Pricing Guide

Wide Area Networking

P2P, MPLS, Dark Fiber & More

Wide Area Networking Pricing Guide

What's inside?

This report discusses the most commonly utilized wide area networking (WAN) products, including:

➢ Point-to-point leased lines (P2P)

➢ Wavelength services

➢ Multiprotocol label switching (MPLS)

➢ Dark fiber

➢ Software-defined wide area networks (SD-WAN)

Each of these wide area network types comes with their own set of **advantages, disadvantages, and pricing nuances.**

This report shares pricing insights that we've gathered from 500+ telecom service providers and over 15k service quotes **so you can make the best decision when it comes to your wide area networking services.**

Lightyear's Telecom Pricing Data
Disclaimer

Let's face it, the telecom industry isn't transparent when it comes to pricing.

And, due to the sheer abundance of service availability, pricing structures can at first blush seem complicated and intimidating.

We know this because we've procured over 15k quotes on behalf of enterprise IT buyers.

We're using our experience to demystify these structures by breaking down the many factors that impact WAN connectivity pricing and what factors should matter most to you.

Notes on our dataset

- Our data set includes ~15k quotes and has been scrubbed for outliers

- In instances where we show a highest and lowest observation, the quotes displayed most often represent abnormal cases that should not be expected and go to exhibit the huge variance and uncertainty you will see in telecom pricing

- We partner with hundreds of providers and our data set includes connectivity quotes from all major carriers, and many regional carriers as well

- Quotes are distributed all around the US, with no heavy regional bias

Key Terms Review

At the end of this report there is a full glossary of various terms that impact the pricing of your WAN connectivity.

But first we'll cover the terms that impact your WAN connectivity, no matter the time, place, or WAN type:

Bandwidth

Bandwidth is the amount of information passed from one point to another over a network at any given time. To increase bandwidth is to increase the amount of information that can be sent in this fixed amount of time, which is increasingly important as the demand for rich media (e.g., streaming movies, video conferences) grows.

Bandwidth is measured in bits per second, commonly megabits per second (Mbps) or gigabits per second (Gbps).

For P2P and Wave circuits, your bandwidth can be offered as fixed or burstable - but it is pretty rare, so we will spare you the details there.

In a Data Center

Connectivity is always cheaper when procured in a data center. This is due to the fact that data centers often have multiple telecom providers you could utilize - and more supply leads to lower costs for you. This is especially true in colocation

service that is on-net means the carrier providing the service already has working connectivity in the building and is able to service your entire project.

Generally speaking, a carrier within 200 feet of your location is considered to be on-net. Off-net, in turn, refers to anything beyond that 200-foot range, typically requiring a telecom construction project (i.e., a build) by your service provider to get their connectivity to this location.

Build costs are exactly that: the amount of money required to physically establish the given part of a network. Once completed, a telecom company will recoup its build costs through the service the network provides to consumers. For this reason, build costs are often prohibitive in remote and rural areas due to the amount of time needed to recoup these expenses.

Service term length

WAN services are procured under long-term contracts. In general, the longer your service term, the cheaper the service will be. In our experience, pricing on terms three years or longer is often similar, whereas pricing on terms less than three years is typically more expensive.

Summary

A quick overview of the services covered in this WAN Pricing Guide:

- **P2P** - Often more expensive than Wave, but costs are dropping due to substantial current growth

- **Wave** - The (semi) new guy in town, providing ultra-high bandwidth at an efficient cost (often cheaper than MPLS and P2P)

- **Dark Fiber** - Coolest name of the bunch, hands down. But be cautious: DF can appear quite cost-effective on the surface, but you have to buy and manage the electronics on your own

- **MPLS** - Traditional, carrier managed, meshed and typically more expensive

- **SD-WAN** - Potentially the most cost-effective, but correct implementation is very nuanced

Point to Point WAN

Point-to-point (P2P) network connectivity refers to A and Z addresses linked by a single connection. The direct linking of this network topology precludes anyone else from sharing its bandwidth. With this form of WAN connectivity, each point in the network acts as both the sender and the recipient of the data.

