Transitioning to IP

Moving full-time channel distribution from satellite to terrestrial IP networks without losing sleep

An introduction to LTN Dynamic Multicarrier Routing™ architecture

A white paper by LTN

August 2023

ltnglobal.com

Table of contents

Executive summary

Technology innovation and shifts in consumer behavior have created a period of rapid change in the media industry. As broadband and wireless internet have expanded worldwide and an increasing variety of connected consumer devices have come to market, the decades-old model of delivering one version of content to a mass audience through traditional broadcast, cable, and direct-to-home satellite channels has become inadequate. Content companies now need to reach consumers through new distribution channels, including social media, direct-to-consumer websites, and over-the-top (OTT) platforms. Companies also need to create multiple versions of their programs tailored to a diverse set of devices, geographies, languages, and cultures.

This new level of content customization has put pressure on the professional video transport chain, which traditionally has relied on satellite or fiber links to deliver content from a programmer to the distribution outlets they do business with. The high cost and limited capacity of satellite transmission restricts the number and quality of program variants that a content company can reasonably provide to consumers. At the same time, satellite frequencies are being sought by regulatory bodies worldwide as a means to deliver new wireless services, which will likely reduce the capacity available to traditional media companies for content distribution.

Given these trends, content companies are exploring IP-based terrestrial transmission — including both private networks and the public internet — as an alternative means of delivering their product. However, private IP networks such as MPLS are expensive, while the public internet has some basic architectural limitations that make it difficult to reliably send high-value content, particularly live programming.

However, the challenges of the internet's underlying architecture can be overcome — and reliable, cost-effective delivery of live content achieved — through the application of unique overlay technology and the development of a professional network of data centers that use proprietary routing protocols. LTN has achieved just that, creating a fully managed service that delivers live content today for some of the world's biggest media companies with equal or better reliability than traditional satellite or fiber paths.

LTN offers the industry's only SLA guarantee for internet-based video transport that covers five-9s reliability, availability, and <200 ms latency in North America; <300 ms worldwide; across the LTN Network. This reliability compares favorably to satellite transport, and is critical for the transport of high-value content distribution. Our patented LTN Dynamic Multicarrier Routing (DMR) algorithms and LTN Rapid Error Recovery™ (RER) protocols enable customer traffic to withstand any carrier failure, congestion, or other network unavailability.

This white paper will examine the content delivery challenge faced by media companies today; describe the potential solution of the public internet and the inherent difficulties it bears; and more importantly, explain the unique, patented technology and powerful data network that LTN has created to solve the problem.

Immunization against any single carrier's failure

All LTN backbone nodes are (a) connected to multiple Tier 1 carrier backbones, and (b) constantly aware of the real-time state of every link of every available path for a customer feed. These capabilities ensure that any single carrier's failure or service degradation results in LTN Dynamic Multicarrier Routing™ (DMR) algorithms rapidly rerouting traffic over another carrier, without interrupting the customer's feed.

Multiple levels of redundancy and resilience

The LTN Transport solution includes multiple levels of redundancy in hardware, software, internet last-mile connectivity, and backbone routes. This redundancy ensures reliable delivery of customer feeds, with no single point of failure in the end-to-end path.

Fully managed service

LTN supports its innovative network solution with a fully managed service, including a always-on technical operations center (TOC) staffed by highly trained technical personnel who are able to monitor, isolate, diagnose, and fix any issues on the LTN Network affecting a customer's feed, and are available as a 24/7 customer help desk.

High-visibility monitoring tools

LTN offers a large portfolio of monitoring and diagnosis tools to both our customers and our TOC personnel. The tools provide detailed visibility and control over the performance of each individual feed on the LTN Network. We monitor video equipment, LTN Leafs and Appliances, internet links, the end-to-end path over the LTN Network, and the metrics and performance of delivered feeds at the receive sites.

Content evolution and transition to IP transport

This is a time of rapid change in the media industry — especially in video transport and delivery. The content chain from playout all the way to the consumer is evolving, driven by significant shifts in content regionalization, technology, regulation, and consumer behavior.

Content providers now reach consumers through a variety of distribution channels beyond traditional cable and direct-to-home satellite outlets, including social media sites, direct-to-consumer websites, over-the-top (OTT) platforms, and international distribution sites. To optimize content monetization, providers are also pursuing consumers and advertisers across diverse geographies, devices, languages, and cultures.

Each of these points of distribution can require varying levels of content customization:

- Closed captioning
- Audio translation
- Graphics
- Advertisements
- Subtitling
- Sometimes even changes in the underlying content

As a result, some content providers are exploring cloud-based content customization to cost-effectively produce multiple versions and efficiently adapt content to suit local markets and various devices.

These changes in content preparation are creating a profound impact on program distribution. The days of broadcasting one version of content to all takers are rapidly disappearing, replaced by sending customized versions unique to each specific destination. Historically, content owners have provided their content on satellite transponders in the geographies where they did business, and have left it to their distribution outlets to downlink the content and deliver it to consumers. The expense of satellite transport limits the number of possible channel variations, and the amount of satellite bandwidth leased by the content provider puts a cap on content quality. These limitations make it difficult for content providers to rely solely on satellite transport to handle the growth in channel variants.

30%

Over the past few decades, last-mile internet bandwidth has continued to grow at the rate of about 30% each year, while cost of bandwidth has been declining at about 30% per year.

Full content monetization

increasingly requires content customization for geographies, devices, languages, and cultures

5G bandwidth appetite

is reducing satellite supply for TV content transport

Regulatory bodies have repurposed part of the spectrum used by C-band satellite operators for 5G wireless carriers, worsening the capacity crunch. Regardless of how the repurposing plays out, overall satellite capacity will decrease at a time when the need is increasing.

