



KubeCon



CloudNativeCon

India 2026

The 900-Second Ghost: Debugging Half-Open TCP

Dhruv Jain - GoTo Group (Gojek), Indonesia



MQTT Infrastructure Operational Scale



KubeCon



CloudNativeCon

India 2026

2M

Concurrent
Connections

15M


MQTT DAUs

3.2M

RPM Throughput

2.2%

Indonesia's GDP



Background: MQTT Basics



KubeCon



CloudNativeCon




India 2026

Lightweight L7 Protocol

3 QoS Levels

Bi-directional & Full Duplex

Pub-Sub Model



Background: Use Cases



KubeCon



CloudNativeCon

India 2026



Customers

150K concurrent connections



Driver Partners

600K concurrent connections



Merchant Partners

450K concurrent connections

Background: Use Cases



KubeCon



CloudNativeCon

India 2026

Key MQTT Communication Scenarios:



Order Updates & Real-time status tracking



Chat for seamless communication



Push Notifications for user engagement



Telemetry & Real-time user monitoring



Dynamic Order Bidding for driver partners



Driver location pings



Search typeahead suggestions



Order for merchant partners

MQTT Setup

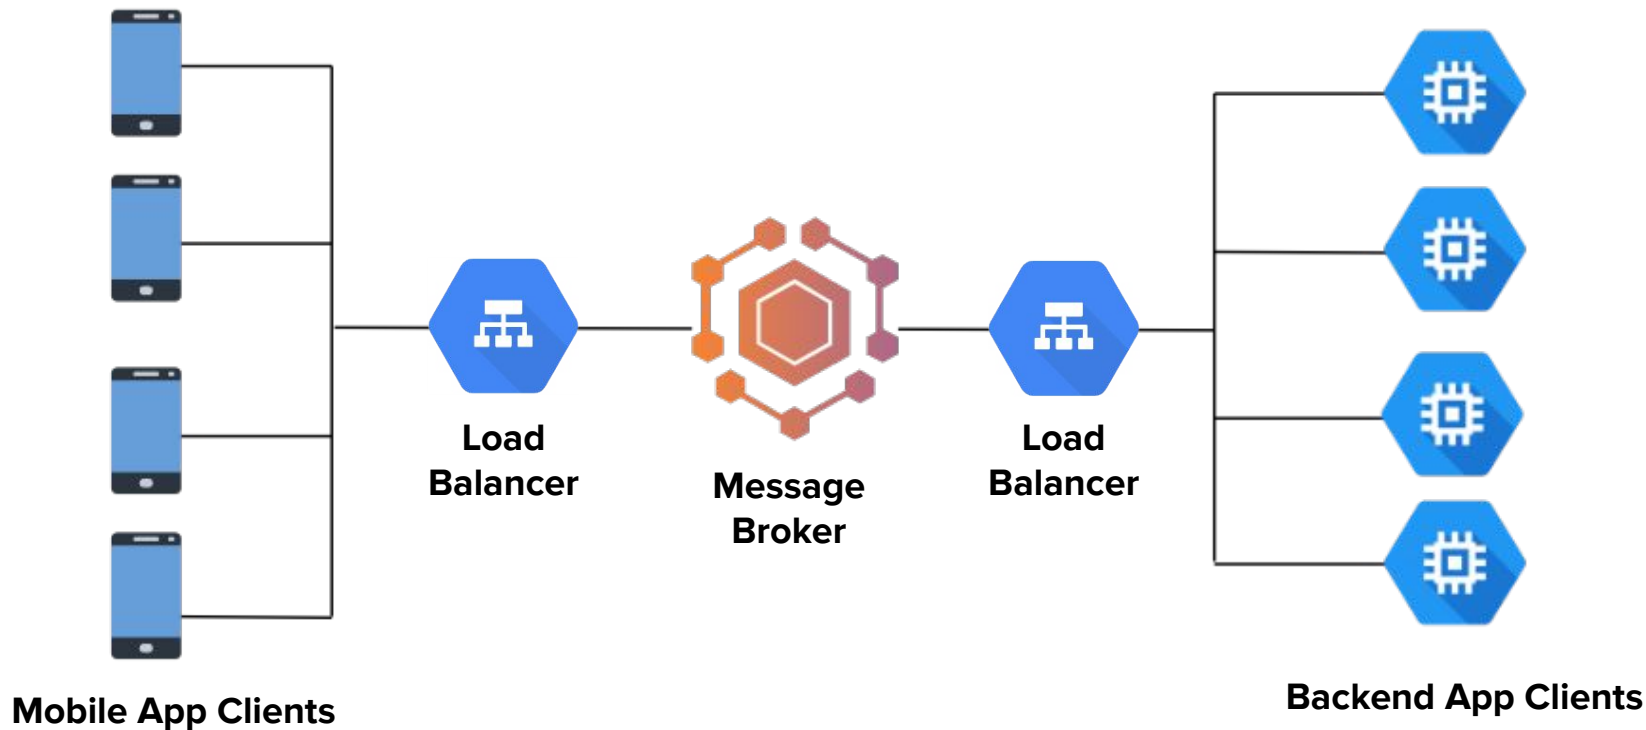


KubeCon



CloudNativeCon

India 2026



The Trigger: Deployments



KubeCon



CloudNativeCon