"Legacy P2P" WAN technology, as it's referred to, consists of T1/E1 and T3/E3 services. And, despite losing share over time to newer telecom technologies, copper still makes up 55% of global broadband infrastructure. Modern P2Ps are of course provisioned over fiber/ethernet and are available in speeds up to 100 Gbps.

P2P WAN Advantages

P2P WAN has been a standard for connectivity for years. It's recognizable, widely available, and easy to understand for most infrastructure teams.

- Available in speeds ranging from 5 Mbps to 100 Gbps
- "Flat" network type that allows end users complete control over routing, traffic queuing, etc. without any intervention from the carrier
- Many providers allow enterprises to opt into special features like dynamic routing, CoS and even data encryption
- Tends to be very low latency

lightyear

Point-to-Point WAN (Cont'd)
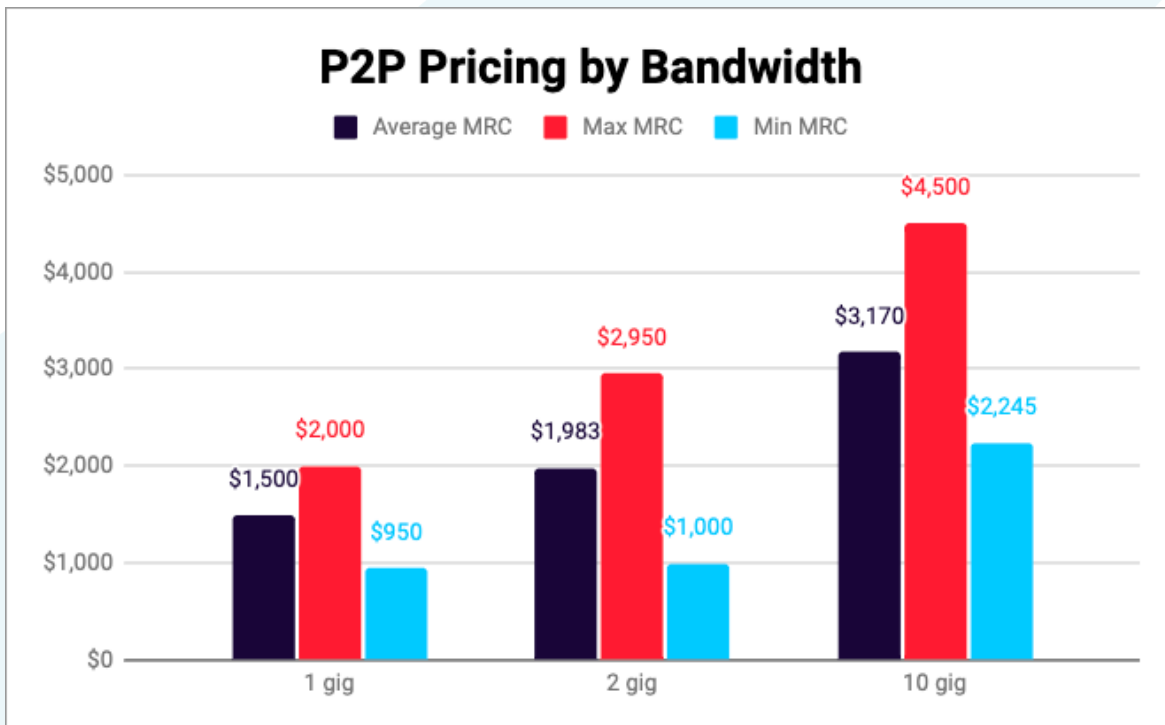
P2P WAN Challenges

As a legacy connectivity type, there are certain challenges with P2P WAN that can make it less attractive for cloud-native brands.

