



An evolutionary leap in networks

Creating a cloud-centric future



Foreword

The future of connectivity is cloud-centric and global organisations are exploring new technologies and processes in a bid to operate successfully in this new environment.

However, taking a truly cloud-first approach is a significant step which involves embracing the scale of cloud connections needed and minimising the complexity of connecting into the hyperscalers and other cloud-based applications. Perhaps unsurprisingly, many organisations are finding this step change more difficult than they expected. Especially as they battle to make current network architectures support the high-bandwidth, low-latency, flexible port choice and on-demand requirements necessary for effective cloud-centric working.

Traditional site-to-site network architectures are rarely set up to embrace and support the connectivity and multi-cloud, multi-edge ecosystems that are the future of operations. Some organisations have managed to adapt their networks to cope, but it's uncertain how far into the future, as demands increase, that these adaptations will hold.

For many more organisations, however, the network is too rigid or costly and is acting as a 'drag'.

To meet this challenge, organisations need a partner who can future-proof their connectivity and deliver network-resiliency confidence, while providing the flexibility and agility needed so they can seamlessly adjust to multi-cloud complexity, and ensure key applications perform even when workloads are increasingly distributed.


After carefully listening to our customers, assessing the market's direction of travel, and consulting leading analysts, we're channelling our investment into driving an evolutionary leap in networking.

The result is Global Fabric - the network for today and tomorrow's needs.

This paper shares our analysis of the challenges facing customers - based on market and analyst insight and primary research with customers just like you. It then lays out precisely how Global Fabric will unite multiple elements to provide the hero networking solution your organisation needs today.

I hope reading this will provide clarity about the role Global Fabric will play in supporting your organisation's strategic objectives and future-proofing your connectivity, well into the future.

Matt Swinden
Director Digital Connectivity, Business



Introduction

86% of organisations are planning to move most or all of their IT infrastructure and applications to the cloud within the next five years. With this knowledge, we set out to create a network ready to meet customer needs. As part of the process, we commissioned independent customer research to give us a granular picture of the global network experience. What emerged was a complex scenario where factors intertwine to create significant blocks to cloud-first networking.

The five issues with current networks that stood out as holding customers back from a cloud-first reality are summarised below.

1. Optimising app performance

Employee productivity is now intrinsically linked to not only application availability but also performance. Wherever an application is hosted, the effectiveness of its connectivity to other workloads across a complex hybrid ecosystem is critical to its performance. This is not just for latency sensitive apps, but core applications that run organisations. Poor performance leads to poor customer experience, reduced engagement and lower productivity. What's required in the event of a network failure is a network that still performs. One that automatically re-routes traffic, so latency is never impacted. And one that delivers coverage of the full estate to quickly identify performance issues and proactively deploy solutions to both mitigate them and ensure applications are always seamlessly connected.



2. Delivering operational flexibility

Manual partner processes with long lead times are a 'drag' on the ability of the organisation to be agile and respond to market conditions. Organisations want their connectivity to change as they change. Whether it's increasing bandwidth as data use increases, leveraging the benefits of SD-WAN intelligence and encryption or connecting to a new site or application. These tasks need to be completed at pace, freeing up time for organisations to focus on their business goals.



3. Managing network costs

As more workloads move to the cloud, complexity goes up and cost visibility goes down. Organisations want greater cost clarity with the flexibility to change, as well as the confidence they're only paying for what they use. As new data policy regulations come into effect (for example the European Data Act), effective spend management will become increasingly vital. The war on OpEx is coming, and organisations need to have all levers at their disposal to manage their costs - from seamlessly switching cloud provider, to reducing equipment and avoiding bill shock from egress fees when taking data out of public clouds.



5. Driving sustainability throughout the evolution

Organisations are increasingly looking to manage their Scope 3 emissions but are finding that understanding their baseline is not straightforward, let alone moving the dial. Therefore, they need the ability to calculate the carbon footprint of products and services and understand how they can be reduced by using modern, energy-efficient networks.



4. Enabling robust security

Connecting to a complex, multi-cloud estate and enabling hybrid remote working practices exponentially increases the attack surface. Organisations want security embedded in the network coupled with the ability to easily layer in additional security controls, facilitating the journey towards Zero Trust without impacting performance and user experience.



Resilient networks, resilient futures

To create an enduring solution to these challenges, we first delved deeper, undertaking detailed research to expand our understanding of how the network is affected when taking a cloud-centric approach.

