Point of View

Are you in the dark about performance?

Exploring overlay and underlay networks and how to avoid blindspots to improve performance and expedite remediation.

verizon business

Foreword

Networking and security are becoming more tightly integrated and networks more aligned to deliver high-availability, high-speed access to applications hosted in the public cloud. This is making the procurement of networking a high-stakes decision.

David Bailey
Global Solutions Executive

David has over 20 years of experience working with large global enterprises and public sector organisations to define and meet their infrastructure and security needs. This has included wired and wireless networks for Internet of Things (IoT) applications and agile, intelligent infrastructure to support demanding applications like artificial intelligence (AI).

Over the last twenty years, I've been involved in the procurement process for networking for dozens of companies, mostly large multinational enterprises, across a multitude of industries. I've seen technologies come and go—like ATM and frame relay, being replaced by MPLS, and then that giving way to the emergence of virtualised architectures that make greater use of the internet for transport, such as SD WAN.

There's been a slew of acronyms, but the biggest change has really been around how companies use their networks. Back in 2011, Marc Andreessen famously said, "Software is eating the world." As It turns out, this was remarkably prescient. What he didn't talk about was the need to deliver those applications in a responsive and highly reliable way. And that's where the network comes in. The growth of artificial intelligence (AI)-based apps has only raised the stakes and made the choice of network provider more critical.

One of the fundamental differences I've seen in how enterprises go about procuring networking is whether the overlay and underlay (see below) are bought:

- Together, with a network provider managing all aspects of networking end-to-end.
- As part of one of several "towers" with underlay often bundled with managed desktop and overlay bundled with managed data centres/cloud services.
- Entirely separately.

Each of these approaches can work and there are many companies around the world following them successfully. Companies often start the procurement process with a strong view of which of these approaches is right for them. This is influenced by a number of factors:

- **History**—The past experience of members of the team
- **Technology**—The company's architecture and/or desired operating model
- **Policy**—Company procurement practices and commercial considerations

On many occasions, customers have asked my team to advise them on which approach to take; in other cases, as part of our response to an RFI or RFP we've suggested why the chosen approach isn't optimal for the company/project.

This paper encapsulates the advice that I have given these companies. I examine the options and share my opinions on their pros and cons, why companies make the choices they do, and the mistakes some decision makers make.

To split or not to split?

Five common justifications for splitting underlay and overlay.

Over the years I've heard a number of justifications for choosing to procure underlay and overlay separately. The table below lists the five that have cropped up most often. These account for the vast majority of cases. I'll explore each of these in more detail over the coming pages.

	Reason	The anticipated pros: What customers say	The potential cons
1.	**"Best of breed/ best value"**	"Separate RFPs will enable us to leverage our buying power to get a better price on the underlay network, while allowing us to choose the best overlay network."	Buying underlay on price can lead to a variety of issues, including performance problems, lack of end-to-end visibility and delays resolving issues.
2.	**"There are better pairings"**	"Overlay provision includes maintaining on-premises equipment, just like LAN management/end-user device management. It makes sense that the same people that configure/fix laptops set up and maintain routers and switches."	The skills and systems needed to maintain end-user devices are very different than those required to manage networks. Whereas in the past, managing LANs was quite separate from managing WANs, increasingly the same tools can be used, enabling simpler, more comprehensive support.
3.	**"Speed and simplicity"**	"Separate processes will mean two RFPs, but each will be more concise and easier to manage, leading to a faster conclusion to the effort."	This approach might mean a shorter RFP, but when it comes to the rollout, you need to coordinate two suppliers. These companies will often have overlapping portfolios and be competitors. This inevitably means that it will take longer and potentially be riskier. There can also be ongoing issues around the delineation of responsibilities.
4.	**"If it ain't broke, don't fix it"**	"We're happy with our underlay provider, why change?"	Stipulating that the underlay provider remains the same runs contrary to the idea of an open and fair procurement process. This is not a good way to get the best performance or best value for the business.
5.	**"Telcos restrict options"**	"A targeted RFP for underlay followed by an RFP for overlay services will mean that we can open the process to all overlay providers in the market."	Opening the process to multiple overlay providers increases the complexity of decision making. There's never really a scenario where the procurement team can compare 'apples for apples'. Telcos are actually very good at evaluating overlay providers and they have compelling reasons to get it right. And there are very good reasons why most telcos work with a few carefully selected vendors.

Definition: Underlay

Definition: Overlay

Examining the reasons

"Best of breed/best value"

Buying underlay on price can lead to a variety of problems, including performance shortfalls, lack of end-to-end visibility and delays resolving issues.

"There are better pairings"

Another reason that underlay networks are split is because the procurement team believes they will be able to reach a much better price point for the overlay. Systems integrators (SIs) often promise a much more economical price for the management of each network device when this is bundled in with the management and operation of other on-site IT services such as desktop.