- Can be expensive and complex to scale especially in site to multi-site topologies.
- Unlike MPLS, telecom providers have limited troubleshooting abilities

P2P WAN Pricing Factors

Traditional P2P WAN services are more or less the standard for connectivity pricing—but their legacy standing does impact price across several axes, including:

- Class of Service (CoS) - superior latency, jitter and packet loss guarantees
- Protected vs. Non-Protected Circuits - dynamic routing across the carrier backbone in order to circumvent network provider issues to ensure availability
- Encryption - additional layer of carrier provided security
- Whether or not your router is managed - nowadays, they are hardly ever managed
- And, of course, the standard four factors we already defined above: bandwidth requirements, if a build is required, if you are in or out of a data center, and your service term length

P2P Pricing by Bandwidth

■ Average MRC ■ Max MRC ■ Min MRC

Bandwidth	Average MRC	Max MRC	Min MRC
1 gig	$1,500	$2,000	$950
2 gig	$1,983	$2,950	$1,000
10 gig	$3,170	$4,500	$2,245

P2P chart data constants: Out of data center, on-net, fixed bandwidth, 36 month term

Point to Point WAN (Cont'd)

Lightyear P2P Pricing POV

It's worth noting that P2P lines and Wavelength Services are very similar services (discussed in the next section), with the mode of transmission being the only difference. P2P services are traditional, while wavelength (or wave) is new school. We see legacy P2P prices going down a ton to compete with Wave.

For example, we know one carrier who is running very aggressive pricing on P2P connectivity if you're in its local exchange carrier (LEC) footprint ("on-net"). As one of the only phone companies offering this type of connectivity in 16 states, this carrier offers $2,245 for 10 gigabit P2P—which is a very good rate.

Wavelength Services

Commonly known as "fiber-optic services" in the U.S., wavelengths transmit data over fiber-optic cables, creating dedicated, large bandwidth connections (often referred to as lit connections). Fiber adoption is growing exponentially, and carriers are aggressively building out network infrastructure to service major metro markets with high capacity, super low latency wave services.

This growth has been significant enough to drive down the costs of legacy P2P technologies.

Wavelength services are wave transmissions over fiber, in most cases using DWDM technology which supports nearly 200 separate waves up to 100 Gbps on a single strand of fiber.

Waves are low latency because they route on the carrier's backbone over purpose-built infrastructure, usually in a static manner. This is good because it's the fastest path. The downside is that if there's a problem on the carrier's network that the wave must cross, there's no ability to route around it. Some carriers provide "protected" wavelength services but these can get pricey and are often limited. That's why solutions such as Lumen's E-Line (which can route around) exist, but this comes at the cost of latency; the E-Line is basically a P2P link.

Wave Advantages

As an ultra high-capacity WAN connectivity type, Wave is one of the most efficient options at your disposal, offering the following advantages:

- Cheaper and faster than legacy P2P
- Regional distance rarely a factor in pricing unless international
- Efficiency of fiber drives down costs
- Designed for high-capacity WANs
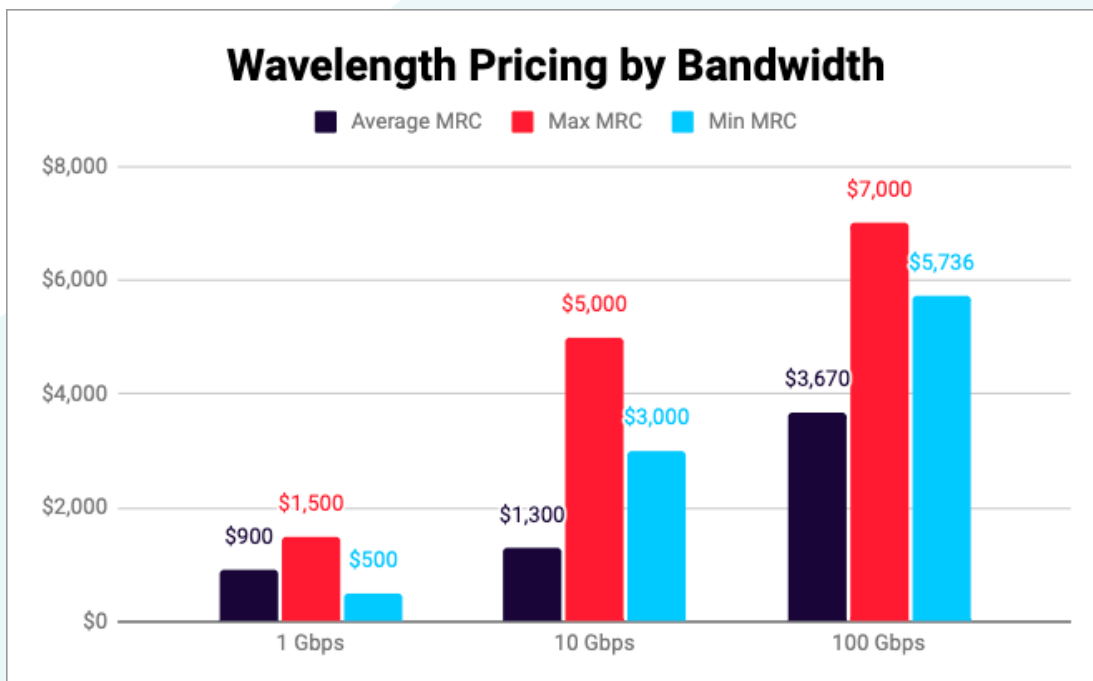
Wavelength Services (Cont'd)

