Enkidu

— MIOVSKI GROUP —

Whitepaper

ENKIDU · PLATFORM ANALYSIS

DGX and IGX at the Edge

A comparative analysis of NVIDIA DGX and IGX platforms for edge AI and high-speed sensor data processing.

Version 1.0

23 August 2026

Prepared by Miovski Group

Classification · Confidential

Document Control

DOCUMENT PROPERTIES

PROPERTY	VALUE
Document ID	ENK-WP-2026-014
Classification	Confidential — approved for client distribution
Status	Released
Canonical language	English (en-CA) — translations are derived artifacts
Governing references	BRAND.md v8 · Brand Guidelines v1.0

VERSION HISTORY

VERSION	DATE	AUTHOR	SUMMARY OF CHANGE
0.1	23 August 2026	Miovski Group	Initial draft — DGX–IGX comparative analysis migrated to the Enkidu whitepaper standard.

VERSION	DATE	AUTHOR	SUMMARY OF CHANGE
1.0	23 August 2026	Miovski Group	First issued version — reviewed and released.

Record every issued version. A version is issued only after the review below is complete.

REVIEW AND SIGN-OFF

ROLE	NAME	SIGNATURE	DATE
Technical review			
Editorial review — voice and claims			
Release authority			

Editorial review verifies: every metric carries a source; no superlatives; canonical language quoted byte-identically; naming per the registry.

DISTRIBUTION

RECIPIENT	ORGANISATION	FORMAT
—	—	PDF
—	—	PDF

Distribution beyond this list requires the release authority’s approval.

Contents

Executive Summary3

Glossary3

Abstract5

1. Introduction5

    1.1 Background and requirement6

2. Platform Positioning6

3. Specifications and Configurations7

    3.1 DGX family7

    3.2 IGX family8

4. The Sensor-to-GPU Data Path9

    4.1 Latency and determinism10

    4.2 Holoscan11

    4.3 Functional safety11

5. Comparative Assessment11

    5.1 DGX Spark and IGX Thor T700011

    5.2 Weighted assessment12

6. Reference Architecture13

7. Indicative Pricing (Canadian Dollars)14
8. Limitations and Scope15
9. Conclusion and Next Steps16
References17
Appendix A — Pricing Assumptions17

Executive Summary

The problem

A planned system must run AI inference against continuous, high-rate sensor streams at the physical edge — the environment Enkidu builds for. The platform decision is between NVIDIA's DGX family, built for data-centre AI, and its IGX family, built for industrial and medical edge deployment. Selecting the wrong tier commits the programme to hardware that cannot meet either the ingest bandwidth or the determinism the requirement demands.

What this work shows

The two families share the Blackwell GPU generation and the CUDA software ecosystem but are engineered for opposite ends of the operational stack. For edge AI with high-speed sensor processing, the IGX Thor T7000 configuration is the stronger platform: it integrates a ConnectX-7 SmartNIC with dual 200 GbE ports, a zero-copy GPUDirect RDMA path into GPU memory, the Holoscan sensor-processing framework, an independent functional-safety domain, and a 10-year lifecycle [1][2]. DGX remains the correct choice upstream as the training and model-development environment, and DGX Spark serves as the per-developer workstation because it shares the ConnectX-7 and Holoscan Sensor Bridge development path [4][6].

Evidence basis

Specifications are drawn from NVIDIA platform documentation and technical publications; pricing is drawn from published list prices, marketplace listings and OEM quotations as of August 2026, converted to Canadian dollars at an assumed 1.37 CAD/USD (Appendix A). Every figure in this document carries its source; quote-based items are marked as estimates.

Recommended next step

Approve a sizing exercise — sensor inventory, per-stream bandwidth and latency budget — to confirm the T7000 configuration per deployment site, and request formal Canadian-dollar quotations from an authorised IGX distributor.

Glossary

Key terms used throughout this document.

DGX

NVIDIA's data-centre AI platform family — systems built for model training, fine-tuning and large-scale inference, from the desktop DGX Spark to rack-scale DGX B300 systems and SuperPOD clusters.