The other thing I often hear from SIs is that they are prepared to work with almost any original equipment manufacturer (OEM). Each OEM has specific skills required for their own specific technology, even if this is a simple activity through a web fronted GUI management platform. Specific knowledge of the OEM platform is essential to the smooth operation of your network.

Decision makers should question the level and depth of the skill sets each overlay provider has with their OEM of choice. In the event of a serious problem, your business will rely on this relationship to get to the route cause and fix the issue. Telcos choose to work with a select few vendors so that they have a wide spread of technical resources to call upon and depth of knowledge of each vendor.

If your business was procuring a fleet of vehicles, would your fleet manager specify the manufacturer of the engine that went into each vehicle? It's more likely that you'd trust the vehicle manufacturer to specify and procure the right engine for the vehicle than try to assess the options yourself. Just like vehicle manufacturers, telcos put a lot of effort into choosing the right partners as getting this decision right is critical to their success.

"Speed and simplicity"

It might be faster, but this is a bit like procuring train tracks and rolling stock separately. You can't do it in complete isolation—otherwise you could end up with trains that don't fit on the tracks. So, then there's the question of how you build an RFP that avoids conflict and ensures clear delineation of responsibilities. And then there's the question of who mediates disputes?

Speaking to organisations that have split overlay and underlay networks, the most common feedback they've expressed to me is not delight about the savings made, but frustration that network issues are taking longer to resolve, leading to increased user dissatisfaction. Some have even said that because there's no single party responsible for the network, some performance issues become too difficult to pin down and resolve and so they end up just having to accept them as "known issues".

If you have an organisation that can see the whole architecture with self-healing tools across the whole architecture, the network can keep pace with the evolution of your applications.

"Telcos restrict options"

There are reasons why telcos have a shortlist of vendors that they prefer to work with.

The relationship between overlay and underlay networks needs to be symbiotic
Network carriers often have the tools and expertise to evaluate networking equipment far more thoroughly than any client company ever could. Because they are making this evaluation for dozens of customers, Tier-1 providers like Verizon invest in hiring world-class networking engineers to put OEM equipment providers through a stringent process of technical and operational testing.

Network carriers also have greater insight into how topologies are likely to change in the future and are able to test this with potential OEM vendors before deciding which ones are most capable of supporting their customers' networks not just today, but into the future.

Avoiding spreading themselves too thinly
Underlay providers enforce strict security and operational compliance on OEM vendors. By focusing on a small number of OEMs they can also create automations between the OEM's management tools and their own systems, rather than trying to integrate this as a one-off after an RFP process concludes. This "industrial scale" integration means that should there be a problem, there's a large pool of skilled people they can engage, so it can be resolved more quickly than if they were relying on a niche team with a shallower depth of expertise.

The importance of visibility

When your business is procuring a network, it's choosing a partner that it will trust its data to. You can encrypt data and purchase a plethora of security technologies, but your data is still passing through a multitude of devices (routers etc.) as it makes its journey. If you have different underlay and overlay providers, if there is a security breach it is much more difficult to work through the breach to understand exactly how and where it occurred. Was it through the overlay or the underlay? Similarly, if the network isn't performing correctly, you can get caught in a game of buck passing—is it because of:

- Bad routing/prioritisation? An overlay problem.
- A misconfigured router or switch? An underlay problem.
- An international peering issue? An underlay problem.

When I explain end-to-end visibility, I often use an analogy with how aircraft are monitored and controlled. For safe, efficient travel, air traffic control must have visibility of where all aircraft are at any one time. Similarly, you need full visibility of the path your data takes for it to be transported efficiently across your network. Imagine the aircraft pictured below is a packet of your very important data.

Field of vision necessary for full visibility and control

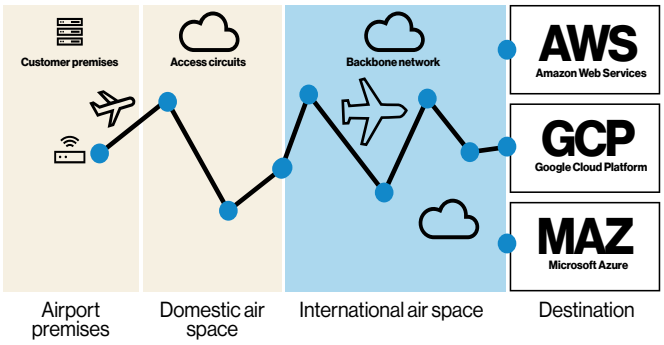

To have a single partner that's responsible for and able to control your data across the whole of your network, the underlay and overlay need to work together. To achieve this, this custodian needs control and management of the underlay to be able to identify and optimise performance along your data's entire journey—all the way from your site to your public cloud, SaaS provider or data centre. Effectively, this means that they must operate the backbone network.