Wave Challenges

Wave connectivity can quickly become expensive if you're a global business, and it is really only a good option if your company has high-capacity requirements.

- International connectivity becomes expensive, and availability is limited.
- Service tiers and throughput are limited to 1, 10, 40 and 100 Gbps in most cases.
- Most carriers are still building out capacity to offer waves in all markets; availability may be limited especially in tier 2 and 3 markets.

Wave Pricing Factors

Wave connectivity is most common for high-capacity WAN transport and is only offered at 1 gig and up. You won't be able to source Wave services circuits under 1 gig for that reason. The following factors impact wave circuit pricing:

- Protection
- Size of market
- And, of course, the standard four factors we already defined above: bandwidth requirements, if a build is required, if you are in or out of a data center, and your service term length

Wavelength Pricing by Bandwidth

■ Average MRC ■ Max MRC ■ Min MRC

Bandwidth	Average MRC	Max MRC	Min MRC
1 Gbps	$900	$1,500	$500
10 Gbps	$1,300	$5,000	$3,000
100 Gbps	$3,670	$7,000	$5,736

Wave chart data constants: *Out of data center, on-net, fixed bandwidth, 36 month term*

≼ lightyear

Wavelength Services (Cont'd)

Dark Fiber

Lightyear Waves Pricing POV

We see P2P prices continuing to decrease as P2P competition increases with wavelength and dark fiber services, and other types of connectivity grow in popularity.

Wave is less expensive in general because it is a new type of product specifically designed for very high-capacity WANs. Demand for these WANs is currently limited to very specific markets, such as big cities, and carriers are building out infrastructure to support very fast, low latency lines there. Another factor driving wavelength service prices down relative to other WAN alternatives is its efficient mode of transport -- a carrier can provision multiple wave circuits over a single strand of fiber.

One fun tidbit on wave pricing - we've recently seen one provider quote $699 MRC for a 10 Gbps wave connection. Mind you, this connection was from data center to data center, but still a good data point to have in your back pocket.

Sometimes these cables are also referred to as "unlit." Many companies look at dark fiber not as a service but rather as an asset or real estate. And the value of dark fiber is greatly determined by its location.

Using these cables can be tempting as they offer enormous bandwidth (800+ gig in some instances), but in most cases, there are more practical alternatives. For instance, buying dark fiber and running it at only 10 gigabytes often doesn't make sense. In this case, wavelength services may make more sense. This is why Lightyear will almost always ask customers who are looking for dark fiber if they would also consider a lit alternative.

Dark Fiber Advantages

Dark fiber has the potential to provide a considerable amount of bandwidth for your business. The larger your connectivity needs are, the more attractive this option will be.

- Outstanding bandwidth potential
- Based on physical location and utilization, some savings are possible
- Enterprises are completely in control of all elements of services.