India 2026




TCP connection terminated but the backend client never knew



MQTT broker



**Load
Balancer**



**Backend
Client**

Impact



KubeCon



CloudNativeCon

India 2026

Queue Buildup

Latency Spike

Relay Rate Drop



Automatic Recovery in 900s

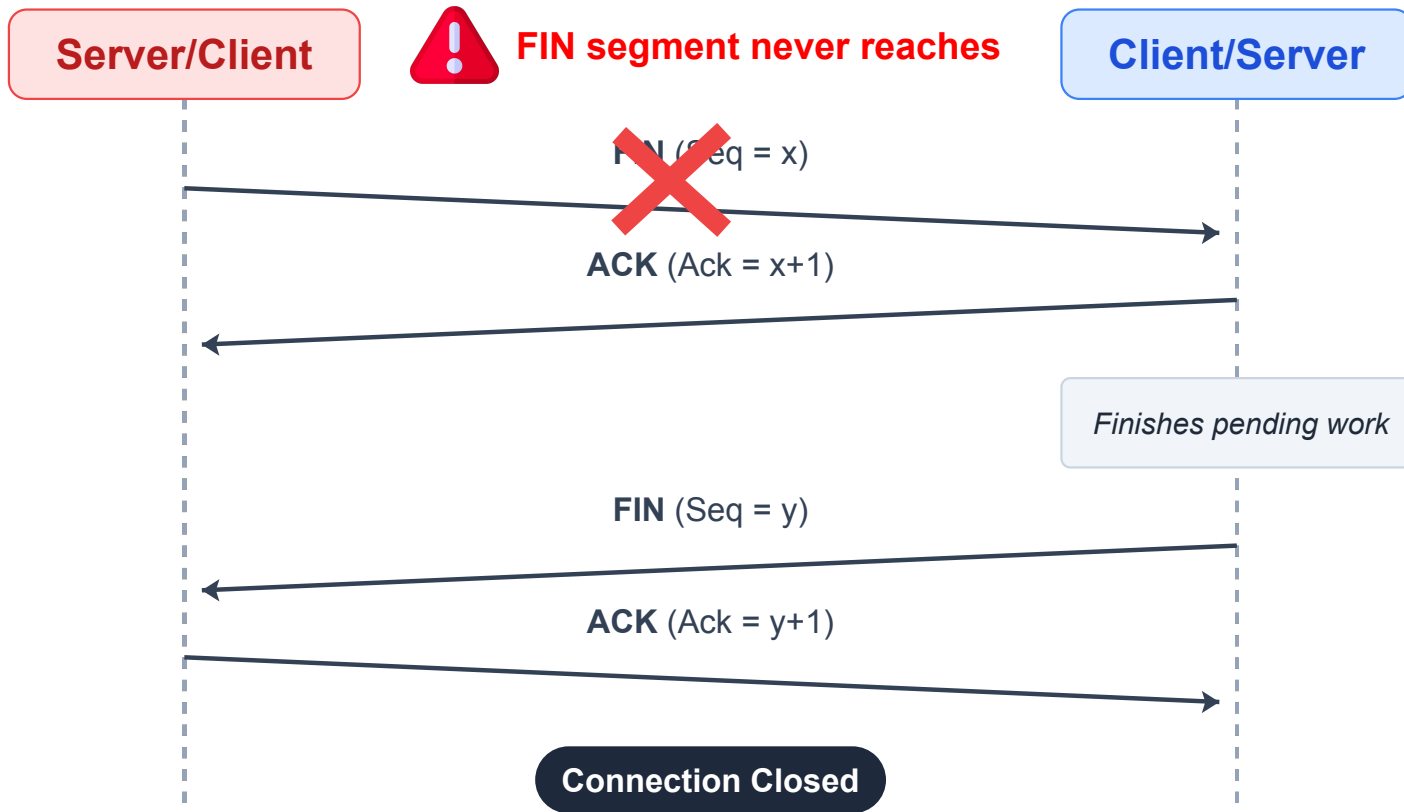
Queue Size



Queue Latency



Problem: Half-Open TCP Connection



Mitigation Technique: Half-Open TCP Connection

