Comparing the Network Performance of AWS, Azure and GCP

Archana Kesavan, ThousandEyes

Agenda

- Report Genesis & Overview
- Research Methodology
- Research Findings
- Summary & Recommendations
- Q&A

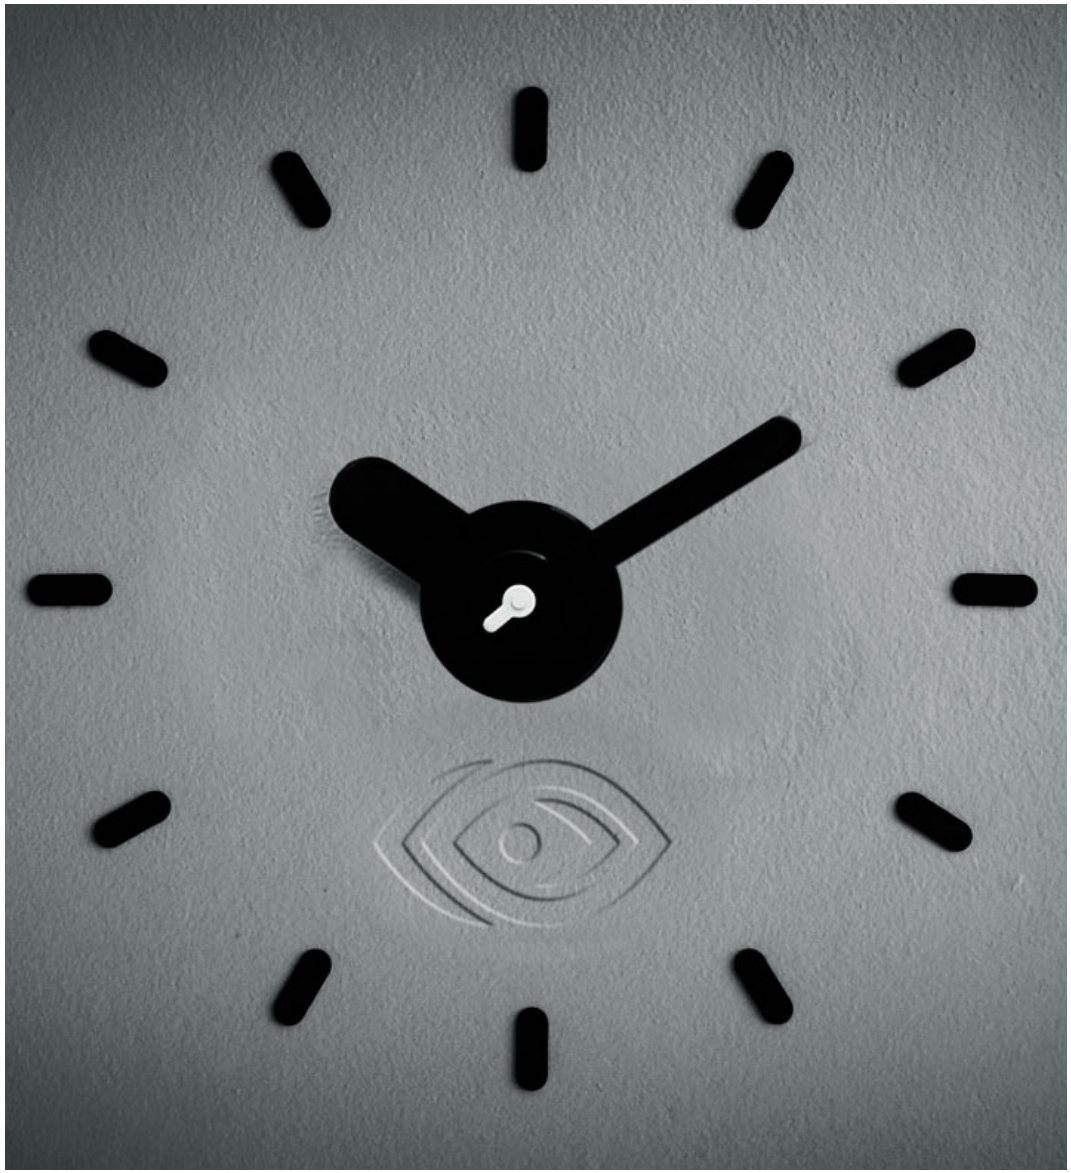

HOW DO WE MAKE DECISIONS ABOUT THE CLOUD TODAY?

Typical Cloud Decision Factors

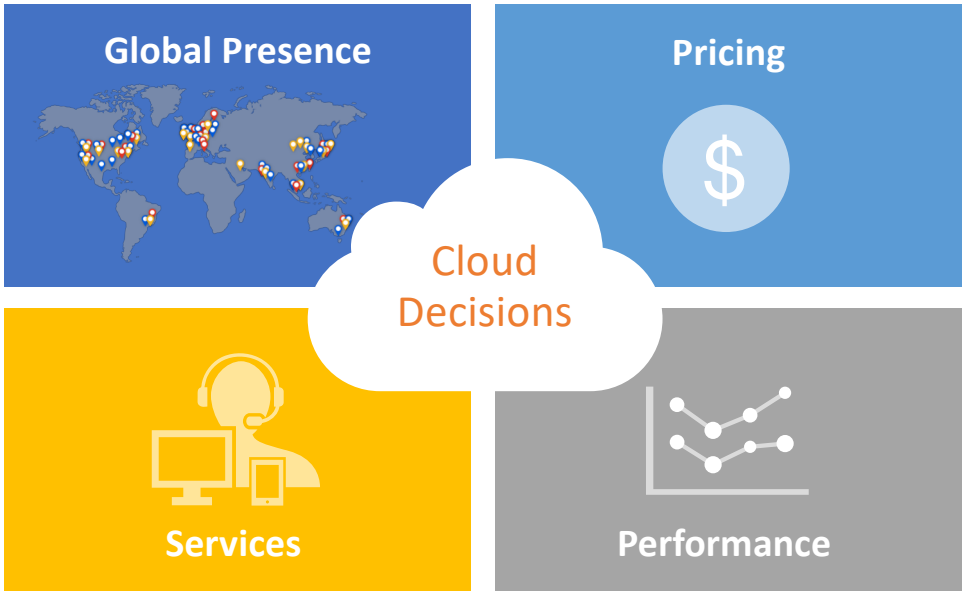

What's Been Missing from Cloud Performance Data?

Single provider
Lack comparisons
Limited geos
No multi-cloud

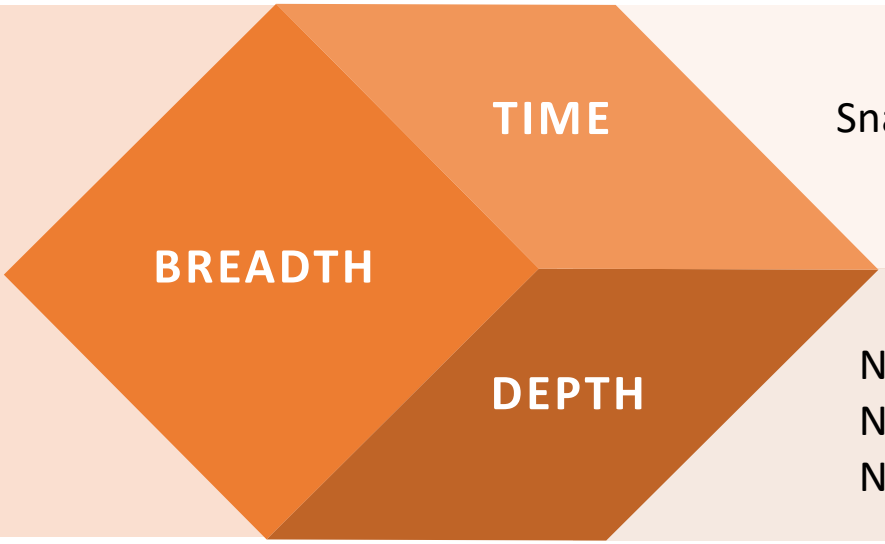

Snapshots not trends

No network architecture
No loss or jitter
No predictability analysis

Data Collection Methodology

- Software probes (monitoring agents) emulate user-data by triggering a “test”
- Customized Traceroute
  - TCP Traces. Resistant to ICMP issues caused by rate-limiting routers.
- Bidirectional
  - Basic traceroute is directed (source-to-destination)
  - Bidirectional tests to account for variation in forward and return routes

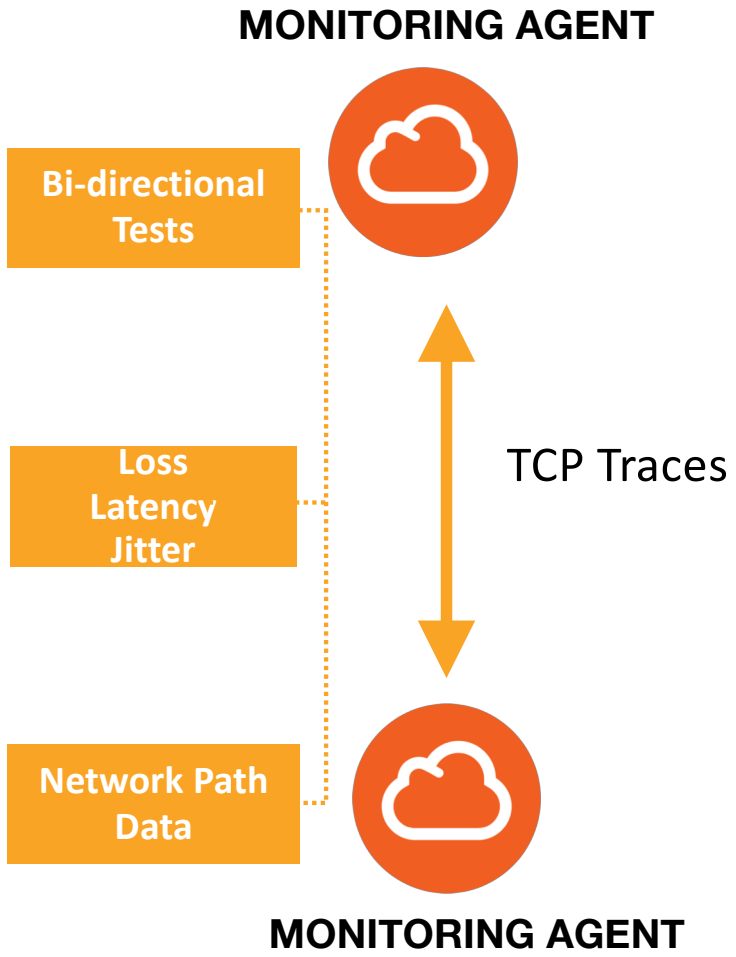

Data Collection Methodology

- A single “test” gathers different types of data
- Network Path Data:
  - Geo-located Layer 3 hops from source to destination including AS paths that allows us to determine connectivity architecture to the cloud providers
  - Per hop loss and latency
  - DSCP re-marking across the path
  - MPLS tunnels
- End-to-End Network Metrics:
  - Loss, Latency and Jitter
  - Helps understand user-experience to cloud providers and within cloud providers