IGX

NVIDIA's industrial-grade edge AI platform family — purpose-built for industrial, robotics and medical environments, combining sensor I/O, functional safety and long-term support. Current generation: IGX Thor.

GPUDirect RDMA

A transfer mechanism in which a network adapter writes incoming data directly into GPU memory, bypassing the CPU and system-memory copies of a conventional network stack.

RoCE

RDMA over Converged Ethernet — the transport that carries remote-direct-memory-access traffic over standard Ethernet links.

ConnectX-7

An NVIDIA SmartNIC providing up to 2 × 200 GbE with hardware RDMA offload; integrated on the IGX Thor T7000 board and fitted in DGX Spark.

Holoscan

NVIDIA's sensor-processing SDK — the pipeline framework for acquisition, synchronisation, pre-processing, inference, fusion and output, integrating GPUDirect RDMA, Rivermax and TensorRT.

Holoscan Sensor Bridge

A hardware/software layer that converges heterogeneous sensor interfaces onto Ethernet for delivery into the Holoscan pipeline; supported on IGX Orin, IGX T7000 and DGX Spark.

FP4 TFLOPS

Trillions of floating-point operations per second at 4-bit precision — the unit NVIDIA quotes for Blackwell-generation AI inference throughput; figures cited here are sparse-math peaks.

Safety Island / sMCU

IGX's independent functional-safety processor and safety microcontroller, isolated from the main AI domain, with IEC 61508 and ISO 26262 design targets.

BMC

Baseboard management controller — out-of-band hardware for remote monitoring, provisioning and recovery of a deployed system.

Abstract

This analysis compares NVIDIA's DGX and IGX platform families against a requirement for edge AI computing and high-speed sensor data processing. It sets out the design intent of each family, details current specifications and configurations — DGX Spark, DGX Station GB300 and DGX B300 on one side; IGX Orin, IGX Thor T5000 and T7000 on the other — and examines the sensor-to-GPU data path that separates them. The IGX Thor T7000, with its integrated ConnectX-7 SmartNIC, GPUDirect RDMA ingest, Holoscan framework and independent safety domain, meets the stated requirement; DGX-class systems serve the complementary role of training and model development, with DGX Spark as the developer workstation. A weighted

assessment, a closed-loop reference architecture, and indicative Canadian-dollar pricing for each tier support the recommendation. The analysis closes with the limitations of the pricing basis and the sizing work required before procurement.

1. Introduction

Industrial and medical operators are deploying larger AI models, more sensors and higher-fidelity data streams at the edge, while safety and regulatory compliance demand deterministic behaviour and verifiable functional safety [2]. Within NVIDIA's portfolio, two platform families are candidates for such a system: DGX, the data-centre AI infrastructure line, and IGX, the industrial edge line. This document evaluates both against a single requirement — edge AI computing with high-speed sensor data processing — and recommends a platform decision and an end-to-end architecture.

CONTRIBUTIONS

- Frame the requirement and the design intent of each platform family.
- Detail current specifications, configurations and the sensor-to-GPU data path.
- Assess the candidates against weighted criteria and recommend an architecture with indicative Canadian-dollar pricing.

1.1 Background and requirement

The requirement has two parts. First, edge AI computing: inference must run at the point of data generation — a machine, facility or vehicle — under the power, thermal and environmental constraints of that location, and in many deployments under functional-safety obligations. Second, high-speed sensor data processing: the platform must ingest continuous multi-gigabit streams from cameras, LiDAR, radar and industrial instruments, and act on them within a bounded latency. These two properties, not peak compute alone, define fitness for purpose. The evaluation that follows draws on NVIDIA's platform documentation [1][2][3], the Holoscan SDK's performance guidance [4], and published pricing [5]–[9].

> Enkidu builds trusted AI for consequential environments — secure, sovereign, auditable, and under human authority.

