DAY ONE GREEN: JUNIPER NETWORKS 2023

Chang-Hong Wu, Kapil Jain, Eswaran Srinivasan, Unmesh Agarwala, Valery Kugel, Peter Fetterolf, Sharada Yeluri, David Owen, Attila Aranyosi, Harshad Agashe, Rebecca Biswas, Yedu Siddalingappa, Gautam Ganguly, Christian Scholz, Paddy Berry, Chris Demers, Nell Triplett, Raja Kommula, T. Sridhar, Samuel Rajeev, Deepti Nene

DAY ONE GREEN: JUNIPER NETWORKS 2023

"Businesses worldwide have started to prioritize their climate efforts. From telco providers, to cloud operators to enterprises, collectively we must act now. This will require research, aggressive solutioning and some true creativity to come up with real answers to the challenging questions we face. The internet has become central to all our lives. It has a role to play in solving the most pressing issue facing this generation." - from the *Preface* by Mike Marcellin

"Networks today need to be efficient, focused and resilient if they are to keep up with the current wave of advanced customer and technological demands, in addition to supporting the global responsibility of long-term carbon reduction. This book contains several engineering papers on these strategies as well as the progress that we've already made on new sustainability best practices."

AE Natarajan, EVP, Chief Development Officer, Juniper Networks

"Experience-First networking is about achieving a state of the network where bandwidth accommodates demand, security is pervasive, and infrastructure conserves energy by design and its usage of sustainable power. Juniper's vision of optimally running a green network will happen within this decade."

Manoj Leelanivas, Chief Operating Officer, Juniper Networks

"This book demonstrates to me that the future is green networking and that designing and developing software and systems for conscientious use of resources is one of the world's greatest challenges. Once power efficient benchmarks are established, new architectures will emerge that will provide step function improvements in power efficiency. It's an exciting time to be a part of this journey at Juniper."

Raj Yavatkar, Chief Technology Officer, Juniper Networks

IT'S DAY ONE GREEN AND YOU HAVE A JOB TO DO:

- Discover the properties of a *Green Network*.
- Get details on how power, thermal design, and silicon can improve your gigabytes per watt.
- Learn how Mist microservices architecture can help minimize energy waste.
- Understand the connection between networking and business carbon targets.
- Discover what Juniper is doing with its supply chains and equipment vendors.
- Configure lower power consumption with Junos® in today's production networks.
- Understand the green potential of Cloud Metro, Paragon, SSR, and other Juniper solutions.

9 781736 316078

Juniper Networks Books are focused on network reliability and efficiency. Peruse the complete library at www.juniper.net/dayone.

JUNIPER NETWORKS

Day One Green: Juniper Networks 2023

Chang-Hong Wu, Kapil Jain, Eswaran Srinivasan, Unmesh Agarwala, Valery Kugel, Peter Fetterolf, Sharada Yeluri, David Owen, Attila Aranyosi, Harshad Agashe, Rebecca Biswas, Yedu Siddalingappa, Gautam Ganguly, Christian Scholz, Paddy Berry, Chris Demers, Nell Triplett, Raja Kommula, T. Sridhar, Samuel Rajeev, Deepti Nene

Juniper NETWORKS

Authors: Chang-Hong Wu, Kapil Jain, Eswaran Srinivasan, Unmesh Agarwala, Valery Kugel, Yedu Siddalingappa, Peter Fetterolf, Sharada Yeluri, David Owen, Attila Aranyosi, Nell Triplett, Rebecca Biswas, Gautam Ganguly, Christian Scholz, Paddy Berry, Chris Demers, Harshad Agashe, Raja Kommula, T. Sridhar, Samuel Rajeev, Deepti Nene

Published by Juniper Networks Books
Editor in Chief: Patrick Ames
Editor: Penny Still
Copy Editor: Jessica Harrington

This book is produced and published by Juniper Networks Books. *Day One* books cover network administration with straightforward explanations, step-by-step instructions, and practical examples.

Download a free PDF edition of this book:
https://www.juniper.net/dayone.

For more about sustainability at Juniper:
https://www.juniper.net/us/en/company/climate-change-and-sustainability.html

ISBN Ebook: 978-1-7363160-6-1
ISBN Print: 978-1-7363160-7-8

Version History: v.1, January 2023
2 3 4 5 6 7 8 9 10

Preface

By Mike Marcellin

Welcome to *Day One Green*, a collection of white papers from Juniper Networks' world-class engineers on how to achieve your company's climate goals using Juniper's technologies and expertise.

Since our inception, Juniper has been working on sustainable networking and in these papers you'll notice references to Juniper's engineering breakthroughs that have exponentially decreased the watts per gigabit in our industry. Fast forward to today, where this book kicks in, and our focus on sustainable networking has never been greater.

In response to the global climate crisis, our industry has responded. Nearly every major Juniper customer has their own climate pledge. And those companies need their networking partners – and all IT companies – to do their part to help meet them. Demand for connectivity and applications continues to grow. So how can we meet that demand while still moving toward a net-zero world?

Last year, Juniper made a large step forward by leading our industry to combat climate change. First, we committed to being carbon neutral in our own operations by 2025, the most forceful pledge of its kind in our industry. (You can learn more about that in Juniper's annual CSR report.) But being a good corporate citizen is not enough. We have a unique role to play in helping all businesses meet their climate goals, thus helping the world much more than any single company can. This book represents some of the areas where we're committed to leading this worldwide initiative.

Each paper in this book was written by a Juniper engineering lead. The topics are fascinating: from Juniper's ASICs to power supplies, to AI to Junos to Paragon and to Cloud Metro, you see how green engineering works and how it can help create more power efficient networks than ever before.

Businesses worldwide have started to prioritize their climate efforts. From telco providers, to cloud operators to enterprises, collectively we must act now. This will require research, aggressive solutioning and some true creativity to come up with real answers to the challenging questions we face. I'd like to be the first to thank the authors of *Day One Green: Juniper Networks 2023* for helping accelerate the discussion and I'd love to hear feedback and ideas on how we can all collaborate on sustainability. The internet has become central to all our lives. It has a role to play in solving the most pressing issue facing this generation.

Mike Marcellin, January 2023
SVP & Chief Marketing Officer, Juniper Networks

Juniper Supply Chain Management

By Chris Demers and Nell Triplett

The following excerpt is from the *Juniper Networks Corporate Social Responsibility Report (CSR) 2022*. It focuses on Juniper supply chains and how Juniper sources and monitors its equipment suppliers. The complete report documents how Juniper conducts its business and treats its customers and employees in a sustainable and responsible way. You can download the PDF here: https://www.Juniper.net/content/dam/www/assets/fact-sheet/us/en/2022/corporate-social-responsibility-report-2022.pdf.

Figure 1 *Juniper Networks Corporate Social Responsibility Report, 2022*

Statement

Juniper's mission is to power connections and empower change – to be a responsible global citizen and influence meaningful differences in the world around us.

In everything we do, we act with a commitment to our customers, employees, partners, and the planet.

At Juniper, we develop trust with our customers, our employees, and our suppliers by being honest, respectful, and reliable in all our business dealings. We expect ethical business practices throughout the value chain and encourage our suppliers to adopt critical corporate social responsibility policies that help ensure working conditions in the technology supply chain are safe and workers are treated with respect and dignity.

Our hardware, software, and cloud solutions are developed and produced within a global network of software developers, contract manufacturers (CMs), original design manufacturers (ODMs), component suppliers, warehousing and logistics providers, and recruiting firms. We are committed to meeting our customers' expectations of responsible sourcing practices and transparency throughout this entire networked ecosystem, and we communicate these expectations to our partners and suppliers through codes of conduct and audits based on industry-leading frameworks. We are a member of the Responsible Business Alliance (RBA) and have adopted the social, environmental, and ethical principles of both the RBA and the Joint Audit Cooperation (JAC). We are committed to working with our suppliers and conducting due diligence to help maintain compliance with these responsible sourcing standards.

Currently, we have 141 active measures identified for suppliers where we seek improvement.

Supply Chain Integrity

Protecting supply chain integrity is part of our commitment to protecting brand integrity through all stages of the product lifecycle. Through our supply chain integrity program, we protect our partners and customers by guarding against the introduction of counterfeit or gray market components and the vulnerabilities they could create in our products, and by facilitating failure analysis on products or processes when quality problems arise.

Our multiyear investment in data analytics focused on component-level risk allows us to predict the likelihood and timeframes of risk impacts and to manage environmental compliance and other sourcing risks. By incorporating an enhanced understanding of key risk factors into our lifecycle approach, we benefit from risk reduction and revenue protection throughout the design and production processes.

We work with the U.S. government and regulatory bodies around the world to meet and exceed security standards and ward off attempts to influence the integrity of our products. We maintain the resilience to quickly adapt, implement, and assure compliance with all new requirements while maintaining business continuity.

To help ensure supply chain continuity, we utilize a third-party risk management platform that offers real-time information on susceptibilities, vulnerabilities, and threats.

Our suppliers are expected to support Juniper's compliance obligations, including trade compliance laws and trade restrictions from sanctioned entities and persons. Our membership in the RBA, and requirement that suppliers follow the RBA Code of Conduct, gives assurance that we observe these rights and expect suppliers to support workers' rights to freedom of association and collective bargaining.

See the *Build Global Resilience* section of the Juniper CSR for more information on environmental sustainability in our supply chain.

Responsible Sourcing Program

We engage with our supply chain partners to work together toward the management of a resilient and responsible supply chain. We encourage the adoption of responsible and sustainable business practices among our direct and indirect suppliers. We aim to manage the majority of our suppliers through a direct agreement, and select our suppliers using sourcing strategies drafted in coordination with our engineering teams. Our Supplier Management Program is based on several key elements, including:

- *Performance Evaluation* - We use our Supplier Excellence Framework to evaluate suppliers, beginning with onboarding and continuing throughout our engagement.
- *Verification and Audit* - We conduct assessments and announced onsite audits of our CMs, ODMs, and critical partners to assess and evaluate their performance compared to Juniper standards and communicate results during business reviews.
- *Certification* - Suppliers must certify that they have read, understood, and committed to complying with Juniper's Business Partner Code of Conduct, which communicates our expectations on important corporate social responsibility standards and is informed by the RBA Code of Conduct and the Ten Principles of the United Nations Global Compact.
- *Accountability* - Both our suppliers and Juniper employees who manage supplier relationships are held accountable for upholding the Juniper Business Partner Code of Conduct and executing on the Supplier Excellence Framework.

If suppliers do not follow Juniper's policies or meet Juniper's performance expectations, we may escalate the matter in the supplier business review process and take the non-compliance into account in supplier performance scorecards. Supplier non-compliance and poor performance on scorecards may result in a determination to suspend, disengage, or take other corrective actions with respect to the supplier.

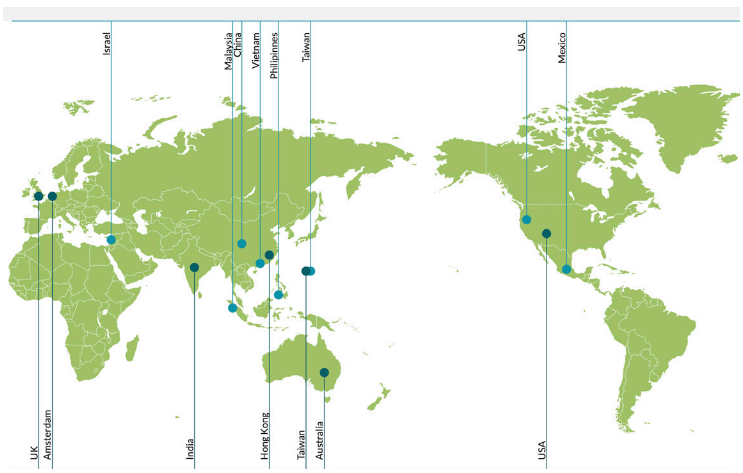

Figure 2 Key Juniper Manufacturing and Distribution Sites

Vetting and Monitoring Suppliers

At Juniper, we work to maintain a culture of integrity and ethics to inspire confidence throughout our ecosystem. It is important that our suppliers are committed to promoting fair labor practices, upholding high ethical and human rights standards, and making a positive impact on society. We screen all new direct material suppliers and manufacturing partners to confirm their commitment to these important principles through our robust vetting process for new suppliers, which includes examination of financials and compliance as well as comprehensive risk assessments and background checks.

We also monitor tier-one suppliers who represent at least 80% of our direct material expenditure in the relevant calendar year (based in part on forecasted spend), 100% of our CMs and ODMs, and all direct material suppliers who provide what are identified as the most critical technologies and can have a broad impact across Juniper. We verify these suppliers' compliance with the RBA Code of Conduct and the Juniper Business Partner Code of Conduct. Additionally, we conduct Customs Trade Partnership Against Terrorism (CTPAT) security audits and business continuity program reviews at critical supplier sites, and use supplier self-assessments, risk assessments, declarations and certifications, and announced onsite audits to ensure supplier conformity.

Scheduled onsite audits at our CM, ODM, and critical component supplier facilities are crucial to the success of our supplier program. Annually, based on risk assessment results and incident and performance trends, we conduct (or partner with a third party to conduct) social responsibility, security, and loss prevention audits at select CM and ODM, tier-one component, and logistics supplier sites.

During the COVID-19 pandemic, Juniper pivoted to virtual audits to continue to monitor our suppliers' performance. This process is aligned with industry standards, including the RBA risk assessment and Validated Assessment Process (VAP), which require an audit of 25% of suppliers categorized as high risk. All audit findings are tracked to closure in accordance with our corrective action process.

In 2021, Juniper saw completion of 35 VAP audits, the highest number completed since joining the RBA in 2015. We found one Priority non-conformance, for which a Corrective Action Plan (CAP) has been approved and is underway. We discovered 34 non-priority findings, for which 23 CAPs were completed and closed. The remaining 11 CAPs are underway.

While suppliers often pass our audits without significant issue, we have noted some serious concerns reported in our supplier audits:

- Workweek for some employees exceeding 60 hours
- Ineffective system of controls for monitoring working hours
- Missing or unavailable inspection reports for buildings or safety equipment
- Inadequate PPE training

- Lack of controls or labeling over waste storage or hazardous material
- Fire escape doors not operating properly or blocked
- Inaccuracies in wage or benefit calculations
- No reasonable accommodation for pregnant mothers

Our approach is to drive change through corrective measures. For all non-conformances, we have instituted actions to mitigate. We, along with other members of the RBA, pressure suppliers to bring corrective action measures to closure.

Accelerants for Customer Sustainability Adoption

By Samuel Rajeev and Deepti Nene

Juniper's vision is to deliver a simplified and secure experience for those who run networks and those who depend on them. We develop solutions to enable customers to build scalable, reliable, secure, agile, and cost-effective networks. Network hardware products have an average life of 4-7 years, and our AI-driven software simplifies operations, enables automation, and extends the life of the network. A reliable and secure network is now mission critical for most businesses including financial services, government, universities, and hospitals. The pandemic and need for remote work has only accelerated this. However, the network infrastructure also adds to the greenhouse gas emissions under scope 3 at purchase (goods purchased) and scope 2 from usage (energy consumed to operate good purchased). With growing awareness, publicly announced commitments on carbon goals and increasing disclosure requirements on sustainability, our customers are looking to reduce their greenhouse gas emissions and optimize their energy costs and consumption.

As new technologies enable energy optimization through innovation in silicon design, smaller form factors, scale-out architectures and automation, the explosion of connected devices, network traffic, and data storage, particularly at the far edges of the network, can risk overwhelming the sustainability improvements these innovations bring.

Rising energy costs are eating into the margins of our customers and represent 12-13% of revenue for certain data centers, or 20-40% of the operating expense of a service provider. These energy costs are projected to keep growing while the mix of renewable and fossil fuels will switch as global fossil fuel demand is projected to peak by 2025. Irrespective of the source of energy, renewable or fossil, our customers want networks that optimize energy use during their operational life (that could last 4-7 years) to help lower their costs and carbon emissions.

This paper highlights what we have heard from our customers, sustainability experts, and industry analysts with respect to these sustainability trends in technology, networking, and the need to combine forces between customers and suppliers to accelerate the industry towards less emissions.

Our efforts around sustainability are being guided by the question: How can we truly empower change that contributes to a meaningful difference to our customers, suppliers, employees, investors, and last and most importantly, the planet?

Accelerants for Greener Networks

Since the 2015 Paris Climate Agreement, when 150 plus countries agreed to limit global warming to 1.5°C by 2050, governments have accelerated investments in renewable energy and have been attempting to encourage businesses to limit greenhouse gas (GHG) emissions. Many businesses have taken an early lead and have made public commitments to reduce their carbon footprint. Consumer and investor awareness around sustainability is also influencing buying and investment decisions.

As multiple stakeholders prioritize for a cleaner planet, we looked at some of the leading accelerants to understand the impact they will have on the network industry.

Energy

The recent rise in energy costs, specifically in Europe, is forcing businesses to pull back their earnings outlook as they see their energy bills go up 3x in the last 2-3 quarters of 2022. The CEO of a large Swedish telecom operator stated, "We basically saw a tripling in the quarter on average and we have got 70% hedge, but still got volatility going on." The telco has had to trim their 2023 profit outlook due to this sharp rise. Additionally, despite purchase agreements and contract hedges, the steep rise is causing utilities to ration the energy allocation to businesses. As companies are preparing to return to the office, some European businesses are encouraging their staff to work from home and companies like Air France are sharing some of the energy savings by paying their ~11,000 employees the equivalent of $4.14 in euros to work from home up to three days a week. These are some drastic measures in response to the recent rise but they convey how volatile operations and profitability can be with respect to energy costs.

For our customers, energy usage and associated greenhouse gas emissions can represent a significant cost in the lifecycle of a network router or switch. From cradle to grave, the biggest contributor of emissions and energy comes from the usage of the product: ~75-95%. The manufacturing, transportation, and end of life only accounts for a quarter or less of emissions and energy consumption primarily due to the length of time the equipment is being used. Juniper is prioritizing energy optimization from usage of the product via innovation in our silicon and hardware as well as AI-enabled automation.

The networking industry is building capabilities to be able to collect and report on the emissions at various stages of the life of equipment, factoring for various usage assumptions. As no industry-wide measure is being broadly accepted, each vendor publishes their own metric with assumptions, leaving the consumer in ambiguity of choosing the most energy efficient product. However, as energy efficiency reemerges as a critical product performance category, an industry-wide metric will take prominence (see *Paper 12: Network Benchmarking*). Juniper is working with various industry-accepted performance measurement organizations, including Energy Star and EPEAT, to help with a widely accepted and fair comparison of energy efficient products.

As most technology infrastructure is a combination of multiple vendors and constant upgrades, one thing we hear from our customers is their need to become aware which part of their network infrastructure is squeezing energy more than required and potentially upgrade or replace legacy equipment or help automate some of these energy intensive processes. Another emerging ask from our customers is to get to "zero watts for zero loads" state where the network can go into sleep mode and use less, or no energy, when traffic is halted.

With the recent spike in energy costs and projections for its continued growth, Juniper is prioritizing how we think about product design and automation with energy efficiency at the heart of our innovation in order to address these growing concerns of our customers.

Infrastructure

Technology advancements and digital transformation have influenced our customers to pursue a significant build and/or refresh cycle of existing infrastructure.

Telecommunications operators are investing heavily in 5G, broadband, and modernizing their networks via automation. Broadband investment has been accelerated by government funding to connect traditionally underserved geographies, as it previously did not make financial sense for the telcos to connect.

In the U.S. alone, the big four wireless operators have spent over $100B on 5G over the last two years. And there is more to come as the 5G capex cycle is projected to be longer due to the various spectrum bands – millimeter wave, 2.5 GHz, C-band, 6 GHz, and more. These different spectrum bands will make for complex deployment with advanced networking equipment. While the U.S. is, or very soon will, peak its investment in 5G, the rest of the world will follow suit in the coming years. The debate is still unsolved if 5G will be more energy efficient versus its predecessors as it has been claimed to be, or if wider bandwidth, high-performance MIMO antennas, and the densification of 5G cell sites will in fact consume more power.

While the investment in 5G is peaking, especially in North America and Western Europe, service providers are dealing with the complexity of managing and operating multiple parallel networks and the inefficiencies that come along with it. Many operators have announced the shutdown of their 3G networks and an accelerated move away from copper to fiber. All these investments and migrations are being built on the service provider's tendency to capitalize these costs and own the infrastructure. While this build and refresh cycle is slated until the next 'G,' or a milestone network advancement, service providers are choosing network equipment that can be optimized for energy use, automated management, reduced maintenance, and provide circular use at the end of its life.

It is a similar situation with data centers. The growth in public cloud is forcing some of the hyperscalers to build 1-2 new data centers a month and that is in the U.S. alone. This growth is also accelerating their need for renewable energy, making them some of the largest corporate buyers of solar and wind energy in 2021. To optimize their carbon footprint, some of these hyperscalers are considering innovative approaches such as using timber instead of metal beams, or investments in sustainable cement and concrete manufacturing companies that remove CO2 from the environment, and other carbon offset examples.

Microsoft's Project Natick has tested submerging a data center into the ocean for two years and claims that it was 8x more reliable than an on-land data center and achieved a PUE of 1.07 (power usage effectiveness) – some of the greenest data centers in the world average 1.2. Given that 40% of the world's population resides near a coast, there is a high chance similar projects gain momentum to reduce the cost of cooling data centers. A key ask from the vendors is how to make the computers and networks as reliable as possible as no engineer intends to swim down to the ocean bed to fix an issue.

However, due to their scale and ability to negotiate larger energy contracts, hyperscalers benefit from multi-tenancy and utilization versus the private data centers that in best case reach 60-70% utilization and average around 10-20% utilization. This under utilization, need for constant upgrades, and skilled teams to manage the private data centers, is also building a business case to move more workloads to the public cloud – even from a sustainability perspective, as they will not be buying all the GHG that comes with the purchased equipment and they can benefit from the efficiencies of a public cloud provider.

Service providers, data centers, and enterprises are all investing in compute and network capabilities that enable new technology improvements and digital transformation. Their ask from the industry from a sustainability perspective is to increase the life of their investments, make them more energy efficient and aware with improved reliability without compromising on product performance and security.

Regulation

Government commitments made at the various climate summits since 2015 are beginning to materialize in disclosure requirements for businesses and RFPs. The United States, United Kingdom, New Zealand, Japan, Hong Kong, and the European Union are planning to formalize their regulation in the 2023-2024 time frame and are aligned to the global goal of limiting temperature increase to 1.5°C by 2050.

In the U.S., the Securities and Exchange Commission has proposed disclosure requirements that are being evaluated now and will be formalized in 2023. The requirement is to have public-listed companies make disclosures in three categories:

- 1. Material Climate Impact
    - Any physical climate related hazard, such as fires or flooding in each of its locations and assets that are exposed to these climate risks.
    - Disclosure of transition risk and managing these climate related risks, which could be regulatory, technological, market or reputational risk.

- 2. Greenhouse Gas Emissions
  - Audited scope 1 (emissions generated by company's own operations) and scope 2 emissions (generated through the energy it purchases).
  - Scope 3 disclosures, if they are material or if the company has a target for scope 3 (upstream and downstream emissions along with the company's entire value chain).
  - Figure 1 shows the various elements and their categorization into scope 1, 2, and 3.

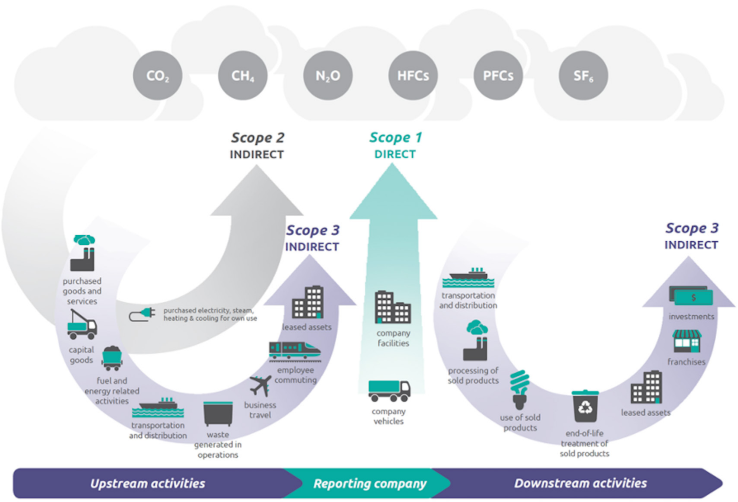

Figure 1 Disclosures of Scope 1, 2, and 3

- 3. Targets and Transition Plan
  - Emission reduction targets, energy use or revenue from low-carbon products.
  - Transition plan to achieve these targets, including the use of offsets or renewable energy.

These disclosure requirements, if adopted, will be required as part of the annual 10-K reporting as well as the 20-F filed by foreign private issuers. Smaller companies will have a year delay to understand the requirements and gather the required information.

