

CARBON SAVINGS OF CLOUD-BASED ENTERPRISE NETWORK FUNCTIONS

Maria Tunberg, Andrew Daly, Darryl Petch, Paidamoyo Mtutu, Jack Potter

September 2023

Contents

1.	Executive summary	3
2.	Introduction	7
2.1	Notes on our analysis approach	7
2.2	Moving network functions to the cloud	8
2.3	Greenhouse gas protocols	10
2.4	Structure of the document	11
3.	Enterprise network functions	12
3.1	Selected enterprise network functions	12
3.2	On-premises network functions used in the calculation	14
3.3	Equivalent Cloudflare products	15
3.4	Additional assumptions regarding the move of functions from on-premises to cloud	17
4.	Estimating the carbon savings of cloud-based network functions	18
4.1	[Calculation stage 1] Traffic demand	19
4.2	[Calculation stage 2] Power consumption of on-premises equipment	21
4.3	[Calculation stages 3-6] Power consumption of cloud-based functions	23
4.4	[Calculation stage 7] Power usage efficiency	24
4.5	[Calculation stage 8] Carbon intensity	26
5.	Results and conclusions	28
5.1	[Calculation stage 9] Carbon savings of moving from on-premises to Cloudflare	28
5.2	A summary of the key assumptions and caveats of the work	30
5.3	Overall conclusions	31
6.	Glossary of terms	32

Analysys Mason has produced the information contained herein for Cloudflare. The ownership, use and disclosure of this information are subject to the Commercial Terms contained in the contract between Analysys Mason and Cloudflare.

Analysys Mason Limited. Registered in England and Wales with company number 05177472. Registered office: North West Wing Bush House, Aldwych, London, England, WC2B 4PJ.

We have used reasonable care and skill to prepare this publication and are not responsible for any errors or omissions, or for the results obtained from the use of this publication. The opinions expressed are those of the authors only. All information is provided "as is", with no guarantee of completeness or accuracy, and without warranty of any kind, express or implied, including, but not limited to warranties of performance, merchantability and fitness for a particular purpose. In no event will we be liable to you or any third party for any decision made or action taken in reliance on the information, including but not limited to investment decisions, or for any loss (including consequential, special or similar losses), even if advised of the possibility of such losses.

We reserve the rights to all intellectual property in this publication. This publication, or any part of it, may not be reproduced, redistributed or republished without our prior written consent, nor may any reference be made to Analysys Mason in a regulatory statement or prospectus on the basis of this publication without our prior written consent.

© Analysys Mason Limited and/or its group companies 2023.

1. Executive summary

Cloudflare is seeking to quantify the potential reduction in carbon emission that could be achieved if on-premises enterprise network functions are replaced with Cloudflare's cloud-based alternatives.

As on-premises equipment tends to involve smaller-scale and under-utilised hardware infrastructure, moving such functions to the cloud has the potential to realise greater energy efficiency and thus reduce the carbon emissions associated with those functions.

Cloudflare has commissioned Analysys Mason to undertake an independent quantification of the potential carbon savings offered by its cloud-based enterprise network products.

Scope

The analysis focuses on 'Scope 2' emissions of carbon dioxide as defined in the greenhouse gas (GHG) protocol¹, i.e. those that are directly related to the generation of electricity required to power enterprise network functions.

Our analysis uses a combination of desk research, industry intelligence and information provided by Cloudflare. Information on the on-premises network functions has been sourced from datasheets published by equipment vendors. All of the information relating to cloud products has been sourced from Cloudflare, and therefore the results are also specific to Cloudflare. However, we have requested and analysed the information as independent and objective advisors, and would employ a similar approach with information obtained from other sources.

The network functions in scope are summarised in Figure 1.1, along with the equivalent Cloudflare products.

Network function	Equivalent Cloudflare product
Network firewall, including Intrusion Detection System (IDS), Intrusion Prevention System (IPS) and Virtual Private Network (VPN) concentrator	Magic Firewall (for network firewall and IDS/IPS) Magic WAN and Cloudflare Access (for VPN concentrator)
Load balancing	Cloudflare Load Balancing
Wide Area Network (WAN) optimisation	Argo Smart Routing
Software Defined WAN (SD-WAN)	Magic WAN
Secure Web Gateway (SWG)	Cloudflare Gateway
Web Application Firewall (WAF)	Cloudflare WAF
Distributed Denial of Service (DDoS) mitigation	Website DDoS Protection Spectrum: Application DDoS Protection Magic Transit: Network DDoS Protection

Figure 1.1: Network functions and equivalent Cloudflare products

[Source: Analysys Mason, 2023]

Different enterprises will need different numbers and combinations of these functions, depending on the size and role of the office and data centre locations, and the number of each.

¹ Please see section 2.3 for more details.

Methodology

We start the analysis by defining the amount of traffic carried by the network infrastructure that is in scope. We calculate the traffic to be served by each network function from each enterprise location at which employee- or customer-related network connectivity demand is created. We consider peak traffic demand – which allows us to quantify the required capacity of on-premises equipment – and average traffic demand – which allows us to allocate capacity on Cloudflare’s global platform of CPU servers.

We then translate the traffic requirements into an energy requirement, both for on-premises equipment and for a cloud-based alternative arrangement.

The capacity and power consumption performance of on-premises equipment is based on an extensive benchmark of information from public datasheets from equipment vendors. At least three vendors are benchmarked for each network function.

The power consumption for serving a unit of traffic via a cloud-based function is based on information from Cloudflare. Cloudflare provided a split of its CPU run time into shared ‘services’, an estimate of how these services are used by its products, and information on the number of customers and traffic served for each product.

The final components of the analysis include assumptions for the power usage effectiveness (PUE) of cloud data centres vs. on-premises data centres or data rooms, and the carbon intensity of electricity generation, based on the mix of fossil fuel vs renewable energy sources in the local grid.

We translate the traffic requirements into an energy requirement, both for on-premises equipment and for a cloud-based alternative arrangement

Results and conclusions

The results of our analysis for a large enterprise² are shown in Figure 1.2.