Content providers recognize these fundamental trends and have started to explore alternative methods of delivering their programming to distribution outlets. IP-enabled terrestrial transmission paths are the most obvious and attractive alternatives, given their flexibility to route traffic easily from anywhere on the network to everywhere without having to create a physical end-to-end path.

The costs of IP-based networks vary significantly by architecture. However, the economics are generally favorable compared to satellite, especially when there's an emphasis on increasing content regionalization or localization. However, the more private the network, the higher the cost of transport.

MPLS networks fall into this category as they require dedicated connections from each endpoint into the MPLS backbone. This becomes very expensive as the number of endpoints grow. Additionally, MPLS service is limited to a single carrier network, which exposes the service to greater risk if the carrier fails. For these reasons, MPLS networks generally aren't practical for broad content distribution.

Traditional linear television and over-the-top (OTT) platforms

will start to merge over the next few years

Analysys Mason European Telecoms Summit 2017

The migration to IP is happening and it affects everything

Potential evolution of TV services across different infrastructure platform, 2015-2030

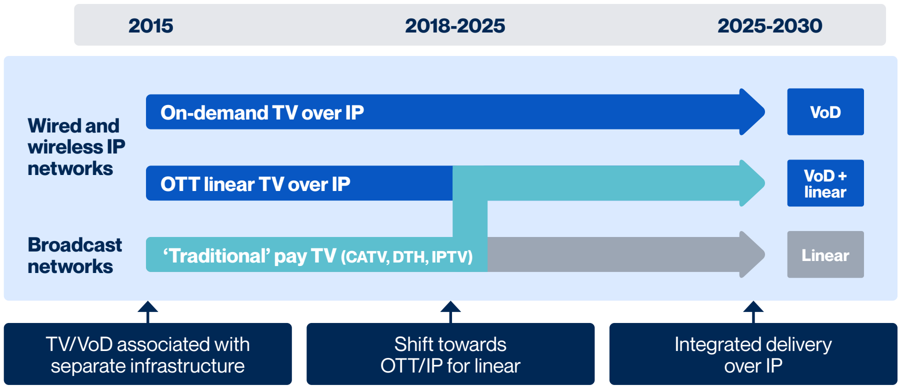

Internet-based delivery holds the most promise for advanced content transport. It's ubiquitously available and is growing worldwide. In fact, the available last-mile internet bandwidth has increased approximately 30% every year over the last few decades, while its cost per unit has declined at approximately 30% per year within the same timeframe.

The internet has supported multiple messaging methods over the last 20 years. Mail has morphed to email, faxes to PDFs, and landline phone calls to Voice over IP. We're now experiencing explosive growth in the amount of video being transported over the internet. Most of this growth over the last two decades has been due to on-demand video, but live and real-time internet video traffic is growing rapidly.

The raw internet wasn't designed to support real-time traffic

However, there are fundamental issues that must be addressed before the internet can be used to reliably and securely transport high-value live or real-time professional content. Analysts are concerned about the ability of the internet to reliably transport live video. The trend toward internet-based transport for even live and real-time video content is clear, but it's also clear that the raw internet wasn't designed to support real-time traffic.

Internet architecture and its limitations

In this context, it's helpful to look briefly at how the internet is architected. What we all think of as a singular internet is really a collection of approximately 60,000 independent internet service provider (ISP) networks worldwide. The ISPs all follow the same general principles in how they build and operate their networks and how they connect to each other.

Each ISP network is built using IP protocols running over a physical network consisting of fiber connectivity with routers and switches that route traffic internally — normally using the open shortest path first (OSPF) protocol. The protocol's job is to select the shortest reliable path from the ingest point to the egress point on the ISP's own network.

One important constraint is that these ISP networks have a form of hierarchy, where the amount of bandwidth offered to customers always greatly exceeds the bandwidth available in the core of the network. Because these internet last miles are only used by a fraction of the total users simultaneously, this normally isn't an issue. However, during periods of unusually high use, customer traffic aggregation points can get overwhelmed, resulting in traffic slowing down in long queues.

Once a queue is formed, all packets in the queue are delayed. Because internet routing protocols are designed to route only on link failure — not congestion — packets that are stuck in a queue simply have to wait until the queue is cleared. Additionally, when queue lengths exceed the capacity of switches, packets are dropped.

When the destination application detects dropped packets and attempts to recover them to support reliable communication, further delays occur. Until the destination application sends a retransmission request all the way back to the source, the dropped packets and all following packets are delayed. Once the dropped packet is recovered, all packets can be delivered in order.