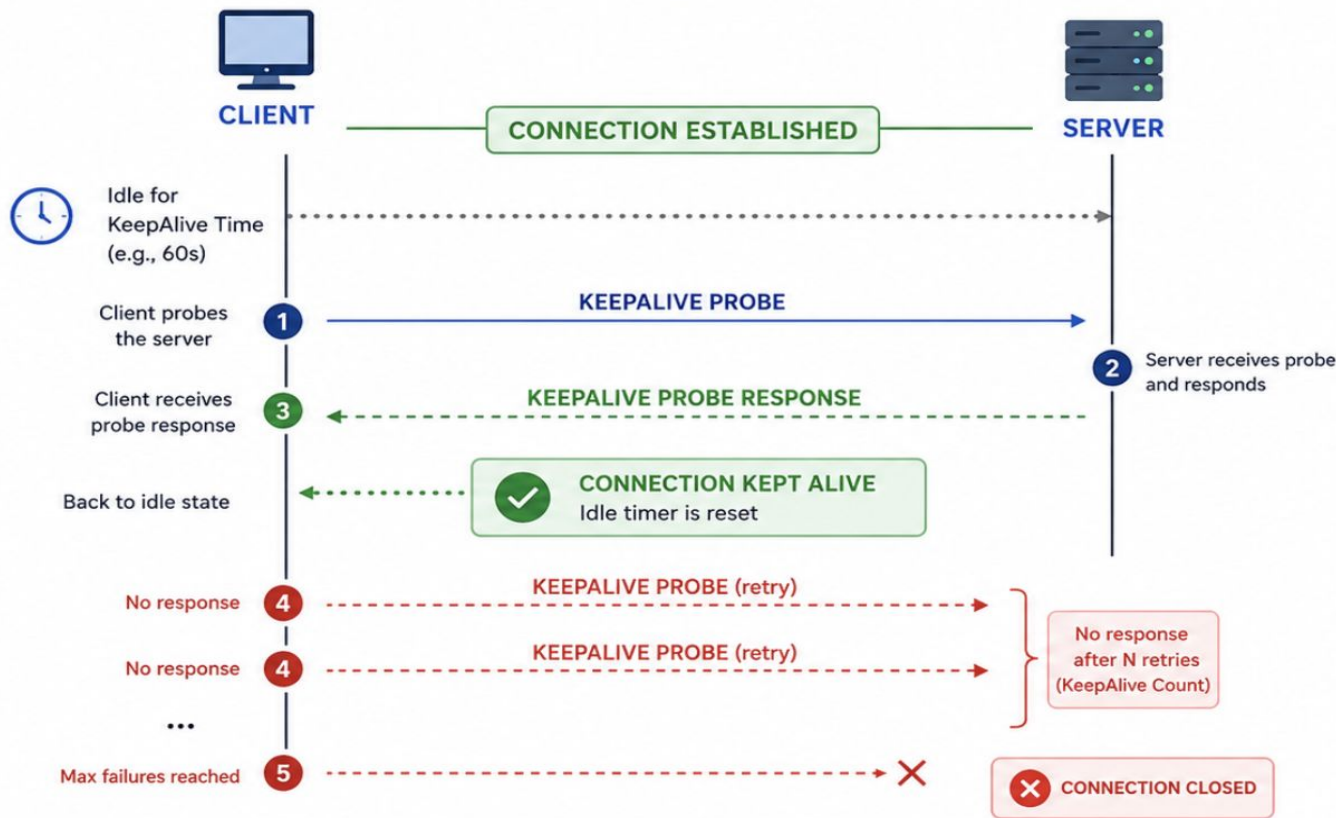


KubeCon



CloudNativeCon

India 2026



TCP KeepAlive Mechanism

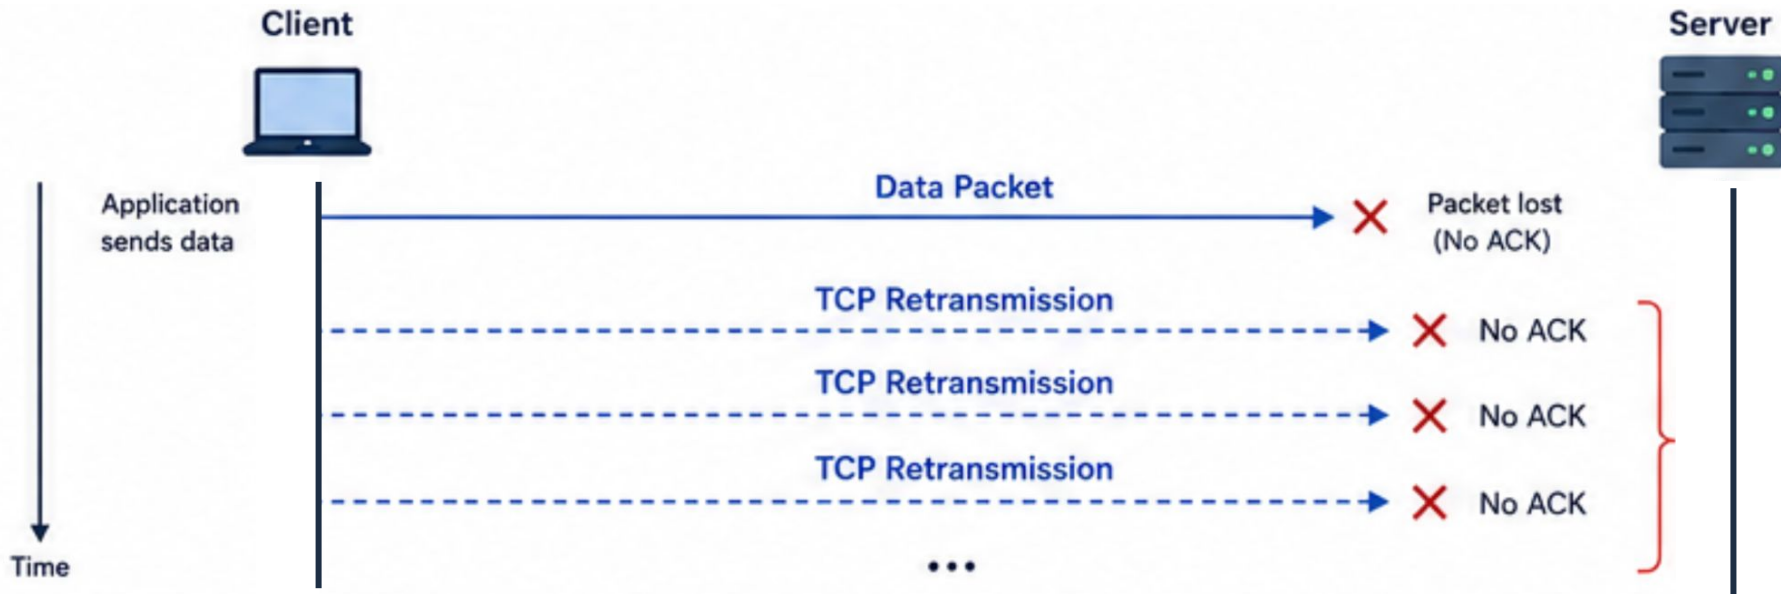


KubeCon



CloudNativeCon

India 2026



KeepAlive mechanism never works

MQTT KeepAlive Mechanism



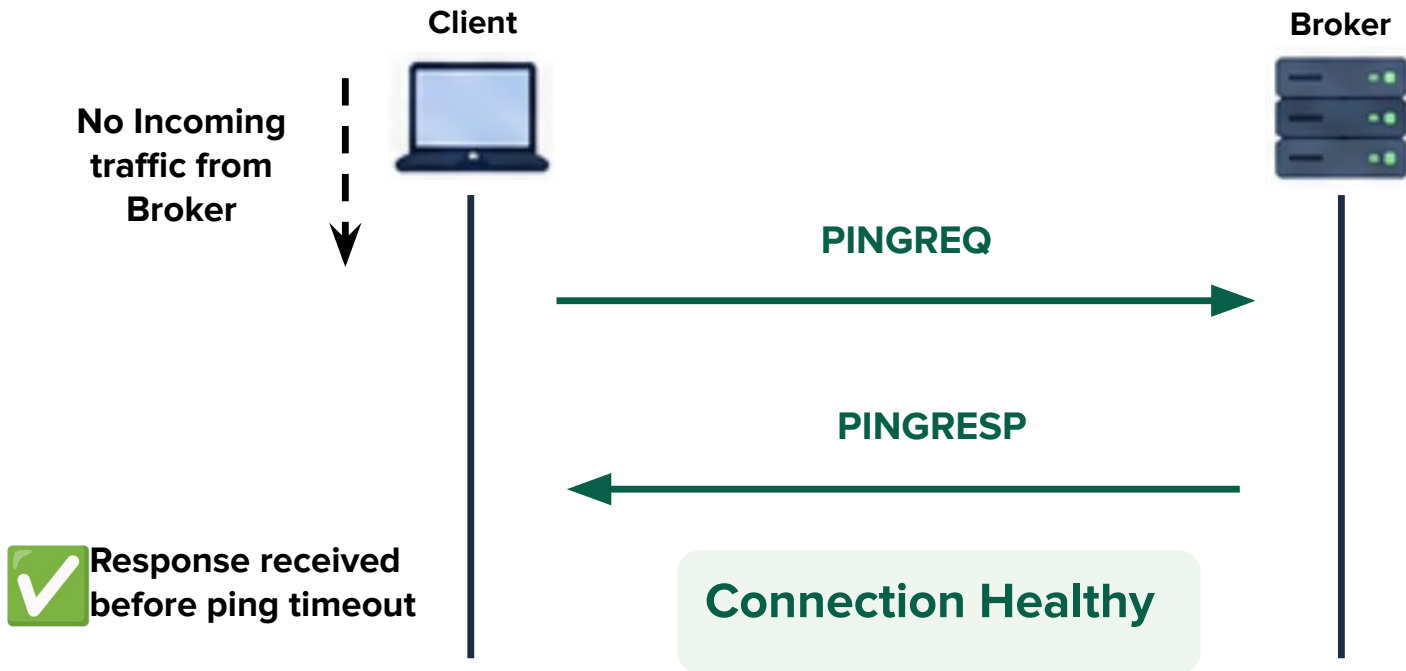
KubeCon



CloudNativeCon

India 2026

Normal Operation (No Incoming Traffic from Broker)



