BROADBAND TESTING

A Network Scorecard For Evaluating SASE Clouds

A Broadband-Testing Report

First published January 2024 (V1.0)

Published by Broadband-Testing

E-mail : info@broadband-testing.co.uk
Internet: HTTP://www.broadband-testing.co.uk

TABLE OF CONTENTS

BROADBAND-TESTING

Broadband-Testing is an independent testing operation, based in Europe. Broadband-Testing interacts directly with the vendor, media, analyst, consultancy and investment communities equally and is therefore in a unique space within IT.

Testing covers all aspects of a product/service from business rationalisation in the first instance to every element – from speed of deployment and ease of use/management, through to performance and accuracy.

Testing itself takes many forms, from providing due diligence for potential investors through to public domain test reports.

Broadband-Testing is completely vendor neutral and independent. If a product does what it says on the tin, then we say so. If it doesn't, we don't tell the world it does what it cannot do… The testing is wholly complementary to analyst-related reports; think of it as analysts getting their hands dirty by actually testing what's on the latest hype curve to make sure it delivers on its claims and potential.

Broadband-Testing operates an **Approvals** scheme which prioritises products to be short-listed for purchase by end-users, based on their successful approval, and thereby short-cutting the evaluation process.

Output from the testing, including detailed research reports, articles and white papers on the latest IT technologies, are made available free of charge on our web site at HTTP://www.broadband-testing.co.uk

EXECUTIVE SUMMARY

- To paraphrase a classic quote, "not all solutions are created equal". In the case of SASE clouds, this is indeed true. The underlying architectures of a SASE cloud and its underpinning network can be radically different from each other, thereby impacting many aspects – performance, resilience, digital experience, security efficacy, data sovereignty and compliance, and more.

- For the potential end-user of a SASE cloud, be it a branch site, remote user or even connected IoT device, the fundamental problem is in understanding and interpreting what a SASE cloud vendor actually means by their description of the infrastructure; topics like coverage, as well as elements within that delivery network. Common terminology might be used: geographical coverage, regional presence, Points of Presence (PoPs), Data Centres (DCs), compute locations, secure edge services, secure web gateways, zero trust approach…. the 'marketingese' goes on and on but beneath the verbose descriptions, each vendor might have a very different definition of the same term. This, in turn, can have a significant impact on the quality of the service being delivered.

- Simple differences in approach, such as a vendor owning its own delivery network, rather than outsourcing – say via a public cloud service provider (CSP) can have a huge impact on the quality of the SASE delivery, from basic performance to downtime issues and the inability to pro-actively manage that SASE network you may be investing millions in.

- How capacity is managed can have a huge bearing on the true scalability and flexibility of a SASE offering. And where that capacity is readily available can make all the difference between a solution that works for a rollout where phase I is in – let's say – a relatively narrow band of regions, but where phases II and III might be far more geographically dispersed and challenging for the solution provider.

- The actual costs of a SASE solution can often be very difficult to calculate, especially over a longer time period. In some cases, it is hard to understand what is actually being delivered for, say, the basic service cost, compared with what is actually required – if it is actually available in the first place. And some SASE suppliers talk about possible "surcharges" and "additional costs" without it being clear what those costs are and when and where they might be levied. No customer wants surprises at the end of month or renewal time!

- This report is an attempt to create a scorecard for comparing SASE clouds and acting as a first line in the vendor evaluation process, that will hopefully eliminate many time-wasting POCs. For a potential SASE customer, the most important advice is to *not* be afraid to ask deep and detailed questions as to *exactly* what the SASE provider is actually delivering and is capable of delivering as your requirements expand and become more complex and what they actually mean by their terminology used within their service descriptions. When speaking with vendors, define *exactly* what you need, when and where you need it – not just for the initial deployment, but where you might need it further down the line and – finally – take **nothing** for granted!

