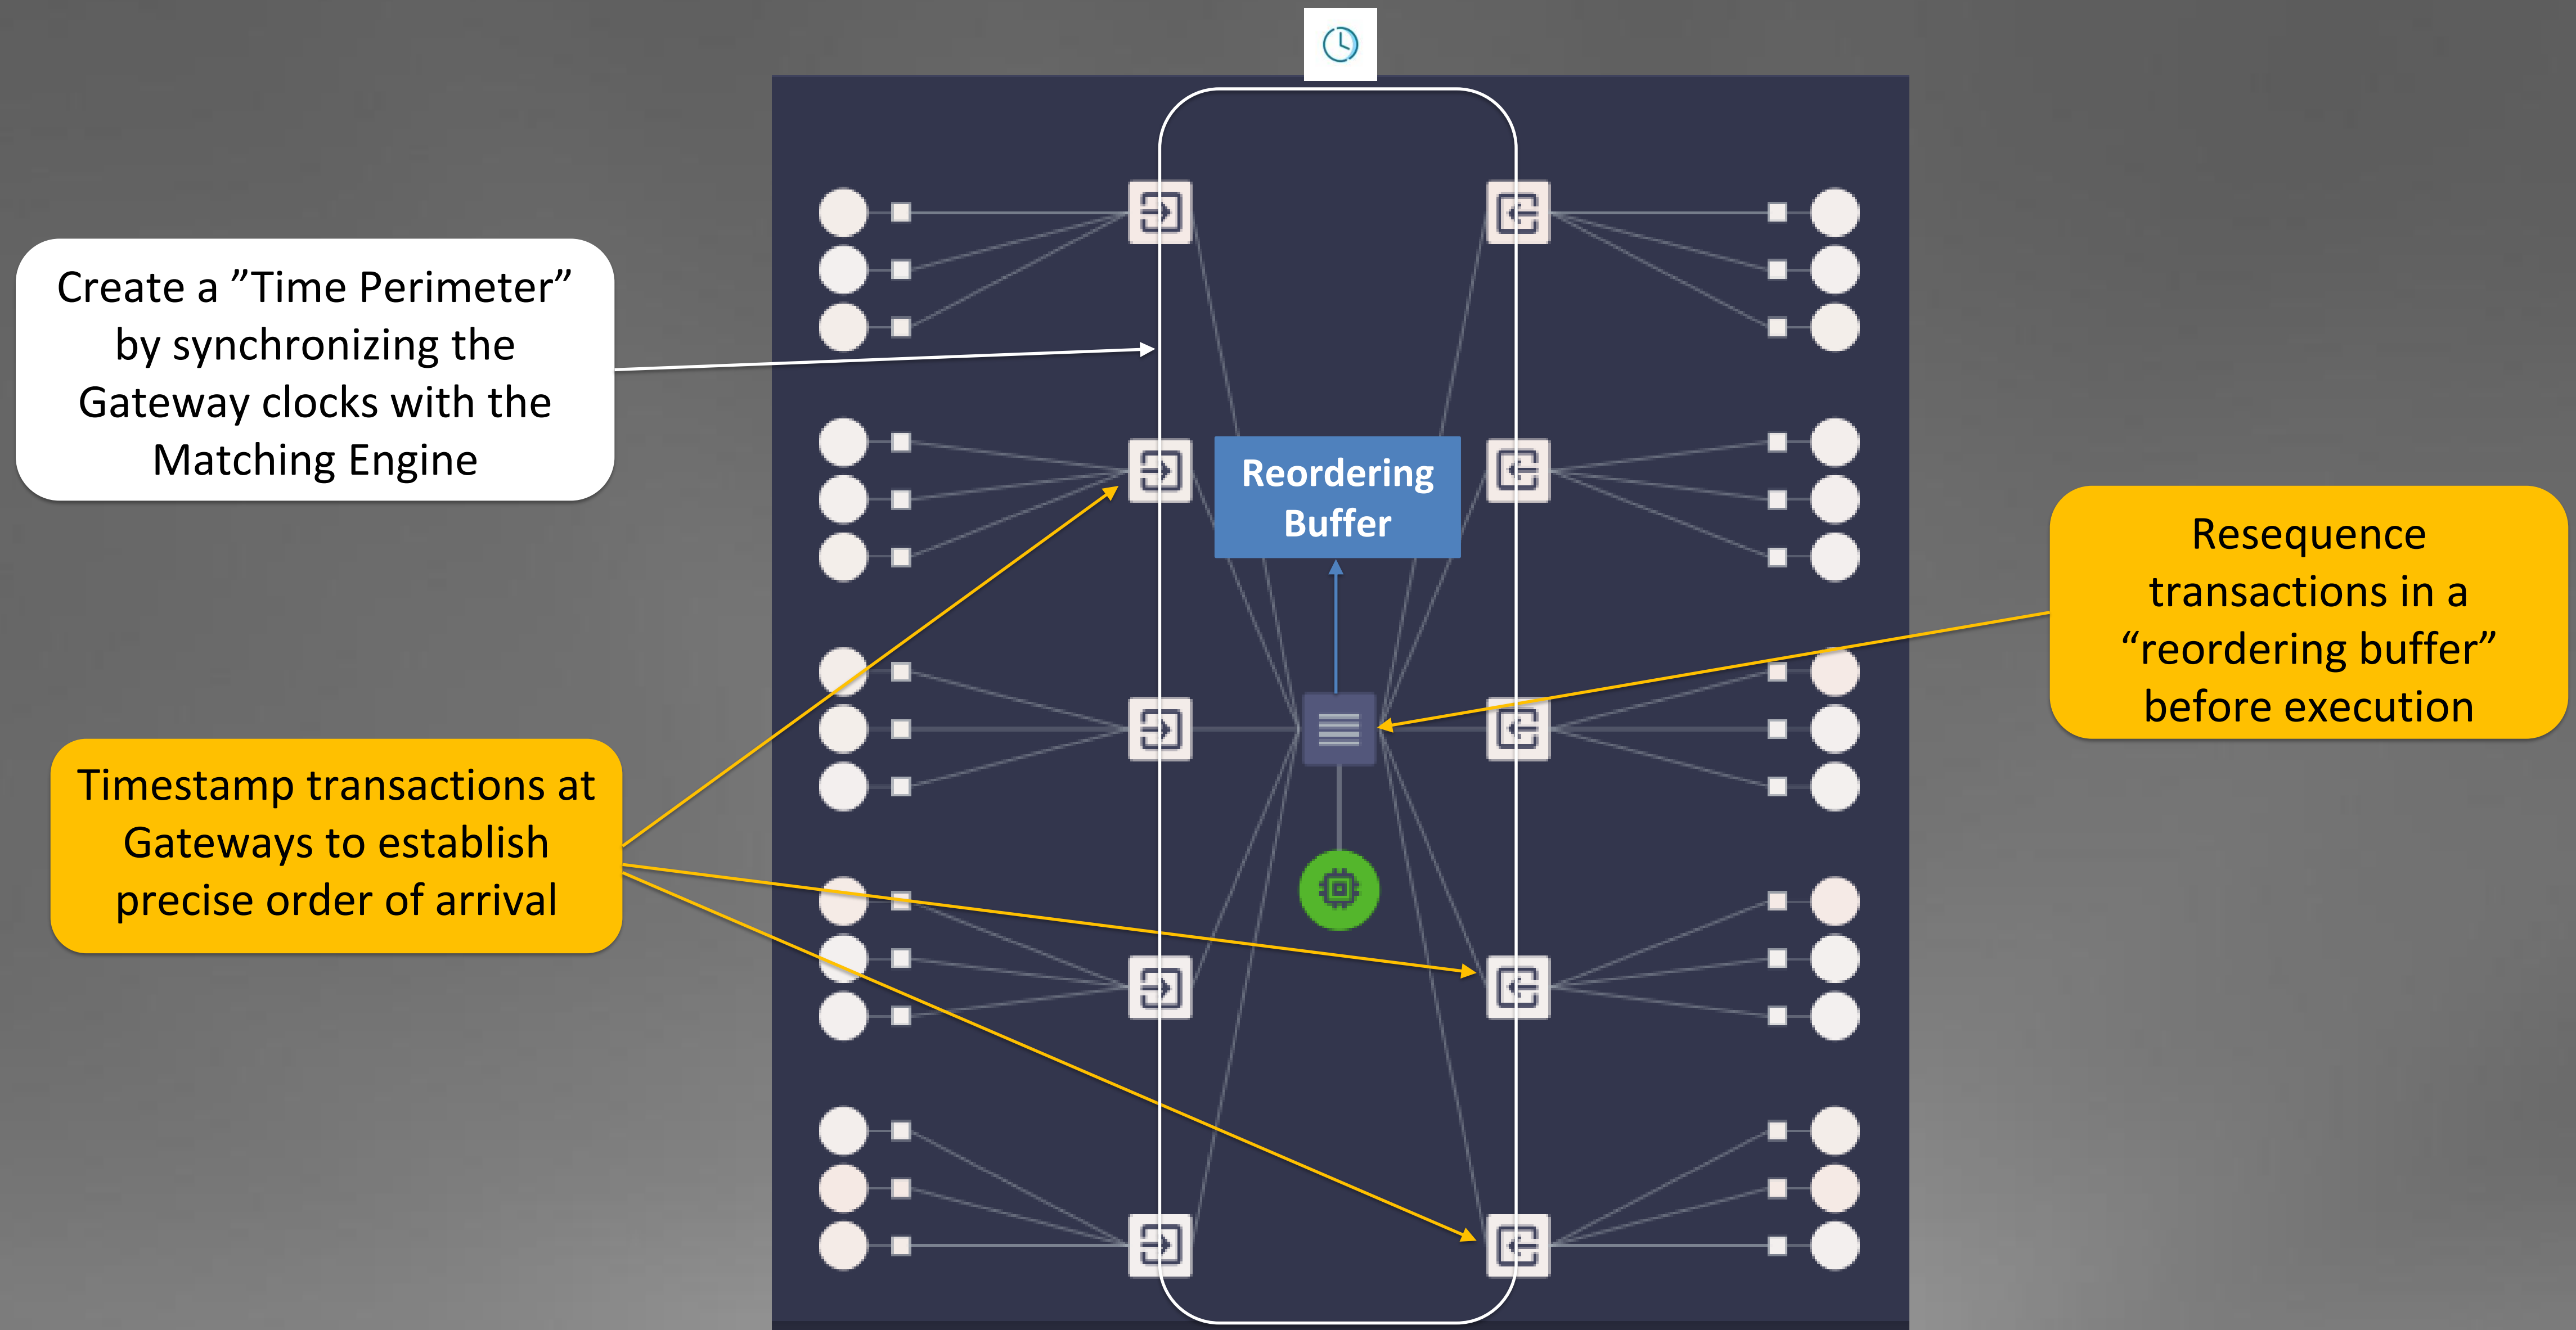
Months after accusing powerful men of abuse, 20 people, including Ashley Judd, above, share how saying #MeToo changed their lives.



