When the cloud fails: Debugging the "undocumented"

Dhruv Jain - GoTo Group (Gojek), Indonesia

MQTT Infrastructure Operational Scale

2 Million
Concurrent Connections

15 Million
Daily Active Users

3.2 Million
rpm
Throughput

2.2 %
Indonesia's GDP

Background: MQTT Basics

1. Lightweight L7 protocol over tcp
2. 3 QoS levels
3. Bi-directional & full duplex long running connection
4. Pub-sub model with topic based routing

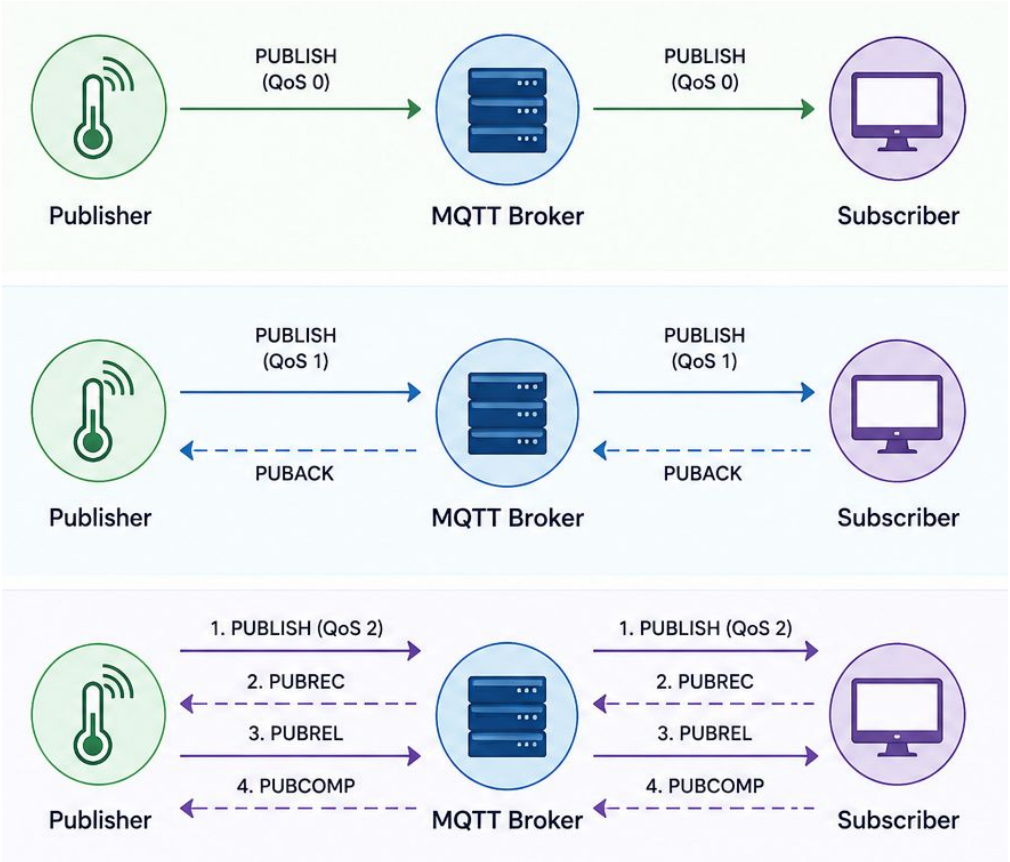

Background: MQTT Setup

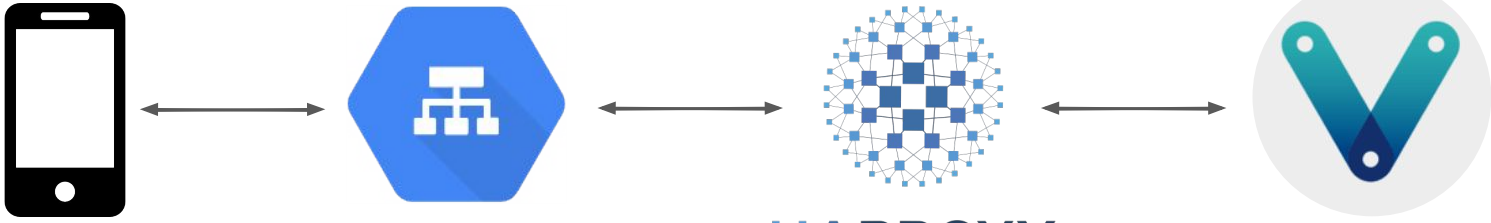

Customers
150K concurrent connections

