AI and the Future of Network Design: A Comprehensive Evolution from SONET to ROADM

Introduction

The past quarter-century has seen U.S. telecommunication carriers transform their core transport networks from legacy time-division multiplexing (TDM) systems into high-capacity, IP-centric optical infrastructures. This whitepaper reviews the **evolution of transport networks** in major U.S. carriers such as AT&T, Verizon, and Lumen (formerly CenturyLink/Level 3) since 2000. It traces the transition from TDM-based SONET/SDH rings to packet-centric networks – first **IP over SONET**, then **IP directly over DWDM** (Dense Wavelength Division Multiplexing), and finally to **IP over ROADM** (Reconfigurable Optical Add-Drop Multiplexer) photonic networks. Alongside the technology evolution, we analyze three major phases of traffic growth that drove these changes: (1) the early 2000s rise of web and e-commerce, (2) the 2005–2015 explosion of video streaming, and (3) the emerging AI-driven traffic wave in the 2020s. In parallel, we examine the evolution of network planning tools – from vendor-proprietary TDM design software to sophisticated IP/MPLS planning and optical design tools – and discuss their capabilities and limitations. The analysis highlights that while planning tools have improved for "what-if" capacity and failure analyses, they have fallen short of true multi-layer, multi-vendor optimization, often serving as aids in vendor equipment deployment rather than holistic network optimization engines. This document is organized into clear sections addressing each era and theme, with key insights and industry citations to support a strategic understanding of how U.S. carriers have modernized their transport networks over time.

Transport Network Evolution: From SONET to IP over Photonics

Legacy TDM and SONET Networks (circa 2000)

In 2000, the backbone networks of major U.S. carriers were dominated by **Synchronous Optical Network (SONET)** and its international equivalent SDH. These systems were designed for TDM voice and private line circuits, providing highly reliable ring architectures and 50 ms protection switching. Carriers like AT&T and Verizon operated tens of thousands of SONET network elements (ADMs, digital cross-connects, etc.) supporting legacy DS0/DS1/DS3 circuits and OC-n rates cisco.com cisco.com. For example, Verizon's network in the 2000s contained "hundreds of thousands of circuits, ranging from DS-0 to OC-192" across **150,000+ TDM-based elements in central offices** cisco.com. These SONET backbones were engineered for the voice era – highly resilient, but static and inefficient for burgeoning data traffic. The early 2000s dot-com boom brought a surge in data demand that strained this legacy infrastructure. New **competitive fiber carriers** (Level 3, Qwest, etc.) emerged around 1998–2001 with fresh nationwide fiber builds,

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

aiming to carry Internet traffic at lower cost. Level 3, for instance, completed a 16,000-mile fiber network by 2001 and laid **excess fiber conduits** to capitalize on future optical technology advances courses.media.mit.edu courses.media.mit.edu. These newcomers envisioned dramatically lower bandwidth prices (Level 3 planned to cut prices ~50% annually to spur 3× traffic growth) courses.media.mit.edu, foreshadowing a shift away from expensive TDM/SONET capacity.

The Shift to IP over SONET (Packet over TDM)

By the early 2000s, incumbent carriers responded to Internet growth by overlaying **IP networks on top of SONET** transport. Routers were interconnected using Packet over SONET (PoS) interfaces or ATM circuits riding the SONET infrastructure. This period saw the adoption of MPLS in core networks to efficiently carry IP traffic with traffic engineering. AT&T, for example, embraced a **converged IP/MPLS backbone** as a strategic direction ("One IP/MPLS network") hit.bme.hu. The IP overlay initially leveraged SONET's reliability – e.g. routers connected via OC-48/192 PoS links protected by SONET rings. However, maintaining parallel networks (IP and TDM) was costly. Carriers started exploring ways to carry legacy TDM services over the new packet backbones. Early efforts included technologies like MPLS-based pseudowires and Circuit Emulation (CEM) to transport DS1/DS3 over packet networks without altering customer endpoints. By the mid-2000s, **IP traffic was growing ~50–100% per year**, pushing carriers to expand capacity continually qz.com. AT&T's data shows its IP backbone traffic was growing about **60% year-over-year** around 2008 convergedigest.com. To keep up, AT&T consolidated IP traffic onto a next-gen MPLS backbone running over a **40 Gbps optical transport layer** – by late 2008 AT&T had upgraded "more than 80,000 fiber-optic wavelength miles" of its ultra-long-haul network to OC-768 (40 Gbps) channels convergedigest.com. This represented the world's largest 40G deployment at that time, underpinning all AT&T Internet and IP services. AT&T's achievement – carrying ~16 petabytes of IP data per business day in 2008 – marked the effective transition to an IP-centric core convergedigest.com convergedigest.com.

Other carriers followed suit. Verizon, after acquiring MCI/WorldCom's UUNET backbone, also migrated to an IP/MPLS core over high-speed optics. Verizon was an early mover in **packet-optical convergence**, trialing 40G and planning 100G wavelengths by 2008–2009 fierce-network.com fierce-network.com. Verizon began **retiring SONET layers** in metro networks by introducing MPLS with circuit emulation. A 2014 Verizon network RFP explicitly aimed to replace TDM infrastructure with packet switches that emulate SONET circuits over MPLS cisco.com cisco.com. Verizon's adoption of CEM pseudowires at OC-192 rates (10 Gbps) by 2018 was key to finally eliminating remaining SONET gear, as noted by its optical network director cisco.com. In sum, the late 2000s saw carriers operating *IP over SONET* – effectively using legacy optical networks as a foundation while an IP/MPLS layer handled exploding data traffic and began to absorb legacy services via emulation.

