

# Gryf: The First Ever Suitcase-sized AI Supercomputer

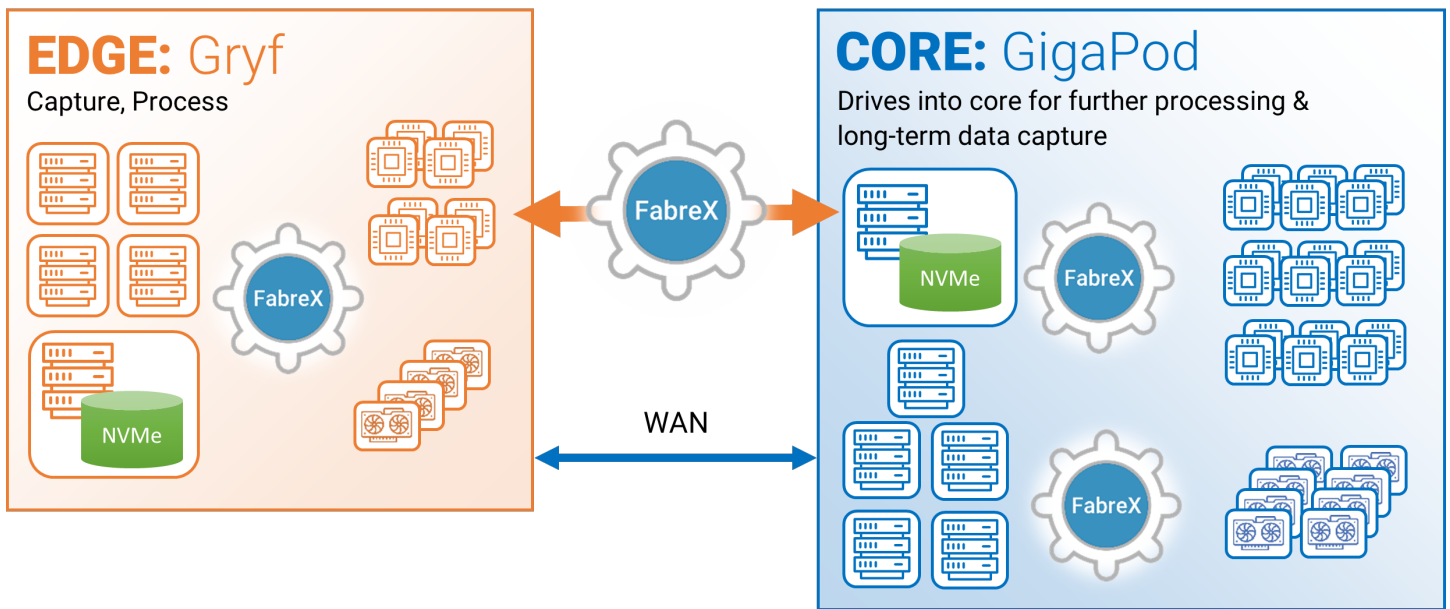
## SOLUTION BRIEF



Co-designed by GigaIO and SourceCode for portability and scalability, Gryf sets a new standard for on-demand configurability in the field. Powered by GigaIO's AI memory fabric, FabreX™, it can dynamically deploy any mission application, for actionable real-time intelligence – anywhere.

It is true that data is generated everywhere, but this phrase has a different meaning for those tasked to protect our country against all threats. Intelligence is only as good as it is timely; data from drones, satellite images, and sensors hold the key to staying one step ahead of our adversaries. Yet the existing approach of using static infrastructure and moving data to compute is no longer effective or timely.

Current efforts to deploy infrastructure at the edge have been plagued by network inefficiencies, lack of process consistency, time lag to critical updates, and disjointed workflows. Now GigaIO's patented AI memory fabric, FabreX, enables data processing in real time, and facilitates actionable intelligence anywhere.



The diagram above shows the ruggedized Gryf, which can be deployed at the edge (left), and a GigaPod, located back in the core datacenter (right). The unique value of GigaIO's AI memory fabric, FabreX, is its ability to combine these into one seamless and dynamic environment.