Figure 1.2: Breakdown of potential carbon savings from moving enterprise network functions from on-premises to Cloudflare products, large enterprise scenario
 (Source: Analysys Mason, 2023)

The analysis shows that moving the enterprise network functions considered in the analysis to Cloudflare products could reduce annual carbon emissions by 86%. There are a number of contributors to this saving, which are shown in the stages of the figure above, and explained here (note, further detail is given in the main body of the report):

Impact of cloud processing efficiency

We identify this impact by allocating the cloud power consumption according to the maximum traffic demand of the business rather than the average. This creates an illustrative result for the cloud functions with a similar level of utilisation as the on-premises functions, and so reveals how the two options compare in terms of processing efficiency. The increase in carbon emissions is because on a 'bit-for-bit' basis (i.e. a bit of information processed by an on-premises device vs. a bit of information processed by a cloud server), and for the traffic loads of a large businesses, on-premises devices are more efficient. This result is intuitive, as the hardware of on-premises devices are typically designed to perform a specific function (or set of functions), whereas cloud servers are generic to any function which is run on them at that time.

² Enterprise with 20 small offices, three medium-sized offices, one HQ and one small data centre with between 151 and 1500 employees

Utilisation gains from cloud

The largest gain in the reduction of carbon emissions comes from the increased utilisation of cloud infrastructure. On-premises equipment consumes power constantly but is only utilised for part of the day and part of the week. By contrast, cloud infrastructure is much more highly utilised, and has less wasted capacity. Cloudflare's global scale and aggregation of heterogeneous demand means that its infrastructure is used by thousands of businesses, and as one business's demand falls away, another business's demand will be picking up.

PUE gains from cloud

There is a modest but material gain from the hardware associated with the network function being located in a modern cloud data centre, instead of a on-premises data room or data centre. Modern cloud data centres make more efficient use of the power they draw from the grid, and so generate lower carbon emissions for a given amount of IT output.

Carbon intensity gains from cloud

The data centres that Cloudflare uses are located in cities, regions and countries where electricity generation generally has a lower carbon intensity than the global average. For an average business, this allows a further modest but material reduction in carbon emissions when on-premises network functions are transferred to Cloudflare.

The overall results (and their composition) vary for individual cases and circumstances, including the size of enterprise: the reduction in carbon emissions ranges from 96% for a small business to 78% for a very large business. For smaller businesses, the 'bit-for-bit' impact is more positive (cloud functions are actually more efficient than on-premises at the low traffic demand from smaller businesses, because of a linear relationship between usage and power for cloud functions) and the utilisation gains are also greater (because smaller businesses are more likely to have over-sized, under-utilised on-premises equipment).

The analysis includes a number of assumptions around how the enterprise manages its move to the cloud, and the wider impact on traffic carried over the associated national and international communications networks. If the specific circumstances of an enterprise do not align with these assumptions, then the overall savings in carbon emissions may be different to those calculated in this analysis.

Overall, the benefits of Cloudflare's cloud-based enterprise networking products, especially the utilisation benefits of a global platform of CPU servers, can make a material contribution to businesses being able to reduce their carbon impact, and contribute to wider sustainability goals.

2. Introduction



Sustainability initiatives are becoming an increasingly decisive aspect of operations across many industries. Driven by the rising cost of energy and increasing scrutiny of greenhouse gas (GHG) emissions, businesses are looking to maximise the efficiency of their operating models. A central theme of sustainability initiatives is the reduction of carbon emissions, which is widely supported by commercial and government policies targeting net-zero economies.

Moving network functions from hardware located on the business premises to cloud-based alternatives is one step that has the potential to offer substantial reductions in carbon emissions

Moving network functions from hardware located on the business premises to cloud-based alternatives is one step that has the potential to offer substantial reductions in carbon emissions. Understanding and measuring how this process can reduce emissions is challenging. There are complex issues associated with making a fair comparison between the typical power consumption of physical on-premises equipment (which are typically dedicated to a single enterprise) relative to cloud-based servers, which are used by thousands of different customers.

Analysys Mason has been commissioned by Cloudflare to conduct an independent analysis to quantify the carbon savings of moving certain enterprise network functions from on-premises equipment to the cloud. Specifically, this study estimates Scope 2 savings, i.e. the reductions in carbon emissions associated with direct consumption of electricity for the delivery of network services.

2.1 Notes on our analysis approach

Analysys Mason is an independent expert advisor to organisations and industry stakeholders in the telecoms, media and technology (TMT) sectors. We are respected worldwide for the quality and objectivity of our work.

We have produced this report for Cloudflare, based on an analysis of Cloudflare's products and Cloudflare's operational data. While the source of data for the cloud functions is almost entirely from Cloudflare, we have

¹ Analysys Mason, DataHub (2023)

independently requested the data required for the analysis and made our own assessment of how to use it in our calculations. While the results of the study are specific to Cloudflare, we would apply a similar approach if we were to perform the analysis for any other provider of cloud-based network functions.

2.2 Moving network functions to the cloud

Network functions can include management (such as switching, routing, acceleration, load balancing) and security (such as firewalls, gateways, and attack protection)

Enterprise networks are made up of different functions which allow the employees and/or customers of the enterprise to connect and conduct business. Network functions can include management (such as switching, routing, acceleration, load balancing) and security (such as firewalls, gateways, and attack protection). Network functions can be implemented in two main forms:

- 1. On-premises network functions** – these are pieces of hardware equipment that are installed at the offices or data centre of the enterprise. Each piece of hardware is typically dedicated to one or a small number of related functions, and also typically only serves the needs of a single enterprise.
- 2. Cloud-based network functions** – these functions are provided via software applications that are run on processing equipment which is agnostic to the function being run. The processing equipment (or server) is located in the cloud (i.e. in a data centre not owned by the enterprise). Each cloud processor may run different network functions for different enterprise clients at different times, depending on the need.

In this study, we consider the implications of moving enterprise network functions from on-premises to cloud-based equivalents. This concept is illustrated in Figure 2.1.

In this study, we consider the implications of moving enterprise network functions from on-premises to cloud-based equivalents




Figure 2.1: Network functions on-premises and moved to the cloud
[Source: Analysys Mason, 2023]

There are two ways in which the move to cloud-based functions are expected to reduce energy consumption:

- Moving network functions to the cloud may mean that the function benefits from the latest software architecture and processing infrastructure, and also from economies of scale, as the functions of one enterprise shares infrastructure with other enterprises.
- Moving network functions to the cloud allows the cloud function to only provide service (and consume power) when the function is needed. This is in contrast to on-premises equipment which is likely to be both:
 - ‘on’ and consuming power, even outside of office hours; and
 - of sufficient size to be able to accommodate a peak traffic load, which may occur relatively infrequently.

