NetOpt: Bridging the Multi-Layer Network Planning Divide

Telecom network planning traditionally relies on separate toolchains for different network layers. IP/MPLS routing is planned with one set of tools, while the optical transport (wavelength/fiber) layer is planned with another. These siloed tools each optimize within their layer – but lack coordination across layers. The result is often suboptimal and over-provisioned networks, where each layer adds its own "safety margin" on top of inflated forecasts. This white paper introduces **NetOpt**, a new multi-layer network planning tool designed to integrate all layers – fiber, wavelength, MPLS, IP, physical infrastructure, and traffic demand – into a joint optimization model. We first survey the landscape of current planning tools in two categories – IP/MPLS planners and optical network planners – and highlight their limitations. We then explain how NetOpt overcomes these challenges, including modeling CDN traffic and multi-period demand patterns, to right-size networks and avoid the costly inefficiencies of today's siloed planning approach.

IP/MPLS Network Planning Tools Today

IP/MPLS planning tools focus on Layer 3 routing and MPLS traffic engineering. They ingest the IP topology and traffic matrix, then simulate routing behaviors under various conditions (e.g. failures, growth) to ensure adequate capacity and performance. These tools typically allow "what-if" analysis of IP link additions or metric changes, helping planners optimize routing or add capacity where needed. Popular examples include:

- **Cisco WAN Automation Engine (WAE)** – A design and planning suite for IP/MPLS networks. It visualizes network topology and traffic flows, simulates routing protocols, and evaluates impacts of failures or traffic growth cisco.com. Cisco WAE can recommend optimal paths for MPLS or Segment Routing traffic cisco.com. (Cisco WAE also has limited Layer 1 integration for IP-over-optical scenarios, discussed later.)
- **Juniper WANDL IP/MPLSView** – An IP/MPLS traffic engineering and management solution supporting multi-vendor networks. It models detailed routing protocol behavior and can simulate single or concurrent failures, design diverse MPLS paths (including Fast Reroute), and perform capacity planning networkscreen.com. IP/MPLSView emphasizes optimizing existing network assets before requiring new infrastructure. It offers extensive multi-vendor support and can handle very large topologies.
- **Riverbed (OPNET) SP Guru Network Planner** – A tool (now discontinued) that enabled "what-if" planning for service provider IP/MPLS networks, including capacity forecasting, failure analysis, and network design validation. It allowed planners to predict effects of new applications or traffic growth on link utilization and optimize MPLS traffic

engineering settings. (Riverbed acquired OPNET's SP Guru; the product has since been end-of-life.)

- **Aria Networks** – An AI-driven network optimization platform. Aria's algorithms compute optimal paths and capacity plans, and have been applied to IP/MPLS routing as well as other domains. Notably, Aria's technology includes a capacity computation engine with AI that can determine optimal traffic paths in a multi-layer network ofcconference.org. In practice, it is often used as an SDN path computation element or for MPLS traffic engineering within the IP layer, rather than a fully integrated IP-optical planning tool deployed by operators.
- **Ciena Route Optimization and Analysis (ROA)** – A planning tool (originating from Cyan's Blue Planet) that performs IP-layer analysis. It can model IGP routing, LSP paths, and simulate failures to ensure the IP/MPLS network is robust and efficient. ROA primarily deals with Layer 3, though it can correlate with optical layer data for visualization of service-to-transport mappings. Cross-layer optimization is not its focus; it treats the optical transport as a fixed underlying infrastructure.

Common capabilities of IP/MPLS tools include simulation of routing protocols (OSPF/ISIS, BGP), MPLS or Segment Routing path computation, and automated capacity planning reports. They excel at identifying IP-layer bottlenecks and suggesting augmentations. However, they treat the underlying transport (optical layer) as static – usually as fixed-capacity links. For example, if an IP link is actually a bundle of optical circuits, the IP tool doesn't model how those circuits are provided or could be reconfigured; it simply sees a link with X Gbps capacity. There is no native awareness of fiber distances, wavelength routing, or optical signal feasibility in these IP-layer tools. If the IP tool decides a new 100 Gbps link is needed between two sites, it cannot itself plan the optical path for that link – that task is handed off to the optical planners. In summary, IP/MPLS planning tools are powerful for capacity planning and failure analysis in the routing domain, and some offer "optimization" like tuning IGP metrics or computing MPLS LSP placements. But their scope is confined to the packet layer. Any impact of optical constraints (fiber delays, regeneration limits, wavelength availability) is outside their model.