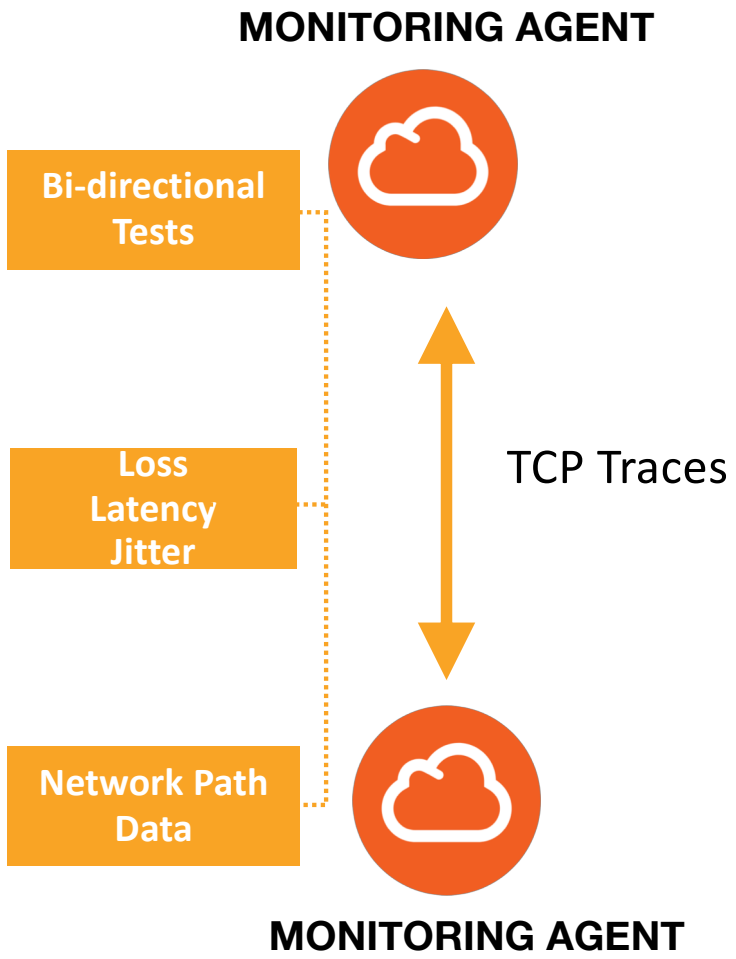

Data Processing Methodology

- Software probes are well provisioned globally with continuous access. Allows easy detection of connectivity issues or local-faults
- Tests are run periodically, at an interval of 10 minutes
- Data gathered over a period of 30 days (07.01.2018 – 07.31.2018) to accommodate for extraneous events such as outages
  - Note: Zero outages occurred during the data collection timeframe
- Data generated is continuously exported to a cloud-based platform for analysis and trend-detection
- Network metrics generated compared across all three providers
  - Latency differences across providers prominent
  - Loss and Jitter were negligible

Test Scope

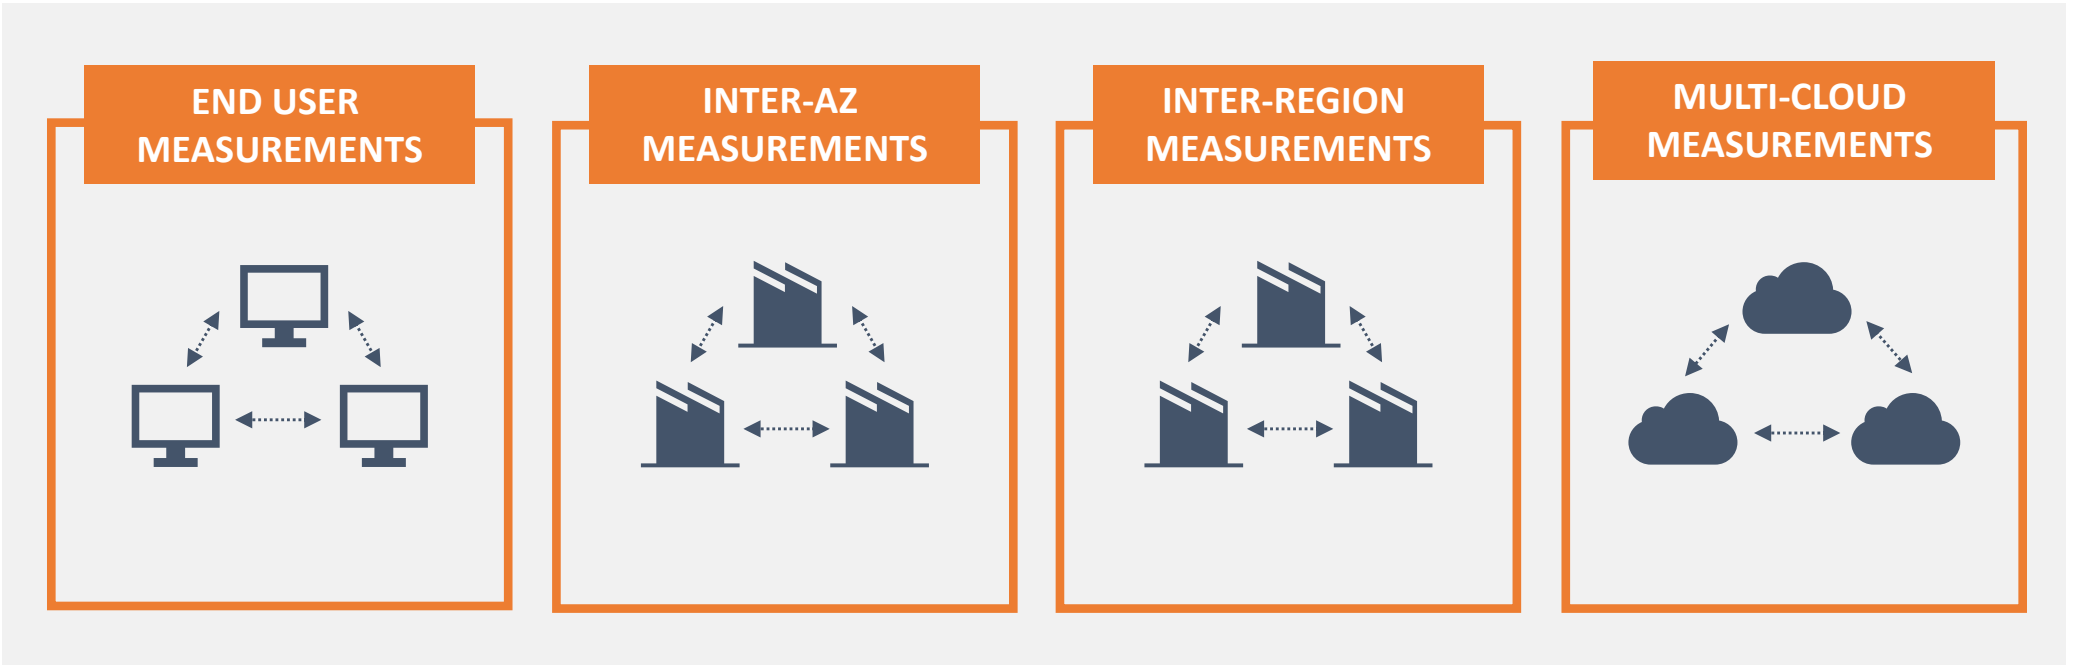
END USER MEASUREMENTS
INTER-AZ MEASUREMENTS
INTER-REGION MEASUREMENTS
MULTI-CLOUD MEASUREMENTS