Here are three key desired outcomes of these disclosures and their challenges:

- Increase accountability: Make companies accountable by increasing disclosure and audit requirements of some of the ambitious climate goals that they have announced. An audit process will also standardize emission reporting. The challenge remains around collecting data from a company's vast supply chain where manufacturing is done in countries where such disclosure requirements are not mandated.
- Risk evaluation: Provide stakeholders information to assess and evaluate the level of risk that climate possesses to the assets, operations, and future growth of these

companies. Investors, employees, and customers will now be able to evaluate the risk and make decisions accordingly. The challenge remains in a company's ability to identify each and every climate-related risk that they could potentially be exposed to and foresee the impact it might have. For instance, how many companies would have foreseen a global pandemic and the severity and length of the impact?

- Tracking of progress: Globally, 2050 is set as a target year, and many companies have been proactive and brought that target forward. However, many of these targets are eight to eighteen years away and these annual disclosures will help track the progress these companies are making to reach their commitments. The challenge here remains the uncertainty of energy prices, availability of and transition to clean energy, global supply chain, and geopolitical factors that could impact the road map of these companies and call into question their long-term commitments.

As governing bodies and companies discuss and debate some of these increased disclosure requirements, there is one underlining truth that is pushing for increased regulation in this area. Climate-related risk is a growing risk to business operations, and greenwashing, or falsifying commitments, will only go so far. Look for more to come in this space in the near future.

Carbon Goals

Over ~8,000 of the global largest companies have committed to the Race-to-Zero Campaign led by the United Nations starting in 2019. The campaign was initiated to mobilize action outside of national governments and enroll businesses, cities, financial, and educational institutes to get to the "starting line" of the race. The "meta-criteria" to get on the starting line are the five P's:

1. Pledge to get to net zero (latest by 2050)

2. Plan to achieve the goal

3. Proceed to take actions

4. Publish progress and

5. Persuade policy and engagement

Of the broader ~8,000 companies with some kind of carbon goals, ~2,300 companies have a science-based target, which validates that the target set by these companies are in line with Paris Agreement of limiting global warming by 1.5°C by 2050.

These commitments take various forms based on how aggressive or early in the Journey these companies are in reducing their emissions, ranging from Carbon Neutral, Carbon Negative, Net Zero, Net Neutral, to even bold enough to remove all historical emissions since the creation of the company. Most of these goals signify the commitment of these companies to offset their emissions with purchase of carbon credits or significantly reducing their emissions from their own operations and in their supply chain. Refer to Figure 1 for scope 1, 2, and 3.

There is a strong correlation between the larger global companies and B2C companies with companies that have some sort of a carbon goal as consumers become more aware. Per research, 85% of the global consumer has shifted their purchasing behavior towards sustainability in the last five years, and a third are willing to pay a green premium. As of October 2022, ~60% of Juniper's top 50 customers have a carbon goal that was attributable to the supply chain. That means they set a goal for their own operations (scope 1 and 2) as well as a goal to reduce emissions from purchased goods or supply chain (scope 3).

As a lot of these goals and the race to zero are at the very early stages – companies have made some ambitious goals and have started to figure out how to proceed with a plan. This is where regulation will make a difference in helping everyone level set on minimum commitments and quantifiable outcomes. Recent research from Gartner (*Sustainability – It's Complicated*, Nov. 2022) finds that some companies with carbon goals might have been too aggressive in some of these target settings, and given the steep rise in energy costs along with the scale and complexity of achieving these goals, will have to reassess their time lines and road maps.

Leading companies have started to incorporate emission reduction asks from their supply chain. As a supplier to many customers, Juniper has received numerous re- quests from our customers' procurement, ESG, or in certain cases, directly from the net- work buyer to provide details of what our carbon goals or commitments are, what is our plan to achieve those goals, and what is our progress thus far. Some of these requests are quite detailed and, in some cases, we have been asked to set our goals based on when the customer has set their deadline. Quoting one letter from a European service provider: "Target dates should not be greater than 2030."

A big challenge is collecting, responding, and tracking the magnitude of requests as sustainability is gaining prioritization across Juniper's various customer groups. We anticipate a sharp increase in the coming years as more and more companies start collecting, rating, and reporting emissions data from their supply chain. Reporting and rating platforms like the not-for-profit organization CDP (Climate Disclosure Project) help bring standardization and save historical records of company emission disclosures. Consolidation of sustainability reporting and rating agencies will help save time and improve comparability across vendors.

Since Juniper has started sharing emissions data with our customers, we found only a handful of customers (incidentally, all European) who have communicated that our response is currently, or will in 1-2 years, influence buying decisions and will contribute between 5-20% weightage to the entire RFP. We have heard the same message from industry analysts but the industry is not there yet and potentially regulation will be the catalyst. Having said that, the handful of early adopters that want sustainability standards to be part of contract decision making are testing out scorecards and measuring and comparing metrics from various suppliers. At an interaction with one such customer, we were able to see how various suppliers faired on the score card and were told that anomalies exists and reporting and audit processes have not been standardized across suppliers. However, the customer has communicated their seriousness of the ask and is looking for supplier support and commitments to improve our overall sustainability score.

Sustainability is also being incorporated within organizations. We heard from certain customers who are leading by example and incorporating these carbon goals in their operations and capital investment to accelerate the desired outcomes. One hyperscaler, an European Tier 1 telco, allocates shadow pricing to assess purchase decisions or uses an internal carbon fee when calculating project viabilities. We also heard from companies that are leading the race, given their early start, and they now include sustainability metrics as part of leadership performance evaluation that translates into how the business unit makes capital and operations investments.

The biggest impact these goals will have are when they expand from their own operations to the supply chain and when they in some way influence buying decisions, by either an explicit inclusion in the RFP or by selecting vendors that have shown leadership on sustainability standards. For now, product performance and pricing are still the top two most significant decision criteria.

Stakeholders

With longer investment horizons and higher visibility of being impacted by climate related risks, younger consumers and /or investors are becoming increasingly aware of sustainability and looking to make investment and purchase decisions based on how green is the company or the investment opportunity. The growth and rise in ESG portfolios is one indication of this trend. Per a recent PWC report, by 2026, $1 in every $5 invested would be in a ESG-focused fund, predicted to emphasize the point that this is not a short-term trend, but an investment philosophy a younger generation of investors are adopting as they see the benefits that AI and automation will bring along with the move away from fossil fuels towards renewables.

There has been a 38% growth in green talent since 2015, as reported by LinkedIn's Global Green Skills Report of 2022, indicating that talent is addressing the demand from businesses. PWC's 2021 Employee and Consumer Survey stated that 86% of employees prefer to work for employers that care about the same issues as they do.

Consumers are choosing with their dollars and selecting cleaner options and are more open to testing and trying out newer technology that reduce emissions (for example, in 2022, electric cars sales were up 2.6%, while total U.S. auto sales fell 8% in 2022) and companies are updating their product lines and product branding to been seen as a more sustainable product or company.

Planet Earth

Juniper's mutual concern is that the risk to the rising global temperature has reached its resounding heights and now, sooner than later, is the time for the world to unite against this eminent threat. We have started experiencing the natural disasters and unusual climate events across the planet. The Global Risk Report of 2022 published by the World Economic Forum, identifies environmental risks as the top three risks over the next ten years and covers 50% of the top ten risks that were identified. Even within the short term (0-2 years), extreme weather and climate action failure stand out as the top two risks to the world. Unfortunately, the report also identifies that the international risk mitigation efforts for climate change mitigation is only at a 2% effectiveness, and we are still at a very early development stage.

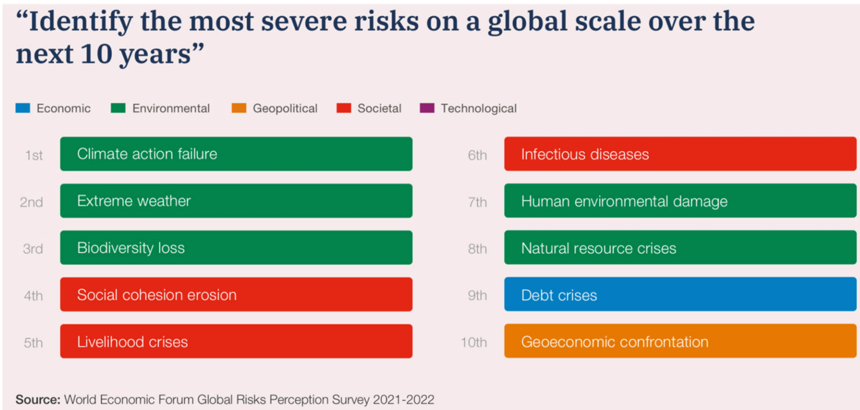

Figure 2 *Source: Global Risk Report of 2022 Published by the World Economic Forum*

The IT industry at large has gathered momentum towards a sustainable planet. Juniper has been inundated with customer requests and the common theme across all of them is that our customers are looking for a partnership to help them in their race to zero and their efforts to meet their carbon goals. We have miles to go but we have great traction and some of our largest customers, who are also the largest employers or businesses in their space, are making serious efforts. Juniper plans on supporting them with products and automation that optimizes energy and meets the objective of a "green network."

Notes

- https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/big-4-wireless-carriers-spent-100b-on-5g-spectrum-was-it-worth-it-68488095
- https://www.fiercewireless.com/5g/mareks-take-2022-year-peak-5g-spending
- https://www.mckinsey.com/industries/oil-and-gas/our-insights/global-energy-perspective-2022?cid=app&cid=app&cid=app
- Gartner: *Sustainability – It's Complicated*, November 2022; https://www.gartner.com/en/documents/4021439
- https://www.simon-kucher.com/en/who-we-are/newsroom/recent-study-reveals-more-third-global-consumers-are-willing-pay-more

Paper 1

Improving Network Efficiency With ASIC Architecture and Technology

By Chang-Hong Wu

Every bit of information that passes through the network needs to be switched from its source to its destination somewhere along its path, usually multiple times by ASICs in networking systems. Therefore, ASIC power consumption and heat dissipation are of paramount importance in keeping the network efficient. Higher power and heat generation from ASICs also require more complex power supply and higher capacity fans for cooling, which themselves consume more power, further compounding the problem.

For many years, improvements in semiconductor fabrication technology kept delivering like clockwork. As captured by Gordon Moore, in the famous Moore's law, semiconductor technology improved by roughly two times every eighteen months or so. If something was too costly, too hot, or too slow, you could just wait for the next generation of semiconductor fabrication technology and essentially get a faster, cheaper, and lower-powered ASIC by riding the technology wave. There was still much hard work to do to better the networking systems in each generation but you also got the benefits from the new technologies themselves.

That's about exactly what happened during the first decade of this century and prior. Each successive generation of networking technology improved the bandwidth, functionality, and power by running the chips at higher speeds and with smaller number of chips, as illustrated in Figure 1. An example ASIC ran at about 156MHz at the beginning of the decade and ended at about 800MHz, greater than a five fold increase, while the number of ASICs comprising the chipsets reduced from ten down to two. It was "the speed era" when we achieved higher bandwidth, better functionality, and lower power by running faster and in less total silicon area. By the end of the decade, Juniper introduced the Trio ASICs in the MX Universal Routing Platforms, a set of ASICs and systems that were so flexible and programmable that they could support features required to not only run many of the networks at that point but for years to come, all at very high throughputs.

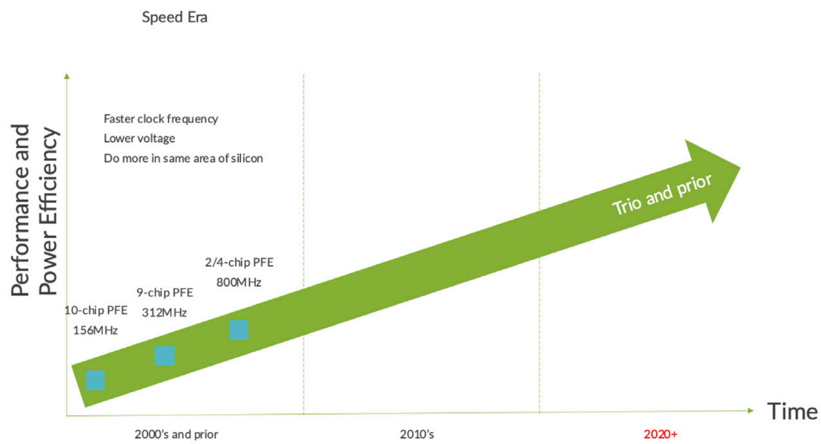

Figure 1 *Decade of 2000s and Prior – Speed Era*

By the latter half of the decade, though, the pace of semiconductor improvement was showing signs of slowing. While the density of logic still increased with each generation of the fabrication technology, the improvements in operating voltage, and therefore, intrinsic power consumption, basically slowed to a stop. By going to smaller and smaller geometries, the intrinsic transistor performance also stopped improving without adversely impacting its power consumption. This trend continued throughout the next decade. In the meantime, with the increased popularity of video and the advent of mega data centers, the demand for bandwidth had never been greater, which put pressure on higher bandwidth and more power efficient networking products.

Here at Juniper, we recognized that off-chip accesses by the ASICs were creating a bottleneck in achieving high bandwidths and they also consumed more power than on-chip connections. From these observations came the Express ASICs and the PTX core and transport routing platforms, with innovations such as Virtual Output Queuing and hash-based longest-prefix-match lookup techniques, which reduced off-chip packet buffer accesses by half and lookup accesses by roughly five times. As a result, by optimizing for the application areas of these networking products to the core and transport, and by using novel architectural techniques, we improved the power efficiency of the end products by roughly two times, in the same semiconductor technology, as depicted in Figure 2.

With the increase in bandwidth and interconnect among the chips, the power consumed by the interconnect made up a larger and larger percentage of the overall chipset power. Working with our industry partners, Juniper ASICs pioneered the use of 3D memories with high-speed serial interfaces, enabling the integration of all the packet forwarding, queuing, and interfacing functions into a single chip. In the latter half of the 2010s, Juniper was among the first in networking vendors to adopt the in-package High Bandwidth Memory (HBM), further consolidating multiple slices and multiple processing cores onto the same dies. These innovative techniques reduced and eliminated much of the high power external interconnects, again producing higher bandwidth and more power efficient networking systems. The improvements made in the decade were mostly due to integration of more functions onto the same die, without much increase in logic speed; that's why I termed this decade as the "SOC era," or System on a Chip.

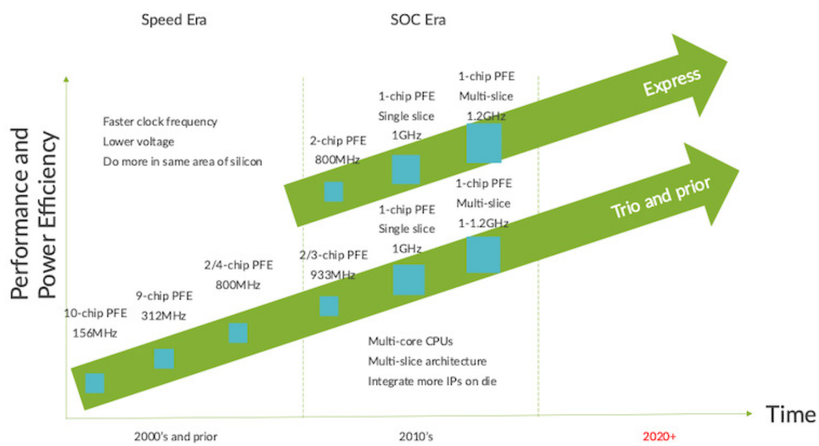

Figure 2 *Decade of 2010s – SOC Era*

The techniques used in the "SOC era" however, resulted in larger and larger die size. In fact, by the end of the decade, many of the networking ASICs were approaching the maximum size that the current processing equipment could handle. In addition, in nanoscale semiconductor manufacturing, defects are naturally occurring phenomenon. Once the die size reaches a certain threshold, the probability of getting a defect-free chip from manufacturing decreases exponentially, thereby increasing the costs of a good product. Clearly, the SOC approach could not continue forever.

In recent years, the industry has embarked on a new approach, sometimes dubbed the "More than Moore" approach. Instead of using separate ASICs on the PCB (thus increasing the interconnect power), or putting everything on a single die (thus sacrificing costs), the idea is to put multiple reasonably-sized dies inside the same package, each in their optimized technology node, interconnected through lower power local interconnects. This way we can continue to increase the functionality of the systems and at the same time optimize the energy costs. This is likely the new approach for the 2020s. See Figure 3.

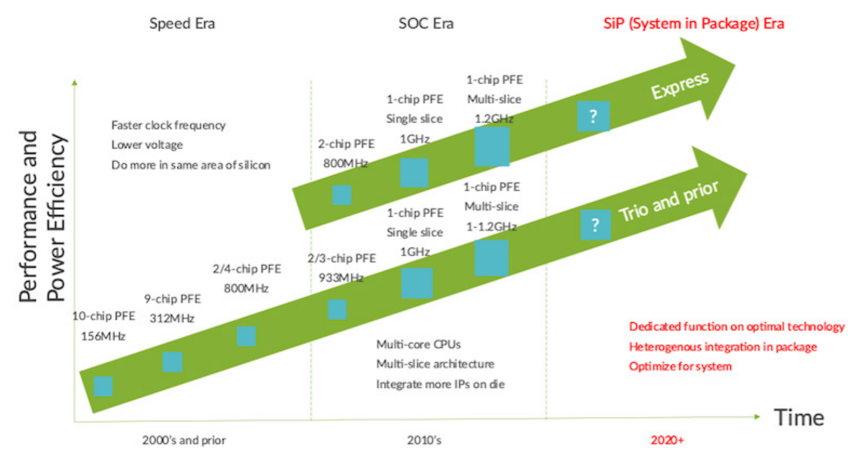

Figure 3 *Decade of 2020s – System-in-Package Era*

Juniper is continuing to come up with new architectural and integration techniques, working with our partners, and improving the economics and power efficiencies of our products in this new "SiP (System-in-a Package) era."

Watch this space for new breakthroughs on how Juniper is improving networking system power consumption with ASIC architecture and technology.

Paper 2

How to Increase Data Center Efficiency With a Lower Carbon Footprint

By David Owen

Juniper Networks has been actively addressing the carbon challenge and its impact on climate since its founding and has focused on leading the industry in power per bit (Watts per Gb) networking bandwidth and features. As networking and computing became pervasive, the projection of power growth and the required ability of the power grid to support it, became known via the work done by Jonathan G. Koomey. During the 2010s, the Juniper power team became actively involved by participating in meetings hosted at LBNL Berkley, joining as an early sponsoring member of the Emerge Alliance (hVDC), participating in Green Grid as a sponsoring member, and attending green energy conferences to understand various power distribution architectures already being adopted in the industry to reduce power consumption and carbon emissions.

While there are volumes of information available, as well as professional services to conduct audits and help get you on the right path to reducing your carbon footprint, a good starting point is the Data Center Maturity Model (DCMM) which is developed and maintained by The Green Grid. The DCMM model is a good resource to audit existing infrastructure and find out where a facility is within the multi-stage roadmap to improved energy efficiency. The model covers all the multiple aspects of data centers such as power, cooling, compute, storage, and network. Evaluating this model can lead to a plan to improve energy while keeping the highest level of performance suitable for the needs of your organization and its available resources.

Okay, given my background in power engineering, let's look at a few best practices at how your power efficiency can be improved. Let's assume you just bought a company and the company came with an existing data center. You're asked to audit the data center to see if it meets the company's high levels of DCMM power efficiency.

You visit the DC site and find it contains multiple isolation transformers due to an infrastructure approach of continuing to add to existing structure and not having the ability to rebuild from a completely new design. This can be very common. The problem is that each isolation transformer has an efficiency loss associated with it and having multiple transformers in a series can lead to compounded loss. Even if the isolation transformers are not in a series connection, often they are operating at low load or unbalanced load. These small details can lead to lower efficiency. It is quite possible to have an efficiency of <90%. Table 1 provides efficiency data for transformers rated with an Energy Star rating or better. Low efficiency can easily occur in a data center that is expanded without careful planning and knowing beforehand the operating load and power balance in each transformer.

Table 1 *Knowing the Expected Loading of the Transformer Will Determine the Final Efficiency*

KVA (Three Phase)	NEMA TP-1 (Energy Star) Federally Mandated	NEMA PREMIUM CSL-3* Not Federally Mandated	DOE 2016 Standards Federally Mandated
15 kVA	97.0	97.90	97.89
30 kVA	97.5	98.25	98.23
45 kVA	97.7	98.39	98.40
75 kVA;	98.0	98.60	98.60
112.5 kVA	98.2	98.74	98.74
150 kVA	98.3	98.81	98.83
225 kVA	98.5	98.95	98.94
300 kVA	98.6	99.02	99.02
500 kVA	98.7	99.09	99.14
750 kVA	98.8	99.16	99.23
1000 kVA	98.9	99.23	99.28

Note: All efficiency values are at 35 percent of nameplate-rated load.
See here for the NEMA TP-1 (or Energy Star labeled).

On May of 2010, the NEMA CSL-3 standards were introduced with higher efficiency ratings than the NEMA TP-1. The benefits of CSL-3 transformers are reduced electrical and heat losses, lower total cost of ownership (TCO), greater energy savings, and green LEED design.

Another best practice and important component in the data center power architecture is the Uninterruptible Power Supply or UPS. There are many types of UPS and the efficiency can range in early generations at 85% and up to 98% for newer generations. The efficiency depends on the UPS type and mode used and the characteristics of the downstream loads connected to the UPS.

Because the UPS has power loss, it is important to have only critical loads operating from the UPS. Basic lighting, air-conditioning, and other electrical needs for facility power and human comfort should not be powered from the UPS to avoid unnecessary power loss. It should only be used for critical loads, such as networking, server power, or equipment cooling.

It's recommended to use an Advanced ECO-Mode UPS, however it may not be possible to take advantage of the Advanced ECO-Mode UPS capability due to critical load transient behavior and ratings. A more modern Advanced ECO-Mode UPS can provide 94% conversion efficiency even in a full conversion mode that protects against voltage dips/sags, transient and interruption on the electrical grid/building inlet. There's more detail concerning Advanced ECO-mode UPS benefits and risks here: https://www.sourceups.co.uk/ups-eco-mode-the-benefits-and-risks/

Up until now my focus has been on a data center with sub-optimal efficiency, so what happens when you find the data center is already designed to be highly efficient. For instance, all the critical loads are powered by dedicated transformers and all non-critical loads, lighting, office air-conditioning, etc., are operating on separate transformers with the ability to minimize power use when not needed. A highly efficient data center is accomplished using networking connected power monitoring and power distribution equipment that is monitored in real-time and can be programmed to be turned on and off based on the need. So with this in place, the focus can be turned back to optimizing the critical loads' efficiency.

There are a few best practices to optimize these critical loads. One is to move to higher distribution voltage from 208/120 VAC to 415/240 VAC. This is well-documented in a white paper at Servertech. The idea is to take advantage of the end equipment power system rating that supports 100-240VAC input voltage. Efficiency gains are typically 2% from eliminating the PDU transformer and an additional 2% to 3% from running compatible IT devices at a higher voltage. Moving to 415/240 VAC can be done by limited engineering resources as discussed in the Servertech white paper.

Taking it a step further, companies with power engineering resources and with the ability to develop their own power solutions embedded into their server racks can use a few different strategies to eliminate the UPS and embed the backup into their server racks, as listed in Table 2. Some popular solutions are to use 480/277 VAC directly to the racks with 12VDC to the IT loads and internal battery backup to replace the UPS. Another is to use 48VDC followed by direct conversion from 48VDC to lower voltages required by CPUs or other chips. And yet another solution is to bring 3ph directly to the rack and into power supplies that convert to 12VDC or 48VDC to the IT loads. And a growing choice is to distribute 240VDC or 380VDC (hVDC) with battery backup at the 240V/380VDC rectifier and then distribute it throughout the data center to the IT equipment.

Table 2 Juniper Networks Advanced Universal Input PSM Use Cases

PSM Input Source		Configuration	Availability	Efficiency	Comments
A	B				
208 VAC	208 VAC	Source A & B both Eco Mode UPS Full Conversion	Highest	UPS 94%, PSM 94% Combined 88%	Lowest Risk
240 VAC	240 VAC	Source A & B both Eco Mode UPS Full Conversion	Highest	UPS 94%, PSM 94% Combined 90%	Lowest Risk
240 VAC	240 VAC	Source A Eco Mode UPS Full Conversion Source B Advanced Eco Mode	Highest High	UPS 94%, PSM 96% Combined 90% UPS 97-98%, PSM 96%, PSM 96% Combined 93.5%	PSMs on Source B will experience a few milliseconds of distorted voltage with grid outage
277 VAC	380 hVAC	Source A Advanced Eco Mode UPS Source B 380 hVAC with Battery	High Highest	UPS 97-98%, PSM 96%, PSM 96% Combined 93.5% 380VDC hVDC 97-98% Combined 93.5%	PSMs on Source A will experience a few milliseconds of distorted voltage with grid outage
380 hVAC	380 hVAC	Source A 380 hVAC with Battery Source B 380 hVAC with Battery	Highest Highest	380VDC hVDC 97-98% Combined 93.5% 380VDC hVDC 97-98% Combined 93.5%	Low Risk 380 hVDC arc'ing resistance detection needs to be provided

Juniper Networks supports these advanced design choices by using power supplies that support universal input voltage and can operate from 200-240VAC, 277VAC (hVAC), and 240/380 VDC (hVDC). By using a universal input power supply, administrators can operate routers and switches based on their choice of power distribution (note uses 80Plus Standard Titanium efficiency: see https://en.wikipedia.org/wiki/80_Plus.

