White paper

The handy guide to enterprise network planning

How to optimise your network for the world of hybrid work

Author: Andy Linham
Principal Strategy Manager for Fixed Connectivity, Vodafone Business

Together we can
vodafone business

Contents

1. Introduction

98% of workers prefer remote working at least some of the time[1]. We have seen a huge growth in hybrid working in the past years, making cybersecurity a major concern as employees work from anywhere and on any device. And though companies encourage the return to the office to drive collaboration, working from home is 5 times more common than 5 years ago[2].

You need to ensure that you understand the implications when deciding how your network needs to evolve to better serve your business.

In this whitepaper, we'll consider the current challenges that organisations are grappling with. These include managing a hybrid of remote- and office-workers, growth in bandwidth and site occupancy levels that fluctuate wildly on a daily basis. We'll then provide some context to help overcome these challenges, looking at how technology can deliver the flexibility that you need to stay ahead.

[1] https://www.forbes.com/advisor/business/remote-work-statistics
[2] https://www.forbes.com/sites/karadennison/2024/01/24/how-the-flexible--remote-work-debate-will-carry-into-2024/

2. Executive summary

There are four key areas to consider when you're planning for the next stage in your network's life:

1. Applications

What applications are you running, and where are they located? Do they have specific performance metrics that impact how they run, or do they have a material impact on your users' ability to be productive? You should categorise your applications into performance tiers, to simplify the task of planning your network.

2. Users

Where are your users going to be based? Will they be flexible, based entirely on-site or working remotely? When you have a view on this, you can map them into personas that reflect their location, the applications they need access to, and the services they consume. These personas, combined with the performance tiers of the applications, give you everything you need to inform the final question. Service providers have expertise in propagation. They can provide appropriate coverage in complex environments and access to external public and private wide-area networks, such as SD-WAN. Service providers can also offer Wi-Fi to meet ad hoc service requirements, such as assisting visitors.

3. Sites

What does your network need to look like tomorrow? Can you move to an Internet-only approach, or do you need to retain some private connectivity? Do you move to edge-based security solutions or do you still have a critical mass of information in a private site? These questions can help you to define what you need from your network when users return, en masse.

4. Security

How can you create an environment protected at the network, application and device layers? Providing cloud and network security, trusted remote access for your users and third parties, as well as secure Internet access are all essential when thinking about cybersecurity. Vodafone Business Secure Access Service Edge (SASE) unifies network and security services for a consistent experience across all of your locations.

The shift to the hybrid world means that an organisation's network needs to be as flexible as possible. Technologies such as SD-WAN and cloud computing are there to support this level of flexibility and help you face the changing world with confidence.

3. Putting things in context

Since the pandemic, we've seen a fundamental shift in our ways of working, with hybrid working very much becoming the norm. Businesses have seen substantial change in how colleagues communicate with each other and how they interact with applications.

In this whitepaper, we'll look at three key contextual changes that will have a significant impact on the network. We'll then consider how you address each of these factors in your decisions moving forwards.

Understanding the hybrid world's networking paradigm

Even before the shift to hybrid work, we were seeing a rise in Internet traffic for business services. This has fundamental implications for the enterprise network, which we'll look at now. As an industry, we spent a very long time designing insular networks. That meant that the networks were focussed inward and not outwards. There was a strong perimeter that formed a key part of an organisation's security barrier; think of it like building a moat around the outside of a castle to keep invaders out. Sites within the network weren't usually allowed to communicate with anything outside of the moat. This approach is shown in the following diagram:

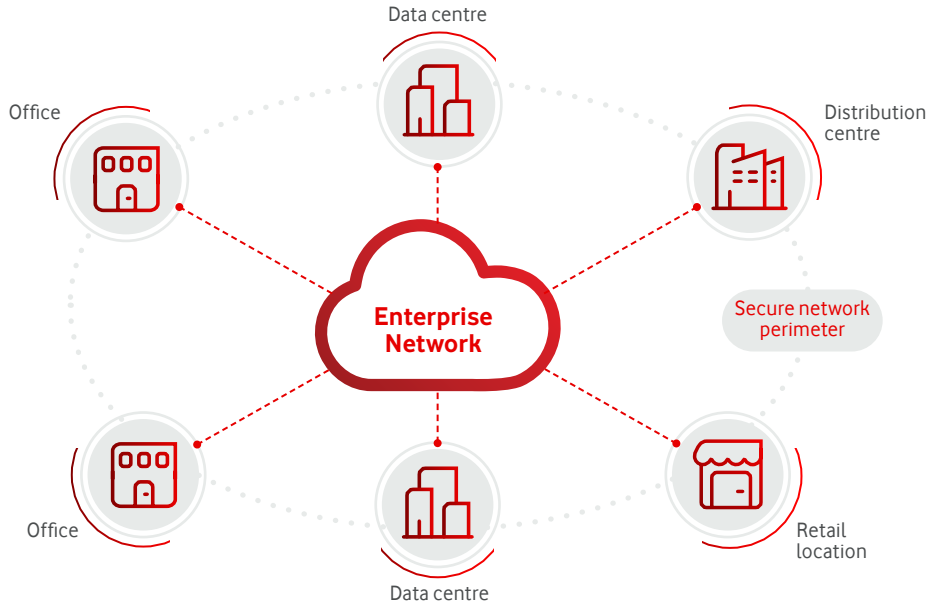

Figure 1 - the traditional enterprise network up until 2010

3. Putting things in context

As an industry, we used routers as the demarcation point between the Wide Area Network (the WAN) and the Local Area Network (the LAN). The router might have some basic protective functions enabled like Access Control Lists (ACLs) but that was often as complex as the on-site security went.

Businesses put their trust in two things:

1. The physical security of their premises preventing people from gaining access to their network.

2. Their service provider preventing attackers from getting onto the core network between their sites.

This model evolved after 2010 to incorporate a significant amount of Internet traffic. This traffic was primarily targeted at cloud services from hyperscale providers such as Amazon, Microsoft, and Google.

The revised network architecture can be seen in the graphic on the right.

The key point to take away from the diagram is that the cloud applications sit outside of the network perimeter. This means we have to think differently about how we access those applications and how we secure them.

The 'secure gateway' is a new addition to our diagram and it's a vitally important component. The gateway might be a single firewall with a policy that prevents anyone from getting in but lets the right sort of traffic out. It could also contain proxy servers, intrusion detection and prevention appliances and a host of other security devices. The size and scale of the gateway was specific to each organisation, but they all had one thing in common – they were capitally intensive purchases.

Organisations needed to buy the boxes themselves, then outfit them with licences and a multi-year support contract. Also, for all but the largest or most tech-savvy organisations, they would need a third-party security company to manage and maintain the gateway. Throw in the fact that the gateway needed a large Internet circuit to connect to the cloud applications and the costs just keep on coming.