The following sections lay out our thinking around the five critical areas.



1. Optimising app performance in a distributed network environment


In today's highly competitive economic landscape, apps are essential for reaching and engaging with customers, optimising operations and staying ahead in the market. With workloads distributed across hybrid infrastructures consisting of data centres, on-premises environments, edge and public and private clouds - app performance and resilience are make-or-break issues.

Every IT team is well aware that, if access to an app goes down or there's a latency issue with an app that requires real-time or near-real-time interactions, it can stop employees working, cause mission-critical parts of the organisation to grind to a halt and impact badly on customer experience.

However, traditional networks struggle to provide the level of support these distributed workloads need, and are unable to guarantee the performance, scalability and resilience required by highly distributed applications. Networks need to change to make sure critical apps aren't impacted in the event of failure, so employee productivity stays high, end-customer experience continues without a break and external business isn't lost.

SD-WAN alone isn't the answer

A few years ago, SD-WAN was seen as the easy option for improving the flexibility of routing and path selections to support app performance. However, today, organisations realise that SD-WAN can't fully meet their expectations of cost savings and transformed performance because performance can only ever be as good as the quality of the underlay it's paired with.



Three pillars of optimised app performance

It's clear networks need to change in order to deliver optimised app performance; the question is how to make this an easily accessible reality. Embedding these three key capabilities will be essential.

1. Wider, deeper, real-time monitoring

End-to-end visibility will enable better control over app performance, shifting monitoring from passive to active. This starts with near-real time reporting of network performance, broken down into latency, jitter and packet loss across the full delivery pathway from site to cloud.

Full-stack observability and app performance insight becomes even more powerful when Artificial Intelligence (AI) and machine learning (ML) capabilities are incorporated - creating the ability to automate the discovery and mapping of network and app structures, for example. On a more macro scale, this insight will enable IT teams to spot apps that aren't hosted in the best places and identify purchased capacity that's no longer needed.

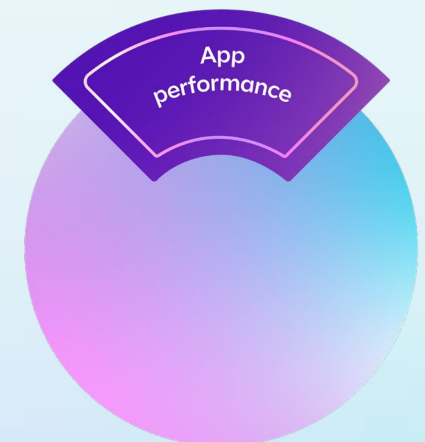
Plus, organisations need to look beyond the standard network SLAs that aren't robust enough to guarantee app performance, asking instead for Experience Level Agreements (XLAs) to measure the entire app journey. These will protect against the performance troughs that impact employee productivity and customer experience, embedding a culture of continuous customer experience improvement.

2. Software-defined intent-based routing

After establishing the visibility requirements for the network, the next priority should be building in the ability to change routing according to business needs. Part of optimising app performance is ensuring latency sensitive apps can follow the lowest latency path - even under failure conditions. This is clearly beyond manual management approaches, so needs to be software defined to give on-demand control over routing policy. Although primarily a tool for app optimisation, software-defined routing is also an effective way to comply with data sovereignty regulations, allowing you to control geographically where sensitive data travels.

3. Pre-emptive fault resolution driven by AI and ML

Advanced visibility and routing capabilities need to be combined with near real-time fault detection and remediation to protect and enhance the customer experience. AI and ML can learn about data traffic, anticipating and pre-empting events likely to impact on app performance. This goes beyond spotting obvious signs of an outage, to looking for early warning signs that, in isolation, may not trigger an alarm. Automated systems can 'see' particular event sequences happening over a period of time that may signify a potential service outage. AI and ML can also deliver a clearer monitoring picture by de-duplicating the 'noise' from automated events, leaving a single abstracted view of alerts, correlated to network elements.



2. Delivering the operational flexibility that traditional networks can't



Almost everything has changed in an organisation's operating environment except the networks that underpin their businesses. This lack of network evolution creates a 'drag' on operational efficiency and an inability to make intelligent, dynamic decisions regarding traffic routing, optimisation and management across the network.

Transformation needs to start from the ground up and should include the ability to digitally manage and maintain the network end-to-end, and scale services up and down as operational needs change.