Challenge #1: In-order Execution of Transactions



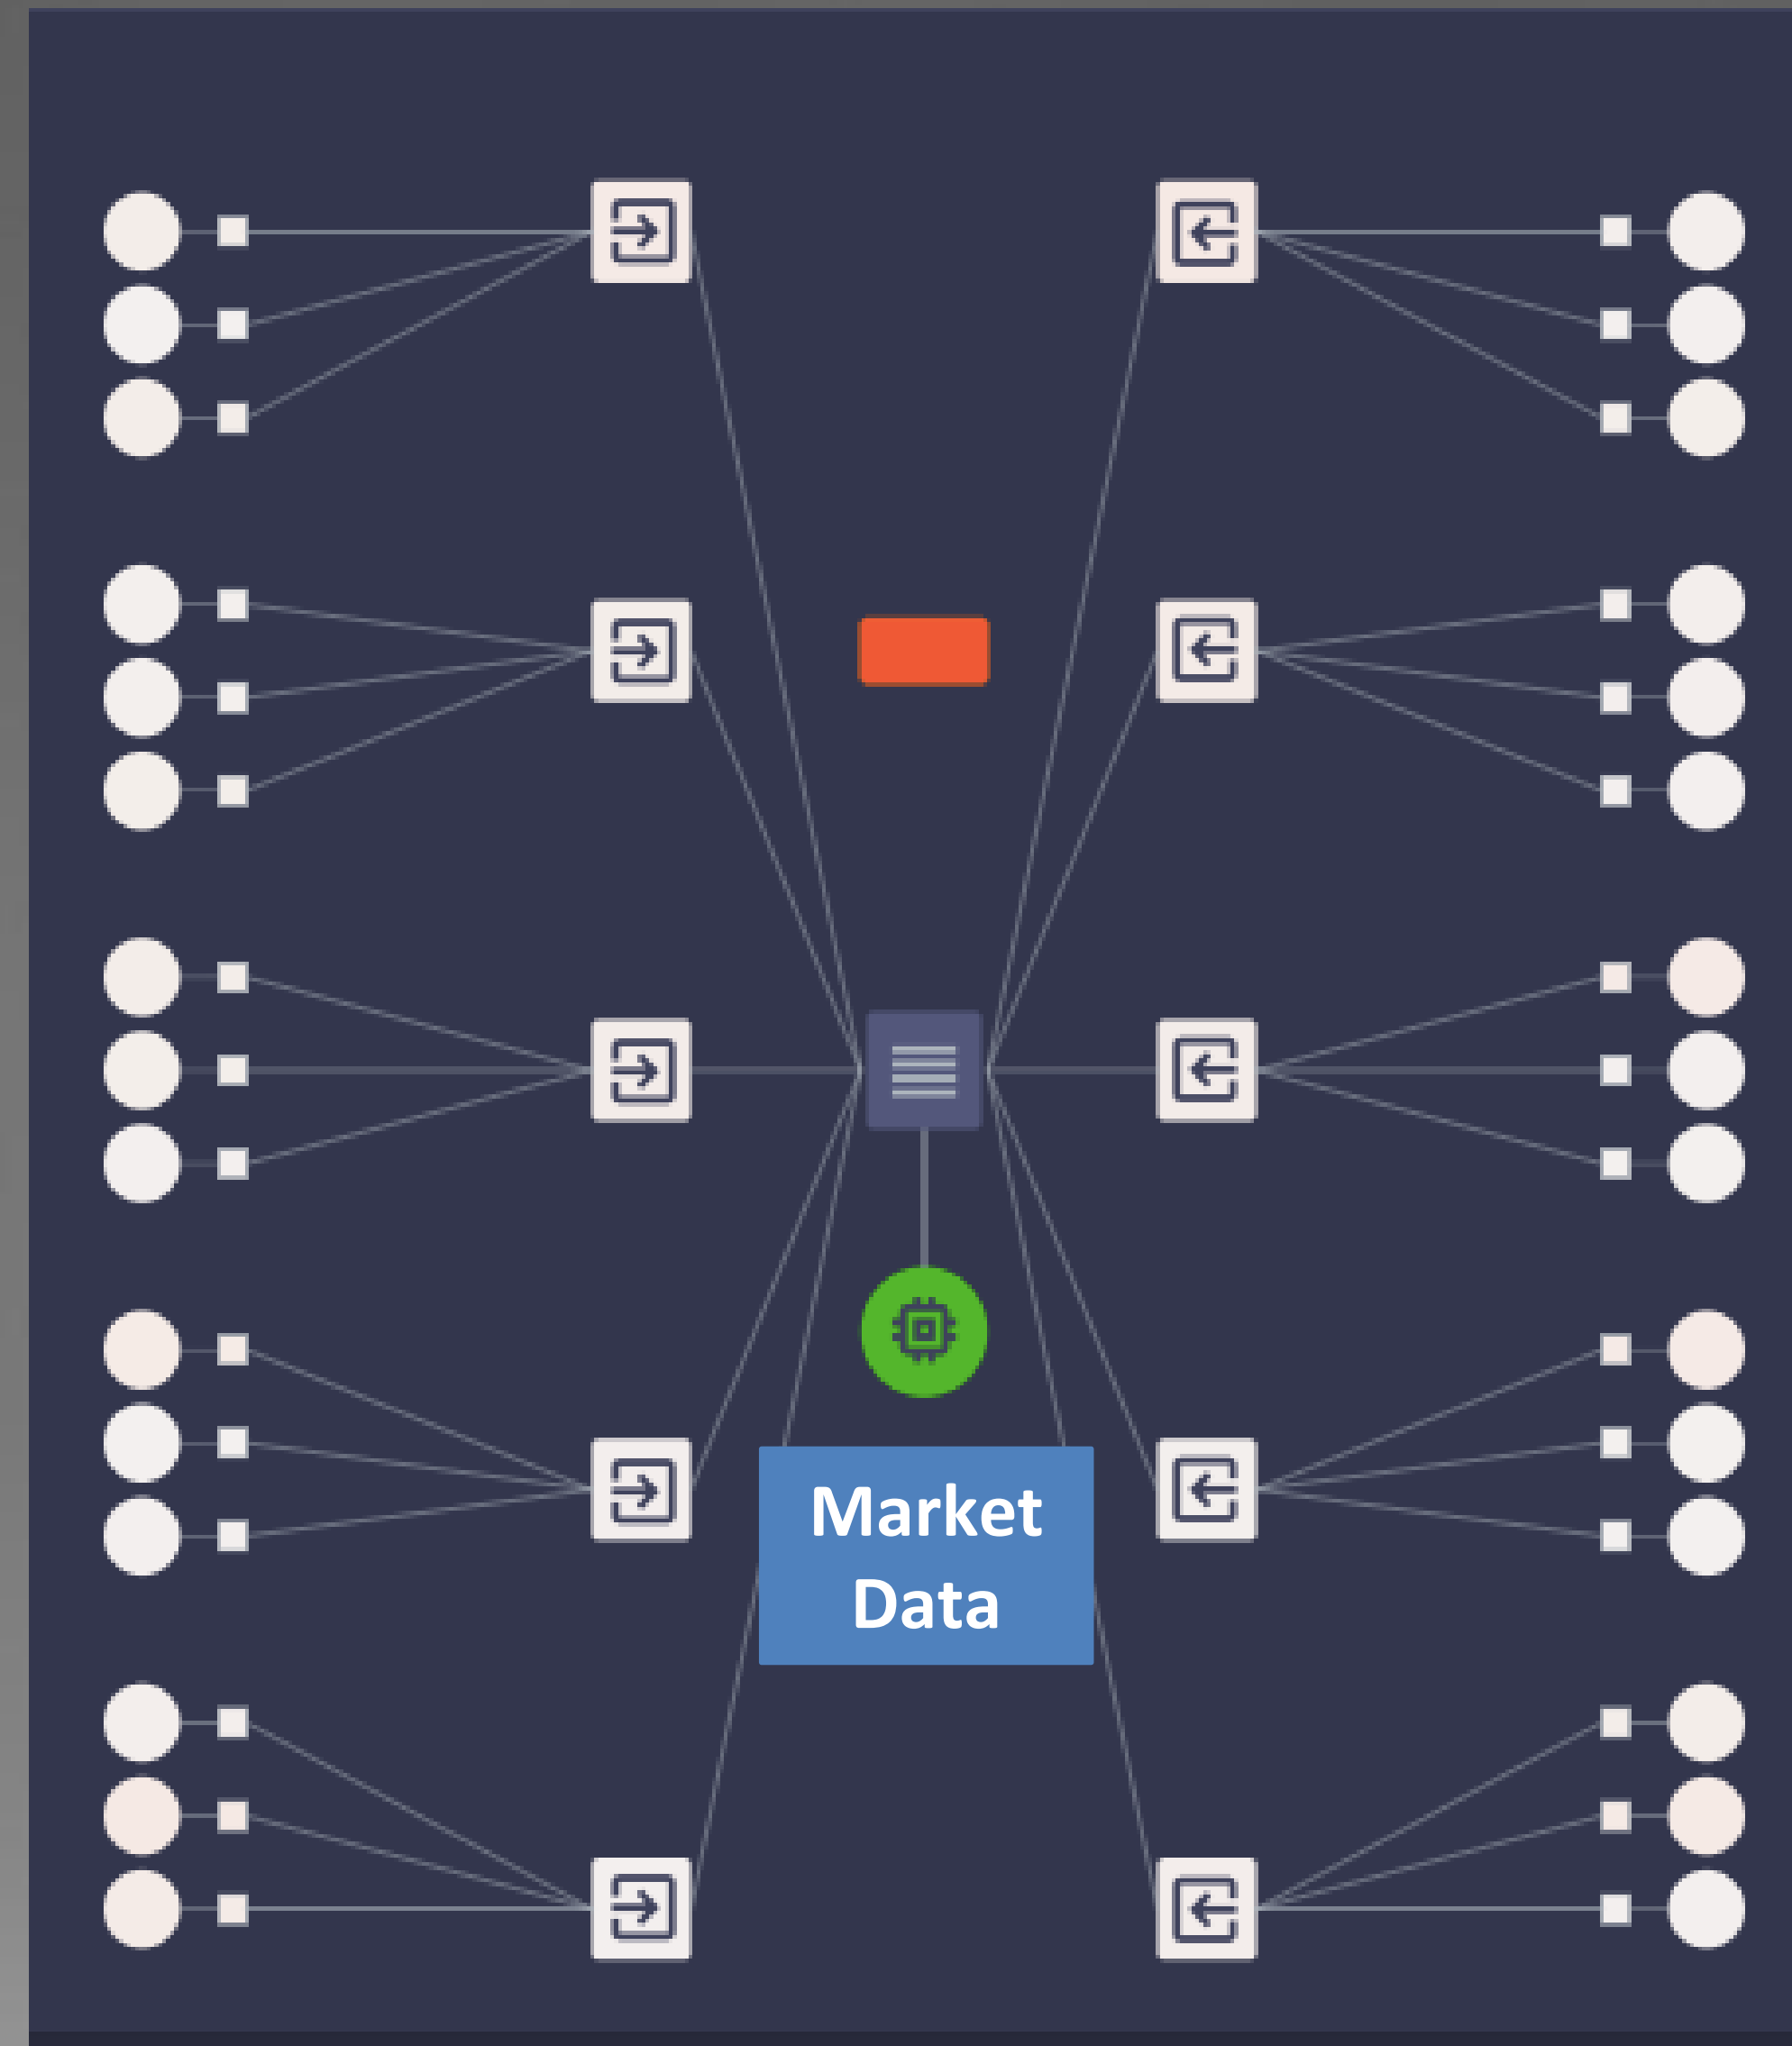
Solution: Create “Time Perimeter” Using Accurate Clock Sync