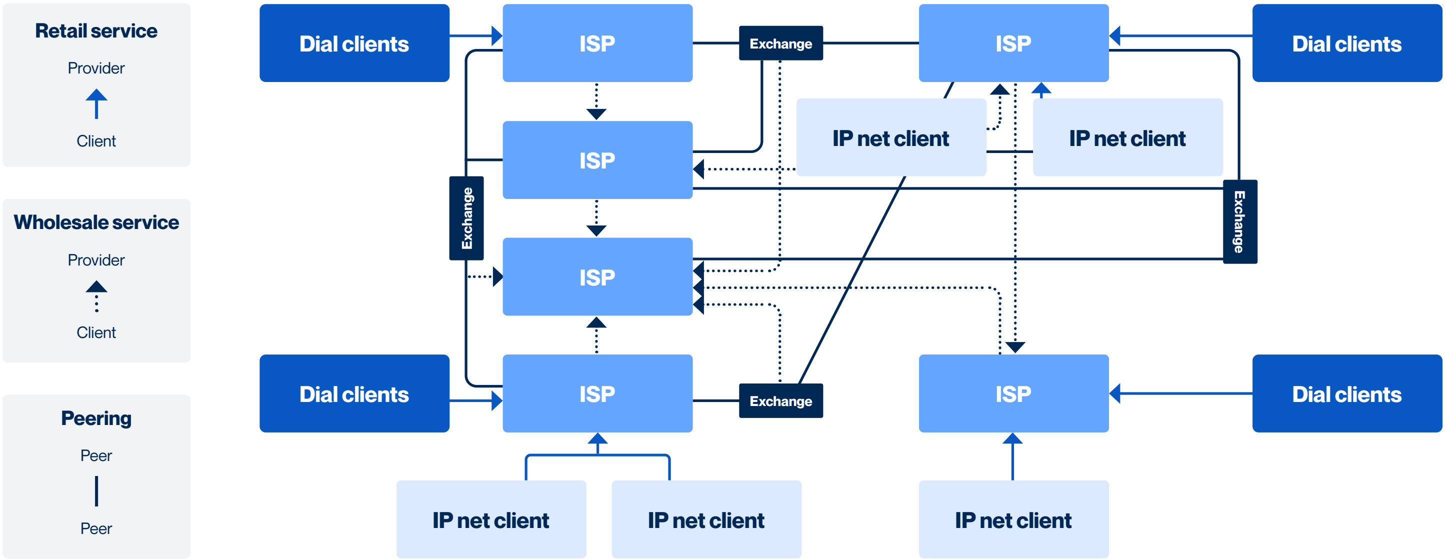

Since most internet users look for connectivity to other users or servers that normally aren't on the same ISP network, independent ISPs have to frequently exchange traffic with each other. The exchange is facilitated by a border gateway protocol (BGP) whose policies are usually commercial. Two equal-sized ISPs may agree to exchange content at no cost to each other as long as they send and receive roughly equal traffic between each other in private or public peering arrangements. Non-equal ISPs frequently sign transit agreements that involve payments.

In either case, traffic arrangements are frequently capped in capacity. While the caps are built with real traffic trends in mind, they are rarely designed to handle peak traffic. During peak traffic conditions, an ISP may deliberately send traffic to non-optimal transit or peering points. Other times, traffic may get stuck behind long queues because the bandwidth is capped. Both types of events cause the delay and possible packet loss described above.

Internet architecture is built around the principle of a simple and efficient core, with most of the internet routing intelligence built on the edge. For example, transmission control protocol (TCP)-based endpoints request lost or missing packets, and wait until those packets eventually show up. This, in turn, allows internet traffic to get from one end to the other reliably, but with an uncertain delay.

A special protocol called user datagram protocol (UDP) was created for applications that require low delay and are loss-tolerant. UDP does not attempt to recover packets on an end-to-end basis, however, which means lost packets remain lost. The protocol trades reliability for low delay.

Internet routing architecture

can't deliver both high reliability and low delay

Neither transport protocol serves the need of live or real-time video applications, which must both be reliably received and stay within a narrow latency window. This challenge is magnified by the protocols' reliance on internet routing with its problems mentioned above. Handling these real-time applications over the internet requires more processing than available within the internet's simple core.

An additional architectural limitation to the internet fundamentally limits its use for reliable content delivery to multiple destinations. The internet doesn't have the native ability to multicast — i.e., to send the same packet from one source to multiple receivers. The internet was built as a unicast system, which means that to send a single video feed to multiple destinations, the source location must send multiple copies of the same content — one copy for each destination. Once the number of receivers grows beyond just a few, this process becomes unusable and impractical.

Inherent limitations of the internet

make it less than practical as a substitute for satellite delivery without some way of overcoming these limitations

Customers considering moving from satellite to an internet-based transport solution must take great care to ensure that their chosen internet solution has the ability to overcome these challenges.

Many of these transport issues occur in the middle of the internet and thus cannot be solved by technology that sits at the edge of the internet.

Without overcoming these inherent limitations, the internet is impractical as a substitute for satellite delivery. Customers considering moving from satellite to an internet-based transport solution must ensure that their chosen internet solution has the ability to overcome these challenges. Many of these transport issues occur in the middle of the internet and cannot be solved by technology that sits at the edge.

LTN's highly resilient Dynamic Multicarrier Routing architecture

To solve these internet limitations, LTN has developed innovative and patented technology, including LTN Rapid Error Recovery (RER) protocols and LTN Dynamic Multicarrier Routing (DMR) algorithms and architecture. LTN has also built a global backbone network around this technology. The two-pronged effort allows us to provide our customers a high-reliability, fully managed video transport service over the internet.

This robust transport service also allows us to offer our customers very exacting service level agreements (SLAs), with metrics that include reliability, availability and latency that match or exceed those available from satellite or fiber service. LTN's service has already been deployed by many of the world's largest and most demanding media companies to distribute high-value, full-time channels worldwide.

The diagram below shows a high-level view of the LTN Network. Client appliances connect to data centers on the LTN Network, and each datacenter has access to multiple carrier networks. Traffic can flow between clients and datacenters or between datacenters over multiple carrier paths.