2. Platform Positioning

NVIDIA's accelerated-computing portfolio spans a continuum from embedded modules to multi-rack clusters. The DGX family — Spark, Station and rack-scale B300/GB300 systems — targets AI development and the data centre [6][7][8]. The IGX family — IGX Orin, and the current IGX Thor generation comprising the T5000 system-on-module and the T7000 board kit — targets regulated, real-time environments: factory floors, operating rooms, autonomous machines and field installations [1][2][10].

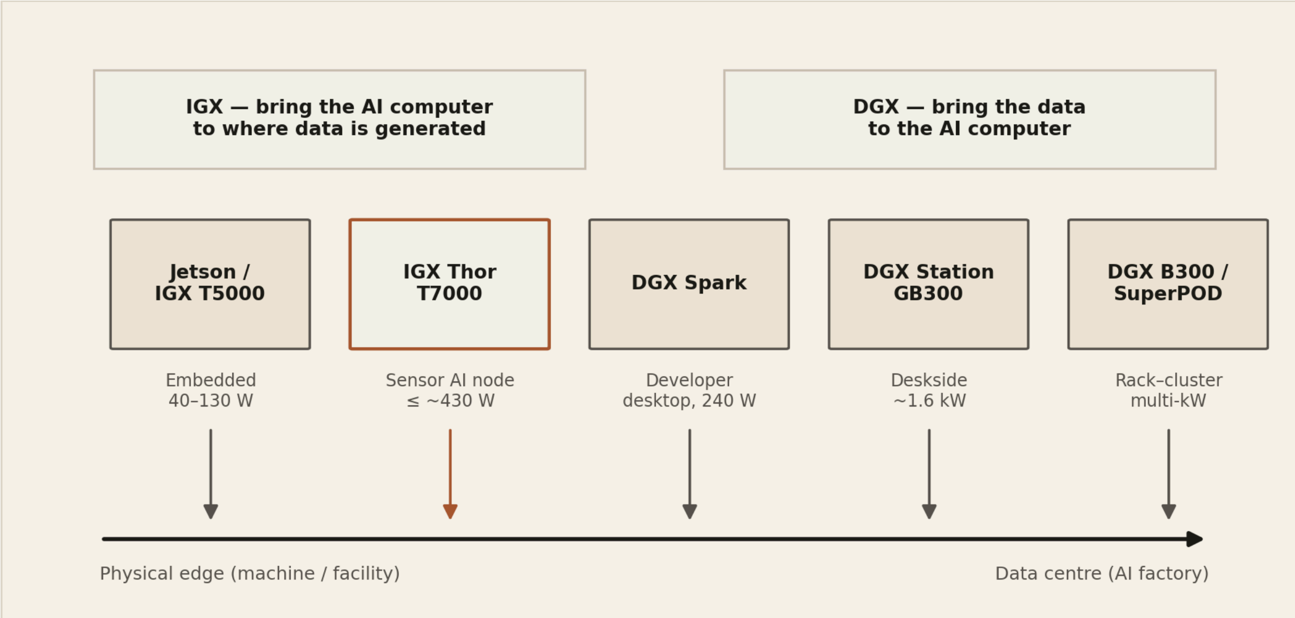

Figure 1 · NVIDIA platform positioning across the edge-to-data-centre continuum · Source: [1][2][6][7].

The design-philosophy difference matters more than any single specification. DGX assumes conditioned power, cooling and a network fabric oriented toward GPU-to-GPU communication; it brings data to a centralised computer. IGX assumes it is installed in a machine or cabinet, ingesting continuous sensor streams and making decisions within a latency budget, supervised by an independent safety domain; it brings the computer to the data.

3. Specifications and Configurations

3.1 DGX family

Table 1 summarises the three DGX tiers relevant to this analysis. Rack-scale DGX systems aggregate further into SuperPOD clusters; the single-system view is the meaningful comparison point here.

