



COMPARING COMPOSABLE OPTIONS

SIMPLE COMPOSITION

Ability to disaggregate storage and accelerators into pooling appliances and to bind them to individual servers to drive up utilization. This is especially handy to upgrade the capabilities of older legacy servers, as well as future-proof your data center and align your upgrade cycles to the life of the asset.

QUESTION	GIGAIO	ALTERNATIVE
Can you compose PCIe-attached components in a disaggregated pooling appliance?	We can compose any standard PCIe device connected to our FabreX Switch. Our fully managed pooling appliances support all devices that are PCIe-compliant, including GPUs, FPGAs, DPUs, network cards, and NVMe drives.	
Can you also compose the PCI devices inside the server?	Yes. Accelerators and/or storage located inside the server (or even in other servers) can be combined with other resources located in pooling appliances.	
What is the scale of simple composition?	Each server root port connected to our PCIe switch can support up to 256 devices in composition mode.	
Can you perform sub-device level composition?	Yes. In particular with NVMe-oF you can scale from fractions of a device to multiples of a device in any arbitrary size. For GPUs that support MIG/SRIOV, each fractional instance of the GPU is composable.	

EXTENDED COMPOSITION

Ability to have resources from one server be available to another server. The benefits include: increased scale, vastly improved granularity of sharing, and the ability to share resources with all the servers in the cluster at the same time.

QUESTION	GIGAIO	ALTERNATIVE
Can you scale out compute clusters over PCIe?	Yes, we run MPI, Libfabric and TCP traffic all natively across PCIe. This capability is unique to FabreX. See internode communication section below. With our newly introduced GigaPod and GigaCluster capabilities, for the first time PCIe fabrics can easily scale.	
What is the scale of extended composition?	Using NVMe-oF and GDR running natively over PCIe, the maximum number of server root ports that can be networked on FabreX (via NTB) is 128. And with 256 devices per server, this means we can connect over 32,000 devices.	
Can you compose PCIe-attached components inside another server over the same network?	Yes. We can compose server-hosted storage through NVMe-oF, GPUs via GDR, and access remote memory today. See explanation of extended composition below.	

EXTENDED COMPOSITION (continued)

Can you use the memory from another server?	Yes. Today you can configure one server chockfull of memory, and other thin servers can share that memory. With CXL, the access performance will improve and the memory will become coherent.	
Can the memory in the other server be Optane pmem instead of DRAM?	Yes.	

INTER-NODE COMMUNICATION

The traditional internode protocols (MPI, Libfabric, TCP/IP) over the PCIe data plane, without suffering the expense and administrative burden of a second network like Ethernet or InfiniBand in the rack.

QUESTION	GIGAIO	ALTERNATIVE
Can you network servers together over PCIe, in addition to composing resources?	Yes. All traffic that would normally run within the cluster can run natively on FabreX including MPI, Libfabric, TCP/IP, and the extended composition protocols of NVMe-oF and GDR. Note this doesn't preclude the use of another network like Ethernet or InfiniBand, but those are not required.	
Can you compose server-to-server over PCIe?	Yes. This capability is unique to FabreX.	
Do you support MPI over PCIe?	Yes. This capability is unique to FabreX.	
Can you run a single PCIe network throughout the rack?	Yes. This capability is unique to FabreX.	
Can you support RDMA over PCIe?	Yes. We also support P2P traffic using DMA between all end points across all servers, and between all servers network-wide since every FabreX port has a DMA engine built in.	
Do you require another network in the rack?	No. FabreX is a routable PCIe network fabric that can connect servers together over PCIe, saving the cost, administration hassle and complexity of a second set of switches (Ethernet or InfiniBand) in the rack, while improving the security profile of the rack.	
Can you use Ethernet or InfiniBand to connect servers?	Yes, but you don't have to. You can avoid the expense and administrative burden of setting up a second network by running all traffic over FabreX.	

QUESTION	GIGAIO	ALTERNATIVE
Do you require a per node host software license?	No. GigaIO's server node S/W is open source without either a one-time or a recurring expense.	
Do you require a management node?	No dedicated management node required. The FabreX Fabric Manager (FFM) runs natively on the FabreX switch(es) regardless of the topology you select.	
Do you require purchasing proprietary orchestration software?	No. FabreX is a fundamentally OPEN platform and has by far the widest support from the industry's leading orchestration and composition software companies. Customers can choose from an array of off-the-shelf orchestration options, or save the cost if they have developed their own package.	
Do you run off-the-shelf composition software?	Yes, the largest choice of any composable vendor: Quali Cloudshell, Supermicro Supercloud Composer, DefineTech Slurm, CIQ Singularity and Fuzzball, Bright Computing Cluster Manager (Fall 2021).	
Does your orchestration software have full enterprise-class features like access control, advanced security, compliance and provisioning?	Yes - our off-the-shelf composition software options are developed by leading vendors of enterprise solutions, names you trust and likely already use.	
Are you an open platform with industry standard Redfish APIs?	Yes. We do not use proprietary REST APIs, but instead are committed to supporting open standards to protect you from vendor lock in. Our support and documentation make integration swift and efficient.	
Can you use automation tools like Ansible, Puppet, Chef, Robot Framework and others?	Yes. Very easy to do with our CLI tools.	

TOPOLOGIES AND ENVIRONMENTS SUPPORTED

QUESTION	GIGAIO	ALTERNATIVE
Can you daisy chain PCIe switches?	Yes. Daisy chain, cascade, interlink modes are all supported allowing you to create whatever network topology you desire.	
Can you create leaf and spine PCIe topologies?	Yes	
Are VDI + 3D Viz use cases supported?	Yes	
Can existing servers be repurposed to add the latest and greatest GPUs or other accelerators?	Yes. FabreX is unique in offering you two capabilities to upgrade existing systems. First, you can use a pooling appliance and compose GPUs to a server that is often incapable from a space, power and cooling perspective of using GPUs. Second, with our FabreX-native GDR capabilities, you can easily group servers and their GPUs to run larger jobs with the combined power of all the GPUs, all at PCIe performance. So with FabreX it is no longer necessary to throw out these older servers - they are now easy to repurpose and bring back to life.	
Can the latest in storage technology be added to older servers as well?	Yes. With FabreX's native NVMe-oF, every server, no matter its vintage, can gain access to the very latest and best storage - from Intel Optane to any NVMe Flash drive or array, and do so at full native PCIe performance.	
What is the recommended Compute Unit or Pod?	GigaIO supports three levels of systems - GigaCell, GigaPod and GigaCluster. A GigaCell is typically about a half rack and can include compute, acceleration, storage, memory and networking in any combination. Up to 6 GigaCells create a GigaPod, and up to 6 GigaPods can be combined to create a GigaCluster.	
Can you provide telemetry information from all PCI enclosures and endpoints?	Yes, our GPU pooling appliance is fully managed and has independent SMBus access to each card in the enclosure - thus insuring complete information no matter where the resource is physically sitting.	
Is the hardware environment fully managed with auto-discovery and full telemetry?	Yes. FabreX insures end-to-end information and control for all resources in the system wherever they are located.	