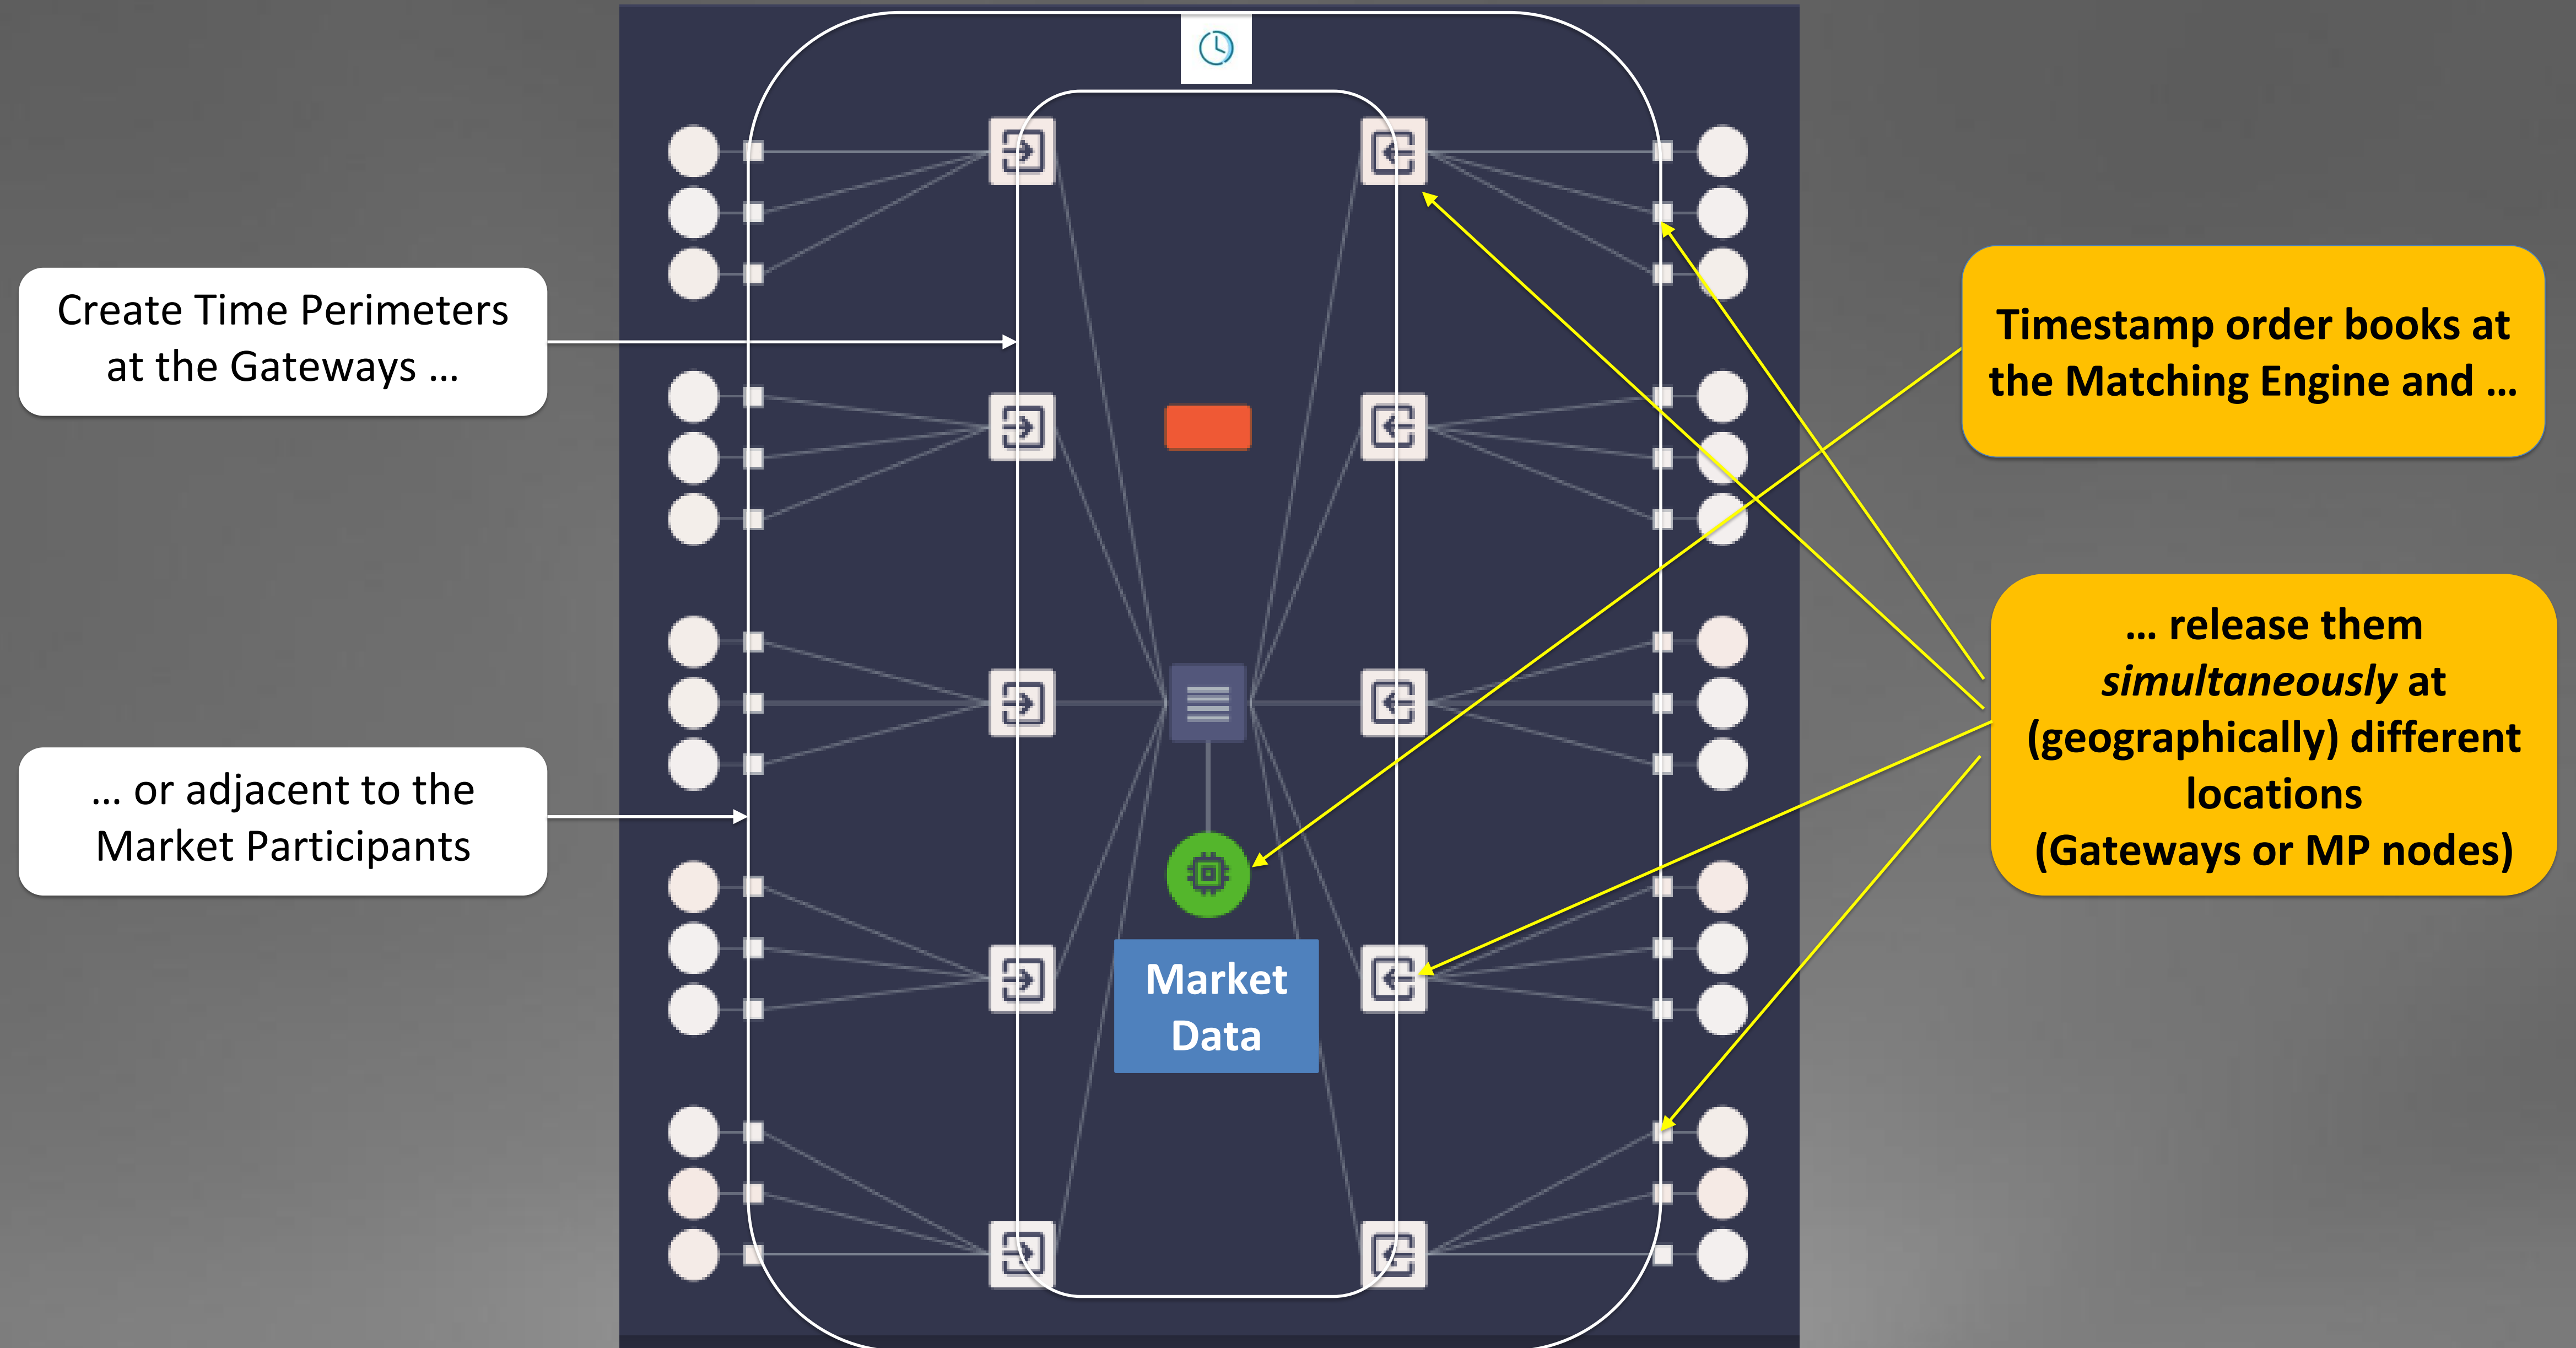
#2: Deliver Market Data Simultaneously

Market data is currently delivered simultaneously to participants by multicasting

Multicasting is not easy in the Cloud



Solution: Time Perimeters + Hold-and-Release Buffers



DEMO of Simulated Stock Exchange in the Cloud

demo-ks2l7x9f

GCP

Exchange

View Alerts

Net: 75 kTPS ▾

CPU: 0% ▾

NTP ▾



Healthy / Total Agents

41 / 41


OOO Tps Before/After Reordering

67k / 67k (100%)

