



DynamicXcelerator AI Photonic Bundle

Key Benefits

- Drut's dynamic attach and detach capability allows workstations to be given additional hardware resources in a non-disruptive manner
- Accelerates a wide range of industry workloads
- Fully disaggregated solution allowing for machines to be created and attached to HPC workloads
- Use of photonics removes physical rack locality requirements
- Better TCO by decoupling GPU resources from server upgrade path. Upgrade GPUs independent of CPU upgrades
- Lower latencies by using a direct connect photonic fabric versus stacked hierarchy of switching
- Efficiency – Grouping resources via direct connect resolves the stranded resource challenge
- Security – Workloads do not transgress over racks, spines and cores

Data centers are fast approaching the power and performance inflection point, where the performance gains from electrical interconnects diminish and the power consumption required to transmit electrons through copper continue to escalate. We are now at the point where photonic interconnects can begin to be introduced to the data center. Over the next few years, the introduction of NxCPo chiplets will allow for the building of denser and lower cost photonic fabrics for the data center. Fortunately for AI users, the initial system proofs are now available from Drut Technologies.

Bringing Brilliance to AI with the Power of Photonics

Drut's disaggregated photonic fabric allows for AI solutions to be deployed to a broader group of users by providing a higher rate of resource utilization in a smaller infrastructure footprint, versus a lower rate of utilization across a larger infrastructure footprint. Moving away from stacked interconnect hierarchies of electrical switches with a photonic fabric allows resources to be directly connected and grouped by nodes. This is how organizations will begin to deploy photonic platforms for AI clusters today. By decoupling resources from the server chassis with a fully disaggregated photonic fabric, users will add the ability to dynamically attach and detach resources to nodes as well as upgrade resources independently of CPU upgrades.

- Compose nodes by dynamically attaching and detaching resources without server reset or reboot
- Eliminate the power and data locality barriers
- Improve utilization of valuable resources such as GPUs
- Deploy modularity options for server and resource upgrades
- Lower Power – no electrical to optical conversions

At Drut, we make AI happen for everyone who is not a hyperscaler leader.

Bill of Materials (BOM) and Pricing Transparency

The sample bill of materials (BOMs) provided on the next page are guidelines, built on some assumptions, but they are accurate. We are trying to be transparent in showing what is and is not included in the BOMs.

What is Included?

- All Fabric Interface Cards (FICs) and Photonic Resource Units (PRUs)
- Photonic fabric to include switches (PXC), fiber cables and optical transceivers (All Active Optical System)
- 1-Year Hardware Support and Replacement
- 1-Year Support and RTU for the Drut Fabric Manager (DFM)
- 1-Year Support and RTU for the Drut Software Platform to include:
 - Bare Metal OS Deployment
 - Openstack Deployment for Cloud Computing
 - Kubernetes Deployment (K8s) for Containerized Applications
 - Management and Application Servers

What is not Included?