MQTT KeepAlive Mechanism



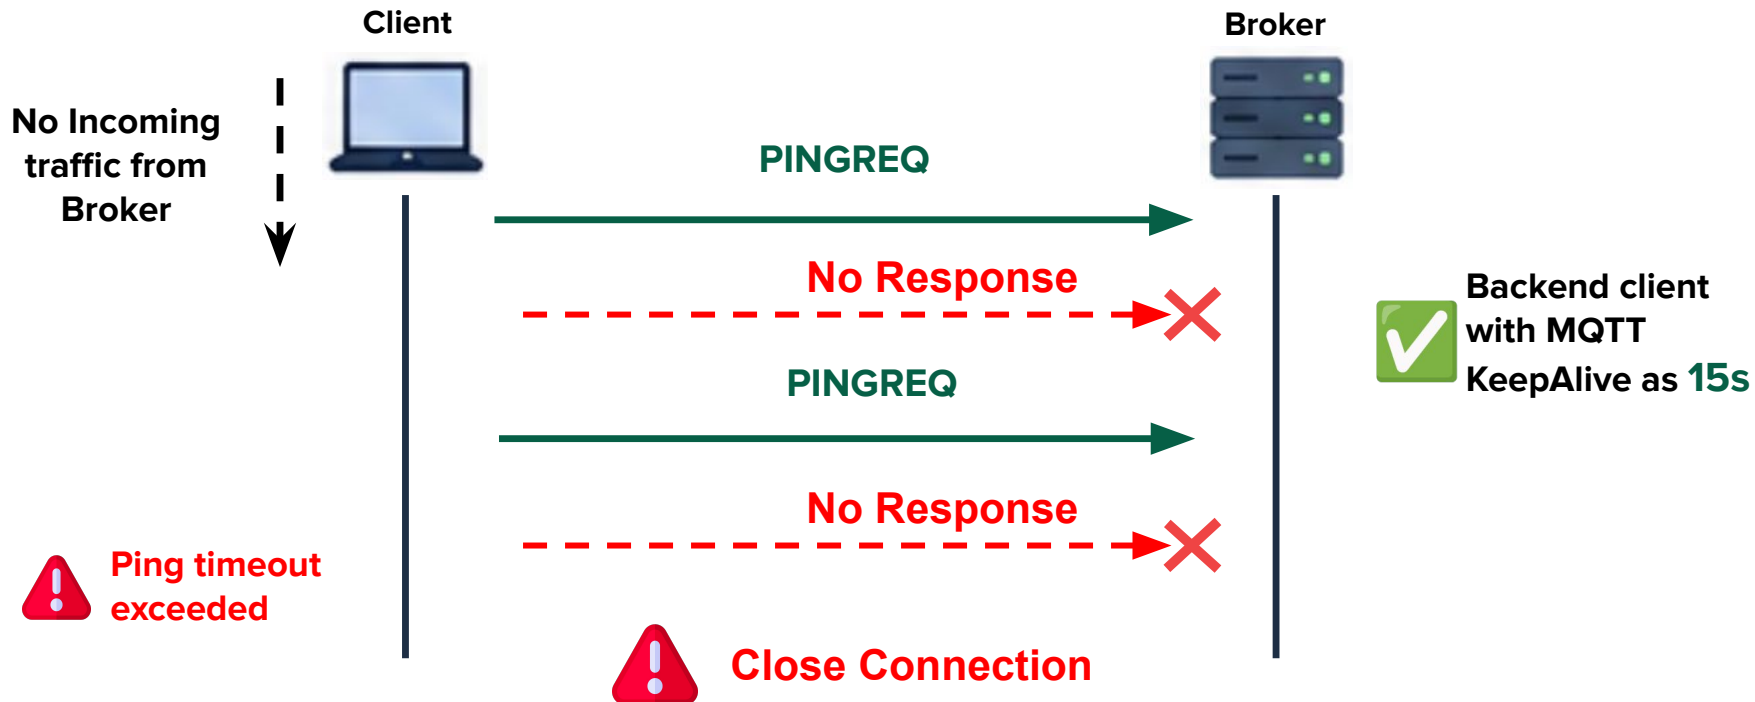
KubeCon



CloudNativeCon

India 2026

Failure Detection (No Response from Broker)



Open Questions



KubeCon



CloudNativeCon

India 2026

01 Why FIN packet drops?

02 Why does not MQTT KeepAlive mechanism work?

03 Despite shorter KeepAlive, why there is a 900s downtime?

Investigation: Replicate on Staging env



KubeCon



CloudNativeCon

India 2026

Attempt #1

Status: Failed

1 MQTT connection

Key Observation

"Backend services with lower throughput never faced this problem of 900s"

Attempt #2

Status: Success

1 MQTT connection with high publish throughput

Next Step

- Capture **tcpdump** for detailed analysis.
- Perform comparison with production network traffic analysis.