Median Clock Offset

774.1 μ s

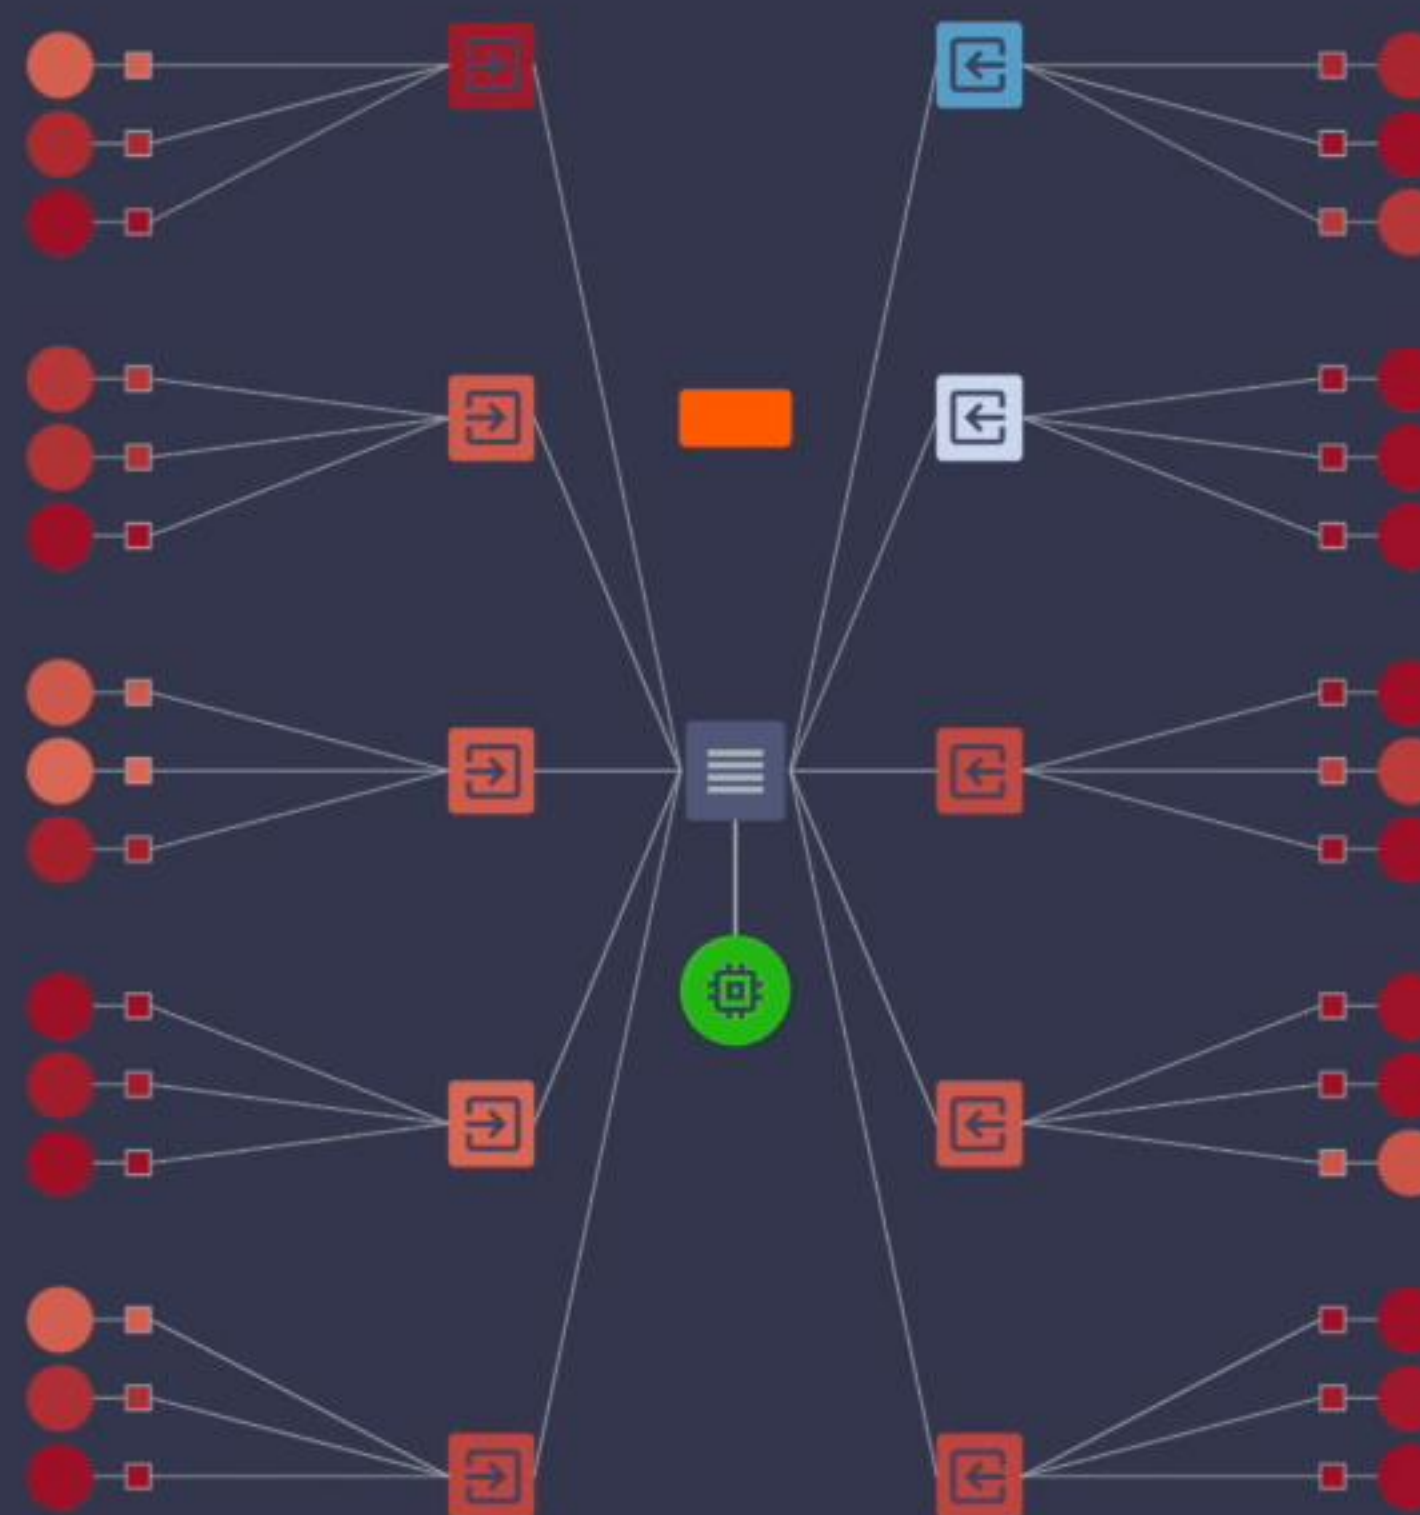

Network



Components

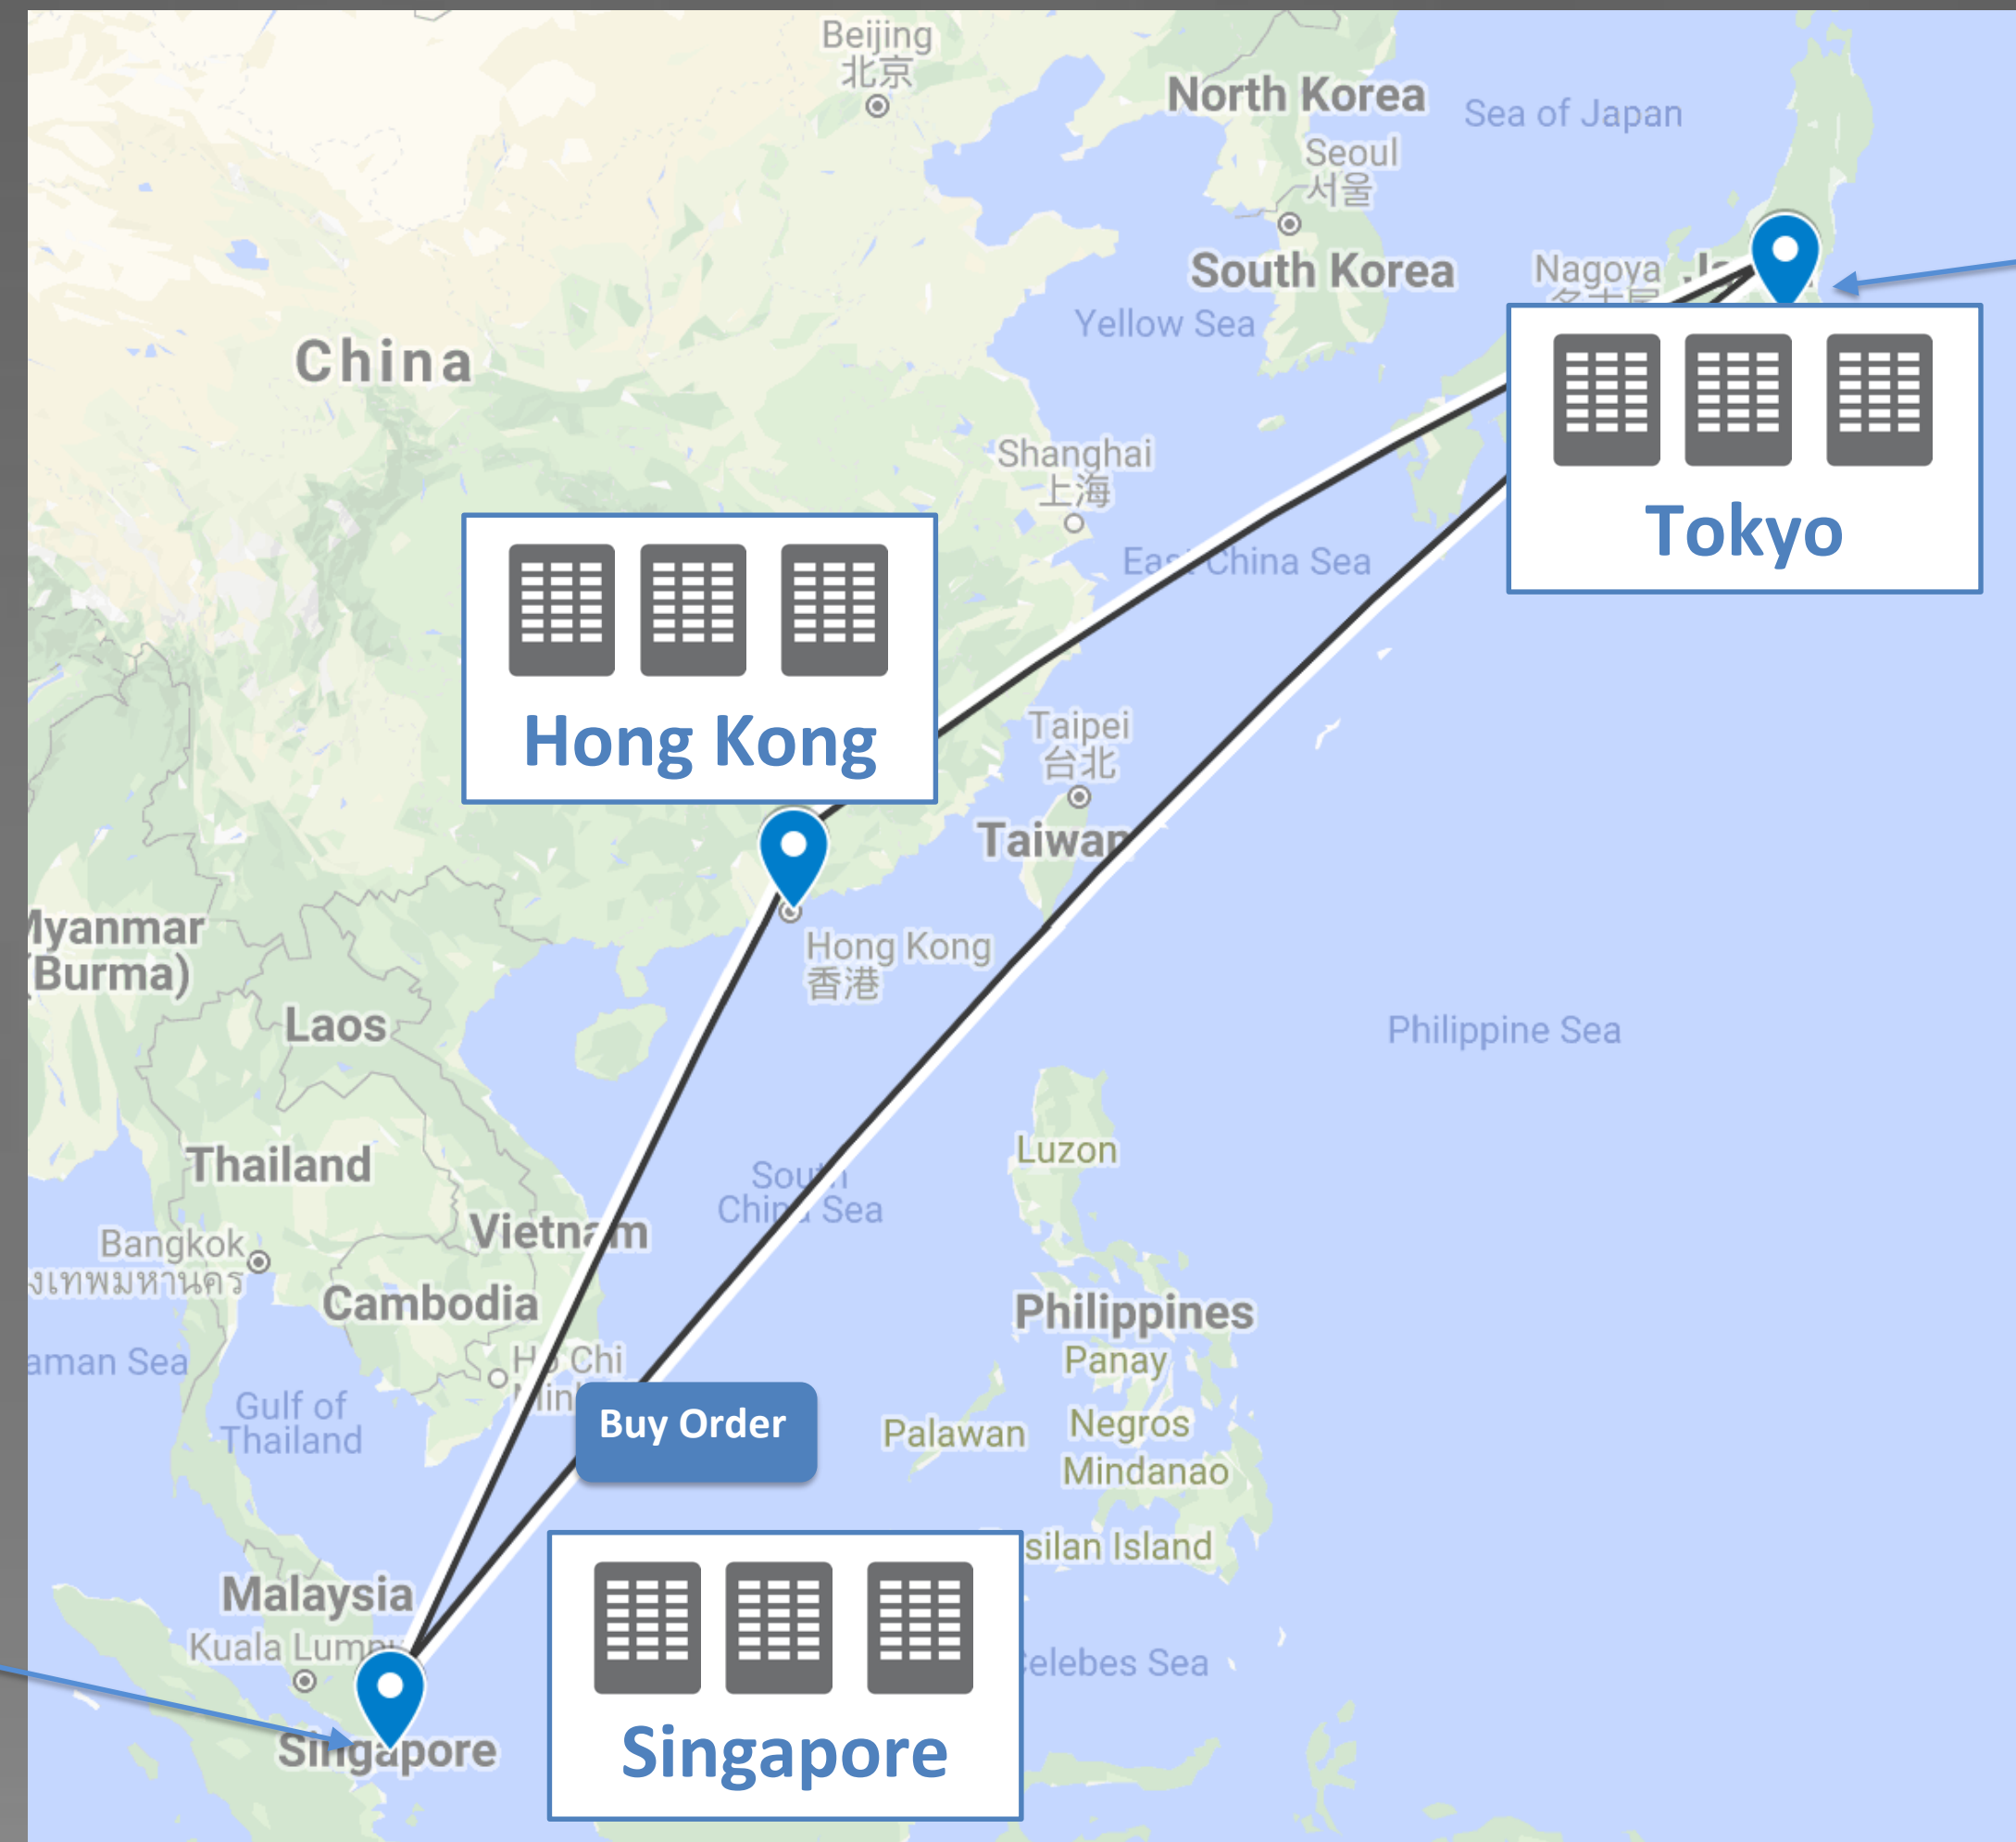
- Coordinator
- Matching Engine
- Gateway
- Reordering Buffer
- Client
- Inactive Client

