ThousandEyes

Thrive in a connected world

eBook

Three Key Performance Factors for Public Cloud Decisions

Data-Driven Cloud Investments for AWS, Azure and GCP

What's the Right Cloud For Your Business?

Enterprises are on the path to digital disruption and cloud computing plays a pivotal role in the digital transformation journey. Public cloud adoption continues to rise, a trend visible across all cloud providers, however, three of them visibly stand out. The three big cloud service providers that dominate the IaaS market are Amazon Web Services (AWS), Microsoft Azure (Azure) and Google Cloud Platform (GCP). Forrester predicts that the Big 3 are well positioned to capture 80% of the total global public cloud revenue by 2020. Even until a few years ago, the necessity of moving to the cloud was questioned and widely debated. Today, the discussion is not about "Should we move our business to the cloud?"
Instead, discussion centers around "Which cloud provider makes the most sense for my business?"

Choosing the right cloud provider is becoming an increasingly nuanced consideration that goes far beyond market share. Enterprises making cloud choices lean heavily on comparative studies predominantly focussed on global data center presence, service catalogs and pricing tiers. While these are all valid selection parameters influencing vendor selection, enterprises often overlook a critical comparative component—network performance. Studies comparing the performance of public cloud providers are rare and the few that exist lack comprehensive scope.

How can network performance be used to make cloud decisions?

Network Performance Factors That Impact Cloud Decisions

Often overlooked, the network is the glue that binds all cloud communication—from the end-user to the cloud and within the cloud. Calibrating network performance across public cloud providers provides insightful information to help guide your cloud vendor selection and shape your cloud strategy. In the cloud, three network performance factors provide a complete view of cloud performance.

User Experience Performance

Understanding end-user performance provides a direct lens into how customers experience your business services. If you are a global enterprise, characterizing regional service delivery allows you to baseline performance while crafting a cloud architecture best suited for your enterprise.

Intra-Cloud Performance

Modern applications are distributed in nature and deployed across multiple cloud regions and availability zones (AZ). Every interaction across micro-service components of an application impacts performance, making inter-region and inter-AZ network metrics critical.

Multi-Cloud Performance

Multi-cloud is a reality in modern enterprise architectures—whether by design or by chance. Multi-cloud is a practical approach to avoid vendor lock-in while leveraging best-of-breed options. Understanding multi-cloud connectivity architecture and its influence on application performance should be as critical as end-user performance.

A Data-Driven Approach to Cloud Performance

To understand public cloud performance holistically ThousandEyes leveraged its fleet of software agents in global cities and within public cloud providers to measure and compare the performance of the Big 3—AWS, Azure and GCP. Network performance metrics such as loss, latency, jitter and connectivity were periodically measured over four weeks across 55 regions of all the three providers. Insights gathered by analyzing 160 million unique data points culminated in the 2018 Public Cloud Performance Benchmark Report.

The 2018 Public Cloud Performance Benchmark Report provides a unique, unbiased third-party and metric-based perspective on public cloud performance as it relates to both user experience and back-end application architecture. For enterprises embracing digital transformation and venturing into the public cloud, the metrics and insights gathered serve as a data-driven guide to best practices and decision-making for operating in the cloud.