Optical Network Planning Tools Today

At the other end, optical planning tools focus on the Layer 1 transport network – wavelengths, fibers, and optical equipment. These tools take a set of sites and required circuit demands (e.g. needed bandwidth between site A and B) as input, and output a detailed optical design: how wavelengths are routed over fiber spans, where regeneration or amplification is needed, and whether the signal quality (OSNR, dispersion, etc.) is acceptable on each route. Key examples include:

- **Cisco Optical Network Planner (CONP)** – An automated design tool for DWDM optical networks, part of Cisco's Optical Automation suite. CONP allows users to input desired capacity between sites and then automatically generates a network design, performing fiber routing and validating optical signal performance cisco.com. It provides standard outputs for fiber span engineering and a bill-of-materials to deploy the optical network. (Cisco also offers an "IP-over-OTN" planning feature in Cisco WAE for multi-layer scenarios, but it relies on external optical data from controllers like Cisco EPNM rather than true co-optimization within a single tool.)

- Ciena OnePlanner Unified Design System – Ciena's design and optimization tool for optical networks. OnePlanner supports photonic layer design for Ciena's DWDM platforms (e.g. 6500, Waveserver) and can also visualize multi-layer associations between client services and underlying optical transports. It performs automated optical route finding and wavelength assignment, checking signal feasibility for each path. OnePlanner is *multi-layer* in the sense of correlating Layer 0 (fiber) and Layer 1 (optical channels) – for example, it can simultaneously design the fiber topology and the wavelength paths on top of it westconcomstor.com. However, its scope is still limited to the optical domain (sometimes including OTN or Ethernet client layers if those are part of the optical system). It requires the higher-layer traffic demands as fixed inputs, rather than calculating IP routing itself.
- Nokia 1830 Engineering and Planning Tool (EPT) and 1390 Network Planning Tool (NPT) – Nokia's optical network design tools for their 1830 PSS WDM/OTN systems. These are used to plan wavelength routes, optimize regenerator placements, and ensure the design meets performance requirements. The Nokia tools allow multi-period planning (i.e. modeling gradual network growth over multiple years) and can optimize utilization of WDM channels, but again, the IP-layer traffic must be provided as an input. The tools do not internally simulate IP routing; they assume the demanded circuits are required at full capacity.
- Huawei OptiPlanner (and similar vendor-specific optical planners) – Used for planning optical transport networks (WDM and OTN). These tools ingest the physical fiber map and requested circuit capacities, and output a detailed equipment plan (transponders, ROADM configurations, fiber routes). They often include impairment-aware routing algorithms to maximize optical reach. From the optical tool's perspective, a "100 G" client signal is just a demand to be carried from A to B – no matter how the IP layer might reroute traffic if a circuit fails, the optical planner doesn't account for that.

Common capabilities of optical tools include route finding across fiber spans (considering distance and loss), wavelength assignment (ensuring no wavelength collisions on a fiber, and often minimizing total wavelengths used), and signal quality simulation (verifying that each optical path meets QoT thresholds given amplifiers and fiber impairments). Some have optimization objectives like minimizing regeneration cost or maximizing spectral efficiency. Crucially, these tools assume the traffic demands are given and static – usually derived from the IP/MPLS layer's requirements. For instance, an optical planner might be told that site A and site B need 200 Gbps of capacity between them; it will then design (say) two 100G waves between A and B. But it doesn't model how IP traffic might fluctuate or reroute; it just ensures 200 Gbps is always available. If extra headroom is needed for protection, the IP planners typically decide that separately and then request (often over-request) those optical circuits.

