


application note



Dynamic photonic infrastructure for AI data centers enabled by Drut and POLATIS®

Introduction

AI infrastructure providers and data center operators manage complex, diverse environments with evolving workloads. To stay competitive, they must deliver high performance, efficiency, and support for next-gen applications. This requires evaluation of current infrastructure and optimization of new technologies without compromising stability or ROI.

Drut makes this possible with a disaggregated, composable architecture powered by a photonic fabric, enabling seamless integration, efficient scaling, and faster innovation.

Key features

- All-optical photonic fabric interconnect with POLATIS® optical circuit switches
 - Disaggregated compute, storage, and acceleration
 - Real-time resource composability
 - High-density architecture with energy-efficient design
 - Interoperable with industry-standard hardware and software
-

Challenge

Traditional compute infrastructure relies on static servers with tightly bound CPUs, GPUs, memory, and storage. This rigid design limits scalability, drives up costs, and leads to poor resource utilization. As AI, HPC, and cloud workloads become increasingly dynamic and diverse, these systems struggle to keep pace, leaving accelerators underutilized and making upgrades costly and inefficient.

- **Inflexibility of traditional architectures**

Servers are monolithic – compute, storage, and accelerators are physically tied together, making it hard to adapt to changing workload needs.

- **Resource contention and security in shared environments**

In shared infrastructure, it is hard to isolate workloads securely and ensure consistent performance.

- **Cost and complexity of scaling AI infrastructure**

Scaling typically requires overprovisioning or forklift upgrades; costly, slow and disruptive.

Solution

Integrating Drut's DynamicXcelerator™ solution with the POLATIS® Optical Circuit Switch (OCS) enables full disaggregation of compute, storage, and accelerators, allowing resources to be independently deployed and dynamically interconnected based on workload requirements. It breaks away from the limitations of fixed configurations, offering greater flexibility and efficiency.

- **Enables modular, incremental scaling of infrastructure**

Add only what is needed.

- **Delivers physical-layer isolation**

For secure, deterministic workload connections, minimizing contention in shared environments.


POLATIS® – single mode low loss optical circuit switch

POLATIS® is a high-performance, fully non-blocking all-optical circuit switch, available in sizes from 8x8 to 384x384 ports. Designed for demanding applications, it delivers low optical loss, fast switching, low power consumption, and all in a compact form factor. Built on patented DirectLight™ technology, it offers ultra-low latency for time-critical, hybrid cloud services. Proven in data center and telecom environments, it is also widely used for automating optical component testing.

- Carrier-class reliability and interfaces from the worldwide leader in optical circuit switches.
- SDN enabled via embedded control interfaces allows reconfiguration of the cross-connect on demand.
- POLATIS' patented DirectLight™ beam-steering technology sets the industry standard for optical switching performance and reliability.
- Optional fully integrated power monitors and variable optical attenuation.



POLATIS 192x192 OCS



POLATIS 48x48 OCS

From fixed to fluid: How Drut redefines data center design

Drut Technologies offers a disaggregated data center infrastructure built to meet the performance and scalability needs of modern workloads. Its PCIe Gen5 x16-based PRU 2500 hardware series, combined with software-defined orchestration, enables flexible PCIe resource management across AI, machine learning, enterprise, and cloud environments. The Drut Software Platform (DSP) adds advanced on-prem cloud capabilities, automation, and integration, allowing PCIe resources to be provisioned and reallocated as needed.

At the core is Drut's photonic disaggregation technology, which creates low-latency, dynamic connections between servers and PCIe devices. This removes traditional resource silos, improves utilization, and allows GPUs, NVME, NICs, and all manner of PCIe cards to be reassigned in real time without disrupting host systems. Drut is also the first company to commercially implement co-packaged optics (CPO) to accommodate PCIe Gen5 x16's lane density in a small, compact form, removing the need for transceivers, all while maintaining compatibility with existing hardware and software.