Electricity consumption is a key source of carbon emissions from the provision of network connectivity

2.3 Greenhouse gas protocols

The focus of this report is on calculating the equivalent carbon dioxide emissions – part of the greenhouse gas (GHG) emissions which include methane and nitrous oxide – associated with the consumption of electricity by certain enterprise network functions.

The framework used for calculating carbon emissions is the GHG Protocol. This protocol was created in 1998 at the World Business Council for Sustainable Development (WBCSD) and the World Resources Institute (WRI)³. It provides the standards for measuring and managing GHG emissions, and categorises the sources of emissions into three scopes:

- **Scope 1:** direct GHG emissions (e.g. from owned boilers, furnaces or vehicles).
- **Scope 2:** indirect GHG emissions from the consumption of electricity (which has created GHG in its generation).
- **Scope 3:** other indirect GHG emissions, such as from the production and transport of purchased materials and other GHG emissions embodied in the wider supply chain.

Electricity consumption is a key source of carbon emissions from the provision of network connectivity, and therefore this study considers only the Scope 2 emissions.

³Greenhouse Gas Protocol –
<https://ghgprotocol.org/about-us>

2.4 Structure of the document

The remainder of this document is laid out as follows:

- Section 3 describes the enterprise network functions included in this study, where they are typically located and the equivalent Cloudflare products.
- Section 4 explains our approach to calculating the carbon emissions from the network function devices, the typical power consumption of these devices, the power usage efficiency and the carbon intensity of the devices.
- Section 5 provides the results and conclusions of the report.

3. Enterprise network functions



3.1 Selected enterprise network functions

Network functions are the building blocks of electronic connectivity for businesses. We have summarised several network functions that are commonly used by businesses in Figure 3.1 below.

Network function	Description	Typical location(s)
Network management functions		
Router	• Routers are used to apply pre-defined logic to direct traffic from the source to the destination and vice versa. Routers are typically found at network ingress/egress points, connecting networks together and/or to the wider internet.	• Traditionally at least one router is found per enterprise location. Two or more may exist in critical locations to provide resilience/redundancy.
Switch	• Switches are used to connect device endpoints to a network and apply segregation where required by introducing virtual local area networks (VLANs). They are also used to aggregate segments of local area networks e.g. floors within a building.	• Traditionally at least one switch is found per enterprise location. Two or more may exist in critical locations to provide resilience/redundancy.
Load balancing	• An Load balancing device or service is used to balance loading of inbound/ outbound network traffic towards a specific platform, service or application⁴.	• Located at a central point where traffic needs to be balanced across multiple separate locations (i.e. a data centre).
WAN optimisation or WAN acceleration	• WAN acceleration employs a device between networks to improve data transfer efficiency by determining traffic types and prioritising / applying bandwidth and security profiles to traffic and making informed routing decisions to determine the most efficient path based on policies or other variables.	• Normally located within a central location, such as a datacentre or head office, in order to control in/ outbound traffic streams and prioritise / optimise data flows. As an example, it may appear in a datacentre to prioritise database traffic over other traffic types as this is deemed business-critical.

⁴ Note that the study considers only the Global Traffic Management (GTM) aspect of load balancing, as this function is offered by Cloudflare's Load Balancing product. Where a business uses Local Traffic Management (LTM), these are often delivered from separate equipment to GTM and are not in the scope of the calculation.

Network function	Description	Typical location(s)
Network management functions		
Software-defined wide-area networks (SD-WAN) platform	- SD-WAN is an approach to connectivity that allows traffic to be routed, controlled and secured via a centrally managed software platform. - Rules such as bandwidth-shaping algorithms, security features and routing decisions can be automated and driven on multiple devices/sites simultaneously from a centralised control function.	An SD-WAN platform will be centrally located in either a data centre or head office. It will however control SD-WAN-enabled devices 'in the field' which could be at any office/branch location that has a connection to the internet.
Network security functions		
Network firewall	- A network firewall is a device or service that protects networks and services from other networks and devices. - It uses rules or policies to determine what traffic is to be allowed or denied into or out of a network. As an example, a firewall may allow all users outbound to the internet, but allow no traffic inbound to protect sensitive data / servers.	Located at the network perimeter or at a sensitive network edge in order to control in/outbound traffic streams and protect users and systems.
Intrusion detection system (IDS) and intrusion prevention system (IPS)	- IDS and IPS are network functions that can detect and take appropriately determined action against specific network security threats or anomalies. - These functions are normally incorporated into network firewall devices, to bolster their effectiveness.	Normally located where sensitive datasets exist, whether this be a data centre or head office location that houses this data, or simply accompanying firewalls for added protection.
Virtual private network (VPN) concentrator	- VPN concentrators are used as the connection point for all remote access users to create a secure, tunnelled network connection to a network from anywhere at any time. - They are also used for site-to-site VPNs whereby a permanent connection between networks is established securely between the different geographic locations. They can terminate thousands of VPN connections at any one time.	Normally only required at a single location (i.e. in a datacentre or head office). However, smaller 'branch-sized' VPN concentrators may be utilised at smaller locations or sites, where permanent site-to-site VPN tunnels need to be formed back to the central VPN concentrator.
Web application firewall (WAF)	- A WAF device is used to explicitly protect web-based systems. - Whilst providing a traditional firewall-based service to protect these systems, these devices provide more granular policy options that can detect and take action on specific web-based protocols such as those aimed at web cookies or website scripting. - In contrast, traditional (network) firewalls cannot sense this level of protocol or application.	Normally located where sensitive web applications exist, whether a data centre or head office location.
Distributed denial of service (DDoS) mitigation	- DDoS mitigation functions specifically protect networks and services from this type of targeted threat. - Utilising sophisticated traffic sensors, the function is designed to deflect and/or clean network attack traffic arising from multiple locations simultaneously.	Typically located in data centres, at the network perimeter or at a sensitive network edge, in order to control in/outbound traffic streams and protect users and systems from DDoS attacks specifically.

