



The Economic Benefits of Automating Capacity Optimization in IP Networks

Peter Fetterolf, Ph.D.



EXECUTIVE SUMMARY

Internet traffic continues to grow at a rapid rate year over year. Traffic growth is driven by recent technologies such as 5G and fiber to the home and also new applications such as AR/VR, cloud gaming, and video conferencing. IP aggregation and core networks carry the bulk of Internet traffic, and they are typically designed as mesh networks with multiple paths connecting an origin with a destination site. In traditional IP networks with shortest path routing protocols, networks are particularly good at rerouting traffic around link or node failures but are very ineffective at optimizing link utilization. Many communication service providers (CSPs) have used a brute force approach to network capacity planning by allocating extra bandwidth to links to ensure there is adequate bandwidth available to support unexpected bursts of traffic or short-term traffic growth. Typically, CSPs will engineer network links such that average utilization is 50% or less.

As network links grow from 100GE to 400GE and larger it is becoming more important to use autonomous capacity optimization to optimize network link capacity. Links that are underutilized can support more traffic; links that are over utilization should support less traffic. In this paper we present a solution to this problem using the Juniper Networks Paragon Automation, which is a network automation suite that includes a Path Computation Engine (PCE) that simplifies traffic engineering, making it possible to leverage benefits provided by transport service paths, such as MPLS/RSVP, segment routing, and network slicing. It enables operation teams to manage strict transport service level agreements (SLAs) more efficiently and dynamically through automated planning, provisioning, proactive monitoring, and optimization of large traffic loads based on user-defined constraints. With this automation, operators can run their networks higher utilization while achieving predictability, resiliency, and SLA guarantees in service providers', cloud providers', and large enterprises' networks. Our study shows that Paragon Automation can help operators increase average link utilization from 50% to 70% or higher.

Autonomous capacity optimization is even more important today because silicon shortages have resulted in supply chain problems. Increasing network capacity requires CSPs to order new components that are delivered via the supply chain.

Delays in the supply chain could result in inadequate network capacity, causing serious network performance problems as well as SLA violations. In this paper we present the results of an ACG business model that compares two scenarios:

- With Paragon Automation
- Using brute-force capacity management

The total cost of ownership (TCO) and return on investment (ROI) model compares the capital expense and operations expenses of a hypothetical network and shows significant savings using a PCE to optimize traffic engineering. The cost of network bandwidth is exceedingly high such that TCO savings in optimizing the network pay for Paragon Automation many times over. Our results show an overall TCO savings of 27%. We also show that even a minor increase in average network utilization of 0.5% will pay for the total cost of the investment in Paragon Automation.

Network Challenges

Over the last 20 years the importance of the Internet has continued to grow, and today the Internet is an essential utility for most businesses, households, and consumers. The Pandemic has only amplified the importance of Internet connectivity as large numbers of businesses, schools, and universities moved to remote work and learning overnight.

Internet connectivity is provided by communication service providers (CSPs) on a set of diverse, interconnected networks. Most CSPs' networks use a hierarchical architecture consisting of access, aggregation, and core nodes to provide network mesh architecture. The mesh provides diverse paths from origin to destination that allows for resiliency and scalability for IP services. Typically, provider edge (PE) routers are used at the edge of the network to provide an interface between a customer's network and the CSP's network. The PE router provides multiple IP services to end customers. PE routers generally connect to core routers that are optimized for high-speed IP transport and scalability. Core routers typically do not provide the same level of services as PE routers. Another critical component of most IP networks are peering nodes. These are routers that connect a CSP's network to other CSPs' networks using the BGP routing protocol. Peering nodes and BGP allow the interconnection of multiple networks into the global internet.

An example of an IP mesh network is provided in Figure 1. The network consists of PE routers, core routers, and peering routers connected in a mesh. The benefits of the mesh are that if a link or node fails, traffic can be rerouted across a diverse path. Additionally, it is possible to use sophisticated traffic engineering techniques to optimize link capacity utilization while maintaining service level agreements (SLAs) for IP services.