There are two ways to achieve this:

- Build your own network. This will allow you to specify and control every component. While this would give you total control, it would be extremely expensive. I've only ever seen governmental or quasi-governmental organisations (such as critical national infrastructure operators) pursue this route.
- Procure services for the whole of your network architecture from an organisation that owns and operates a backbone.

The visibility challenge

When you choose an organisation that owns and operates a backbone network to operate both the underlay and the overlay architecture, its tools, systems and people will have an end-to-end field of vision of your data's journey. This includes visibility into the peering points where they connect to the cloud operators (public and private).

They'll be able give you information on how all of this is performing and advise if components need to be upgraded.

With the right network design, this provider can even provide granular insight that can help identify why a particular application isn't performing as intended.

You might ask the question why the overlay provider cannot have end-to-end visibility of the network. When the underlay is split from the overlay, the overlay provider has very limited visibility and zero control over what is happening in the underlay network.

You have visibility of your data across your LAN, until it reaches your on-site edge router. Once it leaves your premises, that visibility is lost until it reaches your cloud provider or data centre. The same happens with data sent from your cloud providers to your sites.

If a network node (one of the waypoints shown as circles in the diagrams) or one of the peering points into your cloud is running slowly or even offline, the only actor that has visibility of this is the underlay provider.

A partner providing an operating model for your customer premises equipment (CPE) including overlay and then also providing services that manage and operate your cloud or data centres would have a huge blind spot in the middle.

Embedded security

A further benefit of choosing a partner that operates a backbone is that it can incorporate a security fabric into the core of the network, relieving you of the need to procure additional security services from an OEM vendor.

Some Tier-1 providers, such as Verizon, can also offer load balancing and move application workloads into the core of the network. This placement of workloads in the core of the network has the potential to generate major savings on data egress charges from your public cloud providers. This is becoming more important as more enterprise applications are spread across multiple cloud operators.

A false economy?

These direct financial savings on cloud services could dwarf any cost savings that might be achieved by splitting overlay and underlay networks, especially when the total cost of ownership is taken into consideration.

Field of vision of overlay provider

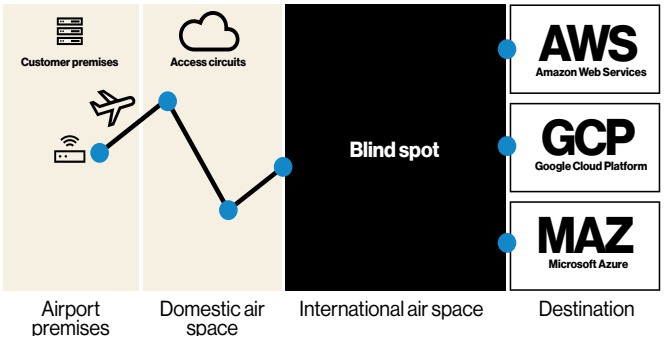

Compare this to the relationship between the control tower at an airport and the airport traffic control system. The tower only manages aeroplanes on the ground. The tower can instruct pilots where and when they can take off, but as soon as the plane leaves the runway, control is handed over to air traffic control (ATC). ATC instructs the pilot which route to take to their destination, which altitude to fly at to avoid collisions etc. It also decides which planes get priority access to runways and which have to queue for a landing spot— are put into holding patterns. The airport tower only regains control when the plane comes into land.

Field of vision of underlay provider

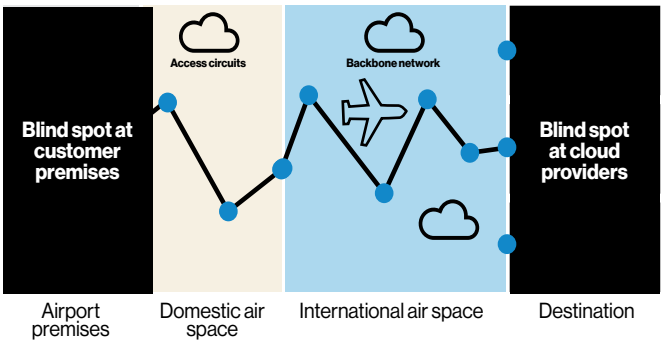

Limited field of control

When underlay and overlay are separated, something similar can happen across your network. The overlay provider can control your data while it's on your premises, but once that data leaves your site, the route it takes and how long it takes to reach its destination is down to the underlay provider.

In reality, the situation can be even worse. If your underlay provider isn't a tier-1 provider, it may not even have full end-to-end visibility of your data's route or performance—in our analogy, this is akin to having separate national and international control.