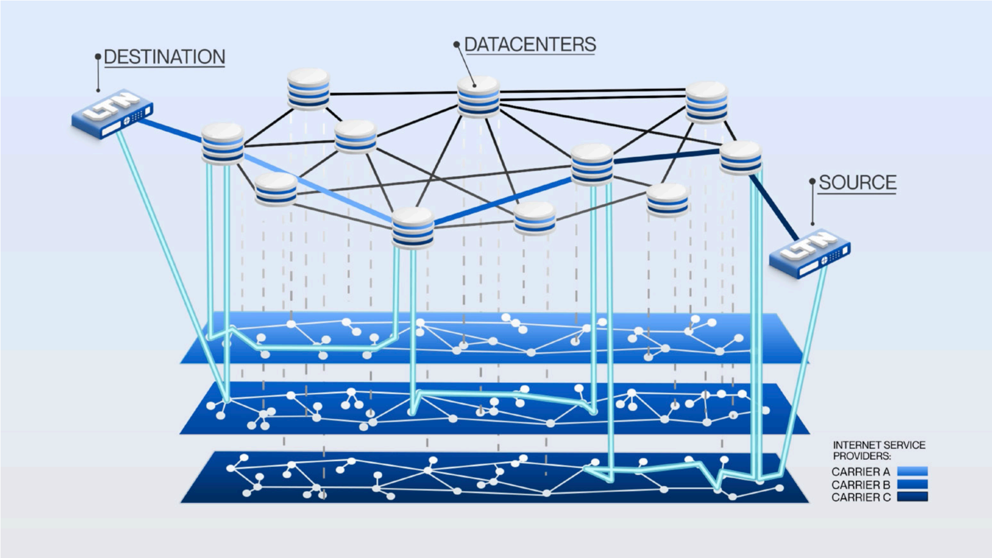

LTN proprietary DMR

Algorithms automatically and losslessly route live customer traffic around congested or delayed carrier paths

LTN DMR network architecture is built on important underlying principles.

LTN puts routing and processing intelligence in the middle of the internet located in high-availability and secure datacenters all over the world. This intelligence allows the LTN Network to process data, monitor traffic, and make decisions that otherwise would not be possible over the internet. For example, LTN runs a proprietary error-correction protocol, LTN Rapid Error Recovery (RER) on each individual segment of the end-to-end path on the LTN Network. LTN RER enables rapid loss recovery in 10-20 milliseconds, rather than normal end-to-end recovery between the source site and a receive site that may take hundreds — or even thousands — of milliseconds. This proprietary feature allows LTN to withstand significant underlying loss on the internet link.

Another benefit of having this intelligence in the middle of the internet is that it allows LTN to decide which datacenter is closest to each receiver out of a group of receivers, and to serve out multiple copies of a feed only at those last-mile datacenters. As a result, all other datacenters act as transit points and simply route a single copy of the feed to the next datacenter on the flow route. This innovation in data routing allows LTN to broadcast live and real-time video feeds reliably and with low latency to hundreds — and even thousands — of end points efficiently and cost-effectively.

LTN DMR

creates high availability, end- to-end connectivity by using multiple carrier backbones

LTN RER

protocols enable rapid loss recovery in 10-20 milliseconds

LTN DMR architecture protects customers against the risk of carrier failure by routing on multiple carrier backbones.

DMR ensures reliable, low latency transport of content across the LTN network, through instantaneous decision making algorithms routing traffic across the optimal paths across multiple carriers.

You can see the LTN DMR in action in the LTN Network diagram above:

01 A feed enters the LTN Network into a datacenter from a client shown on the right.

02 The LTN Network chooses the carrier plane — shown in red — as the primary path to receive and send the feed to the next datacenter on its way to the destination.

03 At the second data center, the Network detects performance degradation on the red carrier.

04 The Network switches the feed to the blue carrier, and sends the feed to the next datacenter in the chain on that blue carrier plane.

05 At the third datacenter, the network detects an issue with the blue carrier leaving the third datacenter

06 The Network reroutes the feed to the green carrier plan, which is also how the destination client receives it.

The LTN DMR enables the customer's feed to be rerouted over three different Tier 1 carrier backbones end-to-end, automatically, with no impact on the users, and within the performance SLAs of 99.999% or higher availability and reliability with <200 ms latency in North America and <300 ms worldwide across the LTN Network.

Without the LTN Network monitoring and intelligently rerouting traffic in real-time around congestion, contention, long queues, or unavailable links, the feed would have suffered from loss, long latency, jitter, or unavailability.

Dynamic Multicarrier Routing principles

01 Routing and processing intelligence in the middle of the internet

02 The LTN Network chooses the carrier plane — shown in light blue — as the primary path to receive and send the feed to the next datacenter on its way to the destination.

03 Flow-based architecture ensures packet protection within each flow

Always-on technical operations center

provides real-time troubleshooting and support, monitoring and managing the network full-time. With full network visibility, the TOC often isolates and fixes any issues before there's any customer impact.

The LTN architecture is a flow-based overlay over the internet rather than a packet-based forwarding system. All packets belonging to a specific flow — like a sporting event, entertainment channel or newsfeed — are routed based on the policies for that flow. LTN deploys a fast, accurate messaging system between its datacenters that allows all servers to share current-state information about all flows, ensuring that if a flow is rerouted over a secondary path, then all affected datacenters on that new path are fully aware of it. This messaging approach allows real-time algorithms to reroute flows in milliseconds with confidence that the LTN datacenters on the new route will have a complete understanding of the policies governing that flow and deliver it without interruption to the end customers.

To be clear, the LTN Network is an overlay on top of the internet network that requires no changes at any layer of the internet infrastructure to operate. All LTN algorithms, protocols, and actions are performed at the overlay level.

There is one more important factor in ensuring LTN's reliability as a high-value video transport service. LTN has invested in an always-on technical operations center (TOC) staffed by highly trained and experienced technical staff. The TOC is supported by proprietary monitoring, reporting, and management tools that allow the staff to quickly isolate and fix issues, including those caused by the customer's own video/LAN infrastructure or by any of the ISPs involved in the flow. This critical investment allows LTN to meet its exacting SLAs to customers and ensures customers always have LTN support personnel to rely on if any issues occur. Our TOC will often isolate and fix issues before there's any customer impact.