Offset



#3: Multi-venue Trades

A trader places an order in Singapore and sends it on to HK and Tokyo

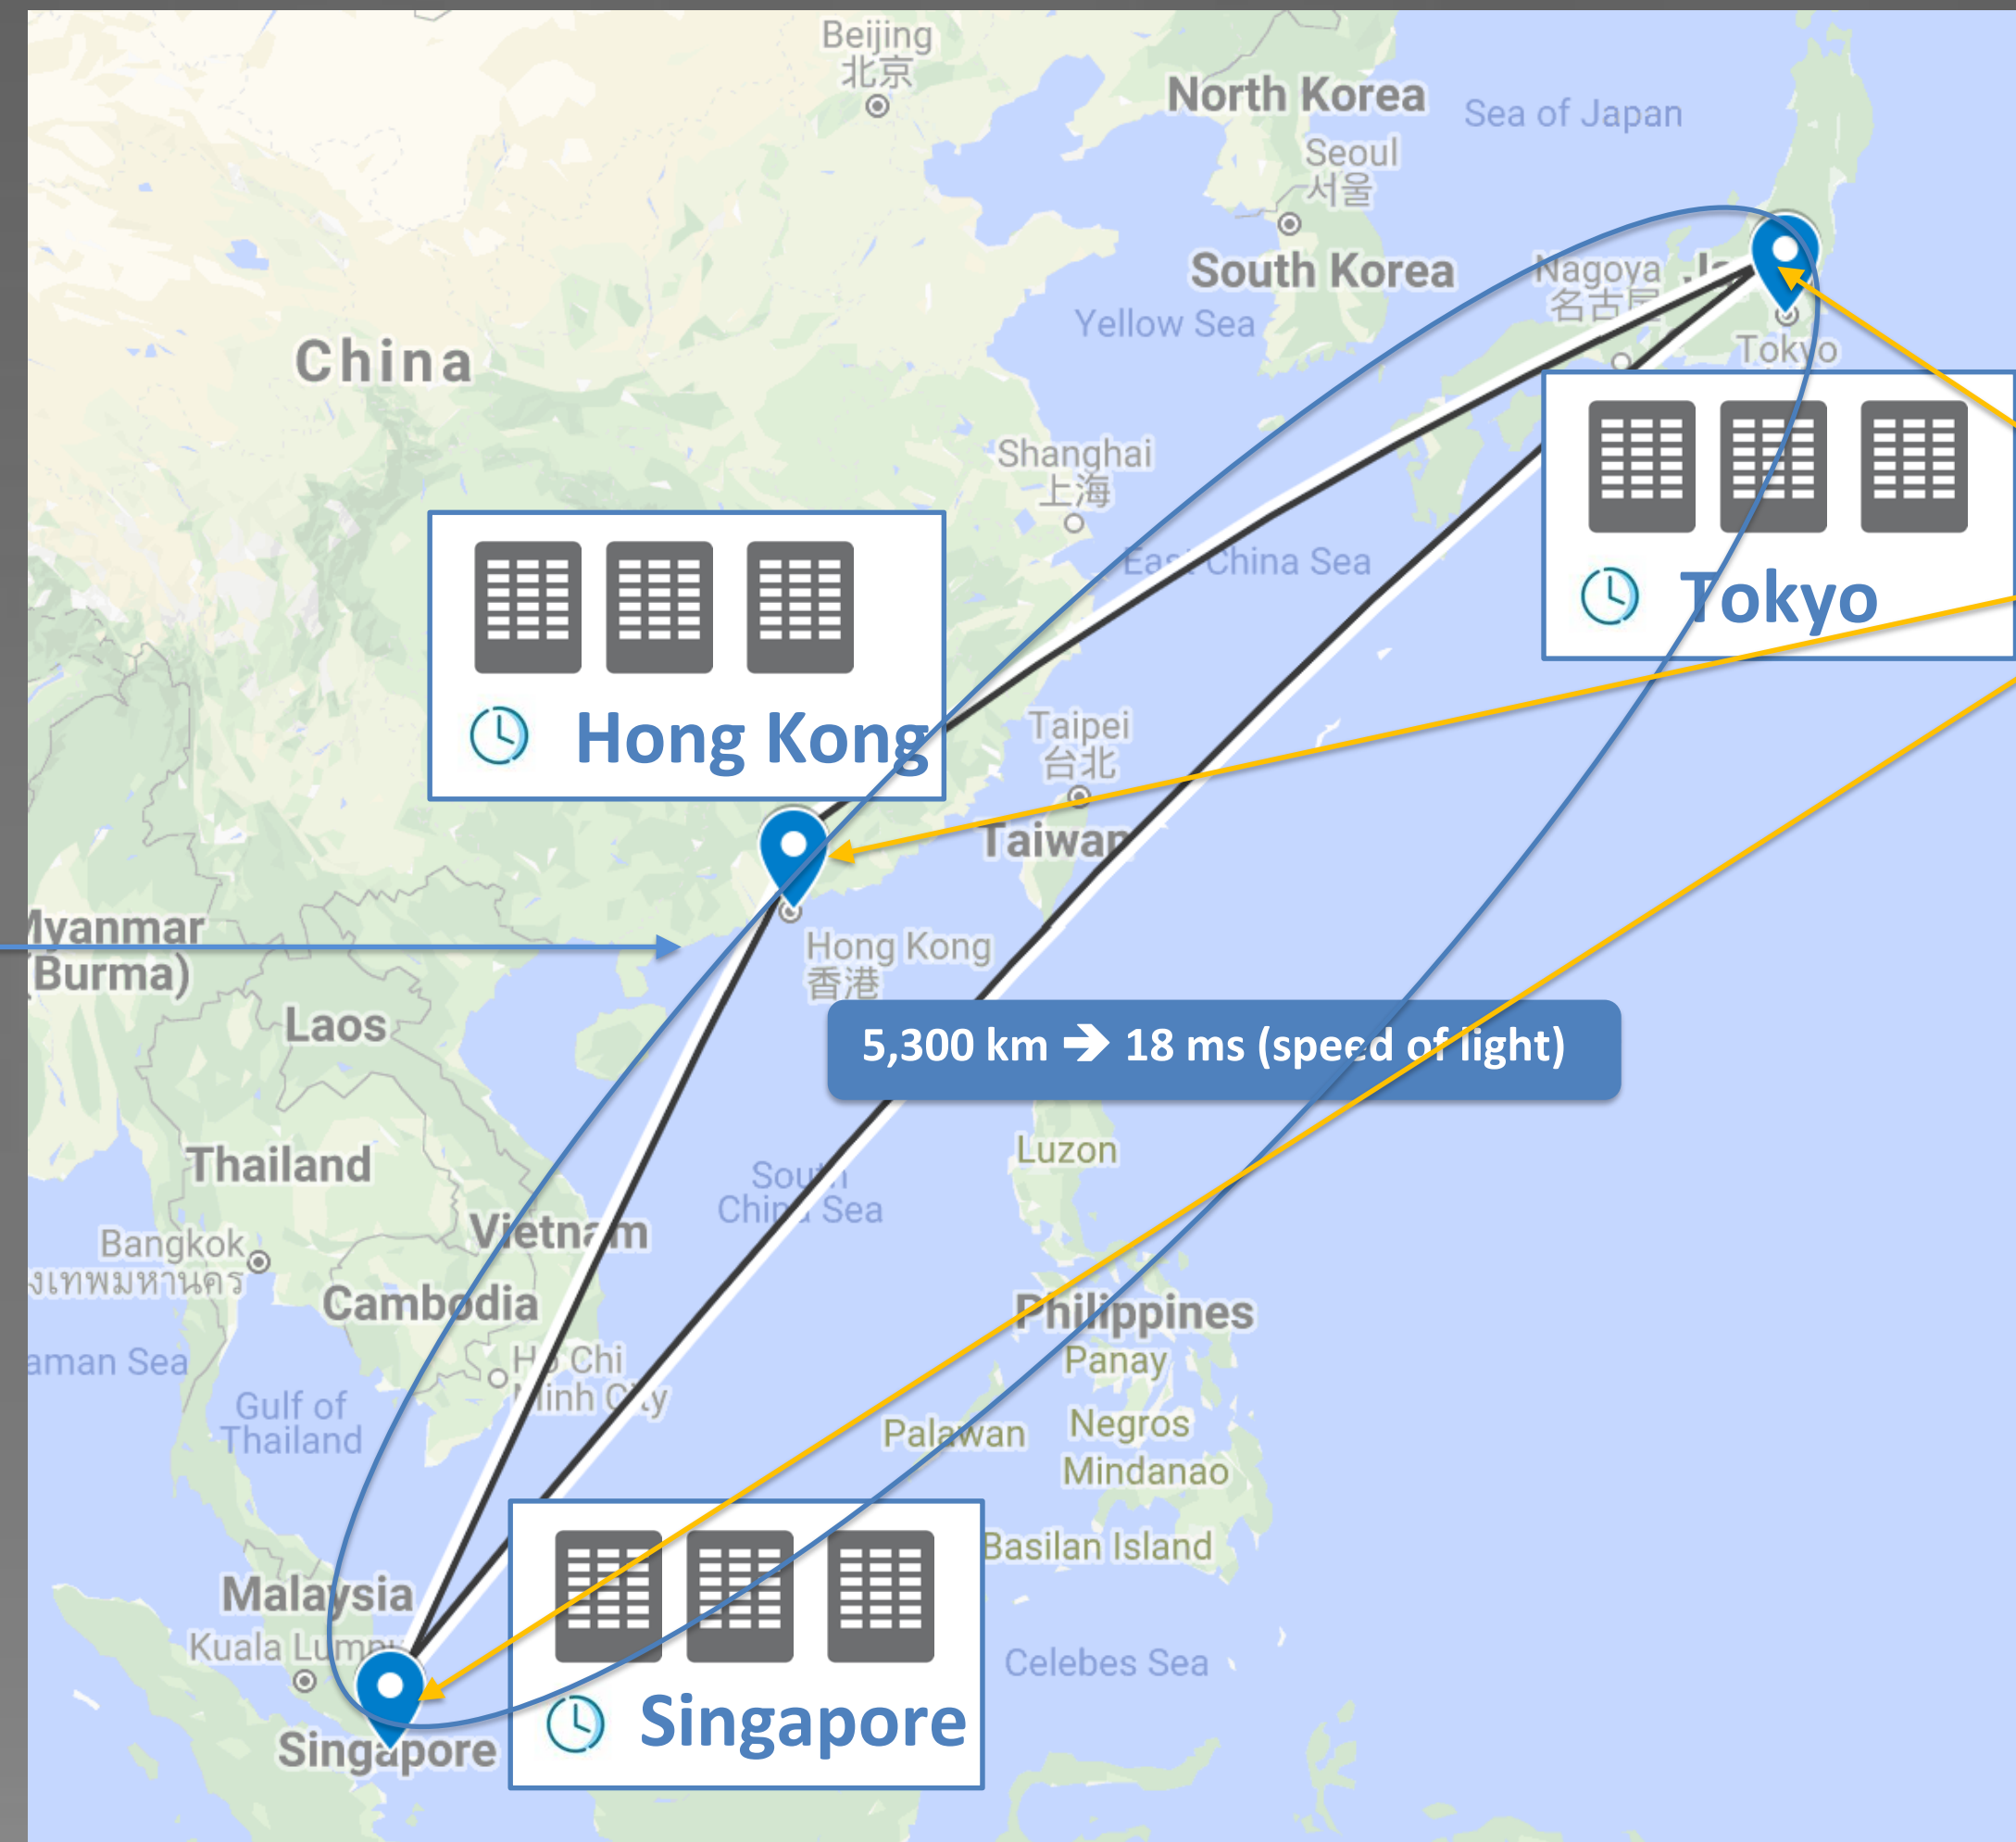


A Market Maker sees the order in Singapore and anticipates it in Tokyo because they've a faster SG → Tokyo link

Results in a sub-optimal deal for the trader!

