

INVENTEC K905G6 + Cornelis CN5000 Omni-Path® 400G Redefining High Performance Computing Accelerate Innovation with Performance, Scalability, and Cost Efficiency

In today's rapidly evolving landscape of High-Performance Computing (HPC) and Artificial Intelligence (AI), enterprises demand more from their infrastructure: faster simulations, lower latency, and smarter capital allocation. INVENTEC and Cornelis have joined forces to deliver a breakthrough architectural solution: the high-performance K905G6 Series server integrated with the next-generation CN5000 Omni-Path 400G high performance end to end networking solution. This strategic collaboration delivers competitive performance empowering users to execute complex workloads with greater speed and efficiency while maximizing Return on Investment (ROI).



Figure 1 - INVENTEC K905G6 and Cornelis CN5000 Omni-Path 400G End to End Networking Solution

Joint Solution Overview

K905G6

The K905G6 is a next-generation high-density 2U server solution, engineered to deliver outstanding performance, scalability, and efficiency for modern data centers. Integrating four powerful, hot-swappable dual-processor compute nodes within a compact 2U chassis, it is ideal for HPC and AI workloads.



Figure 2 - INVENTEC K905G6 node and Chassis with 4 Servers

Feature/Model	K905G6 (AMD)
Processor	Dual-socket AMD EPYC™ 9004/9005 Series (Genoa/Turin)
Max Cores per Node	256
Memory Slots per Node	24 DDR5 DIMMs (1DPC)
Max Memory per Node	Up to 6TB DDR5 @ 6400 MT/s
Max TDP per CPU	400W (liquid cooling)/240W (air cooling)
Expansion Slots	2x PCIe Gen5 x16 FHFL (2U) 2x PCIe Gen5 x16 LP 1x OCP 3.0 Slot 1x M.2 Mezz
Storage per Node	2x U.2 NVMe (1U)/4x U.2 NVMe (2U)
Power Supply	4x 2400W CRPS PSU (Titanium) N+1 redundant
Cooling	Air (up to 240W TDP) Liquid (up to 400W TDP)
Chassis/Form Factor	2U4Nodes – 8 CPUs 2U2Nodes – 4 CPUs, 4 GPUs

Cornelis CN5000 Omni-Path 400G

Cornelis CN5000 Omni-Path is an end-to-end high performance networking solution designed for maximum efficiency and scalability. At its core, the CN5000 400G SuperNIC delivers ultra-low latency and high message rate, purpose-built to empower CPU- and GPU-intensive workloads in HPC environments. From host to fabric to switch, CN5000 ensures congestion-free, predictable performance at scale, making it the premier choice for accelerating next-generation HPC and AI workloads.

Key Technical Advantages

- **Lossless Data Transmission:** Uses hardware-based, credit-based flow control to eliminate packet drops and high-latency retries, ensuring predictable, high performance for massive clusters.
- **Dynamic Lane Scaling & Link-Level Replay:** CN5000 tolerates cable/transceiver failures and downgrades links gracefully without dropping the connection. At the same time, it locally detects bit errors via link-level replay – avoiding disruptive retransmissions that are expensive and negatively impact performance.
- **Fine-Grained Adaptive Routing:** Leverages real-time fabric telemetry to dynamically steer traffic around hotspots, drastically reducing “tail latency” (jitter) for tightly coupled workloads like CFD.
- **Open-Standards Interoperability:** Built on the OpenFabrics Alliance (OFA) Libfabric community project, this vendor-neutral stack ensures seamless compatibility across major CPU and GPU architectures without proprietary lock-in.



Figure 3 - Cornelis CN5000 Omni-Path 400G Portfolio

Joint Solution Benefits

To demonstrate the benefits of the joint solution, we measured performance on the K905G6 equipped with Cornelis CN5000 Omni-Path and compared to similarly configured nodes connected with NDR InfiniBand (IB) from NVIDIA.

Strong Foundational Performance

High performance computing workloads demand infrastructure which enables excellent workload characteristics not only in terms of absolute performance but also regarding performance scalability to maximize ROI. It is therefore critical that a high-performance fabric delivers exceptional foundational performance characteristics; high message rates, high effective bandwidth and low latency, to unlock real world application performance and avoid the scaling limitations associated with other networks.

