consoleconnect

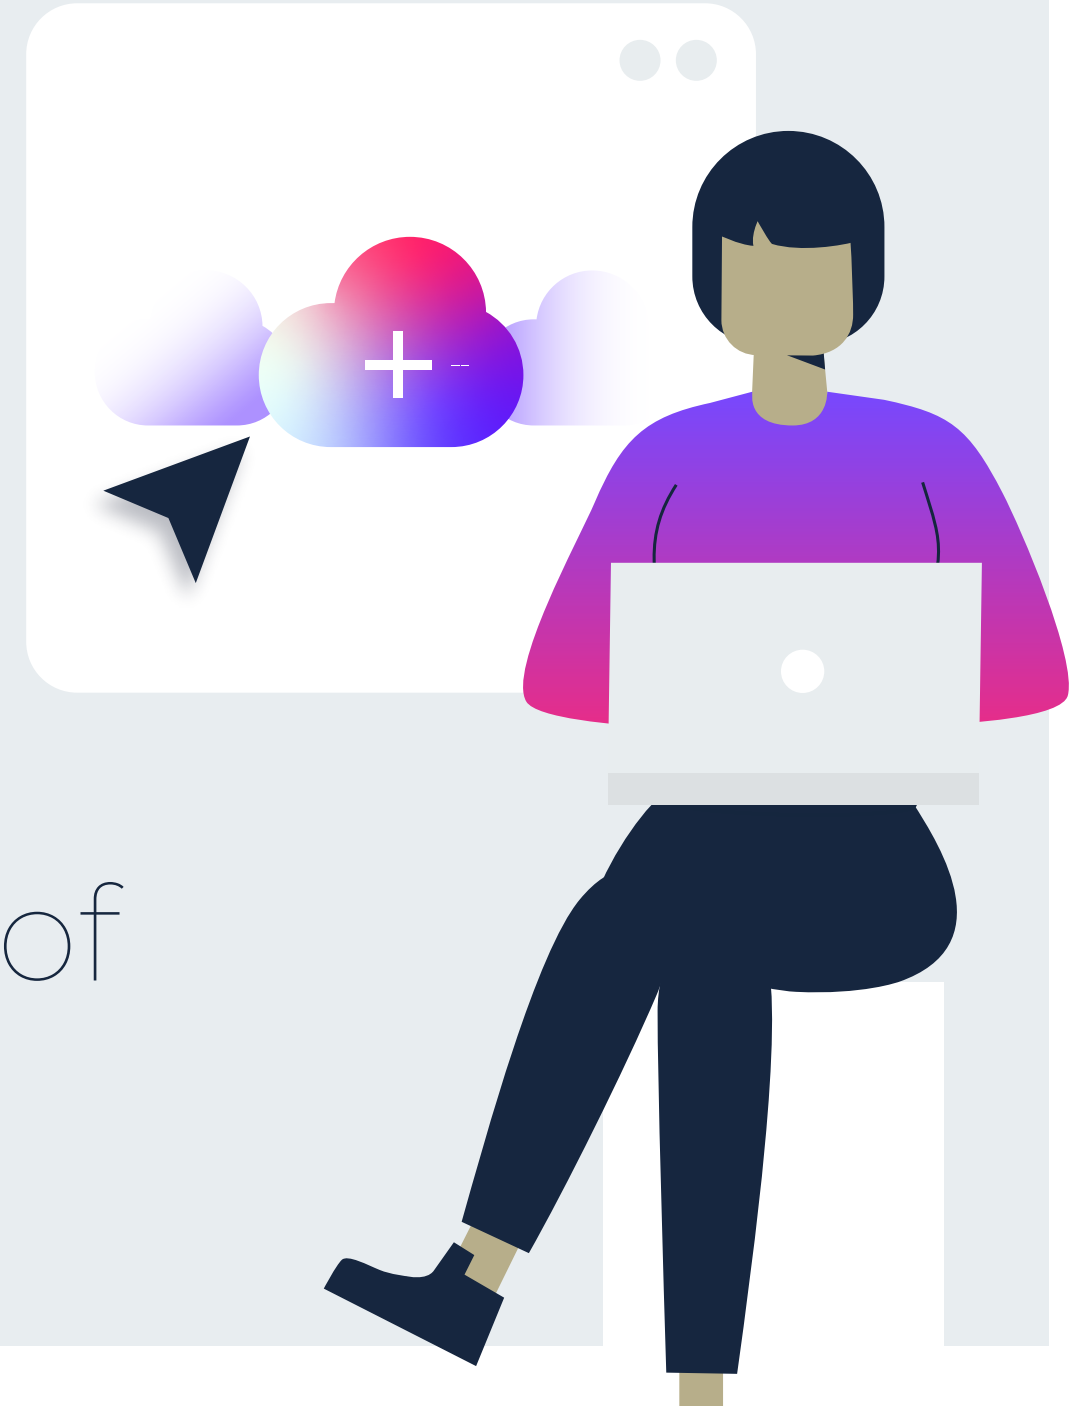

Understanding the challenges of **multi-cloud connectivity**

As individual business functions increasingly drive adoption of cloud-based technology, enterprises today find themselves using multiple cloud providers as much by accident as by design. Public cloud provisioning is no longer the sole remit of the IT department, with organisations seeking a best-of-breed strategy from cloud-based services and making decisions at a departmental level.

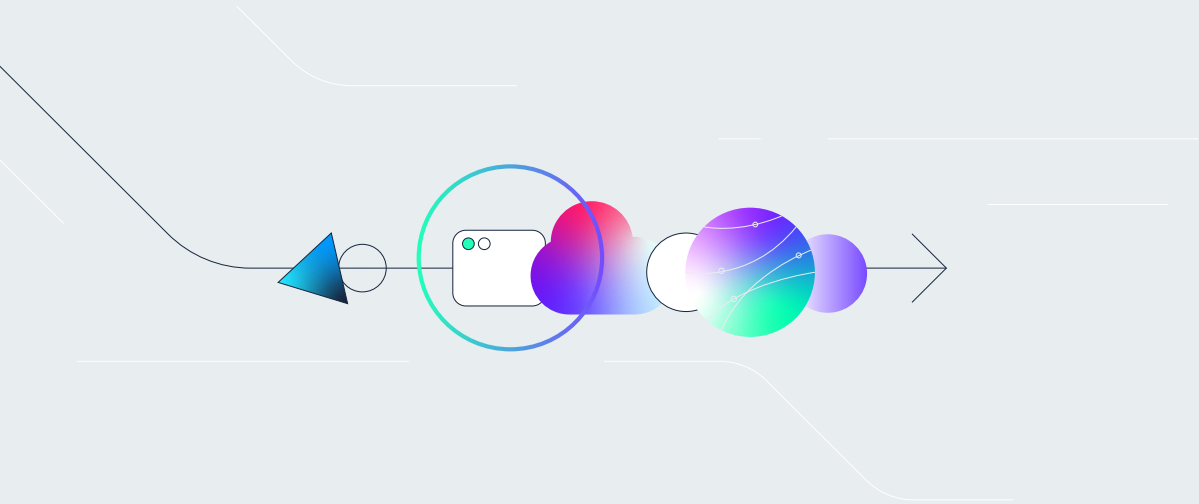

According to Flexera's 2021 State of the Cloud report, over 92% of enterprises now have a multi-cloud strategy, and while IT may not be responsible for commissioning the services, network managers certainly have a role to play in optimising the access and the cloud connectivity experience delivered.

A multi-cloud approach can introduce more network complexity because provisioning direct connections to multiple clouds is no easy feat - and as digital transformation strategies mature, organisations increasingly require the ability to transfer data between multiple clouds or across multiple geographies. Because interconnectivity between different cloud providers can be complex, and moving data between different regions on the same cloud can introduce unforeseen costs, multi-cloud connectivity is a real challenge for even seasoned network professionals.

In its "Predicts 2022: Connecting the Digital Enterprise" report, Gartner states that interest in public cloud networking and multi-cloud networking has increased dramatically in the last 15 months, largely because the native networking capabilities of public cloud providers vary greatly and are insufficient for many enterprise use cases.

In short, the connectivity infrastructure offered by major cloud providers is not designed to work seamlessly together with other providers in a multi-cloud world.

Key multi-cloud **connectivity challenges**

1. Multi-cloud data transfer

When public cloud providers first launched into the market nearly two decades ago, there was a typical land-grab for enterprise customer, much like we saw with databases or Enterprise Resource Platforms (ERPs). It's only in recent years we have seen providers like Google Cloud, Oracle Cloud, IBM, Huawei and others focus their acquisition efforts on winning secondary or tertiary cloud business from enterprise IT departments.

