



LMX COMPOSABLE

APPLICATION DEFINED INFRASTRUCTURE

HPC HAS EVOLVED SIGNIFICANTLY OVER THE PAST FEW YEARS. WITH THE EXPANDING ADOPTION OF AI AND DATA ANALYTICS, WE SEE MORE AND MORE DIVERSE WORKLOADS AND APPLICATIONS NOW BEING RUN ON TRADITIONAL HPC INFRASTRUCTURE. WITH THE BROADENING DEMAND OF THESE VARIED WORKLOADS, THE SOFTWARE AND APPLICATION STACKS REQUIRED TO RUN THEM EFFICIENTLY HAVE BECOME INCREASINGLY COMPLEX.

WHAT IS COMPOSABLE HPC?

Infrastructure needs to be flexible and agile to run the appropriate stack to support the changing workload demands, whether that's SLURM for traditional MPI workloads, Kubernetes for containerized services or Spark for analytics workloads, infrastructure needs to be dynamically reconfigurable to react to the changing demands over time.

With our LMX Composable offering, we've built a platform that offers the flexibility of Cloud software (API driven, scalable and easily reconfigurable) as well as the critical application stacks needed to run modern and constantly evolving workloads.




Fig 1. LMX Cloud delivers a dynamic Infrastructure that can adapt to your changing workloads

LMX CLOUD AND GIGAIO FABREX

In partnership with GigaIO, Define Tech has taken the next step towards a true software defined supercomputer or application-defined infrastructure by enabling hardware composability for HPC workloads. Along with complex software stacks being dynamically configurable on our platform, we can now reconfigure and disaggregate / compose servers to respond to the demands of specific workloads by allocating compute acceleration, memory or IO to systems where and when they are required.

HOW IT WORKS

Historically, a traditional HPC infrastructure would be heterogeneous, with some servers having large memory, GPUs or NVMe resources. The infrastructure admin team would need to estimate how much of these resources would be consumed over the coming years. These resources would have been provided in static configurations within the servers which would invariably result in underutilization at times.

With composable infrastructure the guess work is unnecessary. Compute infrastructure can be provisioned based on the average usage (rather than the peak estimate) with additional resources pooled across the fabric and made available to nodes as and when workloads need them. This pool of PCI attachable resources can be scaled and grow on the fly without any disruption to service and it allows organizations to scale in an appropriate manner adding only the resources required (GPU, CPU, memory, IO etc) see Fig. 1.

With the additional integration of SLURM, LMX Cloud can allow resources to be requested via a job submission script.

THE BENEFITS

- Pool expensive resources to increase their utilization, and drive down TCO
- Allow organizations to scale infrastructure based on demands for specific resources (not having to add more servers every time)
- Seamless upgrade path for PCI devices (drop in latest PCI accelerators without any other hardware modifications required)
- Low latency access to resources (performance comparable to locally attached devices)
- API driven and extensible
- Scale compute, memory or IO (or all)

COMPOSABLE HPC REPRESENTS THE NEXT MAJOR SHIFT IN HOW USERS CONSUME RESOURCES. THE SOFTWARE DEFINED SUPER-COMPUTER IS HERE TODAY.

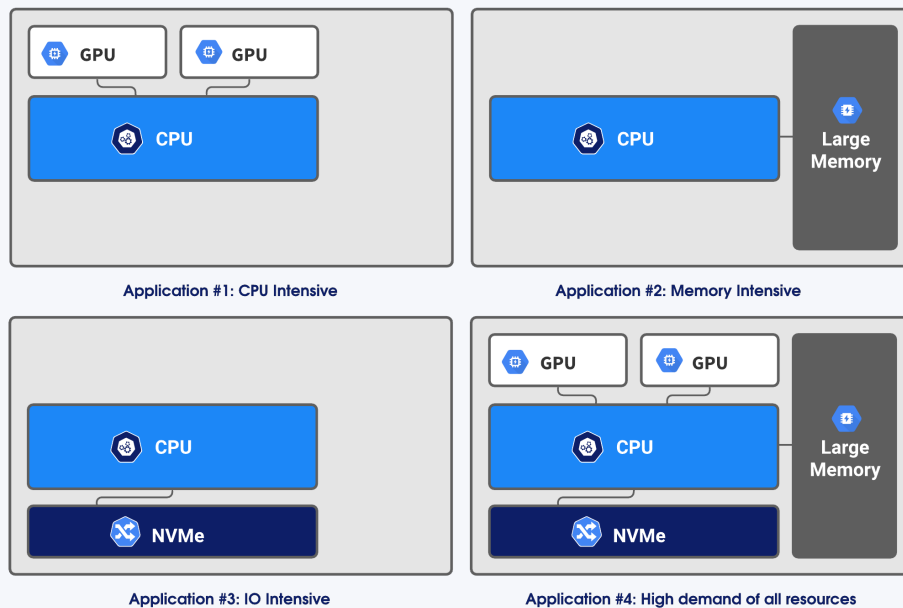


Fig 1. PCI attached resources are assigned and released depending on workload requirements

For pricing or to discuss your requirements:

www.define-technology.com | +44 (0)20 3034 5550 | info@define-technology.com