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

Rise of DWDM and IP-over-Wavelength (DWDM Era)

As traffic demand accelerated, carriers needed more raw capacity and began deploying **Dense Wavelength Division Multiplexing** in long-haul and metro core routes. DWDM allowed multiple optical channels (lambdas) on a single fiber, each at OC-192/768 or 10–40 Gbps, massively multiplying fiber capacity. Initially, SONET circuitry was still used on each wavelength, but soon carriers moved to **"transparent" wavelengths** carrying packet traffic directly. By the mid-2000s, **IP over DWDM** architectures emerged – meaning routers could interface directly with DWDM systems, dropping the intermediate SONET layer for data circuits. This reduced overhead and latency, and allowed more flexible scaling. For example, Level 3 and other backbone providers pioneered IP-over-DWDM in core networks in the 2000s, often using their routers with optical transponders to light wavelengths across DWDM line systems.

A critical inflection came with the introduction of **Reconfigurable Optical Add-Drop Multiplexers (ROADMs)** in the late 2000s. ROADMs added flexibility by enabling remote, software-based provisioning of wavelength paths through optical mesh networks, rather than fixed point-to-point or ring circuits. Verizon was a leader in this regard – in 2008 it announced deployment of ROADMs in an **18-market optical mesh architecture** across its U.S. network fierce-network.com fierce-network.com. This allowed Verizon to more dynamically route high-capacity traffic (for instance, between video hub offices) at the optical layer. Verizon concurrently pushed to **100 Gbps wavelength technology**, reasoning that jumping from 10G directly to 100G (skipping 40G in widespread deployment) made economic sense as video demand exploded fierce-network.com. Indeed, Verizon had been the first U.S. carrier to test 100G in a live network and planned to roll out 100G on major routes by 2009 fierce-network.com. The driver for these upgrades was clear: "the explosive growth of online video and other content" gave carriers confidence to invest in higher-capacity optical transport fierce-network.com. By around 2010, video traffic was becoming the dominant load (discussed further in the next section), and carriers that hadn't already were rapidly overlaying DWDM systems and scaling up backbone links to 40G and 100G wavelengths.

From roughly 2010 onward, U.S. carriers operated **IP-over-DWDM core networks**. In practice this meant IP/MPLS routers (often in a full mesh topology) interconnected by high-speed optical channels provisioned on DWDM fiber systems. AT&T, for example, completed a nationwide deployment of 40G wavelengths by 2008 and was testing *114 Gbps per wavelength* technology in collaboration with vendors (hitting 17 Tbps on a fiber in lab trials) convergedigest.com. By the early 2010s, 100G coherent optics became standard in long-haul, and by mid-2010s 100G moved into metro/regional networks as well. Lumen's network (combining CenturyLink, Level 3, and Qwest assets) by this time consisted of an extensive fiber footprint with DWDM systems supporting both internal IP services and wholesale wavelength services. Level 3's legacy of IP focus meant Lumen was quick to adopt IP-over-DWDM; for instance Level 3 had long connected its PoPs with 10 GbE or 10 Gbps waves since the early 2000s cs.purdue.edu. In 2017–2020, Lumen (CenturyLink) and others began deploying **200G/400G wavelengths** using advanced modulation and flex-grid ROADMs, further boosting capacity in anticipation of cloud and streaming growth.

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

ROADM-Based Optical Core and IP-Optical Convergence (2010s–2020s)

The full realization of the modern carrier transport network came with ubiquitous deployment of **ROADM-based optical networks** in the 2010s. By using ROADMs, carriers like Verizon and AT&T transformed their static DWDM backbones into agile photonic networks where wavelengths can be remotely switched, rerouted, or upgraded without manual fiber patching. Verizon's vision by 2007–2008 was an **"End-to-End All Optical Network"** with a wavelength-centric photonic layer supporting 10G, 40G, and future 100G channels, over which packet and TDM services would ride cse.buffalo.edu. Key features included tunable optics, dynamically switched wavelengths, and a control plane for automated provisioning cse.buffalo.edu. In essence, the optical layer began to emulate some of the flexibility of SONET (restoration, reconfiguration) but for high-speed IP pipes.

At the same time, the **IP layer and optical layer grew more intertwined**. Carriers increasingly adopted integrated IP-optical strategies to improve efficiency. One approach was **IPoDWDM**, where IP routers host DWDM transponders or pluggable coherent optics, allowing direct coupling of router ports into the optical line system cignal.ai. Cisco and Juniper pioneered IPoDWDM capabilities in their routers in the early 2010s, though wide adoption came later with standardized pluggables (e.g. 400ZR). The benefit is fewer intermediary devices, and the IP layer can more directly control optical circuits. AT&T and Verizon have both trended toward simplified "routing + optics" architectures in recent years – sometimes called **routed optical networking** – which minimize standalone optical transport gear in favor of optical interfaces on routers for certain network segments blogs.cisco.com.

By the late 2010s, carriers also prepared for the sunset of legacy networks. AT&T, for example, announced plans to **retire TDM and copper networks** by the 2020s, transitioning all services to IP over fiber/wireless. Regulatory filings show AT&T seeking permission as early as 2017 to discontinue many SONET-based offerings in favor of IP services fierce-network.com. Verizon similarly had largely migrated enterprise services off SONET by the end of the 2010s, using MPLS and Ethernet over its optical network. Lumen (with its inheritance of Level 3's IP backbone) moved to shut down older legacy platforms as well, focusing on an all-IP, SDN-controlled network.