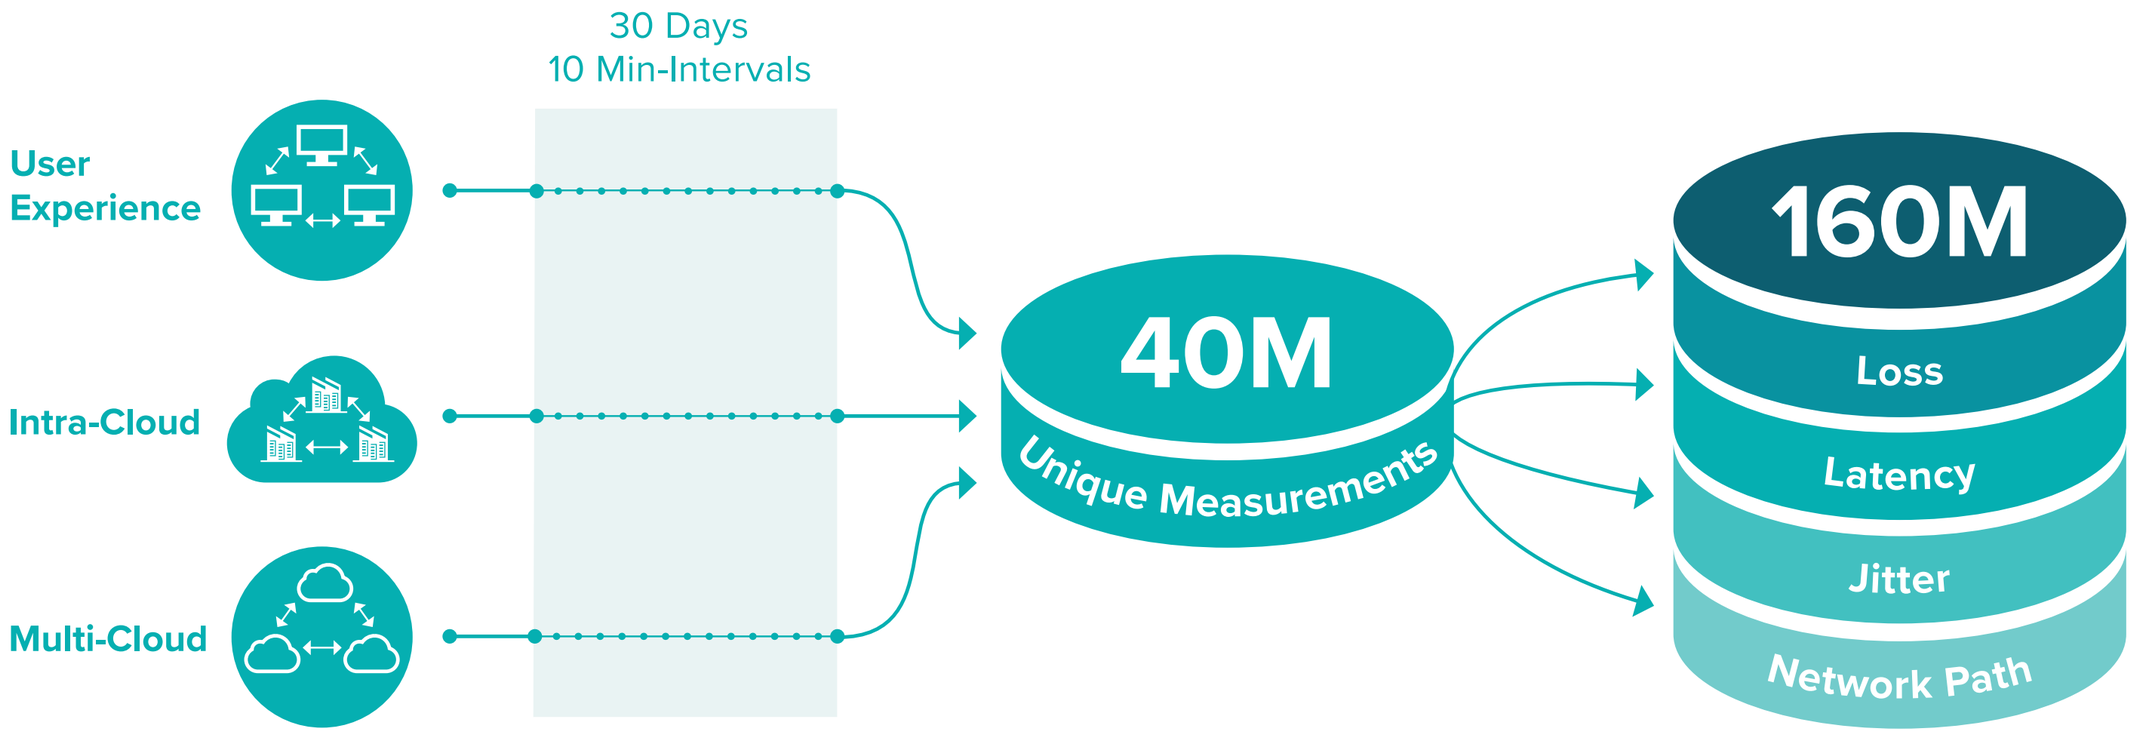

Figure 1: *Data-Driven Decision Framework*

Cloud Connectivity Architecture Impacts **End-User Performance**

Did you know that AWS deployments are subject to lower performance predictability?

AWS deployments have an increased reliance on the Internet—a best effort network, and are subject to greater operational challenges and risks. AWS' network design forces traffic from the end-user through the public Internet, while Azure and GCP rely on their internal backbones to transport traffic.