Typical provisioning lead times of days and weeks are way too slow; these timescales made sense 15 years ago when data mainly moved between sites or to a private data centre. Today, IT teams need flexibility from providers, removing network friction and enabling smooth workload transfers to and between public and private clouds in near real-time, with light-touch management.

Reduce the pressure on IT teams

A cloud-centric network uses digitalisation to transform the management burden and boost efficiency - enabling rapid responses and dynamic scaling to right-size capabilities, for example, increasing bandwidth as data use increases and providing the flexibility to make logical changes in near real-time.

For IT teams, this can dramatically cut the time spent chasing up and fulfilling orders as well as increasing their capacity for value-added activities.

This type of cloud-centric flexibility will allow IT teams to manage their network via their choice of digital interface - be that a portal, an app or using APIs integrated into their service management system. This digitalisation of services streamlines and simplifies day-to-day network tasks by, for example, making it easier to deploy multiple network services onto common infrastructure, like ports. This provides flexibility and reduces complexity and network friction.



Transform cloud workload management

An evolved network opens the door to more flexible workload management options in multi-cloud environments. Whatever the cloud scenario - from moving to the cloud, repatriation from the cloud, moving more to an existing cloud, or switching to a new cloud - IT teams should have the ability to spin up bandwidth paths rapidly when required and reduce or tear them down the moment the migration is complete.

Take the complexity out of DIY cloud connectivity


Connecting to multi-clouds can seem deceptively straightforward. However, the reality of a DIY approach is a fragmented, complex point-to-point system. IT teams need an integrated end-to-end connectivity solution that encompasses all their connection points, cutting the administrative resources required to manage each component in the delivery chain from multiple suppliers.

They need a safety net that a DIY approach cannot provide; the protection of a service wrap or proactive SLAs can vastly reduce the need for external expertise and troubleshooting.

Enhanced operational readiness can only be provided by an end-to-end cloud-centric approach to cloud connectivity, enabling IT leaders to free up their people, reduce delays for time-sensitive projects and decrease risk.

For example, even though private connectivity to the cloud is essential for both security and performance, many fabric providers only support private connectivity based on Layer 2 Ethernet Virtual Circuits (EVCs). From the organisation's point of view, this requires complex meshing designs that are difficult to manage on an ongoing basis. Layer 3 capabilities are sometimes provided, but they're supported by 'cloud routers' that usually have capacity constraints.

The bottom line is that even the most advanced forms of DIY network management often come with hidden overheads in terms of time, expertise, cost and flexibility.



3. Managing network costs in a multi-cloud world

Managing network costs in the cloud is not a one-size fits all process; minimising complexity and maintaining cost visibility is a huge challenge.

This is made more difficult by inflexible commercials, inconsistent cost modelling across different hyperscalers and pricing structures that may be clear at first glance but that get more complex upon deeper examination, hiding extra costs.

Leading analysts widely expect cloud usage levels to jump sharply over the next few years, with today's approximate figure of 20% of all workloads sitting in the cloud growing rapidly to around 60%. Anticipating this, organisations are building a projected view of Total Cost of Ownership (TCO) to identify and tackle any factors blocking cost optimisation now, so they don't hold back cloud progress in the longer term.

This process begins with an investigation of what's behind the current cost pain points in cloud services.

Bill shock wrecks cost optimisation efforts

Complexity and a general lack of transparency around cloud connectivity costs means many organisations could experience bill shock. Partly, this stems from how easy it is for different teams to independently purchase hyperscaler services, without oversight from the central IT team. Only when the invoices come in do financial controllers know the extent of cloud usage and costs. What's missing is a holistic calculation of value and TCO.

However, the biggest bill shock culprits are the charges to move data out of hyperscaler cloud environments. Many organisations eager to implement a cloud-first approach to new projects rush to upload their data into hyperscalers, without fully considering the economic impact of sending data out of the cloud after processing, or moving workloads themselves

between clouds - both of which attract volume-related egress charges.