ALL SASE SOLUTIONS ARE NOT CREATED EQUAL

IT has always been a black art – full of vagaries, black boxes and black holes. And marketing terminology…

Where vendors really excel is in reinventing themselves when it looks like a new and promising bandwagon has just rolled into town. When analyst firm Gartner first coined the term SASE – Secure Access Service Edge – they essentially defined it as "a cloud-based architecture model combining network and security-as-a-service functions, delivered as a single cloud service". Later carve-outs within SASE were created, for example extracting the mostly networking focused SD-WAN capabilities from the more security relevant SSE (Security Service Edge) components.

The problem, from a customer – as in end user – perspective with any analyst, or even government defined solution or requirement is, all that is being defined is a framework or set of guidelines; it's not explicit in its definition. That means any such definition is open to interpretation and here is where the vendors can get "creative" in their solution definitions. In other words, not all SASE clouds are created equal – what one vendor means by "x" and "y" can be very different to what another vendor means, even when using exactly the same terminology.

So, it's essential to be able to differentiate between those vendors who are serious, dedicated players, and those who have simply said "I reckon we can do that as well", as this difference in approach might be readily represented in their respective offerings. If we were to purchase a new car, say, it's very unlikely that simply listening to the sales patter and the website marketing descriptions – and even a detailed technical specification – would result in us buying the car without a test drive first.

But how do you "test drive" the right SASE solution? In an ideal world there would be some fantastic modelling tool available and a set of vendors who were willing to provide explicit technical breakdowns of their entire SASE offering for you to input and compare to see which provided the best solution for your exact requirements – today and tomorrow. And the days, years and decades afterwards.

However, there is no such perfect world out there. In terms of what information is available as a starting point, not only can it be confusing, as we've noted, in terms of definitions of similar terminology from different vendors but, equally, the amount of information on the respective vendor offerings varies dramatically from one to the next.

So, in an attempt to help short-cut the evaluation process when considering investing in a SASE solution, here we've put together our own series of guidelines, in the form of a network scorecard definition which is designed to leave no proverbial stone unturned when creating a shortlist of "go to" SASE vendors for a POC. Read on…

DEFINING A NETWORK SCORECARD FOR EVALUATING SASE CLOUD SOLUTIONS

While it could be argued that there are a near-infinite number of variables – and very few constants – that go into defining a SASE cloud and its underpinning network infrastructure, there are some very obvious – and key – areas to deep dive into when working out exactly what a vendor is actually offering.

Here we've attempted to provide some weighting to each key evaluation category (though obviously this could vary by use case) to help define the target requirements, with 10% as the "default" weighting.

Infrastructure Dependencies	Public Versus Private Cloud?	**10%** - Arguably less important than the functionality and characteristics of the service delivery (uptime, resilience, performance).
Global Footprint	Where Are The Servers?	**15%** - Coverage alignment with customer footprint is critical. Servers need to be adjacent to users for optimal delivery.
Responsibility Model	Who Controls The Network?	**15%** - Direct control over selecting DC locations: e.g., close to IXs – traffic shaping, peering, etc, directly influence service delivery.
Services Menu	What Services Run Where?	**10%** - Not all customers will want to consume all services from day one. For example, they might start with SWG, then add CASB or ZTNA.
Capacity Management	How Is The Capacity Managed?	**15%** - Closely aligned with footprint and server location: available capacity and server infrastructure location are key to both resilience and service availability.
Performance Optimisations	How Is Performance Validated?	**15%** - If customers have a performance issue, this can result in security problems as users will often bypass controls and basic rules, putting apps and data at risk – not good!
Troubleshooting And Remediation	What Happens When Something Breaks?	**15%** - Be honest: service issues *are* inevitable (maintenance, software bugs, cable cuts Internet weather events…), so vendors need the ability to rapidly and directly address issues as, and when, they occur.
Charging And Licensing Model	How Do You Pay Them?	**5%** - Customers will pay more for a superior solution, but understanding the charging and licensing model is still important for budgeting and planning purposes. But first, you need to the service to deliver as required at any cost!