Investigation: TCPDump & Wireshark Analysis - Cloud Load Balancer



KubeCon

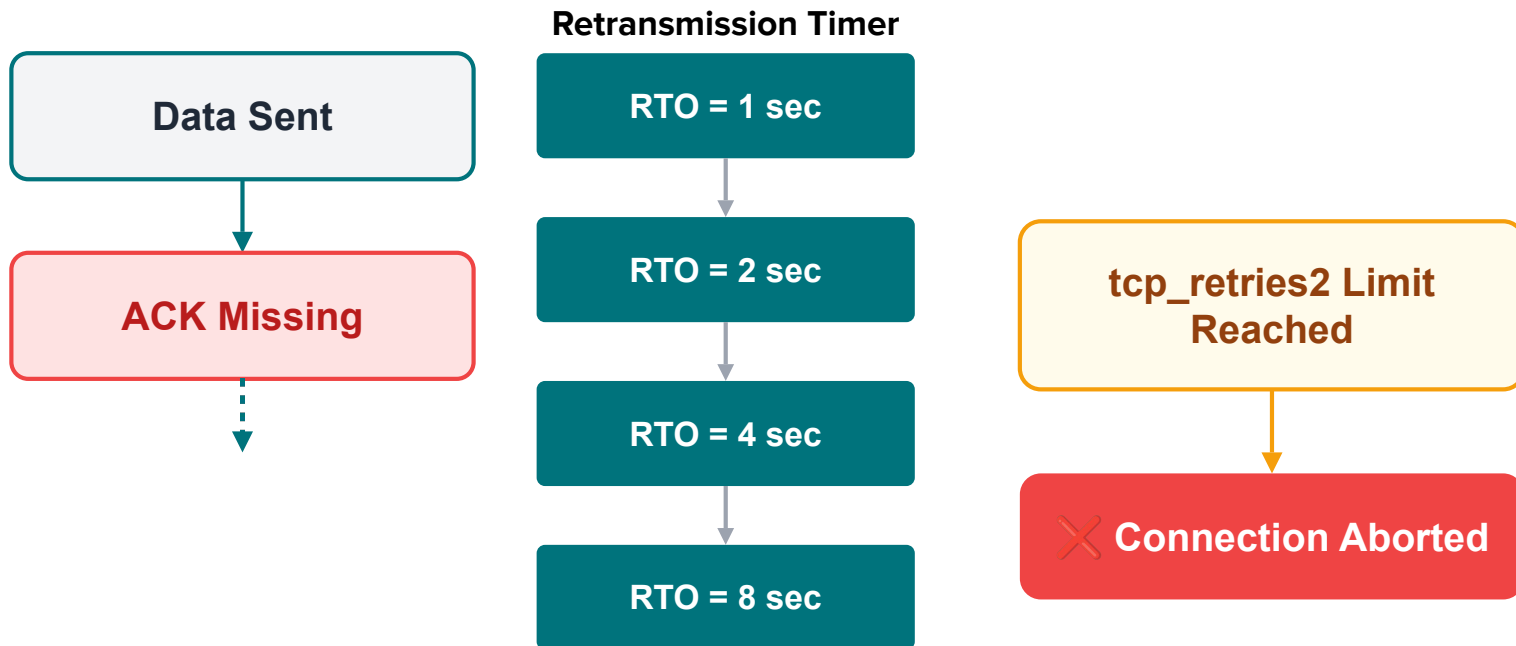


CloudNativeCon

India 2026

78914362	78914416	1377682783	53938	1883	10.108.0.53	[TCP Retransmission] 53938 → 1883 [PSH, ACK] Seq=78914362 Ack=1377682783 Win=65536 Len=54
1377682783	1377682787	78914416	1883	53938	10.108.0.53	Publish Ack (id=8844)
1377682783	1377682787	78914416	1883	53938		[TCP Retransmission] 1883 → 53938 [PSH, ACK] Seq=1377682783 Ack=78914416 Win=73728 Len=4
78914416	78914416	1377682787	53938	1883		53938 → 1883 [ACK] Seq=78914416 Ack=1377682787 Win=65536 Len=0
78914416	78914416	1377682787	53938	1883	10.108.0.53	[TCP Dup ACK 240#1] 53938 → 1883 [ACK] Seq=78914416 Ack=1377682787 Win=65536 Len=0
78914415	78914415	1377682787	53938	1883		[TCP Keep-Alive] 53938 → 1883 [ACK] Seq=78914415 Ack=1377682787 Win=65536 Len=0
1377682787	1377682788	78914416	1883	53938	10.108.0.53	1883 → 53938 [FIN, ACK] Seq=1377682787 Ack=78914416 Win=73728 Len=0
1377682787	1377682788	78914416	1883	53938	10.108.0.53	[TCP Retransmission] 1883 → 53938 [FIN, ACK] Seq=1377682787 Ack=78914416 Win=73728 Len=0
1377682787	1377682788	78914416	1883	53938	10.108.0.53	[TCP Retransmission] 1883 → 53938 [FIN, ACK] Seq=1377682787 Ack=78914416 Win=73728 Len=0
78914416	78914470	1377682787	53938	1883		Publish Message (id=8845) [discard]
78914416	78914470	1377682787	53938	1883		[TCP Retransmission] 53938 → 1883 [PSH, ACK] Seq=78914416 Ack=1377682787 Win=65536 Len=54
78914416	78914470	1377682787	53938	1883		[TCP Retransmission] 53938 → 1883 [PSH, ACK] Seq=78914416 Ack=1377682787 Win=65536 Len=54
78914416	78914470	1377682787	53938	1883		[TCP Retransmission] 53938 → 1883 [PSH, ACK] Seq=78914416 Ack=1377682787 Win=65536 Len=54
78914416	78914470	1377682787	53938	1883		[TCP Retransmission] 53938 → 1883 [PSH, ACK] Seq=78914416 Ack=1377682787 Win=65536 Len=54
78914416	78914470	1377682787	53938	1883		[TCP Retransmission] 53938 → 1883 [PSH, ACK] Seq=78914416 Ack=1377682787 Win=65536 Len=54