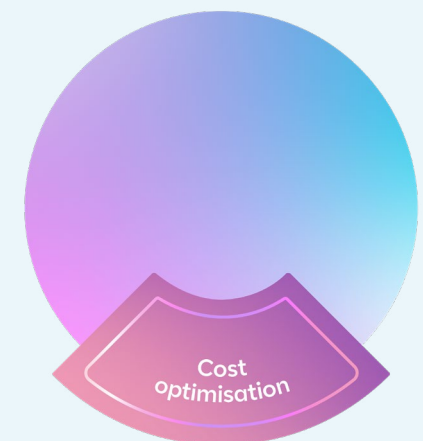
Optimising your design for cloud connectivity can help reduce egress charges – for example, there's a break-even point where using a private connection via ExpressRoute into Azure is more cost effective than the public internet.

We also see changes coming to the way hyperscalers charge for egress, driven by new legislation such as the forthcoming European Data Act, which will enable easy switching between cloud providers. Your network should remove friction, reduce long change lead times, enable the use of lower-cost network paths or services, and allow you to migrate workloads, applications and data quickly and seamlessly between different major cloud service providers.

Network transformation unlocks savings

Modern networks also support cost optimisation by allowing infrastructure to be shared – unlike traditional networks that have a

more expensive structure, needing dedicated physical components (such as a port, access or CPE) for each service (such as internet, IP VPN and Layer 2 Ethernet). This shared infrastructure drives efficiencies and reduces TCO because it supports efficient logical provisioning and in-life management of services. It also enables organisations to run multiple services on the same shared hardware and provides the flexibility to optimise total bandwidth across all these services on-demand.



4. Elevating security for a cloud-centric future

Organisations need to keep pace with the expanding threat landscape and multi-faceted defence challenges of a cloud-centric world.

Traditional perimeter-based models of security are no longer fit for purpose when more data is stored off-premises than on, and more work is conducted outside office sites than inside.

Fully converged networking and security models such as Secure Access Service Edge (SASE) are suitable for particular use cases. However, in many instances, organisations need more flexibility to support their hybrid environments. The pressure is on network partners to deliver networks that can integrate with multi-layered security that's architected with a Zero-Trust approach for working in and across clouds.

Organisations are looking for security outcomes that allow them to protect their users and endpoints, secure their sites and supply chains, secure their access to the internet and to protect their hybrid cloud estates and the data within them, all without compromising on performance.

Realising end-to-end security

As organisations' use of multi-cloud grows, their security focus must shift to protecting vulnerable edge breakout points and delivering defences right up to the carrier neutral facility (CNF) 'doorway'. Part of embedding security will also include a Zero-Trust architecture that assumes all network access is potentially malicious or undesirable. This makes authentication of identity and least-privilege access, combined with multi-factor authentication, essential. Plus, achieving inherent network security must involve building in data protections and providing end-to-end route path controls that comply with international sovereignty standards and legislation.

Three essentials for a new security approach

In a market crowded with options, organisations should focus on three core elements for their integrated network and security architecture.

1. Network design innovation to better secure internet links


Organisations are keen to further leverage the scalability and cost-effectiveness of public internet, however this is inherently insecure. You can navigate this issue by collapsing network services onto a single core using Multiprotocol Label Switching Segment Routing (MPLS-SR), allowing internet traffic to take advantage of private network security, while still using cost-effective internet services. By assessing all hops from origin, segment routing not only delivers the lowest latency path, but can select routes to maintain compliance and data sovereignty. Additionally, by embedding Distributed Denial of Service (DDoS) protection directly into the network, organisations can take advantage of this protection for all their network services - and even expand the service to cover their own sites.

2. Service insertion at CNFs and direct peering with SSE providers

Implementing a comprehensive security posture that meets an organisation's specific requirements can involve a mix of security service insertions into CNFs, either virtually or physically, and integrating direct links to Security Service Edge (SSE) providers, too. This ability to mix and match services while maintaining a single security policy is a crucial defence point against cyber threats for ensuring the confidentiality, integrity and availability of data and services.

3. A programmable network core to align network and security together

This is critical for a dynamic and agile approach to security that can keep pace with the rapidly changing threat landscape. The ideal is a network with a software-defined end-to-end core that can embrace and flex to any security-driven changes the organisation wants to make. And, in turn, the organisation's security stance can adapt to any requirements the network needs, using network telemetry to identify and proactively respond to threats, leveraging the capabilities of generative (GEN) AI.



The journey to a secure network

Creating this programmable network does involve major transformation, but it doesn't always mean a rip-and-replace approach. The overall development emphasis should be on rearchitecting and converging in targeted ways to take advantage of multi-service edge capabilities, drawing on expertise to overlay security in precisely the right spots for the organisation. Organisations should ideally have an option to try out new secure network capabilities on a test site before buying into it on a larger scale. This way they can experience the user interface and digital API integration available, and assess the value on offer, before they invest more significantly.