End User Measurements

27 Global Locations

1485 Tests

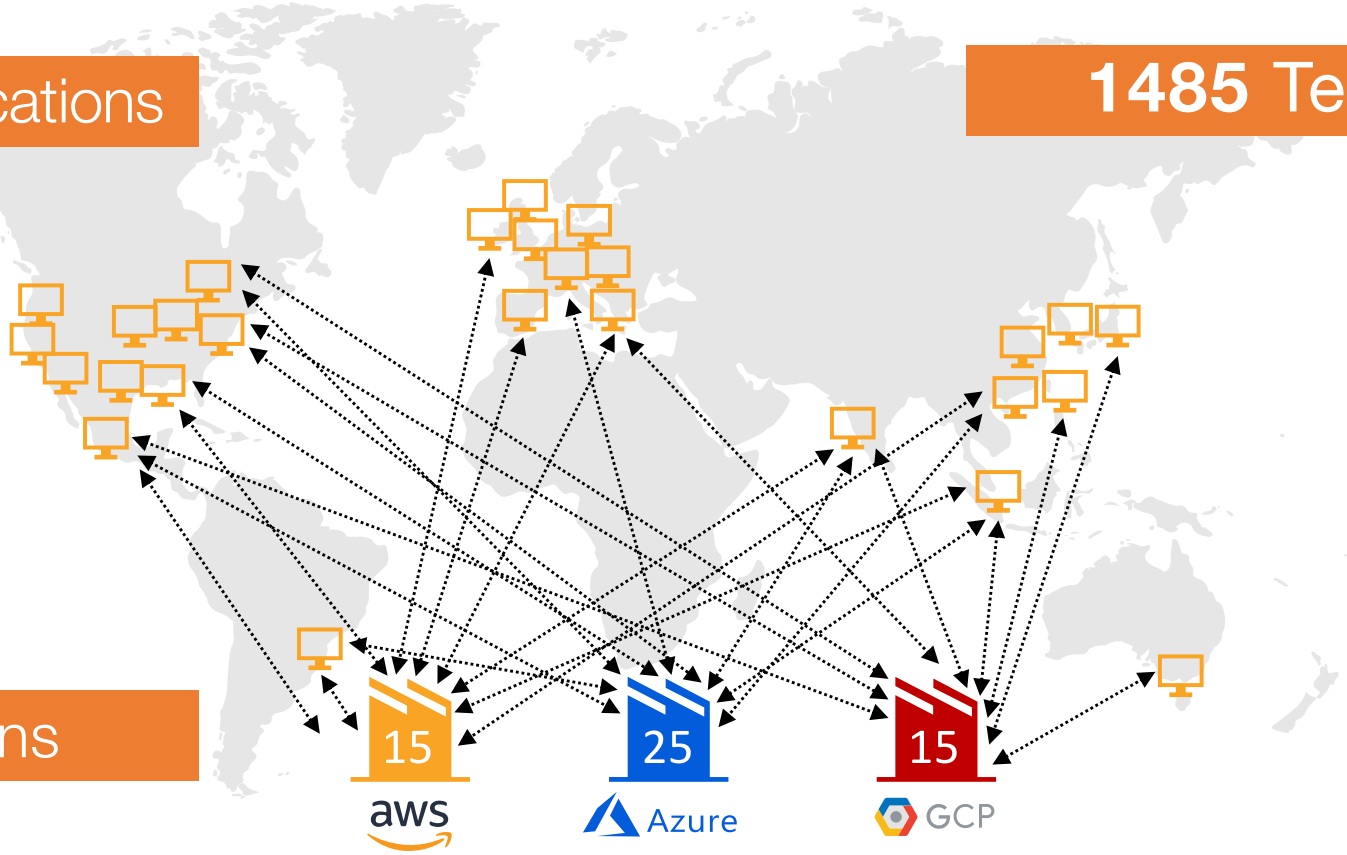

55 Regions

Inter-AZ Measurements

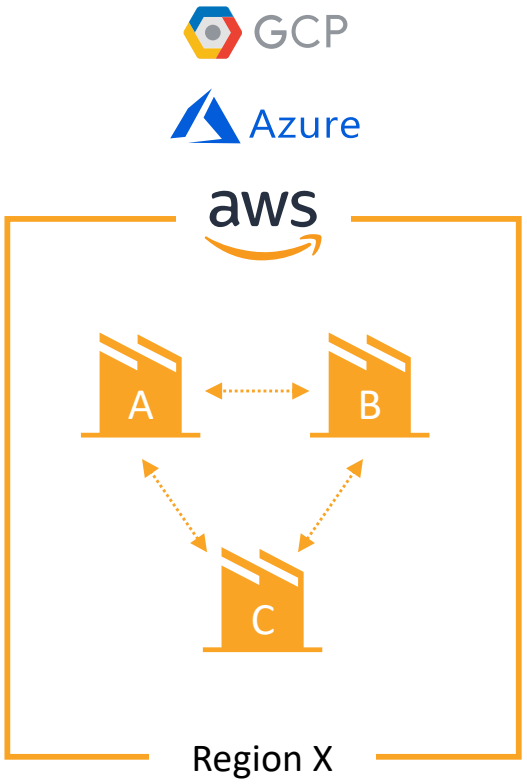

- Inter-AZ tests are measured within a single cloud provider
- Availability zones are mapped independently for each account. Data samples from multiple AZ pairs analyzed to discount for exceptions.
- AWS : 4 regions
- Azure : 1 region (Concept of AZ relatively new for Microsoft)
- GCP : 4 regions