To summarize, for any serious discussion about data center efficiency optimization you should:

- Use the Data Center Maturity Model for evaluating an existing data center.
- Understand the architectures and steps needed to move from a sub-optimal efficiency to a higher efficiency.
- Research the many popular, purpose-built, power architectures that can be adopted for dedicated DC power teams.
- Speak to your Juniper Networks account manager or Professional Services rep about universal input PSM use cases to support your preferred architecture to improve data center efficiency.

Paper 3

Optimized Thermal Design

By Attila Aranyosi, Rebecca Biswas, and Gautam Ganguly

Keeping the power consumption of high-performance, high-bandwidth networking equipment as low as possible is a critical design requirement imposed by Juniper's customers (facility/data center owners) for reduced OpEx (operational expenses) which is another way to say a reduced carbon footprint. We design lower-power systems that can help our customers reach their network's scope 3 carbon targets (see *Accelerants for Customer Sustainability Adoption*). That's because efficient thermal design of our high-bandwidth products contributes to reduced system power consumption and we optimize the thermal management solutions at each level: component, board, and system.

To reduce power consumption at the component level, thermal engineers work closely with the ASIC team to evaluate different floor plan options and identify the best arrangement that meets both electrical and thermal requirements with the lowest possible heat flux levels while mitigating hot spots and reducing leakage currents. Figure 1 shows an example of a thermally inefficient and an optimized MCM (Multi-chip Module) floor plan.

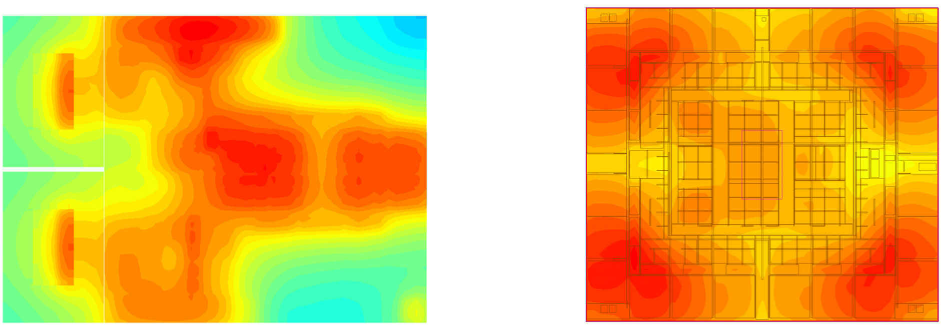

Figure 1 *A Thermally Inefficient MCM Floor Plan Compared to an Optimized One*

ASIC power efficiency is continuously improving (see *Paper 1: Improving Network Efficiency with ASIC Architecture and Technology*) by moving to new technology nodes, but ASIC power density keeps increasing too, because the ASIC and system bandwidths advance faster than the efficiency improvements. To keep pace with these trends and to keep ASIC junction temperatures below their long-term reliability limits with the lowest possible power consumption of the cooling subsystem (in air-cooled systems, fans), Juniper uses lidless, a.k.a. bare die ASICs and MCMs. These eliminate a high thermal resistance element, the TIM1, or thermal interface material between the chip and the lid of lidded packages, from the conduction heat transfer path. Furthermore, Juniper uses the highest-performance TIM2s between the chips and their heat sinks to maximize cooling efficiency.

Besides the ASICs and MCMs, high-bandwidth 400G and 800G pluggable optical modules also consume significant power (between 12W and 25W per module), which in an 1RU 36-port line card, can translate into 900W total optics power or about 40% of the total line card power. We work closely with module vendors to influence the thermal design of modules so that the transceivers can be cooled with the least amount of energy. The main options to achieve thermally efficient module designs are:

- Optimized conduction paths from the main heat-dissipating components to the module case where heat is ultimately removed via integrated and riding heat sinks.
- Tight flatness specification for the top surface of the module housing to reduce the thermal contact resistance between the module and its riding heat sink.
- Choosing DSPs (digital signal processors, the highest-power component of optical modules) that have efficient package thermal design, with low junction-to-case thermal resistance and higher junction temperature limit. Figure 2 depicts the surface temperature maps of thermally inefficient and optimized high-power optical modules.

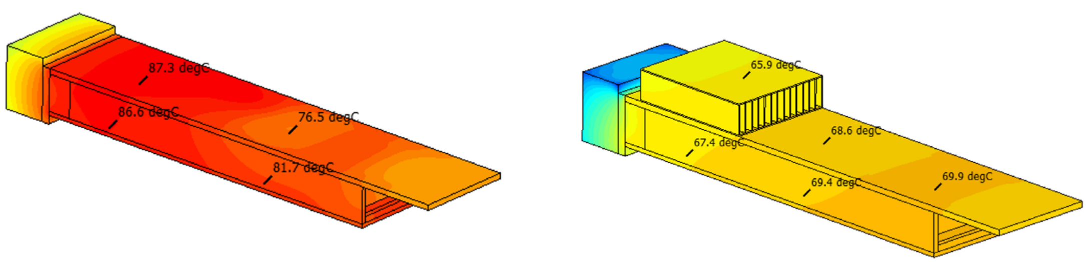

Figure 2 Surface Temperature Maps of Thermally Inefficient and Optimized High-Power Optical Modules Under the Same Boundary Conditions

In many systems, retimers are needed to meet stringent signal integrity performance targets. However, retimers dissipate a significant amount of heat. To reduce system power consumption, flyover cables may be used to replace retimers although trade-offs between power consumption and cost should drive such decisions.

At the board level, Juniper supports HW and SI teams, and carries out thermal feasibility analyses to optimize the board layout and heat sinks and keep component temperatures below their respective long-term reliability limits. We reduce leakage power as much as possible while balancing component thermal margins to keep fan speeds at their lowest levels, further minimizing fan power consumption. Figure 3 shows a vapor chamber main heat sink with a secondary, floating heat sink, which thermally isolates lower-power components with lower temperature ratings from the high-power ASIC. We achieve efficient cooling of the DC-DC power converters (POLs) and a reduced amount of Joule heating in the printed circuit board (PCB) via efficient heat spreading in the power and ground planes.

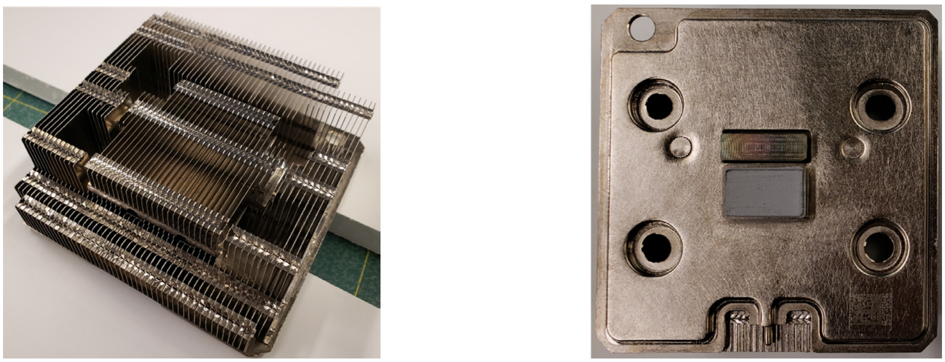

Figure 3 *Vapor Chamber Main Heat Sink With a Secondary Floating Heat Sink*

At the system level, air-cooled equipment is still dominant in the networking industry, and here we select high-efficiency (50-55%) fans which operate in the high-efficiency range against the back pressure imposed by the system. Fan efficiency is defined as the ratio of pumping power (the product of air pressure and airflow rate) and electrical input power. Figure 4 illustrates fan efficiency and aero-performance (P-Q) curves. In the example shown in Figure 4, the maximum efficiency (~48%) is achieved at 105 CFM airflow rate and 5.5 in. w.g. pressure.

With the proper fans and under the worst-case thermal design conditions, power consumption of the cooling system in our equipment, with the fans running at full speed, is typically 10% of the total system power. That's still a significant amount of energy, but with poor fan selection, this number can be twice as high.

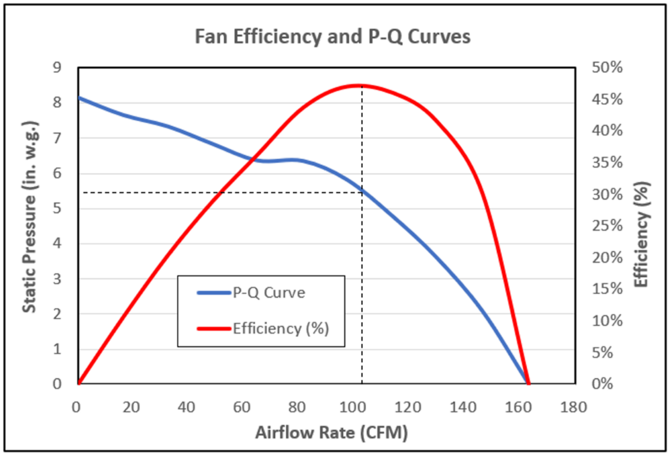

Figure 4 *Fan Efficiency and Aero Performance (P-Q) Curve*

Most of the time our systems operate under nominal/normal operating conditions (that is, the ambient conditions are significantly more benign than the worst-case condition), so we use fan speed control to reduce power consumption. It can result in massive energy savings, as the power consumption of a fan varies with the cube of speed (rpm) ratio. For example, at 50% fan speed, the power consumption is only one-eighth of that at full speed.

Further significant energy savings can be achieved using liquid cooling. Power Usage Effectiveness (PUE) is a ratio between total facility power and power consumed by the IT load. Efficient, liquid-cooled data centers are expected to go below a PUE of 1.10. Although liquid cooling has been used in high-performance computing for quite some time, it is still awaiting acceptance in the networking industry primarily due to reliability concerns. However, realizing that with current power dissipation and power density trends we are approaching the limits of air cooling, there is significant traction in the industry to introduce some form of liquid cooling (cold-plate based or immersion) in the very near future. The other main driver to do so is the potential for huge (50-80%) energy and space savings. In recent years, the Open Compute Project has made excellent progress in establishing a strong liquid cooling ecosystem with standardizations. Juniper Networks has developed several liquid-cooled proof of concept systems (single and two-phase) and currently we have been evaluating multiple newer liquid-cooling technologies in preparation for their deployment.

Juniper's thermal team is continuously innovating to enhance our cooling efficiency at the component, board, and system level. Pay close attention to this space for more *Day One Green* additions.

Paper 4

Connecting Multi-Terabit Packet Processing ASICs Using Multi-Terabit Fabric ASICs

By Harshad Agashe

To build multi-terabit routers, Juniper builds systems using multiple packet forwarding engine (PFE) ASICs. A PFE ASIC supports a few terabits of packet processing capabilities. They are interconnected using cell-based fabric interconnect. Figure 1 illustrates a typical Juniper system. The PFE in Figure 1 could either be Juniper Express or TRIO silicon.

A cell-based interconnect approach provides simultaneous connections to all PFEs without any restriction on flow rates. Given the scale of the bandwidth (BW) requirements, a PFE internetwork is composed of multiple fabric ASICs with multi-gigabit links.

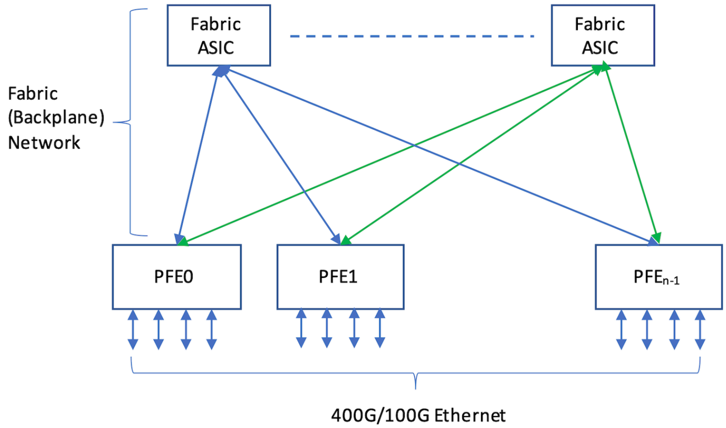

Figure 1 Fabric Network Example

A typical chassis-based system is composed of line cards and fabric cards. Line cards (LC) host PFE ASICs while fabric cards host fabric ASICs. Line cards and fabric cards are connected using high speed serial links. A link segment consists of SerDes on the ASIC, board routing, and connectors and repeaters (for managing signal integrity).

The goal of a fabric network is to provide non-blocking interconnect with almost 90% utilization. Fabric networks should also provide consistent latency behavior with a predictable performance for all packet sizes.

Fabric networks consist of data and protocol traffic along with congestion control mechanisms. Protocol engines are implemented in hardware engines which maintain states per flow.

High-speed links used for fabric networks consume power in PFE ASICs and in fabric ASICs, so a fabric architecture with optimized link count would consume less power over an architecture requiring more links. This implies that protocol traffic needs to ensure performance is achieved along with the power of the system within specified limits. Juniper ASIC architecture optimizes protocols to achieve these goals.

Connecting Multi-Terabit Packet Processing ASICs Using High Throughput Multi-Terabit Fabric ASICs

A multi-terabit system would have multiple PFE and fabric devices. As PFEs are interconnected using multiple fabric devices, PFEs would have to utilize multiple paths simultaneously to achieve performance.

There are two fundamental ways of building fabric-based systems: using cells across the fabric or using packets. A packet-based system may seem obvious since every Ethernet interface of a PFE would be packet-based. However, packets pose challenges in achieving high fabric link utilization.

The Problem with Packets

There are two fundamental approaches to forwarding packets across a CLOS fabric: flow-based and spray.

For flow-based forwarding across a fabric, the flow-identifying header fields of each packet are hashed, arriving at a flow-identifying value. Many actual flows may hash to the same value. This value is then used to select one of the many equivalent paths across the fabric. Thus, all the packets belonging to a flow are assured of following the same path across the fabric and remaining in order. See Figure 2.

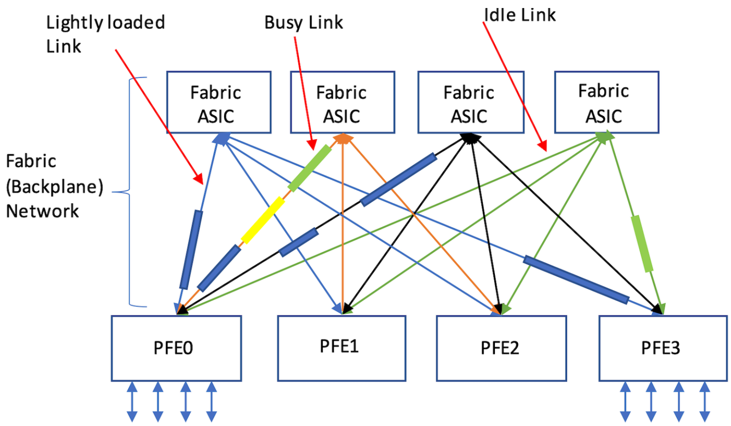

Figure 2 *Flow-Based Packet Distribution*

The problem with this approach is twofold. First, hashing is imperfect, making the distribution of flows to flow-identifying hash values non-uniform. Second, not all flows are created equal: some flows are busy while others are quiet. The hash function has no way of determining how busy a particular flow might be. The upshot of these weaknesses is that some fabric links may reach saturation while others are much less busy. Once one of the links becomes busy, the system has effectively reached its total BW limit.

If Ethernet packets are distributed across fabric using whichever link currently has the shortest queue, the packets within a particular flow can use separate paths across the fabric. This can re-order packets within a flow and that means that packets need to include a sequence number of some kind so that their proper order can be restored. Fabric devices do the store-and-forward and packet size can vary from 64 bytes to 16K bytes. This takes time during which the packet may become mis-ordered and can be quite large, making the re-ordering context within the egress PFE large, complex, and expensive in terms of silicon real estate.

Behavior of the network would be dependent on how flows arrive and could become unpredictable. Fabric utilization can go down to even 60%.

Cell-Based Fabric

A cell-based fabric splits packets into almost-fixed cells and sprays them over all available links. PFE logic will make sure all links are equally utilized for all egress PFEs. This nullifies the problem of saturating a few links over others as in a packet fabric. Cells are stamped with sequence numbers and a protocol engine makes sure oversubscription is handled by granting line rate worth traffic. This ensures fabric devices do not get oversubscribed. Fabric ASICs handle almost-same cell sizes so its design is simple compared to any Ethernet packet switch. The look up engine is simple since it is only switching cells from source to destination. This helps in making fabric ASIC latency behavior predictable and allows a reorder engine design with reasonable complexity and storage.

Overall, a fabric ASIC will consume much less power than a similar capacity Ethernet switch. Cell fabrics can achieve a utilization of almost 90% and they will behave very predictably with consistent latency characteristics.

Cell-based Fabric Components

Let's look more closely at the cell-based PFE internetwork.

Figure 3 shows a system using N PFEs. These PFEs are connected using a NXN crossbar built using fabric ASICs. PFEs communicate over the fabric using protocol engines. Fabric protocol engines enforce non-blocking behavior by exchanging *request-grant* messages. This allows them to compute load on egress PFEs. Engines maintain states for multiple PFEs, or even flows within the PFEs, allowing them to avoid head of line blocking. For example, if PFEs were treated as S streams, then engines maintain state for N*S streams.

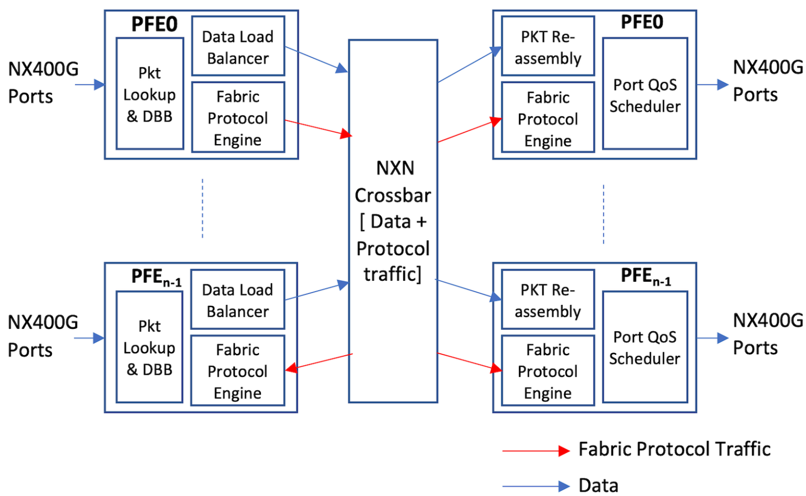

Figure 3 Cell-Based PFE Internetwork

As stated previously the data load balancer will convert packets to cells and spray them over available links by making sure that BW is equally utilized for every egress PFE. Cell reorder blocks put cells coming from various paths back in order and then stitch the packets.

Figure 3's NXN crossbar is implemented using fabric ASICs. Depending on the size of the system (therefore, how many PFEs need to connect), multiple fabric ASICs are used. Each fabric ASIC works internally as a switch which handles cells of almost-fixed size. Lookup engines in fabric ASICs are simple compared to any Ethernet switch and can be statically programmed. Links coming from PFEs are treated as single links allowing large systems to be built.

Fabric Link Utilization and System Power

To achieve full non-blocking behavior, protocol engines need to communicate with each other and that will require BW over and above actual port traffic. BW consumed by protocol traffic will be decided by cell size, cell headers, and request-grant message sizes and their frequency.

A typical fabric network protocol makes sure none of the paths are oversubscribed at any given time. For this to happen ingress-egress communicates using request-grant messages. A request-grant data protocol will have ingress PFEs sending requests to egress. Egress will arbitrate across requests and send back grants and data will flow after grants are received. Egress will grant as per actual flow rate of stream. This makes sure data cells are not oversubscribing any paths.

Data cells would have headers that carry output queue, sequence number, source address, and other control information. Request and grant headers would also carry similar information.

So how much BW is allocated for request-grant traffic and data cell header will be decided by how large a system is built. System power is decided by link count, since SerDes, repeater, and fabric switching power depends on this.

Let's assume data cell size is fixed at 128B and the request and grant header are 8Bytes while each data cell header is also 8Bytes. This implies each cell size would be 128B + 24B = 152B and that would give each cell an efficiency of 128/152 = 84.211%.

Data payload in data cells can be in terms of 4B/8B/16B blocks which can cause wastage when packet sizes don't fall on exact boundaries. In some cases, assuming an 8B alignment, data cell efficiency could be as low as (128-7)/152 = 79.6%. This implies total fabric BW utilization goes down to ~80%.

If cell size increased to 256B and keeping the rest of factors the same, utilization would grow to (256-7)/(256+24) = 88.9%. This implies such system would need ~11% extra BW compared to 20% extra BW.

A typical 50Gbps SerDes consumes approximately 0.5W in 7 nm technology. A system having 4000 fabric links would take 4000 * 0.5 * 2 (SerDes at each end of link) = 4000 W, so a reduction of 10% in links would provide 400W in savings.

A 4000-fabric link system would need more fabric chips over 3600 links. This can roughly save an additional ~200W power. A 10% link reduction will result in ~600W power reduction. Power reduction helps system design (less resources for cooling, power delivery, etc.) and reduces running costs for customers.

The following blog covers the cooling system used by the PTX10008 system. This is an 8-slot chassis with each line card supporting 14.4 Tbps WAN BW. As mentioned in the blog, the number of fans needed to cool down to avoid any component shutting down is significant. Any increase in power further makes cooling costly: https://www.juniper.net/documentation/us/en/hardware/ptx10008/ topics/topic-map/ptx10008-cooling-system.html.

How to Optimize Fabric Protocol BW

There are two major components of protocol traffic: data cell overheads and request-grant messages.

Data cell overhead can be reduced by choosing the data cell size. The size needs to be optimally sized so that the date header overheads are low. A very large data cell size will have low data header overhead, but slow flows may not be able to use it optimally. An approach of slightly variable cell size helps in optimizing such cases.

For request and grant headers reducing its size helps (for example a 12B request and grant header versus 16B) and major improvement can be gained by reducing request and grant traffic. Request BW can be reduced by having a request for a group of cells (1 to N) and protocol engines can dynamically control request size. This is a very powerful tool which allows a fixed allocation for request-grant BW but needs intelligence in the protocol engine to use that as needed.

Summary

Juniper Networks fabric architecture can use all of these design elements to optimize the fabric BW required. Systems built with this approach consume less power, reduce total cost of system, and reduce operating costs. Speak to your Juniper Networks' account manager or professional services rep about fabric architecture optimization.

Paper 5

Optimizing Networks with Efficient System Design

By Kapil Jain, Eswaran Srinivasan, and Unmesh Agarwala

Developing power-efficient networking equipment is a key design metric for Juniper Networks, driven by its customers to minimize their operational expenses and energy consumption. This in turn reduces their carbon footprint and can help them reach their carbon targets during the forthcoming decade. Juniper has been purposefully addressing greenhouse gas (GHG) emissions by making a concerted effort to develop power-efficient networks at both the hardware and software levels.

A typical Juniper modular router platform has the following components:

- Common hardware components including power supplies and fan trays.
- One or two Routing Engines (REs).
- One or more Flexible PIC Concentrators (FPCs).
- One or more Physical Interface Cards (PICs) or Modular Interface Cards (MICs) per FPC.
- One or more Switch Interface Boards (SIBs).

At a high level, the REs provide management connectivity to a router. With dual REs, it also provides hot redundancy for a RE failure scenario. The main purpose of the REs is to have a common routing control plane in a router.

Each FPC contains one or more Packet Forwarding Engines (PFEs) for the processing and forwarding of packets.

The PICs/MICs hosted by a FPC can be fixed as part of the FPC, or can be pluggable. A PIC/MIC supports multiple optics cages with different port speeds for providing WAN connectivity. Certain PICs/MICs have direct connectivity between the optics cages and PFE complex while certain PICs/MICs have gearboxes/retimers between the optics cages and PFE complex.

Each of the PFE complexes in a FPC is connected to all the SIBs in a router to provide a full mesh connectivity without head-of-line blocking.

Figure 1 depicts the common hardware components of a modular router platform.