Dark Fiber(Cont'd)

Dark Fiber Challenges

Dark fiber often comes with increased bandwidth availability that surpasses your needs as a business. And you'll need to spend time working with specific providers to understand exactly what pricing will work for your relationship.

- Availability can be very difficult. Many large carriers do not offer dark fiber as it is not really a "service".
- Completely unmanaged service that requires enterprises to provide all optical gear on both ends.
- Cost/utilization ratio can be prohibitive
- While pricing is based on distance, there can be other arbitrary elements considered when providers quote this service like who is going to use it, is the building commercial or residential, market type, etc.

Dark Fiber Pricing Factors

- Highly variable
- Distance of connection
- Often charged by the strand
- Cost of electronics (not included in cost of circuit)

Lightyear Dark Fiber Pricing POV

We've omitted a pricing chart here as Dark Fiber pricing varies SO MUCH.

We have seen one provider price two strands within the same city for $850 MRC, and we have also seen carriers charge thousands for similar lengths.

The best takeaway here is that pricing for dark fiber is unpredictable – hence where there is no data chart for Dark Fiber in this report.

Multiprotocol Label Switching (MPLS)

A Multiprotocol Label Switching (MPLS) WAN is built on a provider's private backbone - this guarantees bandwidth and, being a virtual private network (VPN), MPLS traffic isn't routed across the internet. It's localized on a carrier's private network, making MPLS a very predictable environment.

MPLS was widely leveraged as a WAN technology to support centralized call processing for distributed VoIP deployments. For this waning use case, there is still nothing better. However, with the migration to cloud based VoIP/UCaaS/CCaaS, you'll see less and less of MPLS.

MPLS Advantages
- Fully managed, meshed network providing any-to-any connectivity
- QoS tiers available to maximize performance of some traffic vs others (aka "express lanes")
- Service Level Agreements ensure guaranteed delivery/performance and packet loss, jitter and availability

MPLS Challenges

As a carrier-backed connectivity type designed to support centralized private data center type networks, MPLS isn't the best choice for cloud-enabled environments. You also face potentially high optimization costs and a lack of customizability.

- Optimized for private datacenter hub and spoke types of networks.
- More expensive than traditional Ethernet options as the same provider must provide the MPLS connections at every site on the network even if they are not optimized to do so in a specific market. I.e. Verizon providing an MPLS connection in an AT&T service area.
- Service provider controls network configurations and provides limited reporting/analytics
- WAN optimization to streamline delivery adds extra costs
- No "paths" to public clouds
- High cost per Mbps vs P2P and wave services
- All "eggs" in a single basket

Read the State of MPLS Report here

Multiprotocol Label Switching (Cont'd)

MPLS Pricing Factors

MPLS connectivity shares a number of pricing factors with Wave connectivity. Make sure to cross-reference your company's specific infrastructure needs as you work out which connectivity type makes the most sense for your business.

• Router management
• Quality of Service / traffic queuing
• Leasing terms
• And, of course, the standard four factors we already defined above: bandwidth requirements, if a build is required, if you are in or out of a data center, and your service term length

Lightyear MPLS Pricing POV

You run into the biggest issues with off-net price gouging when you utilize an MPLS network that requires you to use the same provider at each of your locations; if your enterprise wants to expand into a new geography where your MPLS network carrier is not the LEC, you can expect much higher prices.