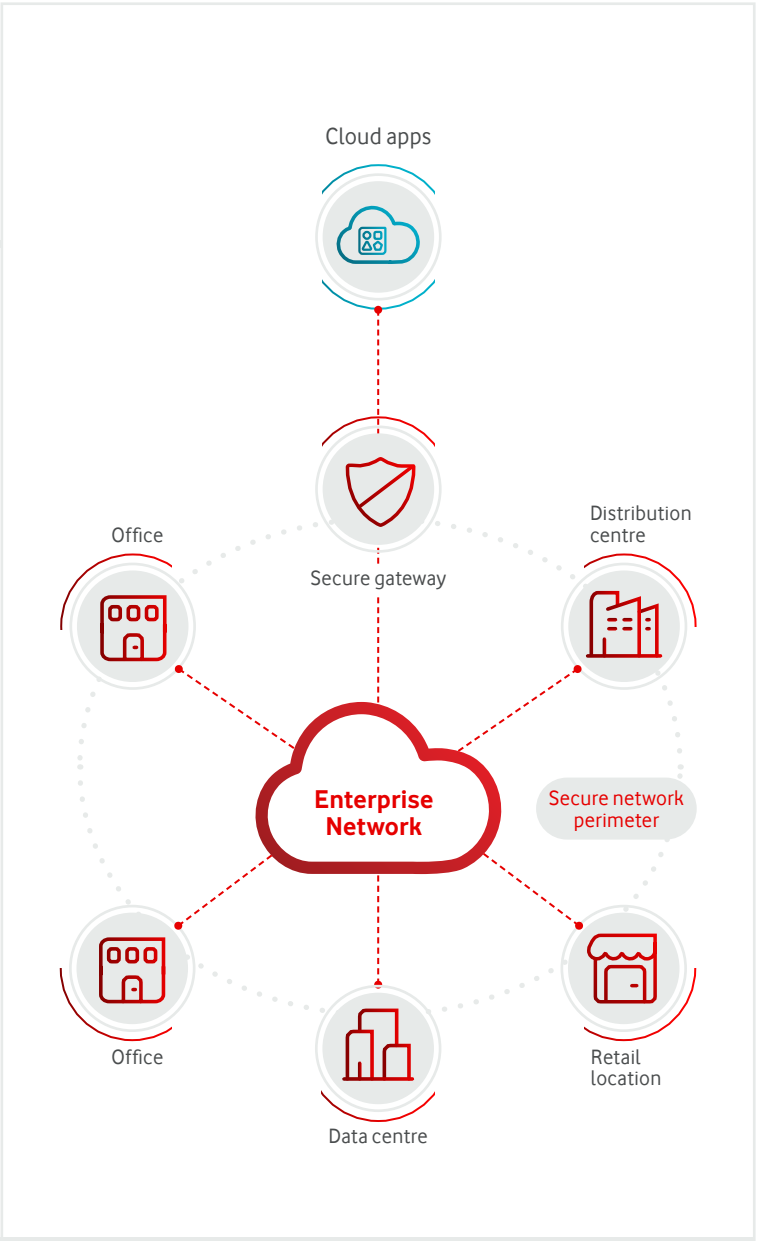

Figure 2 - the emergence of cloud services starts to shift business data outside of the network perimeter

3. Putting things in context

Despite some of the challenges, the Internet has continued to evolve and is now a relatively robust, reliable network. Performance is also acceptable for large percentages of business applications. The combination of improved resilience and speed has seen lots of sites shift to Internet connectivity. Our data shows that about 14% of Internet connections in FY24 were SD-WAN, the rest were stand-alone. We're forecasting this to shift to 25% by FY28. This new network, with a balance of Internet and private circuits is shown in the diagram.

There are challenges too though – not least of which is the security of the site. If you think back to our secure gateway, it can contain lots of equipment with a common purpose: keeping the information in and the attackers out. Replicating all of this functionality for each site isn't practical so other solutions need to be found. SD-WAN is a common answer in use today. Most SD-WAN solutions can provide advanced firewalls and data encryption to mitigate a high proportion of the security threats.

Secure Access Service Edge, SASE, is an even more secure version of SD-WAN and is becoming more popular as it delivers network and security under one solution, securing your cloud, your remote workers, third-party access and the Internet. It's a reliable solution that simplifies your IT team's workload and gives you granular visibility into what's happening in your network.