Driver Partners
600K concurrent connections

Merchant Partners
450K concurrent connections

The Trigger: MQTT Infrastructure Migration

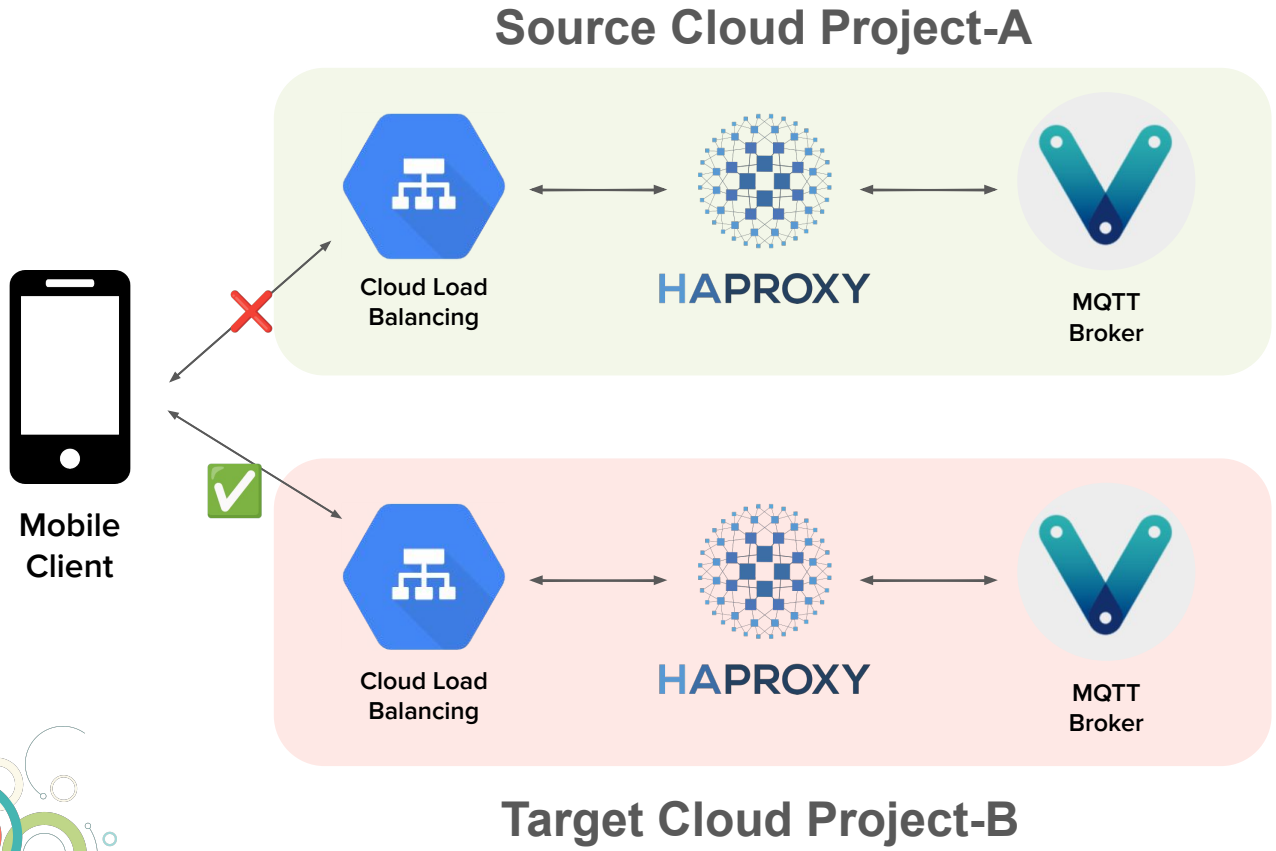

- ✅ Same app versions
- ✅ Equivalent infra specs
- ✅ Network rules & setup
- ✅ Load testing

The Trigger: MQTT Infrastructure Migration

SYSTEM STATUS OVERVIEW — “Everything Looked Green”

MQTT Broker Health	
Connections:	✅ Stable
Publish Rate:	✅ Normal
Receive Rate:	✅ Normal
Broker CPU / Memory:	✅ Healthy
Same-Network Latency:	✅ Low latency
Bi-directional Channel:	✅ Working

