


Networking Leadership Brief

Sponsored by:



Scale Everywhere: How Resolight's Photonic Processing Could Revolutionize AI Networking

Highlights:

- **Israel-based Resolight has developed photonic processing technology that has the potential to disrupt the datacenter networking industry.** Photonic processing can boost bandwidth and lower energy consumption by an order of magnitude.
- **Resolight's photonic processing addresses key bottlenecks and energy economics in AI infrastructure. This breakthrough is key for AI datacenter scaling and profitability.**
- **Resolight is currently working with design partners to develop all-optical switching systems.** These networking systems will work with optical technologies in development, such as co-packaged optics (CPO). **Photonic Processing in the network coupled with CPO will make the network 10X more scalable.**
- **This Leadership Brief dives deep into the challenges of hyperscale datacenter networking and the potential of Resolight's technology to transform the economics and landscape of the optical switching and CPO markets**

1. Intro: AI Drives the Need for a Giant Leap in Networking Infrastructure

With a projected \$1 trillion in capital expense (capex) expected to go into datacenter buildouts in 2026, some serious money is in play. Most of this investment is designed to support the launch of AI and related services. The world has never seen such a scale of technology investment, so the stakes are higher than ever.

AI services require huge amounts of compute, networking, and storage infrastructure. This infrastructure also requires immense consumption of power, increasing operating costs. At the same time, the return on investment (ROI) for these services is unknown, so the economics of operating this infrastructure will be a key factor in determining the success of AI services.

With the deployment of trillions of dollars in capex into AI datacenters over the next few years, the costs of infrastructure and energy consumption are set to skyrocket—and new technology is needed to rapidly lower the costs of operations to drive profitability of this infrastructure. Even modest technology innovation has the potential to save hundreds of billions of dollars and untold amounts of energy.

Networking Remains a Bottleneck

Networking will be the key to it all. The connectivity among GPUs, CPUs, TPUs, racks, and systems remains a large bottleneck in delivering hyperscale AI datacenters. It also contributes to the huge energy and datacenter real-estate footprints.

At the current rate, networking will escalate in cost, even though it is generally held to be 8%-12% of the total datacenter cost. As tens of millions of GPUs are shipped to datacenters around the world, they require millions of electronic switches and hundreds of millions of network connections, which come with huge energy consumption and massive complexity.

The economics and technology for compute are already well developed, with NVIDIA building out state-of-the-art GPU systems and competitors such as AMD and the major cloud providers following fast with advanced solutions of their own. But as massive arrays of GPUs, TPUs, and CPUs are built—in the form of racks, clusters, and even massive datacenters connected over many miles—networking is struggling to keep up, whether that's in the compute racks or out of the rack.

The opportunity costs are high. If you can't scale the network, you can't scale the datacenter.

What's going to be needed is a giant leap forward in the scale and economics of networking. Networking components are starting to run up against the laws of physics. The industry consensus is

that innovative optical networking technology needs to be driven deeper into the datacenter, as soon as possible.

So far, proposed solutions such as co-packaged optics (CPO) or Linear Pluggable Optics (LPO) are only a piece of the solution. CPO technology is exciting as a means to provide a more efficient way for GPUs to communicate. Better LPO will also help improve the economics of energy and space. However, sending massive amounts of data over lightwaves into the networking gear, without matching it with equally powerful processing capability in the optical domain, doesn't provide the entire solution to the scaling challenge. To really move the needle, what's needed is a vast architectural shift to end-to-end optical networking delivering order-of-magnitude gains in bandwidth, power efficiency, and costs.

Powering Up Power Demand

In addition to theoretical limits to electronic processing power for networking, yet another key bottleneck is energy efficiency. Power is not only in short supply, but the costs are significant. By many estimates, power makes up to as much as 50% of the cost of datacenter operations, so delivering power efficiency will be key to building out new datacenters efficiently.

The bottlenecks are connected. Behind the huge energy consumption of datacenters are reams of electronics. Millions of chips and circuits are packed into racks and networking systems to make it all work. If you want to reduce the amount of power being consumed, you need to do three things:

- 1) You need to make each component more efficient.
- 2) You need an architecture that reduces the number of electrical components and devices.
- 3) You need to use light to communicate as much as possible, and refrain from "down converting" it to electronic signals along the way.