MPLS Pricing by Bandwidth

■ Average MRC   ■ Max MRC   ■ Min MRC
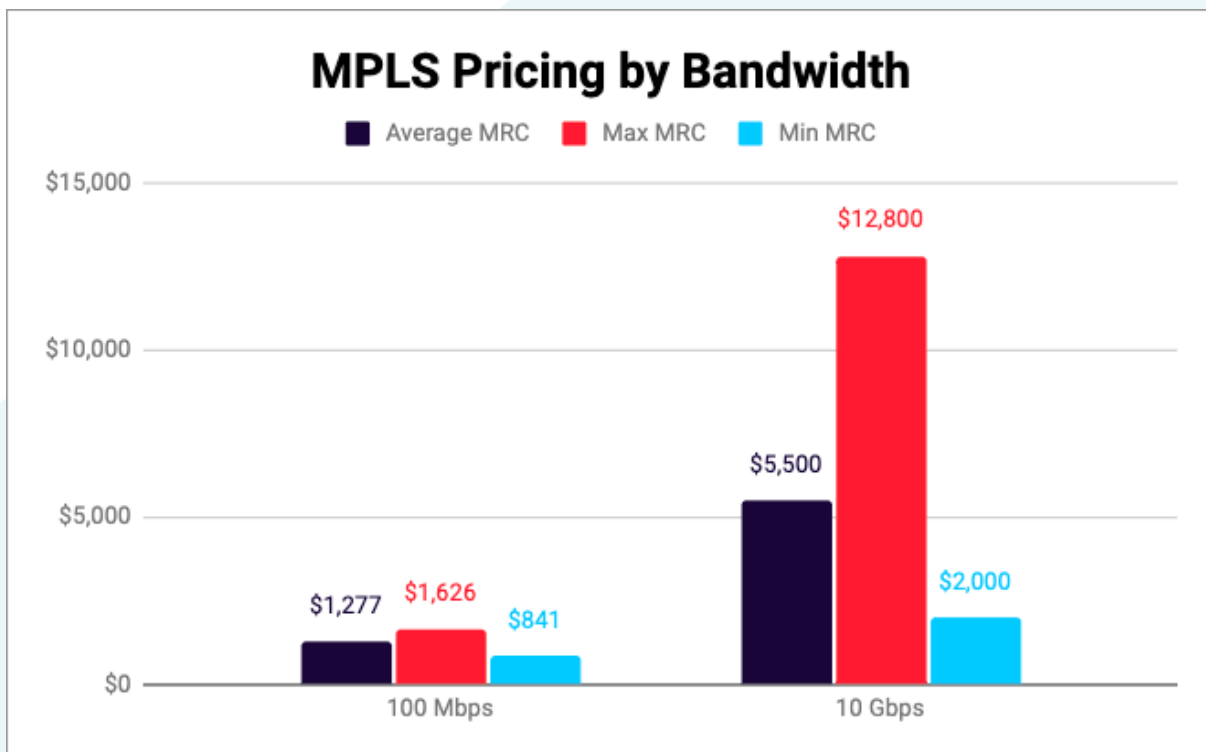
MPLS chart data constants: Out of data center, on-net, 36 month term

Software Defined WAN (SD-WAN)

SD-WAN uses software and a centralized control function to efficiently steer traffic across a WAN based on your business needs. As such, SD-WANs can allow for many of the same quality of experience and security benefits of MPLS or P2P (see our piece on SD-WAN vs MPLS here) with data transmission over the public internet rather than preset or dedicated routes - hence why many people think of SD-WAN as a cheaper WAN option (more on that later).

An important distinction to make is that every other WAN type we have discussed in this report relies on having its own physical form of transport while SD-WAN does not. SD-WAN provides a disconnect from the physical underlay network, so you are completely vendor agnostic, while software manages the overlay network essentially as an "intelligent VPN".

SD-WAN Advantages

As a cloud-native functionality, SD-WAN presents several advantages over other connectivity types. Understanding these advantages helps to determine whether or not your current infrastructure is well-suited for it.

- Improves business agility and ability to scale facilities footprint without reliance on dedicated transport

SD-WAN Advantages (Cont'd)

- Software overlays provide intelligent network management and granular reporting of performance, application utilization, etc.
- Transport-agnosticism often reduces costs
- Designed for increased reliance on cloud services oftentimes with dedicated peering into major public cloud and SaaS platforms
- Ability to massage traffic, perform error correction and set priority queues for traffic
- API integrations on many platforms enabling custom integrations, reporting, monitoring
- Vendor-level guarantees of consistent user experience

SD-WAN Challenges

SD-WAN runs into a few distinct challenges that other types of WAN do not—centered around adoption rates and the overall "newness" of this offering.

