TECHNOLOGY PROVIDER COLLABOARATON

IBM AND EQUINIX PARTNER TO ENABLE APPLICATIONS AT THE EDGE

Interconnected bare metal to power latency-sensitive workloads with IBM Cloud Paks® and Equinix metro-edge as part of a high performing hybrid multi-cloud strategy

Summary

IBM and Equinix extend cloud computing to the edge, with autonomous management capabilities that address new challenges of massive scale, variability and rate of change—in enterprises and telcos. Clients can build and deploy new containerized applications in Equinix locations knowing that they will not be locked in but rather enabled to optimize these deployments to meet current and evolving future needs—whether on metro edge, public clouds or enterprise on-premise edge.

So what is the metro-edge?

The metro-edge is the aggregation point in strategic metropolitan cities around the world that exists in close proximity to the customer premise or in proximity to the endpoints generating data. These locations are high performing, high density, connected facilities with options to network service providers, cloud service providers and more that are parts of a hybrid multi-cloud architecture.

Why the edge matters to you

Most enterprises use multiple cloud service and network service providers as part of their current and future IT landscapes. Running high performing distributed applications and data at the intersection point of multiple service providers and in close proximity to their locations helps to drive innovation and aid the customer on their digital transformation journey. Equinix is optimally positioned as a critical partner having more metro edge locations than any other provider globally. In addition, Equinix metro edge locations address more of the global population than any other interconnected datacenter.

Enterprise Edge-in

Viewing the solution from the perspective of the edge, starting with the remote on-premises edge, then to the network/metro edge, and then to the hybrid/multi-cloud. In terms of movement, edge-in is movement of workload and data from the remote on-premises tiers, to the network/metro edge, and then to the hybrid/multi-cloud. Using IBM Edge Application Manager enables enterprises to optimize what is deployed on the very edge with ability to place data and applications in optimal locations to enhance latency and workload execution performance requirements.

Cloud-Out

Viewing the solution from the perspective of the cloud, starting with the public and/or private cloud, then to the network/metro edge, and then to the remote on-premises edge. In terms of movement, cloud-out is movement of workload and data from the hybrid/multi-cloud tiers, to the network/metro edge, and then to the remote on-premises edge. This enables enterprises to optimize what is deployed across public and private clouds for cost optimization and to move more latency and data privacy related workloads onto country-based metro edge and enterprise on-premise edge locations.

Combining IBM and Equinix gives clients a vision of a fully automated customer experience with provisioning based upon the client's performance requirements.

EQUINIX®
WHERE OPPORTUNITY CONNECTS

Competitive advantage will ultimately depend on your digital infrastructure

50zb
of Data in Real Time by 2025

Data Explosion

Data will continue to grow exponentially, with 30% of global data transmitting in real time by 2025 (50 zettabytes).

8x
reduction in TCP throughput

Distance Kills

Distance kills throughput and contributes to unsustainable ROI. For example, adding 30 ms of latency reduces throughput on a TCP network 8x. Adding2% packet loss to latency reduces overall throughput 25x.

$100k
in transport costs per year

Transporting Cost

Transporting data compounds costs for everyone. Moving 100 terabytes, for example, can total ~$100k/yr. in cloud egress and inter-region transport costs (before network costs).

80%
cost reduction

Future-proofing

Optimizing hybrid cloud for performance vs. optimizing for flexibility can reduce cost by 80%.

Equinix was founded in 1998 with the idea of establishing neutral hubs, where competing networks could physically exchange traffic, quickly scale and maximize the performance of the internet.

From that point forward—we have become the platform for **global interconnection** and solution enablement

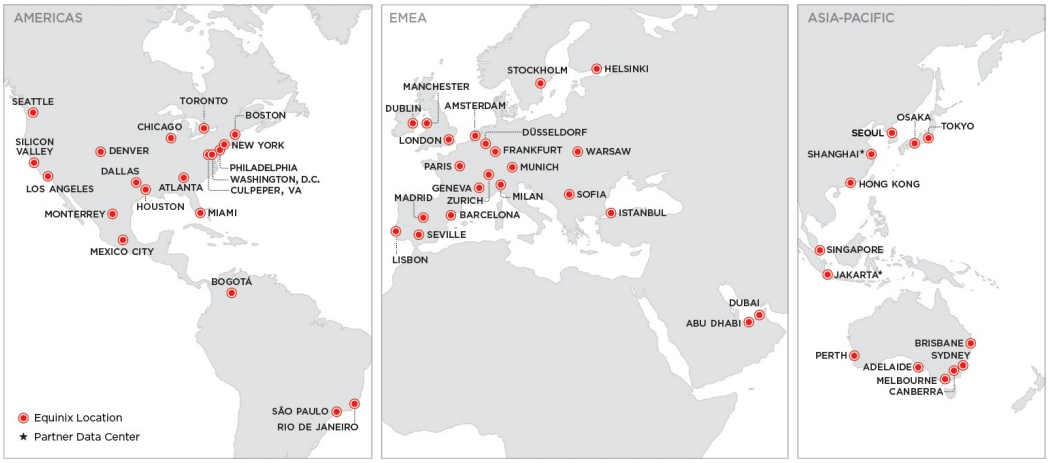

© 2020 Equinix, Inc.