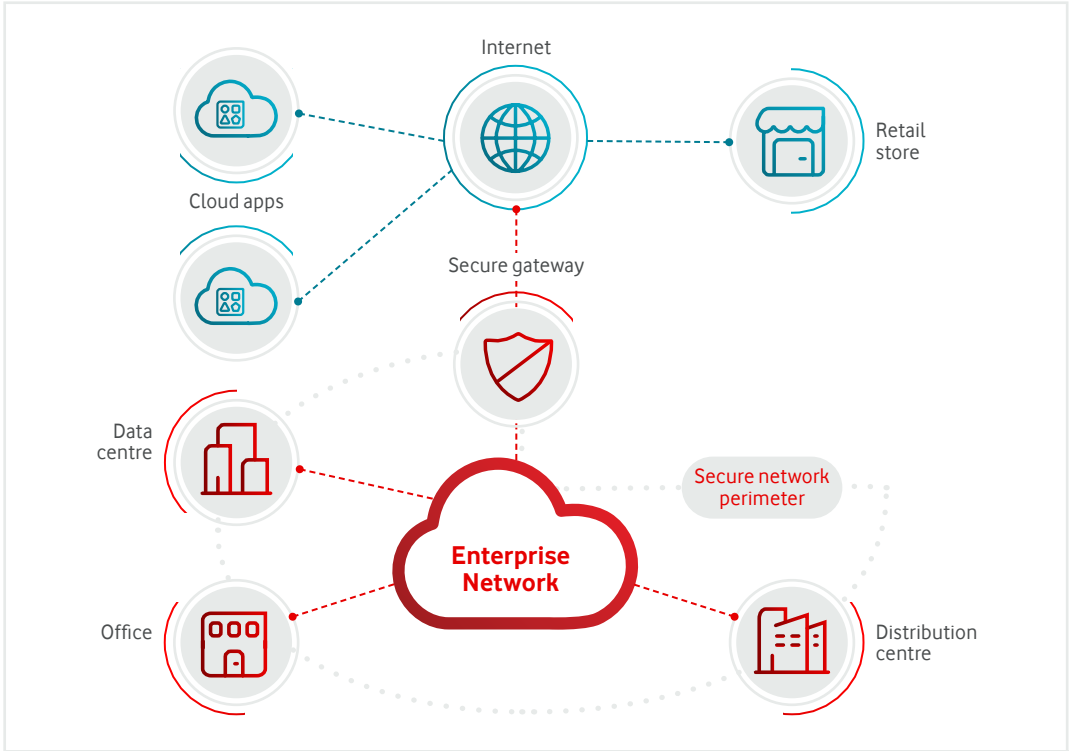

Figure 3 - the evolution of the enterprise network continues with the inclusion of Internet sites

4. The network isn't just about connecting sites together any longer

With a shift to hybrid working, it's not just about the applications. Currently, more than 28% of employees work in a hybrid model[3]. Companies offering hybrid or remote working options are also more likely to attract and retain top talent. Flexible working conditions bring growth, more satisfied customers and employees working from any location.

This changes the balance of the network, as you can see in the diagram.

The diagram really helps to illustrate why we often refer to this approach as a hybrid model. Everything shaded in blue is outside of our network perimeter and theoretically, presents an increased security risk. The increased risk comes from the fact that the Internet is an open network – attackers don't need to break down your physical security to gain access to your data. They just need to sit next to one of your colleagues in a coffee shop...