In summary, today's carrier transport networks are *multi-Terabit*, IP/MPLS dominated, and riding on reconfigurable optical infrastructures. The **layers** have effectively collapsed: whereas 20 years ago a telco backbone might have four separate layers (TDM voice, Frame Relay/ATM, IP, optical), now voice is VoIP, legacy services are tunneled, and **IP is the unified layer riding directly over an optical layer**. ROADMs and optical mesh ensure the photonic layer can flexibly provide bandwidth where needed, while the IP layer (with protocols like segment routing and fast re-route) handles traffic engineering and resilience above. This evolution was driven at each step by massive traffic growth and changing application demands, which we explore next.

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

Three Phases of Traffic Growth and Network Impact

Traffic demand on carrier networks grew in **three distinct waves**, each posing new challenges and driving infrastructure evolution:

Phase 1: Early 2000s – Web and E-Commerce Surge

The first wave accompanied the **rise of web-based businesses and online content** in the late 1990s and early 2000s. As millions of users came online and dot-com companies (portals, e-commerce, online media) flourished, Internet traffic saw exponential growth. From 1997 to 2003, global Internet traffic grew at an **average of ~127% per year**, doubling roughly every 12 months reddit.com. U.S. networks experienced this surge as the general population adopted email, web browsing, and online shopping. By 2000 there were ~300 million Internet users worldwide (from just 16 million in 1995) pingdom.com, and U.S. broadband adoption began replacing dial-up, allowing richer content usage. E-commerce traffic from sites like Amazon and eBay, as well as enterprise data (VPNs, data center connectivity), started to fill carrier pipes. In this period, **carriers upgraded backbone links from ~2.5 Gbps (OC-48) to 10 Gbps (OC-192)** to keep up, and began deploying DWDM to add more 10 Gbps channels quickly. The web traffic surge highlighted the inadequacy of circuit-switched networks for bursty data – this directly led to the widespread deployment of IP/MPLS cores described above.

Notably, **peer-to-peer (P2P) file sharing** also emerged in the early 2000s (Napster in 1999, BitTorrent by 2003), generating significant bandwidth demand outside of the traditional client-server web model. Carriers saw their **backbone utilization skyrocket**, requiring more capacity and prompting innovations like MPLS traffic engineering to better use available links. For example, AT&T in the early 2000s consolidated multiple IP backbones (from its WorldNet, Genuity, etc. acquisitions) into a single MPLS backbone for efficiencyhit.bme.hu. The result of Phase 1 was that by ~2005, most major U.S. carriers had transitioned to a robust IP-over-optical foundation – a necessary baseline for the even larger storm of traffic to come.

Phase 2: Mid-2000s–2010s – Video Streaming Explosion

The second wave – and arguably one of the largest drivers of network growth – was the **rise of online video traffic**. Around 2005, Internet video moved from niche to mainstream. Cisco's Visual Networking Index and others chronicle this era: *"Video traffic exploded in the mid-2000s. It made up 12% of all internet traffic in 2006, but by 2010 it became the largest category of traffic"* qz.com. Several factors converged: YouTube's 2005 launch demonstrated user-generated content's popularity, Netflix introduced video streaming by 2007, and traditional TV began shifting to IPTV and streaming platforms. Early telco IPTV initiatives (like AT&T's U-verse TV, which used Microsoft Mediaroom, and Verizon's FiOS TV) deployed around 2006–2008, delivering multi-channel digital TV over IP infrastructure to consumers. This added *multicast and unicast video* loads onto networks that previously mostly handled web and data. Soon after, over-the-top (OTT) streaming of on-demand content and live events became common – by the

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

early 2010s, Netflix, YouTube, Hulu, and others were consuming a huge share of bandwidth. Facebook and other social platforms also introduced video features (user uploads, live streaming), further contributing to growth.

This video wave had **profound network impacts**. Video traffic is high-bandwidth and continuous, requiring carriers to massively scale capacity and improve delivery efficiency. Between 2010 and 2015, global IP traffic roughly **quadrupled**, largely due to video, with Cisco reporting that by 2016 over **800,000 minutes of video were crossing the Internet per second** qz.com. Content Delivery Networks (CDNs) became essential: companies like Akamai, and the tech giants' own CDNs (Google, Netflix Open Connect, etc.), colocated caching servers inside carrier networks to serve video locally and reduce backbone load qz.com. This led to a **flattening of the internet topology**, as much of the traffic shifted from traversing tier-1 backbones to flowing directly from CDNs to access networks qz.comqz.com. Nonetheless, carriers had to transport enormous volumes of video from ingress points and caching nodes across their core and metro networks.

Carriers responded by upgrading the optical layer to 100 Gbps wavelengths and beyond, as noted earlier. For example, Verizon's decision to move quickly to **100G by 2009** was driven by "the explosive growth of online video and other content" which demanded higher capacities fierce-network.com. AT&T and others similarly deployed 100G tech in their backbones in the early-to-mid 2010s. Core router capacities jumped from 10s of Gbps in the early 2000s to 100s of Gbps or multi-terabit by 2010, with carriers adopting 100GE interfaces and dense port aggregation. The **median traffic per user** also rose with video – instead of a few GB per month, users were soon consuming tens or hundreds of GB, especially with HD and later 4K streaming. By the mid-2010s, multiple sources estimated video comprised **70–80% of consumer internet traffic**.