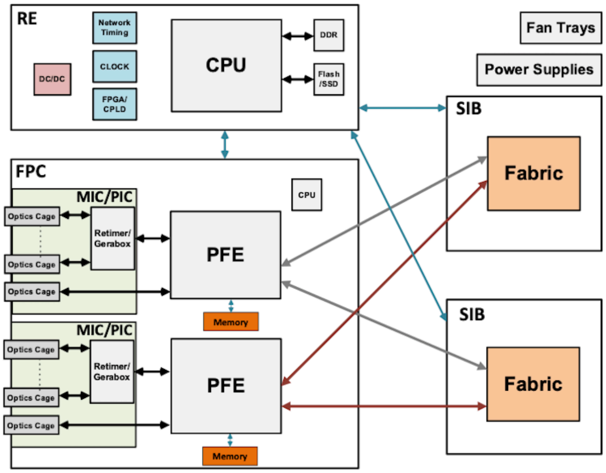

Figure 1 *A Routing and Switching Platform Architecture*

In a networking system, the major power consuming components are: PFEs and their external memories, retimers, switch fabric, CPU subsystem, and fans. Figure 2 details the power distribution of major components as a percentage of total system power dissipation and is based on the measured power dissipation on a Juniper PTX10001-36MR device. Please note that the exact power distribution will vary among different systems but the trend should remain the same.

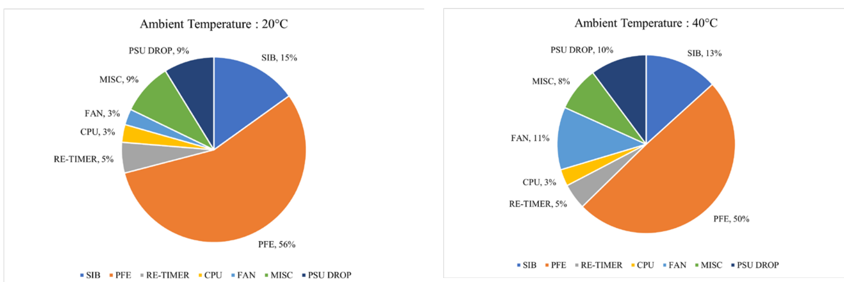

Figure 2 *Power Distribution of Major Components*

To develop a power-efficient system, this paper's focus is to reduce power consumption of these components using the following strategies.

Efficient Thermal Design

Efficient thermal design in our products contributes to reduced system power consumption. (See *Paper 3: Optimized Thermal Design*.) Efficiently applying thermal management solutions can reduce a lot of power consumption in the system. For example, system fans used for cooling the system are the second-highest power consuming component and lowering the fan speed can provide direct power savings. Carefully designed thermal policy makes sure that the fans can run at the lowest speed possible without violating component specifications.

All the Juniper router platforms are designed with efficient cooling systems to support this. As an example, please refer to the following product documentation for more details about the cooling system for PTX10008 and MX10008 router platforms:

- https://www.juniper.net/documentation/us/en/hardware/ptx10008/topics/topic-map/ptx10008-cooling-system.html
- https://www.juniper.net/documentation/us/en/hardware/mx10008/topics/topic-map/mx10008-cooling-system.html

Operational Temperature

Juniper routers are built to operate under different temperature conditions. In general, the power consumption is directly proportional to the ambient temperature. A CLI configuration is available to specify the ambient temperature of a chassis which can help reduce the overall power consumption of the HW FRUs and the provisioned power.

Please refer to the following product documentation for more details about the CLI configuration to specify the ambient temperature of a chassis:

- https://www.juniper.net/documentation/us/en/software/junos/chassis/topics/ref/statement/chassis-ambient-temperature.html

Continuous Monitoring

The temperature of various components are continuously monitored in the router. This includes the PFE ASICs and its external memories. On certain Juniper router platforms, the PFE capacity is dynamically reduced by software when the temperature of the PFE ASIC and/or the external memories goes over a threshold. This mechanism helps to have a better thermal solution with reduced power consumption.

Refer to the following product documentation for more about supporting this functionality for MPC10E-10C-MRATE MPCs on MX240/480/960 router platforms:

- https://www.juniper.net/documentation/us/en/hardware/mx-module-reference/topics/concept/mpc10e-10c-mrate.html

Reset Unused WAN Ports

In many cases, some of the WAN ports are not used in the FPC. It is observed that even though WAN ports are not in use, the connected PFE continues to consume power. From day one, hardware is designed in such a way that each PFE can be kept in reset if not used. Keeping unused PFEs in reset has a significant impact on overall system power consumption. For example, in the PTX10001-36MR, if an unused PFE is kept in reset, then direct power savings of 150W per PFE is achieved. A CLI command is provided in software for users to keep PFEs in reset when the ports are not in use.

Please refer to the following product documentation for more details about the CLI configuration to power ON/OFF a PFE:

- https://www.juniper.net/documentation/us/en/software/junos/chassis/topics/topic-map/chassis-guide-tm-managing-power.html

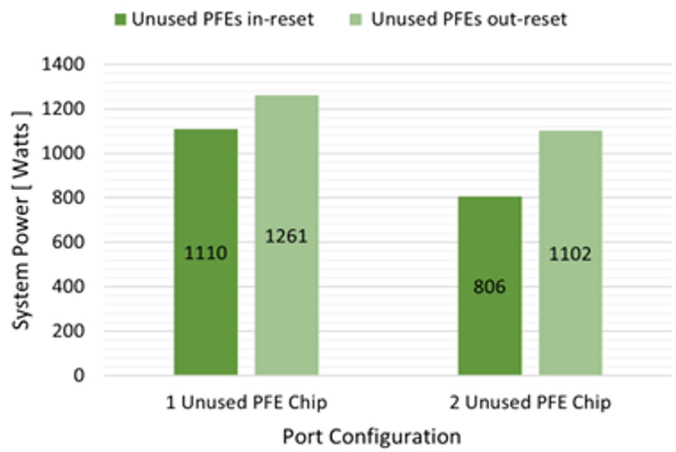

Figure 3 Comparison of Power Consumption Between Cases Where Unused PFEs Are In-Reset and Out-of-Reset

NOTE Similarly, a provision is made in hardware to keep individual Gearbox in reset if not used. This provides power savings of ~18 W per retimer in a few router platforms.

PFE Configuration Flexibility

The external memory connected to the PFEs offer various functionality, including packet buffers and data memory for lookup in the subsystem. The external memory consumption is significant. Juniper ASICs are built with configuration flexibility to use these external memories as needed and the decision can be runtime. The Junos Operating System (SW) is designed to take advantage of this by default and no CLI configuration is required for this. With it, the external memories are used only when needed which significantly reduces the overall power consumption of a PFE complex.

FPCs Support for Pluggable MICs

On certain Juniper router platforms, the FPCs support pluggable MICs. Certain FPC types offer a lot of forwarding features without relying on the traffic on the WAN ports connected to the MICs. Under these conditions, the MICs are not required to be plugged into the FPCs. When no MICs are plugged into the FPC, Junos won't include the power required to operate the MICs and so the overall power consumption of the FPCs is reduced without compromising the forwarding features. This in turn reduces the overall provisioned power required for the routers. This feature is referred to as MIC-aware power management. See Juniper documentation here:

- https://www.juniper.net/documentation/us/en/software/junos/chassis/topics/topic-map/chassis-guide-tm-managing-power.html

When SerDes Links Are Not Initialized

In the case of the FPCs with multiple PFEs, Junos is implemented to initialize links in such a way that if a PFE is not present the corresponding fabric device SerDes links are not initialized.

Similarly, if the fabric chip is not present the corresponding PFEs fabric SerDes side links are not initialized. Using this approach of initializing SerDes links based on the presence of FPC/SIB results in a net power saving of 11W per FPC and 33W per SIB.

Refer to the following product documentation for more details about the CLI show commands used to display the fabric plane status:

- https://www.juniper.net/documentation/us/en/software/junos/system-mgmt-monitoring/chassis/topics/ref/command/show-chassis-fabric-summary.html

When SerDes Links Are Initialized

WAN ports are connected to PFEs using PAM4 SerDes lanes. To save power, SerDes initialization is done only when the optics are inserted into the port. This results in direct power savings as unused SerDes links are not kept in power off. Using this approach ~5 Watts power per port is saved.

Please refer to the following product documentation for more details about the CLI commands that can be used to sanitize the health of the SerDes used for a WAN port and to display the number of SerDes lanes used by a WAN port:

- https://www.juniper.net/documentation/us/en/software/junos/interfaces-ethernet/topics/task/collecting-prbs-statistics.html

Power Off Unused WAN SerDes Lanes Based on the Port Speed

Building on the idea of initializing only those WAN SerDes lanes where optics are present, you can also power off unused WAN SerDes lanes based on the port speed. For example, if a port is configured at 100G speed, then only four lanes are active versus a port configured for 400G speed where all eight lanes are used. This approach results in power savings of 0.6 W per SerDes lane and for each 100G port it would be 2.4W per port. Moving away from static SerDes initialization to dynamic initialization methods, based on optics presence and speed configuration, can result in significant power saving.

Refer to the following product documentation for more details about the CLI commands that can be used to sanitize the health of the SerDes used for a WAN port and to display the number of SerDes lanes used by a WAN port:

- https://www.juniper.net/documentation/us/en/software/junos/interfaces-ethernet/topics/task/collecting-prbs-statistics.html

Clock Gating MACsec Blocks

Juniper's MACsec feature is supported on all ports of a PFE complex up to 400G port speed. MACsec blocks are initialized as part of the ASIC initialization process. MACsec blocks are clock-gated during init to stop this power drain. In the case you want to use MACsec there is a CLI command to enable and disable clock gating of MACsec. Using this approach of clock gating MACsec blocks when not in use can result in net power savings of 20W per PFE.

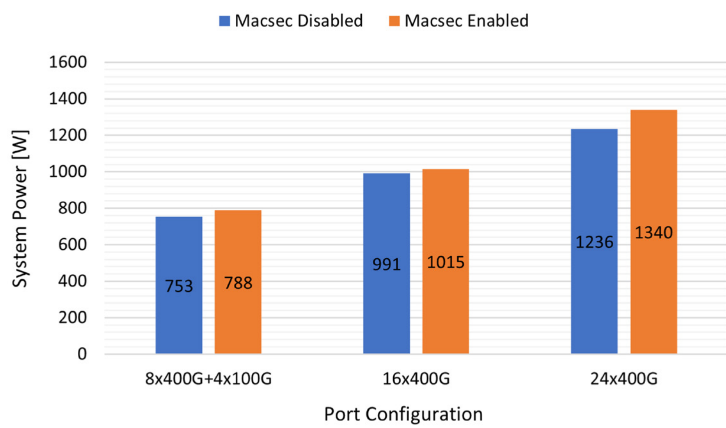

Figure 4 Comparison of Power Consumption Between MACsec Enabled vs. Disabled

Refer to the following product documentation for more details about configuring and managing MACsec:

- https://www.juniper.net/documentation/en_US/day-one-books/DO_MACsec_UR.pdf.

PFE Chip and Switch Fabric Chip

PFE and Switch Fabric are some of the most power-consuming components in the system. So to reduce overall power consumption of the system, the PFE chip and Switch Fabric chip must be addressed. In general, the total power consumed by the ASIC is the sum of static power and dynamic power. Static power is directly proportional to the voltage2 while dynamic power is directly proportional to the voltage2 and frequency.

To reduce the ASIC chip power consumption, core clock frequency reduction and core voltage reduction is done without compromising on chip performance.

For example, MPC8E for MX2008/2010/2020 platforms can be configured in 960G or 1.6T per-slot bandwidth mode. It is worth noting that 960G per-slot bandwidth is the default mode and 1.6T per-slot bandwidth mode can be enabled using a CLI configuration command.

In 960G per-slot bandwidth mode, the PFE ASIC is configured with the core clock frequency of 768 MHz for datapath and 562 MHz for lookup subsystem. In 1.6T per-slot bandwidth mode, the PFE ASIC is configured with the core clock frequency of 862 MHz for datapath and 937 MHz for lookup subsystem.

Refer to the following for more details about configuring MPC8E in 1.6T per-slot bandwidth mode:

- https://www.juniper.net/documentation/us/en/software/junos/chassis/topics/ref/statement/bandwidth-edit-chassis-fpc.html

Summary

You can see there are definitely ways to reduce your power consumption with Juniper devices. You can do these today and weave them into your operational best practices. Always test these techniques in the lab before moving into production environments.

Paper 6

Squeezing Every Last Watt From Juniper Express Silicon

By Sharada Yeluri

In *Paper 1: Improving Network Efficiency with ASIC Architecture and Technology*, Chang-Hong explained how Moore's law started slowing down in the late 2010s. While the density of logic still increased with each generation of the fabrication technology, SRAM densities were not improving at the same pace due to the sensitivity of the SRAM cells to process variations. Improvements in operating voltage, and therefore, intrinsic power consumption also slowed to a stop.

External memory technologies were not doing any better, either. In the past few decades, the gap between the processor and memory performance continued to increase at approximately 50% per year and the gap is now at about 1000 times. External memory density and power improvements have also slowed down significantly.

All of this meant that we could no longer rely solely on the process node advances to double the performance in the same power envelope. There was the need to develop an architecture that relied less on external memory accesses and reduced the data movement within the chip and to the external memories to reduce power consumption.

With that in mind, Juniper set about developing the Express architecture whose main intent is to deliver very high density and power efficiency for transport and core routing applications. The first family of Express chips were introduced in 2012 with the PTX Series (Packet Transport Routers). A decade later, we are currently sampling Express 5 (fifth generation) chips with the best power density (watts/G) one can obtain for this class of chips.

How did we get such high-power efficiency? A forwarding plane architecture that trades some of the flexibility and scale of our Trio silicon and previous M Series architectures in favor of lower latencies thus lower power consumption. It's an interesting angle for a green engineering perspective and the subject of this paper.

Express PFEs

A typical PTX router consists of one or more Express packet forwarding engines (PFEs). When the router contains more than one PFE, they are connected through the Express cell-based fabric as shown in Figure 1. A PFE typically consists of a packet processing complex, WAN interface that receives the traffic from the Ethernet links, fabric interface to connect to other PFEs in the router, and a queueing/buffering subsystem.

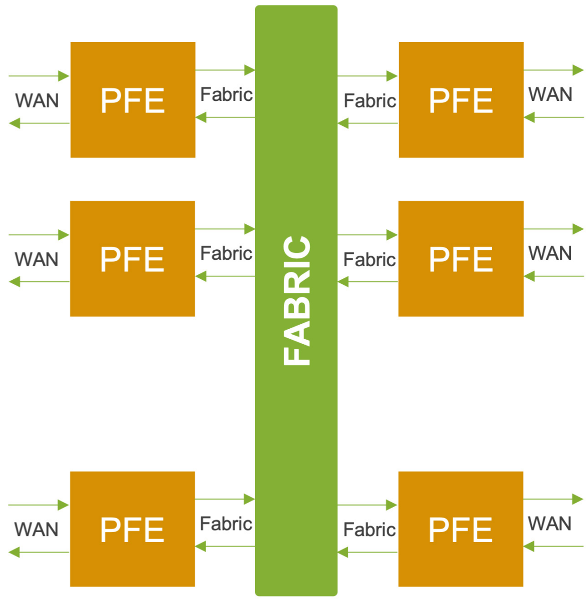

Figure 1 *PTX Router*

Fixed Pipeline Packet Processing

Express uses fixed pipeline packet processing for ingress and egress processing. These consist of a series of subsequent blocks. The packet headers (typically the first 128B-256B of the packet) flow through these blocks (see Figure 2).

Each block performs a specific function on the packet header and passes on that information to the subsequent block and so on. Compared to the network processor cores that are present in merchant silicon switches, or in the packet processing engines (PPEs) in our Trio family of chips, the Express architecture's fixed-pipeline implementation takes a lot fewer cycles to fully process a header. That's because all the functions are either hard-coded or implemented by executing highly specific microcode engines inside each block.

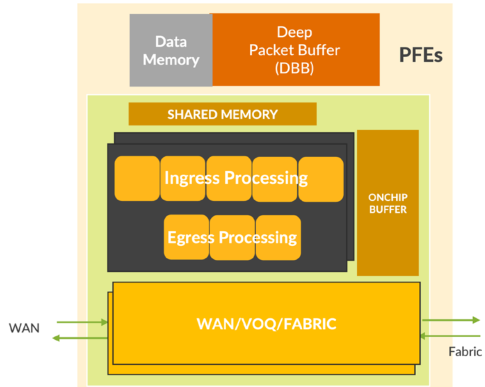

Figure 2 *Express PFE With Fixed Pipeline Packet Processing Complex*

As an example, a typical Layer 3 parsing to identify IPv4/IPv6 headers, perform the length checks, and compute the IPv4 header checksum, takes 1-2 pipeline stages through the parsing block of a fixed pipeline. The same Layer 3 parsing logic would take 2-4x more cycles with NPU/PPE because the software needs to execute a sequence of generic instructions to extract the various header fields from the Layer 3 header and perform the computations.

Fixed pipeline architectures are twice as efficient in the die area and on average take 4-7x less latency to process a header compared to the processor cores.

On-Chip Fungible Data Structure (Shared Memory) and the Caches

The Express architecture also carefully avoids accessing external memory for packet processing data structures like Forwarding Information Databases (FIBs), next hops, tunnels, and encapsulation tables, etc.

Accessing external memory for processing each packet not only consumes a lot more power but also increases the latency of processing.

Even with the advances in memory technologies and with the advent of HBM (high-bandwidth memory), the usable bandwidth from an HBM2E device is around 2.6Tbps. The HBM interface takes up a significant beachfront area of the die edge and Express would be severely limited by the amount of throughput it can pack inside each die if every packet were to access the HBM for lookups. Hence, in Express, most of the look-up data structures are stored in a large fungible on-chip memory (referred to as *Shared Memory*) that can be partitioned between different structures at boot time. Express also allows for some FIB expansion to off-chip data memory that resides in the HBM.

Further, each client implements a lookup cache to store frequently accessed elements closer to where the processing is happening. *These caches additionally reduce data movement which in turn helps conserve power consumption.*

Hash Engines and Bloom Filters

Express reduces the hash/lookup table accesses to the central shared memory by using *bloom filters* that reside within the packet processing blocks. A bloom filter is a space-efficient probabilistic data structure that is used to test whether an element (key) is a member of a set (hash table) or not. Probing a *key* in the bloom filter indicates whether it is present in the hash table that resides in either the central fungible data structure or in external data memory. False positives are possible but there are no false negatives. *Using this approach can cut down on memory accesses by 70-80%, which again in turn saves power consumption.*

VOQ Architecture

The Express data path is based on Virtual Output Queue (VOQ) architecture (see Figure 3) which is a significant departure from the Combined Input Output Queue (CIOQ) architecture used in Trio and in many other high-end routing chips.

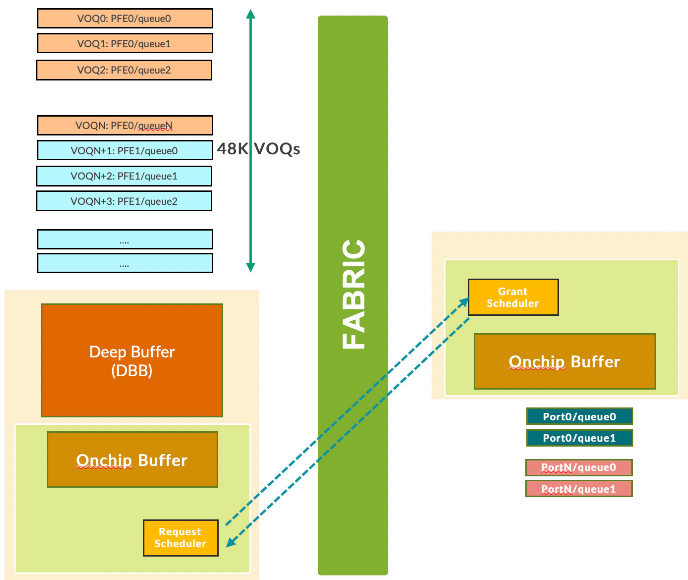

Figure 3 VOQ Architecture

In VOQ architecture, the packet is buffered only once in the ingress PFE after ingress packet processing, in a queue that uniquely corresponds to the final PFE/WAN port/output queue from which the packet needs to depart. These queues in ingress PFEs are often referred to as VOQs or *virtual output queues*. Every ingress PFE has buffer space for every output queue in the system.

A VOQ stays in the on-chip buffer when it is shallow and moves to the deep buffers in the external memory when the queue starts building up.

The VOQ requester waits for a group of packets to accumulate in a VOQ before it sends a request to the egress scheduler (in the egress PFE) for permission to send those packets over the fabric. The egress scheduler grants the access based on strict scheduling hierarchy and the space in its shallow on-chip buffer. Once a grant is given, the packets move to the egress PFE's on-chip buffer through the fabric and leave the PFE through the egress WAN ports.

This architecture is extremely power efficient for a few reasons:

- Packets are queued only on ingress. They reside in the on-chip buffer for shallow queues and move to the deep external buffer only during the congestion – so there's less data movement.
- The egress buffer is very shallow, and packets are admitted to the egress PFE only when it has space in the shallow egress buffer. So, packets never need to be queued in deep external memory buffers on the egress side – and there's less data movement again.
- Once a packet is accepted on the egress, it can't be dropped anymore. Compare this to CIOQ architecture (Combined Input/Output Queued) where the packets are queued in both ingress and egress PFEs, meaning when egress queues are congested, a packet could get dropped after it has moved to the egress PFE. Again, less unnecessary data movement and less power to operate it.

Cell-Based Switch Fabric

Once an egress PFE grants an ingress PFE admission for a group of packets, the ingress PFE 'cellifies' the packets in that group, attaches sequence numbers to these cells, and sprays them over available fabric links. On the egress side, these cells are put back in order and the packets are assembled. By chopping the packets into cells and spraying them across the links, Express can achieve >95% utilization on these links connecting PFEs to the fabric. In *Paper 4: Connecting Multi-Terabit Packet Processing ASICs Using High Throughput Multi-Terabit Fabric ASICs*, Harshad Agashe explores how cell-based fabric is superior in power and performance to Ethernet switch-based fabrics used by many other network vendors.

Multi-Slice Architectures

In the last two generations of Express chips, we squeezed in multiple PFEs inside the same die by sharing the packet processing data structures and the on-chip packet buffering between the PFEs. *This enables us to have a smaller SRAM footprint on the die and not only improves the area efficiency but also saves the leakage power associated with these SRAM structures.*

Power Optimization Techniques During Implementation

While a good packet processing and data path architecture that reduces the data movement and decreases the processing latency can offer significant power savings, power can further be optimized by advanced implementation choices.

The total power consumption of a chip consists of static and dynamic power. Static power is usually the leakage power of the logic gates and the SRAMs, and it is directly proportional to the voltage of operation and the process node. Leakage power is becoming more and more prominent in the latest process nodes, and it is preventing some vendors from lowering the operating voltages on their devices. In Express implementations, we focus on reducing the dynamic power of the chip as that directly relates to the switching activity.

The dynamic power of an integrated circuit consists of switching power and short circuit power:

- Dynamic Power = Pswitching + Pshort-circuit
- Pswitching = a.f.Ceff.Vdd2
- Pshort-circuit = Isc.Vdd.f

You can see that switching and short circuit power are directly proportional to the clock frequency (f) of operation. Supply voltage is often the lowest voltage recommended by the vendor for the process node. Reducing the voltage affects the performance of the SRAMs and the logic gates and can push the minimum frequency of operation to a lower number, which in turn reduces the packets per second and the bits per second you can achieve with a given piece of silicon. In Express chips, we usually keep the operating voltage at the recommended setting by the vendor for that process corner.

Optimizing the Frequency of Operation

While it seems obvious that reducing the frequency of operation reduces power consumption, it can also reduce the performance (power/gigabit). Then, to get the same overall throughput from the PFE or the system, you need to add more logic in the PFE or add more PFEs in the line card/system. Both would add to the power consumed by the ASIC.

A network chip with tens of terabits per second of bandwidth, with central buffers and data structures, has many wide buses that need to be routed across a large die. Operating frequency of the chip decides the width of these buses. A wide bus is often required at lower frequencies to get the same bandwidth. And routing a wider bus involves more repeaters and therefore more power consumption.

Another factor to be considered is the re-use of the existing IP components which might not scale for higher frequencies. Similarly, SRAM performance might not scale with frequencies, so to realize a logical memory you would be forced to use multiple stammer SRAM structures.

For each generation of Express chips, Juniper carefully considers the process node, wiring congestion, IP re-use, and SRAM scaling to select the frequency of operation that reduces the overall system power.

Clock Gating

We provide the ability for software to clock gate (or turn off the clock) for functions that are not used or enabled for the users. By turning off the clock to large chunks of logic, you can save the clock tree as well as the switching power (as the frequency component goes to zero).

For example, the clock network for the logic/functions associated with unused WAN ports is turned on/off by the software as the user attaches/detaches the cables to the WAN ports in system.

We also implement dynamic clock gating. Here, if the output of a flip-flop is not used in a specific cycle, the clock could be turned off for that flop to prevent the output from switching in that cycle. Dynamic clock gating is inferred by the EDA tools during the synthesis (conversion of the Verilog behavioral RTL code to gates) when the designer writes the code for the flip-flops in a specific format. Express uses advanced EDA tools and methodologies to proactively identify and fix all the clock gating opportunities that the designer missed. Our designs achieve >90% efficiencies in turning off the clocks to the flops when their outputs are unused/not changing.

Power Optimization in Placement and Routing

Lastly, Express uses advanced power-driven placement and routing tools and methodologies to optimize power consumption even further. This is the topic of *Paper 9: Juniper ASIC Team Pioneers System-in-Package (SiP) ASICs*.