Network function	Description	Typical location(s)
Network security functions		
Secure web gateway (SWG)	- An SWG is used to protect an organisation against security threats originating from the internet. It enforces company policy, sitting between the user and the internet, filtering web requests against the company policy and blocking malicious attacks and suspicious websites. - It prevents malware from entering the company's internal network and accommodates both in-office and remote workers.	Typically found at head office locations or data centres.

Figure 3.1: Summary of common enterprise network functions

[Source: Analysys Mason, 2023]

Some of the network functions are concerned with providing connectivity to an enterprise's employees, while others are concerned with providing connectivity to an enterprise's customers

3.2 On-premises network functions used in the calculation

Out of the list of network functions listed in Figure 3.1, many on-premises devices could be replaced with cloud equivalents.

However, some functions may not be replaced with cloud equivalents because an on-premises requirement will remain, even if all the others were moved to the cloud. These functions include routers and switches, which are required to aggregate traffic within the office location and provide a physical boundary between the office network and external network connectivity (including the global internet). Therefore, we do not consider routers and switches further in the calculation.

Another dynamic of on-premises network functions is that some are typically combined into a single device. The most common combination is:

- network firewall
- IDS/IPS
- VPN concentrator.

Therefore, we assume that the above three functions would all be provided via a single on-premises network firewall device⁵.

There are two dimensions that are useful to categorise the on-premises network functions:

- Some of the network functions are concerned with providing connectivity to an enterprise's employees, while others are concerned with providing connectivity to an enterprise's customers.
- Different types of on-premises function are likely to be required at different sizes and types of enterprise location.

⁵We note that some next-generation firewall devices also provide other functions, though as these are less commonly implemented, we do not include such combinations in the analysis.

A summary of the on-premises network functions, and their typical location, used in the model is shown in Figure 3.2⁶.

Type of connectivity	On-premises network function	Type of enterprise location				
		Small office	Mid-sized office	HQ office	Data centre (small)	Data centre (large)
Employee-facing	Network firewall, including IDS/IPS and VPN concentrator	✓	✓	✓	✓	✓
	Load balancing				✓	✓
	WAN optimisation			✓		
	SD-WAN			✓		
	SWG			✓		
Customer-facing	WAF				✓	✓
	DDoS mitigation				✓	✓

Figure 3.2: Summary of on-premises network functions used in the analysis and their typical location (Source: Analysys Mason, 2023)

3.3 Equivalent Cloudflare products

To make the comparison between on-premises and cloud-based network functions, we must choose which on-premises functions could be replaced by which of Cloudflare's products. Figure 3.3 shows the equivalent Cloudflare product for each of the network functions we have selected for the study.

Network function	Equivalent Cloudflare product	Description
Employee facing		
Network firewall, including IDS/IPS and VPN concentrator	Magic Firewall (for network firewall and IDS/IPS)	Cloud-native network firewall enforces consistent network security policies across the entire WAN, including headquarters, branch offices, and virtual private clouds. Magic Firewall for IDS/IPS is at Layer 3 and 4.
	Magic WAN (for VPN concentrator)	Magic WAN replaces legacy WAN architectures with Cloudflare's network, providing global connectivity, cloud-based security, performance, and control through a single user interface. Instead of sending all remote traffic through a single choke-point device (such as VPN concentrators at the perimeter of the corporate network), traffic is routed to the Cloudflare edge location closest to the source. Access policies are applied before that remote traffic is sent over optimal secure paths to its destination.
	Cloudflare Access (for VPN concentrator)	Cloudflare Access is a zero trust network access product that creates an aggregation layer for secure access to self-hosted, SaaS, or non-web applications.

⁶We note that some centralised employee-facing functions could also be provided from an enterprise-owned data centre. For clarity of the modelling approach, we have used the assumed locations shown in Figure 3.2.

Network function	Equivalent Cloudflare product	Description
Employee facing		
Load balancing	Cloudflare Load Balancing	Cloudflare Load Balancing distributes traffic across the servers, which reduces server strain and latency and improves the experience for end users. This function operates at Layer 4.
WAN optimisation	Argo Smart Routing	Argo Smart Routing detects real-time network issues and routes traffic across the most efficient network path. These benefits are most apparent for users farthest from the origin server. Argo Smart Routing operates at Layer 3.
SD-WAN	Magic WAN	As described above. This function operates at Layer 3.
SWG	Cloudflare Gateway	Cloudflare's secure web gateway keeps data safe from malware, ransomware, phishing, command & control, Shadow IT, and other Internet risks over all ports and protocols.
Customer facing		
WAF	Cloudflare WAF	Cloudflare WAF protects the customer website from SQL injection, cross-site scripting and zero-day attacks, including OWASP-identified vulnerabilities and threats targeting the application layer. Fully integrated with DDoS protection, Cloudflare WAF blocks millions of attacks daily, automatically learning from each new threat. Cloudflare WAF is at Layer 7.
DDoS mitigation	Website DDoS Protection	Cloudflare automatically detects and mitigates DDoS attacks using its autonomous edge. The autonomous edge includes multiple dynamic mitigation rules.
	Spectrum: Application DDoS Protection	Spectrum provides DDoS Protection at Layers 3-4 of the OSI model, that is against TCP- and UDP-based DDoS attacks.
	Magic Transit: Network DDoS Protection	The Cloudflare network-layer DDoS attack protection managed ruleset is a set of pre-configured rules used to match known DDoS attack vectors at Levels 3 and 4 of the OSI model.

Figure 3.3: Network functions and equivalent Cloudflare products

[Source: Analysys Mason, 2023]

3.4 Additional assumptions regarding the move of functions from on-premises to cloud

In our analysis, we make the following additional simplifying assumptions regarding the move of network functions from on-premises to the cloud:

- The enterprise will move all eligible network functions to the cloud, and not continue to run on-premises devices in parallel, as it seeks to derive the maximum available energy-saving benefit from such an initiative.
- The enterprise will not make the move materially before the natural end of life-cycle of its current on-premises equipment (as such a move could effectively contribute to an increase in Scope 3 life cycle or value chain embodied emissions). In practice, this may mean that different functions are migrated to the cloud at different times, as the useful lifetime of each piece of on-premises equipment comes to an end. This staged migration would mean that any carbon savings are realised more gradually than if all equipment is migrated at once.