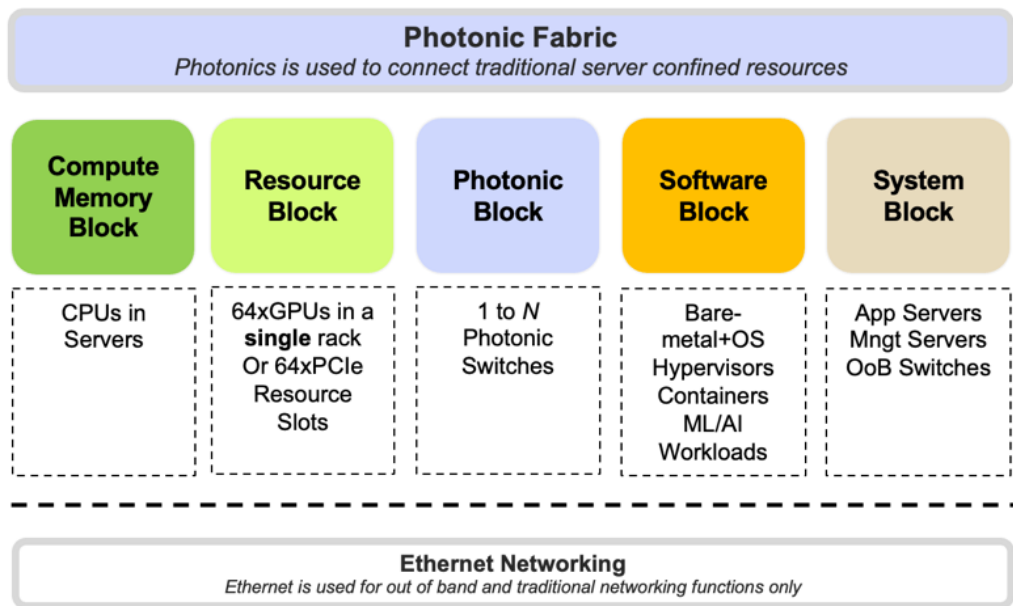
- Compute Nodes (Servers) but we provided a budget figure. Customers can provide their own or we can quote
- The sample BOMs do not include CEPH, but we can add that as part of the DSP and storage hardware
- GPUs but we provided a budget figure
- Installation, but we can provide a quote
- Managed Services, we can install, deploy and manage the complete infrastructure stack as a service in your data center, your colo or in Drut's colo location

What Are the Benefits?

Improved performance for AI/HPC/ML clusters is delivered by high-speed, all photonic interconnects, allowing for ultra-low latency and high bandwidth data transmission. This results in significantly improved performance, enabling faster data processing, reduced bottlenecks, and enhanced user experiences. The underlying hardware can be grouped and isolated around software workloads providing for better performance, higher utilization rates and security.

Scalability and Flexibility: The disaggregated architecture of photonic data centers enables modular scalability. Compute, storage, GPU, and networking resources can be independently upgraded or expanded, providing flexibility to meet changing business needs. This approach allows for efficient resource utilization and cost optimization for GPU intensive applications used by AI, ML and HPC clusters. The DynamicXcelerator takes a different approach to the needs of AI/ML scale out. It focuses on the totality of the system and how that system is used to transition data center infrastructure from the **Hypervisor Era** to the **Accelerator Era**. Drut is focused on optimizing the performance of the entire accelerator cluster across interconnect bandwidth, software integration, power/cooling, latency and total system performance. We begin by

breaking down the legacy static boxes into resource elements. These type 1 elements are hosted in the Photonic Resource Unit or PRU. A PRU is attached to a rate agnostic all photonic fabric via the Fabric Interface Card (FIC). CPUs in the servers are also attached to the photonic fabric via the Fabric Interface Card (FIC).



We use software to control and facilitate the how nodes (i.e. CPU + Resources) are constructed within the DX fabric.

Living with a DynamicXcelerator fabric is different from legacy data center architectures. With a DX fabric, the user can start small (a rack or two) and grow into a large cluster. The photonic fabric is rate agnostic, meaning when the optics are upgraded the fabric is still viable and is upgraded with the optics. The resources (GPU's, NVME, FPGA, etc.) in the PRU's can be upgraded as needed. Powering down a PRU, swapping the GPU's and powering up the PRU is all that is required to upgrade the GPU power of a cluster built with the DX fabric. Compute elements can be swapped as well. It is all software driven.

We still believe in power of ethernet as a scale out technology to the cloud, but for high performance cluster size deployments in the **Accelerator Era**, we believe in the power of a direct connect, dynamic photonic switching fabric. In a climate conscious age, we want to design accelerator clusters using all photonic switching to make a serious improvement in power and cooling budgets as well as performance and latency for efficiency gains.

Energy Efficiency: In terms of micro-processor design as well as large scale data centers, we have a hit the power ceiling. To overcome this barrier, we need to design data centers that leverage photons over electrons. Disaggregated photonic data centers leverage the energy-efficient properties of photonics technologies, minimizing power consumption and reducing operational costs. The decoupled architecture also allows for better thermal management, improving overall energy efficiency that are upgradable and sustainable.