Summary

Express ASICs are all about switching and transporting packets in core/peer and transport routers as fast as possible with the least possible power consumption. With a novel fixed-pipeline VOQ-based architecture, advanced techniques to reduce the data movement within and across the ASICs, 2.5D packages with HBM memories inside the package, and by using the latest EDA tools and methodologies to reduce the dynamic power even further, Express has not left any stone unturned in achieving the lowest power per gigabit of traffic. Express 5 is some of the most efficient silicon in the market with 28.8Tbps of throughput from a single package.

Speak to your Juniper Networks account manager or professional services rep about Express silicon in the PTX Series of Packet Transport Routers.

Paper 7

Cloud Metro Architecture

By Peter Fetterolf and ACG Research

NOTE This paper was originally published as a Juniper Networks blog: The TCO and Environmental Benefits of the Juniper Networks Cloud Metro Network Solutions.

Cloud computing, edge computing, and metro networks are converging as network operators move to 5G networks and evolve to deliver new services. The emergence of disaggregation for the 5G vRAN and separation of the 5G core user and control plane to meet more ambitious service requirements are accelerating this trend. An ACG Research report published in the first quarter of 2022 defined the characteristics, services, and requirements of Cloud Metro networks.

Juniper Networks is leading this network transformation with a new generation of routing systems designed for modern Cloud Metro networks. Juniper's solution is comprehensive with improvements in router architecture, automation, AIOps, service assurance, and security. The Cloud Metro solution allows operators to flexibly deploy next-generation metro edge services while reducing network total cost of ownership (TCO), energy consumption, and CO2 emissions, not to mention added longevity of networking equipment and hardware lifespan. This new generation of Cloud Metro networks can help you reach your own carbon abatement targets. Your results may vary but when combined with other technologies can provide incremental improvements that ad up over time.

Introduction

This study presents a TCO model of a network of 10,000 Cloud Metro routers and compares the Juniper ACX7509 Cloud Metro router with a similar generation router from two competitors with significant global deployment. Specifically, we compare three routers:

- ACX7509 Cloud Metro router
- Competitor A (this is a current generation router from a leading vendor)
- Competitor B (this is an older generation router from a leading vendor with a large global installed base)

Our results show that the Juniper ACX7509 has a TCO benefit of 53% over Competitor A and 71% over Competitor B. Much of this benefit is due to reduced power, cooling, and space. The reduced power consumption of the Juniper ACX7509 in a network of 10,000 nodes results in a CO_2 emissions reduction of 69,765 metric tons compared to Competitor A, and 145,063 metric tons of CO_2 emissions compared to Competitor B.

These savings become larger as the metro network grows. The environmental efficiency of Juniper's ACX7509 helps reduce both TCO and greenhouse gas emissions. A summary of power and cooling, floorspace, CO_2, and total OpEx savings is presented in Table 1.

Table 1 Summary of ACX Savings Over Competitor A and Competitor B

ACX7509 Savings	Competitor A	Competitor B
Power & Cooling	61%	77%
Floorspace	29%	64%
CO_2	61%	77%
Total OpEx Savings (including labor)	53%	71%

Juniper's Cloud Metro Value Proposition

Juniper's Cloud Metro networks are designed to scale modern 5G, edge, and multicloud services while providing a high-availability architecture. The key components are:

- Sustainable high-performance metro networking systems
- The ACX7000 family of routers
- Cloud-delivered Automation as a Service (AaaS)
- AI-Ops to improve network operations
- Embedded active service assurance
- Built-in zero-trust security
- Converged IP services fabric

The ACX7000 family routers deliver a variety of L2 and L3 services at the metro edge of the network:

- L3VPN
- L2VPN
- BNG
- MPLS
- Segment Routing

These routers have leveraged the latest generation chipsets and system architecture design, resulting in lower power consumption and less rack space.

This paper reviews the ACX7509, part of the Juniper ACX7000 family, for its TCO and environmental benefit capabilities (as shown in Figure 1):

- Centralized architecture with orthogonal design and no backplane
- No fabric cards because FEB cards provide both forwarding and fabric capabilities
- Common ports for 1-50GE (SFP)
- Common ports for 10-400GE (QSFP)
- Embeds Broadcom Jericho2c today, ready for Jericho3 for 800GE and beyond
- A smaller 5RU chassis reduces the number of fans required

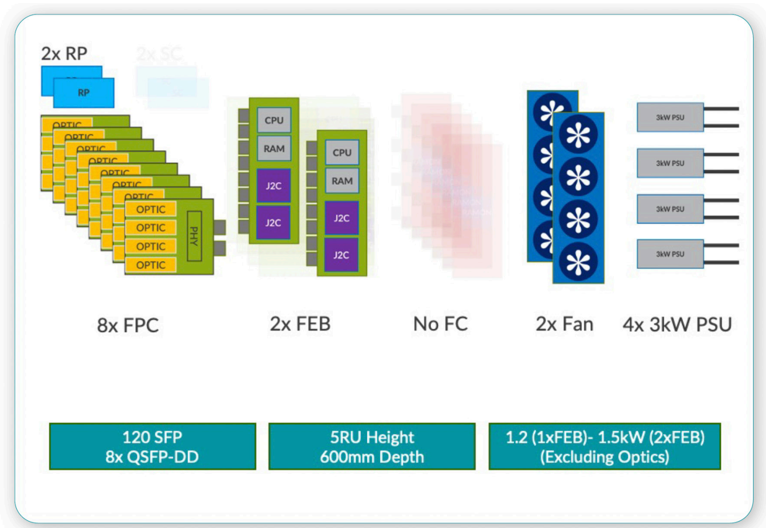

Figure 1 ACX7509 System Architecture

These features lower TCO for the following reasons:

- ACX7509 supports future generation chipsets, which extend system lifetime from 3–5 years to 7–12 years. Extended system lifetime reduces the TCO because forklift upgrades are not required to upgrade router capacity and features with the next-generation chipsets.
- Eliminating fabric cards reduces both power and space.
- Getting rid of fabric cards also improves life-cycle management because FEB cards can be upgraded without also needing to upgrade fabrics or forklift chassis.
- Eliminating fabric cards decreases the number of components that can fail and therefore improves MTBF and increases service availability.
- Common ports for both 1-50GE and 10-400GE allow service instantiation and changes to be done remotely, reducing truck rolls and manually changing hardware.
- Lower power and smaller chassis cut down power, cooling, and floorspace expenses.
- High-density interfaces and lower cost per port reduces TCO.

Paragon Automation as a Service

The goal of Paragon Automation as a Service is to provide a cloud-based service to operators providing network automation. Automation as a Service is a consumption-based SaaS service. Today, most operators are implementing multiple types of network automation, but there are challenges:

- CSPs must create and maintain scripts (Ansible, etc.).
- Automation requires knowledgeable staff to implement and maintain automation software and systems.
- There is a gap in the marketplace for skilled network/automation staff; this was further increased by the great resignation.

Juniper's cloud-delivered Paragon Automation as a Service provides:

- On-boarding new equipment
- Testing new equipment
- Service activation
- Ensuring security

Device onboarding is typically performed semi-automatically with few security and assurance checks. It is a time consuming and error-prone process that requires significant manual effort. In many cases, field technicians must have CLI knowledge, technical documentation, and experience testing network connectivity. The lack of automated oversight translates into costly errors and unacceptable times to market. To ensure secure, fast, error-free deployment at scale, communication service providers need to reimagine their device onboarding process with automation.

Paragon Automation as a Service allows field engineering to perform onboarding quickly, easily, and accurately through its mobile devices. In minutes, hardware and software authenticity is validated, latest software is imaged, secure zero-touch configuration and provisioning is completed, additional device health checks and network performance tests are done before the inventory is updated, resulting in devices that are fully ready for service. This is more than ZTP. Traditional ZTP implementations are limited to automatic configuration while Paragon Automation as a Service provides secure ZTP, device trust validation, device health checks, and network connectivity and performance testing. These features are automated and do not require manual intervention.

The key benefits of Paragon Automation as a Service are:

- Accelerate time to revenue at global enterprise and CSP scale with instant device onboarding.
- Ensure network trust with device integrity, compliance, and health checks.
- Guarantee device performance and service quality.
- Provide error-free deployment, avoiding costly mistakes by getting it right the first time.
- Reduce the skill-sets required by technicians.
- Decrease the labor expense required to install network devices.

Given that there are tens of thousands of routers in a metro network these savings can be significant.

AIOps

One of the main benefits of a SaaS-based AIOps solution is that you can better train algorithms by using anonymized data from multiple situations. The benefit is a better AIOps solution for all.

Predictive maintenance is another benefit of AIOps. Predictive maintenance is a technique that uses data analysis tools and techniques to detect anomalies in operations and defects in equipment and processes so they can be fixed before they result in failure. Ideally, predictive maintenance allows the maintenance frequency to be as low as possible to prevent unplanned reactive maintenance without incurring costs associated with doing too much preventive maintenance.

The Juniper Paragon Automation AIOps solution is based on both device telemetry and active service assurance. Service assurance and customers' experiences are optimized by actively monitoring quality and not just focusing on device telemetry like most other solutions in the market today.

The key benefits of the Juniper Paragon Automation AIOps solution are:

- Change network management from reactive ops to proactive ops.
- Reduce the required skill levels of engineers and technicians managing the network.
- Shorten training times.
- Decrease mean time to repair problems.
- Improve network availability and performance.
- Maintain customer service level agreements.

Reducing labor expense while improving network availability and performance is the main TCO benefit of the Juniper Paragon Automation AIOps solution.

Embedded Active Service Assurance

High-quality service assurance is a key success factor in every network. Although effective service assurance depends on effective fault management and AIOps, it is also important to actively monitor network performance to find problems before users or systems can detect them. Most service assurance monitoring uses either:

- Passive traffic monitoring
- Active probes inserted in the network to generate and monitor traffic and detect performance problems

Juniper's Paragon Active Assurance is embedded in the ACX7000 routers; an active probe is not required. It is a programmable, active test and monitoring solution for physical, hybrid, and virtual networks. Unlike passive monitoring approaches, it uses active, synthetic traffic to verify application and service performance. Service monitoring is delivered throughout the life of the service. Active Layer 2–7 service testing verifies that services are configured correctly the first time and ensures that service changes do not impact service quality. It provides detailed reports and alarms to alert operations of network performance problems. Because Paragon Active Assurance is embedded in the ACX7000 routers, there is no need for additional servers, probes or virtual machines (VMs) to install test agents. Where Juniper ACX7000 routers are not present, software test agents can still be deployed as VMs or as container applications or on bare-metal x86 hardware to meet multivendor testing needs. This reduces CapEx and cuts the cost of integrating and deploying agents. High-quality service assurance is critical because it improves customers' satisfaction and reduces churn.

Built-In Zero-Trust Security

Cloud Metro networks can have tens of thousands of routers deployed. There is a high risk of routers being compromised and the results can be catastrophic. The consequences of security breaches are high. Zero Trust, the security principle of "never trust by default, always verify" has become a best practice across industries.

A recent report from Microsoft, https://www.microsoft.com/security/blog/2022/01/12/microsoft-zero-trust-solutions-deliver-92-percent-return-on-investment-says-new-forrester- study/, has quantified the economic impact of a zero-trust solution. Highlights from the report include:

- Three-year, 92% return on investment with a payback period of fewer than six months
- 50% lower chance of a data breach
- Numerous efficiency gains of 50% or higher across security processes

The cost of security breaches is significant and implementing zero-trust security is critical to reducing the risk of network security breaches. All Juniper ACX7000 routers have built-in zero-trust security. DevID with TPM 2.0 hardware and software attestation validates the identity, authenticity, and integrity of each device. This is especially important in a Cloud Metro network that can have tens of thousands of devices deployed in unsafe locations such as street cabinets. It reduces the risk of counterfeit products or routers without proper software releases being deployed. In contrast, without these security capabilities, routers can be compromised and used to launch DDoS attacks as botnets. In addition to device security, it is also important to ensure data security, for example, protecting data-at-rest with native file encryption and data-in-transit with MACsec.

Converged IP Services Fabric

Cloud Metro reimagines today's siloed, point-to-point metro networks as a versatile IP services fabric that enables "Any Service, Any Place, Any Device" connectivity for distributed edge clouds and applications. It offers the ability to intelligently steer traffic not just to central data centers, but across multiple hubs (Edge Cloud), vaults and caches within the metro domain. These capabilities enable a more intelligent and future-ready metro with improved latency and bandwidth efficiency.

TCO Model Framework and Assumptions

The focus of this TCO model is on the OpEx benefits of the ACX7509 router. The objectives of the OpEx model:

- Compare power and space expense of the ACX7509 with two other industry-leading routers and show the OpEx benefits in a large Cloud Metro network.
- In the same network show the OpEx benefits of AIOps as compared to a similar network without AIOps.

In this model three routers are compared:

- ACX7509
- Competitor A: this is a current generation router similar to the ACX7509 from a leading vendor
- Competitor B: this is an older generation router from a leading vendor with a large global installed base

Table 2 Power and Space Requirements for Each Router

Router	KWatts	RU	Monthly Power Expense	Monthly Cooling Expense	Monthly Space Expense
ACX 7509	1,167	5	118	47	90
Competitor A	3,008	7	303	121	126
Competitor B	4,995	14	503	201	252
ACX Savings vs Competitor A	61%	29%	61%	61%	29%
ACX Savings vs Competitor B	77%	64%	77%	77%	64%

ACG Research used its Business Analytics Engine (BAE)(https://www.acgbae.com/) to model and compare the OpEx of the ACX7509 with Competitor A and Competitor B routers. The BAE is a visual, cloud-based economic simulation engine that calculates TCO and return on investment for many IT and network use cases. Figure 2 presents the high-level input to the BAE. In this analysis we assume a large Cloud Metro network that starts with 2,000 edge service routers and grows to 10,000 routers over five years. We also consider the following categories of labor:

- Change management
- Hardware replacement
- Help desk
- Fault management
- Performance management
- Software upgrades

Table 3 lists these categories of labor and the savings. Most of the savings are due to AIOps network management automation; however, the hardware replacement savings are due to the ACX7509 architecture that has fewer physical components (no fabric cards) and flexible ports that can be configured by software. This leads to reduced truck rolls and hardware replacement costs.

Table 3 *Categories of Labor and ACX 7509 Labor Savings*

FTE Name	ACX 7509 Savings	Notes
Change Management with AIOps	10%	Due to AIOps
Hardware Replacement with AIOps	20%	Due to reduced truck rolls result of HW architecture
Help Desk Trouble Tickets with AIOps	60%	Due to AIOps
NOC Fault Management with AIOps	70%	Due to AIOps
Performance Management with AIOps	70%	Due to AIOps
Software Upgrades with AIOps	10%	Due to AIOps

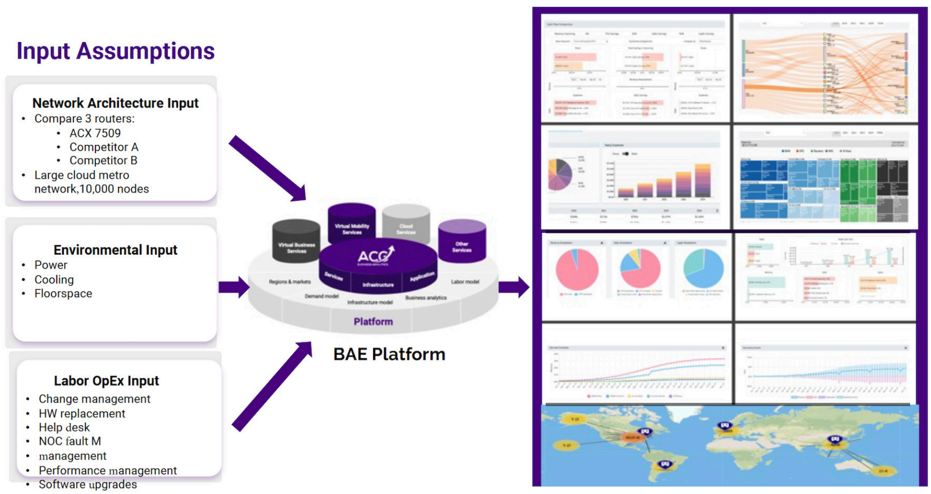

Figure 2 *BAE Input Assumptions and Cloud-Based Economic Simulation*

TCO Results

In this TCO analysis we considered OpEx for power, cooling, floor space, and network management labor expenses. We also examined the environmental benefits of lowering power consumption to reduce CO2 emissions. The results show that the ACX7509 significantly reduces both OpEx and CO2 emissions compared to the competitive products. In a large metro network, the reduction in environmental expenses associated with power, cooling, and floor space are significant. Table 4 presents the cumulative five-year environmental OpEx for each alternative and shows the ACX7509 environmental expense savings.

Table 4 Five-Year Cumulative Environmental OpEx With ACX Savings

Five-Year Cumulative Environmental OpEx		ACX 7509 Savings
ACX7509	$90.9M	N/A
Competitor A	$196M	54%
Competitor B	$341M	73%

In addition to the reduction in power, cooling, and floor space expenses, AIOps reduces network operation labor expenses. Table 5 shows the five-year cumulative OpEx and the ACX 7509 savings. The savings are 53% when compared to a current generation router from a leading vendor and 71% when compared to a legacy platform from a leading vendor with a very large installed base. The benefits of the ACX 7509 architecture with reduced power consumption and lower footprint combined with the AIOps benefits are the drivers of these OpEx savings.

Table 5 Five-Year Cumulative OpEx and ACX 7509 Savings

Five-Year Cumulative OpEx		ACX 7509 Savings
ACX7509	$107M	N/A
Competitor A	$229M	53%
Competitor B	$347M	71%

We compared the five-year annual OpEx spend for the ACX 7509, Competitor A, and Competitor B (see Figure 3). The OpEx grows as the network increases from 2,000 routers to 10,000 routers. The key point is that as the network expands, the difference in OpEx expenses between the ACX7509 and the competitors becomes greater.

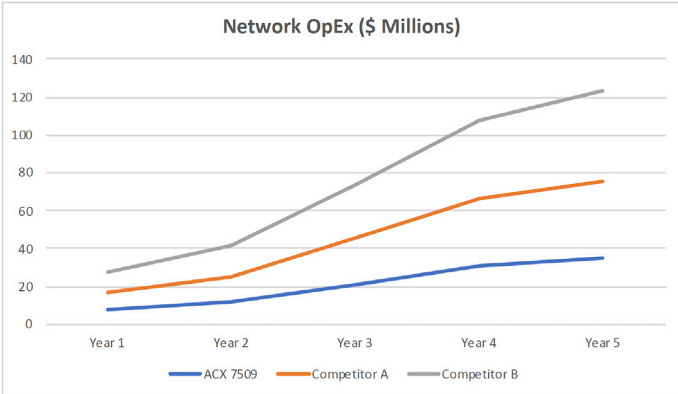

Figure 3 *Five-Year Annual OpEx Spend for Each Alternative*

Next, Figure 4 shows the five-year cumulative OpEx savings comparing the ACX7509 with Competitor A, and a similar OpEx breakdown for Competitor B is presented in Figure 5. The diagrams show the savings both in absolute dollars and as a percentage. In both cases power, cooling, and floor space account for most of the savings because we considered a large and growing metro network where environmental expenses are significant.

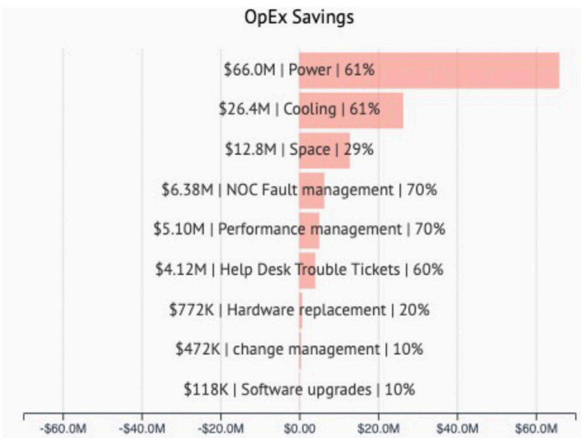

Figure 4 *Five-Year Cumulative Breakdown of OpEx for Competitor A and ACX 7509*

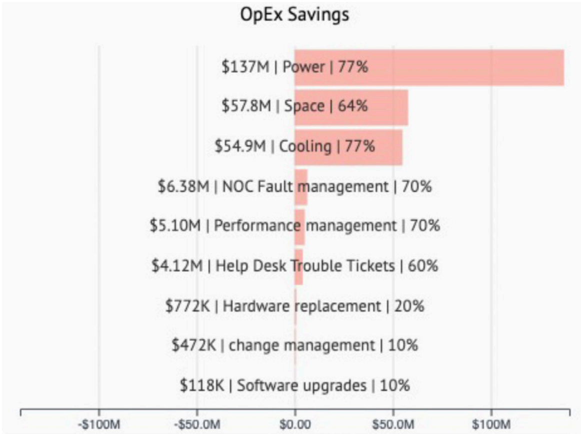

Figure 5 *Five-Year Cumulative Breakdown of OpEx for Competitor B and ACX 7509*

In addition to reducing OpEx, the decreased power consumption of the ACX7509 lowers CO2 emissions. Table 6 shows the total power consumption, CO2 emissions, and the CO2 savings compared to Competitor A and Competitor B. These savings become greater as the metro network grows.

Table 6 *KWatts and CO2 Emissions*

Router	Annual KWatts/Hours	CO² Emissions Metric Tons	CO² Savings Metric Tons
ACX 7509	102,229,299	44,224	N/A
Competitor A	263,500,800	113,989	10.7
Competitor B	437,562,000	189,287	22.2
ACX Savings vs Competitor A	61%	61%	N/A
ACX Savings vs Competitor B	77%	77%	

The results of the TCO model show that the ACX7509 has significant OpEx benefits and CO2 emissions savings over competitive platforms. The dollar value of these benefits becomes larger as the edge network grows. The environmental benefits also increase with the size of the network.

One calculating factor not considered is to take the amount of savings in Table 6 and apply it towards renewable sources of power such as solar, wind, or other. At some point the renewable power can both help pay for existing power OpEx while lowering the need to purchase additional carbon credits.

Conclusion

This paper has provided an overview of Juniper's Cloud Metro solution and its value proposition. Juniper's Cloud Metro is a holistic solution providing:

- Sustainable high-performance systems
- Cloud delivered Automation as a Service
- AIOps to improve network operations
- Embedded active service assurance
- Zero-trust security
- Converged IP service fabric

The ACX 7409 provides a pathway to energy sustainability via its high-performance, long-lived network infrastructure. This longevity on the job, monitored by AIOps, has the ability to outlast the elements and provide a return on the investment by lowering greenhouse gas emissions.

ACG Research developed a TCO model that showed significant TCO savings when compared with two competitive metro routers. The model showed significant reductions in CO2 emissions because of power decreases in the network. As edge computing and Cloud Metro networks continue to increase, minimizing both TCO and CO2 emissions is necessary to improving service profitability and decreasing the environmental impact of networks. This benefits people and the planet while providing connectivity in the decades ahead.

Speak to your Juniper Networks account manager or professional services rep about the new generation of routing systems designed for modern Cloud Metro networks.

Cloud Metro Resources

Here are some Juniper Cloud Metro resources with up-to-date specifications and information:

- Cloud Metro landing page: https://www.juniper.net/us/en/solutions/ip-transport-solution/metro.html
- ACX7509 Cloud Metro Router: https://www.juniper.net/us/en/products/routers/acx-series/acx7509-cloud-metro-router.html
- ACX7509 Cloud Metro Router datasheet: https://www.juniper.net/us/en/products/routers/acx-series/acx7509-cloud-metro-router-datasheet.html
- Paragon Automation as a Service: Cloud-Delivered Network Automation: https://www.juniper.net/us/en/the-feed/topics/network-automation/paragon-automation-as-a-service-cloud-delivered-network-automation.html
- Making the Case for Cloud Metro: https://www.juniper.net/us/en/the-feed/topics/metro/juniper-ceo-rami-rahim-makes-the-case-for-cloud-metro.html
- ACX7000 Family datasheet: https://www.juniper.net/us/en/products/routers/acx-series/acx7000-family-of-cloud-metro-routers-datasheet.html
- Introducing Juniper's Expanded Cloud Metro solutions: https://www.juniper.net/us/en/the-feed/topics/metro/brendan-gibbs-on-introducing-junipers-expanded-cloud-metro-solutions.html.
- Cloud Metro: Reimagining Metro Networks for Sustainable Business Growth: https://blogs.juniper.net/en-us/service-provider-transformation/cloud-metro-reimagining-metro-networks-for-sustainable-business-growth
- Balancing Design Choices for Sustainable Growth: Juniper ACX7000 Family: https://blogs.juniper.net/en-us/service-provider-transformation/balancing-design-choices-for-sustainable-growth-juniper-acx7000-family
- Heavy Networking 639: Juniper Cloud Metro Boosts Metro Performance, Efficiency And Sustainability: https://packetpushers.net/podcast/heavy-networking-639-juniper-cloud-metro-boosts-metro-performance-efficiency-and-sustainability-sponsored

Paper 8

The Sustainable Benefits of AI-Driven Enterprise Networks

By Yedu Siddalingappa

Juniper has brought true innovation to the networking space with the world's first AI-driven full stack network. The Juniper Mist™ AI Platform makes networking predictable, reliable, and measurable with visibility into the user experience, proactive automation, and self-healing capabilities. Apart from the technical and operational benefits, the solution can reduce a network's carbon footprint by way of virtualizations, reduced on-site visits, and automated energy management. All this is achieved through the Juniper's state of the art architecture featuring microservices cloud, the AI-driven operational framework, integrated hardware platforms, and 100% API capabilities. This paper will illustrate some of those capabilities as an initial introduction to the sustainable potential of AI-driven networks.