4. Estimating the carbon savings of cloud-based network functions



This section describes the approach taken to calculate and ultimately compare the carbon emissions from on-premises and cloud-based network functions. Figure 4.1 below describes our methodology.




Figure 4.1: Overview of approach to compare carbon emissions of on-premises vs. cloud-based functions (Source: Analysys Mason, 2023)

Our approach calculates the power required for a unit of service provided by each function

4.1 Traffic demand [Calculation stage 1]

The starting point of the calculation is to define a consistent way to calculate the required size of the on-premises vs cloud-based functions, to ensure a fair comparison. Our approach calculates the power required for a unit of service provided by each function. We define the unit of service as being the relevant data traffic served by each function, in Gbit/s.

There are two dimensions to be considered at this stage:

- peak vs average traffic
- employee-generated vs customer-generated traffic.

4.1.1 Peak vs average traffic

Our approach models the carbon emissions from the perspective of an individual business (including results for small, medium and large businesses), and is based on an assessment of the on-premises equipment needed to support the network functions. The equipment must be able to serve the needs of the business throughout the fluctuating demands of the working day. As such, on-premises equipment should have a size such that it can accommodate a pre-defined *peak traffic* load.

There is limited public data on how businesses estimate the size of their connectivity requirements for peak traffic. Our industry intelligence suggests that the peak connectivity needs for *employee* business processes – web browsing, email, video conferencing, accessing software-as-a-service (SaaS) applications etc. – equate to around 5Mbit/s per business employee.

For cloud-based functions, the analysis is different, because of the way that the infrastructure is utilised. When serving thousands of enterprises, the peaks and troughs of individual customers are smoothed to a collective average, and we understand that the total traffic on the Cloudflare global network does not vary materially according to time of day. Our approach for calculating the cloud-based carbon consumption is therefore based on the enterprise's *average traffic* load, which will be met by allocating a proportion of the total global service provided by Cloudflare. Measures of the total traffic carried on the Cloudflare network are the sum of the average traffic demand across its customers, with differing activity profiles, and in different time-zones. Therefore, to allocate a portion of the global Cloudflare service to an individual business, we use average traffic.

As with peak traffic, there is limited public data on average traffic demand for businesses. Drawing on the small amount of information available, we estimate that average current business data consumption is around 50–150GB per business employee per month. If we take a value at the top of this range (150GB) to reflect the fact that businesses considering cloud services will tend to be heavier data users, and to give a more conservative estimate of carbon emissions savings, this equates to an average bandwidth demand of around 0.5Mbit/s, or around one tenth of the level of peak data requirements.

4.1.2 Employee vs customer traffic

As shown in Figure 3.2 above, some of the functions considered in our analysis are used to serve a business's employees, while others are used in relation to a business's customers. The estimates in the previous sub-section refer to employee related functions.

There is no simple relationship between the size of a business (in terms of staff numbers) and the traffic associated with its customer base. Customer-related traffic is strongly affected by nature of the services supplied to customers, as well as a range of factors, including the age and success of the business.

However, in order to derive a value for the analysis, we have assumed that the *customer-related traffic* is equivalent to the total of the *employee-related traffic*. This is a simplified approach, but what is most important is that there is a fair comparison between on-premises and cloud-based functions that are customer facing.

4.1.3 Summary of traffic demand assumptions

There is huge variation between enterprises in terms of size and composition, number of employees and their distribution across their sites, as well as the intensity of data usage. Network connectivity needs differ substantially, both in terms of the traffic they create, and the network functions they require. In order to be able to advance an analysis of these multifarious entities, we have defined four generalised scenarios for different sizes of business, and their respective office distribution, as shown in Figure 4.2.

	Small office	Mid-sized office	HQ office	Data centre (small)	Data centre (large)	Total employees
Employees per location	10	100	1000			10 150 1500 7500
Number of locations by business size scenario						
Small single office	1	0	0	0	0	
Medium enterprise	5	1	0	0	0	
Large enterprise	20	3	1	1	0	
Very large enterprise	50	20	3	0	1	

Figure 4.2: Summary of business size scenarios (Source: Analysys Mason, 2023)

A calculation of the peak bandwidth demand per location (and thus an estimate of the needs of on-premises equipment) is achieved by combining the scenarios in Figure 4.2 with the assumptions on traffic demand defined above. Peak bandwidth demand per location is shown in Figure 4.3.

Location type	Employees served ⁷	Peak traffic per location (Gbit/s)
Small office	10	0.05
Mid-sized office	100	0.50
HQ office ⁸	1000	5.00
Data centre (small)	1500	7.50
Data centre (large)	7500	37.50

Figure 4.3: Summary of peak bandwidth demand per location for estimating the size of on-premises equipment (Source: Analysys Mason 2023)

When the equivalent cloud functions are estimated, the peak bandwidth requirement is converted to an average bandwidth requirement, before being aggregated across each of the locations where a network function is required.

4.1.4 Additional assumptions regarding traffic demand

In our analysis, we make the following additional simplifying assumptions when moving network functions to the cloud:

- Enterprises will not materially change their traffic demands or behaviour once their network functions are being provided by a third party.
- The move in the location of functions will not create material additions in the distance or number of connections over which the enterprise traffic will travel.

4.2 Power consumption of on-premises equipment [Calculation stage 2]

In this section, we explain our approach to calculating the power consumption of on-premises equipment, including the source of information, the calculation of power consumption itself, and some additional assumptions.

4.2.1 Source of information of capability and power consumption of on-premises equipment

Our calculation of the power consumption of on-premises equipment is based on an extensive benchmarking exercise of the actual performance of current on-premises equipment. We have gathered the information we need from the data sheets published by equipment vendors. For each of the on-premises functions in scope, we collected:

- The maximum throughput of the equipment. This is commensurate with the idea that on-premises equipment must be of sufficient size to meet a peak capacity measure. Where multiple measures of throughput are given, we chose the measure that is typically used for determining the required size of the equipment in question.
- The average power consumption of the equipment. The average power consumption accounts for the fact that these piece of equipment tend to be always 'on', even outside working hours⁹. Where the average power consumption was not provided by a particular vendor, we estimated the

⁷ Note: for customer facing functions such as WAF and DDoS located at data centres, we make the simplifying assumption that the customer traffic is similar to the employee facing traffic

⁸ Note: for functions located only at an HQ office, we assume that statistical multiplexing gains are such that a peak capacity defined to serve the employees at the HQ location is sufficient to also serve employees at other locations

⁹ Some electronic communication equipment has the capability to switch into a low-power or 'sleep' mode during times of low usage. However, this functionality is not found in the types of on-premises enterprise network equipment being considered in this analysis.

average power consumption from the peak power consumption, using the ratio of the peak to average consumption from other vendors of equivalent equipment.