PUTTING A SASE SCORECARD TOGETHER: A STEP-BY-STEP PROCESS

1. Public Versus Private Cloud

The first element of a SASE solution to understand is: does the vendor actually own and operate its own delivery network or is it based on a public cloud/hyperscaler? If it's the latter then it's not simply a case of who the Cloud Service Provider (CSP) is – Azure, AWS, GCP etc, – but how the SASE vendor controls service delivery when it doesn't own the cloud service it is using (more on this later). If the SASE solution is based on a private cloud – i.e., owned and operated by the vendor – then the latter question is readily answered, but it is still a fundamental requirement to learn of the capabilities of that private solution – again more on this later.

It's important to point out here that the service is not simply about performance, but also security and data provenance/governance. It's somewhat ironic that a real cloud in the sky is often grey and the public cloud has potentially more grey areas than a certain famous novel and film…

Key Questions: If a SASE provider is using a public cloud delivery mechanism, ask exactly which service providers make up their end-to-end delivery and what exactly those relationships consist of.

- Is the SASE solution private (owned and operated by the vendor) or operated by a Public Cloud?
- If operated by a Public Cloud, in how many regions are the services configured and are allowed to use the global network for ingress/egress?

2. Where Are The Servers?

Many aspects of a SASE solution can all too easily be taken for granted by a potential customer, not least when dealing with a huge global vendor, but the reality is that even the most basic aspects of a proposed solution should be questioned, starting with: where exactly are the servers?

A service delivery model might be primarily centralised, it might be primarily edge-based, it might be public cloud-based, or a combination of all these elements and more. The important point to understand is which regions of the world are covered locally by a service delivery point. It's far too easy to simply assume that a SASE solution based around a giant CSP engine means that every corner of the globe is being serviced, but that is not necessarily the case. Even where there is a physical presence in a stated region or country, exactly what services are being delivered from that PoP? Is it a real physical PoP with full-service delivery functionality, or merely a virtual PoP that isn't actually running any local compute functions, but simply routing existing traffic? For example, a PoP might be listed as existing in a given location, but what services is it actually performing in said location, such as traffic inspection or policy enforcement?

Where is the DNS hosted and where are the configuration and optimisation services such as load-balancing being carried out? Again, the location is not simply a means of determining performance capabilities, even though that is fundamental to proximity of delivery, as with any remote service. It's important to factor in service delivery elements such as:

- Ensuring customer experience anywhere across the globe remains consistent.
- Optimising the end user experience (low-latency and, again, consistency are key).
- Data locality – is it compliant? Does data remain in the same location?
- Consistent support of regional/local requirements – content and language, for example.

It's not purely a case of what can be delivered and controlled locally but equally, what happens when something goes wrong? Can it be directly fixed or updated, or does it involve going through one or multiple 3rd parties? We all know the potential issues of dealing with tiers of 3rd party support and the escalation processes involved. If a SASE customer is knowingly using a provider that involves 3rd party delivery and support, then it is likely that said customer will want to build its own first line support service on top of that already being offered. This means more expense and more complications – does this mean that outsourcing therefore becomes effectively self-defeating?

Key Questions: Ask the vendor to define explicitly where their servers/service delivery points are located and what they actually consist of. See point 4 for what they are actually delivering.

- Are they in edge DCs? Do they map to your locations?
- Are they in public cloud regions? How many regions are available to you?
- How many regions are included with the various tiers of service?
- How do they get to the edge, where the users are?
- Are the POPs real or virtual?

Example of Hot-Potato Versus Cold-Potato Routing

One well-documented scenario is what is described as hot-potato versus cold-potato routing.