Mist Cloud Architecture Minimizes Appliance Footprints

Juniper's Enterprise networking operations are managed via Mist's open, programmable, and elastic microservices cloud architecture. This delivers maximum scalability, performance, and DevOps agility with reduced equipment footprints at customer premises. The AIDE (AI-driven Enterprise) cloud architecture allows the inherent benefits of power savings to enterprises by eliminating the need to host and manage multiple server appliances. The traditional architecture of networking systems often involves hosting multiple on-prem systems to manage and operate like controllers, master controllers, NMS, assurance systems, and location appliances. With the cloud architecture, all these functions are virtualized and operate as microservices (see Figure 1). Through economy of scale, virtualization, and the elastic nature of the cloud to dynamically scale up and down, both horizontally and vertically, cloud-delivered services can provide sizable power savings for individual enterprises.

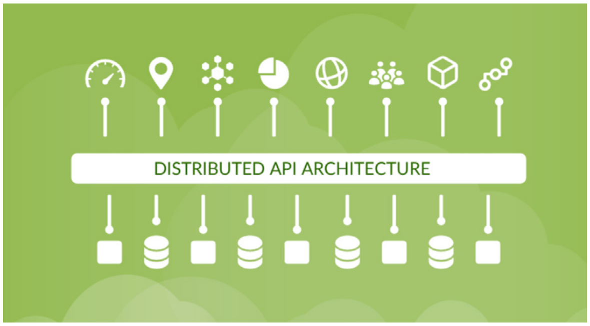

Figure 1 *Mist Microservices Cloud*

Integrated Hardware Reduces Overlay Equipment

Mist Access Point hardware is an integrated service delivery entity that combines Wi-Fi, RF security scanning, Bluetooth LE, and IoT so businesses can increase the value of their infrastructure deployments and reduce environmental impacts. Earlier, these services were available via SILO'ed systems and deployed in an overlay fashion with redundant cabling and power sources. The integrated hardware design of Mist APs means reducing everything from the number of pieces of equipment to cabling, mounting, switch ports, and power consumption without compromising the service quality or performance.

Figure 2 *Mist Access Points With Integrated Wi-Fi, BLE, and IOT Sensors*

Mist AI Minimizes Inter-Site Travels for IT Operation

Per industry research, IT teams managing network infrastructure often spend more than 40% of their time troubleshooting issues, the majority of which is spent on issue replication and data gathering. For enterprises having a distributed presence, troubleshooting serious issues in remote branches often involves on-site visits for effective problem replication and data gathering through debugs and packet captures. If escalations are bad enough, senior IT experts can be forced to travel thousands of miles, incurring additional environmental impacts.

The Mist Cloud uses AI and data science to analyze large amounts of rich metadata collected from Mist Access Points, Juniper switches, and gateway devices. This metadata provides actionable insights, proactive automation, and self-healing capabilities. For example:

- Supervised machine learning correlates events for rapid root cause identification.
- Time-series anomaly detection identifies negative trends and determines the magnitude of their impact.
- AI-driven Radio Resource Management (RRM) optimizes the RF settings in real-time based on changing conditions.
- Natural Language Processing (NLP) is used for making complex queries simple and fast.
- Unsupervised machine learning is used with Mist's vBLE technology to accurately locate users and devices.

Mist AIOps with its dynamic packet capture, seven-day historical record of network states, and data science tools to instantaneously analyze anomalies and identify the root causes, has reduced site visits by IT teams by a factor of 90% for several large, distributed enterprises.

Figures 3-6 show a few scenarios where Mist AI was able to proactively identify the root cause of a user issue and store the historic network state information, including the packet captures. In Figure 3, Marvis, the Mist Virtual Network Assistant, is shown identifying the impacting issues for the full stack network in near real time.

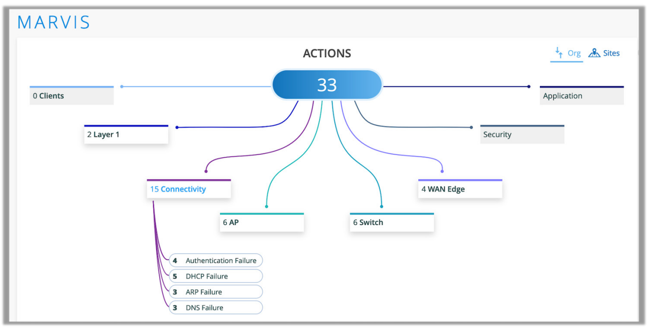

Figure 3 *Mist Marvis Virtual Network Assistant*

There are several scenarios where Mist AI can analyze the client issue and present the root cause and packet captures for the failure events, which previously required lengthy troubleshooting cycles and site visits. Figure 4 is a DHCP timeout issue: the client sends a broadcast discover packet but does not receive an offer packet from the server.

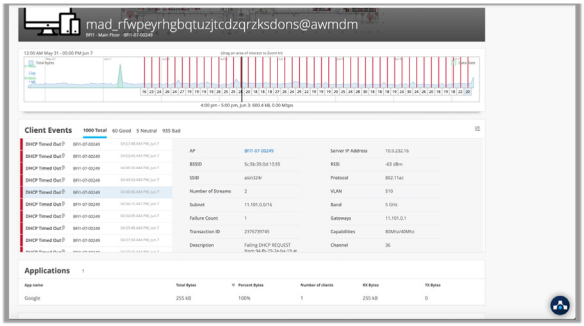

Figure 4 Client DHCP Events With Packet Captures

Figure 5 displays an authorization failure. This could be caused by such various reasons as MIC failure, the Radius server not responding, access reject from the Radius server, the client failing to complete the auth process, etc.

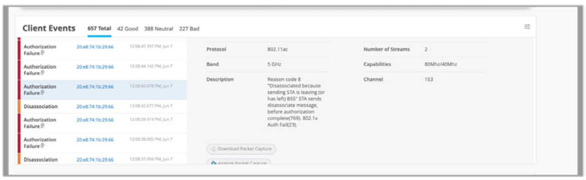

Figure 5 Client Association Failure Events With Packet Captures

And Figure 6 shows an 11r FBT failure, likely caused due to client failing 11r roam.

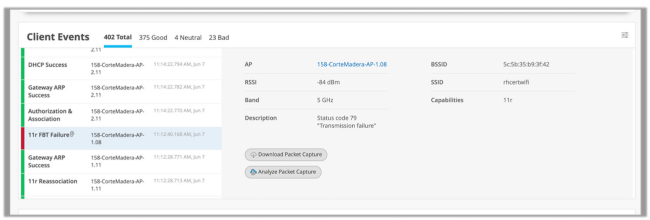

Figure 6 Client Roaming Failure Events With Packet Captures

Automation Eliminates Multi-Hop Shipments and Multi Packaging

The Juniper Mist platform also has automation capabilities for Day0, Day1, and Day2 operations that can save time, travel, and money.

The platform supports zero touch provisioning (ZTP) for the full-stack products of wireless access points, switches, and WAN gateways. ZTP eliminates the two-step shipping process, wherein products had to be shipped first to a staging location to perform initial configuration before being sent to their final destination for deployment. With ZTP and cloud-driven services, products are no longer required to go through the staging process. They can be shipped directly from the vendor to the destination site, where they can be deployed out of the box. The ZTP process reduces carbon footprints and eliminates dual packaging and all the associated wastages.

ZTP and Day0 automation features even eliminate the need for IT personnel presence for the initial equipment setup on-site, thereby reducing travels even further. With the Mist AI mobile app, the technicians involved in cabling and electrical work can onboard the network devices, too. Every cloud-enabled Juniper AIDE product includes a QR code label on the packaging and on the appliance body. The workflow involves scanning the QR code on the appliances and choosing a site for the devices to onboard. The app can even help offload the tasks like replacing or relocating the devices to the installers. See Figure 7 and Figure 8.

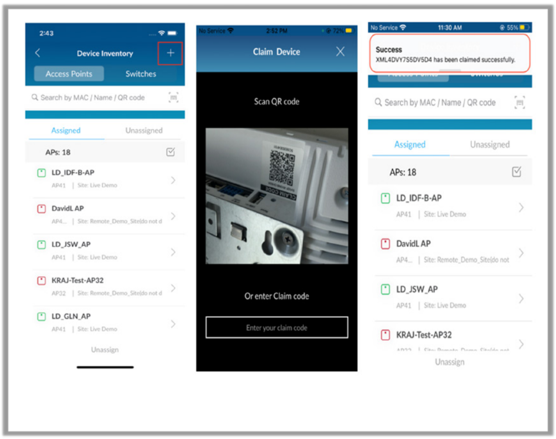

Figure 7 Mist AI App Performing Zero Touch Provisioning of an AP

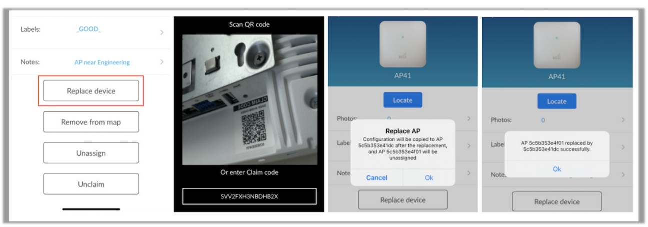

Figure 8 Mist AI App Performing Replace AP Operation

Juniper Analytics Drives Green Benefits

HVAC systems are one of the prime contributors to higher greenhouse gases and their usage in office spaces largely remain unoptimized. In the post-pandemic world, most enterprises are operating in a hybrid work environment with less than half of typical office occupancy. Still, HVAC usage and expenses remain consistently on the higher side because their operations are optimized for the space rather than the occupancy. Juniper AIDE's analytics can help in such situations by gathering the employee presence and occupancy details and feeding that information to HVAC systems to optimize by turning temperatures up or down.

In cases where HVAC systems are old and siloed, Juniper has forged partnerships with specialized technology vendors to offer a joint solution. One such example is a solution by Juniper and KODE Labs.

Juniper wireless architecture converges Wi-Fi and virtual Bluetooth LE (vBLE) to enable high-accuracy indoor location services, allowing occupancy analytics, asset tracking, indoor navigation, and proximity notifications. Mist wireless access points can sense the Wi-Fi and BLE signals emitted by user devices to compute device location and dwell times. Figure 9 shows a dashboard calculation of office space occupancies in real-time.

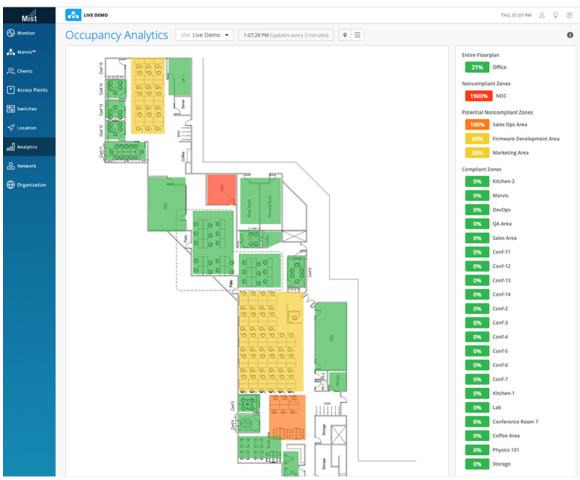

Figure 9 *Mist Dashboard With Space Occupancy Insights of an Office Floor*

The platform's 100% open API support helps to easily export this location intelligence for consumption by other systems, such as KODE Labs.

KODE Labs platform is a data focused, vendor-agnostic, building IoT platform integrating building systems like HVAC, lighting, and fire sensing into a single dashboard. In the joint solution with Juniper, Mist provides data on occupancy, asset tracking, temperatures, and humidity to the KODE labs platform, which in turn leverages the intelligence to adjust building systems to provide an optimal experience for employees, reducing OPEX, and greenhouse emissions in the process as shown in Figure 10.

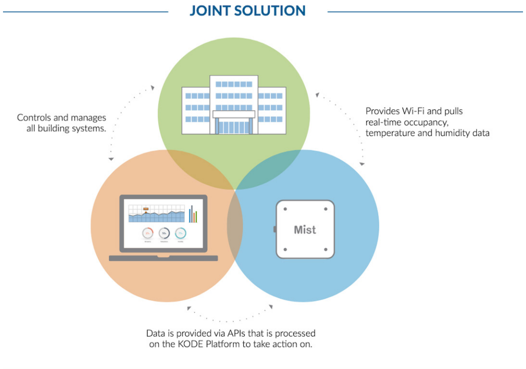

Figure 10 Mist Kobe Joint Solution for Building HVAC Management

For more information visit https://www.mist.com/wp-content/uploads/Mist_Kode_solution-Brief.pdf.

Programmable Network Reduces Energy Consumption

Juniper Mist cloud services are 100% programmable, with all functions (provisioning, monitoring, alerts) available through open APIs. This capability enables enterprises to automate several use cases, including energy-optimizing ones.

In a typical enterprise, the Wi-Fi networks are deployed in a dense mode to handle high-density client environments. Further, all the access points are configured with complete feature sets and operate 24x7, irrespective of the employee presence in the office space and the actual network usage. While this deployment method offers the best user experience during office hours, it will result in significant unwanted energy consumption during non-office hours, holidays, and weekends when the user presence is negligible to nil. Juniper's API based automation can effectively address this challenge by applying time-bound policies for the device operations.

For example, Juniper AP33 and AP32 Access Points require 802.3at power to operate with full functionality but can also use 802.3af power to operate at reduced speeds and minimal features. These APs are powered by POE from switch ports in most deployments. In an open system like Juniper AIDE, a simple automation script executed via an API interface can change the power delivered from the switch ports to the APs, from .3at to .3af, during non-working hours, thereby reducing power consumption on the Wi-Fi network by a significant amount. You can also apply similar policies to other systems like digital signage, collaboration systems, and kiosks to cut down energy consumptions.

The Future

As the world moves towards rapid digitalization, it will quickly see the additions of millions of mobile devices, users, and applications. Not only will this put enormous pressure on IT teams to deliver superior experiences with limited resources, but it will also continue to exert pressure on the environment due to manual and travel intensive IT operations.

Juniper's long-term vision and road map can help alleviate some of these concerns for enterprises. The idea is to incorporate the innovations of Mist AI and microservices cloud architecture into the entire technology stack from client to cloud. The operational and green benefits of cloud architecture, AI-driven insights, self-driving framework, and open architecture will thus extend to multiple layers of the world's IT systems. The reasoning is the same: simplify IT operations and you will either directly or indirectly contribute to a greener planet.

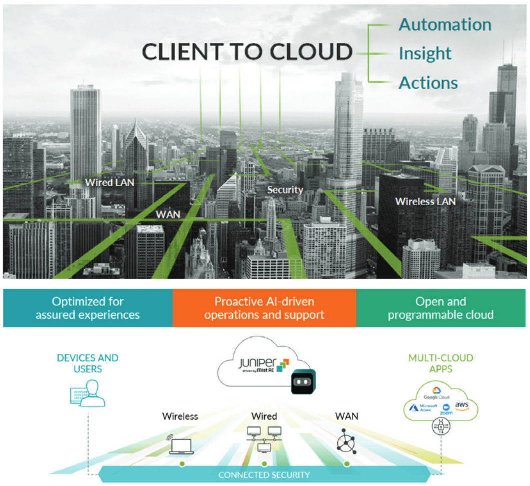

Figure 11 Juniper Mist Client to Cloud AI-Driven Vision

Speak to your Juniper Networks account manager or professional services rep about the future of Mist AI in your sustainable environment. For more information see https://www.mist.com.

Paper 9

Juniper ASIC Team Pioneers System-in-Package (SiP) ASICs

By Valery Kugel

Juniper Networks was founded in 1996 and from day one began designing novel networking ASICs for internet traffic. During the last quarter century, the company has designed more than seventy ASICs starting from the 0.35um CMOS silicon technology (complementary metal-oxide-semiconductor) all the way to the most advanced CMOS nodes available today. This brief paper looks at the future of Juniper ASIC design and the breathtaking abilities of the new System in Package (SiP) design that increases key performance while using less power.

The first Juniper ASICs had ~10 million transistors on a silicon die and a dozen 250Mbps high-speed IOs. The ASICs had a silicon die area around 100 - 150mm and were packaged using small 30 - 35mm on a side flip-chip organic substrate technologies, see Figure 1.

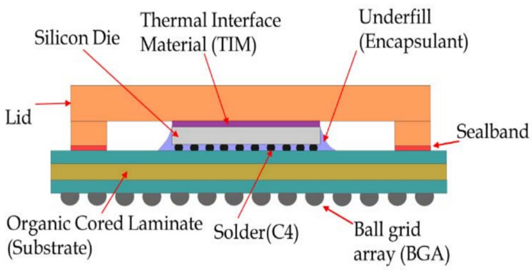

Figure 1 *A Typical Early Juniper ASIC in Flip-Chip Package Assembly*

With the advent of High Bandwidth Memory (HBM) in 2013, and then its second generation HBM2 in 2016, a new era of ASIC packing was introduced – integrating ASIC and HBMx on an ASIC package. It required a new packaging technology development to route very dense ASIC die-HBM interface signals. An example of such a technology is Taiwan Semiconductor Manufacturing Company's COWOS – Chip On Wafer On Substrate shown in Figure 2. The ASIC die-HBM interface signals are routed on a silicon interposer, which is soldered to the package organic substrate using lead-free bumps. Juniper introduced HBM and the new packaging technology in both its Trio and Express family.

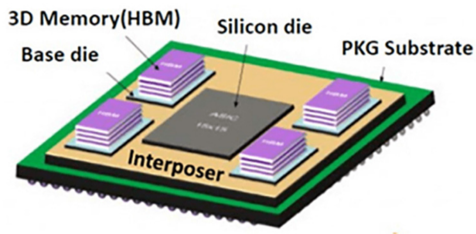

Figure 2 Integrating HBM and ASIC Dies on a Package Substrate Using Silicon Interposer

No doubt, the networking industry has experienced exponential increase in network bandwidth requirements during the last 25 years. The progress has translated to approximately doubling network equipment bandwidth every two years! While CMOS silicon technology was able to support bandwidth-hungry networking ASICs without significantly increasing the area and power of the ASICs for a long time, during the last ten years there's been a gradual slowdown in CMOS silicon technology scaling. As a result, to support the bandwidth scaling requirements, the networking ASIC die's area grew significantly, see Figure 3. Together with the die area increase the ASIC package approached 60 - 70mm on a side.

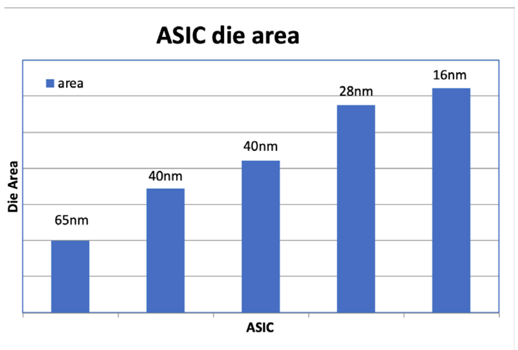

Figure 3 Area Trends for Some Representative Juniper Networking ASICs

And during the last five years we have observed a dramatic slowdown in CMOS technology scaling, so novel architectural approaches were required to keep scaling bandwidth of networking ASICs. Of prime importance has been the Juniper ASIC team's pioneering System-in-Package (SiP) design approach using novel packaging technologies.

In Figure 4, two industry-leading routing Express ASICs are depicted. The ASICs were discussed in detail by Dr. Chang-Hong Wu at the HotChips 2022 Conference (see *References*). Each 85mm x 85mm package has two networking ASIC dies (chiplets) and multiple HBMs; and each ASIC has two silicon interposers. The X-chiplet has 59 billion transistors and the F-chiplet has 35 billion transistors! The ASIC chiplets communicate with each other using low-power SerDes macros designed to the CEI-112G-XSR-PAM4 standard, which Juniper co-sponsored at the OIF (Optical Internetworking Forum). This novel approach allows Juniper to design very complex SiPs in a cost and energy efficient manner.

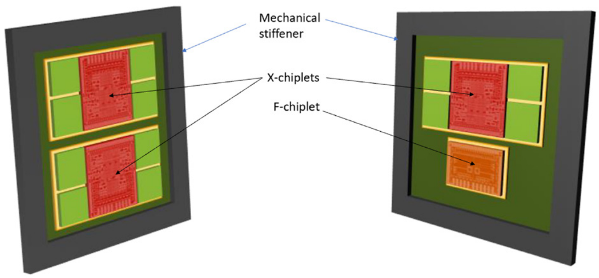

Figure 4 *28.8Tbps Routing ASIC Designed Using Chiplet Approach*

In conclusion, as seen in current trends in CMOS technology, the ability for Moore's law to deliver twice the networking bandwidth at the same cost and for the same die area and power is not there anymore. The Moore's law is de facto "stagnant." Therefore, the Juniper ASIC team has actively pursued novel architectures and technologies to continue delivering cutting edge networking ASICs with the lowest price and least power. We are also actively engaged with relevant standard organizations: IEEE, OIF, and more recently with UCIe consortium to make sure the standards incorporate requirements of the networking industry. Stay tuned for more exciting news from the Juniper ASIC team!

References

Chang-Hong Wu, "Juniper's Express 5: A 28.8Tbps Network Routing ASIC and Variations" https://hc34.hotchips.org/assets/program/conference/day2/Network%20and%20Switches/HC2022.Juniper.ChangHong_Wu.v03.pdf

Paper 10

Using the GHG Protocol Framework to Examine Technology's Role in Reducing Emissions

By Paddy Berry

Introduction

Every business and individual in the world needs to do more to reduce their impact on the environment either through more sustainable supply chains, ecological product choices, or any number of other means. Businesses are becoming more and more conscious and critical of their suppliers' statements and strategy when it comes to dealing with this global reality but are also seeking guidance on how to improve their own environmental agendas.

The challenge here is both quantification of the problem – how do we measure any of this if we haven't done so before – but also in the qualification of what constitutes direct impact versus an imported impact. In other words, what can be controlled by the party in question and what is obfuscated by the suppliers of goods and services.

There are however, standards and frameworks to hold ourselves and our suppliers accountable and put such goods and services under the lens of scrutiny.

One of these standards, which is now being commonly adopted by all major enterprises, is the GHG (greenhouse gas) protocol which defines emissions and energy consumption. Depending on how direct and indirect the impact of emissions might be they are separated into different *scopes*, as shown in Figure 1, from the *GHG Protocol Guidance Document*, page 26, October 17, 2022 (https://ghgprotocol.org/sites/default/files/standards/ghg-protocol-revised.pdf).

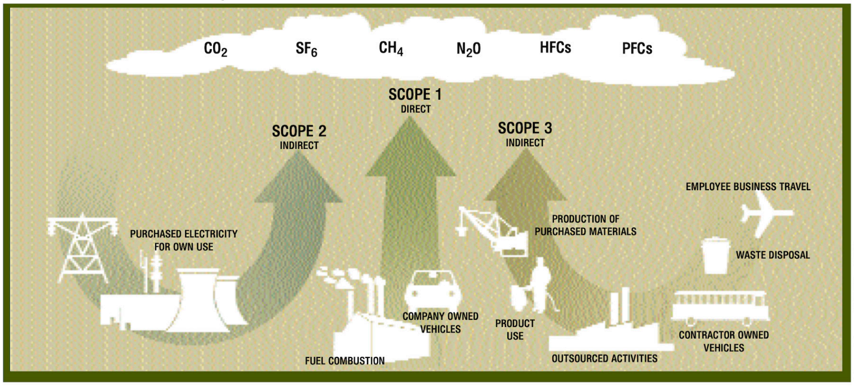

Figure 1 *Overview of Scopes and Emissions Across a Value Chain*

Scope 1 relates to direct combustion and consumption of fuels by company assets; such as company car fuel, gas in boilers in directly-owned facilities, or, in industrial businesses, running production systems such as furnaces, chemical plants, and others. All of these either generate CO2 or another gas which is recognized as a GHG pollutant under the Kyoto protocol.

Scope 2 relates directly to energy that is purchased by the business for use within their facilities and for which is relatively straight forward to account.

Scope 3 is the catch-all of anything that doesn't fall directly into the previous two scopes. It represents all the other emission related activities that a business might undertake, even those as indirect as employee commuting and products that are purchased from their suppliers.