The industry has so far been focused on number 1. But the real scale comes in 2 & 3. The industry's most respected technologists have said that replacing electronic processing with photonic processing would be the way to go if it were possible. As we'll show in this Leadership Brief, the technology developed by Resolight has the potential to change the game, because it shifts networking architecture to massively reduce the networking components and electronics needed to build terabit-scale networking. This could shift the economics by an order of magnitude and transform both the reach and capacity of AI datacenters and provide a leap-frog opportunity for their suppliers

Let's show how all of this could happen.

2. A Major Architectural Shift Is Needed to Scale Everywhere: Photonic Processing

To meet the scale needed for AI datacenters, incremental improvements in networking are no longer enough. What's needed is a quantum leap in networking power that can deliver as much as 10X bandwidth at 1/10th the power. Order-of-magnitude changes like this can only be delivered by big changes in architectural thinking.

Ofer Shapiro, the CEO and Cofounder of Resolight, believes his team has the answer. Shapiro says that huge gains in technology can only come with thinking about a different architecture.

Shapiro has proven this thinking five times before, as the architect behind dramatic shifts in video routing and other technologies. This includes developing a Scalable Video Coding router to deliver multi-party video conferencing to the desktop and mobile, developing the first IP-based MCU (multipoint communication server), along with other disruptive technologies still in use. These disruptive architectures were licensed and OEM'd by some of the largest technology companies such as Cisco, Google, Ricoh, and Nintendo and deployed by Fortune 100 companies, the Department of Defense, and a large share of the world's largest hospitals.

Shapiro with his new company, Resolight, has now turned his sights on networking in AI datacenters, which he believes need an architectural shift that shuns incremental gains for order-of-magnitude gains in bandwidth and power efficiency. The only way to do this, says Shapiro, is to eliminate conventional electronic processing in the core of GPU switching networks and introduce photonic processing to build end-to-end optical switching, which is more efficient and less power hungry.

Using any form of optical technology with electronic switch processing will never be able to compete with the scale and latency advantages of pure photonic processing technology. Shapiro has built such a photonic processing architecture, the processing chip is ready, and the company is now working with design partners.

Where Things Get Out of Rack

Why is this change needed now? It's simple math. With the size of AI infrastructure deployments being discussed, conventional electronic switching can't keep up.

Imagine a standard 64- or 72-GPU rack in a "scale up" scenario. Inside that rack, typically you would require 18 switches to connect the GPUs. These switches, dominated by NVIDIA's NVLink and NVSwitch today, are built in the electrical and copper domain. With a recent \$4 billion investment in photonics partners Coherent and Lumentum, NVIDIA clearly thinks optics are the future. NVIDIA CEO Jensen Huang has said that optical networking will eventually be the best solution, but that copper-based networking in the rack is better for now.

Leadership Brief: Resolight Ready to Scale AI Networking Everywhere


NVIDIA's NVLink and NVSwitch were engineered to solve the tough problems of connecting NVIDIA's massively powerful GPUs at the rack level—the scale-up challenge. In Rubin, these rack-level networking systems will connect as many as 72 GPUs in a rack with 14.4 Tbps of networking bandwidth per GPU (one-way).

“With a recent \$4 billion investment in photonics partners Coherent and Lumentum, NVIDIA clearly thinks optics are the future. NVIDIA CEO Jensen Huang has said that optical networking will eventually be the best solution, but that copper-based networking in the rack is better for now.”

But here's the rub: NVLink is still connected with copper. NVIDIA has expressed the need to move to more integrated optics, and it is investing and designing future solutions including CPO. But CPO is a component-level connectivity technology, not a network switching technology. As you move to “scale out” and “scale across” not only for AI training but for other applications such as inferencing, you will need more powerful switching. To connect datacenter racks, you will need dozens or hundreds of switches, depending on the size of the datacenter. This equates to millions of switches across the globe for AI infrastructure.

The focus on scale-up solutions, including the potential for CPO, shows how the industry is addressing bottlenecks one step at a time. As you scale out, things get scarier. If optical integration solutions such as CPO (or other technologies) are successful at scaling bandwidth at the component level, it only exacerbates networking bottlenecks at the top-of-rack and scale-out level. The networking dependencies grow larger.