From a network architecture standpoint, Phase 2 also drove **peering and interconnection changes**. Large content providers (Netflix, Google/YouTube, etc.) negotiated direct connections to carrier networks or installed on-net caches, altering traffic flows. Carriers had to manage **hotspots of traffic** (e.g., evening video streaming peaks) and ensure quality (low buffering, low latency). Many invested in expanding **metro fiber and aggregation** since video content was often sourced from within metro-area caches or video offices. For instance, Verizon built a dedicated ROADM mesh linking video serving offices to head-ends to efficiently carry FiOS TV streams in metro regions fierce-network.com. In summary, the video era forced carriers to *scale up and smarten up* their networks: capacity upgrades were paramount, and technologies like MPLS-TE, CDN integration, and robust optical mesh were deployed to handle unpredictable, high-volume video flows.

Phase 3: 2020s – AI-Driven Traffic Wave (Emerging)

Now in the mid-2020s, the industry anticipates a third major traffic wave: **AI-driven traffic**. This is a nascent but rapidly growing phenomenon tied to the proliferation of artificial intelligence applications. Two aspects define this wave:

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

- **AI/ML Training Data Movement:** The process of training AI models (especially large neural networks) involves moving *huge datasets* between storage and compute clusters, and often distributing training across data centers. Research and hyperscale cloud providers regularly shuffle petabytes of data for AI training, using high-speed data center interconnect (DCI) links.
- **AI Inference and AI-enhanced Services:** As AI is embedded in applications (from vision analytics to generative AI like ChatGPT), user requests increasingly involve sending images, video, or prompts to AI models and receiving results. Many such services are cloud-based, so inference generates client-to-cloud and cloud-to-cloud traffic. Applications like automated video surveillance, AR/VR experiences with AI processing, and real-time language or image processing all drive high-bandwidth streams.

Concrete projections are eye-opening. According to Omdia's latest data, **all AI-related traffic accounted for 39 exabytes of network traffic in 2024**, and is expected to more than double to **79 exabytes in 2025** networkworld.com. This growth far outpaces traditional traffic. Omdia predicts that by 2031, **AI application traffic will overtake conventional network traffic** in volume networkworld.com. "Net new" AI traffic includes visual apps, sensor and IoT feeds for AI analysis, AI-generated content, etc., while "AI-enhanced" traffic includes things like smarter video streams, content personalization, and augmented cloud app data networkworld.com. In enterprises, the impact is already visible – one report noted a **3,464% increase in AI-related enterprise traffic in late 2024**, driven by usage of tools like ChatGPT, and measured 3.6 PB transferred to/from AI apps in 11 months networkworld.com.

For carriers, the AI wave has several implications. First, **Data Center Interconnect (DCI)** becomes even more crucial – hyperscalers (and some telcos) are investing in 400G, 800G, and multi-terabit inter-DC optical links to support AI clusters that might span sites ciena.com. Ciena notes that the AI boom is "radically transforming network demand" and is comparable to the introduction of the internet or smartphone in its network impact ciena.com. A key trend is the need for **high-bandwidth, low-latency connectivity** between compute nodes, which may drive new architectures (such as hierarchical mesh networks linking data centers, edge sites, and devices with AI capabilities). Second, AI inference traffic – often dominated by large video or image data flows that need processing – will stress the *edge and access networks* as well as the core. For example, autonomous vehicles or AR glasses producing continuous video feeds for cloud AI to analyze would consume enormous bandwidth if widely adopted.

It's worth noting that as of 2025, **AI traffic is still in early days and not yet a dominant factor outside data centers** ciena.com. However, carriers are preparing by ensuring their networks are "AI-ready" – increasing capacity, embedding edge compute nodes to localize AI processing, and leveraging automation for dynamic scaling. The AI wave also renews focus on **network agility**: connectivity may need to spin up quickly between GPU clusters, and on-demand bandwidth services (flexible optical circuits, network slicing in 5G, etc.) could see higher use.

In summary, the AI era is expected to continue the pattern of exponential traffic growth but with new characteristics (more east-west data center traffic, surges tied to AI-driven events, etc.).

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

Carriers that evolved through the web and video eras are now strategizing on how to extend their IP/optical networks – and their capacity planning practices – to handle a potential flood of AI data. This leads us to examine how carriers plan and optimize their networks, and how tools for network planning have progressed (and sometimes lagged) alongside these traffic waves.

Network Planning Tools: Evolution and Use Cases

Designing and optimizing a carrier network through these evolutions is a complex task. Over time, the tools and software used for **network planning** have evolved from rudimentary, vendor-specific calculators to sophisticated multi-layer modeling suites. Each generation of tools aligned with the technology of the day – from TDM circuit planners to IP traffic simulators to modern SDN-aware design tools. Below we detail this progression, highlighting representative tools and their capabilities.

TDM Era Planning – Vendor-Specific SONET/SDH Tools

In the era of SONET/SDH networks (1990s–early 2000s), planning was largely done with **vendor-proprietary software** provided by equipment manufacturers. Each optical/TDM vendor (Nortel, Lucent, Fujitsu, Alcatel, etc.) offered tools to design rings, add/drop multiplexers, and digital cross-connect layouts using their equipment. For example, Nortel's OPNET group developed a product informally known as **Nortel SONET Planner** in the late 90s, which was a GUI tool to design SONET networks (e.g., ring configurations, VT/STS circuits) and plan transitions between architectures collectionscanada.gc.ca. These tools allowed engineers to input traffic demands (number of T1/E1s, DS3s, etc. between sites) and then helped allocate them onto SONET rings or meshes, determine the required Add-Drop Multiplexers (ADMs), and ensure redundancy (e.g., two-fiber vs. four-fiber ring, Automatic Protection Switching paths). They were typically single-vendor – i.e., a **Nortel tool for Nortel OC-192 rings**, a Lucent tool for Lucent gear, and so on. They focused on **TDM circuit efficiency** (grooming lower-rate circuits into higher-rate paths) and ensuring the network met protection and delay requirements. Quality factors in this context were straightforward, like not exceeding the number of hubs in a SONET chain or ensuring diverse fiber routes for protection rings.