For each network function, we collected information from at least three vendors, to ensure that our analysis is not skewed by the performance of one particular vendor.

However, relying on the data collected to derive values we can use in our analysis presents certain challenges. The profile of capacity vs. power varies between vendors, and in some cases, between the ‘series’ of equipment from a given vendor. Further complexity lies in the fact that it is sometimes more power-efficient to run a higher-capacity piece of equipment than is strictly needed.

In order to create a conservative result for the calculation of carbon savings, we typically identified devices with the smallest capacity that can meet the required throughput. We also applied an averaging approach to create a smooth profile for the relationship between maximum throughput capacity and average power consumption that could be used in the model. A summary of the results of the benchmarking exercise is shown in Figure 4.4.




Figure 4.4: Benchmark of relationship between maximum capacity and average power consumption for on-premises network equipment

(Source: Analysys Mason, 2023)

The analysis shows that power consumption tends to grow strongly with capacity at low throughputs, but tends to plateau and benefit from economies of scale at higher capacities¹⁰.

¹⁰ The relatively high power consumption of WAN optimisation is likely due to a lack of development in these types of equipment (as they are being replaced by inherent system- and application-level efficiencies included in modern networks). Where these types of equipment are deployed on premises, they appear to be a good candidate for replacement by a more energy-efficient solution.

4.2.2 Calculating the power consumption of on-premises equipment

Our analysis chooses the right on-premises function at the right capacity for each enterprise location, based on the assumptions stated above. For each on-premises function, the average power consumption is calculated based on the analysis in the previous section. The total power consumption of the enterprise is summed across all the on-premises network functions that it would use.

4.2.3 Additional assumptions regarding power consumption of on-premises functions

In our analysis, we make the following additional simplifying assumptions regarding the move of network functions from on-premises to the cloud:

- The on-premises equipment is on all of the time, 24 hours a day, 7 days a week.
- The average power ratings given by vendors are blended across the whole day and whole week.

4.3 Power consumption of cloud-based functions

[Calculation stages 3 - 6]

Our approach to calculating the power consumption of cloud-based functions is based around allocating a portion of the total power consumption of Cloudflare's CPU servers to the portions of Cloudflare's products used by an individual business.

We calculate the power consumption of each of Cloudflare's products (based on its CPU usage), then divide that power consumption by the total traffic served for that product, to create a power consumption for a unit of traffic served. The approach and data sources used in each component of the calculation are shown in Figure 4.5.

Calculation component	Approach and source data
Total power consumption of Cloudflare CPU servers	Cloudflare provided data on the total number of CPU servers in its network, and the power consumption of each server. The data was provided for three different generations of server. Cloudflare also provided power information on its supporting routers and switches.
Split of CPU resources between Cloudflare's products	Cloudflare's CPU resources are broken down into 'CPU services'. While a small number of CPU services are specific to externally-facing products, most CPU services are shared by multiple products. Cloudflare provided: • a breakdown of the total CPU resources into different CPU services, based on an internal hardware reporting dashboard • an estimate of how different products use different proportions of the CPU services, based on internal data on the number of CPU cores used by each product. In addition, we applied a series of 'equi-proportionate mark-ups' to allocate some shared CPU services to various products: first to Layer 7 products, and then to all products.
Total traffic by each of Cloudflare's products	Cloudflare provided internal dashboard data on the traffic passing through its network. This was combined with information on the take-up of products amongst its customers, and additional assumptions to estimate the total traffic served for each of Cloudflare's products.

Figure 4.5: Summary of approach and data sources for power consumption of cloud-based functions (Source: Analysys Mason, 2023)

We used discussions with Cloudflare’s product and infrastructure experts to sense-check the outputs and ensure the right relativities between the products

We used the above data sources to allocate a share of Cloudflare’s power consumption to a unit of traffic served for each product. We also used discussions with Cloudflare’s product and infrastructure experts to sense-check the outputs and ensure the right relativities between the products. A summary of the unit power consumption for each of the Cloudflare products in scope is given in Figure 4.6.




Figure 4.6: Summary of calculated relationship between traffic served and power used for Cloudflare’s enterprise networking products
[Source: Analysys Mason, 2023]

The relationship between traffic and power for Cloudflare’s products is much more linear than is the case for on-premises equipment. This relationship is a function of the Cloudflare network being driven by a very large number of well utilised CPU servers, which are supporting a wide number of customers. The relationship is also commensurate with the way that Cloud services are provisioned: as more capacity is required, more CPU servers can be added, each of which can support a range of products.

4.4 Power usage efficiency [Calculation stage 7]

Power usage efficiency (PUE) is a ratio used to determine how efficiently a data centre/server/telecoms room uses energy. The concept was introduced in 2007 by the Green Grid (TGG), an industry initiative and affiliate of the Information Technology Industry Council (ITI)¹¹. TGG recognised the impact of increasing power costs on operators and recognised that the pace of customer demand for advanced computing was increasing at a higher rate than the availability of sustainable energy sources. The focus of the TGG is to advocate for energy and resource use efficiency in the data centre ecosystem, to create the tools and to provide the expertise to fulfil this.

¹¹ The Green Grid - <https://www.thegreengrid.org/>

PUE is the industry standard for measuring the energy efficiency of data centres

4.4.1 Understanding PUE

PUE is calculated by dividing the overall data room or data centre electricity consumption by the electricity required to run IT equipment. According to the Uptime Institute, PUE is the industry standard for measuring the energy efficiency of data centres and is used to track the progress of data centre efficiency over time¹². The target PUE for data centres is 1.00 with higher numbers indicating inefficiencies in energy uses usually due to cooling requirements. Since its establishment in 2007, the global average PUE has fallen from 2.50 to 1.55 in 2022. This is attributed to the growing number of state-of-the-art facilities which are close to achieving the 1.00 target.