Because of this separation, planning happens in silos. The IP team runs simulations and determines a set of required circuit upgrades for the next budget cycle. The optical team takes that list and plans the optical build (new fiber routes or wavelengths) to accommodate them. The two sets of tools don't inherently coordinate beyond that offline exchange of demands. Some modern systems offer better visibility – for example, Ciena's Blue Planet Multi-Layer software allows viewing IP flows mapped onto optical paths for troubleshooting, aligning IP and transport data in one UI – but not true joint optimization of capacity planning. Each tool "optimizes" its domain in isolation, potentially missing the big picture.

Siloed Planning: Limitations and Inefficiencies

Operating with separate IP and optical planning leads to several well-recognized issues:

- **Lack of End-to-End View:** Because IP and optical layers are designed separately, planners lack a unified view of how a change in one layer affects the other. As one expert noted, "the two networks are built independently, are expanded independently and maintained by two teams working in complete silos" lightreading.com. An IP planner might add redundant links for reliability, unaware that at the optical layer those links share the same fiber and aren't truly diverse – a miscoordination that could cause an outage if that fiber is cut. Conversely, optical planners might assume worst-case utilization on every circuit because they don't see that the IP layer can reroute traffic dynamically when needed.
- **Over-Provisioning and Margin Stacking:** Each layer adds its own safety margins to handle uncertainty. The IP layer often runs at low utilization – industry practice is to upgrade links when they reach ~40–50% utilization, to leave headroom for failures research.google.com. The optical layer similarly often provisions extra capacity (additional spare wavelengths or even whole fibers) to ensure any future traffic fits. These buffers multiply. Studies have shown average backbone link usage is only on the order of 10% – links run mostly empty except during failures research.google.com. In Faisal Khan's words, "each layer is over-provisioned to cope with uncertain demands and overprotected by using resources on each layer" lightreading.com. All that unused capacity is essentially stranded investment. Ciena notes that leveraging this existing unused margin (instead of building new) could eliminate many inefficiencies of the overbuild model ciena.com.
- **Iterative, Slow Planning Cycles:** When IP and optical are planned in sequence, it often requires several back-and-forth iterations to converge on a feasible solution. For example, the IP team might request a very large new circuit that the optical team finds impractical (due to distance limitations or equipment constraints), so the IP plan must be adjusted to use intermediate sites or additional hops. This back-and-forth slows down the planning cycle and may yield suboptimal routes. The IP topology might end up being constrained by where optical paths exist, rather than what's ideal for IP routing. It would be more efficient to consider such constraints together from the start, rather than in a sequential trial-and-error process.

- **Inability to Model CDN and Modern Traffic Patterns:** Today, a majority of internet traffic is driven by content delivery networks (CDNs) for video and rich media. However, current planning tools do not explicitly model CDNs or caching infrastructure. IP planners treat traffic matrix inputs as static demands between routers, not accounting for the fact that much traffic may originate or terminate at distributed caches within the network. In reality, CDN traffic has distinct patterns – it is highly diurnal (video streaming peaks in evenings), often localized (served from a cache in the same metro), and can shift rapidly with content popularity. Ignoring these nuances can lead to overestimating external traffic demand (planners might assume content traffic comes from distant transit peers or data centers, when a good portion is actually served locally). This further contributes to overbuilding – operators might build out capacity for peak OTT video demand network-wide, not realizing that improved caching or off-peak usage patterns could ease the load. Without modeling CDN deployment and hierarchy, planners miss opportunities to optimize around the real dominant traffic sources. (For instance, Sandvine's Global Internet Phenomena report shows Netflix and YouTube alone comprise over 26% of global internet traffic ccianet.orgccianet.org, and video streaming as a whole represents well over half of all traffic. Yet planning tools typically lump this into generic "internet" demand.)
- **Fixed Layer Boundaries Restrict Optimization:** In many networks, certain design choices are effectively locked in per layer. For example, the optical layer might be designed with fixed routes for all wavelengths (optimized once and not easily changeable), and the IP layer then has a fixed topology over those wavelengths. The siloed tools reinforce this: once an optical plan is set, the IP tool won't consider alternatives that deviate from it, and vice versa. Some advanced networks use SDN to allow more dynamic reconfiguration (e.g. remapping an IP link to a different optical path on the), but planning tools have not caught up to exploit this flexibility. As a result, suboptimal routings persist – for instance, an IP link might take a longer fiber route even when shorter fiber is available, simply because that's how the optical layer was initially built and the IP planning tool treats it as given. The rigid layer boundaries prevent holistic trade-offs.