Photonic processing, which keeps switching in the optical domain, is the only way to provide an order-of-magnitude reduction in the devices and power needed. In the diagram below, we show both the problem and the solution. Resolight says that its photonic processing technology can be used to scale the capacity of switching to 1,800 Tbps on each switch, reducing the number of switches needed by more than a factor of 10. This would also reduce additional infrastructure needs such as the number of transceivers, cabling, and power.



To Recap those numbers, they look like this:

Resolight Photonic Processing Scaling vs. Electronic Switching

	Tier	Packet Switches Needed	Resolight Photonic Switches Needed
# Switches	Tier 1	128 X 18 = 2,304	128
	Tier 2	64 X 18 = 1152	64
# Connections	Tier1	64 X 18 X 128 = 147,456	64 X 128 = 8,192
	Tier 2	1152 X 128 = 147,456	64 X 128 = 8,192
# Transceivers per connection		8	2

Imagine the power and space reduction in this scenario. In the diagram above, optical processing powering the switching layer could reduce the number of needed connections by the thousands, both in the rack and top-of-rack. With higher bandwidth on each connection, it would also reduce the number of transceivers by a fourth – providing significant relief on painful supply chain concerns.

With datacenters now scaling to 100s of thousands of GPUs being networked together in scale-out clusters, the numbers are huge—potentially hundreds of billions of dollars. The cost of networking and energy could mean the difference between profit or loss for many AI infrastructure companies.

Why CPO Isn't Enough Without Photonic Switching

Why now? The industry consensus is that innovations are needed in the optical domain to achieve end-to-end, all-optical nirvana. Leading industry luminaries such as NVIDIA CEO Jensen Huang, Broadcom CEO Hock Tan, and Arista Networks Cofounder and Chief Architect Andy Bechtolsheim have all stated

publicly that innovations in all-optical networking technology are the goal of the networking industry. Cisco bought optical component provider Acacia in 2001, understanding the significance of optical systems. The challenge has come in delivering an all-optical networking architecture.

Let's look at one example. Broadcom has remained a key networking innovator by unveiling Tomahawk 6, a single chip with the capacity to switch at 102 Tbps and the option to include CPO. But Tan says the goal is to eventually go all-optical. See the quote below, from a recent Broadcom conference call (December 2025):

When you talk about your next question of silicon photonics as a means to create basically much better, more efficient, lower power interconnects in not just scale-out, but hopefully, scale-up. Yes, I could see a point in time in the future when silicon photonics method is the only way to do it. We're not quite there yet ... try to do scale-up within a rack in copper as long as possible and in scale-out, non-pluggable optics.

– Hock Tan

The industry is currently investing in CPO and LPO as potentially more efficient ways to integrate the switching electronics and optics, thereby increasing bandwidth and reducing power consumption for in-rack connections between GPUs. Some experts believe that more efficient LPO technology will be a near-term solution by reducing power consumption, latency, and the need for as many racks (by consolidating optical component space), but it doesn't provide photonic switching.

CPO integrates optical connectivity with an electronic ASIC, bringing the optics closer to the heart of processing. But CPO technology is not yet a mature technology, and many experts believe it has challenges (such as in manufacturing). When and if it becomes widespread, CPO would only be part of the solution, because even though it enables optical connections to be more easily embedded in ASICs, it still requires processing packets in electrical switches in the core of the network. "CPO is very important but it's not addressing the full problem," says Shapiro. "It's more efficient but it just gets it to the next bottleneck. What really helps is an all-optical switch powered by photonic processing that has 10X more capacity.

"CPO is very important but it's not addressing the full problem. It's more efficient but it just gets it to the next bottleneck. What really helps is an all-optical switch that has 10X more capacity."