It's derived from the classic phrase: "drop something like it's a hot potato" since in hot-potato routing, an ISP hands off traffic to a downstream ISP as quickly as it can. Cold-potato routing is – unsurprisingly - the exact opposite of hot-potato routing, where a service provider carries the traffic as far as possible on its own network before handing it off to a downstream ISP, as is necessary. The two policies directly reflect the different priorities for the ISPs. An ISP or CSP can choose the carry the traffic to and from their network, either on their own backbone, or force their customers' provider to carry the traffic to the required destination. Key to this are cost savings and, importantly, eliminating ownership of performance.

Example of Hot-Potato Versus Cold-Potato Routing (cont)

In other words, with the hot-potato approach, the goal is to offload the traffic as quickly as possible, minimising the usage of the ISP/CSP network. In contrast, the goal of cold-potato routing is to carry traffic on the providers network as far as is possible, so as to maximise the control that provider has on the end-to-end QoS. However, a hot-potato approach can provide performance gains – especially when combined with route control and intelligent traffic management, in tandem with extensive peering and premium transit options to optimise performance further and expand the service delivery footprint. In contrast, using a cold-potato approach, a public CSP would see benefits from using its own network as far as possible to control the customer to its own requirements and, inevitably, cut transit costs by eliminating 3rd parties as far as possible.

3. Who Controls The Network?

Directly related to the question of server location is "who is controlling the delivery network?" Moreover, how much control of that network does the SASE supplier have? If it's a case of said provider peering with other network providers, then this opens up several areas to query:

- How – if at all – is optimal connectivity and service delivery assured? After all, it's not a simple case of optimising locally hosted applications but also connectivity to major cloud, SaaS, CDNs, and other commonly accessed services. For example, probing questions about whether transit selection is based on cost or performance, as well as redundancy of transit links and connectivity in every location may shed valuable light on the underlying resilience of the vendors' infrastructure.

- Is the service being prioritised over other customer services being delivered by that 3rd party peering service? This is particularly important when the "platform" is disaggregated from the SASE solution, for example if based on public cloud or hosted in a 3rd-party ISP network, essentially outside of the SASE vendor's sphere of control.

- How many 3rd parties are involved? For example, first mile, backhaul and last mile might all be different suppliers working together. Customers need to understand every hop in the path of their critical traffic, as it not only impacts security and performance, but also data sovereignty and regulatory compliance. In addition, a SASE provider using peering might choose least cost routing over optimal performance, at least for part of a 24-hour cycle. Or is it a single path solution, meaning a single point of failure is introduced to the delivery network?

- If it does involve multiple network providers, then how does the SASE solution provider, let alone the customer, have complete visibility over the end-to-end delivery? In which case, again, what happens in the event of a network failure or any form of unscheduled downtime? Are we back to the classic finger-pointing exercise that has been the bane of an IT managers life, when outsourcing, for decades? Everyone in IT knows that the "middle mile" is where no provider makes any money, but is where the majority of problems occur – so who actually wants to take responsibility for fixing them? Most problems experienced in delivering end-to-

end traffic are directly Internet-related events. But who resolves them in a scenario involving one or more 3rd parties?

The point is, if a SASE provider does not see your traffic as being critical in terms of prioritisation, why outsource - if it's not a critical requirement - in the first place? A direct analogy here is when ordering a product from an online retail vendor; for example, one that might also be a major CSP 😊.

We all know the scenario where – when tracking the package delivery – we can see it has arrived at the local depot, but then sits there for 48 hours before it goes out for delivery, usually not on the requested day. But the problem is, there is absolutely nothing we can do about it.

Imagine if that's your data and applications in a similar scenario? Even reducing 48 hours to a few seconds is enough to kill real-time applications stone dead, for example. So, who do you call – Ghostbusters?

Key Questions: What happens – support-wise – in the event of an inevitable network issue? Who is accountable, who ultimately owns the issue and what SLAs are in place to ensure responsive support and remediation of issues quickly?

- Does the provider interconnect with other networks?
- Does the provider peer with cloud and SaaS providers?
- If there is a routing problem in the cloud provider network, how does the vendor troubleshoot?