- Fragmented landscape with less certainty on performance metrics vs MPLS/P2P
- Not as capable as MPLS in some centralized hub and spoke environments
- Technically public network in many cases, which raises security concerns for some customer profiles
- Despite the power of SD-WAN, underlying quality circuits still matter

SD-WAN (Cont'd)

SD-WAN Pricing Factors

It should be made clear that SD-WAN pricing is so nuanced that it really deserves its own report on "how SD-WAN is priced". In the meantime, here's the skinny of it....

- There are numerous SD-WAN deployment models that affect pricing, including managed, outsourced, and self-serve
- The number of nodes in your organization's network
- Capability of the SD-WAN hardware appliances (# of appliances, # of power supplies, # of WAN interfaces, etc.)
- The level of security you require between locations and any remote workers. You can pair up your SD-WAN solution with Secure Access Service Edge aka SASE to ensure the tightest security standards
- Bandwidth aka throughput

Lightyear SD-WAN Pricing POV

SD-WAN is a powerful alternative to other WAN solutions because of its ability to route across various provider networks and the network visibility it provides. It could very well be the answer to your networking prayers, however proper scoping and implementation of an SD-WAN is nuanced and complex.

SD-WAN Pricing by Bandwidth

■ Average MRC

Bandwidth	Average MRC
50 Mbps	$228
100 Mbps	$300
250 Mbps	$500
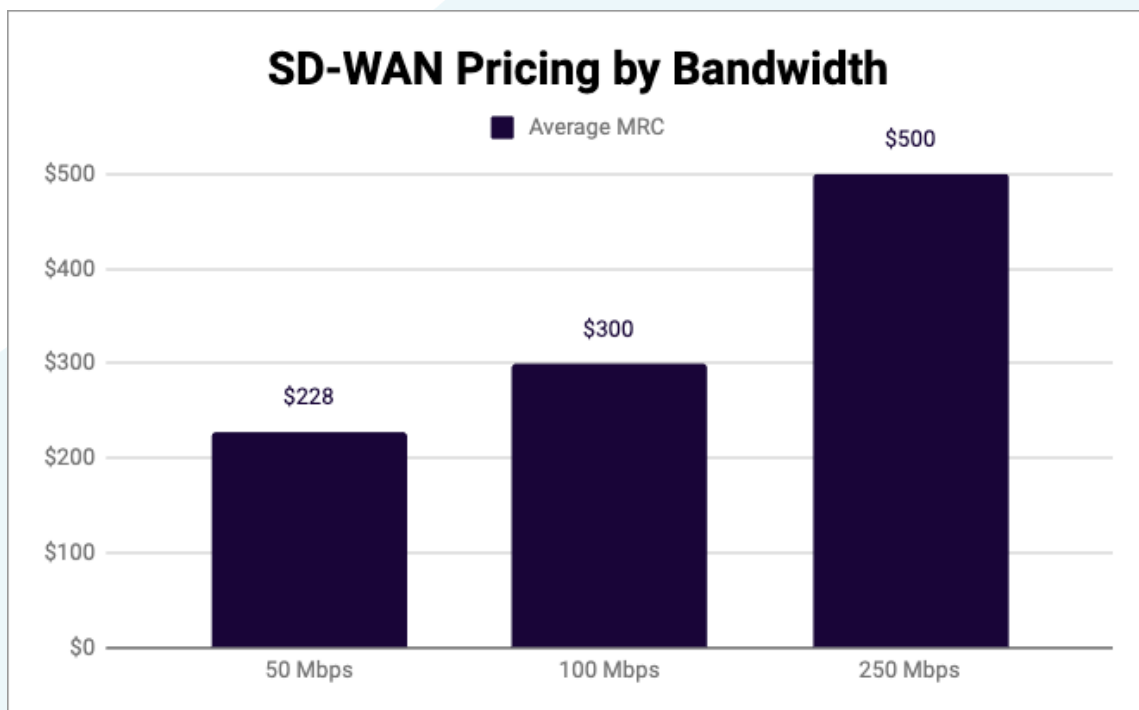
SD-WAN chart data constants: *36 month term*

Wide Area Networking Glossary

We've now discussed the pros and cons of each WAN connectivity type and how they are priced. For the rest of this report, we're going to define and discuss the key terms that impact your pricing (and some additional definitions, for fun).