Figure 4.7: Understanding the PUE calculation (Source: Analysys Mason, 2023)

There are various factors which affect the PUE of a data centre, including its utilisation, the efficiency of its cooling systems, and the efficiency of the IT equipment contained within it. Age and design has an overall impact: more recent data centres are typically more efficient (and more easily upgradeable to keep up with new efficiency developments).

4.4.2 Benchmark values for PUE and use in the model

In the calculation, we use the PUE to multiply the power consumption of the network function(s) to the power actually consumed from the grid. A benchmark of PUE values is shown in Figure 4.8.

¹² Uptime Institute Global Data Centre Survey 2022 - <https://uptimeinstitute.com/resources/research-and-reports/uptime-institute-global-data-center-survey-results-2022>




Figure 4.8: Benchmarks of PUE¹³ [Source: Analysys Mason, 2023]

While there are examples of datacentre PUEs getting close to 1.00, for the cloud functions we use the global average data centre PUE value of 1.55. This is commensurate with the fact that Cloudflare does not own its own datacentres, and works with a range of datacentre partners around the world. For the on-premises results, we use an estimate of 2.00¹⁴, as on-premises data rooms and data centres are typically less power-efficient than modern large-scale data centres.

4.5 Carbon intensity [Calculation stage 8]

In order to calculate the carbon emissions from electricity usage, the carbon intensity multiplier is required. Carbon intensity is a measure of the amount of carbon dioxide equivalent released per kilowatt hour of electricity produced by the main electricity generation grid, usually in a specific country. This varies depending on the energy source used in the generation of electricity (i.e. the mix of renewables vs. fossil fuels). Therefore the location of the electricity consumption affects the associated carbon emissions. For the carbon intensity multipliers included in the analysis, we used different values for on-premises and cloud:

- For the on-premises calculation, we used the global population-weighted average carbon intensity of 445gCO₂/kWh¹⁵. This choice is based on the assumption that an average company receives an average carbon intensity, or alternatively, that a business considering a move of their network functions to cloud services could be based anywhere in the world. We considered whether the fact that a business considering moving their on-premises network functions to the cloud would mean that the business

¹³ 3JcMQ [minkels.com] U.S. facility may have best data center PUE (techartest.com Sustainability - Lefdal Mine Datacenter) Alibaba : AliCloud : Launches New Energy-Efficient Qiantao Lake Data Center | MarketScreener Rethinking data center design for Singapore - Engineering at Meta (fb.com) Green Data Centers Around the World 2021 | Sunbird DCIM; Microsoft, Google data centres

¹⁴ <https://www.minkels.com/files/3JcMQ>

¹⁵ Carbon intensity of electricity: <https://ourworldindata.org/grapher/carbon-intensity-electricity>

was more likely to be located in more developed countries. Therefore we also calculated a carbon intensity average which excluded many less developed countries¹⁶, and as an alternative, the carbon intensity of OECD countries. As shown in Figure 4.9 below, these two additional values sit either side of the global average, suggesting that more developed countries do not necessarily have systematically higher or lower carbon intensities than the global average, so we have used the global average for simplicity

- For the cloud calculation, we used a Cloudflare-specific value of 340gCO2/kWh. This is based on analysis of the energy usage at each of Cloudflare’s actual locations¹⁷ around the world, and the local carbon intensity of generation at that location, to create a weighted average carbon intensity specific to Cloudflare.




Figure 4.9: Comparison of 2022 average carbon intensity values (Source: ourworldindata.org, Analysys Mason, 2023)

It should be noted that the figures above represent a blended average. The exact carbon intensity of electricity generation in the location of an enterprise transferring its network functions to Cloudflare may be higher or lower than the global average, and thus the individual resultant carbon savings may be different to those shown in our analysis.

¹⁶ We excluded countries with a GDP per capita below USD8000 at Purchasing Power Parity

¹⁷ Cloudflare uses carbon intensity which is state-specific for locations in the USA, but based on country averages for non-USA locations.

5. Results and conclusions



5.1 Carbon savings of moving from on-premises to Cloudflare [Calculation stage 9]

The results of our analysis for a large enterprise are shown in Figure 5.1.




Figure 5.1: Breakdown of potential carbon savings from moving enterprise network functions from on-premises to Cloudflare products, large enterprise scenario (Source: Analysys Mason, 2023)

The analysis shows that moving the large-enterprise network functions considered in the analysis to Cloudflare products could reduce annual carbon emissions by 86%. There are a number of contributors to this saving, which are shown in the stages of the figure above, and explained in more detail here:

Impact of cloud processing efficiency

We identify this impact by allocating the cloud power consumption according to the maximum traffic demand of the business rather than the *average*. This creates an illustrative result for the cloud functions with a similar level of utilisation to the on-premises functions, as so reveals how the two options compare in terms of processing efficiency. The increase in carbon emissions is because on a 'bit-for-bit' basis (i.e. a bit of information processed by an on-premises device vs. processed by a cloud server), and for the traffic loads of a large businesses, on-premises devices are more efficient. This result is intuitive, as the hardware of on-premises devices are designed to perform a specific function, whereas cloud servers are generic to any function which is run on them at that time.

Utilisation gains from cloud

The largest gain in the reduction of carbon emissions comes from the increased utilisation of cloud infrastructure. On-premises equipment consumes power constantly but is only utilised for part of the day and part of the week. By contrast, cloud infrastructure is much more highly utilised, and has less wasted capacity. Cloudflare's global scale and aggregation of heterogeneous demand means that its infrastructure is used by thousands of businesses, and as one business's demand falls away, another business's demand will be picking up.

PUE gains from cloud

There is a modest but material gain from the hardware associated with the network function being located in a modern cloud data centre, instead of a on-premises data room or data centre. Modern cloud data centres make more efficient use of the power they draw from the grid, and so generate lower carbon emissions for a given amount of IT output.

Carbon intensity gains from cloud

The data centres that Cloudflare uses are located in cities, regions and countries where electricity generation generally has a lower carbon intensity than the global average. For an average business, this allows a further modest but material reduction in carbon emissions when on-premises network functions are transferred to Cloudflare.

Different sizes of enterprise are likely to realise different savings, as shown in Figure 5.2.