While adequate for static voice networks, these early tools had **limited scope**. They did not account for IP packet routing or statistical traffic patterns – traffic was assumed deterministic and circuit-switched. Each vendor's planner optimized the use of that vendor's hardware (often recommending additional ADMs or DCS gear, unsurprisingly). Optimization meant minimizing stranded bandwidth on SONET rings or ensuring fully utilized STS channels. As packet data demands grew, the inherent inflexibility of these tools (and SONET networks) became apparent. Network planners increasingly needed tools that understood **IP topologies, routing protocols, and bursty traffic**.

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

Emergence of IP/MPLS Planning Tools (2000s)

As carriers deployed IP and MPLS networks, a new class of planning tools arose to tackle the complexity of **routing, traffic engineering (TE), and failure analysis** in packet networks. Unlike TDM, where circuits are fixed, IP networks have dynamic routing that can shift traffic when failures or congestion occur. Planning an IP backbone requires modeling how traffic flows between every pair of nodes (the traffic matrix) and how routing protocols (OSPF/ISIS, BGP, MPLS-TE) will carry that traffic over the physical links. By the mid-2000s, industry-leading tools in this space included **Cariden MATE**, **WANDL IP/MPLSView**, and OPNET's **SP Guru** (later acquired by Riverbed). These tools were typically multi-vendor and software-centric, often founded by tech startups and later absorbed by larger vendors (Cisco acquired Cariden in 2012; Juniper acquired WANDL in 2013; Riverbed acquired OPNET in 2012).

Cariden MATE (Multi-layer Traffic Engineering) – later known as Cisco WAE (WAN Automation Engine) – was widely used by carriers and large ISPs. MATE could ingest network topology (either via router configs or via routing protocol databases) and build a model of the network's nodes, links, and IGP metrics blog.packetsource.net blog.packetsource.net. Planners would input a **traffic matrix** (volume of traffic from each ingress to each egress) – if not known, MATE had tools to estimate it from link loads blog.packetsource.net. The tool could then simulate routing, showing which paths traffic takes and where congestion might occur. Crucially, one could perform **failure simulations** (single or multiple link/node failures) to see worst-case loads when traffic reroutes blog.packetsource.net blog.packetsource.net. This allowed planners to identify weak points and engineer the network to avoid overload. Cariden MATE and similar tools also supported **MPLS traffic engineering**, where explicit Label-Switched Paths (LSPs) could be designed to balance load. They could recommend optimized IGP metrics or LSP placements to better spread traffic – in effect, performing **offline optimization of IP networks**. For example, Cariden could suggest moving traffic off an overused link by adjusting OSPF weights, and quantify the improvement.

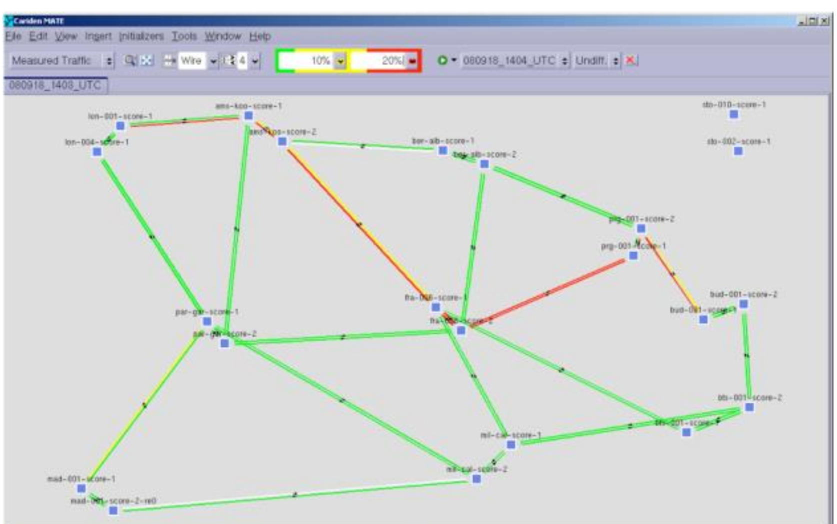

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

Figure: Example output from Cariden MATE (Cisco WAE) – an IP/MPLS network topology with color-coded link utilization. Such tools model multi-vendor networks and simulate traffic flows under various scenarios to aid capacity planning and traffic engineering.

Likewise, **WANDL IP/MPLSView** (now Juniper's product) offered comprehensive multi-vendor support and scalability up to thousands of routers. WANDL distinguished itself with detailed modeling of routing behaviors and an extensive device libraryjuniper.net juniper.net. According to Juniper, IP/MPLSView could perform *"exhaustive single and concurrent failure"* analyses, traffic matrix estimation, MPLS diverse path design (for redundancy), fast reroute planning, and even suggest **network optimizations** juniper.net. It effectively combined **traffic engineering and traditional FCAPS management**: it had modules for fault and performance management, and could import live network data (via SNMP, CLI, etc.) to keep models accurate juniper.net juniper.net. For planners, this meant a single platform could both monitor the network and simulate changes. WANDL IP/MPLSView emphasized ensuring every bit of capacity was used profitably – *"every bit in the network leads to a dollar of revenue"* was a guiding principle juniper.net. By optimizing existing assets before adding new ones, these tools helped delay capital expense. They also provided **reports** like needed capacity augments, traffic growth trending, and hardware inventory for upgrades.