Inter-Region Measurements

Google Cloud Platform

15 REGIONS

Multi-Cloud Measurements

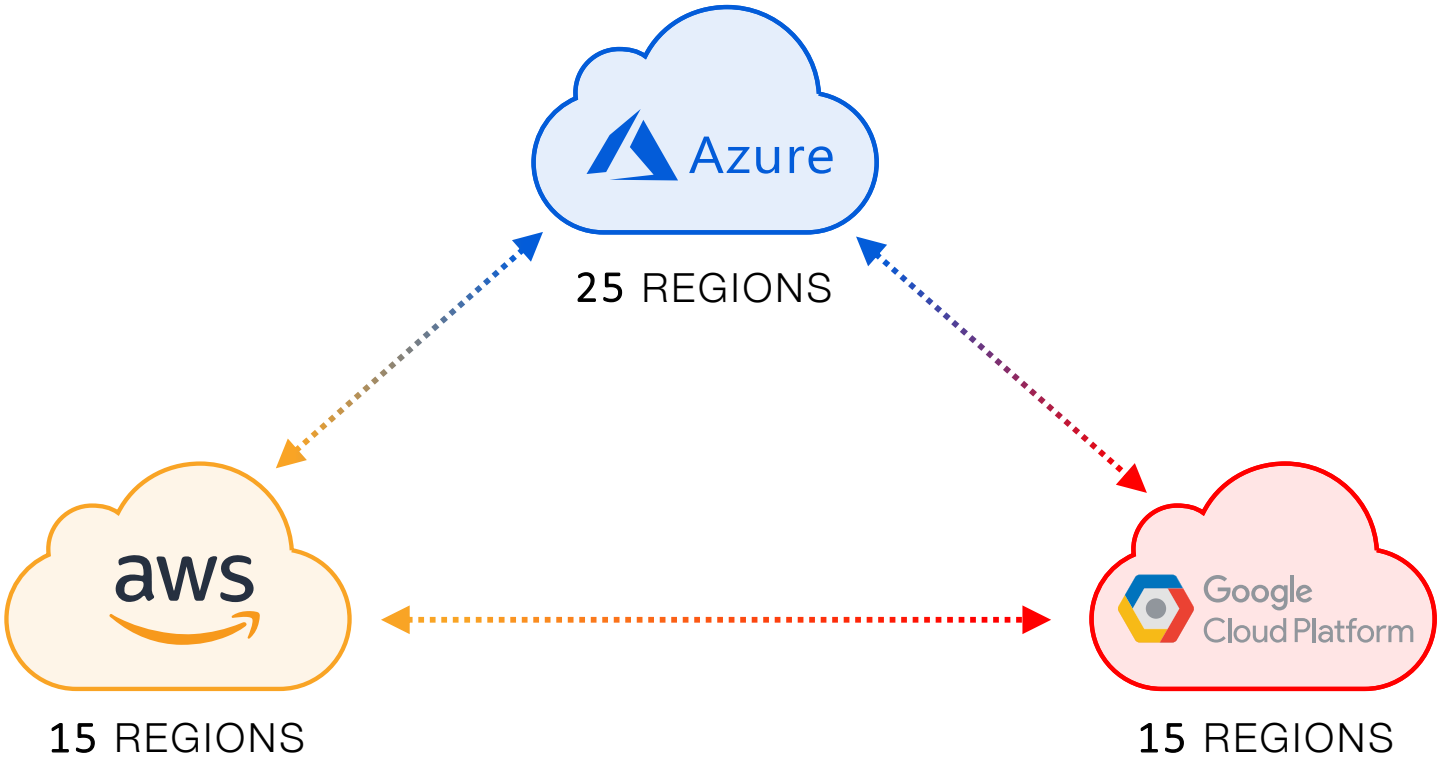

Data Points Gathered

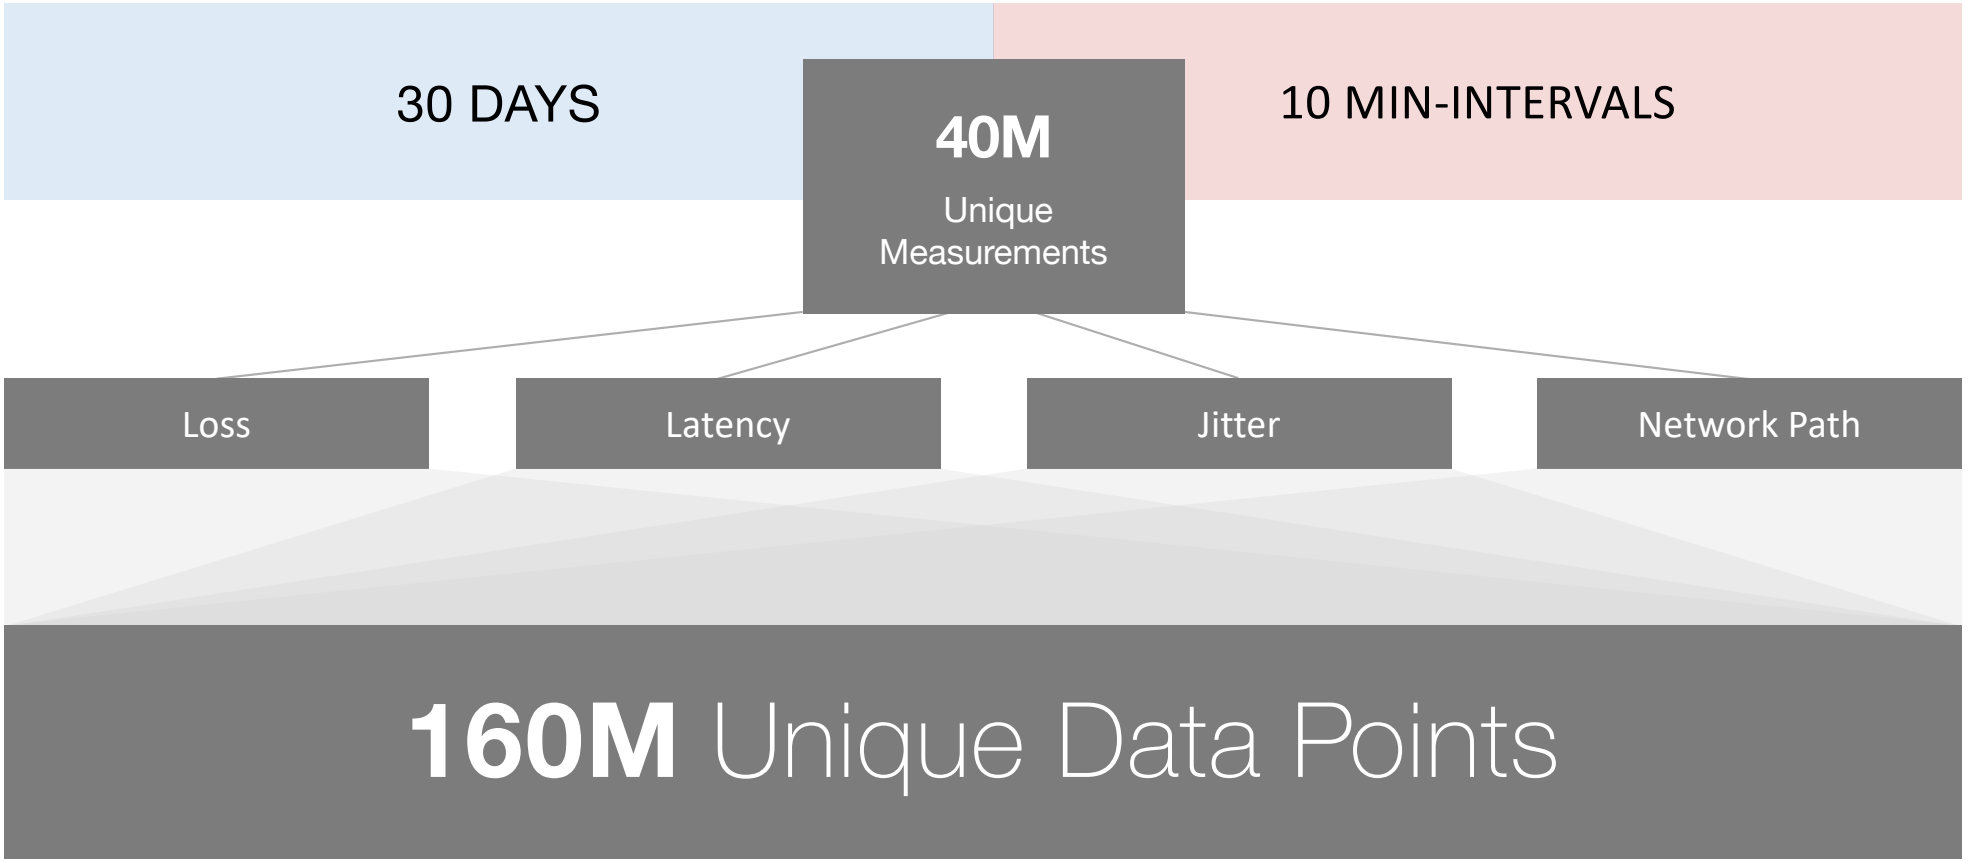


GOOD
BAD
UGLY

Inter-AZ Performance

EXPECTATION

< 2 ms

aws

0.82 ms

us-east-1	0.92 ms
ap-south-1	0.72 ms
eu-west-2	0.61 ms
sa-east-1	1.13 ms

Azure

1.05 ms

centralus	1.05 ms

Google Cloud Platform

0.79 ms

asia-south-1	0.92 ms
asia-southeast1	0.93 ms
southamerica-east1	0.60 ms
europe-west-2	0.84 ms

Inter-AZ performance is reliable and consistent

Indicates robust regional backbone for redundant multi-AZ architectures

Inter-Region Performance

- Inter-region network connectivity stays “within” provider network
- Performance baselined with Internet averages to reflect relative performance
- Most region pairs perform well but exceptions exist

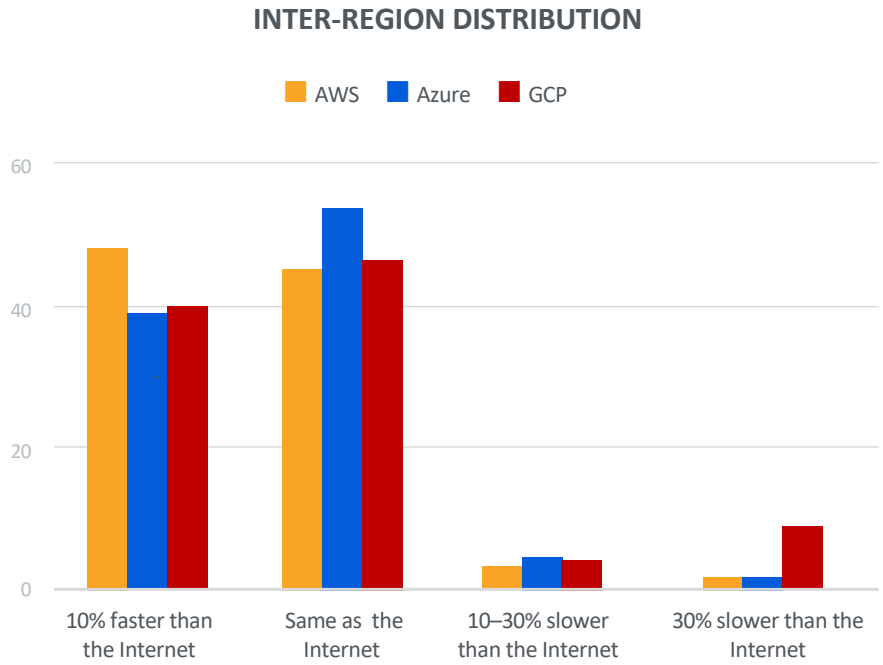

How to Choose Inter-Region Pairs?

Region Pair	Bi-directional Latencies (ms) from Sydney, Australia (Primary Region)		
	AWS	**Azure**	**GCP**
Tokyo	109.59	107.32	104.34
Singapore	174.65	108.36	168.49
Mumbai	228.48	162.15	228.92

10% faster than Baseline Baseline 10%–30% slower than Baseline 30% slower than Baseline

Where is End User Performance Strong?

- Comparable performance in North America and Western Europe
- No significant packet loss or outages in a 4-week period

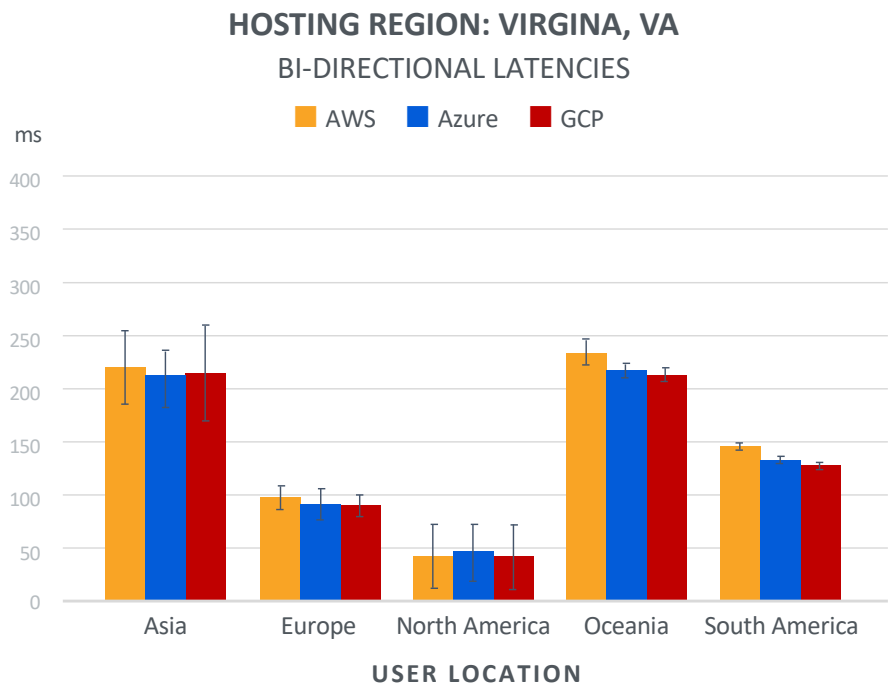

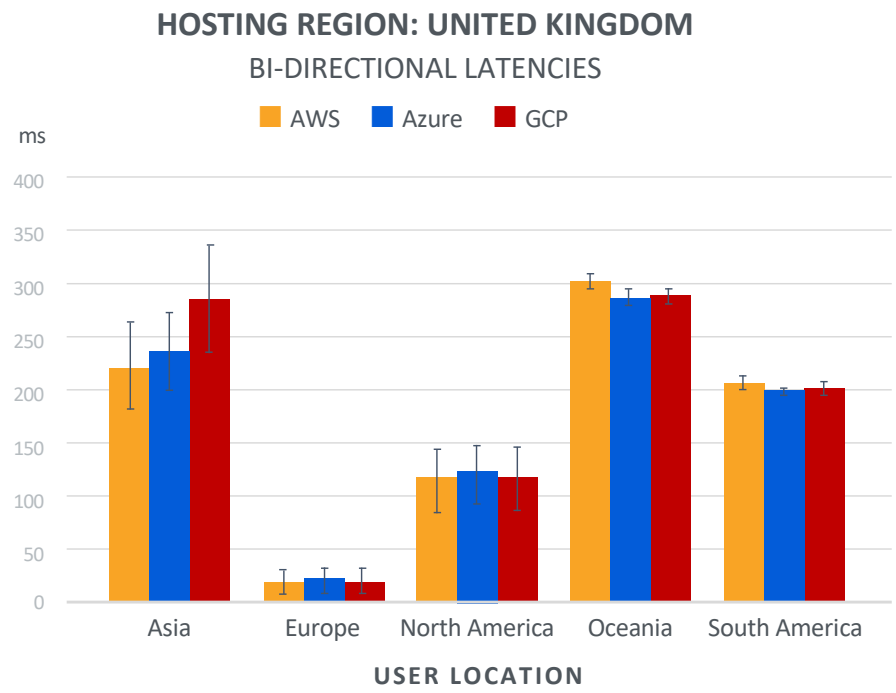

Exception: GCP is 3X Slower from Europe to India

- GCP exhibited 3x the network latency in comparison to AWS and Azure
- Circuitous route from Europe via North America to India on the GCP backbone