Figure 5.2: Potential carbon savings from moving enterprise network functions from on-premises to Cloudflare products, different scenarios for enterprise size [Source: Analysys Mason, 2023]

Across the scenarios, the impact of PUE and carbon intensity is similar. For smaller businesses, the reduction is much larger than for larger businesses. This is because for smaller businesses, the ‘bit-for-bit’ impact is more positive (cloud functions are actually more efficient than on-premises equipment at low demand, because of the linear relationship between usage and power for cloud functions) and the utilisation gains are also greater (because smaller businesses are more likely to have over-sized, under-utilised on-premises equipment).

5.2 A summary of the key assumptions and caveats of the work

The results of the analysis show that an enterprise could potentially reduce the carbon emissions associated with some of its network functions, by moving from existing on-premises equipment to Cloudflare’s products. However, it is worth reiterating the key assumptions we have made in the analysis. Our assessment assumes that:

- The enterprise will move all eligible network functions to the cloud, and not continue to run on-premises devices in parallel.

- The enterprise will not make the move materially before the natural end of life cycle of its current on-premises equipment.
- Enterprises will not materially change their traffic demands or behaviour once their network functions are provided by a third party.
- The move in the location of functions will not create material additions in the distance or number of connections over which the enterprise traffic will travel.
- The on-premises equipment is on all of the time, 24 hours a day, 7 days a week, and that the average power ratings given by vendors are blended across the whole day and whole week.
- The enterprise is located in a region with a carbon intensity of electricity generation similar to the global average.

If the specific circumstances of an enterprise deviate from these assumptions, the precise savings in carbon emissions may differ from those calculated in this analysis.

5.3 Overall conclusions

Overall, the analysis has shown that a move of certain enterprise network functions from on-premises equipment to Cloudflare's products could create a 78–96% reduction in the associated carbon emissions. While the analysis includes a number of assumptions (and therefore may have scope for refinement in the future), the results are encouraging.

The benefits of Cloudflare's cloud-based enterprise networking products, especially the utilisation benefits of its global platform of CPU servers, can make a material contribution to businesses being able to reduce their carbon impact, and achieve wider sustainability goals.

A move of certain enterprise network functions to Cloudflare's products could create a 78–96% reduction in the associated carbon emissions

6. Glossary of terms

CEE	Central and Eastern Europe
CO₂	Carbon dioxide
CPU	Central processing unit
DDoS	Distributed denial of service
GHG	Greenhouse gas
GTM	Global Traffic Management
HQ	Headquarters
IDS	Intrusion detection system
IPS	Intrusion prevention system
ITI	Information Technology Industry Council
kWh	Kilowatt hour
Large enterprise	Enterprise with 20 small offices, three medium offices, one HQ and one small data centre with between 151 and 1500 employees
LTM	Local Traffic Management
Medium enterprise	Enterprise with five small offices and one medium sized office with between 11 and 150 employees
MENA	Middle East and Africa
OCED	Organisation for Economic Co-operation and Development
OSI	Open Systems Interconnection model
OWASP	Open Worldwide Application Security Project

PUE	Power usage efficiency
SaaS	Software as a service
Scope 2	Carbon emissions that are directly related to the generation of electricity required to power enterprise network functions
SQL	Structured Query Language
SD-WAN	Software-defined wide-area networks platform
Small enterprise	Enterprise with one small office and up to 10 employees
SWG	Secure web gateway
TCP	Transmission control protocol
TGG	The Green Grid
UDP	User datagram protocol
Very large enterprise	Enterprise with 50 small offices, 20 medium offices, five HQ offices and one large data centre with between 1501 and 7500 employees
VPN	Virtual private network
WAF	Web application firewall
WAN	Wide-area network
WBCSD	World Business Council for Sustainable Development
WRI	World Resources Institute

7. About the authors



Maria Tunberg (Partner).

Maria is a Partner at Analysys Mason and head of the Lund office in Sweden. She has extensive theoretical and practical experience of organisational change management. At Analysys Mason, Maria manages projects for several authorities and municipalities, including the Nordic Council of Ministers and the Swedish Agency for Economic and Regional Growth. Her expertise lies in digital transformation, both at sector and firm level, and in smart-city applications. She is especially skilled in issues located in the intersection between digitalisation and sustainability.



Andrew Daly (Principal).

Andrew is a Principal at Analysys Mason. He joined Analysys Mason in 2007. Andrew is an expert in the capabilities of next-generation fixed and wireless broadband networks, and the associated commercial, economic and policy implications. Andrew advises operators, regulators, local authorities, and governments on these issues. Andrew has extensive experience of working in the UK, and on-site in many other countries around the world, covering Europe, Africa, the Middle East, South America and Australasia. Andrew delivers projects which provide a wide range of analyses for his clients, including network modelling, cost modelling, market modelling, stakeholder interviews and primary research, benchmarking, defining regulations, audit, production of expert reports, due diligence, and strategic recommendations.



Darryl Petch (Associate).

Darryl is an expert in enterprise network technical solutions. Darryl's track record includes delivering network solutions for businesses, including cloud hosting, cyber security, DDoS protection, virtualisation and WAN/LAN network design. His expertise covers hardware and software support and management in an IT environment.



Paidamoyo Mtutu (Consultant).

Paida is a Consultant at Analysys Mason. She has worked on a wide-range of projects from commercial and technical due diligence projects to supporting infrastructure strategy projects. She has helped deliver projects for a range of customers including operators, regulators, governments and private equity firms. In addition, Paida has supported multiple utility clients in identifying and understanding the communications technology available in deploying smart grid networks.



Jack Potter (Associate Consultant).

Jack recently joined Analysys Mason and has worked on range of projects, including appraisals of technical evidence for fibre networks, and market research and analysis supporting network investments. Jack has a sustainability background by both qualification and employment experience. Jack was previously involved in the design and implementation of a major carbon reduction project for Pernod Ricard whisky distilleries, for which the designs have subsequently been made open source.



Stay connected

You can stay connected by following Analysys Mason via LinkedIn, Twitter and YouTube.



[linkedin.com/company/analysys-mason](https://www.linkedin.com/company/analysys-mason)



[@AnalysysMason](https://twitter.com/AnalysysMason)



[youtube.com/AnalysysMason](https://www.youtube.com/AnalysysMason)



[analysismason.podbean.com](https://www.podbean.com/analysismason)