So, as multi-cloud becomes more common, multi-cloud use cases also proliferate, yet different clouds still do not play nicely together.

For organisations that use multiple different cloud providers, it is a common requirement to transfer data between different clouds. This may just be for the purposes of aggregating data but could also be due to different teams working in different environments across distributed cloud applications.

For some applications there are also considerations that may require hosting in an on-premises data centre or a cloud instance located in a specific geographic location. These might be business-critical apps, those that have high throughput, require low latency, strict security needs, or have geographic stewardship requirements such as data that needs to be hosted in a specific location due to GDPR.

So, for example there might be the need to move stored sensitive personal data from an Azure environment in a European country, where it is required to rest for GDPR purposes, to a Google Cloud Platform environment in another country to be computed, before pushing the results back to Azure.

Or, we might have a situation where an organisation has its ERP application installed in IBM Cloud, but it needs to call data from a storage lake in an Amazon S3 environment.

In these cases, the native networking offered by one cloud provider is not typically consistent with that of the second cloud provider, causing issues with performance and reliability

2. Multi-region cloud transfer

Interconnecting different brand cloud providers may be a frustrating but perhaps predictable headache. What often comes as a surprise though is that moving data between different geographic instances of the same cloud platform is not as straightforward as you might think.

Confusingly, all the cloud providers seem to have different approaches to cross-region data transfer and their own specific lexicon and definitions to go with it.

A common pitfall is that an organisation transferring a large amount of data from a public cloud environment in one region (Asia for example) to the same public cloud in another region (Western Europe for example), can be hit by hefty charges for traversing the cloud provider's backbone between regions.

Furthermore, this often isn't so easily solved by deploying dedicated connectivity to public cloud instances, because the network router will 'think' the data is staying within the one cloud instance and leave it be, when in reality it's travelling across the world on the cloud provider's native connectivity (as a chargeable transfer!).

Understanding this can be a headache, not least because the rules differ from provider to provider, but also because some data transfer types are charged only one-way (either in or out), while others have one fee when going in and another fee for outgoing transfers or between different zones or regions.

Let's look at the different approaches of the top three cloud providers.

AWS

Given its size and footprint, AWS has a more complex infrastructure than its closest competitors. For AWS, a Region is a physical location around the world where it clusters data centres, and each group of logical data centres is an Availability Zone. However, each AWS Region consists of multiple, isolated, and physically separate Availability Zones within a geographic area, unlike other cloud providers, which often define a region as a single data centre.

AWS also goes one further with the concept of Local Zones, which place compute, storage, database, and other select services closer to end-users and are specifically designed for highly-demanding applications that require single-digit millisecond latencies to end-users such as media & entertainment content creation, real-time gaming, electronic design automation, and machine learning.

Each AWS Local Zone is an extension of an AWS Region, providing a high-bandwidth, secure connection between local workloads and those running in the AWS Region.

Microsoft Azure

For Microsoft Azure, a region is a set of data centres, deployed within a latency-defined perimeter and connected through Azure's own dedicated regional low-latency network. An Azure geography is a segmented market containing one or more regions, that preserves data residency and compliance boundaries. This enables customers with specific data-residency and compliance needs to keep their data and applications in specific locations.

Azure Availability Zones are unique physical locations within an Azure region made up of one or more data centres equipped with independent power, cooling and networking.

Google Cloud

For Google Cloud, regions are independent geographic areas that consist of zones. Typically, a Google Cloud Platform region will have three or more zones allocated to it, allowing organisations to distribute apps and storage across multiple zones to protect against service disruptions. These zones are physically located in the same or a nearby data centre.

Essentially, zones and regions are logical abstractions of underlying physical resources provided in one or more physical data centres. The data centres themselves may be owned and operated by the cloud provider, or they may be leased entirely or in-part from specialist third-party data centre operators. However, the idea is still the same in that they are there to provide a uniform level of performance, security, and reliability.

But, because zones or regions may exist in physically separate locations, perhaps even in a separate geography or country, relying on a cloud provider's native connectivity to move data between instances can turn out to be costly for the reasons described above.

3. Disaster recovery

To avoid vendor lock-in and single points of failure, many organisations rely on multiple cloud infrastructure providers for automated data backup and recovery. This process requires reliable connectivity to and between different clouds, with native redundancy.