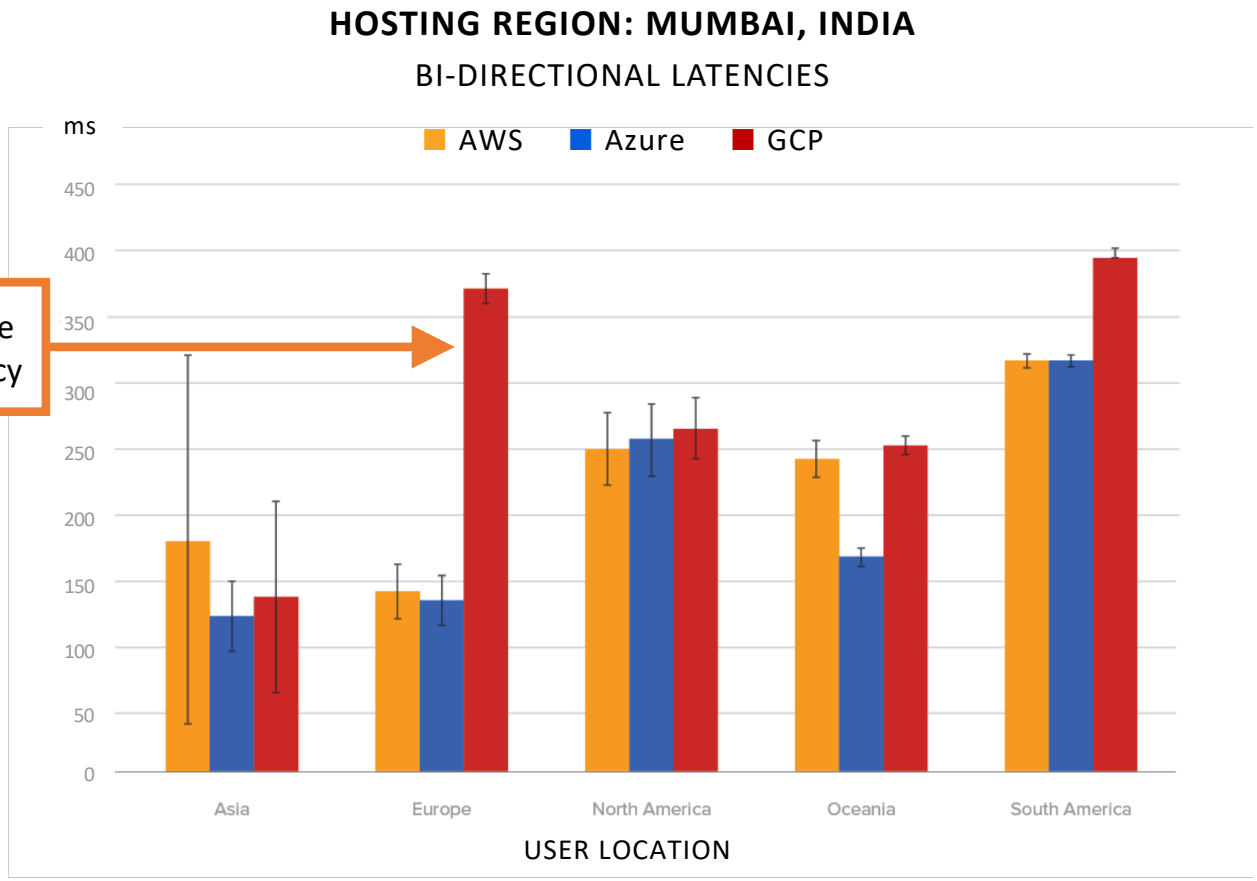

Why is GCP Slower to Mumbai?

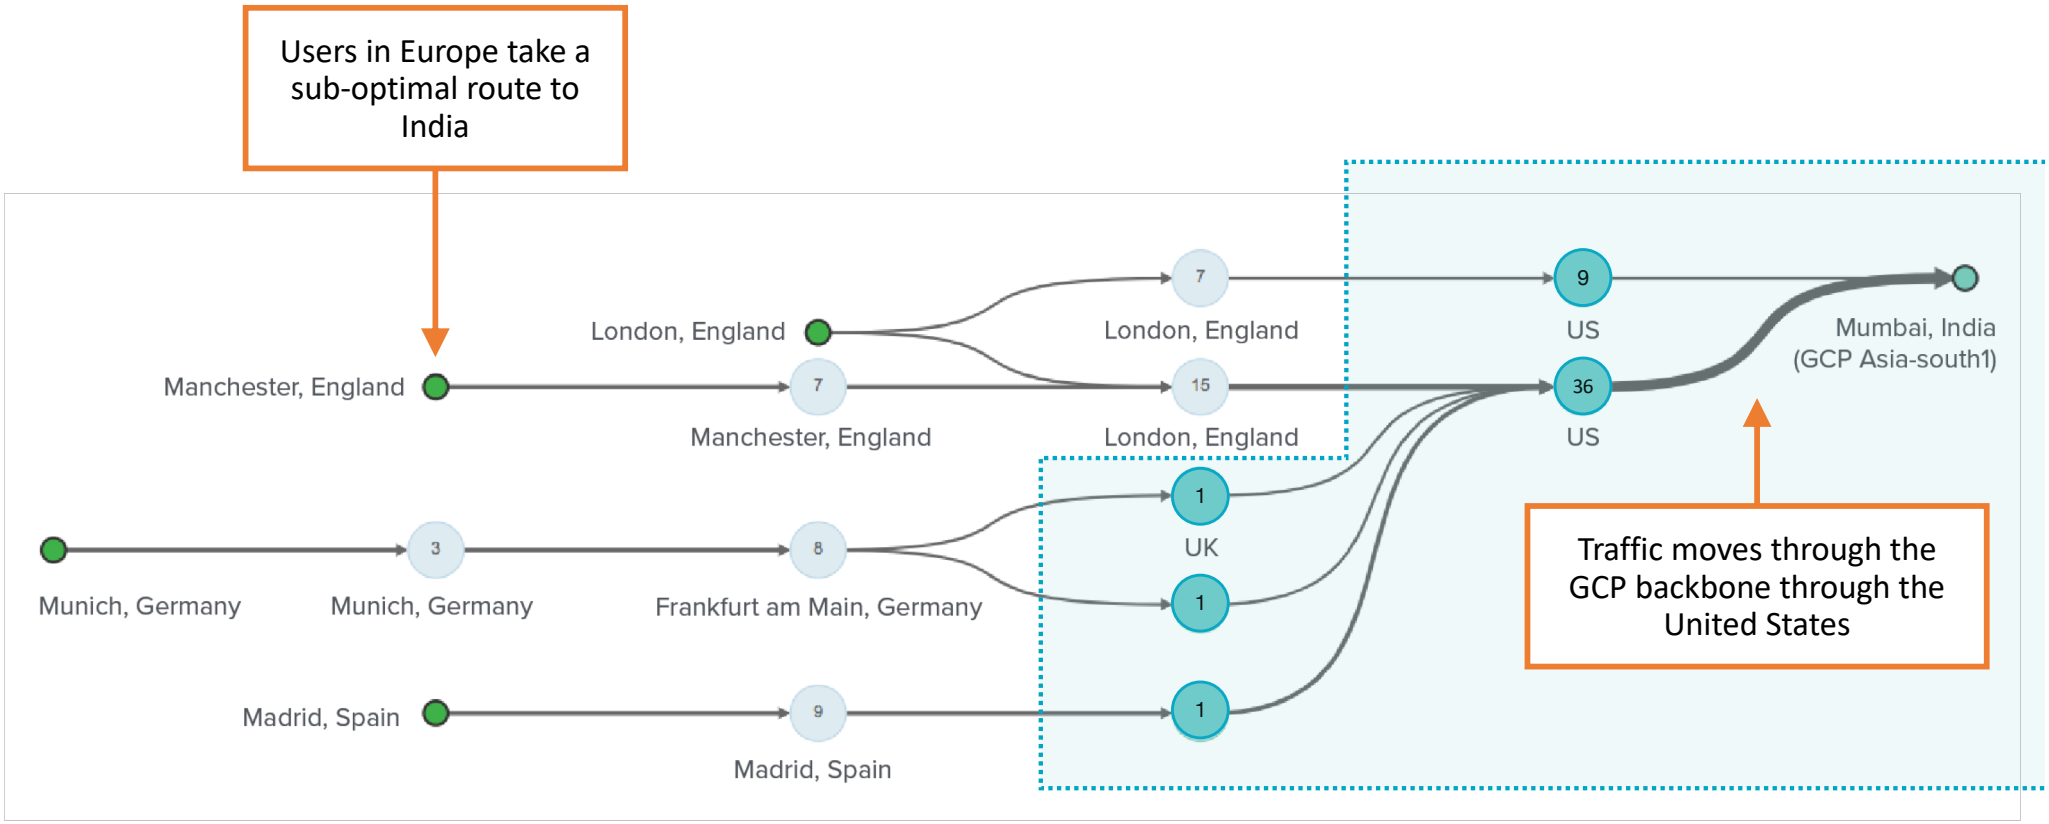

Why is GCP 3x Slower to Mumbai?