Solution: Time Perimeter + Hold-and-Release Buffers

Create Time Perimeter
around SG, HK and Tokyo



1. Hold SG trade in buffer
2. Send it to HK and Tokyo
3. Release it "simultaneously" in SG, HK and Tokyo at future time

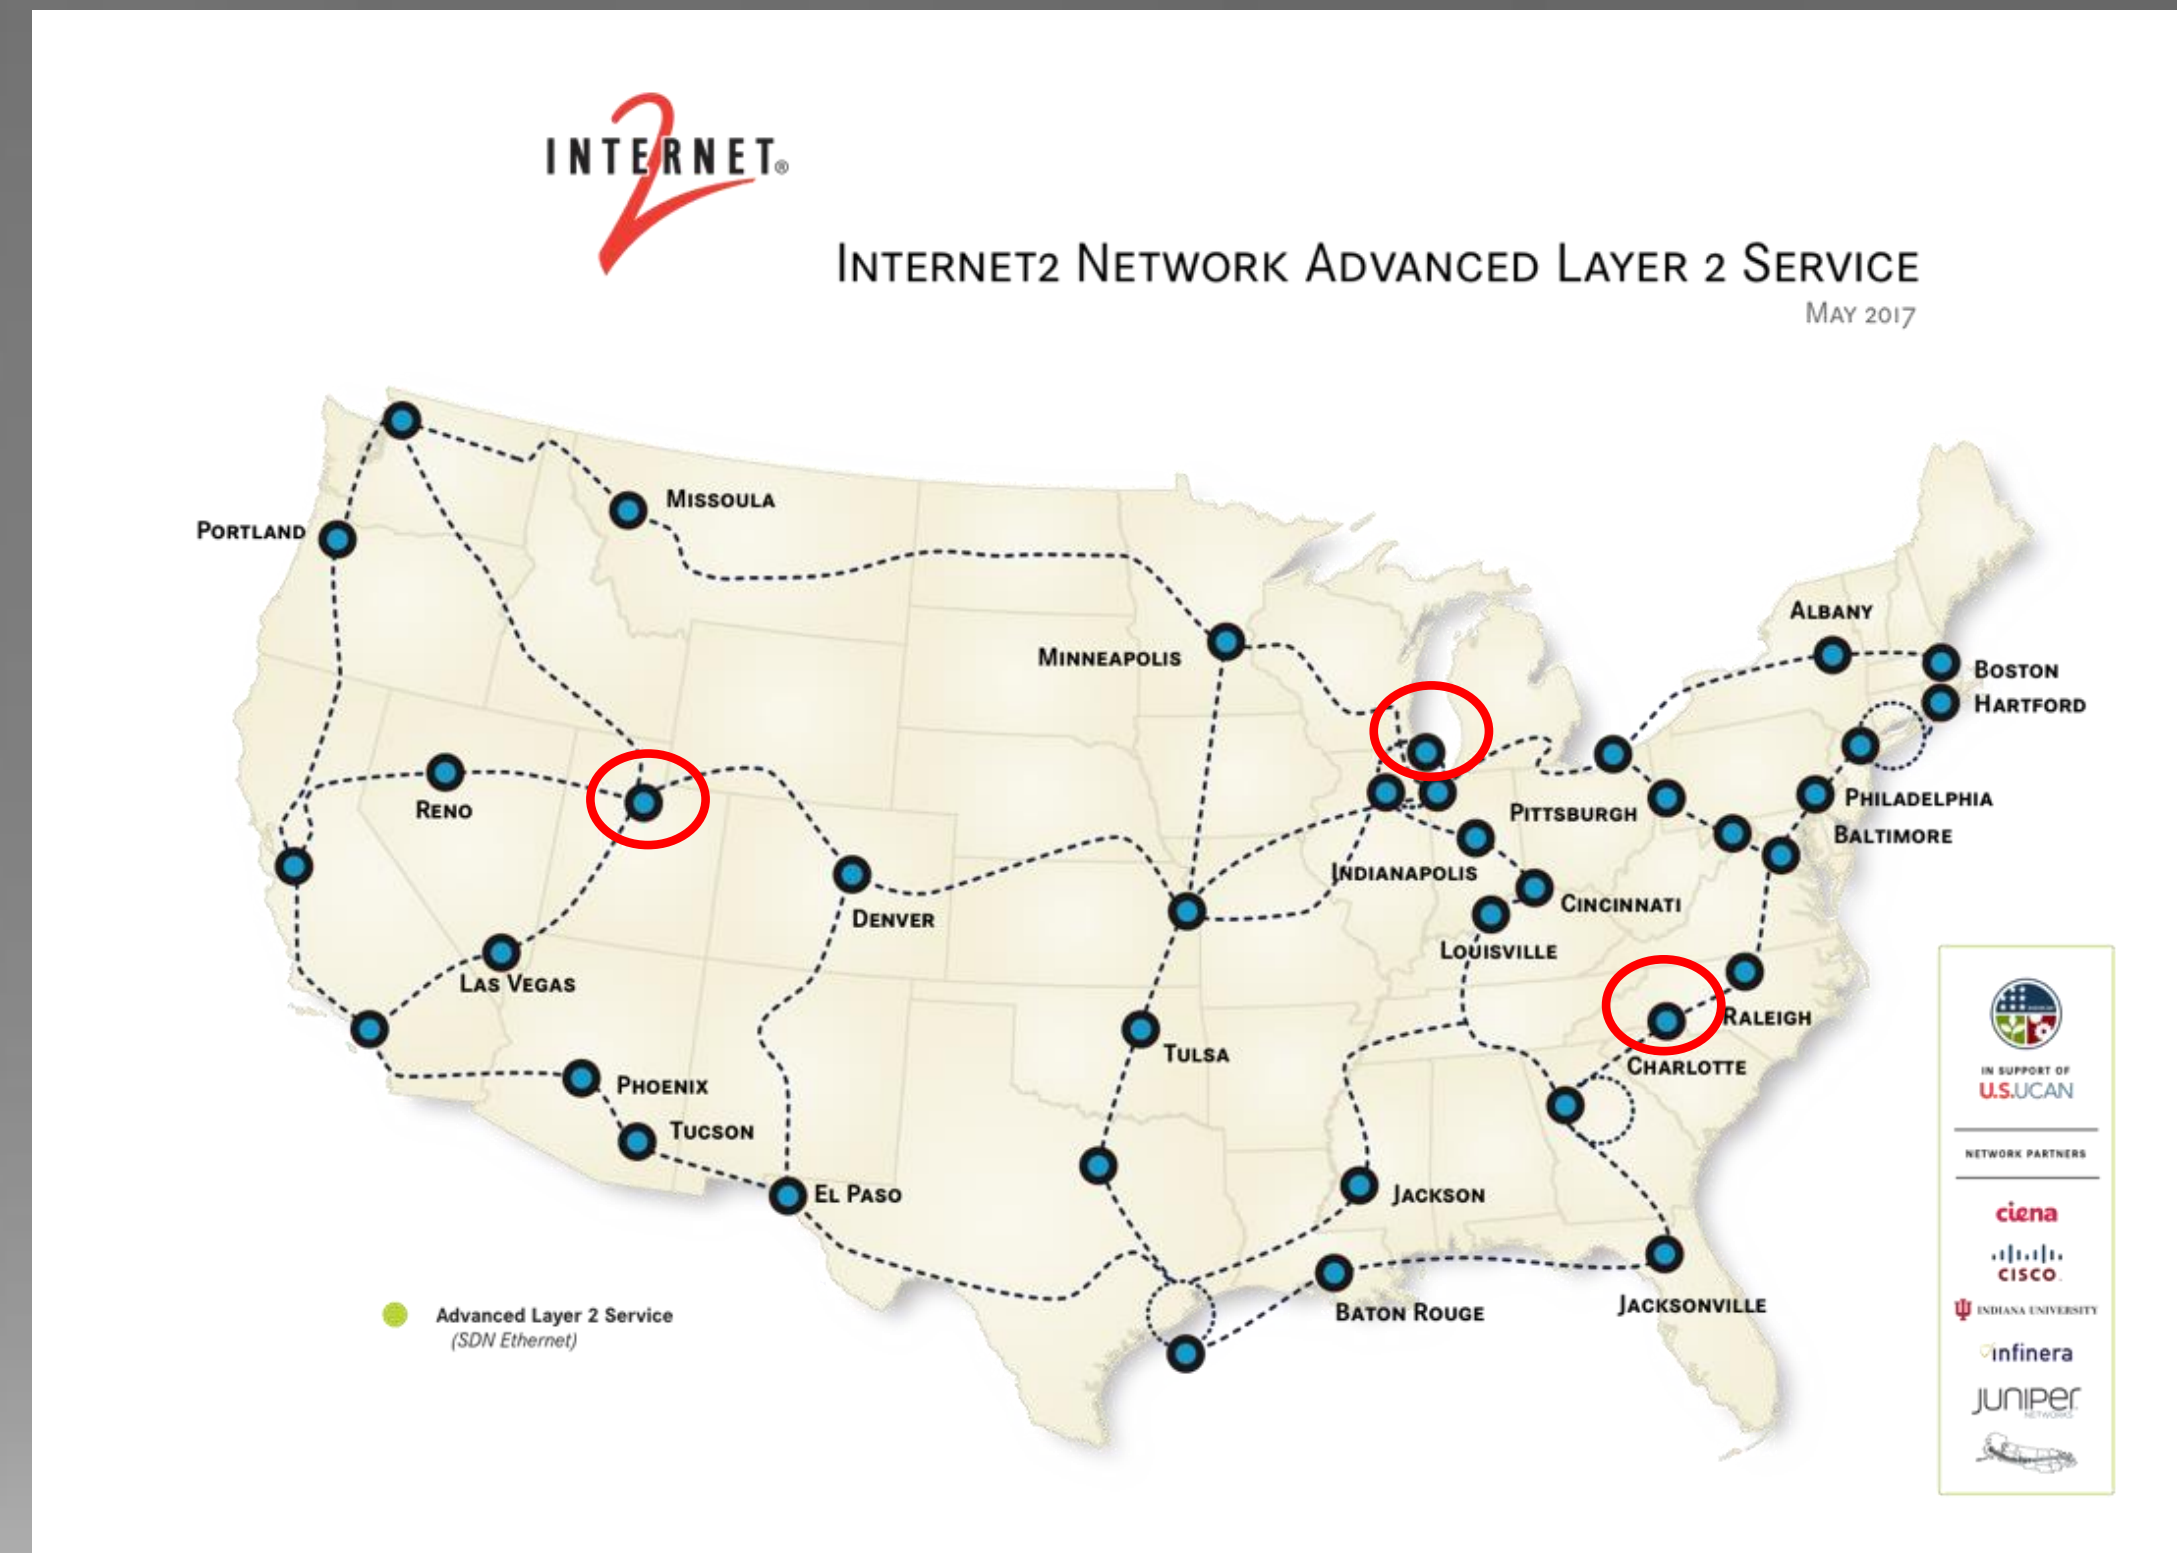
Note: Clock sync accuracy only needs to be better than speed of light distance between venues.

