FOR AI, HPC AND CONVERGED CLUSTERS

Technical evaluation worksheet

The questions to answer before you compare fabrics.

The interconnect is part of most HPC-AI system decisions. To choose it well, your team needs to agree on what it needs before testing starts. Answer these questions first, so you can compare options faster, ask vendors sharper questions and avoid testing the wrong things.

Key takeaways

- **Start with the business goal:** Know what the cluster must deliver and how you will measure it.
- **Know your workload:** Message sizes, collectives and scale shape fabric needs more than port speed does.
- **Know your limits:** Power, cooling, budget, timeline and existing systems narrow the options before any test.
- **Agree on success first:** Decide what "better" means before you see any vendor results.

Five areas to work through

Work through five areas in order, from what the fabric must carry to what has to be true on day one. Type your answers into the form fields, or print the worksheet and write them in.

1. **Workload discovery**
2. **Architecture fit**
3. **Performance goals**
4. **Integration**
5. **Deployment readiness**

Answer it as a team

Bring your platform lead, network and systems architects, operations and procurement. Gaps you find here become the agenda for your first vendor conversation.

92% of HPC-AI buyers make the interconnect part of the system acquisition decision. Source: Intersect360 Research, 2026.

1 Workload discovery: what must the fabric carry?

Your workloads decide which network traits matter. A cluster that runs many small, frequent messages needs different strengths than one that streams large files. Start with why the project exists, then describe what will run on it.

A. Why now?

What outcome must this cluster deliver for the business or mission?

What is driving it: ☐ New AI or HPC workloads ☐ Growth in demand ☐ Current system end of life ☐ Jobs run too slowly ☐ Power or cost limits

B. What will run on it?

☐ Mostly HPC ☐ Mostly AI ☐ Both HPC and AI

Name your main applications or models, if you know them.

Application or model	Share of cluster time
1.	%
2.	%
3.	%

C. Do you know how your jobs communicate between nodes?

You don't need exact numbers yet. Get a rough sense of whether jobs trade many small messages, run large collectives or move big blocks of data.

☐ We have profiled our main jobs ☐ We have a rough idea ☐ Not yet

What seems to slow jobs down today?

WHY IT MATTERS

If jobs trade many small messages or run frequent collectives, message rate and latency will shape results more than port speed. That is the traffic the CN6000 Omni-Path SuperNIC is built for. You will describe each workload in detail later, with the RFI/RFP checklist.

Area 1 status: ☐ **Answered** ☐ **Partly answered** ☐ **Not yet** Owner:

2 Architecture fit: what must the design allow?

Fabric choices set your cost and risk for years. Agree on how far the cluster must grow, which network rules apply and how much one failure can take down.

A. How big, and how fast will it grow?

Endpoints at first build: _____ Endpoints in 3 years: _____

Per node: ☐ 400 Gb/s ☐ 800 Gb/s ☐ Not sure Build plan: ☐ All at once ☐ In phases

Oversubscription OK? ☐ No ☐ Yes, ratio: _____

B. Which network rules apply?

☐ Open to the best-fit fabric ☐ Must run on standards-based Ethernet

☐ Must reuse an existing CN5000 or Omni-Path fabric ☐ Storage and checkpoint must share the fabric

Reason for any rule: _____

C. How much can a failure take down?

Longest job run time: _____ Largest fault domain you accept: _____

How often you checkpoint: _____

D. Root questions to answer for any fabric

- **Congestion under load:** How does it handle incast and hot spots at full scale, not only in a quiet test?
- **Rerouting:** Can the fabric reroute flows around congestion or a failed link or switch?
- **Open software:** Which open interfaces does it support, such as libfabric, MPI, NCCL and RCCL? What would your applications need to change to use new features?

E. Evaluate the CN6000 Omni-Path SuperNIC and Ethernet SuperNIC

	Omni-Path SuperNIC	Ethernet SuperNIC
Best for	Large-scale AI training and HPC codes bound by latency or message rate	Fine-tuning and inference on Ethernet networks
Switching	CN6000 Omni-Path Switch, or existing CN5000 switching	Standards-based Ethernet switches
Lossless traffic	Credit-based flow control and link-level retry	Priority Flow Control
Software	Open-source OPX on libfabric	Standard RoCEv2; Ultra Ethernet compliant

WHY IT MATTERS

Need 800G per node? Compare designs by the adapters, PCIe slots, cables and switches each one takes. A CN6000 Omni-Path SuperNIC carries two independent 400 Gb/s links, so one adapter and one cable can reach two switch planes. Multi-plane designs repeat the switch tree in each plane, so ask every vendor for total switch and cable counts at your scale.

Area 2 status: ☐ Answered ☐ Partly answered ☐ Not yet Owner: _____

3 Performance goals: what will "better" mean?

Agree on what success looks like before you see any vendor results. You don't need targets yet. You need a shared view of what matters most, so no single headline number decides for you.

A. Which outcomes matter most? (pick up to three)

- ☐ Faster time to result
- ☐ Jobs that keep scaling as they grow
- ☐ Lower power per completed job
- ☐ Higher accelerator utilization
- ☐ More tokens per second, or lower cost per token
- ☐ Steady run times, with fewer slow outliers

B. Where do you think time goes today?

☐ Computing ☐ Waiting on the network ☐ Reading and writing data ☐ We don't know yet

C. How will you compare options?

Which jobs or models will you use as your yardstick?

Evidence your team will trust: ☐ Our own codes ☐ Standard benchmarks ☐ OEM or partner results ☐ Published case studies

WHY IT MATTERS

For jobs heavy with small messages and collectives, message rate and latency often decide time to result. You will turn your goals into specific targets later, with the RFI/RFP checklist. For reference, here are the published CN6000 figures.