Exception: AWS Performance Anomaly in Mumbai

AWS exhibits the largest standard deviation in network latency in Asia

Considerably higher latency variation with AWS

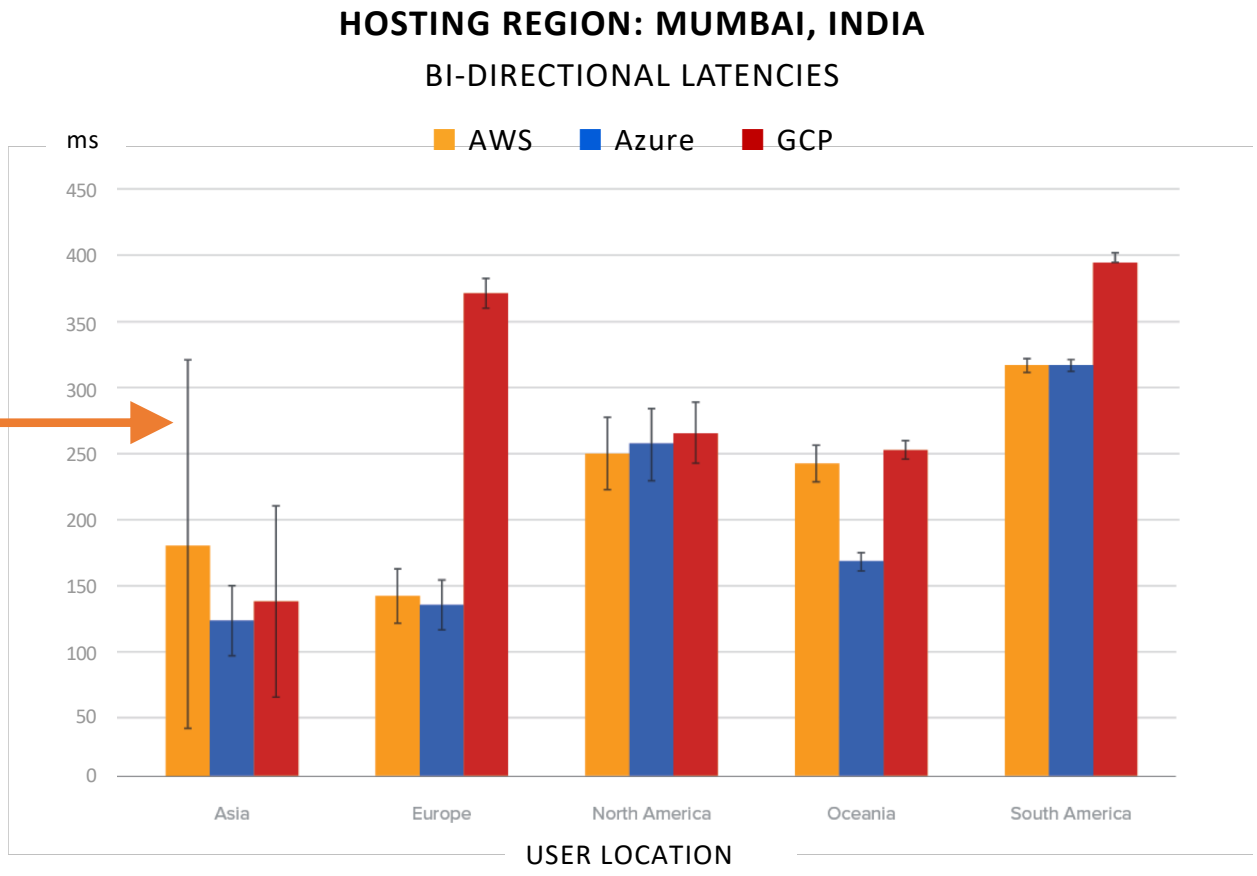

AWS in Asia is Less Predictable

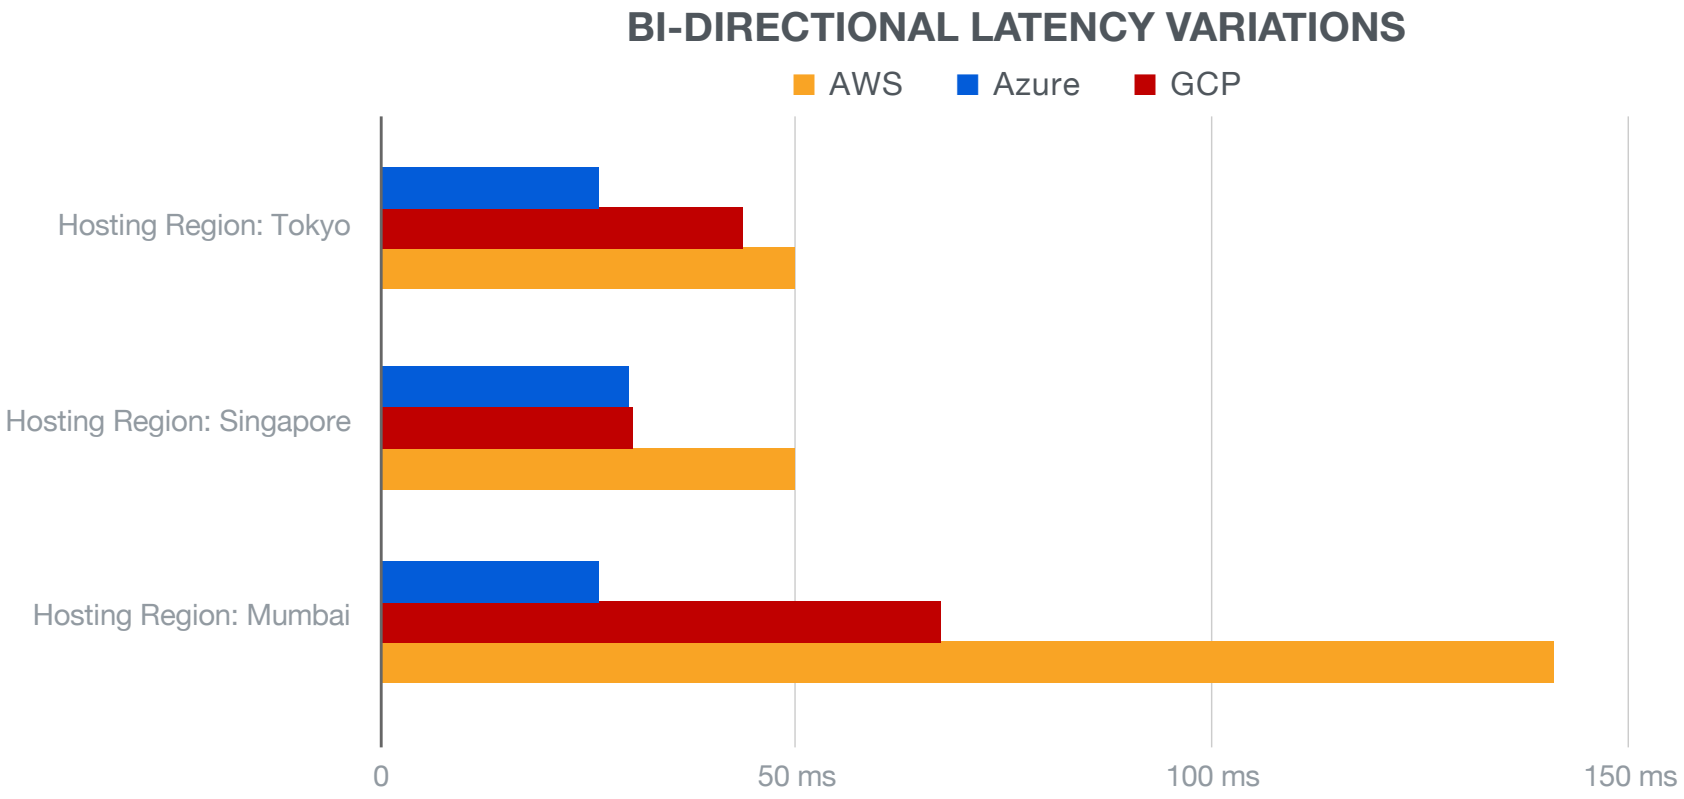

Why AWS Exhibits Less Predictability

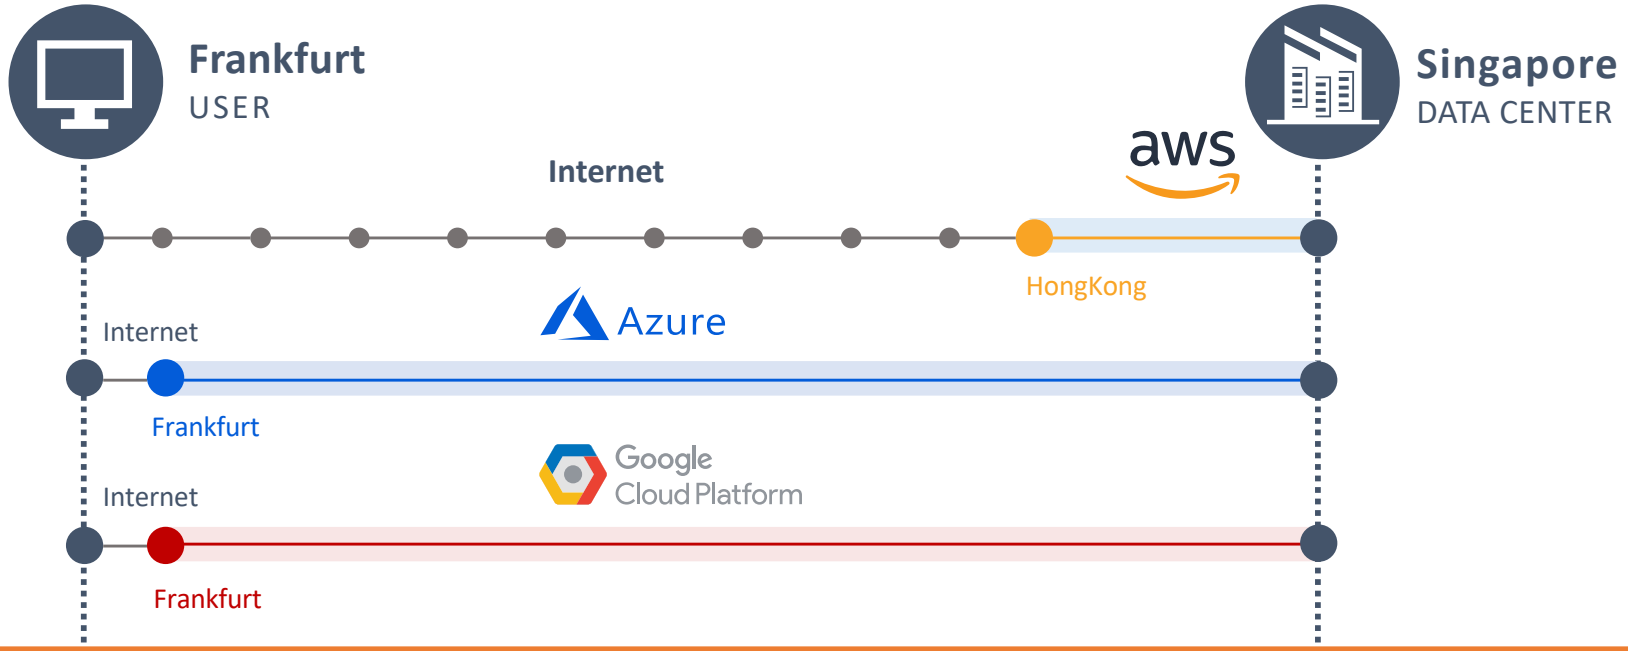

Not all clouds are created equal

Architectural and connectivity differences have an impact on performance

AWS Loves the Internet

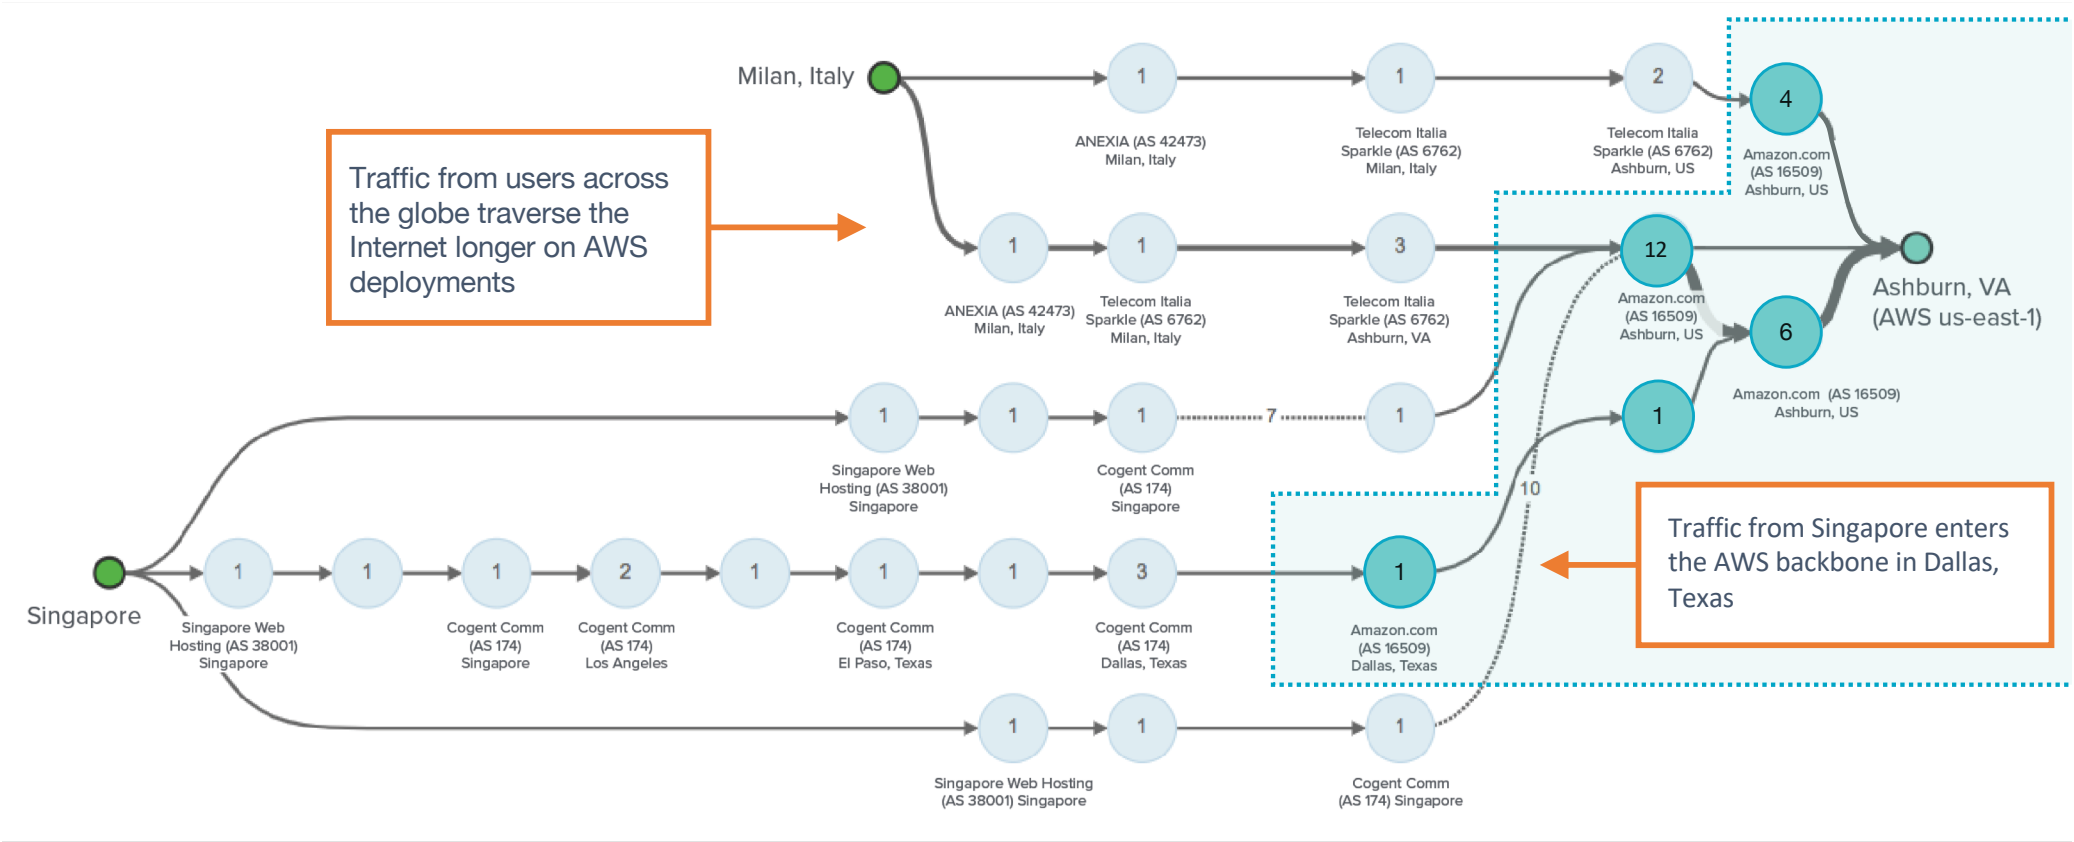