The Problem: Drop in Delivery Rate for Merchants

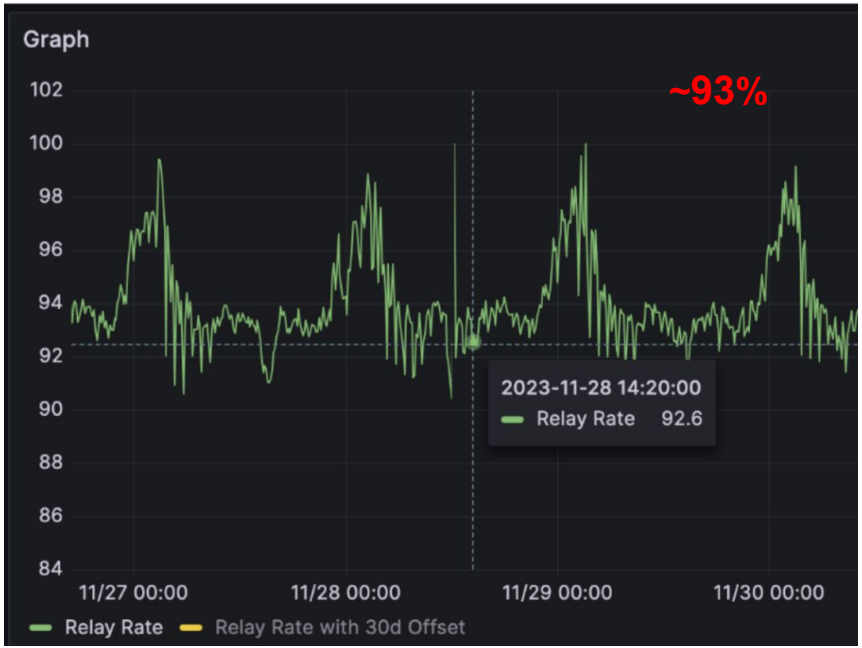

Pre-cutover

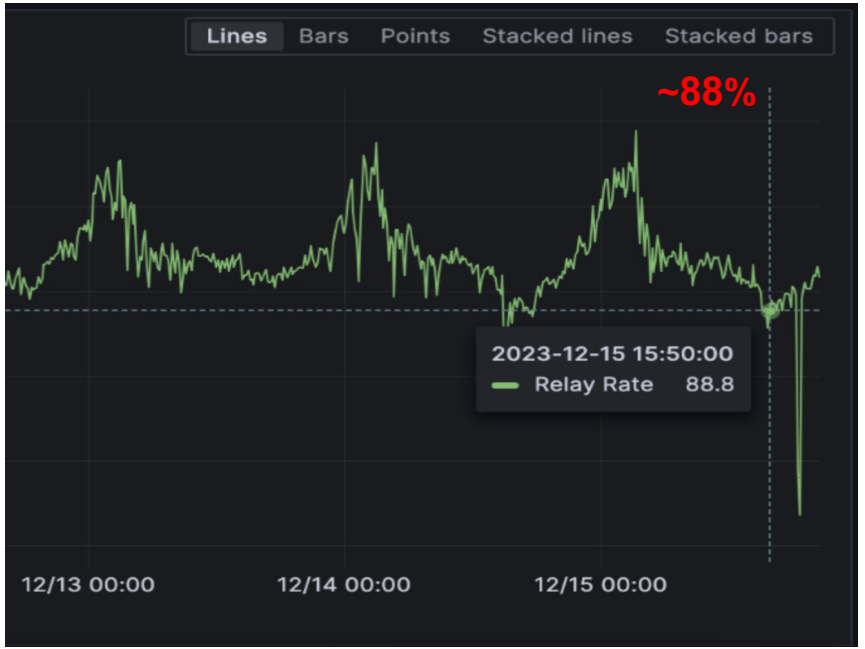

Post-cutover

~5% Drop Observed in Push Path

The Problem: Drop in Delivery Rate for Drivers

Pre-cutover

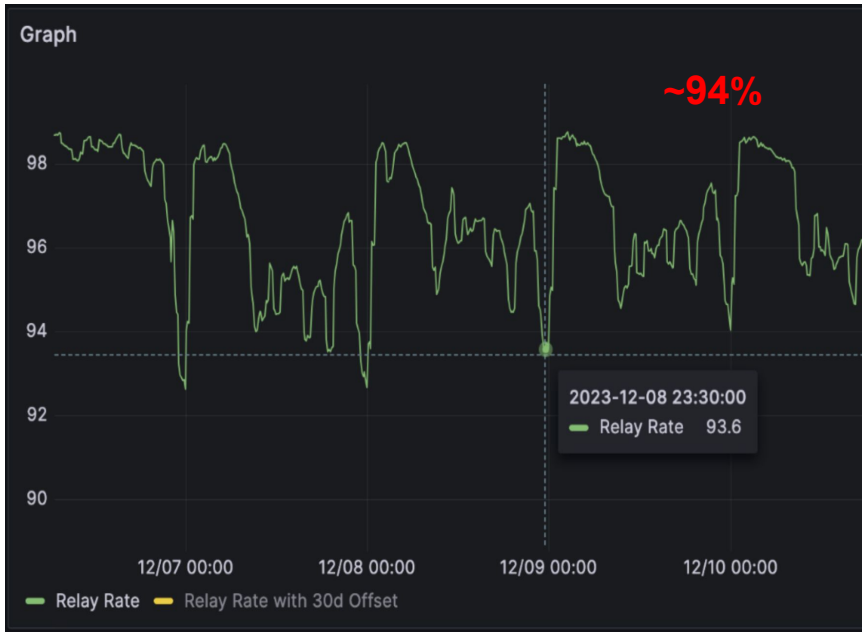

Post-cutover

~4% Drop Observed in Push Path

Investigation: Reduce Surface Area

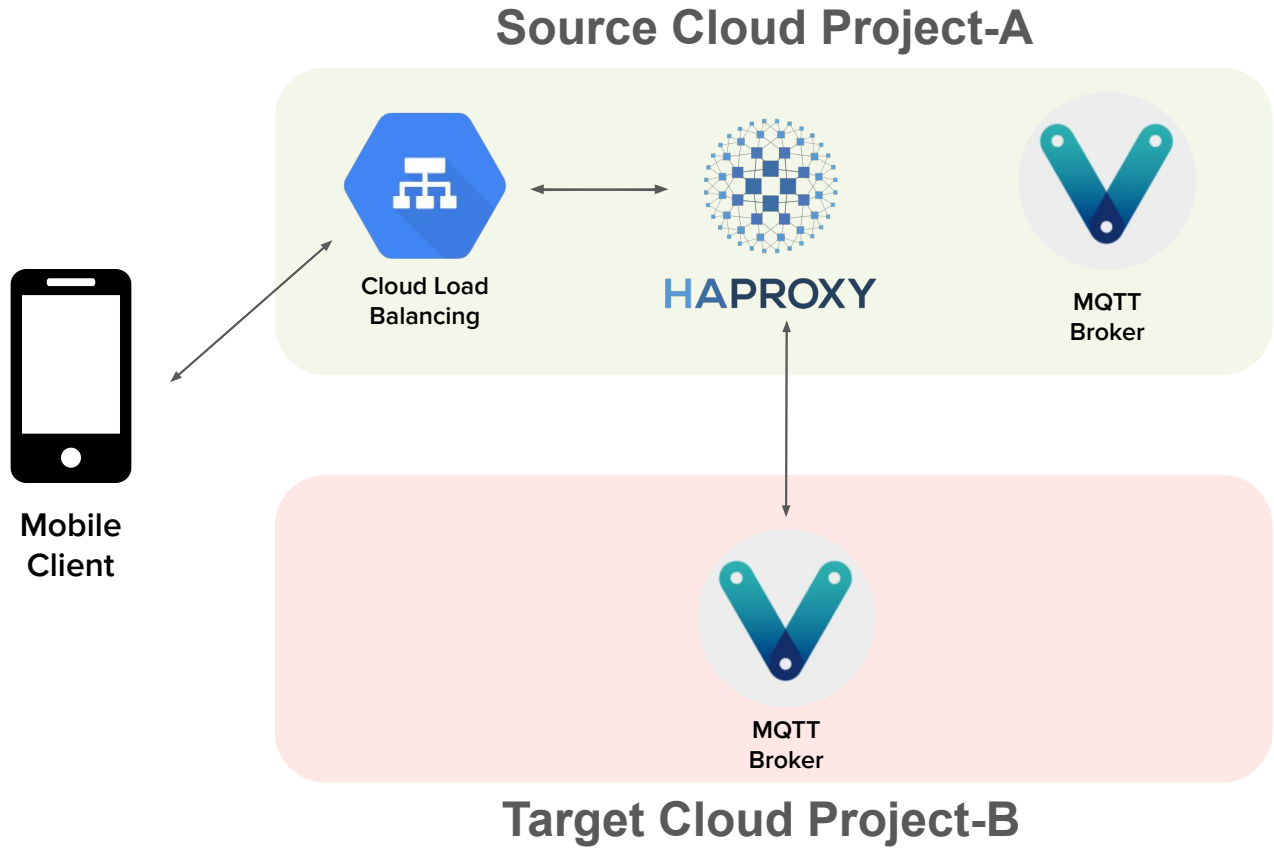

Investigation: Reduce Surface Area

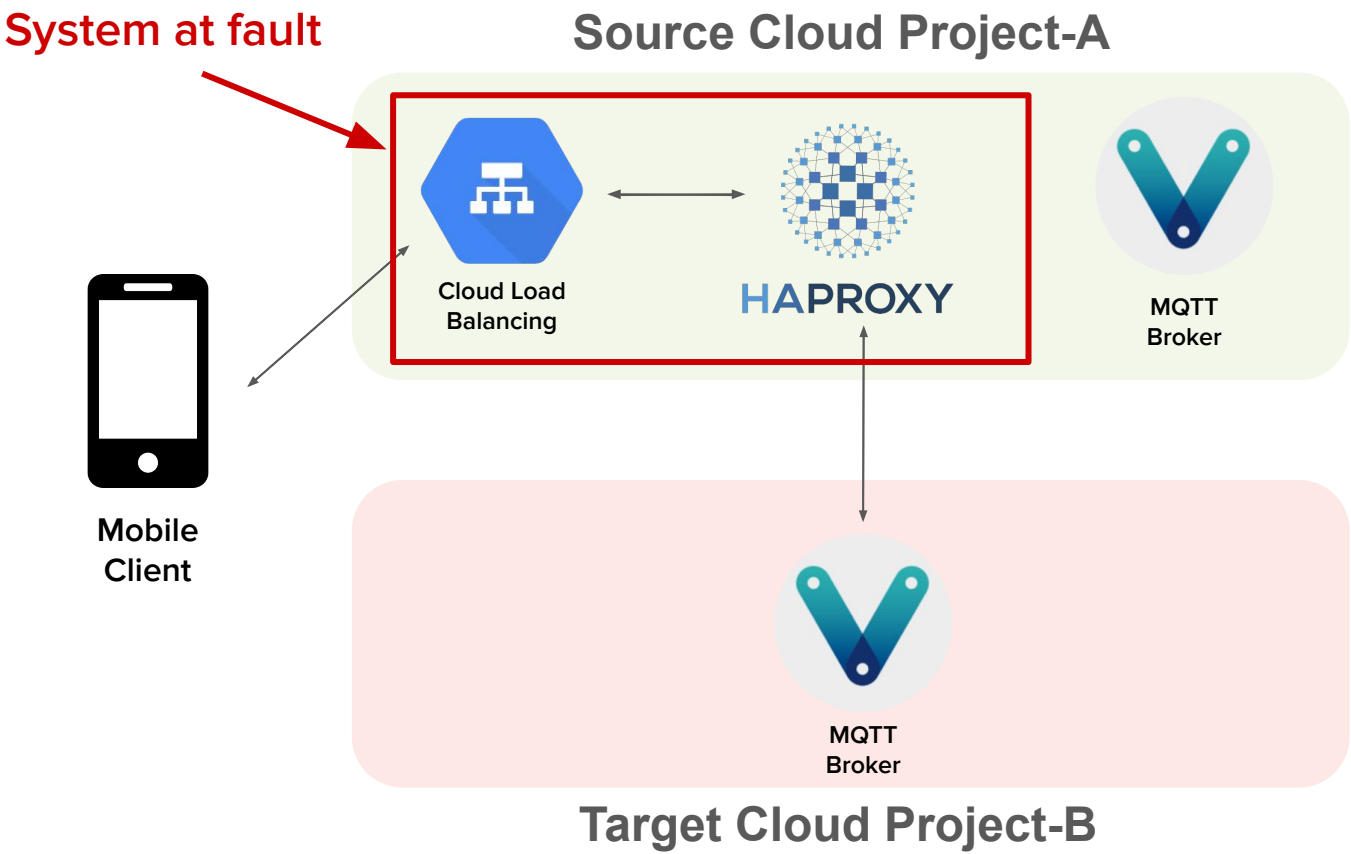

Investigation: VM Network Metrics

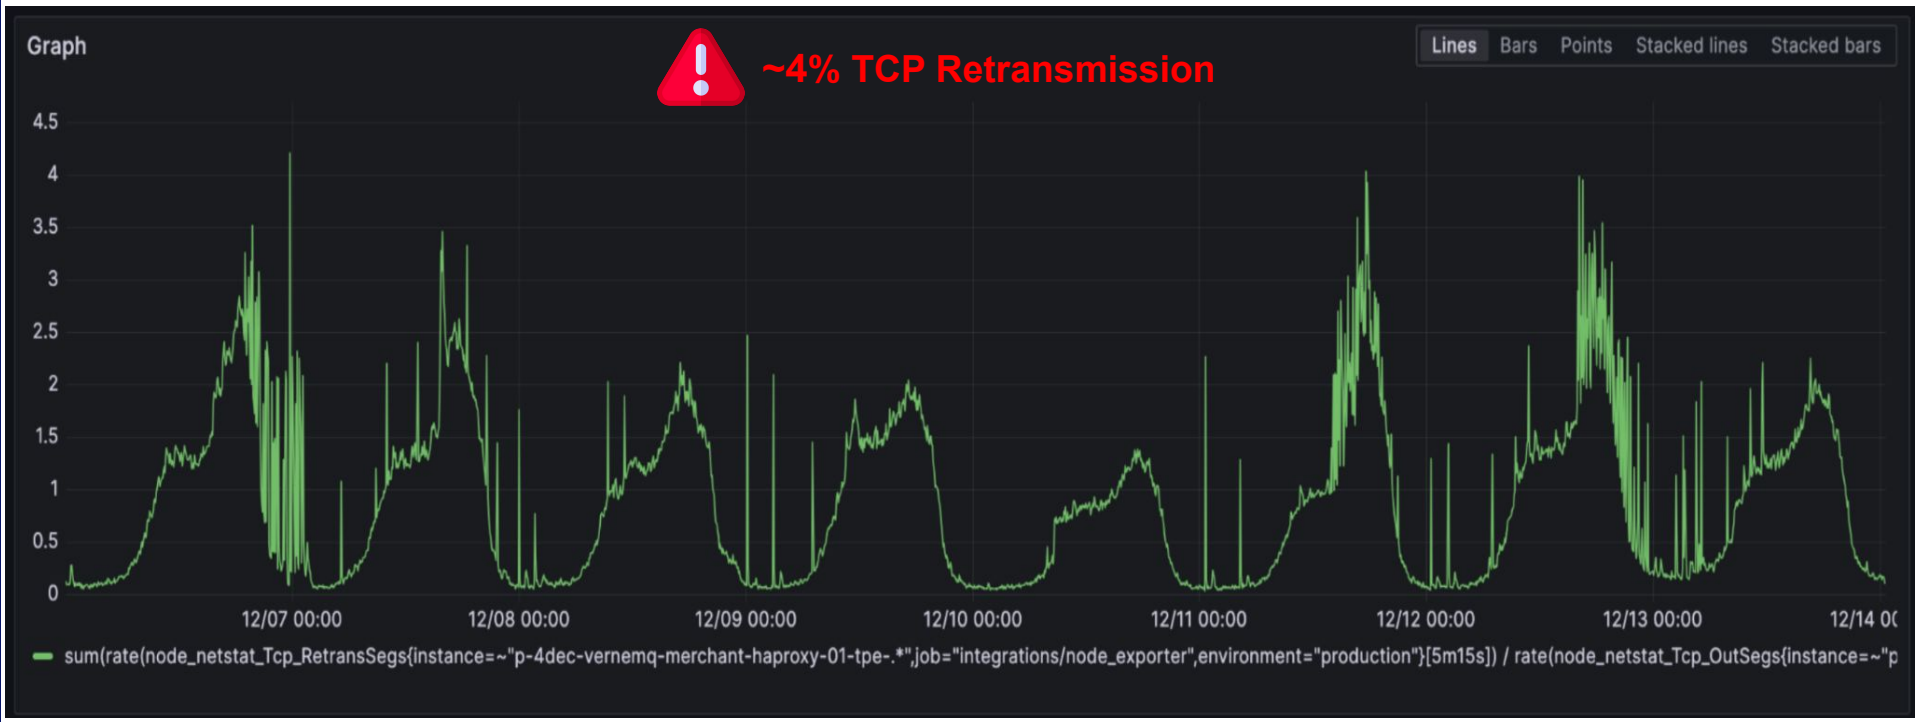

Investigation: Load Test & Wireshark Analysis