Although the vast majority of problems addressed by our TOC turn out not to be related to the LTN Network, our extensive, end-to-end view of the flow allows us to identify problems all the way to customer sites and video networks. LTN also makes monitoring tools available to customers so they can see the health of their traffic over the LTN Network.

LTN Network routing data architecture

Because of the LTN Network architecture and placement inside the core of the internet, we're able to collect real traffic data on internet congestion and document how LTN algorithms have rerouted flows to avoid loss or delay that would have been caused by congestion on the primary path.

To understand the data, it's helpful to view the flow of customer traffic over the LTN Network as shown in the diagram below.

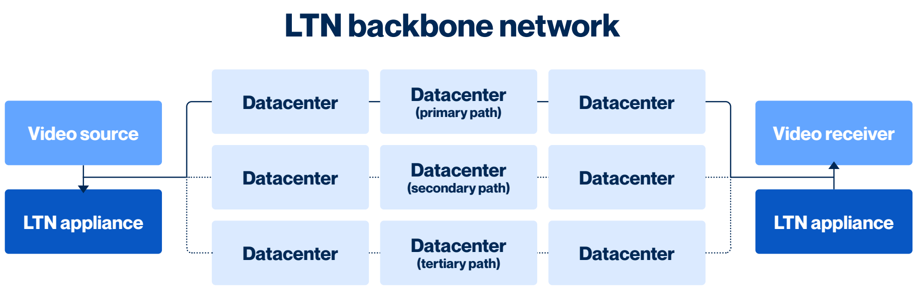

LTN multiple path resilient architecture

LTN measures and monitors the health of each link every 70 milliseconds.

This diagram shows a simple example of how flows can be routed from a source to a destination over different paths. Although it shows only three physically separate possible paths, a greater number of paths is possible. Additionally, all paths are available for rerouting, allowing LTN to deliver customer feeds with unprecedented resilience and reliability.

LTN provides customers with an on-premise appliance that connects to multiple ingress points from the customer's location for redundancy. The appliance also has the ability to terminate on multiple ISP backbones within each datacenter. This allows traffic from the source site to reach the LTN Network over multiple paths. The LTN Network automatically chooses the optimal path as a primary route and measures and monitors the health of each link every 70 milliseconds. If the health of a link on the primary route shows signs of degradation, the LTN Network will automatically reroute the traffic to a second path. It's sometimes necessary to reroute the traffic to a third, fourth, or greater path to ensure that the flow metrics stay within our SLAs. These reroutes are made without impact to customer-received traffic. This process exists for the last-mile connection to both the sender and receive site(s).

For the purposes of this white paper, we show flow data collected for two consecutive months in the tables below. Our experience is that this data is consistent month-over-month, and this representative sample was collected to show detailed analysis of traffic patterns over the LTN Network.

Month 1	Full-time channels	Occasional use feeds
# of end points	1,780	5,704
% of feeds rerouted	29.94%	8.63%
Average # of paths	2.58	2.29

Month 2	Full-time channels	Occasional use feeds
# of end points	1,808	6,602
% of feeds rerouted	33.57%	9.02%
Average # of paths	2.57	2.25

The data is divided into two columns, one for full-time channels and the other for occasional use or ad-hoc traffic. The number of end points shows the number of clients for each service during that month. The '% feeds rerouted' refers to feeds that were routed away from the primary path for more than 5% of the time; for full-time channels, that equates to at least 36 hours per month. Reroute data for less than that duration wasn't counted in this data. 'Average # of paths' shows the average number of alternative routes that were used more than 5% of the time. Traffic that was sent only on a secondary path would register as having had two paths. Tertiary path traffic would be recorded as three paths. Additionally, to register as a tertiary path, the traffic would have had to be on the tertiary path for more than 5% of the time.

Without the LTN overlay network,

full-time channels would experience more than 36 hours of delay or loss per month over the internet

LTN real traffic routing behavior analysis

There are some important takeaways from the data in the tables above. Reroute percentages are quite high across the board, but especially for full-time channels. This is an important finding since the LTN Network reroutes traffic only when packet loss is unavoidable on the primary path — after all attempts at LTN RER packet loss recovery have been deployed.

The LTN first line of defense is to recover lost packets within a preset delay budget. LTN RER protocols for this loss recovery are very advanced and, more importantly, operate on short hops between customer premise appliances and the nearest data centers, as well as between data centers, which are normally 10– 20 milliseconds apart. This results in rapid and efficient packet loss recovery, something that isn't possible if a customer deploys boxes at its premises without a network in the middle.

However, despite the hop-by-hop packet loss recovery effectiveness of the LTN RER protocols, our second line of defense, the LTN DMR algorithms, still switches away from the primary path for a quarter to a third of the full-time channels in any given month. Had the LTN Network not rerouted the traffic, it would have certainly experienced loss as well as uncertain and variable delay. Internet routing would have waited for a disconnection to reroute traffic, which is accompanied by a long wait — from approximately 40 seconds to minutes — for convergence. Without the LTN Network, content flows would have to depend on internet routing, which isn't designed to route around congestion, loss, increased delay, or any of the normal issues that occur frequently over the internet's traffic aggregation, transit, or peering infrastructure. Internet routing protocols like OSPF and BGP will only route around links which are broken or otherwise not available, and there's a waiting period for reconvergence.

Edge technologies attempt to counter this reality by either deploying very large buffers — several seconds or even tens of seconds — sending multiple packets of the same feed, using methods like forward-error correction (FEC), or simply deploying end-to-end selective packet recovery protocols. These methods all have serious drawbacks. They are not effective since they rely on internet routing in the middle and therefore have no recourse if the traffic is blocked at a congested aggregation router, peering point, or transit connection. Sending multiple copies of the same feed or retransmitting delayed packets doesn't help, and can even worsen the situation if there's already a long queue at the congested point.