Azure and GCP Network Differently

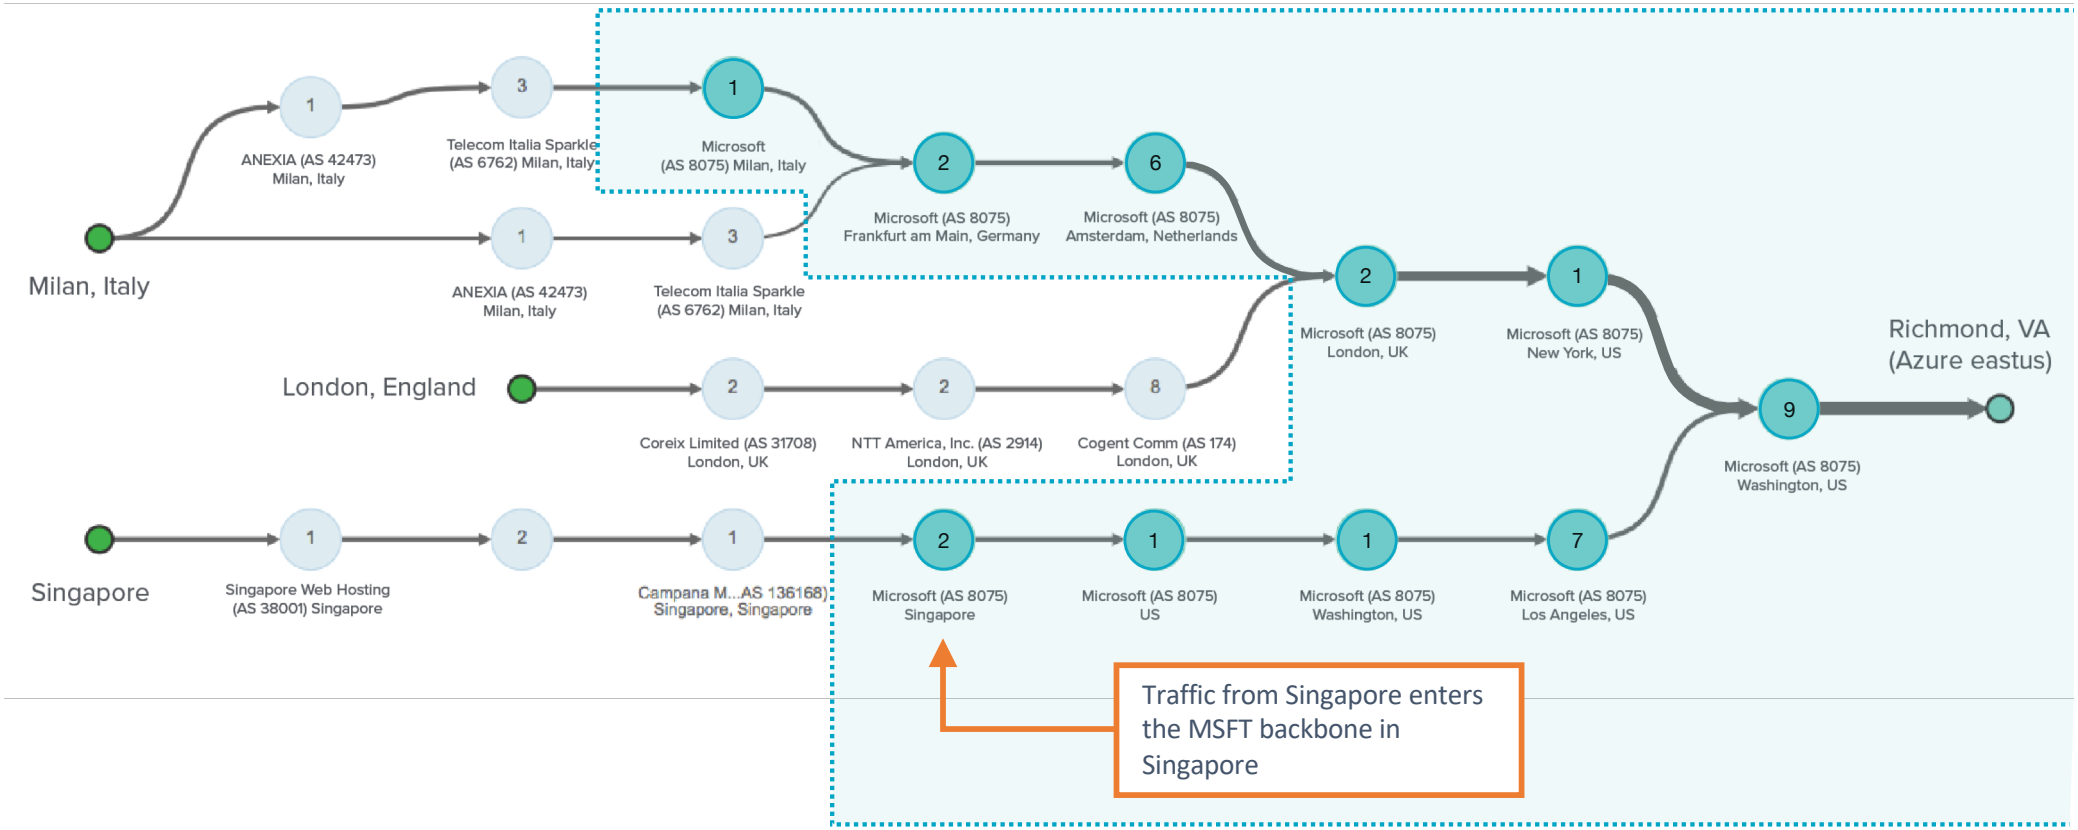

AWS Address Performance Instability With New Service

- AWS Global Accelerator for improved performance
- Pay AWS more $$ to ride the AWS backbone instead of the Internet
- Cloud Trend Alert: Monetization of the backbone

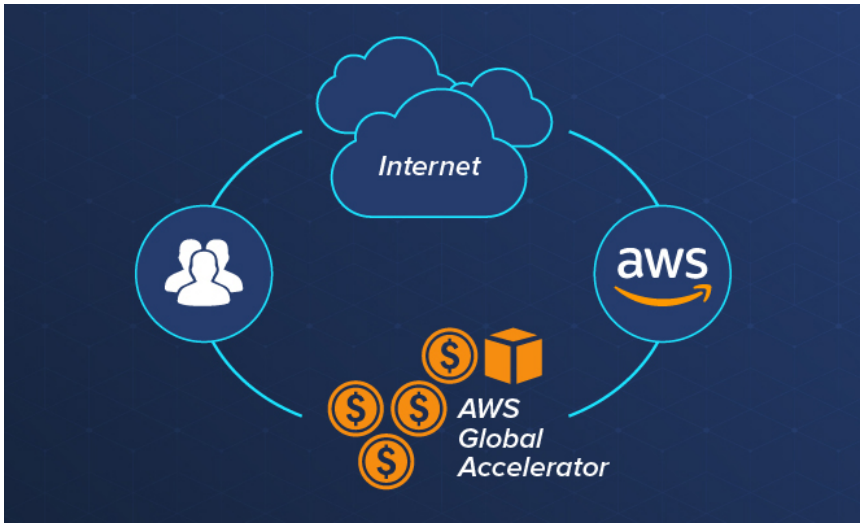

Multi-Cloud Performance

- AWS, Azure and GCP directly peer with each other
- Negligible packet loss and jitter
- Traffic between clouds does not leave their backbones