800G	1.6B[1]	<1 µs[1]
Bandwidth per SuperNIC	Messages per second, Omni-Path	MPI latency, Omni-Path

Area 3 status: ☐ **Answered** ☐ **Partly answered** ☐ **Not yet** Owner: _____

1. OSU microbenchmarks between two nodes connected through one switch. Test configuration: AMD EPYC 9006 (Venice); OPX 14.0.0.27; Open MPI. Results vary with configuration, workload and topology.

4 Integration: what must it work with?

Deployment risk often hides at the interfaces. Parts can each pass a test alone and still fail together. List the software, systems and tools the fabric must fit, then judge every option against your real environment.

A. Which software do your jobs depend on?

- ☐ Open MPI
- ☐ MPICH
- ☐ MVAPICH2
- ☐ Intel MPI
- ☐ OpenSHMEM
- ☐ NCCL
- ☐ RCCL
- ☐ UCX
- ☐ PyTorch
- ☐ GPUDirect, NVIDIA
- ☐ GPUDirect, AMD
- Other: _____

B. What is already in place?

Operating systems and versions: _____

Scheduler and containers: _____

CPUs and accelerators: _____

Server platforms (OEM, model): _____

Current fabric: ☐ CN5000 ☐ Omni-Path 100 ☐ Other

Storage: ☐ Parallel file system ☐ NVMe-oF

C. What must management and security cover?

- ☐ Redfish, PLDM and MCTP
- ☐ Fabric telemetry
- ☐ In-band and out-of-band
- ☐ Hardware root of trust
- ☐ SPDM device attestation
- ☐ Existing monitoring tools

D. Root questions to answer with your server and fabric vendors

- **Qualified combinations:** Which server, firmware, BIOS, OS, driver and fabric software versions were tested together?
- **Management coverage:** Does validation include health reporting, monitoring and telemetry?
- **Published proof:** Is there a reference architecture or solution brief for your platform?
- **Upkeep:** How does support stay current as firmware, operating systems and processors change?
- **Existing fabric:** Can new adapters join the fabric you run today, or is it a full replacement?

WHY IT MATTERS

Open, shared software lowers integration risk. CN5000 and CN6000 SuperNICs use the same open-source OPX driver, so both can share one Omni-Path fabric, and MPI, OpenSHMEM, NCCL and RCCL run unchanged. Cornelis also works with OEM partners to qualify server configurations before they reach your site: more than 50 in the past year.

Area 4 status: ☐ **Answered** ☐ **Partly answered** ☐ **Not yet** Owner: _____

5 Deployment readiness: what has to be true on day one?

Your facility, budget, schedule and people limit the options as much as performance does. Answer these early, so you only evaluate designs you can install, power and run.

Area	Question to answer, with your answer	Known	Need to find out	Owner
Power	How much power per rack can the network use?	☐	☐	
Cooling	Which cooling can your facility support: air, conduction or liquid?	☐	☐	
Cabling	What cable lengths and paths does the room allow?	☐	☐	
Rack space	Where can switches sit, and is there room to grow?	☐	☐	
Budget	What is the budget, and which buying process applies (OEM or integrator)?	☐	☐	
Timeline	When must the system be in production?	☐	☐	
Migration	What must keep running while you cut over?	☐	☐	
Operations	Who will run the fabric, and what skills do they have today?	☐	☐	
Support	What support level, spares and response times do you need?	☐	☐	
Expansion	What is the next growth step, and when?	☐	☐	

WHY IT MATTERS

Expansion is where fabric choices pay off or cost you. CN6000 Omni-Path SuperNICs can connect to existing CN5000 switching and share the same open-source OPX software, so teams add bandwidth without learning a new protocol.

Area 5 status: ☐ **Answered** ☐ **Partly answered** ☐ **Not yet** Owner: _____

Ready to evaluate? Close the open questions first

Copy each area status here. When most areas are answered, start comparing options. Bring the areas marked "not yet" to your first conversation with Cornelis.

Area	Answered	Partly	Not yet	Biggest open question
1. Workload discovery	☐	☐	☐	
2. Architecture fit	☐	☐	☐	
3. Performance goals	☐	☐	☐	
4. Integration	☐	☐	☐	
5. Deployment readiness	☐	☐	☐	

Choose your next step

IF YOUR WORKLOAD PROFILE IS UNCLEAR

Talk to a performance expert

Find where your jobs spend time on the network, and which results to target.

IF THE DESIGN IS UNCLEAR

Schedule a design session

Work through scale, topology, cooling, cabling and phasing with Cornelis specialists.

IF YOU ARE READY TO TEST

Reserve a test drive

Run your reference jobs on partner-hosted CN5000 clusters, with Cornelis engineers to help.

IF YOU ARE WRITING AN RFI OR RFP

Use the RFI/RFP checklist

Turn these answers into requirements you can compare across vendors, or contact Sales.

Further reading

- **Measure communication bottlenecks:** How to tell whether the network limits accelerator use
- **Reduce deployment risk:** Beyond the benchmark: OEM-validated AI and HPC fabrics
- **Design for scale:** A network-centric approach to cluster architecture
- **Market research:** Intersect360 Research, State of the HPC-AI Market: Networking and Interconnects
- **Open standards:** Ultra Ethernet Consortium and OpenFabrics libfabric

Email your worksheet before your call

Email it to **sales@cornelisnetworks.com**, and a Cornelis solutions architect will review it with you. Or keep it offline and bring it to the call.

Email worksheet

Save your filled-in PDF first, then attach it to your email.