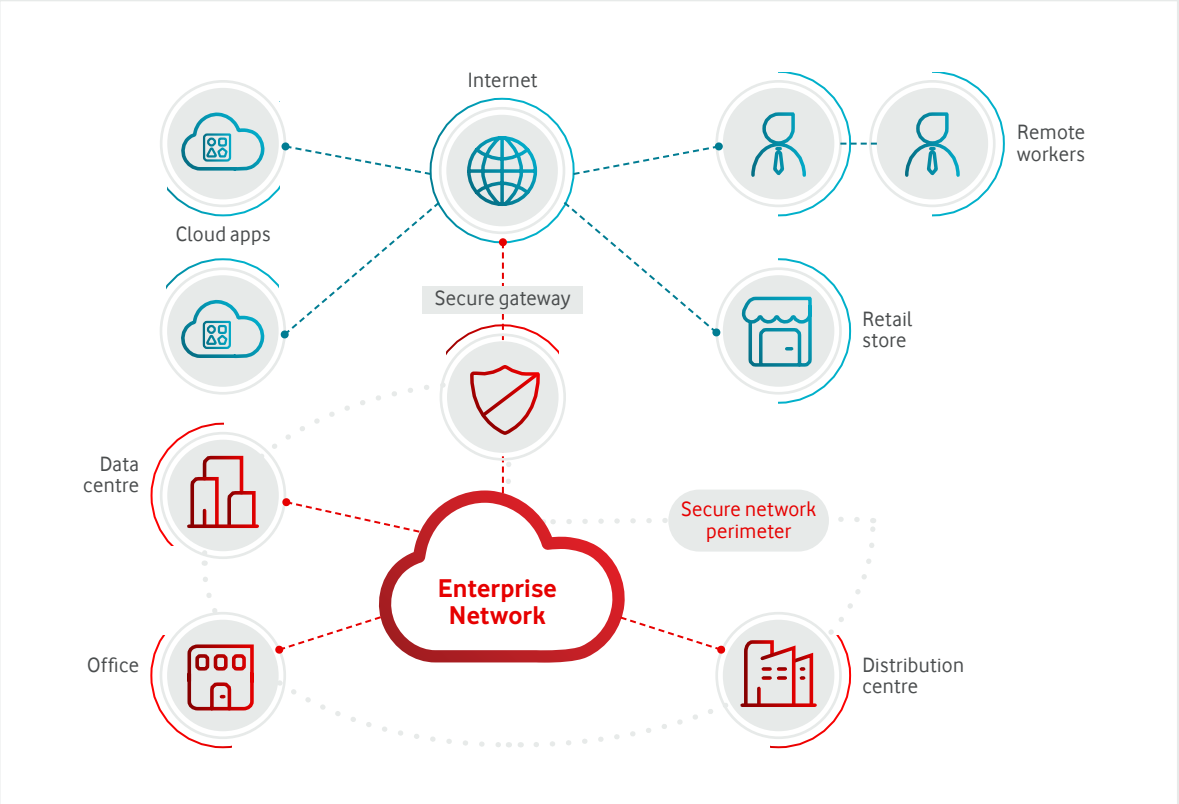

Figure 4 - the hybrid work model with users and applications now sitting outside of the traditional network perimeter

[3] https://www.forbes.com/advisor/business/remote-work-statistics/

4. The network isn't just about connecting sites together any longer

For remote work to be successful, the users need secure access to the same information they would have in the office. Given the hybrid network model we've been looking at, this means you will either need to give your users the software and/or hardware to allow them access to your data centres, or you need to shift the applications to the cloud.

According to a recent Gallup study, the biggest advantage of hybrid working for employees is improved work-life balance (reported by 76%), while for organisations, it's reduced employee burnout/fatigue (58%)[4]. Workers can use their time more efficiently and businesses can improve their retention rates. Management teams must think about long-term hybrid working strategies to support productivity and employee wellbeing.

[4] https://www.gallup.com/workplace/511994/future-office-arrived-hybrid.aspx

5. The data centre is no longer the centre of the network

We've already highlighted that applications are moving away from private data centres. The hyperscale cloud providers have re-invented the hosting industry in the space of a decade and show no signs of slowing down now. According to a recent study, cloud migration will be growing in 2024 as well, with 58% of respondents stating that they will migrate more workloads to the cloud, compared to 44% last year[5].

Note that this is the increased use of cloud, and not the death knell for the data centre. Some applications are better suited for running in shared infrastructure – think of websites as an example that need to be Internet-facing and able to scale elastically in the event of a surge in demand. Other applications are designed to run in private data centres – this might include a trading floor application that uses real-time stock market data to ensure traders make the most informed decisions.

In fact, according to Gartner's Multicloud Adoption Survey, 81% of customers use multiple cloud providers[6]. This is where organisations have some of their applications in the cloud, and some hosted privately. The private applications may not be in their own building but might be hosted in a multi-tenant data centre on physically separate hardware. There are hundreds of combinations, and the reality is, it can be different for each organisation.

What is clear though is that the data centre is no longer the hub of the network. There is a balance to be struck between the public (i.e. the Internet) and the private (i.e. the corporate network). The diagram on the side shows the next evolution of our network architecture, and this is where lots of organisations find themselves today.