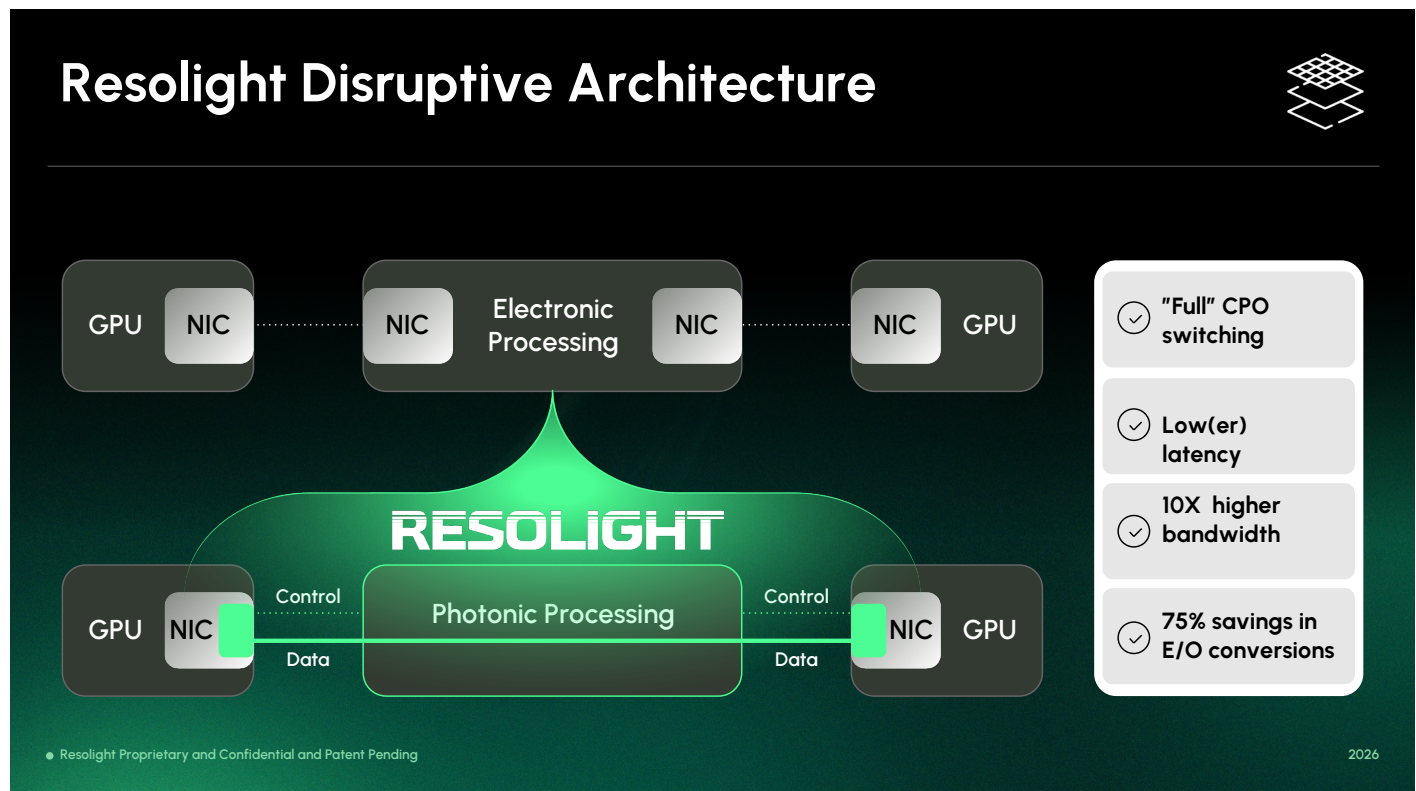
– Ofer Shapiro

Leadership Brief: Resolight Ready to Scale AI Networking Everywhere

Resolight is working on exactly that problem. The diagram below, courtesy of Resolight, shows how CPO can solve the problem of connectivity between the GPUs or CPUs, but it doesn't solve the problem once you go outside the network—you still need an electrical switch. To provide real scaling, you need to eliminate the step of electronic conversion. Even the latest and greatest switching platform such as the Tomahawk 6 can only handle 102 Tbps. Compare this with GPU two-way bandwidth of 28.8 Tbps (14.4 x 2) and you will still need one switch for every four GPUs—clearly not a scalable networking proposition.

What is needed is what Resolight calls “full CPO” switching, in which the full bandwidth coming out of the GPU can be handled and switched in the optical domain, without being converted to an electrical switch. This can only be offered through photonic processing.