## Gryf includes disaggregated GPUs, NVMe storage, and CPU/memory. It is designed to:

- Expedite data analytics at the edge. Configurations can be optimized for each application's workload demands.
- Make the captured data and data analysis output at the edge available over GigaIO's AI memory fabric, FabreX. Workflows are optimized where the data is located.
- Combine all resources once Gryf is in the same facility as GigaPod and make them dynamically available to any application. The NVMe drives in the Gryf become available for processes running in GigaPod without moving data.
- Move data over Ethernet from the edge to the core and vice versa, if bringing the Gryf to the core is not an option.



Deploying this GigaIO Edge-to-Core solution reduces time to results and eliminates stagnant intelligence. When it comes to keeping adversaries at the gate, it is critical to have the very best, most time-sensitive intelligence. This is only possible when the mission's applications can run at the edge and process data in real time in order to get actionable results. Optimize workflows by reducing the time required to get relevant intelligence into the hands of those in the field.

This solution has been designed to deliver maximum return on investment. With Gryf and GigaPod, numerous possibilities for server configurations are available without forklift upgrades, and with a smaller footprint and lower power draw, depending on application need.



|                         | Gryf                       | GigaPod               |
|-------------------------|----------------------------|-----------------------|
| <b>CPU Cores</b>        | Up to 384                  | Up to 640             |
| <b>Memory</b>           | Up to 2TB                  | Up to 4.6TB           |
| <b>Accelerator</b>      | L40S-48GB                  | Up to 16 L40S-48GB    |
| <b>Storage</b>          | Optional sleds up to 120TB | Up to 307TB           |
| <b>Network</b>          | 100/25/10 GbE and OOB      | 100/25/10 GbE and OOB |
| <b>AI Memory Fabric</b> | 256/128 Gb/s               | Multiple 128Gb/s      |
| <b>Composable</b>       | ✓                          | ✓                     |
| <b>Weight</b>           | 15-65 lbs                  | 200-800 lbs           |
| <b>Dimensions</b>       | 9" x 14" x 22"             | 24U Rack              |




# Choose Your GigaPod Configuration

Back at home base, connect Gryf to your data center GigaPod or GigaCluster for post-field intelligence analysis.

## GigaPod

The perfect entry point and preconfigured for easy deployment.



SUPPORTS UP TO:



640 CPU Cores



4.6TB Memory



307TB Storage




16 Accelerators

- 1 GigaPod kit – includes FabreX™ composable switches, network adapter cards, and cables
- Up to 2 GigaIO Accelerator Pooling Appliances – each up to 8 mix-and-match GPUs
- Up to 2 storage servers – each up to 154TB
- Up to 4 compute nodes – pick the processor – dual AMD EPYC™ 7713, 75F3, or 7543 – each with up to 1TB memory and 128 cores
- NVIDIA AI Enterprise Essentials

## GigaCluster

Expand your system to increase compute and storage performance.



SUPPORTS UP TO:



1,280 CPU Cores



9.2TB Memory



614TB Storage




32 Accelerators

- 1 GigaCluster kit – includes FabreX™ composable switches, network cards, cables, and expansion pack
- Up to 4 GigaIO Accelerator Pooling Appliances – each up to 8 mix-and-match GPUs
- Up to 4 storage servers – each up to 154TB
- Up to 8 compute nodes – pick the processor – dual AMD EPYC™ 7713, 75F3, or 7543 – each with up to 1TB memory and 128 cores
- NVIDIA AI Enterprise Essentials

## 3Pod GigaCluster

When you need the ultimate in compute power.



SUPPORTS UP TO:



1,920 CPU Cores



13.8TB Memory



922TB Storage



48 Accelerators

- 3Pod GigaCluster kit – includes FabreX™ composable switches, network adapter cards, cables, and expansion pack
- Up to 6 GigaIO Accelerator Pooling Appliances – each up to 8 mix-and-match GPUs
- Up to 6 storage servers – each up to 154TB
- Up to 12 compute nodes – pick the processor – dual AMD EPYC™ 7713, 75F3, or 7543 – each with up to 1TB memory and 128 cores
- NVIDIA AI Enterprise Essentials

GigaIO's AI memory fabric, FabreX, unifies edge-to-core infrastructures to dynamically deploy any mission application, for actionable real-time intelligence – anywhere.

Learn more at [gigaio.com/products/gryf](https://gigaio.com/products/gryf)