Although multi-cloud is often championed as offering best-of-breed, it can increase the amount of resource and infrastructure overhead needed to support from a redundancy perspective.

In many cases, multi-cloud redundancy requires that you set up a primary version of the application and data set on one cloud, then you do it all over again on a secondary cloud as a 'standby'. But connecting these clouds together is no easy feat.

Unfortunately, as we discussed above, there's no incentive for different cloud providers to provide a single co-defined open API for connectivity, which means a lot of work to access different clouds and even more work to get different clouds to talk to each other.

So, whether an organisation is doing a backup or a failover into Azure or a backup into Alibaba there are different API endpoints to contend with and different workflows happening behind the scenes.

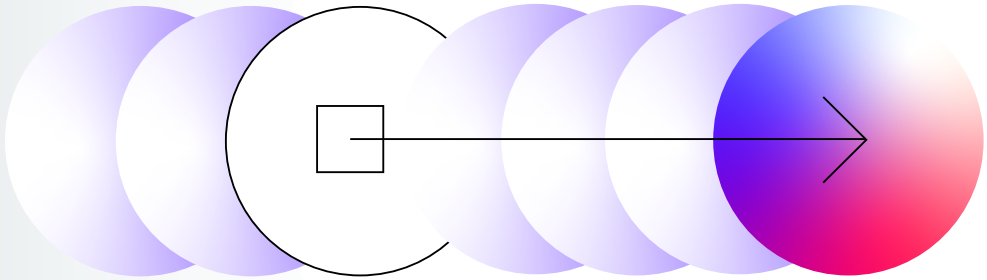

Introducing **Cloud**Router®

As businesses run more and more critical services in the cloud - from ERP and CRM (Customer Relationship Management), to billing systems - they are looking to take this traffic off the internet, with a view to improving security, resilience, and reliability as the data moves between data centres, public clouds, and different regions around the physical world.

This is where Network-as-a-Service (NaaS) has stepped in, as a solution for taking the complexity out of managing direct access to a multitude of different cloud providers.

Console Connect's new CloudRouter® is a multi-cloud interconnection service that simplifies inter-cloud connectivity, providing private Layer 3 connections between cloud providers and, by extension, cloud-based applications without the need for additional hardware.

Here's how CloudRouter® can address the key multi-cloud connectivity challenges highlighted in this e-book.