Overall, today's siloed approach leads to wasteful capital expenditure and operational rigidity. Analyses by industry experts have compared legacy multi-layer architecture versus a converged approach, highlighting issues like manual cross-domain coordination and duplicative "shadow" capacity in each layerfile-wog2boktstmr3n7hpj4hh4. The bottom line: operators often deploy far more equipment (routers, transponders, fiber) than needed to meet actual traffic demands with reliability. One joke is that backbone networks are so overbuilt they could handle the traffic of several Internets – yet the operator still worries about capacity because the tools make it hard to see how much headroom truly exists across layers. This is the gap NetOpt aims to fill.

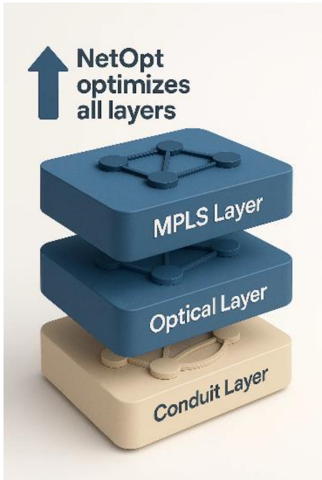

Figure 1: Traditional network planning silos vs. NetOpt's integrated multi-layer optimization. Separate IP/MPLS and optical tools today plan in isolation, whereas NetOpt operates across all layers (physical fiber, optical, and IP) in one model.

This diagram illustrates how NetOpt provides a unified planning model spanning all layers, in contrast to the segregated approach of conventional tools.

NetOpt's Integrated Multi-Layer Approach

NetOpt is built to break down these silos by treating the network holistically. It simultaneously considers the IP/MPLS layer, the optical/WDM layer, and even physical infrastructure constraints, under a single optimization framework. The core idea is to have one planning model to optimize the entire stack, rather than sequential per-layer optimizations. Key capabilities and innovations of NetOpt include:

- **Joint Fiber, Wavelength, and IP Topology Design:** NetOpt doesn't assume the optical topology is fixed. It can decide not only how to route IP flows, but also how to route the underlying fiber paths that carry them. Fiber routes, regeneration sites, and wavelength assignments are variables in the model. This means NetOpt can make trade-offs between adding a router hop vs. using a longer optical path, etc., to find a truly optimal end-to-end design. All layers are optimized together – a single objective function (e.g. minimizing total cost or maximizing utilization efficiency) considers both IP and optical equipment costs. This is in contrast to existing tool workflows that might optimize the IP network given fixed transport capacity, or separately optimize the optical network given static traffic demands, but not both at once.
- **Incorporating Physical Infrastructure Constraints:** NetOpt is aware of physical layer details like conduit routes, site locations, and even rack space or power availability at sites. This is important in practice – an "optimal" network design on paper might require installing a new device in a location that has no room or power. NetOpt can include such constraints so that it only proposes solutions that are actually implementable. For instance, if one data center is out of space, NetOpt might route additional capacity through an alternate site that does have space (even if it's a slightly less direct path). Traditional planning would likely ignore such practical constraints until very late in the process, often forcing manual redesigns. NetOpt bakes them in from the start.
- **Modeling Multi-Period Traffic Growth:** A breakthrough in NetOpt is modeling traffic demand not as a single static matrix, but as a time series (multiple demand scenarios representing different periods such as current load, future years, busy hour vs. off-peak, etc.). This allows *optimization over time* – ensuring the network meets today's demands and can be efficiently upgraded for tomorrow. It avoids overbuilding for an overly distant future. NetOpt can optimize a multi-year plan: for example, deploying just enough new fiber or wavelengths each year to meet actual growth, rather than a huge one-time build

based on an aggressive 5-year forecast. It can even suggest the timing of upgrades (e.g. which year to light a new fiber pair) to balance performance and cost. Because these are computed in one model, NetOpt finds synergies – perhaps re-routing some traffic in Year 3 is cheaper than laying new fiber, etc. These are decisions that layer-specific tools would never consider.