As demonstrated in the diagram below, Resolight technology enables any optical connection to the GPU, via CPO or otherwise, to connect through photonic processing and keep all traffic in the optical domain, without electronic processing.



Scale up, Scale Out... or Scale Everywhere

What does this mean for AI infrastructure architectures? The answer is that photonic processing could provide a more scalable, all-optical switching environment for massive GPU clusters, in any configuration: Scale out, scale across, or ... scale everywhere.

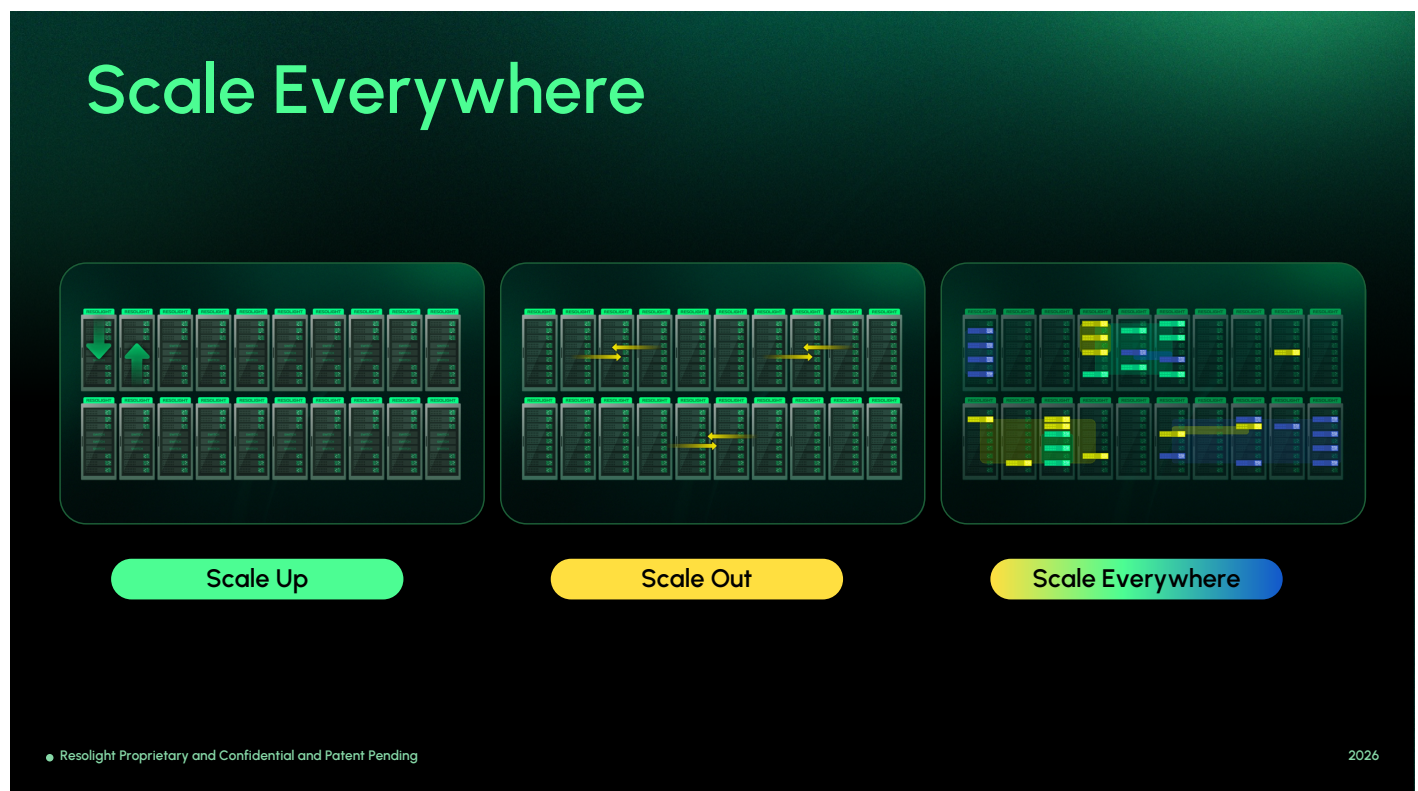
Leadership Brief: Resolight Ready to Scale AI Networking Everywhere

After all, the industry terminology “scale up,” “scale out,” and “scale across,” simply describe different architectures that were developed when different parts of the network hit scaling bottlenecks. But even when these fragmented solutions are pieced together, they still run into electronic switching bottlenecks.