Speed of light SG → Tokyo: 18ms

Clock Sync at Distance: The Stanford CloudLab Experiment

CloudLab

- 3 sites – Utah, Wisconsin and Clemson
 - Connected through Internet2's 100Gbps network
- RTT between the sites
 - Utah <-> Wisconsin: 36ms
 - Utah <-> Clemson: 52ms
 - Wisconsin <-> Clemson: 26ms
- Don't know the routes



CloudLab: A Heterogeneous Environment

- 2-layer fat-tree network in each site
- Utah: 10Gbps
 - Mellanox ConnectX-3 Pro NICs
- Wisconsin: 1Gbps
 - Intel I350 NICs
- Clemson: 1Gbps
 - Intel I350 NICs
- All NICs support hardware timestamping

Clock Sync Accuracy

Single-site: Utah

p50	Huygens		NTP	
	0% load	50% load	0% load	50% load
HW timestamp	2ns	9ns	7.9us	6.7us
SW timestamp	95ns	80ns	61us	67us

p99.9	Huygens		NTP	
	0% load	50% load	0% load	50% load
HW timestamp	26ns	54ns	21ms	14ms
SW timestamp	486ns	384ns	9ms	54ms

Multi-site: Utah, Wisconsin, Clemson

	Huygens		NTP	
	p50	p99.9	p50	p99.9
HW timestamp	2.8us	10.0us	52ms	57ms
SW timestamp	2.7us	10.2us	52ms	73ms

DEMO of Tick Tock: Multi-Site

demo-dlqiyyya

AWS [OR, OH, VA]

Datacenter

View Alerts

Net: 0 Mbps

CPU: 0%

TTCS



Healthy / Total Agents

120 / 120

90%ile Clock Offset

425 ns

Median Clock Offset

180 ns

Network



Components

- Coordinator
- Reference
- Agent
- Inactive Agent

Offset

- 1 ms
- 50 μ s
- 10 μ s
- 0 ns
- 10 μ s
- 50 μ s
- 1 ms



An Accurate Clock System is an Excellent Measurement System

Effort required to maintain accurate clock sync over network

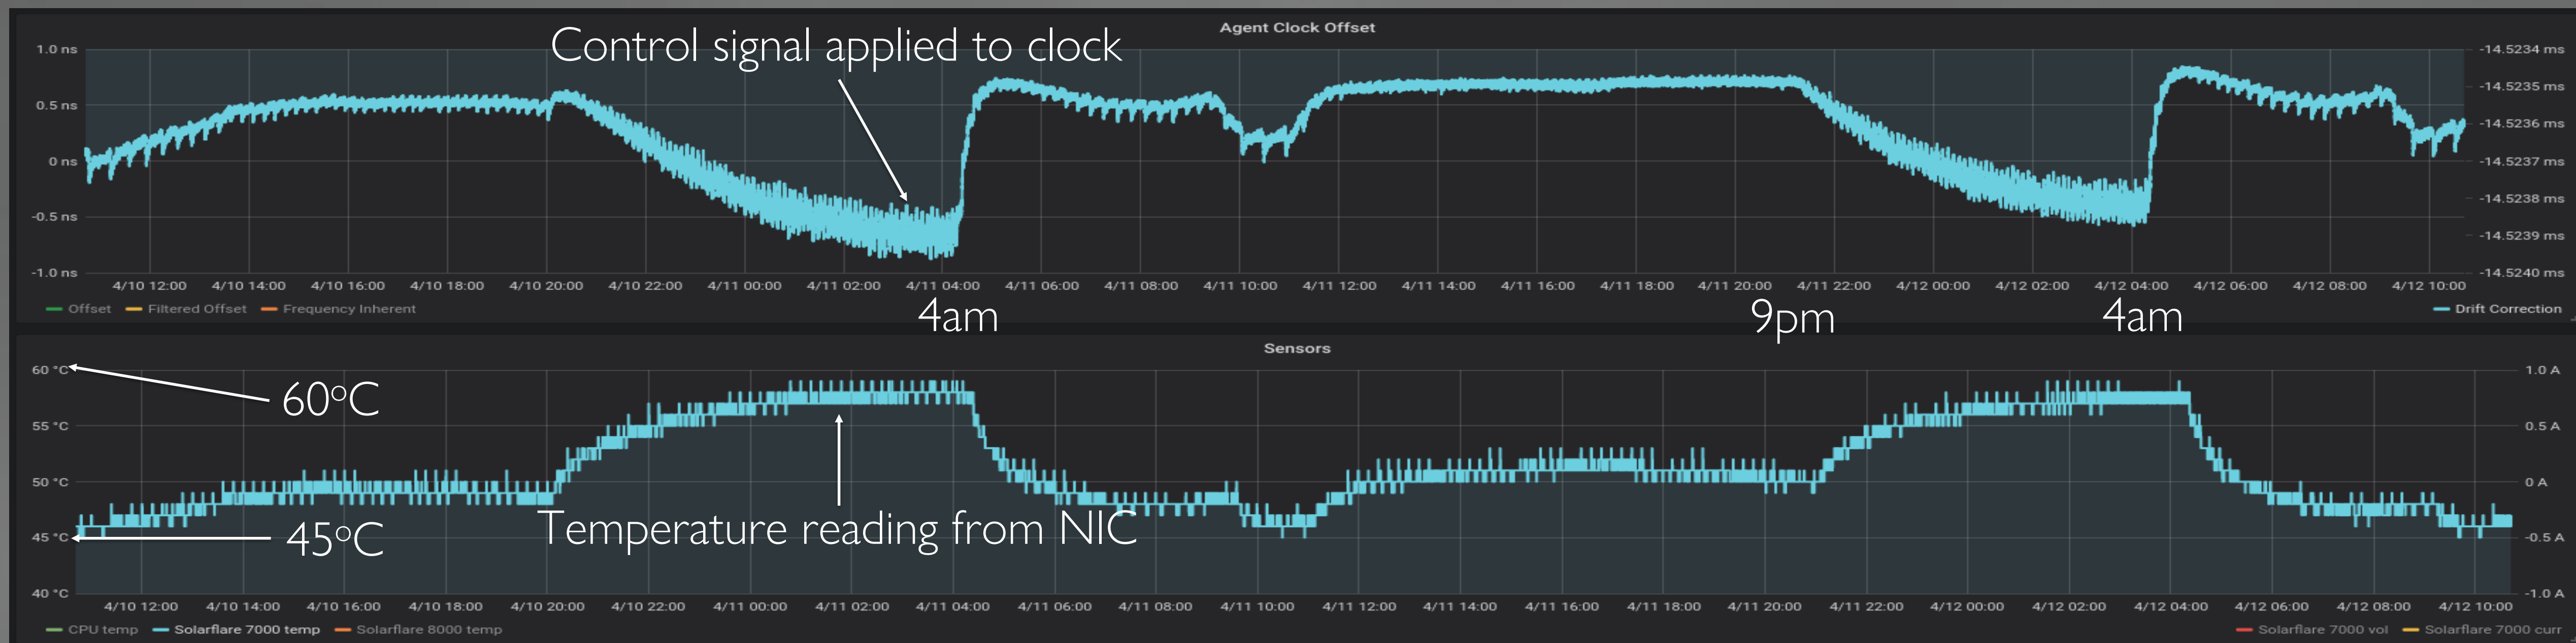
α

Amount network and system “badness”