4. What Services Run Where?

It's all too easy to assume that, when a vendor cites the existence of a PoP or compute location, that it will support the complete scope of available SASE services. But this might easily not be the case. On a simple level, it could be that the security and networking services are served from different locations, and therefore not all available locally. Or even different SSE services like SWG, ZTNA, DLP or Threat separated into different regions and not centrally located in a unified service-delivery fashion.

One question is, quite simply, does the presence of a PoP mean anything at all? Point being (no pun intended), what services are actually running on that PoP – whether real or virtual? Even if they have compute functionality, how scalable is that local compute resource? A PoP might simply not have any local compute resource or incomplete services, but merely act as a routing node, having to "phone home" for any information and backhaul some or all traffic to, and from, remote locations.

It's important therefore to understand how the layers of applications work. For example, if you have a SASE provider with an in-line or edge-based service, what exactly do they need to backhaul, all their control traffic or just minimal traffic or traffic for specific services? What applications are served locally, and what remotely (or even extremely remotely)? While it may be that a provider has edge locations defined, equally it might be that, for most functions, it has to talk back to a central location which might be several router hops away, with all the implied performance issues associated with multi-hop routes.

There are also potential problems with outages relating to the inability to access end or main services, wherein you could run out of edge compute first, thereby causing service requirements from other locations – more traffic being routed on less bandwidth and to more distant locations, so a domino effect in terms of performance drop could easily happen.

Even where there is compute at the local edge or PoP, performance is further impacted upon by whether or not the delivery networks use a single-pass or multi-pass architecture for functions such as data security scanning; a multi-pass architecture can readily increase latency and impact negatively on application performance and the user experience.

Similarly, with respect to the policies, are they unified policies across all SASE services, do they have to replicate and, if so, again, is this inducing further latency or creating a point of vulnerability in the architecture? Equally, from a management perspective, it's important to know how many different interfaces you are going to have to manage, or is it all done via the classic "single pane of glass"? Obviously, the former implies training costs as well as simply being more complicated. More prosaic: how easy – if at all – is it to see the status and health of every location?

Finally, from a user experience perspective, a lack of local compute resource might not only impact on application performance, but also the availability of local content (such as web search data), local language related content and, even more importantly, geo-fenced applications where security and compliance may be compromised.

Key Questions: What exactly does each stated PoP and/or compute location actually provide, in terms of localised processing and services?

- Are the vendor's security and networking services available in all locations?
- What is the architecture; for example, is there a phone home service or main brain?
- Do all POPs have compute resources for service delivery (and, for example, processing traffic to protect data, stop threats, inspect encrypted flows, etc.) or is there a need for backhauling traffic to another location?

5. How Is The Capacity Managed?

Of all the IT buzzwords of the past decade, "scalability" surely has to be right up there with the most overused. And in isolation it actually means nothing. But in SASE delivery terms it is fundamental in two ways: what is the absolute capacity of a vendor's SASE offering and where and what can it scale to?

It is critical therefore to understand how a SASE vendor operates its infrastructure. Again, it is all too easy to take it for granted that a vendor has true global coverage, along with unlimited capacity and limitless scalability. After all, that's how the major CSPs have brought customers in, with such implied promises. And there's the key word "implied".

So, it's vital to provide your absolute capacity and expansion plans to a potential supplier and get specific examples of how their solution is working in environments that are as close to your requirements as possible.

One regular gauge for performance delivery, and therefore the underlying capacity to deliver, is what is known as the p99 usage or latency. Put simply, p99 latency is the 99th latency percentile. This means that 99% of requests will be faster than the given latency number. Or put another way, only 1% of the requests will be slower than your p99 latency. Being able to eliminate latency as a major cause of lowering the user experience, should be a foremost requirement in identifying a potential SASE provider, especially with the reliance on video and VoIP traffic nowadays.