Cornelis CN5000 Omni-Path demonstrates leading performance across several industry standard microbenchmarks. Figure 4 demonstrates the advantage CN5000 has over NDR IB in terms of MPI Message Rate, measured between two K905G6 servers through a CN5000 switch. As the number of core pairs increases, we see superior Message Rate scaling up to the maximum of 412M messages per second, which translates to over 2X the max message rate of NDR IB.

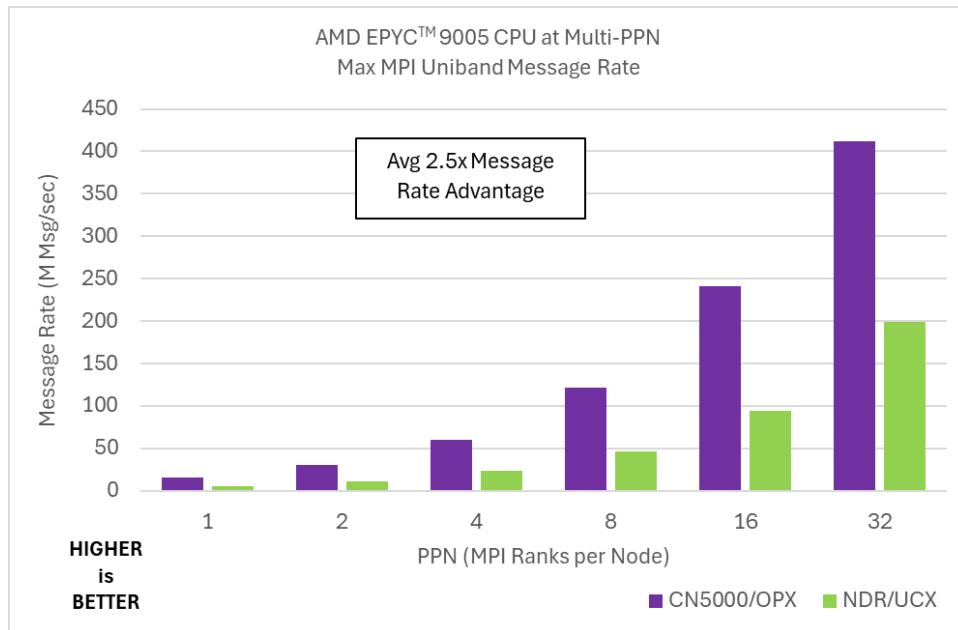


Figure 4 - Message rate (unidirectional) vs core pairs for Cornelis CN5000 and NDR IB

Figure 5 demonstrates the effect of CN5000’s higher message injection rate resulting in higher effective throughput across the bandwidth curve up until we see convergence at the maximum line rate. For HPC, message sizes typically drop as workloads are scaled to more nodes, where CN5000 has an average 1.7x bandwidth advantage over NDR IB.