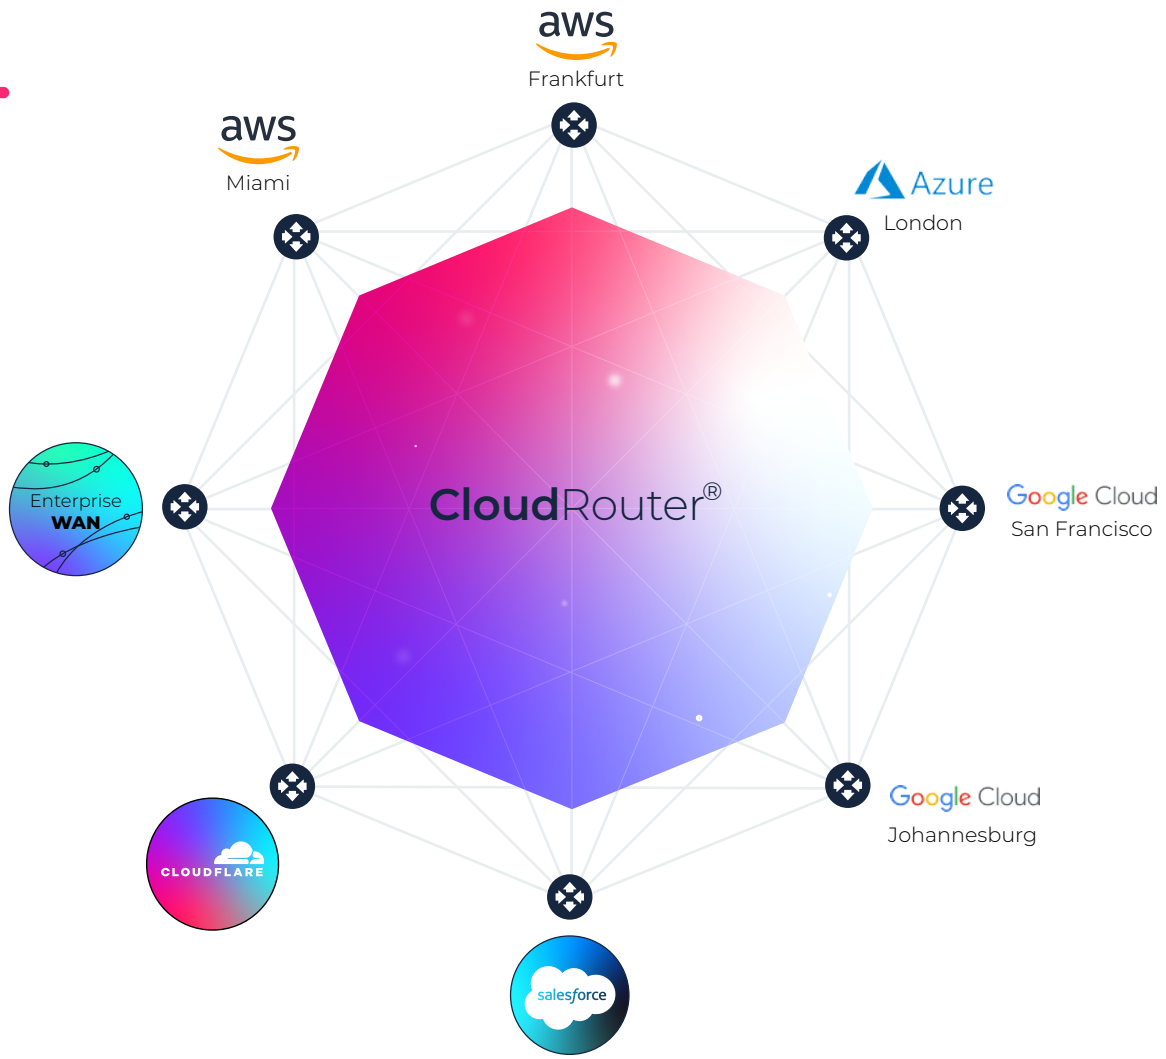

SCENARIO:

Multi-cloud data transfer

Use CloudRouter® to instantly create a private and secure meshed network for your multi-cloud environment - connecting any cloud to any cloud or any cloud to any location without the need for dedicated hardware.

Access all the world's largest cloud platforms, including AWS, Google Cloud, Microsoft Azure, Oracle Cloud, IBM Cloud, Alibaba Cloud and more, with real-time network provisioning and high availability of bandwidth.

But as well as connecting between different cloud platforms, you can also connect directly to SaaS ecosystems and enhance the performance of your applications with private connections to your SaaS provider.

Although doing this on a Layer 2 network is already possible, it requires access to a physical port. CloudRouter® enables data transfer between cloud provider virtual private clouds (VPCs) on a Layer 3 network without physical access.