Resolight proposes an architecture that can boost the bandwidth, reduce energy consumption, and provide low-latency for any deployment, both for scale-up and scale out. Only all-optical networking technologies can get you to that level.

Imagine photonic processing technology powering switching that keeps the data in the optical domain the entire time. At the same time, these switching solutions scale to 100,000s of GPUs (rather than capping out at hundreds, or just 72 in a scale-up scenario); and they connect the chips in an any-to-any fabric, circumventing the need for separate silos, providing full scale-up bandwidth across the cluster.

On the left and center in the diagram below, we have scale-up and scale-out scenarios, respectively. In those silos, GPUs are bound by the networking architecture in each scenario. In the colored diagram on the right, a GPU or GPU array in an all-optical architecture can be networked with any other GPU in cluster of 1,000s of GPUs. That is the concept of scale everywhere.



3. Inside Resolight: The Technology Behind Photonic Processing

Optical switches have been around for decades, and they have been one of the great innovations that drive the Internet. The introduction of dense wave division multiplexing (DWDM), changed the game in long-haul networks, enabling orders-of-magnitude gains in bandwidth. But most commercial optical switching still requires some form of electronic packet processing at the core, with optical-electronic-optical (OEO) equipment. In an OEO system, electronic switching occurs in the middle of optical transmission, like a tollbooth on a highway.

The missing piece of all-optical networking has been the capability to process networking packets as fast as conventional electronic switches while keeping all the data in the optical domain, avoiding expensive conversions to electronic circuitry. It's the latency of all-optical systems that's been a key challenge. Photonic processing is the key to solving that.

Past Shortfalls in OCS

All-optical switching, or optical circuit switching (OCS), has never quite fulfilled all the requirements to be adopted in broad datacenter applications. The primary reason, according to engineering sources, is the need for very low latency. MEMs and other OCS systems historically have lagged electronic switching (performing at one millisecond) in providing the adequate latency—especially for AI systems.

Google is currently using all-optical MEMs in specific datacenter use cases, and it says that it has achieved sub-100-microsecond latency per switch. But when you examine how many switches are needed to cover a modest 100-port system, you come up to a millisecond-level latency—a thousand times slower than the scale-up switches. This may explain why this technology is not more widespread. Much of Google's datacenter switching is still OEO, and few datacenters worldwide employ MEMs or other OCS technology. The technology is applicable for reconfiguration or fail-overs, but not for general Ethernet switching.


The bottom line is that despite predictions of the use of photonic OCS for two decades, it has yet to penetrate the datacenter or enterprise infrastructure at a mass-market level. The main reason, as we have described, comes in challenges to match the low latency of electronic switching platforms.

Resolight: Key Innovations in Photonic Processing

How does Resolight make this happen? In the Resolight architecture, a photonic processing device controls networking decisions on an all-optical switch. The company has pending patents on this technology.

Photonic processing, as proposed by Resolight, would remove the need to process packet header information and instead would embed the information in a photonic processor that could route traffic using direct control of the optical circuits in latency performance that surpass even that of the electronic switches.

Resolight has developed prototype devices using cutting edge TFLN (Thin Film Lithium Niobate) & Silicon Nitride to enable optical circuits to be manipulated directly within the characteristics of light, rather than requiring conversion to electronic circuitry, a process known as photonic processing. By keeping the network processing in the optical domain, the network can scale by 10X while at the same time reducing required energy for switching by 90% and the number of switching servers by a similar amount. Photonic processing as a replacement for a data switch was considered the holy grail yet was unachievable by experts. Only now with the availability of new manufacturing technologies and Resolight's unique patent-pending server technology and proposed standards-based approach is Photonic Processing a reality that will transform the networking space.