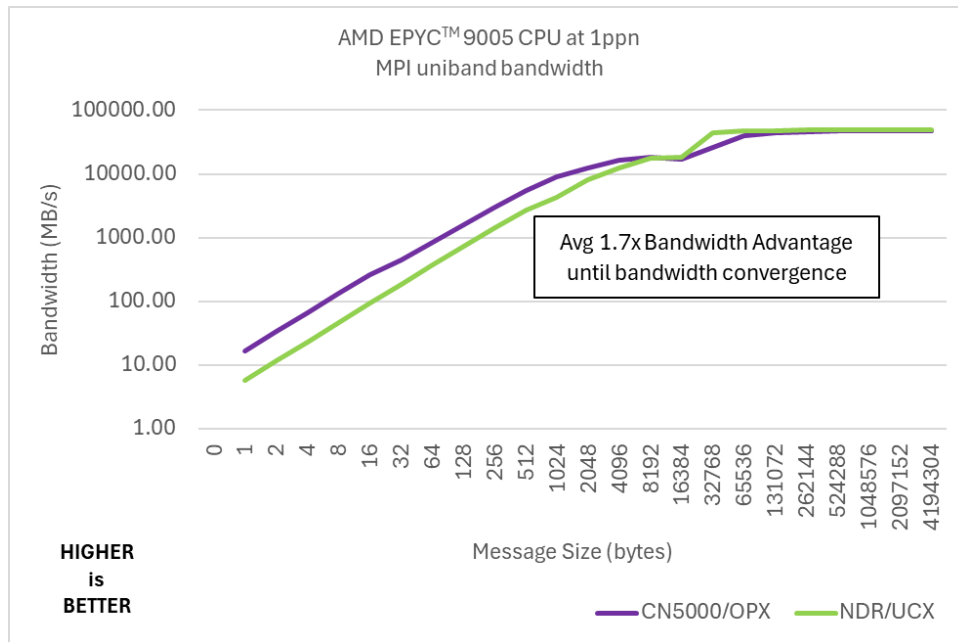


Figure 5 - Bandwidth (unidirectional) vs message size for Cornelis CN5000 and NDR IB

Finally, CN5000’s latency advantage is shown below in Figure 6 in an MPI AllReduce collective operation (used in many HPC and AI training workloads), consistently matching or demonstrating lower latency across the message size range than NDR IB. CN5000 is 29% faster on average across all message sizes.

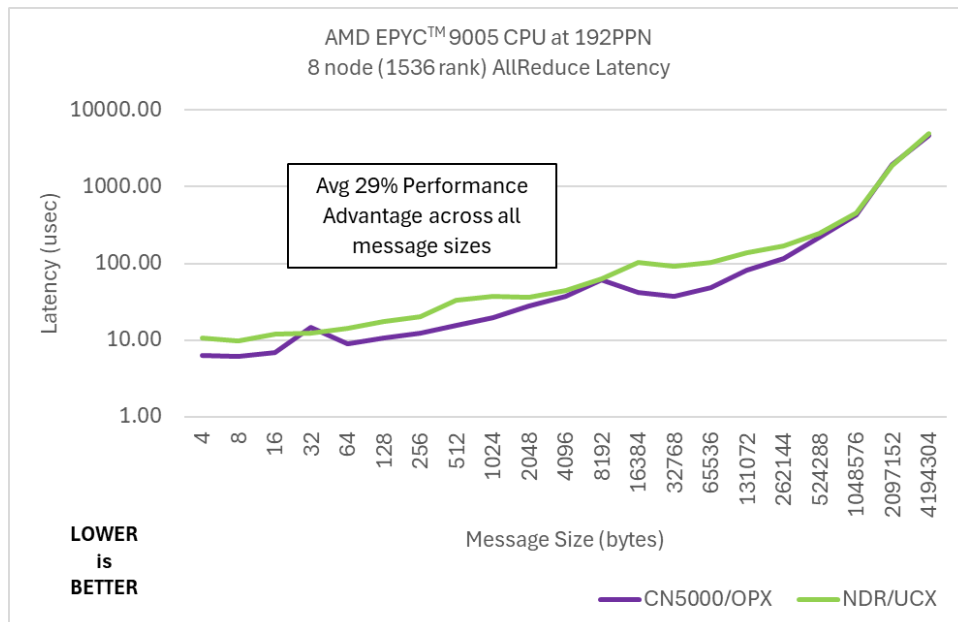


Figure 6 - Latency (AllReduce) vs message size for Cornelis CN5000 and NDR IB

Scaling Efficiency

Scaling Efficiency measures how effectively a cluster performs as the number of compute nodes increases. Ideally, increasing the number of nodes should deliver a corresponding improvement in workload performance, but in reality, most HPC applications will exhibit plateau or inflection point on the scaling curve and face diminishing returns, often due to communication overhead and synchronization delays. Maintaining high scaling efficiency is therefore a critical indicator of how well the system’s network, compute architecture, and software stack work together under real-world workloads.

We evaluated the scaling performance of All-Reduce MPI collective, which has relevance to both HPC and AI training workloads. In the test results, the K905G6 equipped with Cornelis CN5000 fabric delivers superior AllReduce performance at 192ppn (MPI ranks per node) when scaling from 768 ranks to 3072 ranks compared to similarly configured nodes connected with NDR IB. Additionally, it was observed that CN5000 demonstrates superior AllReduce performance at double the rank count compared to NDR IB. This outstanding performance demonstrates that the K905G6 and CN5000 platform minimizes typical communication overhead, demonstrating superior AllReduce collective performance resulting in shorter time to results.