Observability Summit North America

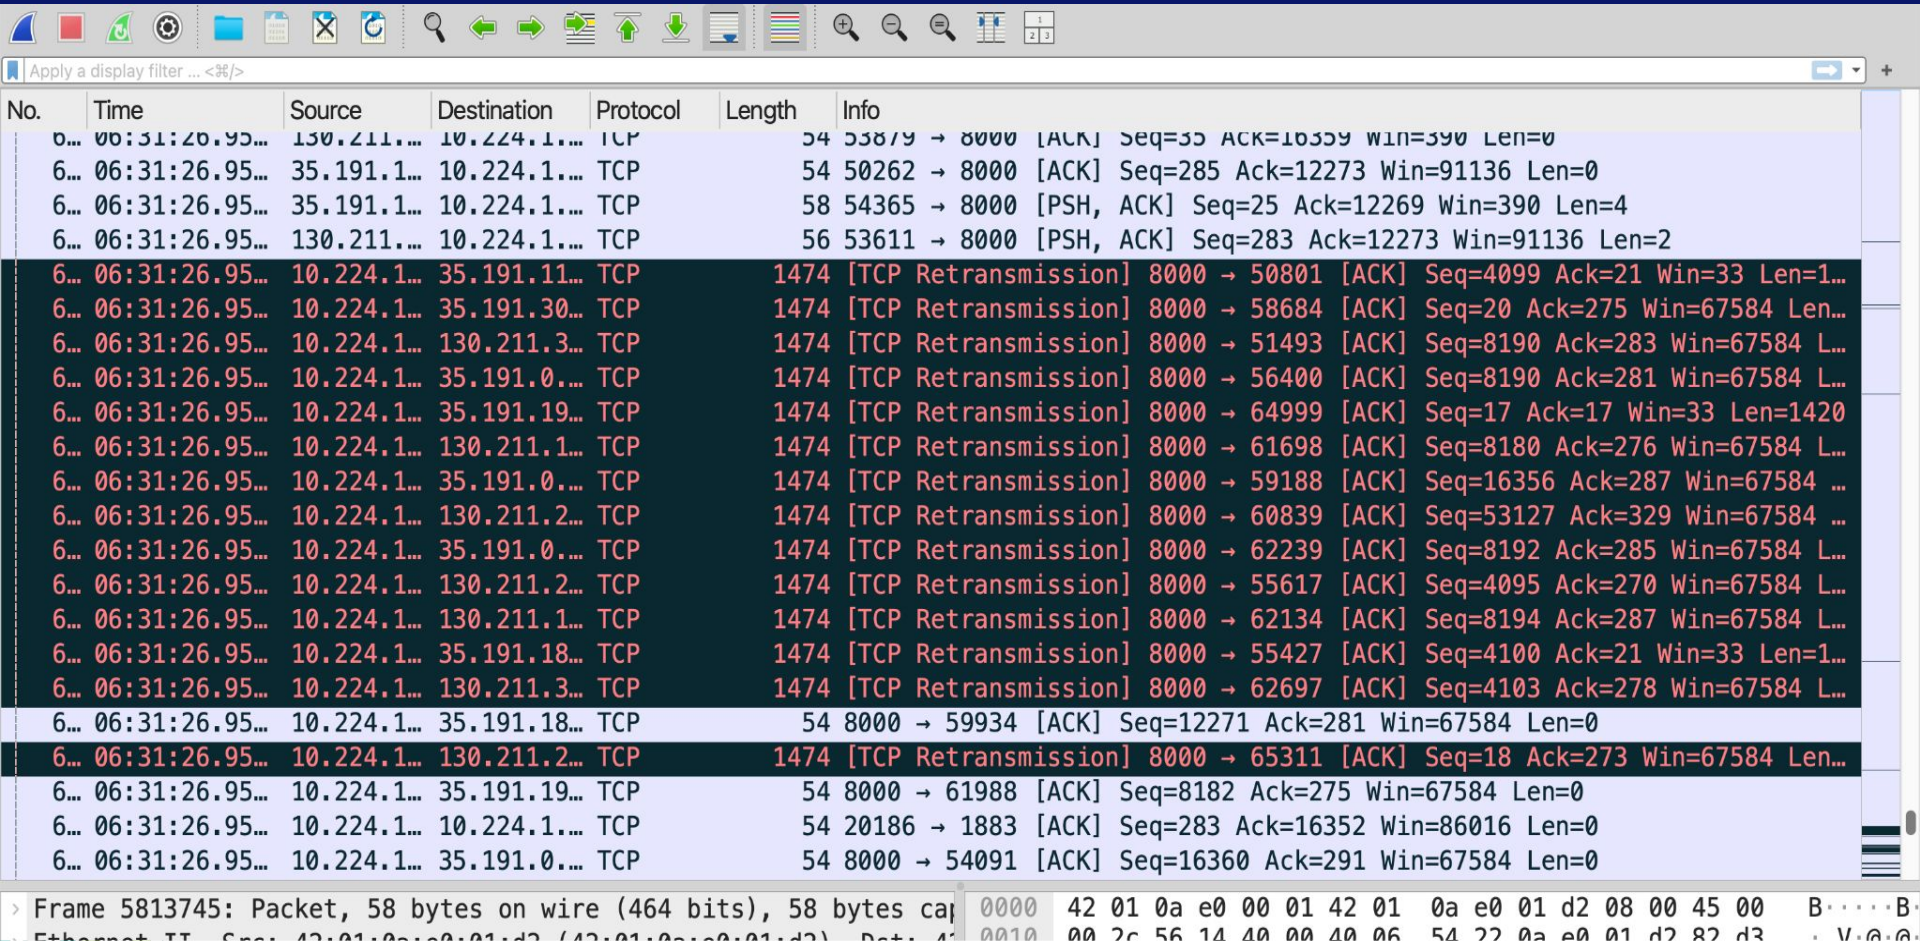

Investigation: Cloud Provider VM Metrics

Observability Summit North America

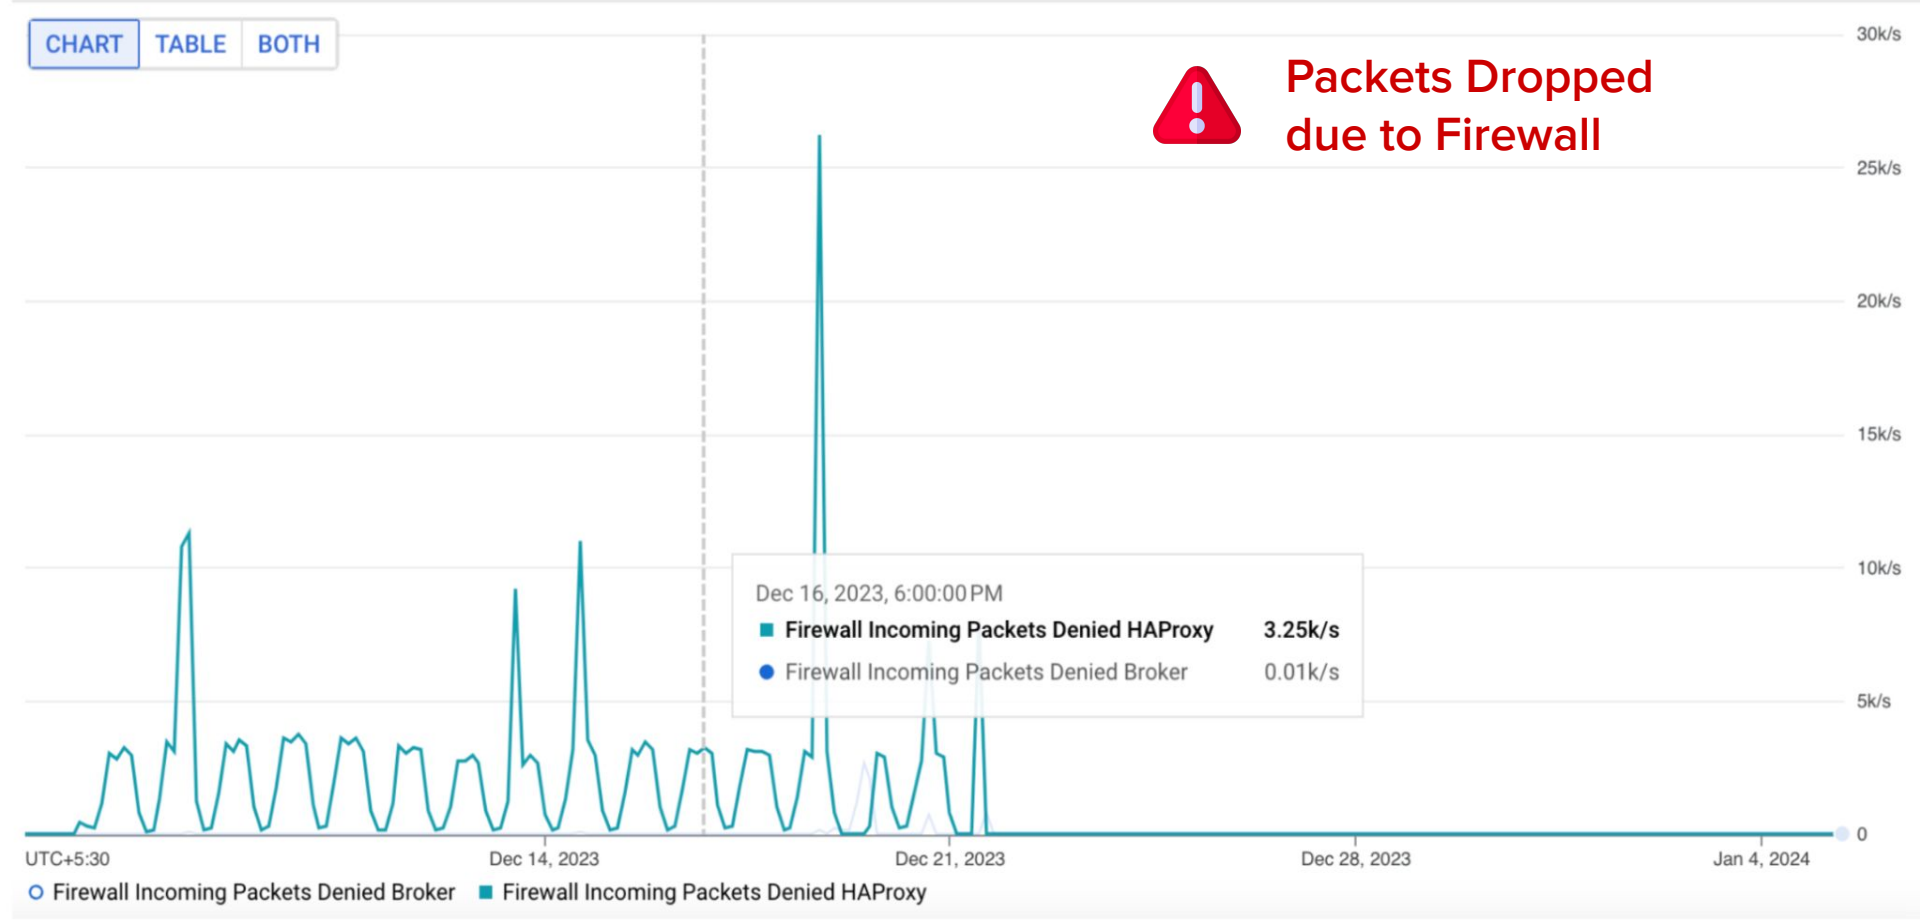

Investigation Results & Next Steps

1. No impact on MQTT setup for customer user type
2. Impact during peak hours & push path only for merchant & driver partners
3. High tcp restrans on HAProxy resulting in increased latency in push path
4. Despite firewall correctly configured on HAProxy VM, incoming packets getting randomly dropped with the reason as "denied"

HAPROXY

Allow outbound tcp traffic to VerneMQ's port 1883