OPNET's **SP Guru** was another notable tool, particularly for MPLS networks. Renamed Riverbed Modeler / SP Guru after acquisition, it similarly allowed modeling of large IP networks, tuning of OSPF/BGP, and capacity planning. These IP-era tools often could handle **multi-layer modeling to an extent** – for example, Cariden could represent both IP links and underlying optical circuits (explicitly or abstractly) if provided, to ensure resilience plans account for shared risk groups, etc. However, their primary focus was the IP/MPLS layer. By automatically highlighting overloaded links, suggesting alternate routing, and evaluating "what-if" scenarios (like *what if traffic grows 50% next year, or what if a fiber cut happens on this route*), these tools became indispensable for network engineering teams. They enabled a shift from reactive over-provisioning to **proactive planning** – finding weaknesses *before* they caused outages or customer impact eetimes.comeetimes.com.

Crucially, these platforms were **multi-vendor and multi-protocol**. A carrier's network often included routers from Cisco, Juniper, etc., and these planning tools integrated via open data (config files, SNMP, routing table dumps). For the first time, carriers had vendor-agnostic software to optimize across the whole network. Still, they were often used in a consultative manner – run by a dedicated planning group or even by vendors as a service to the carrier – rather than tightly integrated into daily operations. They provided recommendations (add capacity here, tweak this metric, etc.) that humans would then implement.

Optical Planning Tools – Ensuring Light Path Viability

In parallel to IP tools, the optical layer evolution gave rise to specialized **optical network planning tools**. These tools focus on the **Layer 0/1** domain: designing DWDM wavelength routes, amplifier sites, and ensuring signal quality (QoT) on fiber links. Examples include vendor-specific tools like

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

Ciena OnePlanner, Nokia's **1830 Engineering Planning Tool (EPT)** and **1390 Network Planning Tool (NPT)**, as well as vendor-neutral simulation software like **VPI Photonics** and RSoft's OptSim for physical layer modeling.

Optical planning is a different art – it must account for fiber characteristics (loss, dispersion), amplifier noise, nonlinear effects, and so on to guarantee that a 100G or 400G signal will reach its destination with error rates within spec. Early on, optical vendors provided static tools: essentially spreadsheets or simple programs to budget power and OSNR (optical signal-to-noise ratio) across a link. As ROADMs and mesh networks emerged, this advanced to full topology design. For example, **Ciena's OnePlanner** is described as an *"advanced multi-layer network design and optimization tool"* combining Ciena's expertise in Layer 1 planning and photonic simulation scribd.com. It correlates data from different layers, allowing planners to see how services map onto wavelengths and equipment scribd.com. OnePlanner can simulate **what-if failure scenarios in the optical domain** – e.g., multiple fiber cuts – to ensure the network can be restored scribd.com. It supports designing both the photonic layer (e.g., setting up ROADM nodes, fiber routes, and optimizing amplifier placements) and the higher-layer grooming (SONET/OTN or packet paths) on that optical layer scribd.com. Ciena's tool, of course, is aligned to Ciena's hardware (6500 Packet-Optical platform, CPL photonic layer, etc.) scribd.com, providing accurate modeling and Bill-of-Material outputs for that equipment. Similarly, **Cisco Optical Network Planner (ONP)** was introduced as an automation-era tool for DWDM design. Cisco ONP is a *"multilayer network design and planning tool"* that can take input data (site locations, traffic demands) and produce an automatic DWDM network design cisco.com. It performs **fiber transmission validation** (ensuring the proposed links meet optical requirements) and outputs bills-of-material so customers can purchase and install the correct hardware cisco.com cisco.com. Such tools often include **network import and what-if planning** features, enabling planners to model upgrades or new wavelength deployments on an existing network cisco.com. Nokia's EPT/NPT serve similar roles for Nokia's optical gear, adding automation and intelligence to design optical transport networks efficiently customer.nokia.com scribd.com.

There are also independent photonic simulation packages like **VPItransmissionMaker by VPIphotonics** and RSoft's tools. These are used for deep physical-layer R&D and planning – for example, to evaluate a new modulation scheme or to design a custom link. They can simulate propagation of optical signals through fibers and components, accounting for impairment accumulation. While primarily used by equipment designers or researchers, carriers have used them in planning ultra-long-haul or exotic links (e.g., undersea cables) where detailed physics matter to decide regenerator spacing, dispersion maps, etc. lightwaveonline.com lightwaveonline.com.

Overall, optical planning tools are geared towards **Quality of Transmission (QoT)** assurance and cost optimization in the optical domain. They answer questions like: how many wavelengths can we run on this span? Will a 100G signal reach 1000 km with these amplifiers? Where do we need signal regeneration? They output network designs that meet reliability and performance targets, often including **protection routes** (optical mesh restoration paths or equipment redundancy).

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

Modern optical planners also account for **multi-period planning** – e.g., designing a network that can grow in phases, adding wavelengths as traffic increases vpiphotonics.com.

One important trend is the attempt to bridge IP and optical planning. Some tools (like Ciena OnePlanner, Nokia NPT, or even academic open-source tools like Net2Plan) advertise **multi-layer optimization** – for example, grooming IP traffic onto wavelengths optimally, or evaluating IP rerouting in tandem with optical switching. In practice, however, this integration has been limited (as we discuss in the next section on limitations).