A key element will be rearchitecting the network so it's easy to incorporate security configuration tools for comprehensive monitoring and visibility, as well as proactive security control, across hybrid and multi-cloud environments.

What does this approach mean for multi-cloud security?

An integrated secure network will protect diverse attack surfaces, meet data governance and compliance requirements, support a Zero-Trust approach and enable visibility and monitoring to drive intelligent defences. It'll be capable of protecting complex environments using consistent enforcement policies - securing users, endpoints, sites and internet access across supply chains and any cloud configurations.

With their network as a true enabler, organisations can build a service that fits their specific needs and is simple to deploy and manage, easing the impact of the cyber security skills shortage. Their network provider should also offer flexible co-managed and managed options, and integrated security monitoring and reporting services, so organisations can build intelligent, tailored and holistic protection as they move into the multi-cloud world.



5. Actively contributing to sustainability goals

Organisations committed to science-based net zero targets need to minimise the impact of their digital network infrastructure as it evolves.

Although sustainability has had a seat in the boardroom for some time, it now needs to be operationalised. CIOs must play a significant role, if organisations are to achieve their carbon reduction targets and this must translate into clear requirements for IT teams.

This can be broken down into three steps: defining goals, accurately identifying and monitoring sources and volumes of carbon emissions, and concrete actions to reduce the organisation's carbon footprint.

1. Define and disseminate sustainability goals

In a network context, infrastructure providers are a core part of most organisations' supply chains, and that puts them front and centre in the organisation's mission to reduce their Scope 3 carbon emissions. Any network partnership should involve the provider aligning with the organisation's sustainability objectives and supplying effective routes to achieve them.

As a result, establishing how a network provider's strategy and operations can deliver low carbon-usage levels must be part of any selection and contract process. Purchasing teams should look for providers who have woven sustainability into their end-to-end planning and can actively support the organisation's strategy to reduce their carbon footprint.

2. See the true carbon emission picture

An organisation can't change what it can't measure and, increasingly, they need to be able to baseline and track network carbon footprints in a standardised way, so they can quantify and demonstrate how they're reducing their carbon emissions. Incoming regulations such as the Corporate Sustainability Reporting Directive in Europe and the Climate Disclosure Act in the UK strengthen this imperative.

IT teams should challenge network providers to share detailed information about the energy use and overall carbon emissions of their services, as well as thinking ahead by asking what investments they're making into continuous innovation to drive future sustainability improvements. In particular, they should look for network infrastructure providers that build renewable energy use into their services.

3. Utilising effective sustainability tools

Reducing the organisation's carbon footprint requires effective tools to benchmark carbon emissions today, and to measure and monitor them on an ongoing basis.