Field of vision of underlay provider without backbone

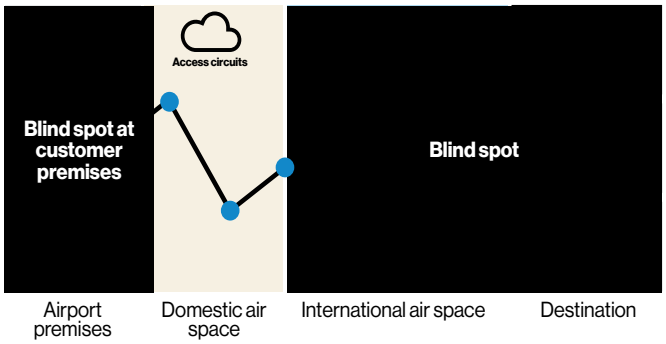

Stuck in the middle?

Where the overlay provider has a blind spot on the underlay, the opposite is true for the underlay provider. Back to our analogy, the underlay provider is able to see where all of the aircraft are in the sky. But they don't know where they took off from or what their destination is.

Customers often expect the underlay provider and overlay provider to communicate to identify and resolve bottlenecks and other issues, but in my experience this rarely happens effectively.

Usually, the customer sits in the middle and has to engage both organisations to help identify and iron out any issues across the network fabric. As well as putting additional burden on the customer's team, this can increase the time it takes to resolve issues. Given the rapid pace at which applications are now developed, evolved and moved around, by the time that one issue is resolved, another may have appeared.

Conclusion

Faced with these problems of managing separate underlay and overlay providers, many CIOs may look to re-negotiate. But by the time they reach this point, it's likely that the business has become reliant on some of the tools and systems that have been adopted. This means that it can be very difficult to sever managed desktop, underlay and the cloud operating model. And of course, there may be contractual issues, including penalty clauses. The situation can be further complicated by the provider placing minimal cost on managing on-premises devices (CPE)—so that the cost of managing just the underlay is practically the same as managing the underlay and the CPE.

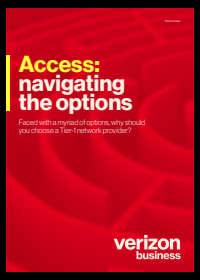

Navigating the options for enterprise access

This other paper from this series gives insight into how the choice of provider can have a huge impact in terms of performance and outcomes for your business. This is especially important for latency-sensitive applications, such as many AI use cases.

15-minute read

Supercharge your AI applications with a better network

This other paper from this series discusses the rapid pace at which AI is evolving and how supporting technology layers need to keep pace with AI software and hardware. This includes high-performance, low-latency transport (networking), which can be crucial to successful outcomes.

15-minute read

Let's talk

Verizon can help you quantify the total cost of ownership of your infrastructure. The cost savings that companies manage to realise from buying underlay separately often don't live up to expectations and may be outweighed by the cost of lower-than-expected network performance and increased fault resolution times.

We could work with you to look at the total cost of ownership of your infrastructure. There are often much bigger savings to be achieved through network architecture improvements than haggling over the price of access circuits.

If you still have questions after reading this paper, get in touch with us at:

verizon.com/business/en-gb/contact-us

The Network Procurement series

This paper is one of a series exploring the growing demands on enterprise networks and important questions companies should ask during the procurement process to help ensure that the solution they chose are truly enterprise-grade and will meet their current and needs.

Something big is coming

Data | IoT | AI

This paper explores some of the key drivers behind the explosive growth in the volume of data enterprises are gathering and what that means for network planning.
verizon.com/business/resources/articles/iot-genai-data-explosion.pdf

Access: navigating the options

Performance

There are many decisions to make when buying networking. Understanding the three tiers of the internet is critical to thoroughly evaluating the options. This paper explains what they mean for network performance and security.
verizon.com/business/resources/articles/tier-1-isp-enterprise-connectivity.pdf

Network peering

Cloud | Performance | Reliability

Peering is fundamental to network performance and consequently enterprise applications, particularly ones based in the cloud. Despite this, it's rarely discussed during procurement. Read this short paper and put that right.
verizon.com/business/resources/articles/network-peering.pdf

Are you in the dark about performance?

Data | Performance | Manageability

Read this paper to learn how the decision to split the procurement of physical (underlay) and logical (overlay) networks can affect network performance, visibility and manageability.
verizon.com/business/resources/articles/overlay-underlay-network-procurement.pdf

Better together

Security | Performance | Manageability

Cyberthreats continue to grow in volume and sophistication. This short paper offers six reasons to consider greater integration between cybersecurity and networking to improve protection while reducing workload and cost.
verizon.com/business/resources/articles/unified-network-security-services.pdf

Supercharge your AI applications

AI | Performance

Artificial intelligence (AI) promises to be the most disruptive technology since the internet became mainstream around 30 years ago. This paper explains why network performance is critical to the performance of many AI applications and realising the anticipated benefits.
verizon.com/business/resources/articles/network-infrastructure-ai-platforms.pdf

verizon business