This architecture is ideal for high-performance workloads that require frequent scaling, such as scientific research, animation, and geophysical modeling. It simplifies GPU cluster management, improves peer-to-peer connectivity, and supports both horizontal and vertical scaling. Users can even switch between GPU vendors – like NVIDIA® and AMD – giving more flexibility in matching compute resources to specific tasks.

By reducing latency and server sprawl, Drut helps organizations scale more efficiently and extend the life of existing infrastructure. The system supports advanced PCIe workflows like GPUDirect Storage (GDS) while keeping data within the PRU 2500 backplane and out of the host memory path.

The PRU 2500 chassis supports up to 12 PCIe devices (including 10 double-wide GPUs). The backplane provides 288 lanes of PCIe Gen5 switching, supporting four x16 upstream and twelve x16 downstream PCIe slots with full peer-to-peer communication.

Drut's platform enables adaptable architectures that lower ongoing costs while supporting multi-vendor environments and future upgrades.

Key benefits

- **Dynamic resource pooling and orchestration**

The Drut Software Platform (DSP) provisions and reallocates PCIe resources in real time across AI, ML, enterprise, and cloud workloads.

- **Optical switching supports low-latency data movement**

Drut's photonic disaggregation and co-packaged optics in tandem with POLATIS® OCS deliver ultra-low latency connections between servers and PCIe devices.

- **Secure, locked mapping of devices to hosts**

Ensures stability and protection when GPUs, NVME, and NICs are reassigned dynamically without disrupting host systems.

- **Scalable design using single-mode fiber**

PRU 2500 backplane and photonic architecture supports horizontal and vertical scaling for HPC, scientific research, and other compute-intensive workloads.


- **Independent upgrades of compute, storage, and accelerators**

Disaggregation eliminates server sprawl and extends infrastructure lifespan while reducing costs.

- **Vendor-neutral architecture supporting heterogeneous environments**

Allows flexibility to switch between NVIDIA®, AMD, and other PCIe devices, ensuring users can tailor compute resources to their specific tasks.

PRU 2500 system overview



Server connectivity

- Servers connect externally to shared resources instead of hosting accelerators locally.

POLATIS® optical circuit switch

- The central POLATIS switch provides all-optical, ultra-low latency connectivity between servers and PCIe devices.

PCIe device pooling

- GPUs, FPGAs, NICs, and storage devices are hosted in PCIe slots in the PRU 2500 and made available as a shared pool of resources.

Dynamic resource allocation

- Orchestration software assigns PRU 2500 resources to servers on demand, enabling flexible workload-driven composition.


Composable photonic-native infrastructure

- The architecture supports independent scaling of compute and accelerators, delivering true photonic-native composability.

Software-defined photonic infrastructure

Traditional data centers lock resources into static configurations, creating silos, stranded capacity, and costly inefficiencies. Drut's Software-Defined Photonic Infrastructure transforms that model into a fluid, dynamic fabric where compute, storage, and accelerators can move freely to meet workload demand.

By combining the PRU 2500 hardware, POLATIS® photonic switching, and the Drut Software Platform, organizations gain a disaggregated, software-orchestrated environment that delivers flexibility, efficiency, and scalability for AI, HPC, and enterprise workloads.



Share pools of resources

- Disaggregation with photonic switching makes GPUs, NVMe, NICs, and other PCIe devices part of a common pool, available to any host when needed.

Break static data center constraints

- Move beyond fixed server configurations with dynamic optical connections that reassign devices across hosts in real time.

Simplify AI growth

- Scale GPU clusters horizontally or vertically, mix vendors like NVIDIA and AMD, and expand deployments with ease

Improve infrastructure utilization

- Reduce stranded resources and maximize efficiency with orchestration that ensures devices are fully used instead of locked to single servers.

Save on CAPEX, OPEX, and TCO

- Extend the life of existing infrastructure, eliminate server sprawl, and reduce ongoing operating costs with independent upgrades of compute, storage, and accelerators.

More to explore

Please visit:

hubersuhner.com/en/markets/communication/data-center

drut.io

Email: info@drut.io