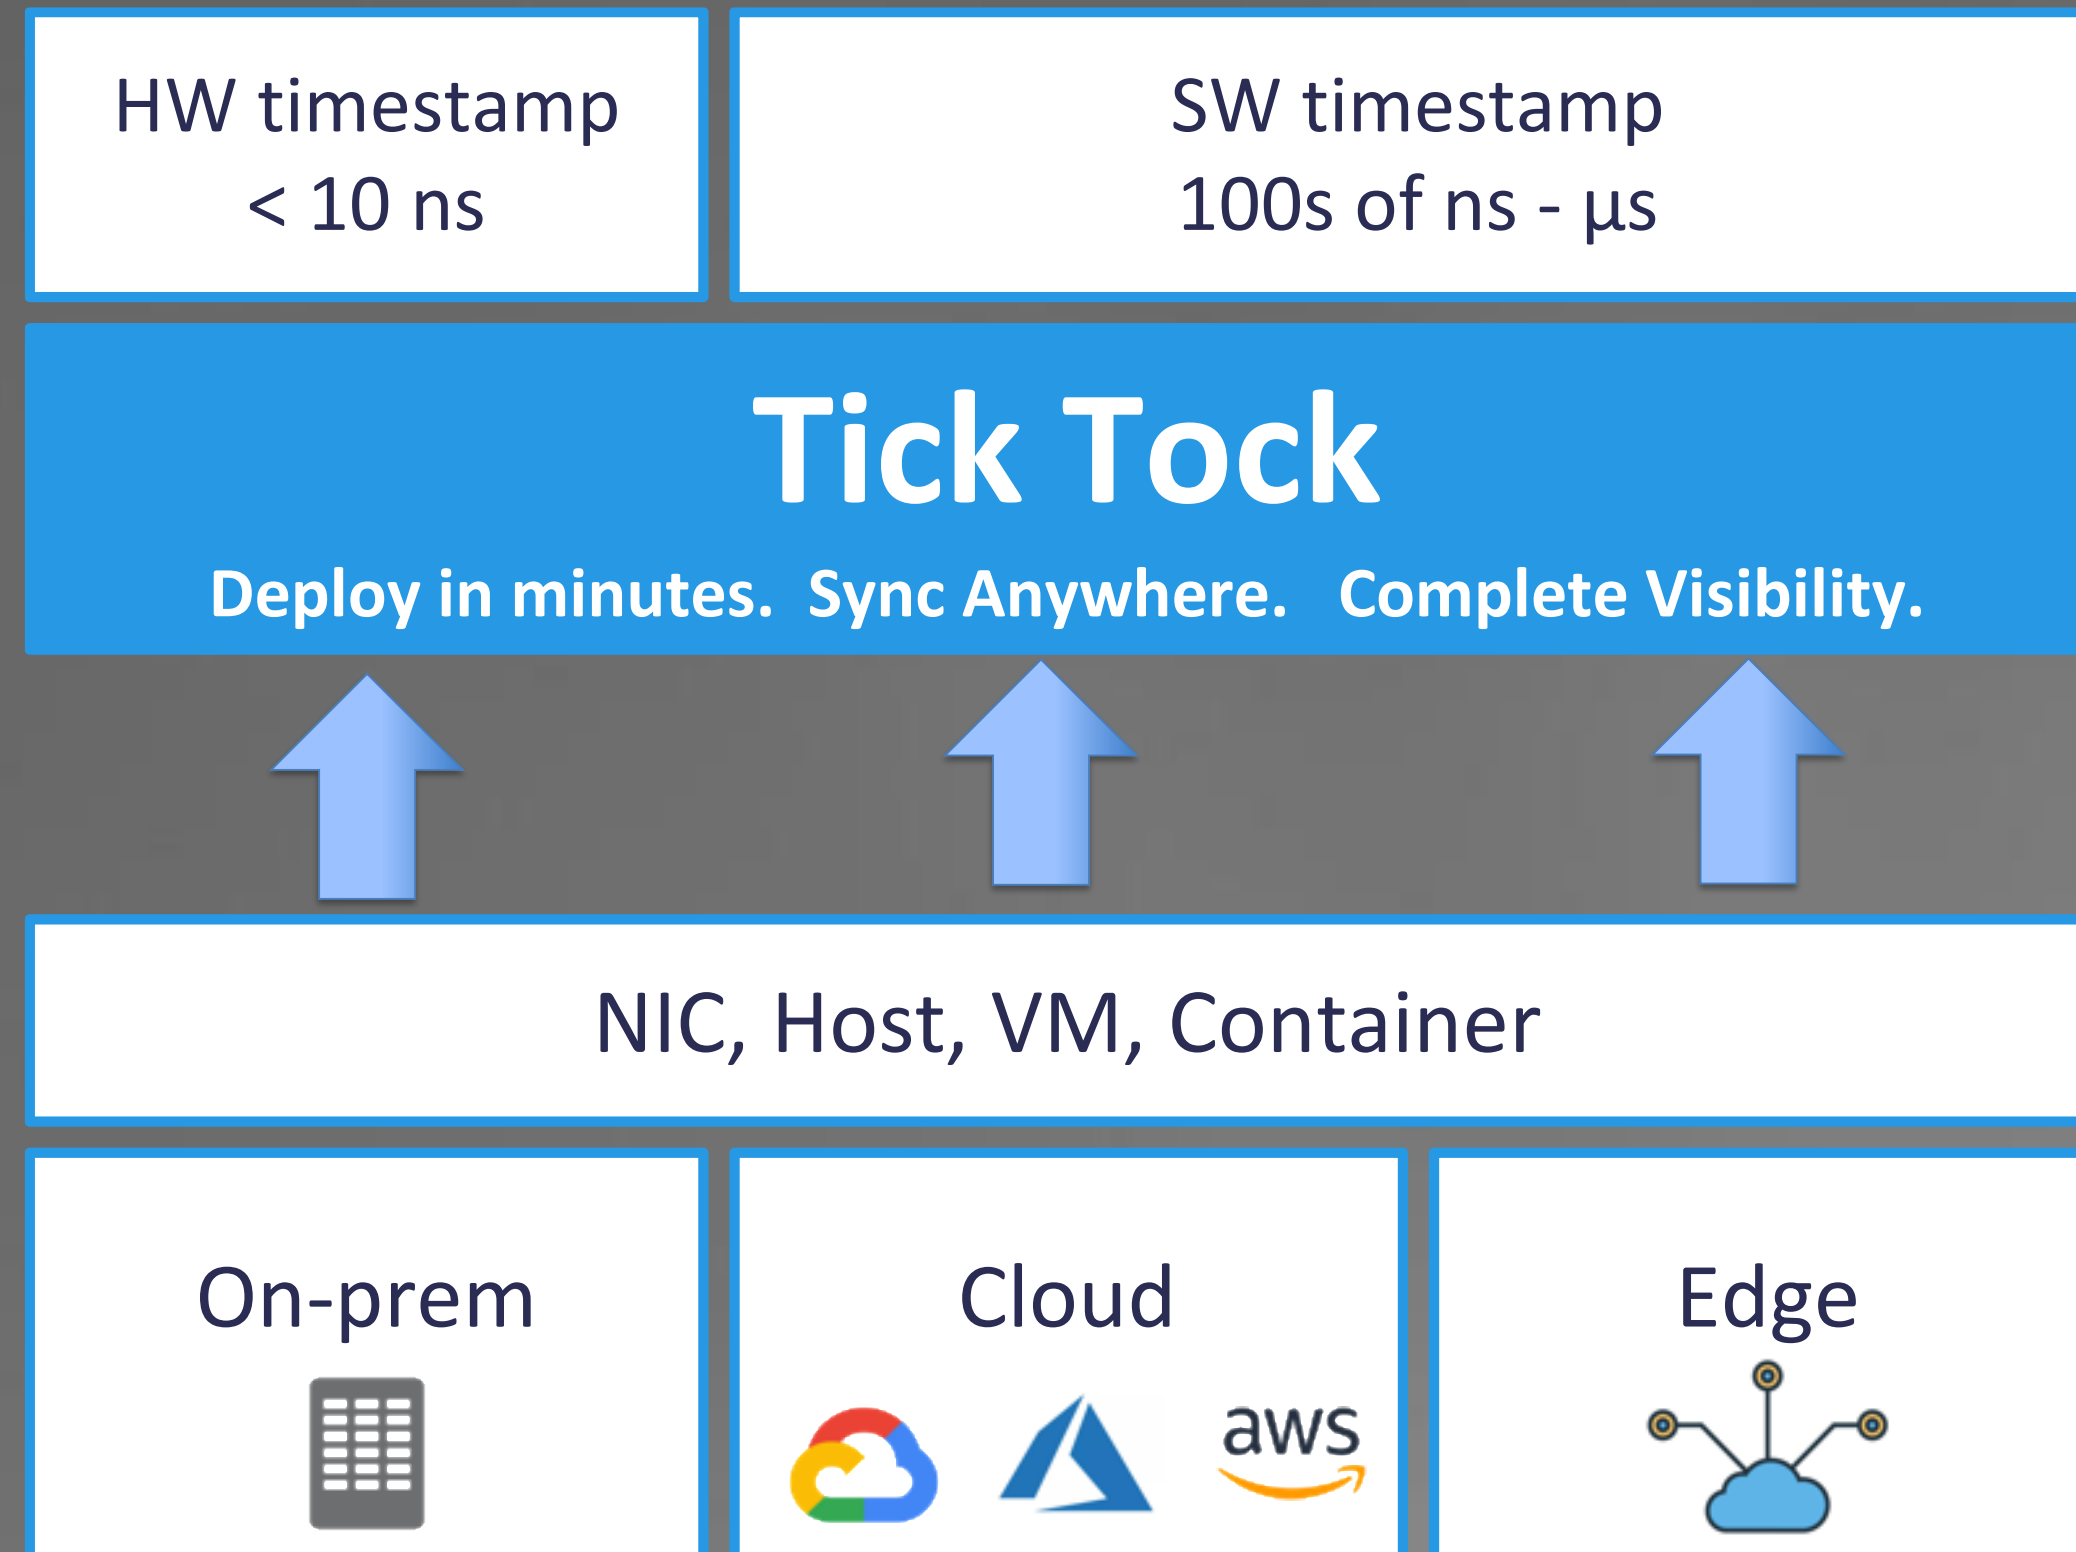
Network badness: Path congestion, path asymmetry, link/node failure, ...

System badness: NICs/CPU getting hot, room temperature change, vibration (e.g., fans), ...

→ Therefore, an accurate clock sync system = great telemetry system



Tick Tock: A Modern, Software-based Solution for Accurate Time Sync & Infrastructure Monitoring



High accuracy and high performance clock sync

- Nanosecond-level with NIC hardware timestamps
- Support single-site sync and multi-site sync (regional and global)
- Scale up to 10s of 1,000s of servers

Complete visibility and insights across clocks, servers and network

- Monitor and visualize clock sync performance
- Correlate clock and network performance to pin point and fix issues, including one-way delay measurements and analytics
- Inter-site connectivity: identify path asymmetry and path delay changes.

Easy to deploy and manage

- Install, config, and run in 30 minutes
- Admin console, APIs, and reporting

Not STAC Benchmarks