



CnuasSHMEM, Datasheet

Partitioned global address space library over the CnuasLink fabric

Document DS-CNU-018 • **Revision** A • **Issued** 03 August 2026 • **Status** Preview, not implemented

Item	Value
Part	CnuasSHMEM
Type	Partitioned global address space communication library
Models	NVIDIA NVSHMEM
Status	Preview, specified but not implemented
Repo	PacketFive/CnuasGPU

Preview datasheet

CnuasSHMEM carries **no implementation** in the tree. No symmetric heap, no one sided operations and no signalling exist. This datasheet records the specified design.

1. Overview

CnuasSHMEM is the intended one sided communication library for CnuasGPU. Where CnuasCCL provides collective operations that every rank enters together, CnuasSHMEM lets one accelerator read and write another accelerator's memory without the remote side taking part, which suits irregular and latency sensitive access patterns.

2. Specified operations

Group	Operations
One sided data movement	Put, get
Atomics	Remote read modify write
Signalling	Set and wait on a remote flag

3. Transport

Item	Specified value
Primary transport	CnuasLink, accelerator to accelerator
Address space	Symmetric heap across participating accelerators, not specified in detail
Underlying path	The accelerator fabric rings described in Kernel Modules section 5

4. Dependencies

Requirement	Status
CnuasLink transport	Implemented



Requirement	Status
Remote memory addressing model	Not specified
Device side calls from within a kernel	Not implemented, needs CnuasCC
Atomic operation support on the fabric	Not implemented

Note that one sided operations issued from inside a kernel, which is the usual NVSHMEM pattern, need a device compiler. Host initiated operations could land earlier, and that split has not yet been decided.

5. Interim position

Applications needing one sided semantics today should use RDMA verbs directly through the verbs provider, which supports remote write, remote read and atomic operations over the network fabric.

6. Revision history

Revision	Notes
A	First publication as a preview. Operation list taken from the CnuasGPU design document, section 7.6.

PacketFive, Packet Five Networks Ltd., Dublin, Ireland.

Cnuas Virtual AI/HPC Infrastructure is published at <https://github.com/PacketFive/cnuas> under the Apache License 2.0; read it at <https://github.com/PacketFive/cnuas/blob/main/LICENSE>.

Cnuas Virtual AI/HPC Infrastructure is emulation software. It is not affiliated with, endorsed by, or derived from any hardware vendor, and every device it models is a software artefact. The work is published by its authors in a personal capacity and is not sponsored or endorsed by any employer.

Specifications describe the referenced revision of the software and may change without notice. Contact info@packetfive.com.