Modern Vendor-Aligned Planning Platforms

As networks became more software-driven, vendors have consolidated planning functions into broader tool suites, often tied to their SDN controllers. Today, we see **Cisco WAE**, **Juniper IP/MPLS View**, **Ciena OnePlanner/Planner+,** and **Nokia NSP/EPT** being positioned as parts of an end-to-end management ecosystem. These platforms retain the core capabilities discussed (simulation, capacity planning) but are evolving to be more **continuous and automated** rather than one-off offline exercises.

For example, **Cisco WAE (WAN Automation Engine)**, born from Cariden MATE, now integrates with live networks via streaming telemetry and interfaces with Cisco's SDN controllers. Planners can use WAE Design offline to test changes, but the same engine can feed into real-time traffic engineering adjustments via an SDN approach github.com. **Juniper's IP/MPLS View** similarly is described as both a planning and a management tool – it has modules for performance monitoring, inventory, etc., meaning it can be part of day-to-day network operations juniper.net juniper.net. This blurs the line between "planning" and "management," ideally enabling a closed-loop where the network's state is always known and optimization is continuous.

On the optical side, **Cisco ONP** is part of an "Optical Automation Suite" along with controllers (Cisco ONC) cisco.com, indicating its role in a vendor-aligned automation framework. **Ciena MCP (Manage, Control, Plan)** now includes Planner functionalities (PlannerPlus) for designing services that can then be provisioned through the same interface ciena.com ciena--integrate.sandbox.my.site.com. Nokia's Network Services Platform (NSP) likewise incorporates planning for IP/optical networks as part of an integrated software suite nokia.com.

Despite these advancements, the fundamental usage of these tools by carriers has often remained **scenario-based**: perform capacity planning for the next year, simulate a major failure to verify resiliency, or evaluate a network design for a new customer/service. As the next section explores, these tools – while feature-rich – often serve to justify or fine-tune network upgrades (frequently aligned with the vendor's hardware proposals) rather than to globally optimize a multi-vendor network for minimal cost.

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

Limitations of Planning Tools and Industry Challenges

While network planning tools have grown more powerful, they have notable limitations that have constrained their impact on *holistic network optimization*.

One core issue is that **planning tools are typically used for "what-if" analysis and validation, rather than automated optimization of live networks**. In practice, operators use these tools to run scenarios: *If traffic grows 50%, where will congestion occur? If this fiber fails, can the network handle it?* They then use the insights to add capacity or adjust configurations manually. The tools excel at ensuring capacity plans are sufficient (so-called **capacity management**) and that failure survivability is met – for instance, an optical planner ensures a new 200 km wavelength meets QoT, or an IP planner confirms that after adding a new link, utilization stays under 70% on all paths under N-1 failure. However, they rarely perform **true multi-layer optimization** where, for example, IP routing, optical routing, and capacity upgrades are jointly optimized to minimize total cost or maximize efficiency across layers. Each tool tends to operate in its silo or vendor context.

A related limitation is **vendor specificity and bias**. Many planning tools (especially modern ones) are tied to a particular vendor's product line – they are often provided "free" or at low cost to carriers as part of large equipment deals. This inevitably means the tools are geared toward making that vendor's gear look essential. For example, a vendor tool might readily identify a capacity shortfall and suggest adding a certain card or chassis (from that vendor) to fix it. Rarely would it suggest that a competitor's solution or a novel network design could obviate the need for more equipment. Even ostensibly multi-vendor tools, once acquired by a vendor, may prioritize features that align with that vendor's architecture. This has led to an industry cynicism that **planning tools support equipment sales more than cost-saving optimization**. Instead of enabling a lean multi-vendor network, they often assume a single-vendor network for modeling ease, nudging operators toward one ecosystem.

Moreover, **multi-layer coordination remains rudimentary**. True optimization would involve, say, re-routing IP flows in real-time when an optical impairment degrades a circuit, or choosing an IP link topology that minimizes regen equipment in the optical layer. In practice, IP and optical planning are still largely separate processes, often handled by different teams. Tools like OnePlanner that correlate layers 0-1, or attempts to incorporate packet layer in optical tools, have mostly been used to verify designs rather than to dynamically optimize across layers scribd.com scribd.com. The complexity of converged planning is high – each layer has different constraints and timescales – and commercial tools so far haven't truly cracked multi-layer optimization beyond incremental steps.

Another limitation is that many tools operate **offline with stale data**. They take snapshots of the network (e.g., nightly dumps of routing tables and traffic stats) and perform analyses, but they are not continuously updated with real-time network conditions. This means rapid changes (say a sudden traffic shift due to a viral event or an unexpected fiber cut) are not accounted for until

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

after the fact. Operators then rely on over-provisioning and protocol fail-safes in real time, and use planning tools only post-mortem to analyze and adjust for next time. While modern integrated tools are improving real-time visibility, most carriers have been slow to trust any automated optimization that isn't human-reviewed – especially in multi-vendor scenarios where an error could cause widespread outages.

Importantly, **optimization goals of tools are often narrow**. For example, IP/MPLS planning tools might aim to minimize max link utilization or ensure no packet loss, but they don't directly factor $$ costs or future-proofing unless manually modeled. Optical tools ensure signal quality but may not account for, say, the cost benefit of routing a wavelength on a slightly longer path to avoid a new regen site. Planning teams often use these tools to generate options and then apply business judgment. There's also a lack of **truly unified multi-vendor optimization** software in the market – no off-the-shelf tool can take a heterogeneous network with Cisco, Juniper, Nokia IP gear and Ciena, Infinera optical gear and spit out an optimal plan that spans all. Each vendor pushes their own platform, and independent tools (like academic open-source projects) haven't gained industry traction due to support and trust issues ieeexplore.ieee.orgciena.com.