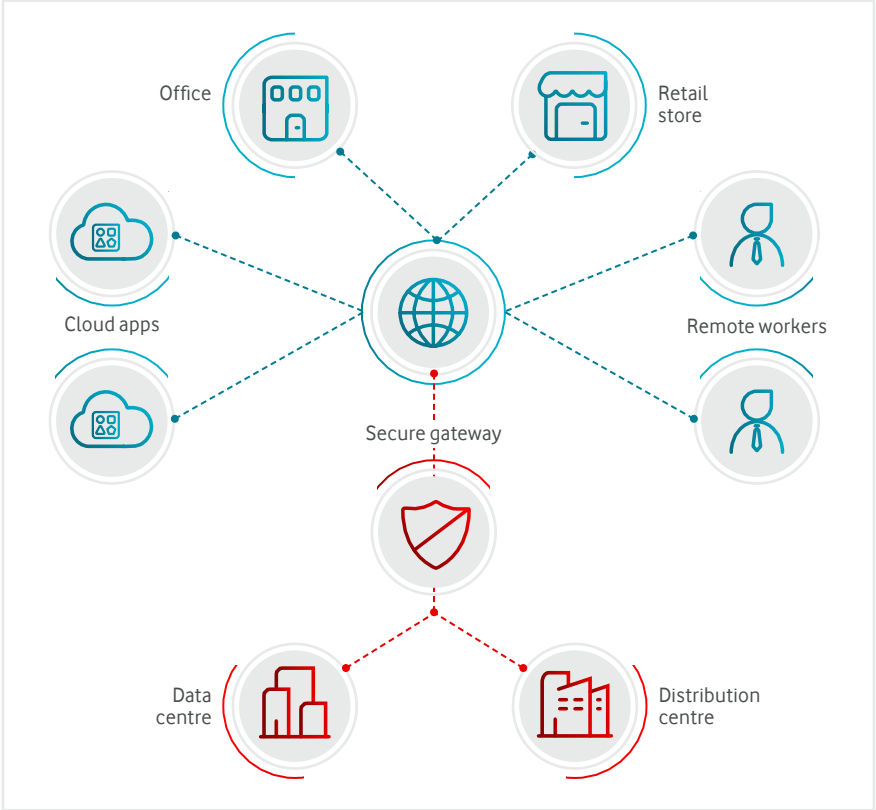

Figure 5 - the Internet evolution will continue with fewer and fewer sites having private connectivity

[5] thenewstack.io/cloud-migrations-pick-up-the-pace-in-2024/

[6] Multicloud Adoption Survey, Gartner, 2023

5. The data centre is no longer the centre of the network

There are large numbers of remote workers and locations connecting over the Internet and a concentration of cloud services complemented by the private connectivity to the largest sites. The balance of applications that the remote workers and smaller sites are using on a regular basis are running in the cloud – this is the defining reason why they're using the Internet. It doesn't make sense to install private circuits to get access to information that's available on the Internet.

The data centre and distribution centres are in constant contact though. Exchanging information in real-time to ensure that the mission critical stock applications are constantly up to date. Resilience is also a key concern here so there is likely to be more than one circuit going into each site. Fundamentally, the applications and the usage profiles are what drives the choice of connectivity.

6. What does this mean for tomorrow's enterprise network?

Where are you going to run your applications?

This is probably the single most important thing to focus on optimising in the hybrid world. Applications are the lifeblood of an organisation – if they don't perform, your users will feel the pain. The location of the applications underpins the types of connectivity you'll need and where you'll need it.

There are four things to consider when you're thinking about each application:

1. **Performance.**
   What are the performance requirements that drive the user experience for the application? If it's a Unified Communications service, latency will be key to the effectiveness of the platform. If you're running a public-facing website, then availability will be a really important metric to ensure you're not disappointing your customers.

2. **Scale**.
   What are your scalability requirements? Cloud hosting is designed to scale elastically as and when needed. There is spare capacity ready to be turned on at the touch of a button and this means you can react almost instantly to spikes in usage.

3. **Sensitivity.**
   What sort of information is available within the application? If the information is business critical, or particularly sensitive, you may need to reconsider any plans to host it in a public cloud.

4. **Data security.**
   Linked to the information, what are the data regulations you may be subjected to? If you're storing personal information, there are likely to be data sovereignty rules that mean you need to ensure it doesn't leave the country.

It's also important to understand that you don't need to put all your applications in the same place. Some will have a taxonomy that lends itself to a public cloud, some may need to be kept on-premise. You may also find that you need more than one public cloud provider. You might have some off-the-shelf Software-as-a-Service (SaaS) applications that run independently of your Infrastructure-as-a-Service (IaaS) deployments which are typically more tailored.