The only effective solution is to route around these points of congestion much like motorists use apps like Waze to find alternative paths around congested routes.

Without LTN RER hop-by-hop error recovery protocols and LTN DMR rerouting algorithms, content traffic will experience loss and uncertain delays, especially for full-time flows. This is non-negotiable for live and real-time high-value content delivery.

By guaranteeing arrival time with consistent reliability, LTN goes one better than road traffic rerouting systems. This is, in part, because traffic on the road is restricted to a single two-dimensional system of possible paths to the destination while LTN adds the third dimension of multiple carrier planes so our traffic can move from one plane to another rapidly in response to issues. This three-dimensional mesh — and the intelligence to gather information and make routing decisions in real time — allows LTN to make its reliability, availability, and latency guarantees.

An interesting contrast in the data is the difference between full-time channels and occasional-use traffic in terms of the percentage of rerouted feeds. At LTN, most of the occasional-use traffic we see is less than six hours in duration. There are exceptions, such as multi-day events, but they represent the minority of traffic. Internet congestion is lumpy — it's not evenly spread across ISPs, geographies, or connections. Full-time traffic has a 100% guarantee of seeing all such events that occur across their primary path, while the probabilities are much lower for ad hoc traffic since it has a much shorter duration. Nevertheless, for high-value content like sports, esports, concerts, summits, breaking news, or other critical events, the risk of serious disruption implied by the 8–9% reroute data is often too high.

High-value, live, and real-time content distribution has exacting requirements for consistency, reliability, and delay. Guaranteed reliability requires an intelligent overlay network that can track, route, and reroute flows over multiple carrier backbones in milliseconds to ensure delivery under any internet conditions.

In fact, as the data shows, more than one reroute is often necessary to ensure proper delivery metrics. These requirements can't be met by deploying smart boxes at the edge, connecting them to the internet, and relying on internet routing to get the traffic delivered across the network. Customers looking for high reliability alternatives to satellite transport need service providers to make the necessary infrastructure investment, including installing intelligence at datacenters worldwide, purchasing hundreds of gigabits per second of bandwidth from multiple Tier 1 ISPs around the world, and building always-on TOCs to manage each customer's traffic individually — all of which LTN has already implemented successfully.

Benefits of LTN Network

SLA guarantees for performance

LTN provides customers with the internet-based video transport industry's only SLA guarantees on reliability, availability, and latency. Our patented and innovative LTN DMR algorithms and LTN RER protocols enable LTN to deliver guarantees to customers that are essential for high-value content distribution — content that's expected to be delivered on time, on spec, and with consistent reliability.

Multiple levels of redundancy and resilience

The LTN Network includes multiple levels of redundancy in hardware, software, internet last-mile connectivity, and backbone routes, ensuring reliable delivery of our customers' video feeds with no single point of failure anywhere in the end-to-end path.

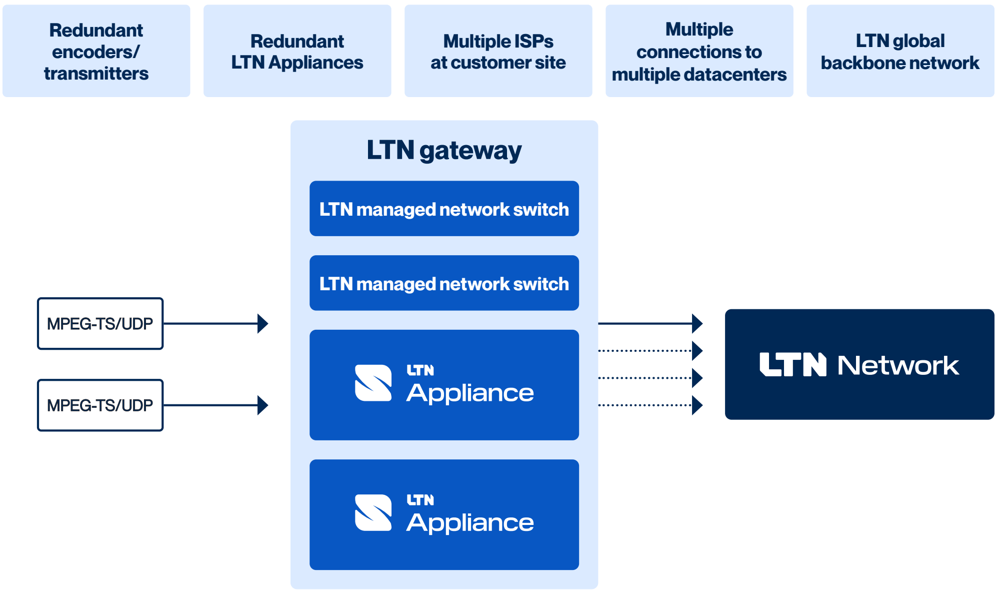

Automatic failover, timely recovery, and seamless rerouting

Immunization against any single carrier's failure

All LTN backbone network nodes are a) constantly and fully aware of the real-time state of every feed, regardless of whether or not they are on the primary path of that feed's delivery; and b) connected to multiple 10 Gbps Tier 1 carrier backbones. This combination ensures that any single carrier's service degradation or failure will result in LTN Dynamic Multicarrier Routing algorithms rapidly rerouting traffic over another carrier without interrupting the customer's feed. Private fiber and MPLS networks, on the other hand, are single-carrier networks with the attendant risk of failure.

Universal access from anywhere in the world

Customers access the LTN managed service via commercially available internet connectivity from any site, anywhere in the world, enabling them to easily, rapidly, and cost-effectively transmit or receive live and real-time video feeds under high performance and reliability SLAs.