- **Global Time-Zone Traffic Patterns:** For networks spanning large geographies, traffic peaks occur at different local times (e.g. East coast evening vs. West coast evening). NetOpt explicitly accounts for this by using time-dependent demand matrices. Research has noted that networks covering multiple time zones see non-coincident peaks, and capacity can be shared if managed cleverly patents.google.compatents.google.com. NetOpt can exploit this diversity. For example, in a transcontinental backbone, East Coast video demand peaks at a different hour than West Coast's. A wavelength could carry East Coast traffic at one time and West Coast traffic a few hours later. Current planning tools typically dimension each segment for its absolute peak as if all regions peaked simultaneously. NetOpt's optimization recognizes the time-staggered peaks – it might schedule or route traffic such that no segment needs to carry all regional peaks at once. By doing so, it safely reduces total capacity needed. Essentially, NetOpt enables *time-of-day aware* network design, leveraging the fact that not all cities or services hit their max demand at the same moment.
- **Explicit CDN and Cache Modeling:** Unlike traditional tools, NetOpt treats CDN caches and content servers as first-class elements in the network model. If an operator has, say, Netflix or YouTube caches at certain peering or aggregation points, NetOpt can model the portion of traffic that will be served locally versus hauled across the backbone. In effect, the traffic matrix becomes partly an output of the optimization – NetOpt might even suggest where to deploy additional CDN nodes or caches to reduce long-haul traffic, if that yields a cheaper overall solution than upgrading backbone capacity. Sandvine's data showing video dominates traffic csimagazine.com can thus be accounted for in the plan: NetOpt ensures the network is built to carry that video in the most efficient way (ideally, serving it as locally as possible). This approach is novel – today's planners usually treat CDN placement as outside their scope, but NetOpt embraces it as part of holistic network planning.
- **Brownfield Optimization Using Existing Assets:** NetOpt is not only for greenfield (new) design. It can ingest the current network state (existing routers, fibers, wavelengths in place) and then optimize *additions* or reconfigurations – a **brownfield optimization**. The software will try to reuse and repurpose existing capacity first. For example, if there are dark fiber strands or unused wavelengths available, it may choose to light those up or reroute traffic onto them, rather than recommending a purchase of new capacity. This addresses the common situation where operators have plenty of latent capacity that is untapped due to siloed planning and poor visibility. NetOpt's global view can identify and utilize those "hidden" resources. The result is a plan that maximizes ROI on what's already deployed before requiring major new investment.
- **Policy Constraints and Legacy Requirements:** While NetOpt optimizes freely across layers, it also allows the user to input constraints to respect real-world policies or legacy requirements. For instance, operators sometimes have fixed MPLS paths (certain high-

priority traffic must follow a specified path for regulatory or security reasons), or fixed wavelength routes (perhaps due to fiber lease obligations or undersea cable routing restrictions). NetOpt lets planners designate these as fixed elements that the optimization must honor. The tool then finds the best solution around them. In other words, NetOpt can optimize around cross-layer constraints that represent business decisions or legacy commitments. (Traditional tools also allow some fixed constraints, but only within their layer – e.g. an IP tool can fix an LSP path, but wouldn't know why an optical route might need to be fixed.) NetOpt can handle cross-layer rules – for example, "this service must go via subsea cable X" – and still optimize everything else jointly.
- **Integrated Resiliency and SLA Compliance:** NetOpt natively handles reliability requirements like latency and availability SLAs. Planners can specify, for instance, that between site A and B, at least one path <50 ms latency is required (for low-latency applications), or that the connectivity must survive any single fiber cut with at least 99.99% availability. NetOpt will incorporate these in the design – perhaps choosing shorter fiber routes for latency-critical sites, or providing extra path diversity to meet the high availability target. Importantly, because it's multi-layer, it can enforce diversity in a smarter way: ensuring that protected IP paths truly go over disjoint *fibers*, not just disjoint routers. If the SLA calls for extremely high availability, NetOpt might even decide to place an additional regeneration site or an extra fiber spur to create a fully diverse ring – choices an IP-only tool would not consider. Meeting SLAs while minimizing cost is a delicate balance; multi-layer visibility lets NetOpt avoid overkill. For example, rather than simply duplicating everything (which is what margin-stacking often does), it might find a set of partially shared resources that still satisfy the SLA mathematically. This precision avoids the "overprotecting" noted earlier.