Figure 1. Example of an IP Mesh Network

Network traffic engineering is becoming increasingly important because of the tremendous growth in traffic driven by innovative technologies and applications. Many CSPs have converged their IP networks to provide transport for mobile, business, and residential traffic. Some of the key drivers for traffic growth are:

Mobile traffic growth

- LTE migration to DSS: average cell site traffic increases from 300Mbps to 700Mbps
- LTE migration to 5G: average cell site traffic increases to 2.8Gbps

Business traffic growth

- Video conferencing continued growth due to the pandemic
- Video training
- AR/VR and other new applications are driving bandwidth growth
- Edge computing driving new traffic for Industry 4.0 applications

Residential traffic growth

- Smart TVs and video streaming
- 4K/8K TV
- Work at home with video conferencing
- Cloud gaming
- Diverse mix of devices: laptops, smart phones, tablets, gaming consoles, and smart TVs

ACG projects average household traffic of 14.2 Mbps in 2022 growing to 20.1 Mbps in 2025¹ . However, not all traffic is created equally. Diverse sets of applications have different requirements:

- Delay- and jitter-sensitive applications
- High-availability applications
- Bandwidth-intensive applications
- Best-effort applications

Most networks do not have the capability to differentiate services for these applications, but moving forward differentiated services with SLAs will grow in importance, especially for business and Industry 4.0 applications and services.

Traffic growth is driving network capacity growth. CSPs have two options, brute-force capacity and intelligent traffic engineering and traffic optimization, to manage network capacity growth.

Brute-Force Capacity Management

CSPs can continue to use shortest path routing with minimal traffic engineering and use a brute-force approach to adding capacity. This requires CSPs to manage capacity such that average link utilizations are 50% or lower. This allows networks to manage traffic bursts and events with larger than expected demand. Some of the downsides to this approach are:

- Adding network capacity can take months; setting link utilizations to 50% or lower allows CSPs time to upgrade the network
- Underutilized links create large and unnecessary expenses in network bandwidth

Intelligent Traffic Engineering and Traffic Optimization

Alternatively, CSPs can use network intelligence to optimize network capacity and routing, reducing the need to upgrade network capacity and providing significant savings in network capital expense (CapEx) and operation expense (OpEx). An example of this approach is the Juniper Paragon Automation solution.

Given the growth of network traffic, autonomous capacity optimization provides CSPs with opportunities to:

- Minimize link capacity CapEx and OpEx
- Support traffic growth and seasonal traffic bursts
- Expand to new markets

¹ <https://www.acgcc.com/reports/middle-mile-networks-capacity-requirements-for-fix/>

Autonomous Capacity Optimization

Juniper Paragon Automation provides a comprehensive solution to network optimization and capacity management. The key financial benefits of Paragon Automation that are considered in the paper are:

- Optimizing network link capacity, which reduces both CapEx and OpEx
- Simplifying network operations, which reduces labor OpEx

Paragon Automation enables CSPs to simplify, automate, and optimize traffic engineering using a centralized cloud-native controller. Network path design, provisioning, and management are fully automated using centralized path calculation with a complete view of the network topology and real-time traffic. With this automation, operators can increase significantly the utilization of their networks while achieving predictability, resiliency, and service-level guarantees in service providers', cloud providers', and large enterprises' networks. The workflow in Paragon is depicted in Figure 2. The key processes implemented by Paragon are:

- Deploy: Automatically configure and provision the network using segment routing and/or MPLS-TE to optimize transport while maintaining SLAs
- Monitor: On-going monitoring of network performance and SLAs
- Analyze: Discover network topology, routing, traffic, and service requirements
- Optimize: Optimal path computation based on network topology, traffic, and service requirements




Figure 2. Autonomous Capacity Optimization Workflow

The next sections provides an overview of a business model that shows the financial benefits and return on investment (ROI) of deploying Paragon Automation in a CPS's medium-size network.

Business Model Framework and Assumptions

ACG Research developed a detailed TCO and ROI model to analyze the cost savings and ROI for deploying Paragon Automation. The key benefits are reducing CapEx and OpEx associated with network capacity and reducing labor expenses, which are requirements for network operations and management.

In our analysis we compare two scenarios:

- With Juniper Networks autonomous traffic optimization
- Using brute-force capacity management

In a network with brute-force approach to capacity planning has an average node link and peering link utilization of 50%. This is because network traffic is highly bursty. To guarantee SLAs for high-priority traffic there needs to be extra capacity to allow for traffic bursts as well as unexpected traffic growth. In a network with autonomous traffic optimization we assume central traffic engineering and optimization allows links and peering points to run with higher utilization. This is because Paragon Automation will automatically reroute and optimize traffic to guarantee SLAs while also running the links with higher utilization. In our analysis we consider several scenarios for link capacity improvement, depicted in Table 1. Each network is unique, and some networks will achieve greater improvements in link utilization than others. For this reason we consider six scenarios where link utilization improves from 5% up to 30%.

Base Link Utilization with no Optimization	Link Utilization with Autonomous Capacity Optimization	Utilization Improvement
50%	55%	5%
50%	60%	10%
50%	65%	15%
50%	70%	20%
50%	75%	25%
50%	80%	30%

Table 1. Scenarios for Link Utilization Improvement

We model a hypothetical mesh network consisting of 35 nodes, depicted in Table 2.