NOTE This paper's intent is not to be a guide on how to interpret these sometimes elusive guidelines but rather aims to map the way our IT systems and processes can seek to reduce emissions both directly or indirectly through a technical lens. These are often my subjective opinions because I am passionate in our necessity and capability to do more and I hope they can be useful both in understanding the subject matter but also opening up and normalizing the conversation further.

MORE For further information on the GHG protocol standard and some of the terminology used within this paper, please refer to the GHG website for more details at https://ghgprotocol.org/.

What Can Juniper Do?

Juniper provides a wide variety of network equipment and software which, when combined together, can facilitate some of the environmental IT outcomes which businesses are striving to deliver.

I will outline some of the latest technologies that Juniper provides and demonstrate how they can be used to either reduce emissions defined within the scopes or at least act as an enabler for said change. To reiterate, this is a subjective interpretation of the framework above and others might come to different suggestions or conclusions; something which this paper advocates in the spirit of discourse.

NOTE Advocating throwing a solution away that is working perfectly well and within a serviceable life goes against the whole purpose of defining these scopes and process since e-waste and product waste in general is something that generates emissions and is detrimentally impactful in other ways.

Let's look at a few key Juniper Networks solutions that can facilitate some of these favorable environmental IT outcomes.

AIDE/Mist

Mist is the engine behind Juniper's AI-Driven Enterprise (AIDE) solution and provides both proactive and intelligent troubleshooting with unprecedented scale out functionality all whilst being delivered with the minimal amount of on-premises footprint.

Traditionally, trying to troubleshoot client issues, especially on a wireless medium, has been both operationally expensive in terms of the business being impacted but also the opportunity cost and travel expenses that are accrued because of lack of data points to isolate problems.

The AIDE aims to alleviate such concerns by providing administrators with both the data they need and the guidance to perform remedial actions without having to understand the physics or dive into a controller debug.

These are some of the intrinsic benefits of the Mist architecture as well as the enablers of the AIDE solution as a whole that could assist with reductions in the defined scopes discussed previously. For further information about the technology and an expansion on some of these points below, please refer to *Paper 8: The Sustainable Benefits of AI-Driven Enterprise Networks*.

Mist Microservices

Juniper Mist is built from the ground up and is unique in the market for exploiting a truly microservices-based architecture in the public cloud and this means minimum energy waste as compute resources are only used when required and requested. This is more efficient than legacy IaaS-based platforms or physical controllers that need to retain spare capacity for onboarding new tenants and services.

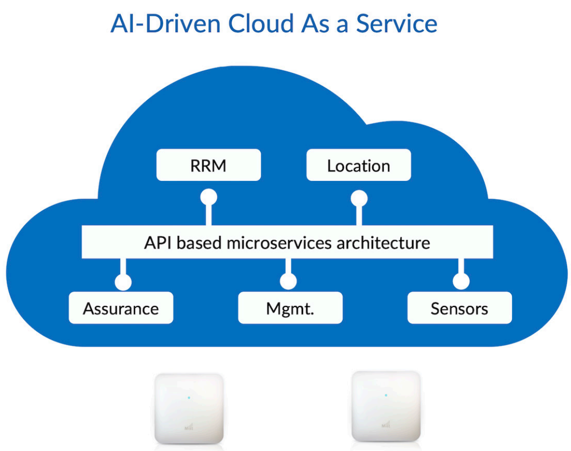

Figure 2 Each of the Mist Functions Is a Microservice in the Cloud

Furthermore, a microservices architecture brings advantages in terms of flexibility such as being able to add new services without having to think about controller software and hardware dependencies or scale considerations which come with older solutions. This reduces the need for events such as forklift upgrades which not only cause disruption during migration but require the hardware to be disposed of and new hardware to be purchased; all of which contribute to an emission and impact from both production of equipment and e-waste.

Although this might be attributed as an indirect saving, most green agendas are concerned with more than just scope 1 models (direct emissions), and this relates to either a scope 2 or 3 saving depending on how it is measured; either as an electricity-saving compared to physical infrastructures or a reduction in waste for operations.

Building Optimization Through Intelligent Location Services

Energy costs and the environmental impact of real estate including lights, heating/cooling, and IOT, has great potential for optimization and thus could have the biggest impact on a business' net CO_2 emissions.

The unique hardware design of the Juniper Mist APs with in-built hyper-accurate vBLE antenna array and software ecosystem is an enabler to wider building optimization through precise location services that can interface into a variety of smart building management software. This is relevant to current and future building planning and strategy by reducing utilization of traditionally always-on services and replacing it with point-in-time activation based on occupancy and activity.

We believe that Mist is the only infrastructure provider that provides the precision necessary, and the uptime required, for applications that enable use cases such as:

- Broader understanding of building use and occupancy for planning and consolidation
- Automated temperature management with built-in sensors
- Automated power management of the wider IOT and smart device estate based on BLE
- Occupancy-based lighting and power management automation

Juniper works in tandem with an ecosystem of vendors upstream in the software stack and provides a consolidation of sensor information that these platforms need. Many enterprises have building management applications and hardware platforms which could benefit from such consolidation.

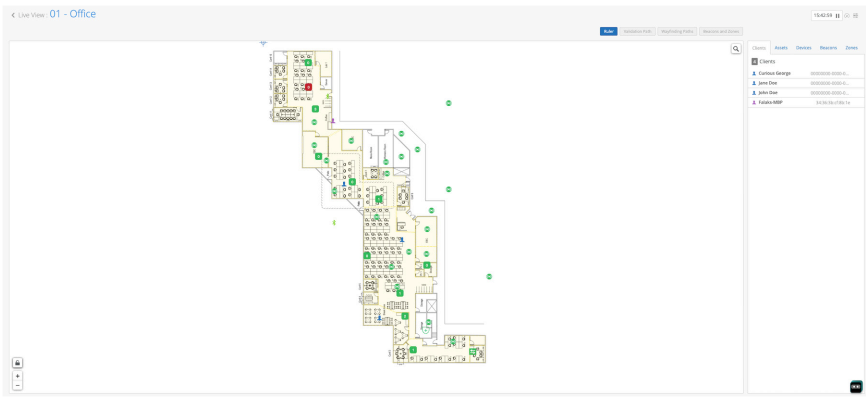

Figure 3 *An Example of Office Monitoring*

The consolidation is driven by reducing the demand for server appliances and removing the need for wasteful battery-driven beacon systems for legacy location services: which has been one of the most compelling factors for large enterprises both from a green and ROI perspective.

Whereas other vendors are trying to sell location services as a nice-to-have feature to show client count, Juniper believes that this should form the foundation of truly transformational smart building and real estate rationalization as well as a way to engage future clients and retain existing workforce.

All of these use cases can inspire savings within scope 1, 2, and 3; the scope of which will depend on the way that businesses power, heat, and cool their real estate. More importantly this is likely the area that can have the most direct effect but requires the largest amount of planning and cross collaboration from NetOps and facilities management.

Powerful Templating, More Consistency, Less Travel

The Mist platform with integrated AIOps can reduce GHG emissions through reducing the need for engineers to physically attend sites. Creating site templates, automated firmware upgrades, and integrated packet capture means that non-IT based employees can simply just plug new devices in or assist with the troubleshooting process without needing to be a subject matter expert (SME). Where further investigation is required, Mist provides the data points to the teams that require it to correlate events to the fault conditions and can coordinate local resources rather than needing to travel themselves.

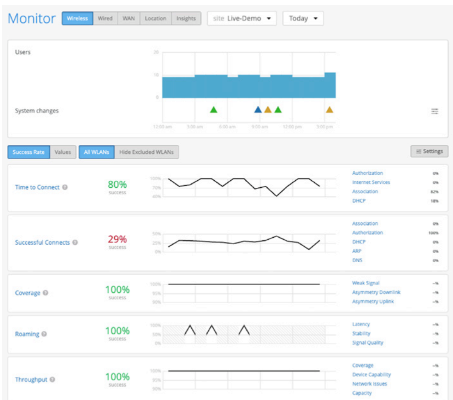

Figure 4 *Mist Service Level Expectations*

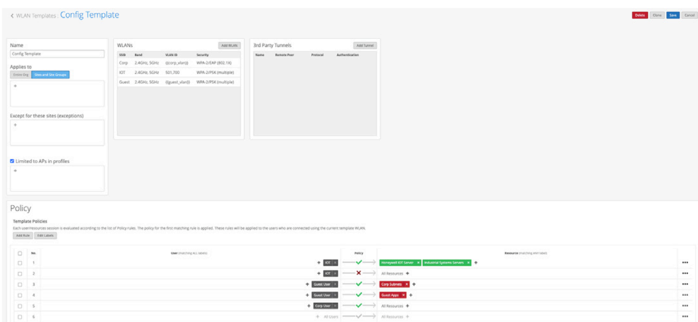

Figure 5 *Mist Configuration Template With Security Policy*

Depending on whether the transport assets are company-owned or leased, rather than employee-owned, might represent either a scope 1 or scope 3 reduction. However, since the activities are to do with company activities and employee transport may be used in lieu of available company transport, it might better represent a direct emission.

Cloud Metro 2.0

Juniper launched their new Cloud Metro portfolio in September 2022 with the purpose of providing security, performance, and efficiency enhancements to the previous generation.

However, it is becoming more common nowadays for service providers, especially those who are new to the marketplace in altnet or regional high-performance competitors, to require simple scale out as well as off the shelf products to manage and monitor their WAN estate. Whereas larger carriers have invested in OSS tools and processes for many decades, the ability to take a software product direct from the vendor, that can provide the majority of functions on Day1 is a significant benefit and allows these newer companies to remain competitive and agile against their larger competitors.

Paragon Automation as a Service (PAaaS)

Using the knowledge gained by the development of the Mist platform and the assurance capabilities therein, the Cloud Metro portfolio brings with it an evolution in the way network assurance and provisioning is delivered. Beforehand, the Paragon portfolio for WAN lifecycle was an on-premises offering comprising many different virtual machines or running physically. Not only does that require businesses to find a place to run these machines, which due to criticality would probably run on dedicated appliances or cluster, but the footprint is significant.

Evolving this offering makes it more accessible for businesses who want the capabilities delivered as a Service for an operational expenditure saving, and being cloud-operated should help reduce inefficiencies.

Paragon provides the following benefits to service providers that could provide some form of emissions savings:

Paragon Device Onboarding

As new provider networks scale, so does the need to install new hardware to connect consumers with the backbone. In the past, this may have required different skill sets of people; one team to design and plan, another to build out and cable the service, and lastly, one to configure and check. Requiring more than one set of hands to be physically present at the onboarding and installation of an asset is an obvious inefficiency that grows greater as the scale increases. With Device Onboarding, a standardised configuration can be pushed to the device from a cloud and clear instructions provided, regarding cabling and physical checks to be performed, to the on-ground teams. Those who need to perform post-installation checks can be remote within fewer operations centers and rely on the advanced telemetry streamed back to the cloud to provide teams on the ground further instruction if need be.

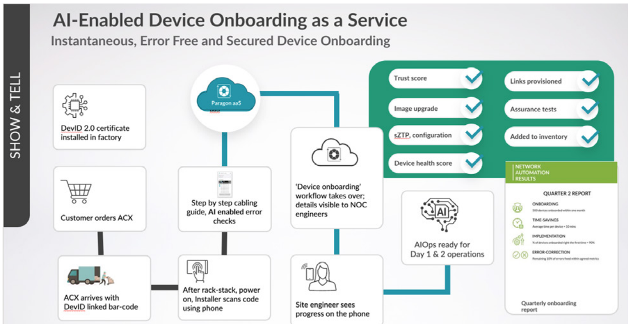

Figure 6 *AI-Enabled Device Onboarding as a Service*

This directly links to scope 3 in terms of contractors' travel emissions by right-sizing the skills and the resources where they are required rather than sending groups of people around the globe for a task that could be made more efficient.

Paragon Network Optimization

Using some of the tools within the Paragon portfolio, network administrators can plan how best to optimize their network to cater to the needs of the consumers rather than investing in higher-speed circuits even when speed may not overcome the challenges and inefficiencies that are present. Furthermore, the healthier our networks are in terms of forwarding productivity, the less time the traffic spends traversing slower paths and investing CPU/ASIC cycles in unnecessary additional processing overhead.

Without the data points being fed back to a central location and being able to visualize all of this information, the only tools that most businesses have to go on is monitoring utilisation of individual ports or appliances which is a poor indicator of efficiency or a functioning system. Being fully apprised of conditions across the network, planning ahead for demand, and only brining online what is required, doesn't just make good commercial sense but also ensures that the network hardware is using the energy it requires in the most efficient way.

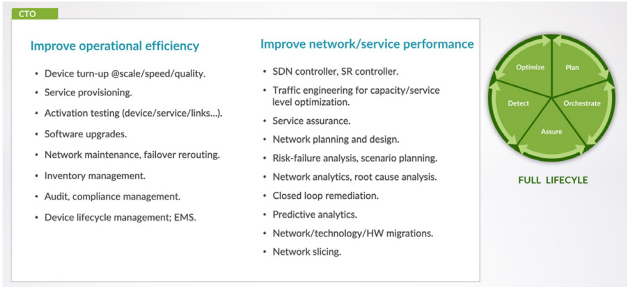

Figure 7 Paragon Automation Benefits

All of these benefits link back to both scope 2 and scope 3 savings. Scope 2 relates to ensuring that the energy that we do consume is effectively utilized, and scope 3 relates to making a more efficient product that is sold to customers as well as reducing the time to troubleshoot inefficiencies, which all add up to wasted emissions.

There is an argument that could be raised relating to whether the business in question has a responsibility for the emissions related to the electricity consumed within colocation facilities, since they are a third party and outside of the direct or indirect control of the business. However, we need to be mindful that we should be regularly questioning the efficiency of all partners that we engage with and who provide services to us, by building the true cost of consumption which reflects an environmental agenda: therefore, that emissions should play some role in the commercial argument when it comes to procurement.

Connected Security (CSEC)

Many customers are considering how best to secure their assets within the public and hybrid cloud and for that you need a broad selection of technology tools. As much as people might want to completely redesign their software from the ground up as a disaggregated microservices block accessed using API gateways and other more modern security techniques, the fact is that IaaS infrastructure and diverse security controls are still necessary.

Figure 8 *SRX Effectiveness*

As an example, Juniper vSRX and CWP products provide a familiar look and feel of a high-performance firewall with a more modern dynamic application code recognition and remediation suite. Furthermore, the vSRX has been recognized as the most efficient and effective firewall for use within the public clouds from a software perspective as it makes the best use of the underlying cloud computing power that it runs on compared to its competition.

Any customer that is familiar with public cloud billing will know that right sizing and ensuring the most efficient underlying infrastructure for the performance it offers is key to a good cloud strategy; otherwise, costs tend to spiral out of control.

Now, advocating wholesale change from one piece of software to another, without design review and planning, can introduce short term complexity and inefficiency which will lead to more indirect emissions. So when considering new solutions or planning this move, you want to ensure that the advertised figures as promoted by the vendor deliver in the most efficient manner based on the cost paid and the emissions generated.

	Juniper Networks	Fortinet
Security Effectiveness	100% Block rate, Zero False positives	100% Block rate, Zero False positives
Rated Throughput	974 Mbps	946 Mbps
TLS Throughput	948 Mbps	892 Mbps
SSL/TLS	13/13 Ciphers & use cases	12/13 Ciphers & use cases
TCO	$56.66 per Protected Mbps	$66.06 per Protected Mbps

Figure 9 *CyberRatings Efficiency Sample*

Returning to our GHG scopes, we can see that this clearly falls within scope 3, as anything run and purchased from a third party doesn't represent imported emissions through electricity or direct emissions. However, like the previous sections, we have to be mindful of outsourcing what was originally classed as an easily attributed indirect cost to the public clouds. We should still consider ourselves responsible for being efficient with the resources we contract through good control but also ensuring that whatever is running is right-sized and does what is advertised.

SSR (Session Smart Router)

SD-WAN originally entered the market with a clear goal in mind and that was around cost efficiency. Within all markets, MPLS was used as the defacto standard that securely interconnected offices with their hosted applications over a L2/L3VPN after replacing legacy technologies such as Frame Relay or ATM. Regardless of whether the applications were hosted via managed service providers, self-hosted, or collocated, they represented a simple way for businesses to connect what they needed without having to purchase dedicated circuits and manage their own WAN routing protocols.

However, with simplicity came cost and therefore the first SD-WAN products aimed to create the same secure experience over commodity circuits using IPSec VPNs between hubs and spokes. This certainly reduced the costs in the short term as commodity circuits were inherently cheaper at the time of launch but over time, especially in the EMEA region, MPLS was rarely significantly more expensive and therefore the cost argument started to wane. Vendors refocused on visibility to hold providers to account, however this only served to place the burden on the consumers to show back issues to the service providers.

Alongside this was the issue of performance for certain types of applications. For those more modern applications using either HTTP/HTTPS client-server models or those that were directly web accessible, there was little issue using circuits without a defined latency/jitter commitment as the protocol is inherently "tolerant" to such conditions. However, traditional voice applications and those applications which retained protocols which were both intolerant to latency or jitter, faced significant onboarding challenges. SD-WAN vendors weren't up to speed with looking at the user experience or requirements gathering sufficiently in the early design phases to account for these applications and their quirks.

Therefore, without significant application migration efforts or re-engineering, something that small-medium businesses were incapable of doing themselves and smaller application providers were not prepared for, MPLS was required and retained alongside some commodity service for general use. The nirvana of cost efficiency for removing MPLS and also performance and simplicity wasn't realized either.

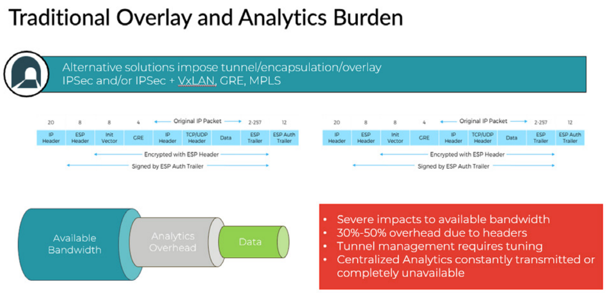

Figure 10 *Traditional SD-WAN Inefficiencies*

At the same time, a company called 128 Technologies was developing a new protocol and standards for the WAN which could ensure path and session quality for voice applications over commodity circuits. Instead of using IPSec to secure between WAN endpoints which added even more latency and overhead to an already best effort service, 128T created session vector routing (SVR) which ensures a secure path across the WAN before transmission as well as calculating backup paths without the need for IPSec end-to-end. Using these mechanisms meant for a significantly improved performance but also dramatically decreased the overhead and therefore the wasted bandwidth as a result of IPSec.

Furthermore, aside from tunnel overhead, the software can be configured to only encrypt those packets which are not encrypted as standard; there is no benefit to encrypting traffic which is already encrypted and represents even more overhead.

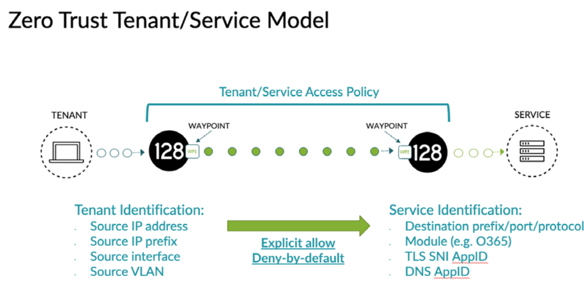

Figure 11 *Per-Service, Per Waypoint Secured Transport*

Juniper acquired 128T in October 2020 and it has since rebranded the platform as the *Session Smart Router* (SSR). Combining the no-overhead approach to SD-WAN alongside the MIST platform’s rich data analytics and telemetry means that regardless of the application you are running or the circuit types that are used, you can be assured that the end-to-end experience is ensured on a per session basis and no other solution can claim to do the same.

But how is this important from an environmental standpoint? Well, it relates quite closely to the previous example but also to more traditional understandings of efficiency.

From a scope 3 perspective, we can quite clearly see that for the same performance across a given WAN there is a reduction in emissions both in terms of the overhead from a customer perspective but also throughout the WAN backbone and carrier; the further the session travels the broader the savings.

On the other side, if we take the same principle and apply it to the cloud, we get a different sort of savings. If you had a public cloud environment, you are generally limited to the amount of bandwidth you can effectively use once secured using legacy technologies. Additionally, the business is directly charged for ingress and egress traffic which is naturally more with IPSec being used as an overhead. Therefore, here we can demonstrate more efficient use of the cloud onramp and therefore right sizing the circuit/service to reduce emissions through unnecessary waste but also a normal cost saving.

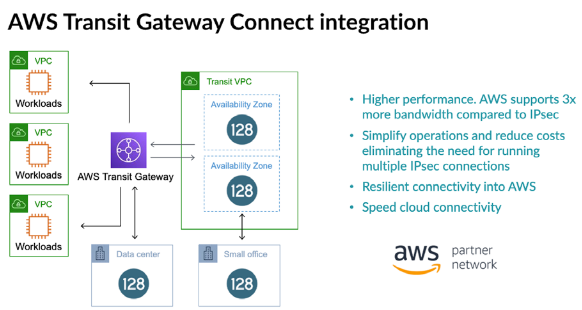

Figure 12 Integration of SSR Into AWS Direct Connect and Transit VPC

Lastly, consolidating the management and analytics from SD-WAN and the rest of the campus means that there is both the inherent management efficiencies discussed previously in the Mist section, but also another SaaS/on-premises platform isn't running due to lack of integration. These both demonstrate scope 3 reductions in emissions either through product/vendor selection or outsourced efficiencies.

Conclusion

The aim of this paper has demonstrated how to qualify the role that our products can play in reducing emissions through intelligent uses of new technology and the efficiencies they provide. Once qualified, each business can quantify the savings they could make by examining how relatable the use cases are with any of the technologies discussed. Businesses are still getting to grips with how to define their green agendas with the ever-increasing necessity to provide a higher level of service and productivity to their consumers and end-users. There isn't a quick fix but rather a journey that must be followed.

As with all new initiatives, there is a requirement for subjective opinions to pave the way forward to a more universally held objective view on the role that vendors, and the consumers of their solutions, can play in the reduction of harmful emissions that continue to endanger the environment.

Although some of the conclusions that I have reached here may be debatable, it is exactly this sort of debate that needs to be encouraged to bring the subject matter into the forefront of conversation when it comes to the decision-making process of future technology selection. Without asking the sort of questions raised herein, we will continue to be bound by traditional high-level cost-performance metrics which feed into current procurement processes.

To find out more about Juniper products and the benefits that they can bring to your business, please contact your account manager or reach out via your partner network for more information. Juniper is committed to assisting your navigation of this emerging and necessary reality and will continue to promote these and other ways of helping our customers along the road to net-zero or other environmental targets.

Further Reading

- Juniper Day One Green
- Optimising Microsoft 365 with Juniper AI-Driven SD-WAN
- Juniper Smart Session Routing
- Juniper Smart Session Router SD-WAN Hero
- Paragon Automation as a Service
- Enabling the AI-Driven Enterprise
- Juniper Connected Security
- Juniper Cloud Metro 2

Paper 11

Optimize Your Lab With Energy Savings Via Virtualization

By Christian Scholz

Energy grids, environment crisis, global warming – you've read these headlines at least a hundred times but can you make a difference? The fact is that you can save resources and energy consumption in your lab setup by using the most powerful feature that every IT system has to offer: virtualization! By virtualizing your lab, or parts of it, you can save money, energy, and a reduction in your carbon footprint while being more flexible and agile. Most importantly you can help to slow climate change down a bit. Your lab savings may not be enough to save the planet, but climate scientists tell us that every reduction helps, every thousandth of a percent lower can and does have a global impact.

In the Lab

My fellow network engineers all know this: when testing new protocols, new features, new designs, you need a lab. A Proof of Concept (PoC) is great but it's only inside a lab that you will know if the technologies being tested will really work in your environment. You must test every bit to verify the desired impact and confirm nothing results in a disaster.

In the past you had to purchase multiple routers, switches, and firewalls similar to your production boxes to achieve this goal, resulting in an enormous amount of heat (which you need to cool down by using more devices), a huge electricity bill, and a *lot* of rack space. And that wasn't even the worst part. Lab changes were mostly done manually and required a technician to travel to the site to patch in new lines and if you forgot a cable, you had to travel all over again, resulting in a lot of CO_2 when traveling. Each of these lab devices consumes electrical energy and needs cooling. Wouldn't it be awesome if just one device that already consumes the energy and that already produces the heat would be able to run multiple devices or even a complete lab, site, or data center? Guess what? This dream *is* possible, and it's called virtualization (see Figure 1). Virtualization is the ability to run multiple *hardware* appliances on shared hardware by separating the hardware from the OS. It's what Juniper did with the RE and PFE separation so that the *brain* can run independently from the hardware.