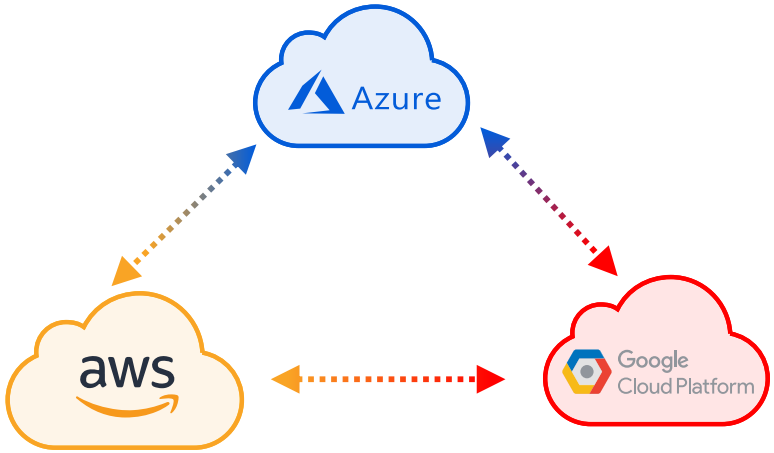

Metric	AWS - Azure	Azure - GCP	GCP – AWS
Jitter	0.43ms	0.29ms	0.50ms
Packet Loss	0.01%	0.01%	0.01%

Summary of Learnings

NOT ALL CLOUDS ARE CREATED EQUAL
Architectural and connectivity differences have an impact on performance

AWS has less performance stability in Asia

GEOGRAPHICAL PERFORMANCE VARIES
Do not assume uniform performance across the globe and use data to test, validate and optimize

GCP is 3x slower from Europe to India

INTRA-CLOUD PERFORMANCE IS STELLAR
Robust internal backbone supports redundant and distributed application architectures

Inter-AZ performance between 0.5 ms – 1.5 ms

MULTI-CLOUD IS SAFE
Symbiotic coexistence between the Big 3 and stable multi-cloud performance

AWS, Azure and GCP peer directly with each other

Recommendations

- Use data, not your gut to guide your cloud investment decisions
- Find your “best” and ”worst” region pairs before deciding on your cloud architecture
- Consider your organization’s tolerance to Internet exposure and evaluate risk vs $$
- Factor in inter-AZ and inter-region performance. Your app performance and user-experience relies on it
- Trust, but verify. Avoid assumptions in the cloud

Our Research Vision

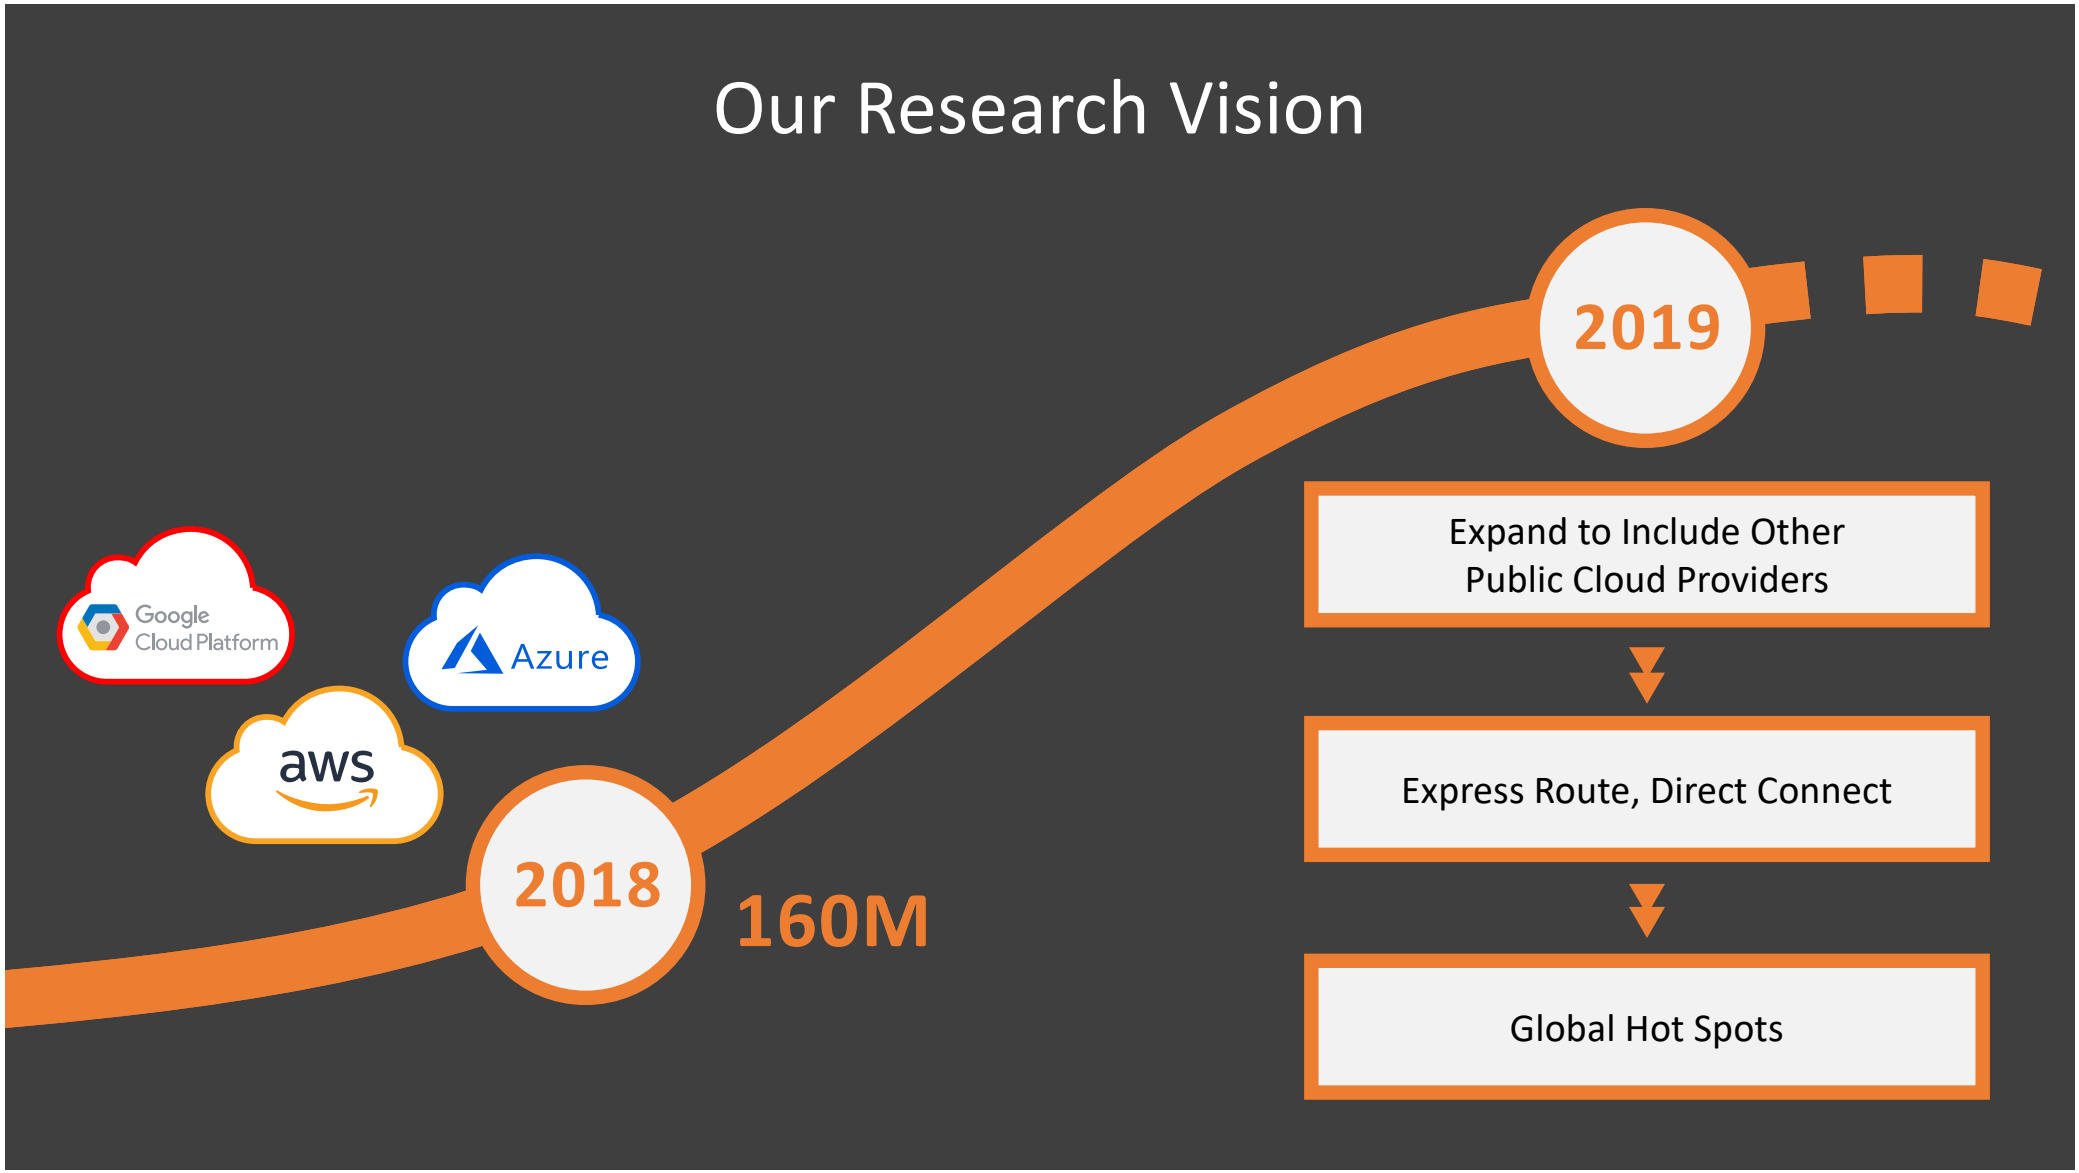

Thank You

@archana_k7

For more PCPBR information, please go to
https://www.thousandeyes.com/research/public-cloud