In essence, NetOpt behaves as a unified "brain" for network planning. It constructs a digital twin of the entire network (IP + optical + physical) and then runs optimization algorithms to find the best possible design or evolution of that network under the given demands and constraints. The outcome is a plan that *right-sizes* the network – often revealing that far less new capacity is needed than siloed tools would indicate. NetOpt targets elimination of the ~10× over-provisioning by tightening the design to actual requirements, with intelligent sharing of capacity across layers and over time.

Comparison of Planning Tools

The table below compares the characteristics of traditional IP/MPLS and optical planning tools versus NetOpt's integrated approach:

TOOL / CATEGORY	LAYER FOCUS	FUNCTIONALITY	CROSS-LAYER COORDINATION
CISCO WAE DESIGN	IP/MPLS (Layer 3)	IP topology modeling, traffic routing simulation, MPLS/Segment Routing path	Limited – can import optical link info for documentation, but no

		optimization, failure impact analysis cisco.com. Produces capacity upgrade plans for the IP layer.	joint optimization; treats optical layer as fixed bandwidth pipes.
JUNIPER IP/MPLSVIEW	IP/MPLS (Layer 3)	Multi-vendor IP/MPLS network planning; detailed IGP/BGP simulation, traffic engineering (LSP design, FRR), capacity planning under growth networkscreen.com. Emphasizes optimizing existing assets.	None – focuses on IP layer only, assumes given link capacities; no awareness of optical feasibility or reconfiguration.
ARIA NETWORKS	IP/MPLS primarily (multi-layer capable)	AI-driven route optimization and capacity management platform. Can compute optimal traffic routing and has patented capability for multi-layer path finding ofcconference.org, but typically used one layer at a time (often as an IP/MPLS TE optimizer or SDN controller input).	Partial – the underlying engine can model IP+Optical theoretically, but integration is ad-hoc. Not commonly deployed for full co-optimization of live IP and optical networks together.
RIVERBED SP GURU PLANNER (DISCONTINUED)	IP/MPLS (Layer 3)	OPNET-based planner for service provider networks. Offered what-if analysis for IP traffic, topology design, and MPLS TE optimization (including capacity forecasting and failure analysis).	None – did not integrate optical domain; any new IP link it recommended had to be provisioned separately in an optical tool.
CIENA ONEPLANNER	Optical (Layer 0/1)	Photonic network design for DWDM systems. Automates wavelength routing, amplifier/regenerator placement, and simulates optical-layer failures westconcomstor.com. Can co-design fiber and wavelength layers in tandem.	Minimal – can correlate IP services to optical paths for documentation, but uses static traffic demands from IP layer; no feedback loop to adjust IP routing or jointly optimize capacity.
CISCO OPTICAL PLANNER (CONP)	Optical (Layer 1)	DWDM network design and validation. Creates optical topology and equipment plan from capacity requests cisco.com, ensuring signal quality for all designed paths.	None – operates after IP capacity is decided; does not inform or adjust the IP design.