As we have shown, the savings of this technology could come in many forms:

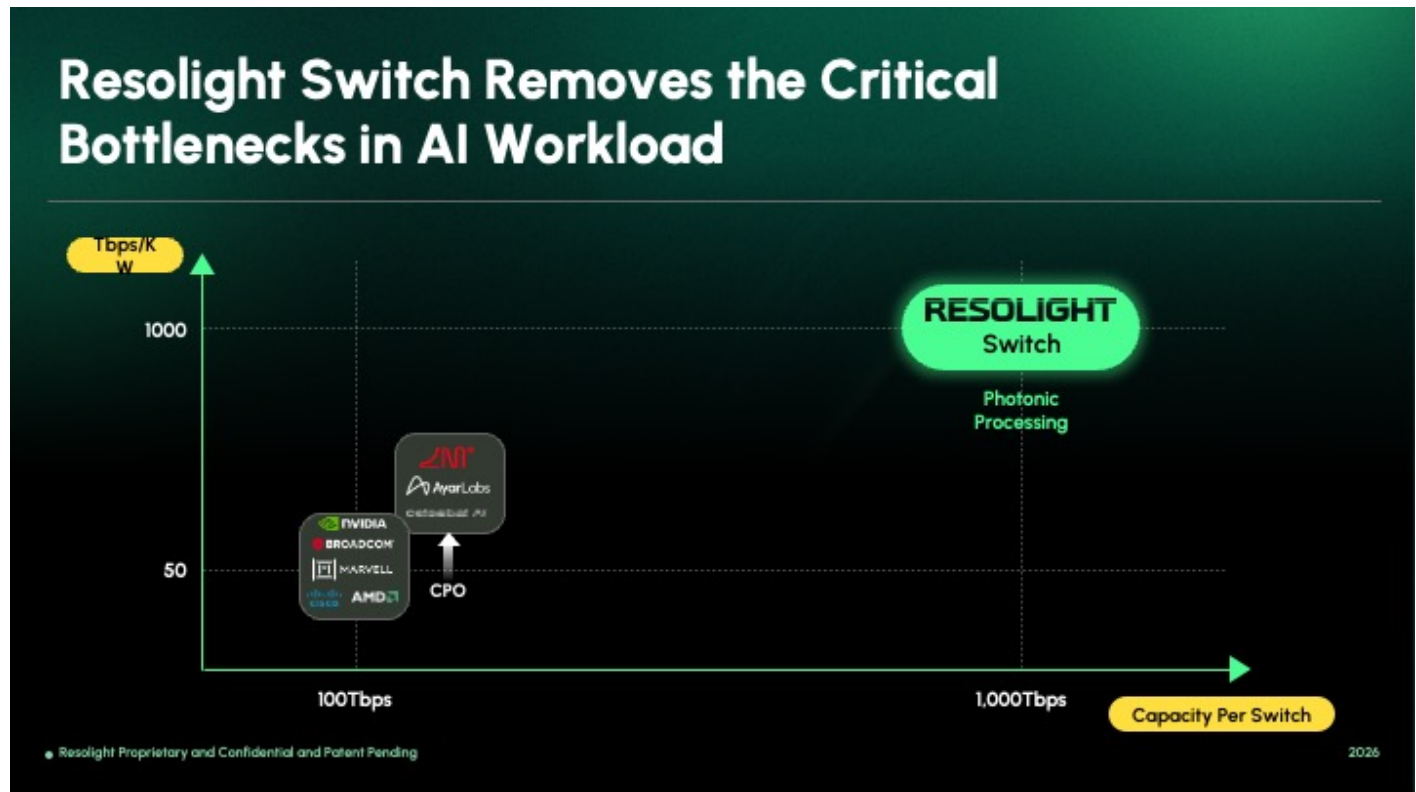
- A dramatic reduction in the number of switching devices required, because the all-optical switching environment can handle more bandwidth and switching requirements per device.
- Any-to-any connectivity of GPUs, whether that's inside a GPU rack, connectivity to an adjacent rack, or even between datacenters.
- A massive reduction in the overall energy consumption by the network.

Leadership Brief: Resolight Ready to Scale AI Networking Everywhere

Resolight says that its tests have shown this technology can deliver 10X higher bandwidth and reduce electrical-optical conversions by as much as 75%. It would also deliver lower latency and reduce the number of components, such as transceivers.