Highly scalable in number of channels and destinations per channel

The LTN Network infrastructure is built with the capacity to transmit thousands of simultaneous feeds to tens of thousands of sites. It is also architected to grow and scale beyond that capacity, taking advantage of a distributed system of datacenter capacity and bandwidth supported by robust messaging systems, protocols, and routing algorithms.

> **True multicast network**
>
> LTN solutions scale to thousands of endpoints for a channel.

Always-on network operations center

engineers monitor and manage the network, providing real-time troubleshooting and support.

Always-on expert customer support

LTN supports its innovative network solution with a fully managed service, including an always-on technical operations center (TOC) staffed by highly trained technical personnel who monitor, isolate, diagnose, and fix any issues on the LTN Network affecting a customer's feed, and available as a 24/7 help desk for our customers.

High-visibility monitoring tools to control the performance of flows

LTN has developed a large portfolio of monitoring and diagnosis tools that allow our TOC, as well as our customers, to have an unprecedented degree of visibility and control over the performance of each individual feed on the LTN Network. This includes monitoring video equipment on the customer's premises, LTN Appliances, internet links, the end-to-end path over the LTN Network, and the metrics and performance of the feed at the received site. The LTN TOC spends the majority of its time solving issues outside its own network domain.

The diagram below is a screenshot from the LTN last-mile internet monitoring tool that gives its TOC personnel's customers visual and metrics-based feedback on the performance of the underlying internet link, and allows LTN to work with the ISP to resolve any issues on their own infrastructure.

Standards-compliant encoding and transport

LTN technology is fully compliant with broadcast standards for transmission, such as MPEG-TS streams (ISO/IEC standard 13818-1), use of MPEG2 (ITU-T Rec. H.262), MPEG4/AVC (ISO/ IEC 14496-10 and ITU H.264), and HEVC (ITU H.265) video compression, and it interoperates with all standard broadcast video equipment. This interoperability enables customers to connect their traffic to the LTN Network in an open and flexible manner, without any concerns about proprietary walled gardens. Customers also have the freedom to choose any encoding equipment or format as long as it is MPEG-TS compliant.

Secure and protected flows

LTN provides secure and protected feeds over its Network. LTN has the ability to encrypt flows which can only be decrypted by authorized receiving Appliances. The LTN content protection uses standard, well-understood, and respected cryptographic algorithms such as AES-128 for encryption and RSA-2048 for public-private keys.

LTN fully manages authorization and authentication of each channel to ensure customers control which receivers are allowed to receive specific channels, giving them dynamic control of their content's security. Additionally, LTN backbone nodes are located in highly-secure data centers with secure entry, back-up power, and redundant HVAC facilities.

LTN project management

LTN has accumulated substantial experience over the last decade in configuring, installing, activating, and supporting customer feeds in networks comprising hundreds of locations. This mature process ensures timely activation of new channels or new sites in a reliable, consistent, and repeatable manner.

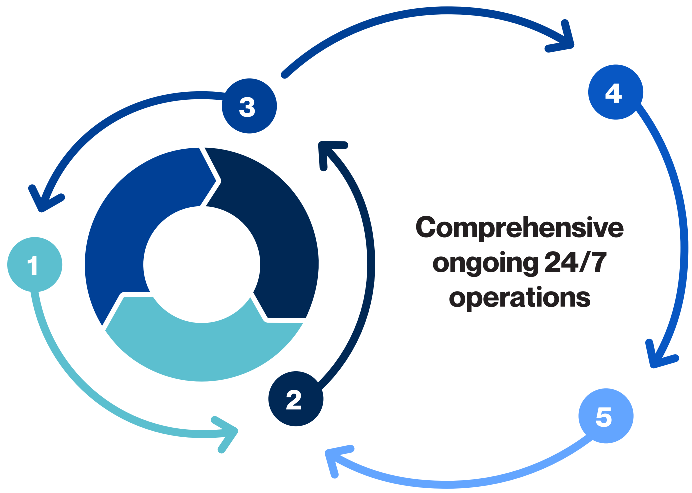

Detailed, meticulous, and holistic project execution.

1. Requirements
2. Design
3. Implementation
4. Operational readiness
5. Ongoing ops and monitoring

Flexible and cost-effective pricing model

Many of our customers use LTN for more than a substitute for existing fiber or satellite infrastructure — but also as a provider of next-generation video transport technology and service that enables new business models. LTN pricing isn't just more cost-effective than traditional transport services, it also has the flexibility to enable successful customer business models in a world where content customization and regionalization is critical to business growth. Our pricing models enable customers to send multiple geography-, device-, or market-specific versions of their channels, with a cost in proportion to the number of receivers rather than to the number of versions of the channels being sent. Our business success is built on our customers' success.