900s Mystery: Understanding tcp_retries2



Application may wait minutes before TCP declares the peer dead.

Why does not MQTT KeepAlive work?



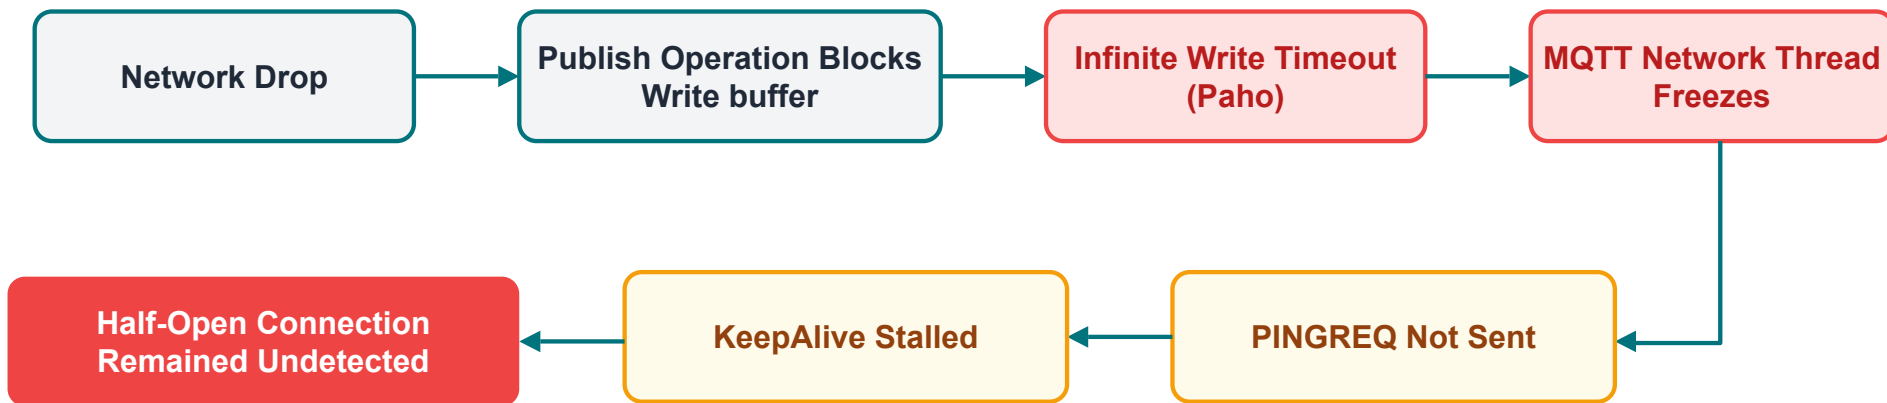
KubeCon



CloudNativeCon

India 2026

The Failure Chain



Make KeepAlive Ping Timeout
Detector Process **Non-Blocking**

Why FIN packets dropped?





KubeCon



CloudNativeCon

India 2026



Bug in cloud provider's fully managed load balancer

Workaround was to enable forceful termination setting

Lessons Learned



KubeCon



CloudNativeCon

India 2026

01. MQTT in Mobile Apps

02. Kernel Parameter Synergy

03. L7 vs L4 KeepAlive

04. Packet-Level Analysis




Questions?

Dhruv Jain - GoTo Group
(Gojek), Indonesia

www.linkedin.com/in/dhruvjain99



gojek.github.io/courier



The 900-Second Ghost: Debugging Half-Open TCP

Connecting Your Laptop to Present

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 \square 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

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>