Node Type	Quantity
PE Nodes	9
Core Nodes	21
Peering Nodes	5

Table 2. Breakdown of Node Types in Hypothetical Network

Provider Edge (PE nodes are the points demarcation where the CSP’s network interfaces with a customer or enterprise’s IP network. The PE routers provide edge IP services, and the cost per port of PE routers is typically higher than the cost per port of core routers, which are primarily used for IP transport. Peering nodes are used to interconnect with other CSPs and the global Internet using the BGP protocol. Peering nodes are typically scalable high-capacity nodes similar to core nodes.

We assume the network supports mobile, business, and residential broadband services with the expectations for demand, presented in Table 3. Traffic is driven by the number of endpoints (base stations, business services, and broadband subscribers and the traffic per endpoint. Our model assumes traffic growth driven by the inputs in Table 3.

Demand Input	Year 1	Year 2	Year 3	Year 4	Year 5
Number of Base Stations per Node	500	550	600	650	700
Mobile Traffic per Base Station (Mbps)	300	800	2000	2400	2600
Number of Business Services per Node	150	200	250	300	350
Business Service Traffic (Mbps)	200	300	350	400	450
Number of Broadband Subscribers per Node	15000	18000	20000	22000	25000
Average Traffic per Broadband Subscriber (Mbps)	13	14.5	16	18	20.1

Table 3. Traffic Demand Assumptions

To calculate the cost of network capacity with and without autonomous capacity optimization we also use assumptions for the cost of router ports, optical transport, and monthly peering expenses. Specifically, we account for:

- 100GE and 400GE PE router port expenses
- 100GE and 400GE core router port expenses
- 100GE and 400GE optical underlay expenses
- Monthly peering expenses

In addition to network capacity expenses we also consider labor operational expenses. We examine the cost of network capacity planning and operations full-time equivalents. The financial model calculates the TCO (CapEx and OpEx) of a network with and without Paragon Automation and also calculates the ROI of an investment in Paragon Automation.

Business Case Results

For the network configuration and demand specified in Table 2 and Table 3 we have calculated the TCO savings for six utilization scenarios with and without autonomous capacity optimization as specified in Table 1. The cumulative five-year TCO savings for each scenario is depicted in Figure 3. This analysis shows that regardless of the level of link utilization improvements significant TCO savings can be achieved. As link capacity utilization improves TCO savings continue to grow.

The savings of link capacity improvements are extremely high compared to the cost of deploying Paragon Automation. We determined that with an **average link capacity improvement of 0.5%** the TCO savings will pay for the cost of Paragon Automation. The cost of Paragon Automation includes:

- Paragon automation software licenses
- Juniper professional services to deploy Paragon Automation
- Operator's labor expenses to deploy Paragon Automation