The Internet is a best effort medium, a constellation of ISP networks that is vulnerable to security threats, DDoS attacks and ISP outages. Relying on the Internet increases unpredictability in performance, creates risk for cloud investments and raises operational complexity. For instance, AWS connectivity architecture results in poor performance predictability, especially in parts of the world that are known to have a less reliable Internet, such as Asia.

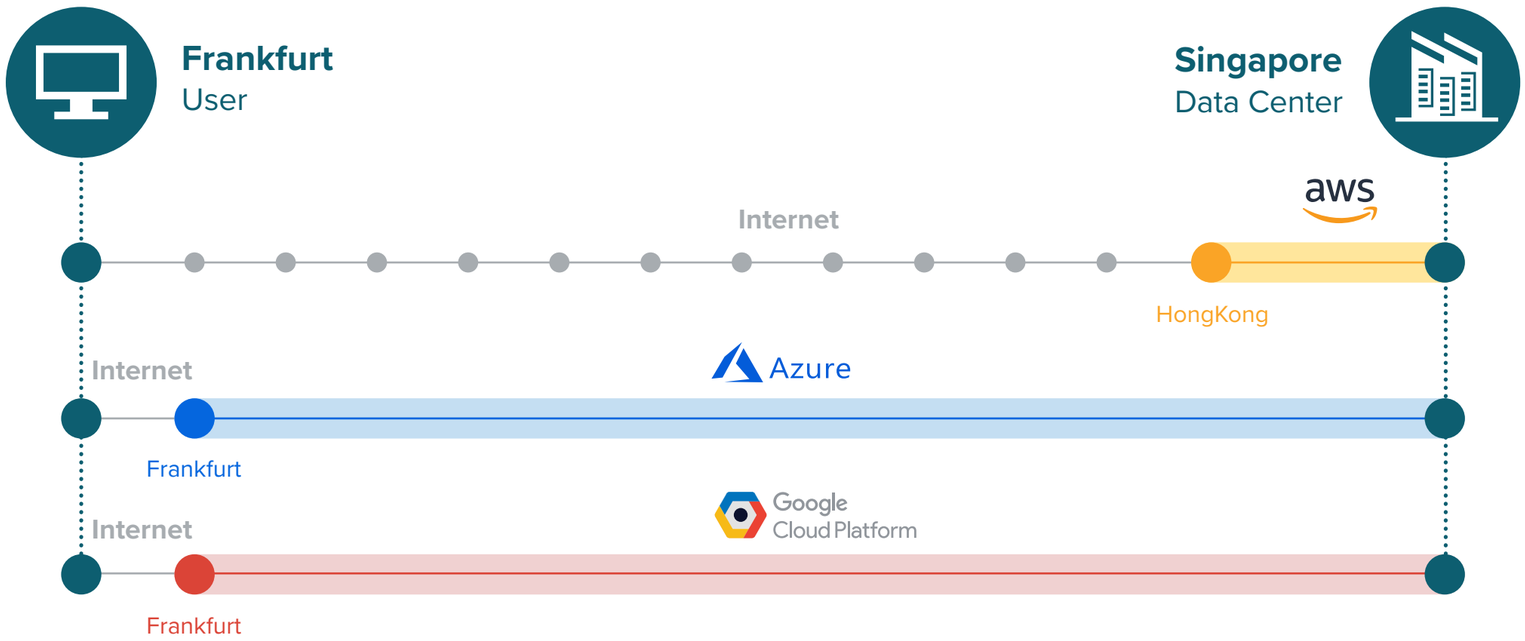

Figure 2: *Connectivity Architectures of the Big 3*

In Asia, **AWS** demonstrated

35% less

network performance stability than **GCP** and

56%

less than **Azure.**

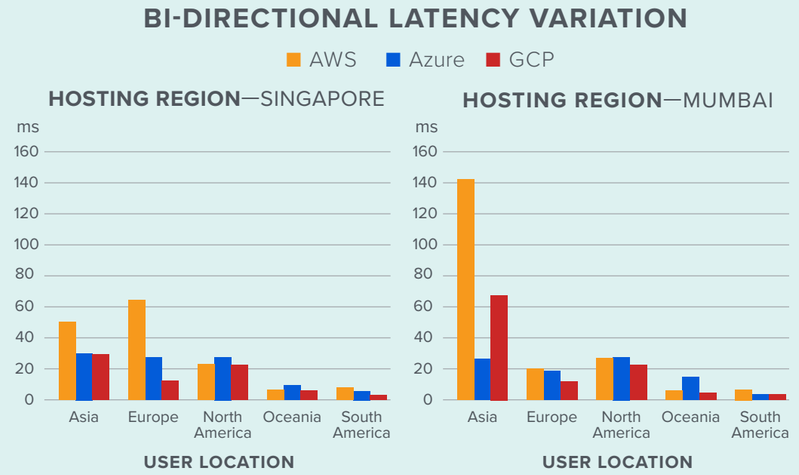

Intra-Cloud Performance:
Your Guide to Picking Cloud Regions

Inter-region performance is critical for enterprises adopting a tiered multi-region architecture. A common practice across cloud architectures is to distribute compute or workloads across multiple global regions but centralize common functions and services such as storage or databases in a single region. Tiered architectures are often essential, but geographical expanse can incrementally affect network latencies and thus impact end-user experience. What if you could use data to make these cloud architectural decisions?

Inter-region latency is expected to be better than regional Internet averages, but quantifying and baselining measurements can give you an understanding of your best and worst regions. If your primary hosting region is Australia, Singapore and Mumbai are not the best options for a secondary redundancy regional pair on AWS and GCP.

Region Pair	Bi-directional Latencies (ms) from Sydney, Australia (Primary Region)		
	AWS	Azure	GCP
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