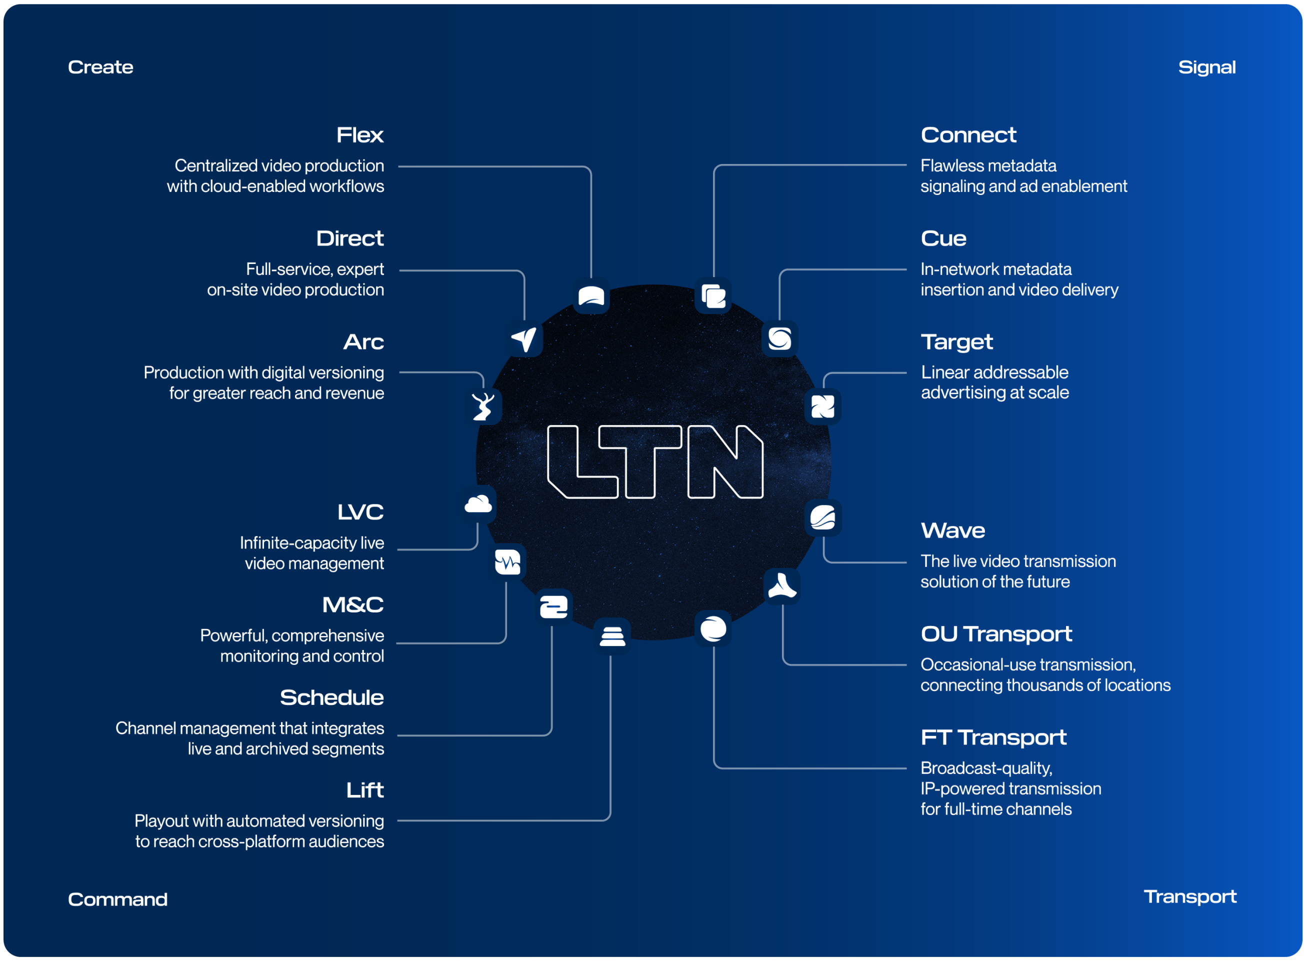

The LTN difference

LTN was founded in 2008 by three experts with nearly a century of combined experience in the IP and internet world. Malik Khan led the division at Motorola that built the first commercially deployed cable modems and CMTS systems and spurred the advent of residential broadband internet access. Yousef Javadi was the head of Sprint International and led all of Sprint's Global entities outside the US, including IP networking, hosting, and other infrastructure. Professor Yair Amir[1] is the former Chairman of the Computer Science Department at Johns Hopkins University and is recognized as one of the leading researchers in large IP systems architecture. This founding team assembled a uniquely qualified research, development, operations, and customer solutions team at LTN to creatine the industry's first terrestrial IP network for live video and audio transport at a global scale.

1 - Participation by Dr. Amir as an officer does not constitute or imply endorsement by Johns Hopkins University.

LTN was founded on a vision that the future of video content distribution would take place over the internet, using new technology to build an overlay network on top of the internet through which high-quality content could be distributed anywhere in the world, to any device a consumer wished to use — quickly, easily, consistently, reliably, and with high quality. LTN has created the essential technology to fulfill this vision and has been granted patents in network architecture, algorithms for lossless traffic routing around internet issues, and protocols to recover lost or out-of-order packets. LTN innovation has resulted in a network that achieves both extremely high reliability and very low delay to transport live and real-time audio and video traffic anywhere in the world. The LTN Network treats each customer's program or event as an individual flow that is tracked, monitored, and managed end-to-end, allowing us to control the routing of that flow in real time.

Over the past decade, LTN has worked with some of the largest media companies in the world in news, sports, and entertainment program acquisition and distribution. Customers like Disney, ABC, ESPN, CBS, NBC, CNBC, MSNBC, Fox, Sinclair, Turner, CNN, and many others have taught us what features and functions are needed for them to transition to the future of live television programming and pursue new live-event business models and distribution strategies. Because our core strength is our research and operations capabilities, and since we own the technology we created our network and fully-managed service with, we have been able to provide a bridge to the future to our customers.

We would welcome the opportunity to work with you and add you to our growing list of connected companies and sites.

Contact us

Technological innovation is driving growth and creating new opportunities for the broadcast and media industry. New competitors are vying to capture viewers by leveraging new technologies in unique and creative ways. If you're ready to explore how you can reach a greater global audience, connect to media partners around the world, and achieve more efficient workflows, LTN is your guide.

Connect with us at **sales@ltnglobal.com**