Consistent Vision to Cultivate Ecosystems

1,800+	675+	1,250+	2,900+	8,000+	356,000
Networks	Content & Digital Media	Financial Services	Cloud and IT Services	Enterprises	Interconnections

We've helped multiple industry ecosystems exchange traffic at scale and are experts at helping companies enable global network and application performance. This began by working with the world's leading networks to build and connect their infrastructures inside Equinix data centers to grow the internet. It extended to content providers who wanted to become more proximate to their consumers.

In the years following, Equinix saw its financial ecosystem grow as the stock market moved from manual to electronic trading where the value of a millisecond is enormous. Today, Equinix is in the world's 20 leading financial markets, and the world's top banks do business with our company.

Cloud service providers then populated our ecosystem as the consumption of their cloud services exploded through the value of private interconnection. Today, Equinix data centers house 40% of all the world's public cloud on ramps. This led to the focus on the enterprise and the digital transformation journey as they seek to build interconnection-first architectures to overcome legacy IT constraints on digital transformation.

Interconnection helps you compete in a digital world.

When companies connect directly with each other via fiber or switched connectivity within a data center campus, they increase performance by reducing distance and reduce cost by eliminating networking expense. This is the power of interconnection.

Physics of proximity

Network traffic cost can be reduced by 60% when two parties connect directly versus connecting by internet or MPLS.

Economics of aggregation

A 30% reduction in latency can be achieved when two parties connect directly instead of connecting over a distance using internet or MPLS.

The Network Effect

The value of interconnection is increasing with the number of available counterparties, creating a network effect.

© 2020 Equinix, Inc.

Demand for enhanced AI architectures

It's best to host AI infrastructure at an interconnection hub where businesses can integrate with distributed IT systems via high-speed and secure networks.

Due to the data residency, compliance, performance requirements (of moving large datasets to far-off core clouds) and cost reasons (for backhauling large datasets to core clouds) mentioned, it becomes critical for businesses to place their AI compute infrastructure co-resident with their data, following the mantra of "train where your data lands."

Furthermore, AI platforms cannot exist in isolation. They need to be integrated via secure and high-speed networks to an enterprise's corporate IT systems, which can exist in private data centers and public IaaS and SaaS clouds. Thus, there is a need to connect AI systems with the rest of an enterprise's IT infrastructure.

Finally, AI platforms need to be situated where they can ingest data from multiple sources to fuel model prototyping and improve accuracy. In many instances these datasets reside in multiple public clouds, data brokers, private data centers and streaming from distributed devices. Thus, it is desirable to host the AI infrastructure at an interconnection hub that has high speed and secure access to these different data sources. This allows for the training of AI models at the interconnection hub enabling autonomous distribution of models to run on devices where data is created and where actions need to be taken. This creates an AI architecture that can continually train, learn and deploy analytics at the most optimal edge locations needed with full automation.

Use Cases for This Solution

AI at Equinix is specifically designed for use cases where:

- Enterprises want to do AI training or inference in their private data centers instead of in public clouds for control, cost, performance and privacy reasons. But many are finding it difficult to host AI in private data centers due to their inability to handle high power-density requirements, and the complexity of managing AI hardware and software infrastructure. Thus, enterprises want to access AI as a Service at a colocation data center.

- Datasets need to get processed at the edge instead of hauling the data back to a remote core data center, for cost, latency, and privacy and compliance reasons.

- AI applications need to integrate with enterprise IT systems or need to access external data from multiple sources such as clouds, private data centers, data brokers and edge locations. In these hybrid multicloud use cases, it's best to host the AI infrastructure at an interconnection hub where businesses can integrate with distributed IT systems via high-speed and secure networks.

- More than 9,800 enterprises and providers already have their infrastructure at Equinix and are interconnected to each other. We want to make it easy for both existing customers and new customers to do AI processing while leveraging data from this ecosystem.

© 2020 Equinix, Inc.

Let's Build Together

Interested to discuss further?
Email us at IBM@Equinix.com

Components of a Potential Solution

- AI at the metro edge: Perform AI at 200+ securely connected Equinix International Business Exchange™ (IBX®) data centers worldwide in 55 global metros.
- IBM Cloud Pak for Data: A fully integrated data and AI platform that modernizes how businesses collect, organize and analyze data and infuse AI throughout their organizations.
- IBM Edge Application Manager: An autonomous management solution designed to enable AI, analytics and IoT enterprise workloads to be deployed and remotely managed, delivering real-time analysis and insight at scale -up to 10,000 devices simultaneously.
- IBM Cloud Pak for Multi-Cloud Management: Provides consistent visibility, governance and automation from clouds, on-premises and the edge.
- Data Controls: Avoid cloud lock-in, maintain data residency requirements and keep control over sensitive persistent data when locating data stores on your own edge premise equipment or with private connections into distributed cloud storage.
- Interconnection: Secure high-speed access to data sources in hybrid architectures and partners in financial services, media networks, clouds and enterprise ecosystems.
- Scale on demand: Build capacity to handle surges in data and step it down to save when business slows.

AI Stack at Global Interconnected Metro Edge

In addition to having an industry-leading fully managed AI stack, it is important to host this stack at an interconnection hub that is close to the edge and public clouds for performance, cost and compliance reasons.

Hosting an AI stack at an Equinix colocation data center provides the following benefits: I at the metro edge: Perform AI at 200+ securely connected Equinix International Business Exchange™ (IBX®) data centers worldwide in 55 global metros.

© 2020 Equinix, Inc.