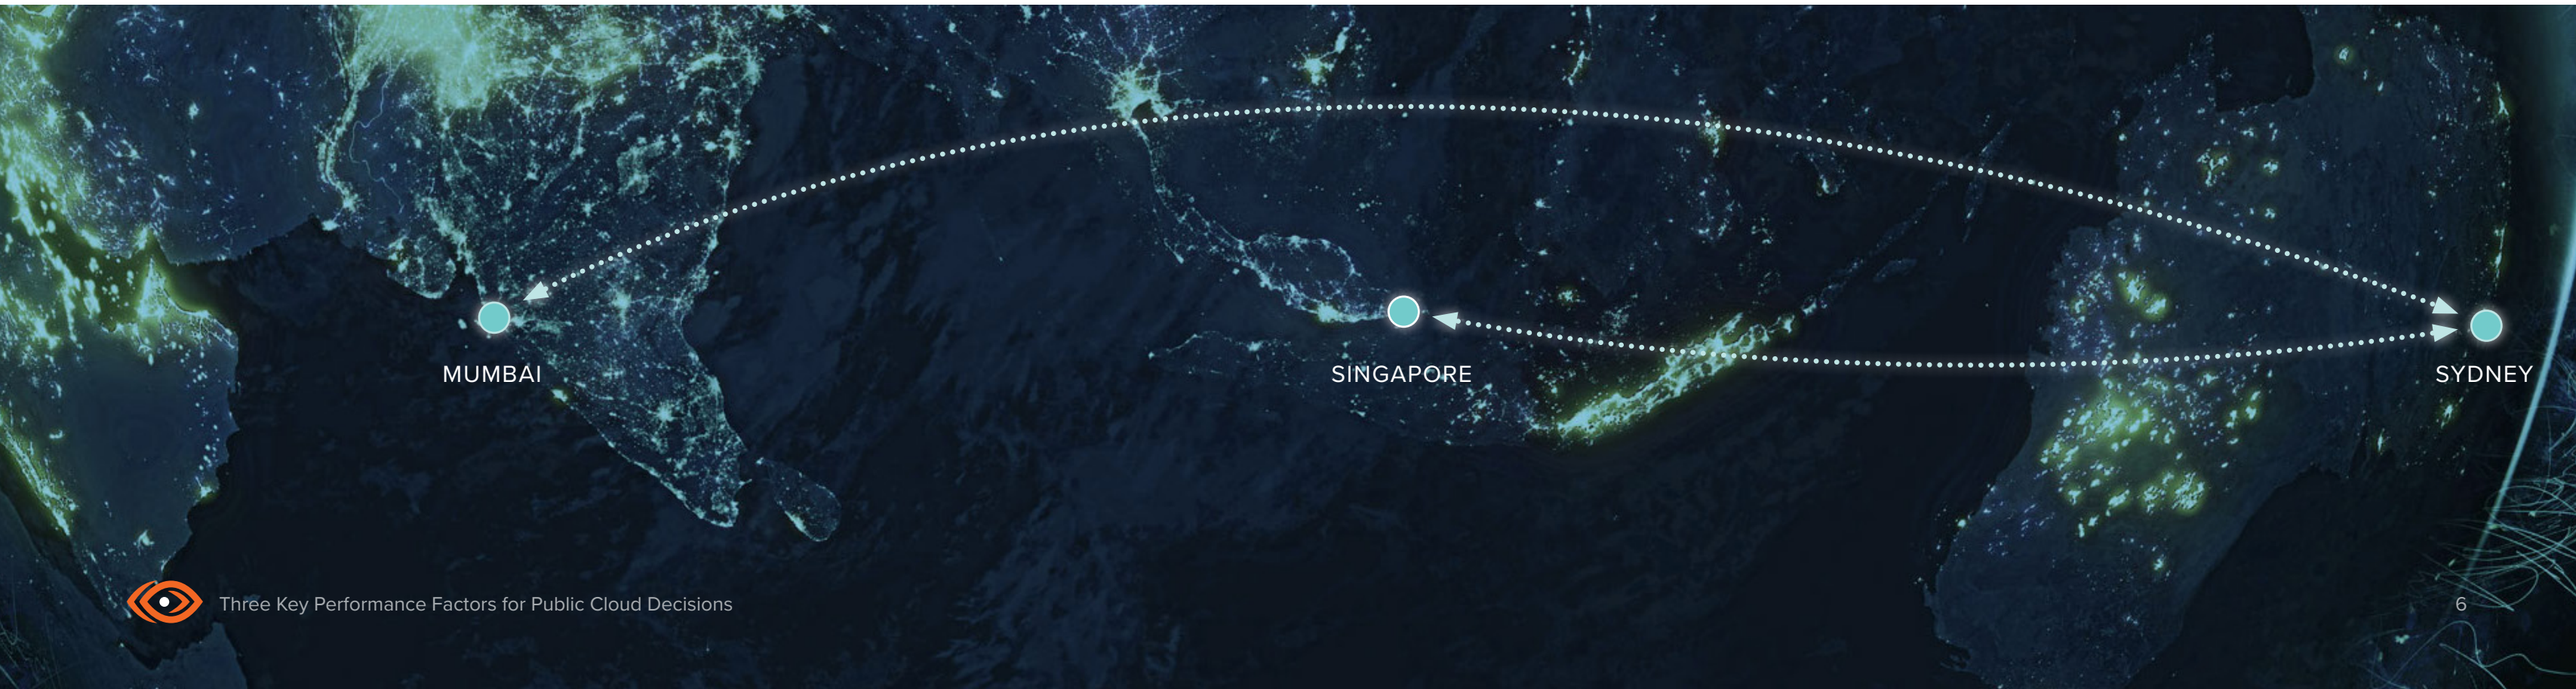

Multi-Cloud:
The Cure to Performance Variations between AWS, GCP and Azure

Multi-Cloud is quickly becoming the default cloud architecture template—by choice or by chance. Enterprises consider multi-cloud strategies to reduce vendor lock-in, mitigate risk, leverage best-of-breed of cloud provider innovation. In many organizations, different teams select different cloud providers only to recognize later an intersection resulting in multi-cloud communication.

Network performance has not been a traditional metric to consider in the formulation of a multi-cloud strategy. However, global performance variations presented above reinforce the case for multi-cloud. Enterprises lack the ability to influence cloud provider architectures that cause performance variations. Multi-cloud is an intriguing option for enterprises to circumvent these performance inconsistencies, making it imperative to understand if multi-cloud deployments are ready for prime time.

Research data indicates a symbiotic network relationship between AWS, Azure and GCP. Traffic between cloud providers almost never exits the three provider backbone networks, leading to negligible loss and jitter in end-to-end communication between providers.

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

> “Multi-Cloud is no longer a matter of **IF**, it’s a matter of **WHEN**”

Conclusion

Ultimately, it is imperative for enterprise IT leaders to understand that cloud architectures are complex, and not rely on network performance and connectivity assumptions or instincts while designing them. Enterprises relying heavily on the public cloud or considering a move to the cloud must arm themselves with sound performance data on an ongoing basis to guide them through both planning and operational stages. Every organization is different—enterprise cloud connectivity architectures are highly customized and hence the results must be viewed through the lens of one's own business in choosing providers, regions and connectivity approaches.

Download the 2018 Public Cloud Performance Benchmark Report for detailed insights and performance comparison of AWS, Microsoft Azure and GCP.

ThousandEyes

201 Mission Street, Suite 1700
San Francisco, CA 94105
(415) 231-5674

www.thousandeyes.com

About ThousandEyes

ThousandEyes delivers Network Intelligence—immediate visibility into experience for every user and application delivered over any network, so companies can deliver superior digital experiences, modernize their enterprise WAN, and successfully migrate to the cloud.