ATTRIBUTE	DGX SPARK	DGX STATION (GB300)	DGX B300 (RACK SYSTEM)
Role	Developer desktop and prototyping	Deskside AI system	Data-centre training and inference node
Compute	GB10 Grace Blackwell Superchip	GB300 Grace Blackwell Ultra Desktop Superchip	8 × Blackwell Ultra (B300) GPUs + dual x86 CPUs
CPU	20-core Arm (Grace)	72-core Arm Neoverse V2, NVLink-C2C to GPU	Dual-socket x86
AI performance	~1 PFLOP FP4 (sparse)	Up to ~20 PFLOPS FP4	Tens of PFLOPS FP4 per system
Memory	128 GB unified LPDDR5X, 273 GB/s	748 GB coherent (252 GB HBM3e + 496 GB LPDDR5X)	HBM3e-class GPU memory pool

ATTRIBUTE	DGX SPARK	DGX STATION (GB300)	DGX B300 (RACK SYSTEM)
Networking	ConnectX-7, 200 Gb/s class + 10 GbE	ConnectX-8 class, 800 Gb/s	InfiniBand / Ethernet fabric, 800 Gb/s class
Power	~240 W	Up to ~1.6 kW	Multi-kilowatt; data-centre cooling
Functional safety	None	None	None — enterprise IT reliability only
Lifecycle	Standard product line	Standard enterprise cycle	Standard enterprise cycle

Table 1 · DGX family specifications · Source: [6][7][8][12].

Two points bear on this requirement. DGX Spark is the only DGX member with an edge-plausible power envelope, and it carries a ConnectX-7 NIC with official Holoscan Sensor Bridge support — a legitimate sensor-pipeline development platform [4][6]. The DGX Station GB300 is a compute-only superchip without video output (a separate RTX PRO card is required for display) and is designed for deskside model development, not deterministic sensor ingest [12].

3.2 IGX family

Within the current IGX Thor generation, the critical distinction is between the T5000 system-on-module and the T7000 board kit: only the T7000 integrates the ConnectX-7 SmartNIC that enables the accelerated RoCE/GPUDirect sensor-receive path [2][4]. IGX Orin remains available and supported for designs already committed to it [3].

ATTRIBUTE	IGX ORIN	IGX THOR T5000 (SOM)	IGX THOR T7000 (BOARD KIT)
Role	Prior-generation edge platform	Embedded module for custom carrier boards	Sensor-aggregation and AI edge node
CPU	12-core Arm Cortex-A78AE	14-core Arm Neoverse V3AE	14-core Arm Neoverse V3AE
GPU (integrated)	Ampere-class, 248 TOPS (module)	Blackwell, up to 2,070 FP4 TFLOPS	Blackwell, up to 2,070 FP4 TFLOPS
Optional dGPU	RTX 6000 Ada — to ~1,705 TOPS aggregate	Via carrier-board design	RTX PRO 5000/6000 Blackwell — to ~5,581 FP4 TFLOPS aggregate
Memory	64 GB LPDDR5 class	128 GB LPDDR5X, 273 GB/s	128 GB LPDDR5X, 273 GB/s (+ dGPU GDDR7)
High-speed networking	2 × 100 GbE QSFP28 (dev kit)	Up to 4 × 25 GbE; no integrated ConnectX-7	ConnectX-7: 2 × 200 GbE QSFP112, RDMA/RoCE
GPUDirect sensor path	Yes (dev kit)	Not via ConnectX receiver configuration	Yes — primary design pattern
Functional safety	Safety MCU + extensions	Safety Island + sMCU	Safety Island + sMCU
Power	System-dependent	40–130 W (SoC)	40–130 W SoC + up to ~300 W dGPU

ATTRIBUTE	IGX ORIN	IGX THOR T5000 (SOM)	IGX THOR T7000 (BOARD KIT)
Lifecycle	10-year	10-year, long-term support	10-year, long-term support

Table 2 · IGX family specifications · Source: [1][2][3][11].