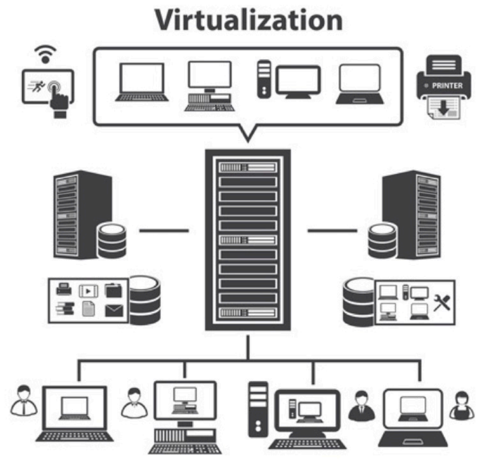

Figure 1 *The Virtualization Concept*

vLab Solutions

Let's have a look at some numbers and begin proving the point that vLabs can conserve energy better than a lab full of racks. Before we begin, note that the focus is on a general calculation since every lab, like every network, is different. Also note that each vDevice has a use case consumption of its own that is not covered nor calculated into our basic assumptions.

Each networking device has a unique power consumption number that you can find within the Juniper data sheets. This consumption, times 3.41, results in the heat emitting from the device in terms of BTU/h. The rack space needed can also be found in the data sheet. Our next few examples will look at the following devices as they are very close to the performance numbers that a vDevice can run on:

- The vSRX compared to an SRX345
- The vMX compared to an MX150
- The vQFX (technically a QFX10k) compared to a QFX5120-48Y

The server running our sample vDevices is an HP DL360G8 with 1HE rack space (see Figure 2). If you use a 2HE, or a 4HE server, with four or six possible CPU slots, the savings can be way bigger as you get more slots that share the same power footprint. Carefully check your use cases and size your vLab accordingly. You want to have neither too few but also not too many resources.

Figure 2 The HP DL360G8 Server

NOTE Obviously there can be certain use cases that cannot be use a vLab. In this case you need the real hardware. But even here, look for ways that you can at least partially virtualize. Remember, every percentage counts.

Here's the basic numbers in our simple math comparison:

SRX345 apc (average power consumption): 122W
SRX345 ahd (average heat dissipation): 420BTU/h
SRX345 rack space: 1HE

MX150 apc: 140W
MX150 ahd: 480BTU/h
MX150 rack space: 1HE

QFX5120-48Y apc: 550W
QFX5120-48Y ahd: 1880BTU/h
QFX5120-48Y rack space: 1HE

DL360G8 apc: 520W
DL360G8 ahd: 1773BTU/h
DL360G8 rack space: 1HE

The server described here has a dual Xeon CPU (six cores each, so twelve cores when using HT) and the max possible RAM (768GB). HT, or *hyper-threading*, is a technology which enables you to double your CPU cores. But this topic is obviously way more complex than intended. If you want to maximize your vLab capability, then learn more about the vConcept. There are a ton of resources on Google.

The aforementioned DL360G8 can use a virtualization solution like EVE-NG, GNS3, or similar. They enable you to run the vendors' software on your server in parallel. So instead of powering up multiple QFX or MX devices, you can simply power up your EVE-NG Server and run multiple devices on it – immediately resulting in less rack space. Let's have a look at how many devices can run (not mixed) per server.

NOTE The numbers for the CPUs and RAM that you are about to see are taken from the official data sheets and reflect the corresponding settings that the real device would use for the same performance.

The vSRX on a DL360G8: 2vCPUs per device, server has 24 logical cores so that makes 12 devices (not looking at technologies to further improve the number of devices):

12 SRX345s:	1464W	5040BTU/h
DL360G8:	520W	1773BTU/h
Savings	*944W*	*3267BTU/h*

The vMX on DL360G8: 4vCPUs per device, server has 24 logical cores so this makes six devices (not looking at technologies to further improve the number of devices):

6 MX150s:	840W	2880BTU/h
DL360G8:	520W	1773BTU/h
Savings	*320W*	*1107BTU/h*

The vQFX on DL360G8: 4vCPUs per device, server has 24 logical cores so this makes 6 devices (not looking at technologies to further improve the number of devices). Let's also look at some costs associated (in Euros, my lab currency).

DL360G8	520W	1773BTU/h,	5,62€/day,	2051,3€/yr
6 QFX5120-48Ys	3300W,	11280BTU/h,	35,64€/day,	13008,6€/yr
Savings	*2780W,*	*9507BTU/h,*	*30,02€/day,*	*10957,3€/yr*

NOTE These savings ignore the needed devices to cool down your lab – the savings are purely the electrical power costs assuming the devices run at the specified wattage 24/7.

NOTE2 And now imagine running the topologies on Juniper's vLabs – you have zero, yes, zero costs! Juniper will deal with the power consumption and cooling for you. And they provide this for free – that's not a typo. Nice of them, isn't in? Check out the various topologies at: https://jlabs.juniper.net/vlabs/.

As you can see, the savings from using a vLab are quite significant. And we just covered the *pure* approach by not mixing the vDevices. Normally in a lab this is exactly what would happen. And the more powerful your server is, the more rack space, power, and cooling you can save. I've seen real life scenarios where customers could save up to two racks worth of equipment and instead use a 2HE Server to achieve the very same – including more agile setups and quicker redesigns.

Example Topology: OSPF - Multi-Area

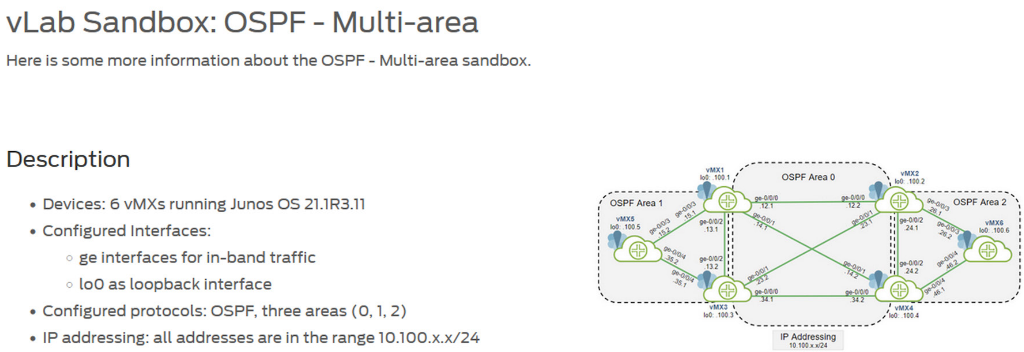

Figure 3 *OSPF Topology*

Let's assume you are a new engineer and need to learn about OSPF. Or maybe you are an OSPF veteran and want to verify a certain behavior that you rarely deal with. This topology runs six vMXs – remember that the MX is *the* router that powers the Internet! It's also heavily optimized to get the most throughput per watt out of these ports – that alone could potentially fill a book (and if you want to learn more on the MX and vMX see this highly recommended book: https://www.juniper.net/documentation/en_US/day-one-books/DayOne_vMX.pdf).

Let's assume that we run this topology in hardware, then on your EVE-NG Server and lastly on the vLabs (where this topology is actually from). Let's also assume that our lab runs eight hours a day:

Hardware

6 MX150s (140W each) = 840W
Total = 840W
8h running = 6,7kWh
Price per Day (assuming 0.29€/kWh) = 1,95€

Virtualization

1 DL360G8 (520W) = 520W
Total = 520W
8h running = 4,2kWh
Price per Day (assuming 0.29€/kWh) = 1,21€
Power Savings per Day = 0,74€

vLabs

Power Savings per Day = 1,95€

While you might think "*C'mon Christian, it's not that much*" the MX listed here needs optics. Depending on the speed of said optics this adds additional power requirements and you also might need to have a license for the hardware depending on the scenario. And again – we just looked at the power consumption costs – the costs for the 6HE rack space and cooling are in addition and not calculated. That makes this little number a bit more scary – and that's just one topology. Consider the world over and the number gets very large, very quickly.

Example Topology: NAT on vSRX

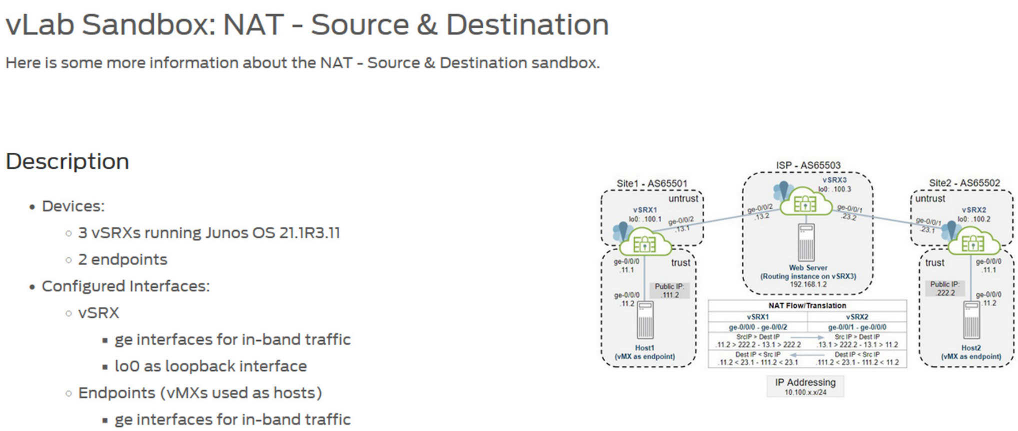

Figure 4 *NAT Topology*

Security is my favorite part because there can't be enough of it in every topology, and it becomes more important each and every day. In this example, we assume that you want to implement NAT (hideous I know – just deploy IPv6 and be happy). But let's say your ISP does (sadly) not yet support IPv6 and we have to lab up some NAT. We use three vSRXs for this (in hardware that's three SRX345s). We also need two servers.
Let's assume that we run this topology in hardware, then on your EVE-NG server and lastly on the vLabs (where this topology is actually from). Let's also assume that our lab runs eight hours a day.

Hardware

3x SRX345 (122W each) = 366W
2x DL360G8 (520W each) = 1040W
Total = 1406W
8h running = 11,2kWh
Price per Day (assuming 0.29€/kWh) = 3,26€

Virtualization

1x DL360G8 (520W) = 520W
Total = 520W
8h running = 4,2kWh
Price per Day (assuming 0.29€/kWh) = 1,21€
Power Savings per Day = 2,05€

vLabs

Power Savings per Day = 3,26€

Savings of just 3,26€ for power does not sound scary – but do you get your kWh for 0.29€? In some parts it goes to 0.90€ and higher – and that's just for the power consumption. As you can see, the vLabs can really save here. The next example is an even scarier number…

Example Topology: EVPN-VXLAN + Apstra on vQFX

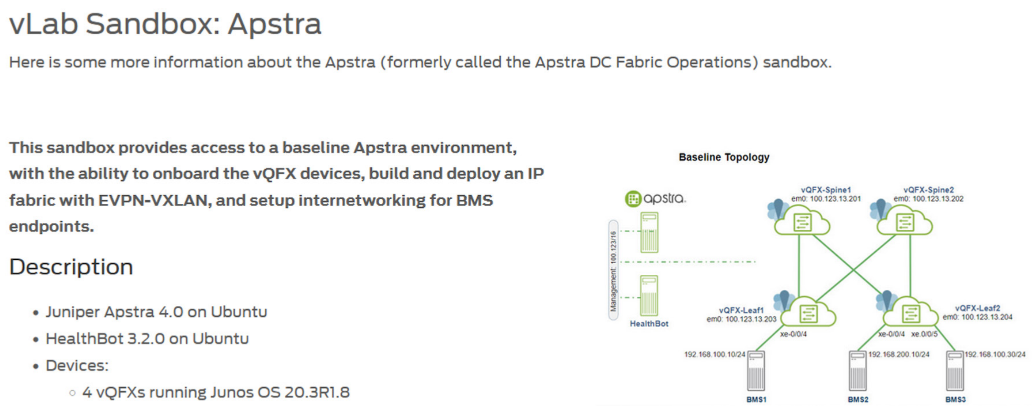

Figure 5 Apstra Topology

Let's have a look at a EVPN-VXLAN spine leaf architecture and add an Apstra Management as a "cherry" on top. This topology runs four vQFX devices (each running the RE and the PFE image) and of course we need an Apstra Server. Apstra can be deployed as a VM or bare metal (since it's basically like EVE-NG just a Ubuntu OS with additional packages). Let's assume that we run this topology in hardware, then on your EVE-NG server and lastly on the vLabs (where this topology is actually from). Let's also assume that our lab runs eight hours a day.

Hardware

4x QFX5120-48Ys (550W each) = 2200W
1x DL360G8 (520W) = 520W
Total = 2720W
8h running = 21,8kWh
Price per Day (assuming 0.29€/kWh) = 6,31€

Virtualization

1x DL360G8 (520W) = 520W
Total = 520W
8h running = 4,2kWh
Price per Day (assuming 0.29€/kWh) = 1,21€
Power Savings per Day = 5,10€

vLabs

Power Savings per Day = 6,31€

So, for running your test fabric for eight hours a day you can save 5,10€ by switching from hardware to a server that runs the devices as vQFX and Apstra-VM, or save even more, the full 6,31€/day, by switching to vLabs. This might sound like a low number but remember, that's each day and it adds up. Wanna know the price per year? 2303,15€ - doesn't sound that cheap anymore, right? And that's just this topology, usually you are running multiple topologies in a lab with way more devices. And in addition, prices are skyrocketing and you are emitting more carbon than you need to.

Summary Thoughts

Obviously, a lab is not just about power consumption. It's also the space needed, the power needed to cool the emitting heat, and even the noise pollution. While the devices usually run in a separate environment it's still pollution, resulting in a *dirty* environment.

We should take into account that test labs should not run 24/7 all the time. You can leverage powerful Junos features to automatically power off the devices via API when no longer needed. With a CI/CD Pipeline and EVE-NG it's even possible to create, spin up, test, and later tear down the vLab all the while sending you a test report to confirm that your config worked.

Another point to be made is that in the world of carbon targeting, where corporate owners are paying to become carbon neutral, when you save power you not only save on the initial costs but you also save on the need to buy extra carbon credits to offset your usage. That means double savings! Energy you didn't use in your labs and carbon credits you don't have to buy. The possibilities are compounding. We are still at the very beginning of a series of best practices for the IT industry and its tens of thousands of labs, to save and reduce energy and advance climate resources. Remember, every incremental percentage counts.

Paper 12

Networking Benchmarks

By Raja Kommula and T. Sridhar

The internet consists of thousands of data centers that could be considered the internet's brains and the network infrastructure can be thought of as its nervous system. That's because these data centers process, store, and communicate data across thousands of services that we rely upon every day. All the information that we access on the internet is stored in these world-spanning data centers. A typical data center consists of servers, storage arrays, power supplies, cooling equipment, and networking devices, and they all need electricity to run. In addition, according to the International Energy Agency (IEA), for every bit of data that travels between an end user and a data center, five bits of data are transmitted within and among data centers.

So, it's not surprising that when internet traffic went up 440% between 2015 and 2021, data centers (excluding crypto) used around 0.9-1.3% of the global electricity demand. Table 1, from the IEA 2021 report, provides more details on the growth in data center energy usage.

Table 1 *Global Trends in Digital and Energy Indicators 2015-2021 (source IEA 2021)*

	2015	2021	Change
Internet users	3 billion	4.9 billion	+60%
Internet traffic	0.6 ZB	3.4 ZB	+440%
Data centre workloads	180 million	650 million	+260%
Data centre energy use (excluding crypto)	200 TWh	220-320 TWh	+10-60%
Crypto mining energy use	4 TWh	100-140 TWh	+2 300-3 300%
Data transmission network energy use	220 TWh	260-340 TWh	+20-60%

Sources: Internet users [ITU (2022)]; internet traffic [IEA analysis based on Cisco (2015); TeleGeography (2022); Cisco (2019), Cisco Visual Networking Index]; data centre workloads [Cisco (2018), Cisco Global Cloud Index]; data centre energy use [IEA analysis based on Malmodin & Lundén (2018); ITU (2020); Masanet et al. (2020); Malmodin (2020); Hintemann & Hinterholzer (2022)]; cryptocurrency mining energy use [IEA analysis based on Cambridge Centre for Alternative Finance (2022); Gallersdörfer, Klaaßen and Stoll (2020); McDonald (2022)]; data transmission network energy use [Malmodin & Lundén (2018); Malmodin (2020); ITU (2020); Coroama (2021); GSMA (2022)].

These are eye-opening numbers. To make matters worse, new companies are working towards making emerging services and technologies mainstream, such as streaming, cloud gaming, blockchain, machine learning, and virtual reality. Crypto mining alone uses nearly 50% of the energy of all the other data centers combined.

The energy consumption in a data center can be divided into four categories: compute, storage, networking, and cooling, but the energy share of networking is growing each year and it is expected to reach 20% by year 2030.

It should be clear to most readers that network energy innovations have not been able to keep up with the explosive growth in networking energy usage. There is a dire need to develop new benchmarks, new techniques, better protocols, and industry-wide standards to improve networking's energy efficiency. We, within the Juniper CTO team, have begun working on two such benchmark proposals, called *Green Quotient* and *Bit Cost*. This paper explores these benchmarks in more detail but do realize that we are still working on them. We are expecting the year 2023 to be a year of discovery on several fronts in networking energy usage.

Green Quotient

Green Quotient (GQ) attempts to quantify the "green-ness" or "green-factor" of a data center (GQ_{dc}) and an application (GQ_{app}) running in a data center. The following are some of the properties of GQ:

- GQ is dynamic. It can change based on the incoming sources of power. For example, during the day, when the majority of the data center can be powered by solar power, resulting in a higher GQ.
- GQ_{app} identifies the energy usage of an application.
- GQ_{dc} identifies the energy usage of a data center.
- GQ assumes that a data center/application is green when running at *g%* capacity, where *g* is the percentage of energy procured through green sources like solar or wind.
- GQ ranges from 0 to 100. Higher the number, the greener an app/DC is.
- GQ goes up exponentially as the energy usage goes up from g% to 100%.

Calculating GQ

Streaming telemetry (see Figure 1) is key to calculating GQ, which can be obtained from, for example, the following sources in a data center:

- CPUs
- Kubernetes infrastructure/Hypervisors and VM infrastructure

- Networking devices
- HVAC
- Storage
- Electric meters

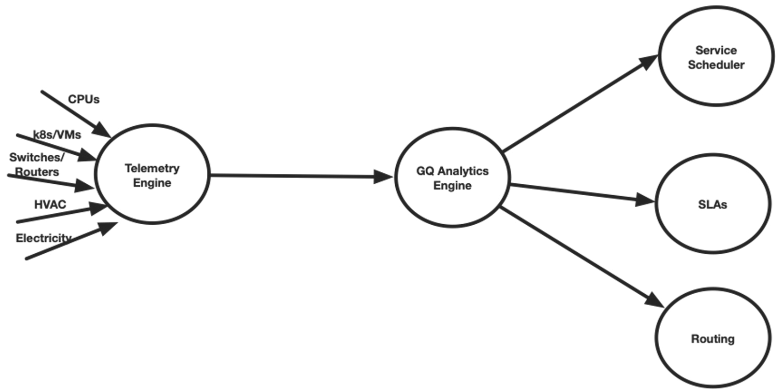

Figure 1 *Calculating GQ*

Telemetry data from these sources can be fed into the following piecewise function to compute GQ_{dc}:

E_T = Total energy usage of a datacenter running at 100% capacity

E_C = Current energy usage of a datacenter

E_Q = Energy Quotient

$$E_Q = \frac{E_C}{E_T} * 100$$

$$GQ_{dc} = \begin{cases} 101 - 101^{\left(\frac{E_Q - g}{100 - g}\right)}, & E_Q > g \\ 100, & E_Q \leq g \end{cases}$$

The function works by keeping the GQ value at 100 if the incoming power is green or carbon neutral. After that, GQ goes from 100 to 0 exponentially. For example, if g is 30%, which means 30% of the energy is coming from renewable sources, GQ will remain at 100 as long as power usage does not hit the 30% threshold. The following figure shows GQ when $g=30$ and when $g=80$.

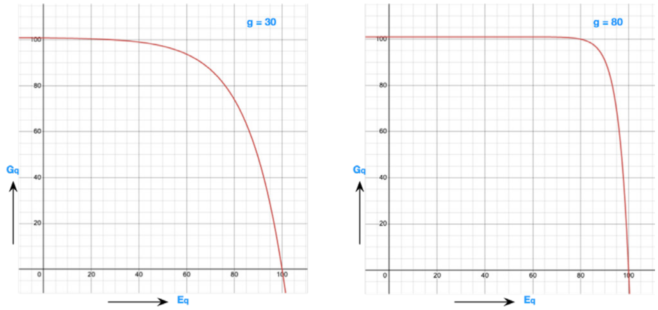

Figure 2 GQ_{de} Graph When g=30 and g=80

Due to the nature of modern applications, which share resources across a distributed set of nodes, finding GQ_{app} is a bit more complex than GQ_{dc}. We can calculate GQapp by starting with a stable data center and run the app over a period of time and take the average energy usage during that time:

- $GQ_{dc1} = GQ_{dc}$ of a stable data center
- GQ_{dc2} = average(GQ_{dc}) of the application started in a stable data center
- $\mathbf{GQ_{\Delta}} = GQ_{dc2} - GQ_{dc1}$

$$GQ_{app} = \begin{cases} GQ_{\Delta}, & GQ_{\Delta} > g \\ 100, & GQ_{\Delta} \leq g \end{cases}$$

Bit Cost

Currently there is no benchmark to measure the cost to transmit a bit between two points on the internet. For example, how much energy in picojoules does it cost to do a simple Google search for the phrase "Google"? There is no way to calculate it at this time. We are describing a framework below as a step in this direction starting with a data center.

Bit Cost should include protocol and acknowledgement overhead. Bit Cost (BC) should also take into consideration the active path the packet traverses. Figure 3 gives more details on Bit Cost for various kinds of transports.

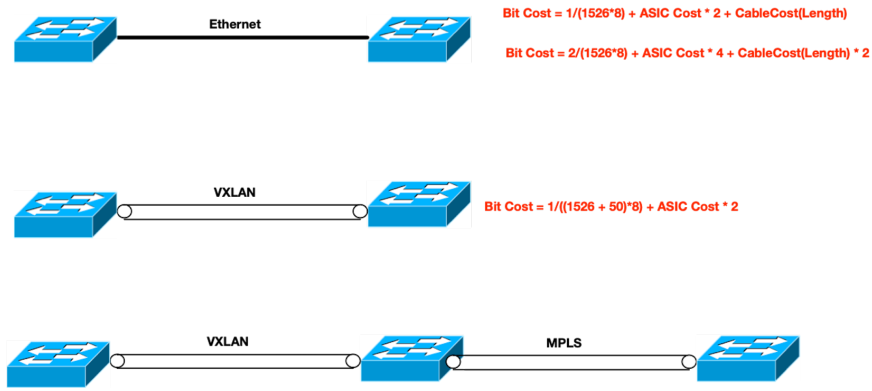

Figure 3 Bit Cost Protocol Overhead

Thus the BC between two directly connected switches should be:

1/(1526 * 8)	+	// Ethernet cost
Silicon Cost * 2	+	// Forwarding cost on both switches
Cable Cost(length)		// Cost to transmit signal in the cable, which is proportionate to the length of the cable

And the BC with ACK between two directly connected switches should be:

2/(1526 * 8)	+	// Ethernet cost of request + ACK
Silicon Cost * 4	+	// Forwarding cost of request + ACK
Cable Cost(length) * 2		// Cost to transmit signal in the cable, which is proportionate to the length of the cable

Similarly, BC can be calculated when data is tunneled inside another protocol header. Figure 4 shows an example of various active and inactive forwarding paths present in the network. BC should be calculated using the path in which data is traversing. Juniper's Paragon Active Assurance product can be used in finding active paths.

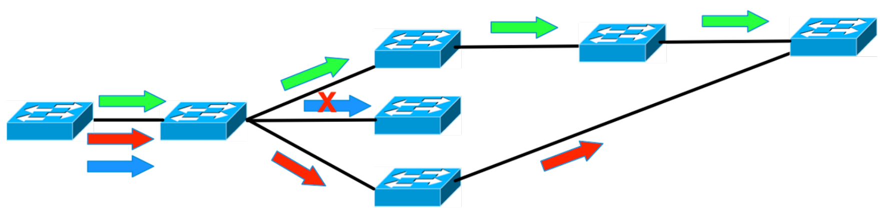

Figure 4 Bit Cost Measured Using Active Paths

Juniper Paragon Active Assurance provides programmable, active testing and monitoring for physical, hybrid, and virtual networks. Unlike passive monitors, it uses active, synthetic traffic to verify performance throughout the life cycle of each application and service.

NOTE The Paragon Active Assurance test agent is natively embedded into Juniper ACX Series Routers. It turns your Juniper Cloud Metro network into the "experience sensor" that enables you to deliver differentiated service levels across your 5G, Cloud Metro, and other networks.

Conclusion

It is evident from this book that we need to benchmark our networks with intelligent metrics to learn how to lower our energy usage. GQ and BC are two methods among several that we're looking at within the Juniper CTO team. There are things we can do today to establish benchmarks for tomorrow and new technologies we can then develop with the results of reliable benchmarking. Let's do this.

3100063-001-EN- Feb 2023