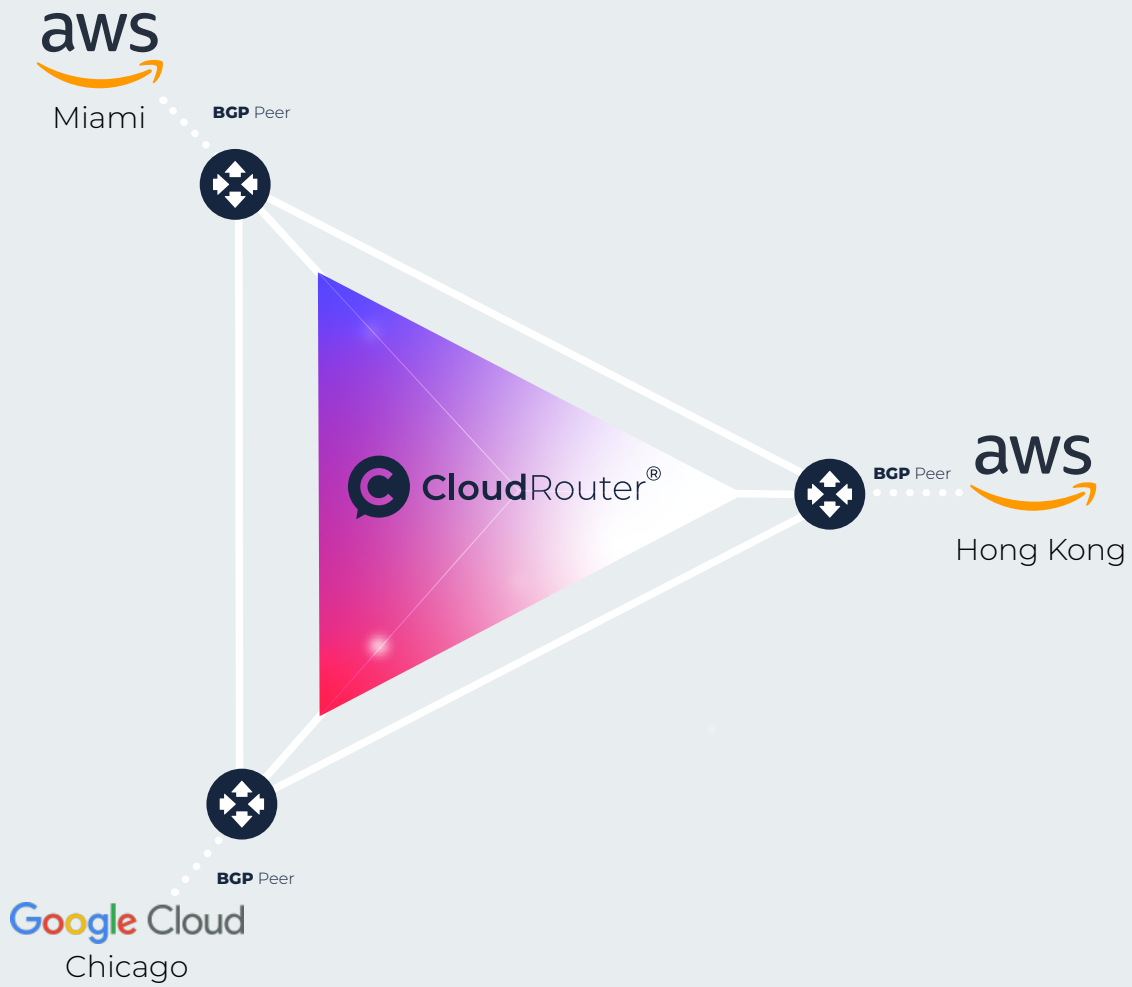

SCENARIO:

Multi-region cloud transfer

Avoid backhauling your network traffic between data centre locations and the cloud, even those locations that are operated by the same provider.

CloudRouter® is smart enough to take the burden off enterprises and avoid sending traffic over the cloud provider's backbone, sending it over PCCW Global's MPLS network instead.

What makes CloudRouter® different is that it creates a virtual 'full mesh' between network endpoints, ensuring that network traffic benefits from enhanced routing between data centre locations and the cloud. Simply use the Console Connect centralised management portal to add or remove new network edge locations and let CloudRouter® do the rest.