One way of making this daunting task a little easier is to use performance tiers. If you categorise each application into a tier, you can quickly see what your footprint looks like. You should aim to have no more than five or six tiers and experience tells us that this should be achievable. Private networks have been using Quality of Service, or QoS, for years to determine how we prioritise information when it crosses the network. QoS has typically been limited to five or six tiers, so this is a good target number to aim for.

7. What connectivity do you need for each user?

There are two key things to consider when it comes to user connectivity requirements:

1. **What applications do they need access to?**
   If the applications they're using have specific requirements in terms of performance, then these impact where the user can realistically work from. The chances are that there won't be any such limitations, so the choice becomes linked to their role and ability to be productive. An easy way to categorise your users is to segment them into personas. You can create as many as you need to map to your organisational profile. Some common examples of personas include executive, remote knowledge worker, on-site knowledge worker, and contact centre worker.

2. **Where are they going to be?**
   If the user is going to be in an office, then their needs must be considered in parallel with their colleagues on site. However, if they'll be remote, things get a little more complicated... and the detail you need now depends on the sensitivity of the information the user has access to. If it's of a very sensitive nature, then you may need to get into the specifics of how they're connecting. From that point of view, you need to understand whether they'll be using one location or multiple locations. Will they have access to fibre broadband as a minimum, or will they be relying on Guest Wi-Fi and the limitations that presents?

For users who can work remotely, you need to plan for their access over the Internet. There are lots of ways you can achieve this, and each has their benefits and challenges. The current technology trend in terms of remote access is Zero Trust networking.

This is an approach to security that ensures users can only access the applications and services that they are supposed to. This sounds like a basic security concept but starting from the principle of least privilege is one that has been overlooked historically by many organisations in favour of speed and simplicity.

7. What connectivity do you need for each user?

Forrester introduced the term Zero Trust in 2010 and Gartner introduced the term Secure Access Service Edge (SASE) in 2019. SASE talks about the convergence of network and security to provide a joined-up approach to connecting and protecting your users and your information. The key principal is this: security enforcement hosted at the edge of the network and not within the private data centre. You can see this represented in the diagram below:

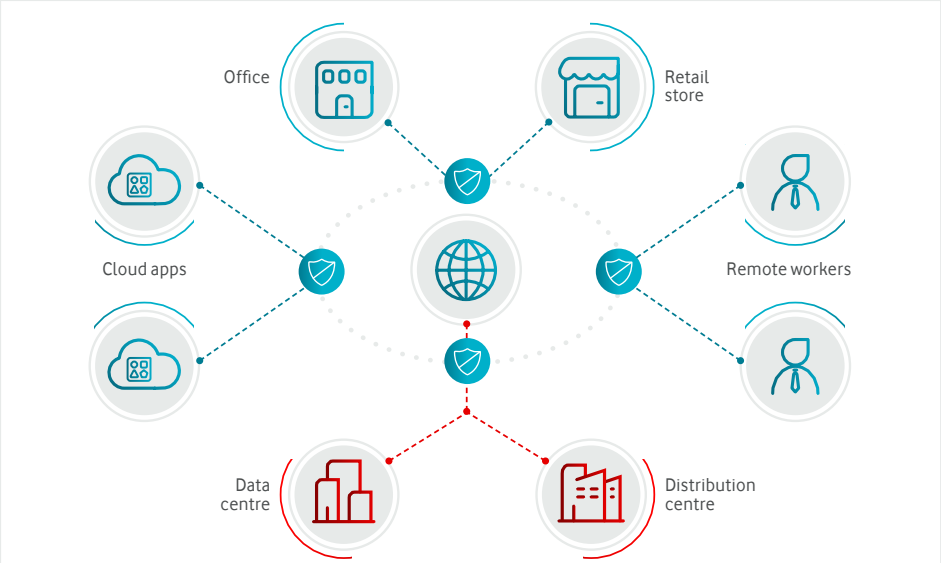

Figure 6 - an Internet-first enterprise network with security enforcement delivered at the edge

There are many benefits to adopting a SASE approach.

The key ones are:

1. **Homogeneous security for all your users and devices.**
   Previously, users and devices on a site connected to the private network might not have been subjected to the same levels of security inspection as their remote counterparts. That changes with edge-delivered security with all connections having access to the same level of inspection and enforcement.