Similarly, when probing the SASE provider's solution, it is important to ask about the triggers for adding capacity and what is the utilisation at steady state to allow for headroom when there are the inevitable traffic spikes, Internet weather events or the addition of more users or expanded infrastructure demands - for example, if the customer consolidates more services with that same SASE provider.

Naturally, any service provider is going to want to maximise its capacity usage – that's simply being realistic – but, given the bursty nature of Internet traffic especially, it's important to know how it is managing and paying for its capacity and availability. Basically, there are two ways to pay for cloud services: reserved instances (RI) or on-demand instances (ODI). The latter can be essential when burst capacity is needed but, at the same time, it is more expensive to provision. So, it may be that a vendor wants to reduce costs by going down the RI route only, as the cost differences can be significant. However, it also means that a provider with ODI capabilities might be passing that cost onto the customer, so your SASE bills could reflect this.

Another favourite trick of a vendor – and this is not simply SASE-related – is to talk about its increases in capacity and reach over a given period: say the past 12 months. However, don't mistake statistics for what is actually available and out there right now. For example, a vendor might talk about doubling its capacity over that 12-month period, without stating what the original capacity was. In other words, it's a meaningless statement. Find out therefore, exactly what capacity is where now and what is absolutely planned, in terms of expansion, in the near and mid-term, in very specific terms – what and how?

A lack of anticipated expansion could easily result in heavy investment in one SASE provider having to be written off with a move to another provider that has the required capacity. It's an expensive and potentially damaging mistake to make…

Key Questions: What are your absolute capacity limits now and into the future, in terms of being able to plan for the unexpected in what is a very uncertain world?

- What is the scaling metric used by the provider for adding new servers/PoPS?
- How many new locations did the provider add in the last 12 months?

6. How Is Performance Validated?

Performance validation is fundamental in choosing a SASE provider.

We've already touched on the importance of minimising latency and ensuring capacity is available, as and when needed, and with the same levels of performance that are normally available.

A vendor needs to be able to demonstrate exactly how its performance is validated and guaranteed; are they even monitoring the performance and how does this tie into capacity increases and decreases? Is it even true end-to-end validation? If 3rd parties are involved in the delivery process, how is performance monitored and guaranteed on and via those 3rd party network elements?

As a customer, it is important to be able to see, in real time, the current status and performance of the network, ideally in a way that measures the user experience – latency levels, uptime levels and the ability to see longer term trends, historical usage; basically, any metrics that validate performance and service delivery. A NetOps person is answerable to their C-level bosses; they will demand this level of detail as validation for investing in that SASE vendor.

In the event of a network problem and failover requirements, where are the closest failover PoPs – within how many hops and are they within acceptable latency levels? And, in a worst-case scenario, where several infrastructure elements go down (such as several PoPs), to what extent can the failover PoP or PoPs sustain an acceptable level of service delivery until everything is back online? Again, there are also potential local content issues here too: if a location in one country goes down and the backup is in another country – or even continent – are local language, content and compliance maintained or not? If not, what are the implications for the users?

Key Questions: What optimisation technology and features are being deployed in order to offset the inevitable latency and packet loss issues that will arise at times, and how does that relate to the geo-location of the end users and guarantee the best user experience, wherever they are located?

- Does the latency integrate into the product?
- Is there end-to-end performance management?
- What is the closest PoP to fail-over if the primary centre goes down? Is there a 3rd location?

7. What Happens When Something Breaks?

In the previous section, we touched on some specifics where there are network problems and outages but, on a general level, it's important to understand a SASE vendors' failover, redundancy and recovery strategies.

So, what does happen when something breaks? It's understandable if a NetOps guys doesn't want to think about network failure scenarios, but it's vital to do so. So, is there automated failover built into the plan? This needn't be in the event of a total backout, but in the situation of a brownout or simply performance deterioration as a result of network congestion - what kicks in and from where? How is performance maintained or not? And what are the notification and fault management processes? How pro-active is the notification?