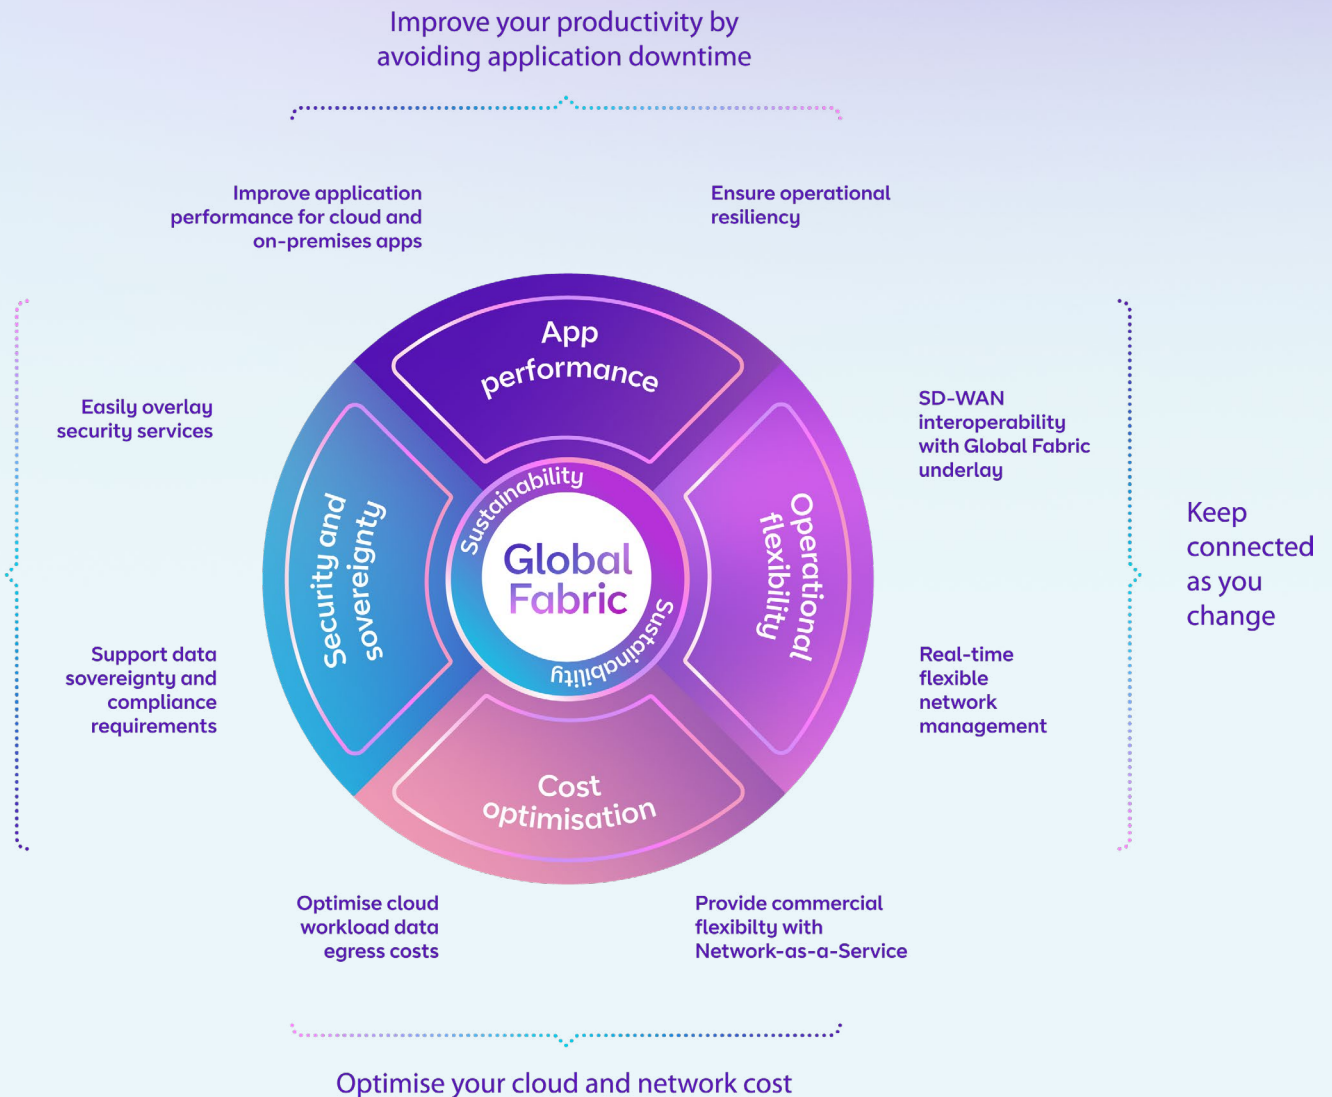
These should be combined with optimising device and data centre management to both minimise carbon emissions and to accurately track and allocate emissions. Providers should offer easy access to insights into network performance and emissions and be challenged to supply market-leading low-carbon services, simple carbon calculator tools and clear dashboards that summarise deep insight into carbon usage.



Introducing Global Fabric – connecting you to everything

Our new Network as a Service (NaaS) platform, Global Fabric, is a next-generation network that transforms connectivity for modern organisations.

Global Fabric is the key to improving your productivity by avoiding application downtime, removing network friction and maintaining resilient connectivity as your network evolves. It also offers clear routes to optimising your cloud and network costs, reducing your cyber risk and improving sustainability throughout your estate. Global Fabric is the network solution for today and tomorrow.



Improve your productivity by avoiding application downtime

Global Fabric delivers the redundancy and resiliency that highly distributed and latency sensitive applications need - which lets you boost your employees' productivity, provide uninterrupted end-customer experiences and ensure you don't lose any external business.

What's behind this?