3 Factors That Don't Impact WAN Pricing

1. Service Level Agreements (SLAs)

Service level agreements are legally binding documents that outline network performance and usage metrics for ISPs and TSPs. They define the relationship between two business parties and guarantee certain types of service across the board. The presence of an SLA shouldn't impact your WAN pricing as it is seen as a normal component of the package nowadays.

2. Transport Type

In today's WAN connectivity landscape, virtually everything is transported over fiber. Given the prevalence of fiber lines (especially in the U.S.), the type of transport (copper vs. fiber) your WAN utilizes will not substantially impact your pricing.

3. Handoff details

Handoff details concern where the connection from the carrier's network goes to the customer hardware. Fiber or electrical—electrical is a fancy way to say Ethernet. Handoff rarely has an implication on price. What you do see is sometimes there's inflexibility; sometimes electrical just isn't an option.

So, if you're an old network with old boxes, your electrical transport might not be able to peer with the carrier.

Other WAN pricing factors

Class of Service (CoS)

Class of Service (CoS) is a way to manage traffic on a network. Similar types of traffic, like email or streaming content, are grouped together. Each of these groups then gets a specific priority in the network. While CoS prioritization can't guarantee bandwidth or delivery times on a network, it can improve/maintain processing time for critical groups of information when the network is congested.

Managed router

Managed routers are routers that are accessed, monitored, and manipulated remotely by an independent entity, which may or may not be related to the owners of a network or those who use it.

Protected circuits

Protected circuits are components of network path protection, the goal of which is to ensure failure at certain points on a network mitigates or minimizes performance issues for users. Protected circuits are circuits that contain backup paths. Meaning, if one circuit path fails, the circuit itself does not fail.

Wide Area Networking Glossary (Cont'd)

While this is a great option to ensure service availability, protected circuits often cost more than non-protected circuits.

Quality of Service (QoS)

Quality of Service (QoS) is no single product or application. Instead, QoS refers to the concept of shaping and prioritizing network traffic as a whole. And the application of QoS involves a set of technologies that intelligently group and prioritize similar types of traffic, similar to CoS. But QoS places all non-priority traffic in a queue, ensuring priority traffic makes it to its destination before anything else is sent. This finer degree of control more directly reduces packet loss, latency, and jitter on a network.

Type 1 vs. Type 2 providers

Type 1 WAN providers own both the infrastructure the network is built on and the service it provides. Type 2 WAN providers don't own the infrastructure their service provides. Instead, Type 2 WAN providers will lease network infrastructure from a third party, which should be clearly identified in their SLAs.

When procuring Type 2 circuits, be sure to diligence the cost. As Type 2 providers are essentially resellers, there's bound to be an MRC mark-up.

Other WAN Terms to Know
Latency

While bandwidth is the amount of information that can be passed between points on a network (i.e., how much), connectivity speed refers to the rate that information can travel (i.e., how fast). Speed expectations will be defined in your SLA in terms of latency.

Latency is a measure of the time it takes one piece of information (a data packet) to make a round trip between two points. This time is often measured in milliseconds (ms). And, while zero latency (i.e., instantaneous transmission) may be impossible, ideal network configurations produce as little latency as possible.

Network topology

A network is the connection between two or more points, referred to as nodes. Network topology, in turn, is the term for how all the nodes in a network are arranged in order for it to function.

Redundant circuits

Unlike multiple routes within one circuit (i.e., protected), redundant circuits are failovers designed to automatically take over if the primary circuit fails completely. Whether your network is redundant or not does not impact the cost of the primary circuit you are procuring, but we wanted to point out that sometimes your secondary aka redundant circuit can be cheaper than your primary circuit.

lightyear

Thanks for reading!

Would you like to learn more about how Lightyear can help you with your telecom procurement and contract management needs?

Talk to us

lightyear