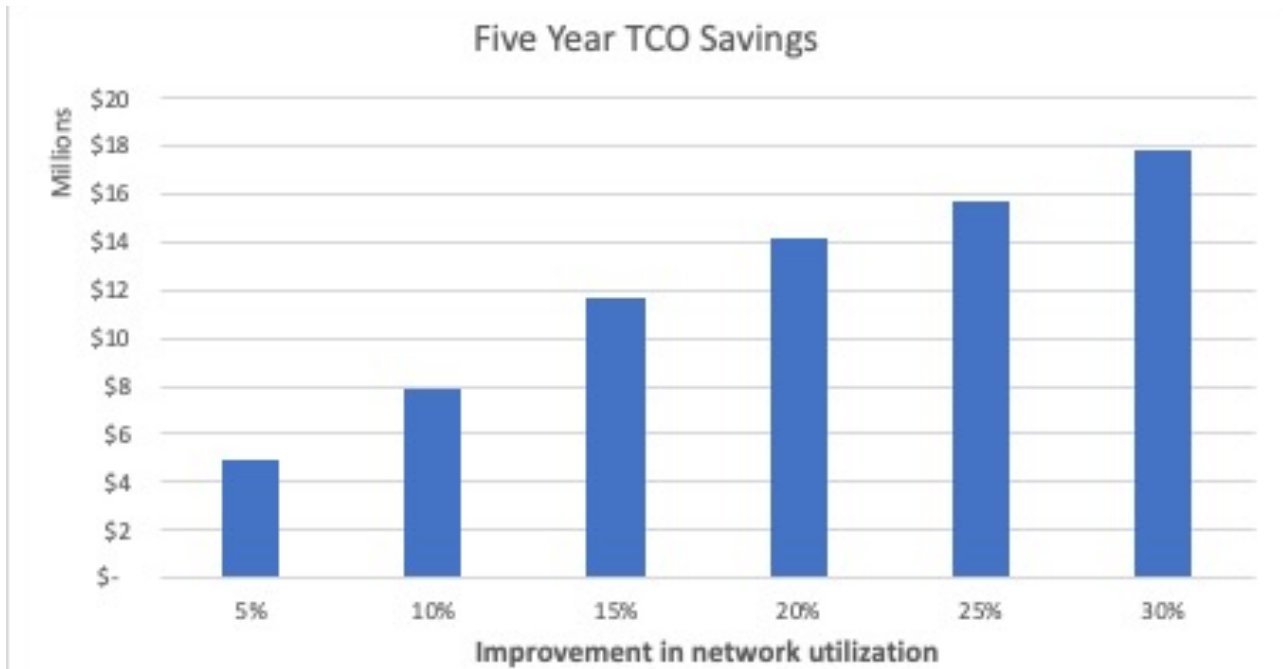


Figure 3. Five-Year TCO Savings for Different Levels of Link Capacity Improvement

In the specific scenario where link utilization improves from 50% to 70% more detail is presented on the TCO results. Specifically, the cumulative TCO results comparing the two scenarios with and without Paragon Automation are summarized in Table 4.

Five-Year Cumulative Results	CapEx	OpEx	TCO
With Paragon Automation	\$ 17,507,864	\$ 19,292,103	\$ 36,799,967
Without Paragon Automation	\$ 24,518,712	\$ 26,014,899	\$ 50,533,611
Savings	\$ 7,010,848	\$ 6,722,796	\$ 13,733,644
Savings Percentage	29%	26%	27%

Table 4. Five-Year Cumulative TCO Results with and without Paragon Automation

The annual TCO spend comparison for networks with and without Paragon Automation are presented in Figure 4. The increase in TCO from Year 1 to Year 5 is driven by the growth in network traffic specified in Table 3. As traffic grows the benefit of autonomous capacity optimization becomes increasingly more important. This means that the benefit of Paragon Automation will be greater in the future as traffic and network capacity continue to grow.

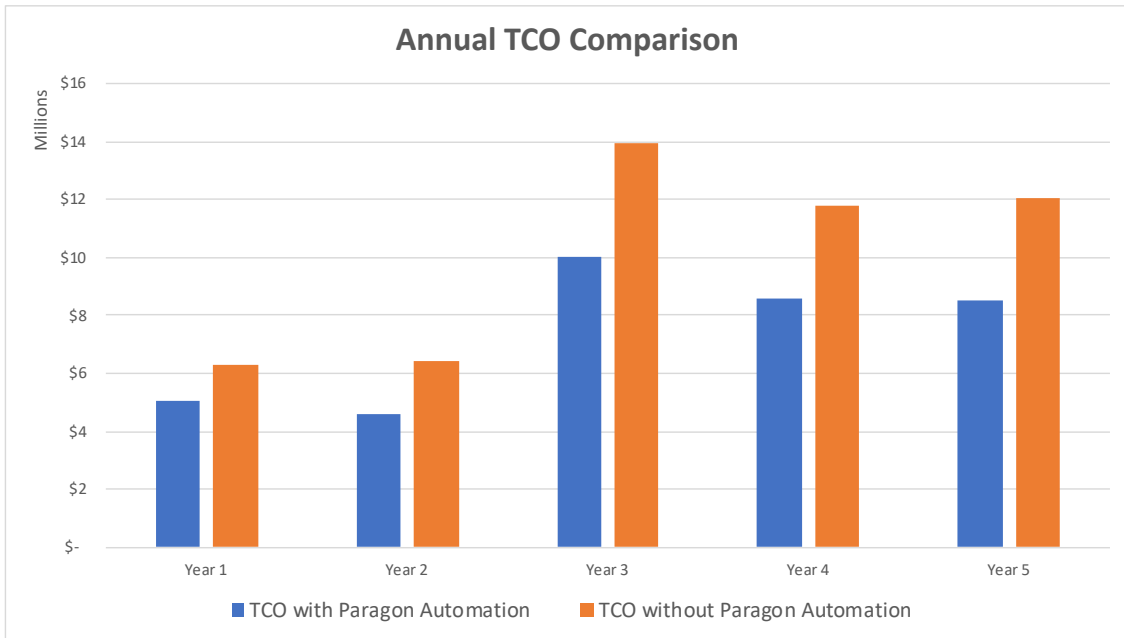


Figure 4. Annual TCO Comparison of Networks with and without Paragon Automation

A breakdown of TCO expenses is presented in Figure 5. This breakdown shows the key drivers of expense and subsequent TCO savings are PE and core router 100GE and 400GE port CapEx and peering link transport OpEx. The router port CapEx savings are a direct result of running the links with higher utilization, and the peering node OpEx savings result from optimizing traffic distribution to peering sites. The total cost of deploying Paragon Automation is small in comparison with the savings.