Global Fabric seamlessly integrates global coverage, dense metro architecture (with at least two Points of Presence (PoPs) in each metro zone), and software-defined intent-based routing. This combination delivers a resilient, high-performance network that easily adapts to changing conditions and optimises traffic flows to meet specific requirements.

Designed with a high-capacity core, Global Fabric handles large volumes of data or traffic easily. It also provides real-time historical performance metrics that allow you to smoothly manage and optimise the performance and reliability of in-life business-critical applications. It also protects the health of your end-to-end application environment, by using continuous tracking, and analysing metrics such as latency, jitter and packet loss. Driven by comprehensive synthetic monitoring capabilities covering site-to-cloud and cloud-to-cloud routing, Global Fabric, simplifies performance observability as you move more workloads to and between clouds, enabling the proactive identification and resolution of issues.

To this we add robust, end-to-end connectivity and full-stack observability enhanced by our AI and ML capabilities. All these services are backed by a wide-range of rigorous SLAs - to the cloud, and between both public and private clouds - as well as XLAs to measure user sentiment, helping you to improve your users' overall experiences and making sure you get the best possible service from us. Taken together, these capabilities allow you to deliver a seamless user experience and derive maximum value from your organisation's cloud investments.

Get connected as you change

Global Fabric is an advanced, flexible network infrastructure that allows you to shape your connectivity experience effortlessly, meaning you can reduce the time you spend chasing up orders and increase your capacity for value-added activities. The days have gone where you had to wait and chase your provider while lengthy physical changes took place. Now, you can make a rapid sequence of logical, digital transformations.

What's behind this?

With Global Fabric, you can make near-real-time logical changes, via an easy-to-use digital user interface and cutting-edge API-first integration. This means your team can configure and control the network logically, making modifications or adjustments, such as creating new connections and tweaking bandwidth, with just a few clicks. This flexible network adaptability speeds up changes and ensures that your network integrates seamlessly with your existing systems and workflows.

Global Fabric is also architected to expect your connectivity requirements to change, so is able to provide direct connections at the click of a button into 74% of hyperscalers' global PoP locations. It's also pre-connected to 700 global data centres and 630 cloud product and partner services and has global connectivity that covers 92% of our customers' network location needs. You can run different services concurrently, without requiring dedicated ports and access for each service. This gives you greater adaptability and enables you to seamlessly connect with a diverse range of services and partners.

The combination of an optimised global footprint and our performant business internet and private services means you can leverage the latest cloud developments without compromising any ongoing operations. And, because every organisation's evolution is different, Global Fabric has been designed to mould to your ways of working with a choice of flexible management options - ranging from DIY to fully-managed and hybrid management approaches.

Optimise your cloud and network costs

Choosing a pay-as-you-use model with Global Fabric gives you the flexibility to pay only for the services or resources your organisation needs, scaling services up or down based on requirements. As a result, you can control your costs and resource allocation more closely, especially in situations where usage may vary over time.

What's behind this?

With Global Fabric, you can run multiple services on the same port, increasing the efficiency of how you use network resources. This also allows you to reduce your customer premises equipment (CPE), and fewer devices mean lower capital and operational costs, as well as simplified maintenance. This all helps streamline your network infrastructure, resulting in easier troubleshooting and upgrades, that further contribute to a more agile and cost-effective network.

Global Fabric also tackles the issue of egress costs, as our pre-connected links to hyperscalers offer a stable and cost-effective alternative to the high-cost and unpredictable variable egress charges you're subject to over the public internet. We'll help you to cost-model egress charges and advise you on how to optimise your cloud costs, eliminating any potential bill shock. With transparent and controlled expenses, you can align your data transfer needs seamlessly with your budgetary requirements, giving you both reliability and financial predictability.


Reduce your cyber risk

Global Fabric is designed to support your Zero-Trust journey, without compromising your network performance. The combination of multi-service edge capabilities and decades of security expertise, allows us to overlay security in precisely the right spots for your organisation.

What's behind this?

Global Fabric has DDoS detection and mitigation embedded as standard and can also support advanced DDoS protection services. Plus, the end-to-end route control it offers across its segment-based, flexible routing makes data sovereignty compliance simple and reliable.

Global Fabric offers a mix of security service insertion into CNFs - either virtually or physically - and pre-integrated direct peering links to our SSE partners. Being able to mix and match services while maintaining a single security policy like this is a crucial defence point against cyber threats, ensuring the confidentiality, integrity and availability of your data and services. You can also layer on further security elements with Global Fabric. It's designed to support quick and straightforward deployments of consistent, unified, end-to-end overlay functions such as SASE and endpoint and identity security.