It's telling that when Juniper bought WANDL after Cisco bought Cariden, analysis pointed out the moves were more about integrating these tools into each vendor's SDN strategy (Contrail and Cisco's SDN) than about solving multi-vendor optimization blog.cimicorp.com blog.cimicorp.com. In essence, vendors saw planning tools as valuable for **analytics and path computation** to enhance their own controller software, not necessarily to empower carriers to mix-and-match vendor equipment freely.

Finally, we must note that **human factors and organizational silos** limit tool efficacy. A tool might be capable of a multi-layer design, but if a carrier's IP engineering team and optical engineering team use different tools and processes, the coordination may not happen. Planners have often been using tools for *verification* (i.e. "does this upgrade plan meet requirements?") rather than generative optimization ("find me the best possible network design"). The result is sub-optimal networks persisting. As an EE Times commentary noted, *"planning tools should simplify multi-layer, cross-vendor implementation… But that is not the case. Legacy tools are retrofitted into next-gen networks, and operators struggle to find the right support tools for convergence."*eetimes.comeetimes.com. Many next-gen operators migrating to all-IP found that their old tools didn't truly meet new needs, and new tools weren't fully mature, leading to a gap where planning was partly guesswork and vendor consulting.

In summary, today's planning tools provide valuable assistance in capacity forecasting, failure impact analysis, and design validation – they **increase efficiency and reduce OPEX** by saving time and preventing mistakes eetimes.com. However, they often stop short of delivering *optimal* networks. They are constrained by vendor alignment, lack of multi-layer insight, and use more as advisory systems than real-time decision-makers. The competitive nature of vendor offerings has also meant no neutral industry-standard optimization software has emerged. As traffic continues to surge (with AI, etc.), these limitations become more apparent: carriers may need to embrace more open, perhaps **AI-assisted planning tools** or open-source platforms to truly coordinate

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

across layers and vendors ieeexplore.ieee.orgnet2plan.com. There is ongoing research and calls for open multilayer planning solutions that could unify the process and prevent vendor lock-in on network design researchgate.neteetimes.com.

Conclusion and Outlook

The journey from 2000 to 2025 has seen U.S. telecom carriers completely reinvent their transport networks. They have moved from rigid TDM SONET infrastructures carrying primarily voice, to agile IP-over-ROADM networks delivering petabytes of data for video, cloud, and soon AI applications. Each phase of traffic growth – web, video, and now AI – forced key upgrades: **the web era prompted IP/MPLS adoption**, **the video era drove DWDM and 100G scaling**, and **the emerging AI era is pushing toward 400G+ and edge/cloud re-architecture**. Carriers like AT&T, Verizon, and Lumen navigated these changes with massive investments and by leveraging new technologies (e.g., AT&T's early 40G backbone convergedigest.com, Verizon's fast adoption of ROADMs and 100G fierce-network.com fierce-network.com, Level 3's IP-centric buildout).

Network planning tools have been an essential part of this story, enabling operators to plan and deploy upgrades intelligently. We saw the progression from **vendor-specific SONET planning tools** to **multi-vendor IP/MPLS planning suites** as networks shifted to IP, and then the parallel development of **optical planning tools** for complex photonic networks. In recent years, planning functions have started converging with network automation, hinting at a future where planning is not a separate offline task but an ongoing network self-optimization practice.

However, the current state reveals that most carriers have not yet achieved *fully optimized, multi-layer, multi-vendor networks*. Planning tools are powerful, but their **limitations** – siloed operation, vendor bias, and usage primarily as offline simulators – mean that networks are often run conservatively with plenty of over-provisioning and manual engineering. There is opportunity in the next decade to harness advances (perhaps even AI-driven network optimization) to fulfill the unmet promise of these tools: true end-to-end optimization that cuts across IP and optical layers and across vendor boundaries.

For a strategic business audience, the key takeaway is that **transport network evolution is driven by traffic demand**, and success depends on both adopting the right technologies and using planning intelligently. Carriers that recognized early the rise of packet traffic invested in MPLS cores and retired TDM, those that anticipated video built out fiber capacity and CDN strategies, and now those preparing for AI are extending fiber deep and automating networks. Planning and optimization tools, while imperfect, have provided the data and confidence to undertake these network transformations with manageable risk. They have helped avoid failures by predicting them and have justified capital expansions by quantifying needs.

As we stand in 2025, major U.S. carriers have robust, modern transport networks that are a far cry from the SONET rings of 2000. They carry orders of magnitude more traffic with a fraction of the cost per bit. Looking ahead, the next challenges will be integrating **compute and network (edge computing for AI)** and **leveraging software intelligence** to manage complexity. The hope is

© 2025 NetOpt.Design – All rights reserved | www.netopt.design

that planning tools – perhaps evolving into AI-empowered network “brains” – will finally deliver on holistic optimization, allowing carriers to squeeze maximum value from every fiber and wavelength. If the past is prologue, the continuous collaboration of network engineers, tool developers, and equipment vendors will remain critical to navigate whatever new demands the future holds. The evolution is ongoing, but the lessons of 2000–2025 make one thing clear: adaptable networks and informed planning go hand in hand in meeting the world’s insatiable demand for connectivity.

© 2025 NetOpt.Design. All rights reserved. **NetOpt.Design™**, the NetOpt.Design logo, and all related marks are trademarks or registered trademarks of NetOpt.Design. Unauthorized reproduction or distribution is prohibited. Other trademarks mentioned in this paper belong to those entities, information about the other entities based on publicly available sources as mentioned. **Contact:** info@NetOpt.design