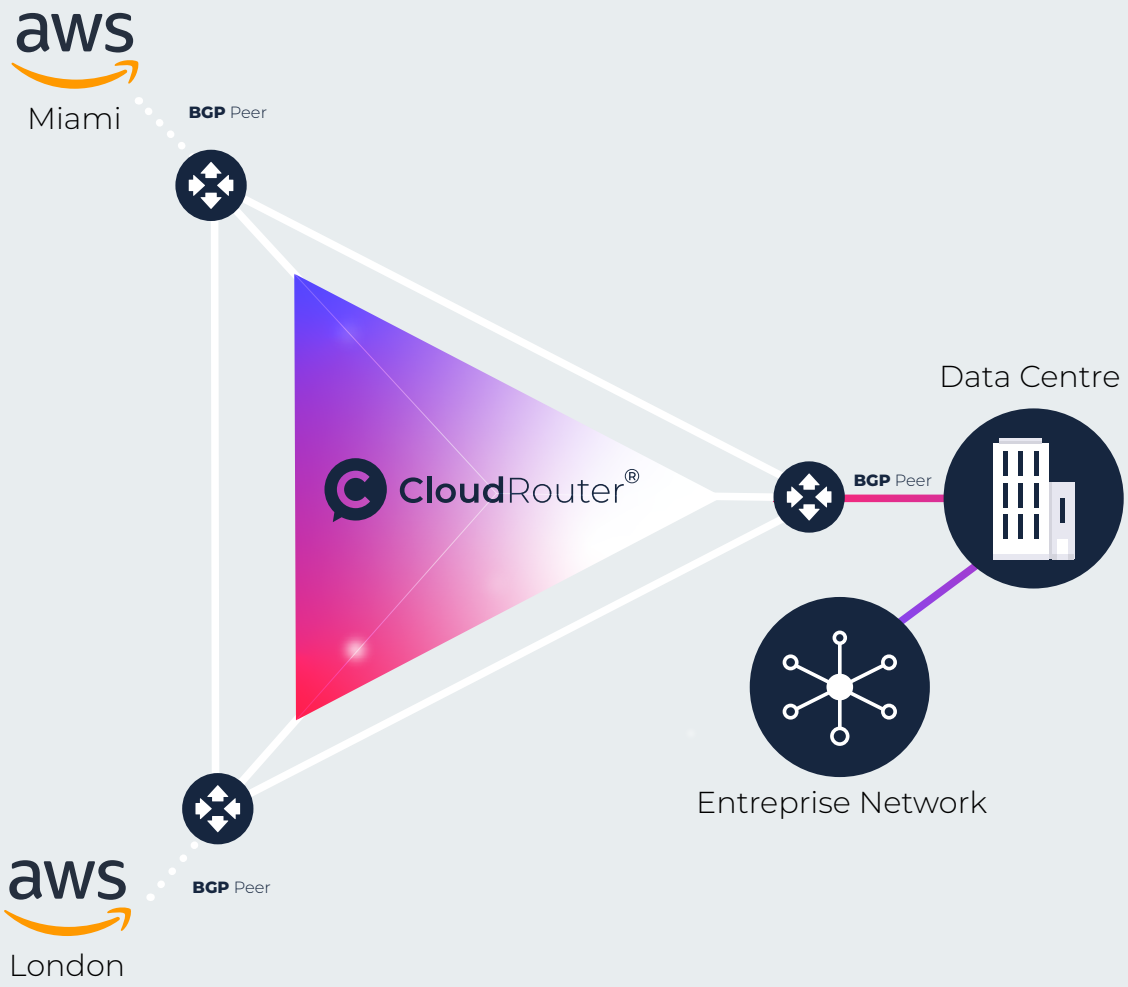

SCENARIO:

Disaster recovery

Because CloudRouter® provides a way to securely and cost-effectively interconnect multiple cloud providers or regions, the fact that CloudRouter® operates on a mesh network reinforces the potential for disaster recovery versus similar offerings from other providers.

The CloudRouter® mesh allows for rejoining the network at different points if a node goes down, whereas the hub and spoke model typically used by other providers can take a lot longer to rebuild in the event of a spoke failure.

Automate your data backup and recovery between clouds with reliable and redundant connectivity. Ideal for organisations that run primary systems on one cloud platform and backup or secondary systems on another.