Allow inbound tcp traffic from HAProxy

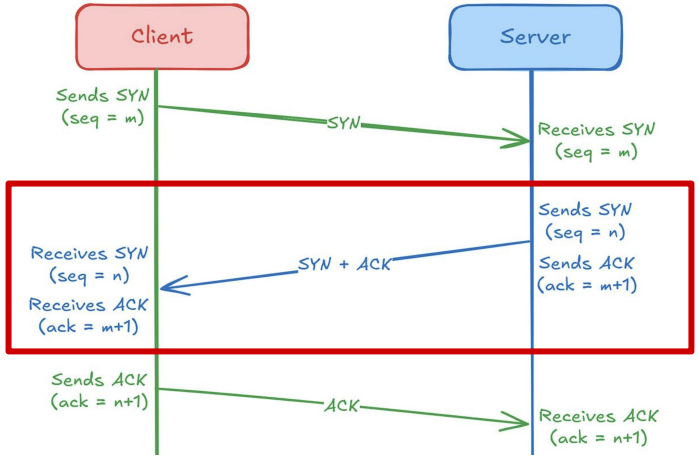

Workaround: Additionally Whitelist Reverse Rules

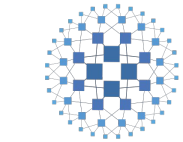

Allow outbound tcp traffic to VerneMQ's port 1883	Allow inbound tcp traffic from HAProxy
Allow inbound tcp traffic from VerneMQ's port 1883	Allow outbound tcp traffic to HAProxy

New Rules

Workaround: Result

Recovered

Root Cause & Lessons Learned

Root Cause

HAProxy VM conntrack table hits the limit.

✅ Updated documentation

✅ Exported connection tracking table metrics graph

Lessons Learned

1. Observability is the key