The best way to measure the current state of networking efficiency is by comparing bandwidth to energy consumption. Most of the advanced switching solutions for AI infrastructure deliver 100 Tbps/Kilowatt, with CPO possibly advancing these numbers incrementally but unlikely attaining an order-of-magnitude improvement (and still leaving the challenge of requiring electronic switching). Resolight claims that with photonic processing, it can attain an order-of-magnitude leap: 1,000 Tbps/Kilowatt with total switch capacity in excess of 1,000 Tbps.

“Every box will do 1,000 Tbps with one kilowatt,” says Shapiro. “That’s an order of magnitude increase in both dimensions”



Inside Resolight's Technology

The Resolight team and board is comprised of employees and advisors with photonic, routing, quantum computing, semiconductor, and deep-tech experience at companies that have licensed technology or sold companies to Cisco, Google, Nvidia (Mellanox), and Intel. They have decades of experience in building companies from pre-seed to IPO.

Key members of the staff include:

CEO and Co-founder Ofer Shapiro has a decades-long track record in bringing innovative architectural changes to industries that result in leap frog performance gains and unmatched advantages in cost and time to market. He is an expert in bringing to market new server architectures that leverage new standards-based communication approaches and cross layer server architectures that leverage it to provide a 10x advantage. He co-founded or was the initial lead investor in 12 companies and is a named inventor on 52 patents. He built and licensed disruptive architectures that secured nearly 100% market share enabling an IPO, was CEO, co-founder and the architect of the first multi-party video conferencing on desktop and mobile devices which was licensed to dozens of companies. Shapiro was named a World Economic Forum Technology Pioneer, and he received the Wall Street Journal Innovation award for economic disruption that transformed an entire industry, in the category of Internet, Networking and Broadband.

VP of Research and Development and Co-founder Raanan Ben Shahr was CEO, VP of R&D and Director of Innovation at technical/defense and deep-tech companies.

Chief Architect and Co-founder Tamir Weinstock was a team leader and strategic project manager of cutting-edge multi-disciplinary systems in electro-optics RF, algorithms, and electronics in the Israeli Defense Forces.

VP of Corporate Development Mari Mineta Clapp has led GTM efforts with Silicon-Valley based networking and AI start-ups including Arista, Apstra (Juniper), Big Switch (Arista), Alkira, and Otter.ai. She has also served as a strategic advisor to the CTO of AT&T and Managing Director of ONUG.

The company's Intellectual Property includes technology developed over many years and exclusively licensed from Tel Aviv University as well as proprietary patents developed in-house around the company's photonic processing architecture. This includes prototype devices using cutting edge TFLN & Silicon Nitride to enable optical circuits to be manipulated directly within the characteristics of light, rather than requiring conversion to electronic circuitry, a process known as photonic processing.

4. Conclusion: It's Time to Scale Everywhere with Photonic Processing to Deliver AI Networking Infrastructure

All-optical switching has been the holy grail of networking for decades, but it's always been elusive. Resolight now believes it has the photonic processing innovation that will finally enable all-optical, low-latency switching with 10X better economics than electronic switching.

The question of whether this is needed is not debated—it's just been a matter of when it arrives. The key needs for all-optical switching are already accepted by the industry. Here's why it's needed:

- **Speed & Latency:** To achieve maximum bandwidth transmission at maximum distance, networking signals need to move at the speed of light, drastically reducing the time it takes for data to travel across a chip.
- **Energy Efficiency:** Unlike electrons, photons do not experience electrical resistance, meaning they generate an order of magnitude less heat and require significantly less power for data movement. Moreover, handling signals in the optical domain presents an opportunity to provide order-of-magnitude improvements in all problematic dimensions: power, cost, number of boxes, number of connections, reliability, and datacenter space.
- **AI Acceleration:** Modern photonic processing technology can be used to speed up not only compute functions but also those that control networking, delivering speeds up to 10X higher at 90% lower energy consumption.

Debates about CPO and potential innovations at the GPU, rack, and component levels are a moot point if the core of the network can't scale. The move to photonic processing and all-optical switching has the potential to revolutionize the AI industry by lowering the costs of infrastructure and accommodating huge amounts of bandwidth.

Resolight's architectural vision presents a refreshing perspective in an industry that has been regularly bumping up against bandwidth bottlenecks. We look forward to watching this development as Resolight builds systems with partners in 2026.

For more information, go to www.Resolight.ai.