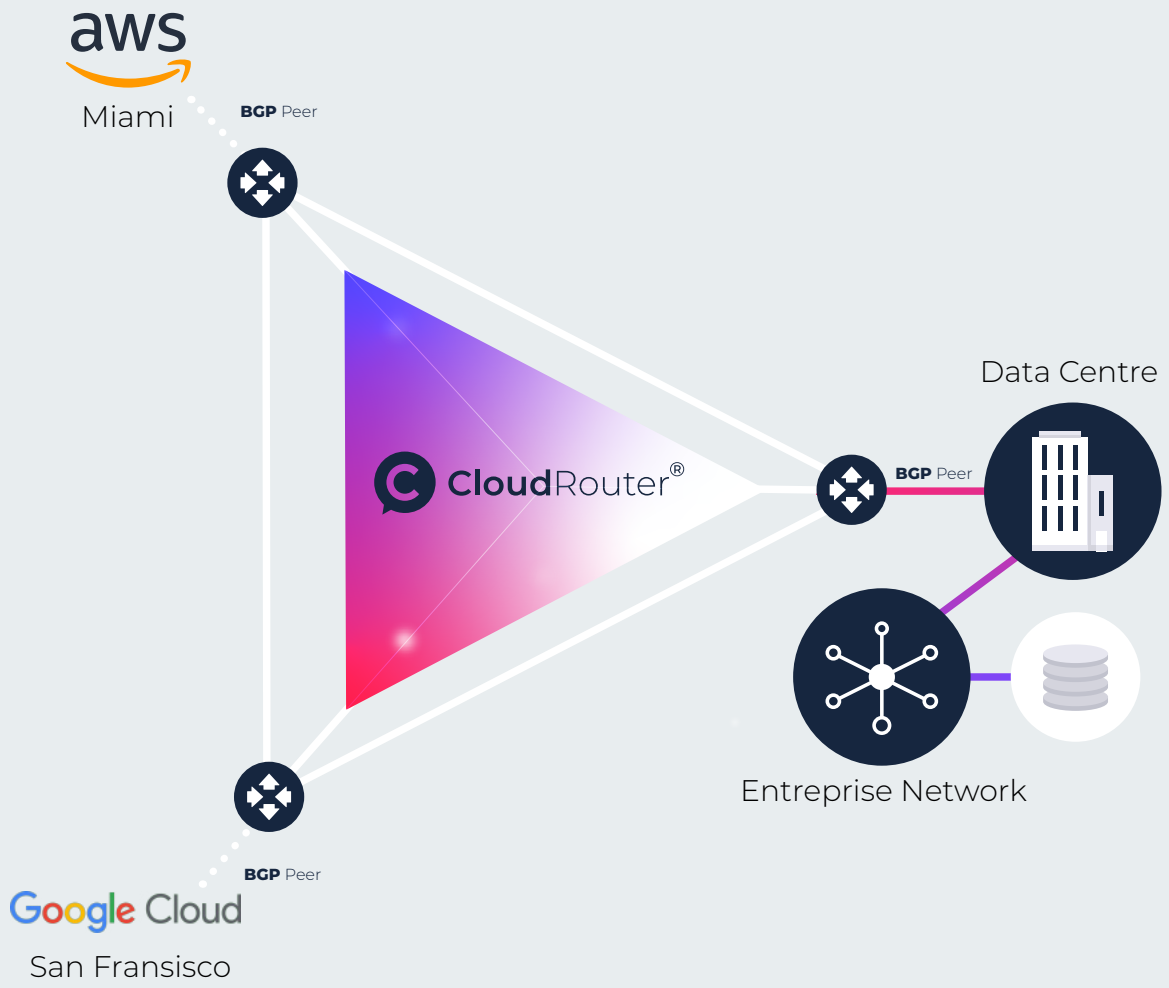

Predictable, scalable performance **for intercloud connection**

Built on PCCW Global's extensive IP network, CloudRouter® provides high performance, private data transfer between all leading hyperscale cloud providers across multiple global regions, and opens the door to direct connections to the entire ecosystem of SaaS providers.

With carrier-grade class of service and SLAs you would expect from a global MPLS network, SaaS regions are treated as remote instances and don't need to be collocated with your other cloud assets, the private Layer 3 network overlay created by CloudRouter® can get the traffic to and from wherever it's needed.

Furthermore, moving traffic in a very controlled private network like this also removes the need to route traffic back to an on-premise environment, avoiding associated latency and performance issues.

Once you are all set up, CloudRouter® takes care of the rest – as your traffic dynamically flows between all your network endpoints over our own private global network infrastructure. Which means better network performance, speed and security – and ultimately means an enhanced online experience for you and your customers.

consoleconnect

Australia
Level 3 | 200 Mary Street | Brisbane QLD 4000 | Australia

United Kingdom
7/F 63 St. Mary Axe | London EC3A 8AA | UK

France
2/F 16 rue Washington | 75008 Paris | France

Greece
340 Kifisias Avenue/340 Olimpionikon | Neo Psychiko 154 51 | Athens | Greece

Germany
Schillerstr. 31 | 60313 Frankfurt/M. | Germany

United States
475 Springpark Place | Suite 100 | Herndon | VA 20170 | USA

Singapore
6 Temasek Boulevard | #41-04A/05 | Suntec Tower Four | 038986 | Singapore

Hong Kong
20/F, Telecom House | 3 Gloucester Road | Wan Chai | Hong Kong

Japan
3/F Marunouchi Mitsui Building | 2-2, Marunouchi 2-chome | Chiyoda-ku | Tokyo 100-0005 | Japan

South Africa
Building 12 | 1 Woodmead Drive | Woodmead | Johannesburg 2191 | South Africa

UAE, Dubai
Office 504 & 505 | Level 5 | Arjaan Business Tower | Dubai Media City | Dubai

CloudRouter®

Easy as a click! Try it for free **HERE**

Have other questions we didn't cover?
Join our community of experts **HERE**

www.consoleconnect.com
TALK TO US: sales@consoleconnect.com