Driving sustainability

Helping organisations to operate more sustainably is part of Global Fabric's DNA. Our dashboards and tools help you to make smarter, more sustainable decisions on how to run your workloads and applications.

What's behind this?

Our dedicated Carbon Network Dashboard allows you to baseline, monitor and manage devices across your entire estate, providing full carbon footprint data for your Scope 3 reporting. Plus, our Digital Carbon Calculator supplies ongoing carbon tracking, providing accurate, reliable and consistent data that allows you to measure and reduce your carbon emissions across your network infrastructure.

Global Fabric is not only powered by 100% renewable energy, but it also demonstrates an environmental commitment that resulted in a 79% reduction in carbon footprint, when compared to our current MPLS network¹.



Future-ready network resilience

Global Fabric, our revolutionary NaaS platform, is waiting to empower IT leaders to embrace their cloud-centric future. This transformative solution, built in new global locations using the latest hardware and software technology, allows us to redefine connectivity for modern global organisations.

As your strategic partner, we'll guide you through the complexities of modern networking, highlighting how Global Fabric can deliver a comprehensive and future-proof solution. Get ready for unparalleled efficiency, resilience and success in the dynamic landscape of global connectivity.



Our credentials as your partner

Leading organisations choose to partner with us because we have one eye on the present, and one on what's to come. As a proactive technology partner, we scan the horizon for technology trends, craft them into effective tools, and then help our customers incorporate them into their digital infrastructure to achieve their business ambitions.

Strengths of our partnership include:

Our global reach and scale

We can serve customers in over 200 countries, with access to 700 data centres satisfying data and application requirements from a performance, security and regulatory perspective.

Digital connectivity control

We provide choice, flexibility and control, from DIY to fully managed, with real time in-life performance visibility.

Resilient performance

Our resilient, high bandwidth, global core network, has a dense metro PoP architecture to maintain network performance even under failure.

Global Fabric

By 2030, we aim to become the most trusted connector of people, devices and machines. Leading the way to a bright, sustainable future through responsible, inclusive and sustainable tech.

Greater TCO control

Our PAYG model means there's no contract lock, so bandwidth and costs can be managed according to your business needs, and network designs optimised to avoid egress bill shock.

Outstanding security

We have the breadth and experience of capability to secure how our customers use the multi-cloud, with no compromise on performance.

Pre-integrated ecosystem

With our extensive hyperscaler partnerships and peering agreements, we help you simplify your supply chain and network management.

Connect to Global Fabric, and connect to everything

Global Fabric will open the way to mastering the multi-cloud and staking your claim in the new digital universe.

Talk to your account manager today to [find out how to start your evolution to Global Fabric.](#)



1Sustainability facts and figures:

We estimate that when fully rolled out, Global Fabric will use 79 per cent less electricity than its current global networks.

Using the Greenhouse Gas Protocol ICT Sector Guidance by Gesi, we have generated the following estimates:

- In use, Global Fabric will consume 8,326 MWh/year versus its existing international networks at 39,890 MWh/year — a 79 per cent reduction. Use stage gross carbon, including PUE will be 2,964 tonCO₂e/year for Global Fabric versus 13,596 tonCO₂e/year for our existing international networks — a 78 per cent reduction.
- Global Fabric's embodied emissions will be 363 tonCO₂e/year versus its existing international networks 2,185 tonCO₂e/year — an 83 per cent reduction. This gives a total carbon in use plus embodied carbon figure of 3,327 for Global Fabric versus its existing international networks at 15,781 tonCO₂e/year — a 78 per cent reduction.
- We estimate the average power consumption of Global Fabric will be 787 Watts per device, versus 2,201 Watts per device for its existing international networks. Moreover, Global Fabric will be built using 1,326 devices versus 1,571 devices — a 16 per cent reduction.

Offices Worldwide

The services described in this publication are subject to availability and may be modified from time to time. Services and equipment are provided subject to the respective British Telecommunications plc standard conditions of contract. Nothing in this publication forms any part of any contract.

© BT Communications Ireland Ltd Registered office: Grand Canal Plaza, Upper Grand Canal Street, Dublin 4. Registered in Ireland No. 141524