NOKIA 1830 EPT / 1390 NPT	Optical (Layer 1)	Optical transport network engineering tools. Plan wavelength routes on fiber, perform optical power/budget calculations, support multi-period upgrade scenarios.	None – IP traffic is an input, not dynamically linked; no IP routing simulation or re-optimization if optical changes.
NETOPT (MULTI-LAYER)	IP + MPLS + Optical + Physical (All layers)	**Joint optimization of IP routing, MPLS paths, wavelength routes, fiber paths, and placement of new equipment (routers, transponders, fiber) under one model. Simulates traffic patterns over time, accounts for CDN caches, and optimally provisions each layer to meet demands with minimal cost.**	**Full integration** – optimizes across layers simultaneously. IP and optical decisions are coordinated (e.g. may reroute IP traffic *and* light a new wavelength in tandem if needed). Feedback loop is internal – no separate iterations between disparate tools. All layers share a unified view of constraints and resource usage.

Table 1: Comparison of current network planning tool categories vs. NetOpt.

As shown above, existing tools each excel within their domain but do not coordinate across domains. Some vendors have started to market "multi-layer optimization" in the context of SDN – for example, adaptive control systems that adjust IP routing when optical links change – but those are operational techniques, not planning solutions. In planning, true multi-layer **co-optimization** has remained elusive. NetOpt's value is in finally providing that unified planning and design capability, fulfilling a long-sought goal in the industry.

Business Impact and Value of NetOpt

The telecom industry has long understood the drawbacks of siloed network planning. A converged planning approach promises to unlock efficiencies by treating the IP and optical layers as two parts of one whole. NetOpt is a realization of this promise – it bridges the gap between packet and transport network planning, incorporating even physical infrastructure and modern traffic dynamics into its model. By doing so, it enables right-sizing of networks without the typical over-conservative stacking of margins that lead to massively underutilized assets. It accounts for the rise of CDN/video traffic and time-varying demand, so operators plan for *actual* usage patterns (e.g. nightly video peaks that sweep across time zones) rather than static worst-case assumptions. NetOpt also preserves the engineer's intent where needed – honoring locked paths or required routes – while optimizing everything else around those constraints.

For an operator, the end result of using NetOpt is a network that meets all reliability and capacity requirements with significantly lower capital outlay. Optical networks need not be highly oversized (as mentioned before, some networks are at 10x) when multi-layer optimization ensures each layer's capacity is utilized and shared intelligently. Because NetOpt plans across layers, it inherently designs a more *resilient* network as well: IP and optical protection mechanisms back each other up, rather than duplicating and stranding capacity. Planning cycles also speed up, as one integrated tool replaces numerous manual iterations between departments.

In summary, NetOpt offers a new paradigm: plan the network as a coordinated whole, not as isolated pieces. The following references highlight both the limitations of current tools and the benefits a multi-layer approach can achieve – benefits which NetOpt is uniquely positioned to deliver. For CTOs and network planners, NetOpt is a strategic tool to improve ROI and agility, reducing overbuild and Opex waste while ensuring the network is ready for the bandwidth-hungry, dynamic traffic of the future.

References:

1. Cisco Systems – *WAN Automation Engine (WAE): Network Design and Planning*. Cisco product page and data sheet for WAE Design, detailing IP/MPLS planning features cisco.comcisco.com.
2. Juniper Networks – *WANDL IP/MPLSView Data Sheet*. Describes IP/MPLSView's traffic engineering, failure analysis and multilayer visibility capabilities networkscreen.comjuniper.net.
3. Faisal Khan (Light Reading) – *"Are IP/Optical Integration Initiatives Moving to Silos?"*. Industry article discussing the silo problem and over-provisioning at each layer lightreading.comlightreading.com.
4. Google Research – *"Network Utilization: The Flow View"*. Study reporting that backbone links are ~10% utilized on average, due to low normal usage and extra headroom for protection research.google.comresearch.google.com.
5. Ciena – *"What are advanced analytics in optical networking?"*. Explains leveraging unused margin in optical networks to avoid overbuild inefficiencies ciena.com.
6. Sandvine (via AppLogic Networks blog) – *Global Internet Phenomena Report 2024*. Shows CDN/video services dominating traffic mix (e.g. Netflix ~15%, YouTube ~11% of global bandwidth) ccianet.orgccianet.org, underscoring the need to plan for modern traffic patterns.