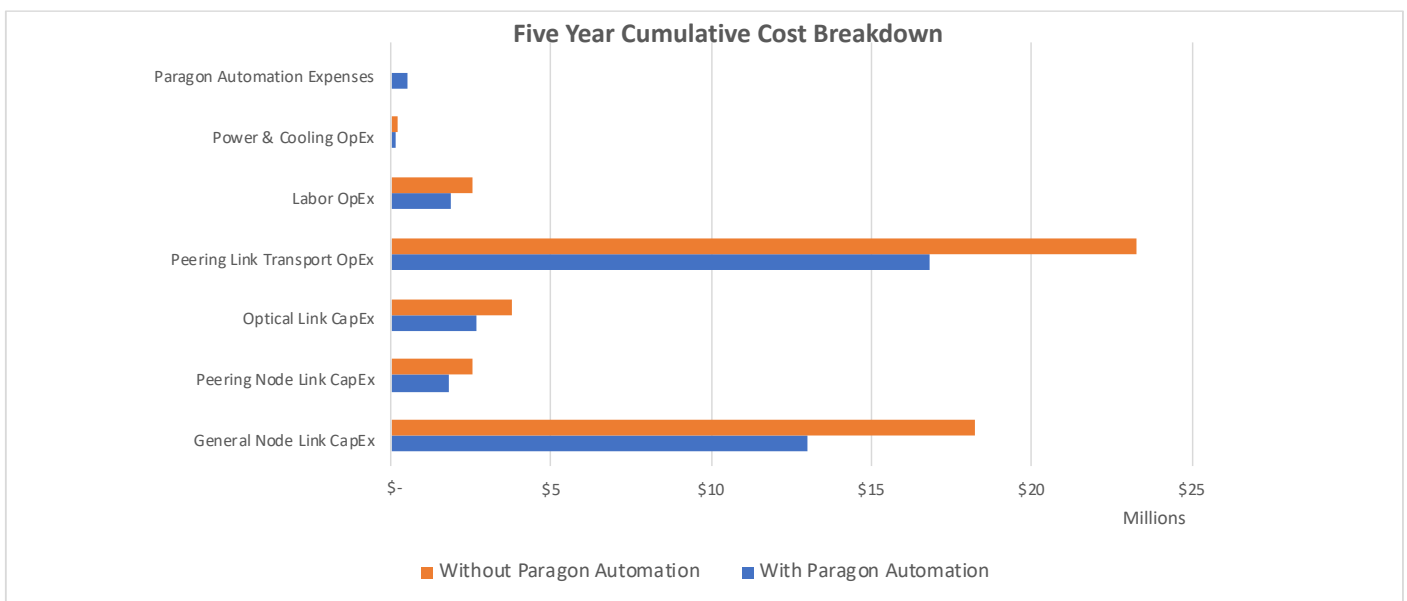


Figure 5. Five-Year Cumulative TCO Breakdown of Expenses

Conclusion

Traffic in IP networks is continuing to grow at a rapid rate, driven by increases in 4G and 5G mobile networks, business traffic, and residential broadband traffic growth. Traffic growth drives expansion in network capacity. The Juniper Paragon Automation can mitigate the expenses related to traffic and broadband, which is expensive. ACG has modeled the TCO of 100GE and 400GE router ports, optical network transport capacity, and peering expenses and show TCO savings of 27% with Paragon Automation. This study has demonstrated that the cost of bandwidth is orders of magnitude larger than the cost of implementing the Juniper Paragon Automation to optimize network path provisioning and traffic engineering. Optimization of traffic engineering results in significant five-year cumulative savings ranging from \$5 million to \$17.8 million for the hypothetical network considered in this paper.



Peter Fetterolf

Peter Fetterolf is an expert in network technology, architecture and economic analysis. He is responsible for financial modeling and white papers as well as software development of the Business Analytics Engine, a platform for simulating 5G networks, SD-WAN, NFV, and general cloud services for service providers, vendors, and enterprises. pfetterolf@acgcc.com

ACG Research delivers information and communication technology market share/forecast reports, consulting services, and business case analysis services. Copyright © 2022 ACG Research. The copyright in this publication or the material on this website (including without limitation the text, computer code, artwork, photographs, images, music, audio material, video material and audio-visual material on this website) is owned by ACG Research. All Rights Reserved.