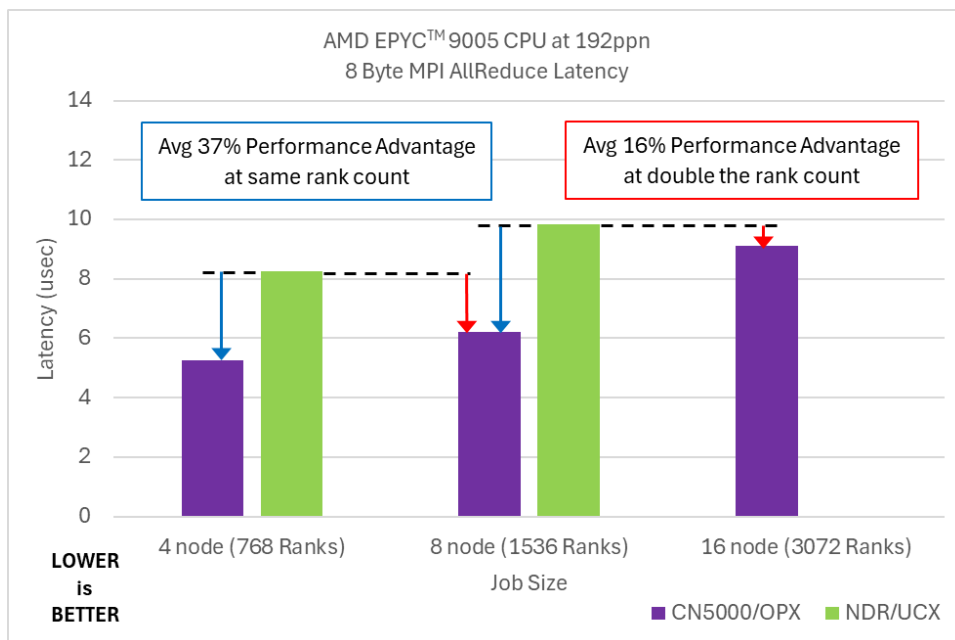


Figure 7 - 8 Byte AllReduce latency vs node count for Cornelis CN5000 and NDR IB

Performance per Dollar

In today's competitive landscape, every dollar invested must deliver maximum impact. This joint solution focuses on achieving an optimal ratio of performance to cost, ensuring that each dollar works harder for the business. Through precise resource allocation and smart technology choices, the investment becomes a growth engine rather than merely an expense. This architecture delivers exceptional performance while enabling seamless scalability, allowing expansion without unnecessary cost.

Considering the CN5000 performance measured above and leveraging published pricing for switches, adapters and cabling, the K905G6 plus CN5000 solution could deliver up to 1.4x advantage in performance per network dollar, proving that high performance and cost efficiency can coexist. For customers deploying dozens or hundreds of nodes, this advantage compounds into lower total cost of ownership and significantly reduced cost per simulation or design cycle.

Conclusion

Why Choose INVENTEC + Cornelis?

By choosing the K905G6 server platform with the CN5000 Omni-Path 400G high performance end-to-end networking solution, organizations are not just investing in high-performance hardware—they are securing a strategic edge in the era of AI and advanced scientific computing. This solution's industry-leading performance and scalability empowers enterprises to accelerate innovation, reduce time-to-insight, and confidently tackle the most demanding workloads. The robust hardware integration featuring top-tier processors, high-speed memory, and advanced networking—ensures maximum reliability and flexibility for diverse enterprise and research environments. Most importantly, the outstanding scaling efficiency and optimized collective performance means users benefit from more performance for every dollar invested. By maximizing infrastructure utilization and minimizing bottlenecks, this platform delivers outstanding cost-effectiveness, helping organizations achieve superior results without overspending.

Appendix

Test Configuration

Dual socket AMD EPYC™ 9655 96C/384MB/400W/C1 Processor, 24x32GB DDR5 6400MT/s; Samsung PM9A3 MZQL21T9HCJR-00A07 U.2 NVMe

1x 400G Cornelis CN5000 Omni-Path SuperNIC or Nvidia ConnectX-7 NDR HCA

Red Hat Enterprise Linux 9.5 (Plow) - Kernel 5.14.0-503.11.1.el9_5.x86_64

Cornelis OPX version 12.0.3.0.49 (Libfabric 2.5.0); NVIDIA HPC-X 2.25.1 (UCX 1.20.0, Libfabric 2.2.0); Intel MPI 2021.17