It's important to ask if a SASE vendor can guarantee given levels of service availability in the event of the kind of different levels of network failure described above. Otherwise, how can you plan for any kind of network issues and create contingency plans? Do you (and the SASE provider) fully understand and control every hop in the packet's journey for triaging

issues when and if something indeed breaks? In the event of 3rd parties being involved in the service delivery, when a problem arises, where is the first line support and what are the escalation processes – at 11 in the morning, four in the afternoon and three in the morning? In other words, is the same level of support available 24x7?

Key Question: If 3rd parties are involved in the delivery network, what is the accountability of each provider in the event of a problem arising? Is there always the same single point of first line support contact or not?

8. How Do You Pay Them?

And now we get to the crux of the matter – how is a vendors' SASE solution priced and what exactly do you get for your money?

So, how exactly does a vendors' bill work? What comes as standard (and what exactly is "standard"?) and what are additionally charged as premium services? For example, there might be a basic centralised configuration service included, but configuring at the edge may involve a premium. It's all too easy for vendors to hide potential additional costs, especially with something as complex as a global SASE deployment over an extended period of time. It's essential therefore to work out what coverage is free and what will incur surcharges.

For example, some SASE vendors are merely quoting that for expansion into certain world regions: "additional costs may be involved". Obviously, this is not a workable scenario – it is essential therefore to highlight exactly where your SASE requirements are initially and could be over various phases of proposed or potential expansion. Otherwise, there could be some nasty financial shocks awaiting. So, what PoPs are included as standard and what PoPs require additional fees in order to be available? In terms of traffic levels, are there burst fees or usage-based fees on top of the standard contracted delivery levels and are there other potentially hidden annual costs that you need to be aware of?

What SLAs are available and at what cost? Where 3rd parties are involved, if a vendor can only provide an SLA on its own network element, then it means they simply don't have control over the rest of the delivery network, with obvious implications for performance and downtime issues. Can – and/or will – a vendor build-in penalties in the event of not matching an SLA? If not, then why not? It's important too, to understand exactly what is being honoured in an SLA. For example, guaranteed performance levels may relate in some instances to simply a one-way, end-to-end delivery, rather than a true bidirectional network service delivery guarantee.

Key Questions: Can you give me absolute, explicit bottom-line costs on my initial deployment plans and those of future, phased expansion, including service level guarantees?

- Are there burst fees or usage-based fees?
- Are there billing "true ups" in the contract that the customer needs to be mindful of?

TO SUMMARISE...

As we've noted several times – take nothing for granted when pre-evaluating a SASE vendors' cloud offering, starting with what do you get for what you pay? And what does any and every addition and expansion to the default service cost?

There's also the issue of localisation. We've noted the importance of local content availability but there's also the data sovereignty aspects: what runs where, and does data really go via and into the countries as stated? Moreover, can you – as a customer – control where data and applications reside and where they are routed? In some cases, it simply optimises delivery, but in other cases it could present costly compliance issues or, even worse – think medical scenarios - lives being potentially at stake,

In terms of your dedicated service, just how dedicated is it? For example, in a scenario where there are shared IP addresses between subscribers, just one customer could create problems for every customer sharing those IP addresses, such as a shared, dedicated egress IP. And how transparent is that SASE service? Is end-to-end connectivity visible via a public facing portal, for example? Are the status of services and uptime and performance statistics available 24/7? Or even at all?

How "global" is a proposed global coverage? For example, you could present a vendor with an initial deployment of 20 sites in one part of the world and expand that basic deployment over one or more phases to, say, 100 sites in many different regions of the world and ask if they can provide a given level of services to all those locations, with guarantees.

Don't be afraid to ask what might seem like dumb questions – you might get dumb answers! Put simply, don't part with any of your budget without being 100% sure of what you are being offered and will receive now and further down the line.