2. **Cloud scalability for your security perimeter.**
   Service providers like Vodafone are investing in edge services to ensure the capacity is there for customers to use. This means the responsibility isn't on each individual organisation to predict their usage patterns and foot the bill for new hardware if it needs expanding. You treat your security like any other cloud workload and increase and decrease your subscription as and when you need to.

3. **No more technological debt.**
   When organisations had large appliances running in the data centre, these would be subject to lifecycle management every four to five years. This would lead to a large bill for the organisation. As these edge-delivered security services are typically part of a managed service, that challenge now falls to the service provider to manage and not the customer.

When you adopt an edge security model, the user-based connectivity becomes simpler to understand. You need to ensure all your users are covered with a licence or subscription for the security component, and that they have access to the right network service mapped to their persona.

8. What connectivity do you need at each site?

Now that you understand what the footprint of your users and your applications should look like, you can begin to plan the physical connectivity.

This used to be the most complex part of a network architecture, but now it's one of the simplest. The bulk of the work is already done by categorising your applications into performance tiers and segmenting your users into personas. It now becomes a cross matching exercise where you can look for the specific requirements that drive the demand for connectivity.

Among the things to consider are:

1. **Volume of users on site.**
   You need to ensure you have enough bandwidth for the number of users you expect to have on site. The question you need to ask yourself is – do you plan for the peak usage or can you build capacity for your average usage and accept that there may be challenges at peak occupancy? One of the benefits of SD-WAN is the ability to use a mix of private and public circuits for access. If the application requirements dictate, you can put in a smaller private circuit and a larger Internet one for the balance of your applications. In this way, you can manage the capacity and flexibility of the circuits into a site without breaking the bank. The remote working shift introduced by the pandemic means you may need to flex your bandwidth on a much more regular basis than before – service providers like Vodafone are now offering services that allow you to do just that through an intuitive portal.

2. **Importance of the site.**
   There are many ways of categorising importance and the route you take will depend on your organisation. For some businesses, the most important locations are those where the front-line staff are looking after the customers; for others, it's where the executive team sit as they're making decisions based on real-time information. Try and rank your site's importance and then use this to inform the levels of resilience and assurance that you need.

3. **Location of the site.**
   If you operate sites in well developed areas, access to network services should be plentiful. There should be fibre broadband passing your premises and Ethernet services can be relatively affordable as well. You can also benefit from strong mobile signal for either rapid deployment of a new site, or as a primary or backup connectivity option. Conversely, if you're running remote locations, your choices will be more limited. In these scenarios, the speed and reliability of the circuits may be affected. In this case, you may need to use a mobile service for your link to the outside world, or you can invest in satellite or microwave services to save the expense of a lengthy fibre dig.

4. **Applications in regular use on site.**
   We've mentioned the application requirements previously, and here's where they come in to their own. If you know who is on site and what they are accessing on a regular basis, you can understand the performance parameters they place on your network. As an example, if you have lots of users on conference calls for large parts of the day, you may need to allocate up to 2Mbps per user depending on your calling platform. You then need to understand concurrency levels – as a guide, the industry standard used to be 10:1. In other words, for every ten people on site, one will be running an application at any given time. With the new ways of working, this could now be as low as 3:1 for your organisation. It may also flip the other way if your sites become collaboration spaces with fewer people on site each day. Those that come into the office may be spending more time in face-to-face sessions and not on conference calls.

9. Bringing it all together

There is a lot of work to complete in order to understand exactly the right profile for your network. The shift to hybrid working has added an extra dimension in terms of the volume of additional remote workers to consider. It's far from an impossible task though and the key to it is understanding your applications.

The applications are the lifeblood of your IT estate and underpin the experiences of your users. Once you have categorised them into a range of performance tiers, the next job is to think about the users. If you can map these into a set of personas, you can create a new set of parameters to determine your network requirements. The final step is to use your new understanding of your applications and users to decide what this means for the connections into your sites.

So, it can be a challenging piece of work but one that will help you build a more optimised and responsive network for the for the hybrid world.

www.vodafone.com/business

Together we can

vodafone business