   a. Export cloud provider's metrics
   b. Always plot tcp retrans in %
2. Load test using exact same setup as production
3. Tools like TCPdump & wireshark for packet level analysis are quite effective.

Observability Summit North America

Questions?

gojek.github.io/courier

Dhruv Jain - GoTo Group (Gojek), Indonesia

www.linkedin.com/in/dhruvjain99

Connecting Your Laptop to Present

Observability Summit North America

If you plan to use speaker notes, please follow the steps below to ensure your laptop is set up correctly. This will allow you to view your notes while the audience sees only your slides.

Before Connecting

- Close all unnecessary apps & tabs
- Turn off notifications
- Disable Night Mode / TrueTone
- Don’t enter full-screen presentation mode yet

At the podium

- Plug in the HDMI cable
- If needed, use your USB-C to HDMI adapter
- On Mac, click “Allow” if prompted

USE “EXTENDED DISPLAY” MODE (NOT MIRRORED)

Using MAC

- Go to System Settings > Displays
- Click the “Arrangement” tab
- Make sure “Mirror Displays” is unchecked

Using PC

- Press ⊞Win + P, then select “Extend”
- Or: Right-click on the Desktop > Display Settings
- Under Multiple Displays, select “Extend these displays”

Connecting Your Laptop to Present

Observability Summit North America

When Using PowerPoint

- Start slideshow: ⌘Cmd + Enter (Mac) or F5 (PC)
- Or: Click Slide Show > Play from Start
- If notes show on the wrong screen, click “Swap Displays”
- **https://www.youtube.com/watch?v=gQ3D4m-5pww**

Google Slides

- Use Google Chrome as your browser
- Click the dropdown next to ‘Slideshow’ > select Presentation Display Options
- Check “Presenter View” and “Full Screen”
- If prompted, select "Allow"
- Choose the external display (may have a strange name like PanasonicXYZ)
- Click Start Slideshow
- **https://www.youtube.com/watch?v=-GT7WCvPcys**