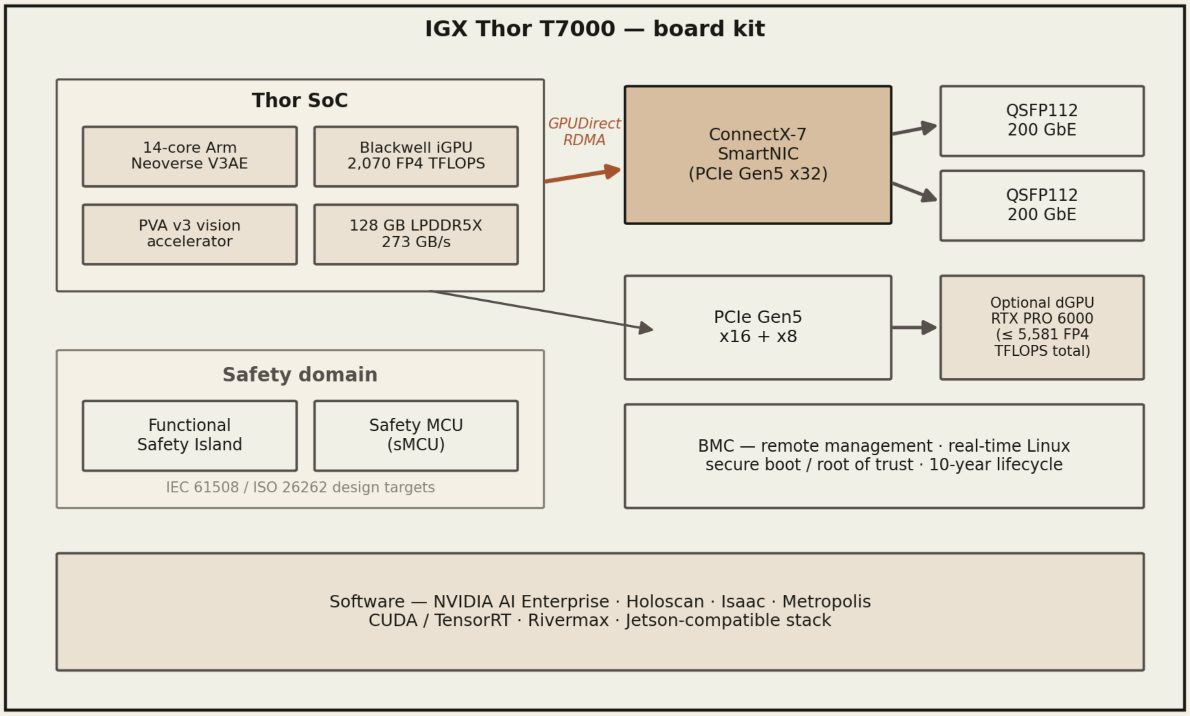

Figure 2 · IGX Thor T7000 functional block diagram — compute, sensor I/O and safety domains · Source: [2].

Configuration guidance within the family: select the T7000 for high-bandwidth multi-sensor aggregation — its dual 200 GbE ports give roughly 400 Gb/s of aggregate physical capacity with GPUDirect RDMA ingest [2]. Select the T5000 SoM where the priority is deep embedded integration on a custom carrier board at 40–130 W and sensor bandwidth stays within a 4 × 25 GbE envelope; the ConnectX-based accelerated receiver configuration is not supported on T5000 [4]. Add the RTX PRO 6000 Blackwell dGPU to a T7000 when model size or concurrent stream count exceeds the integrated GPU [9].

4. The Sensor-to-GPU Data Path

High-speed sensor processing is governed less by peak TFLOPS than by how sensor bytes reach GPU memory. In a conventional CPU-centric design, every packet traverses the NIC, the kernel network stack, CPU buffers, system memory and a PCIe copy before reaching the GPU. Each stage adds copies, CPU utilisation, memory pressure, scheduling overhead, latency and latency jitter.

IGX integrates the ConnectX-7 SmartNIC into the board so that high-rate streams — industrial camera arrays, LiDAR, radar, ultrasound — are written directly into GPU memory over RoCE using GPUDirect RDMA, with the transfer offloaded from the CPU [2][4]. Combined with the Holoscan Sensor Bridge, which converges heterogeneous sensor interfaces onto Ethernet, this forms a zero-copy, deterministic ingest pipeline.

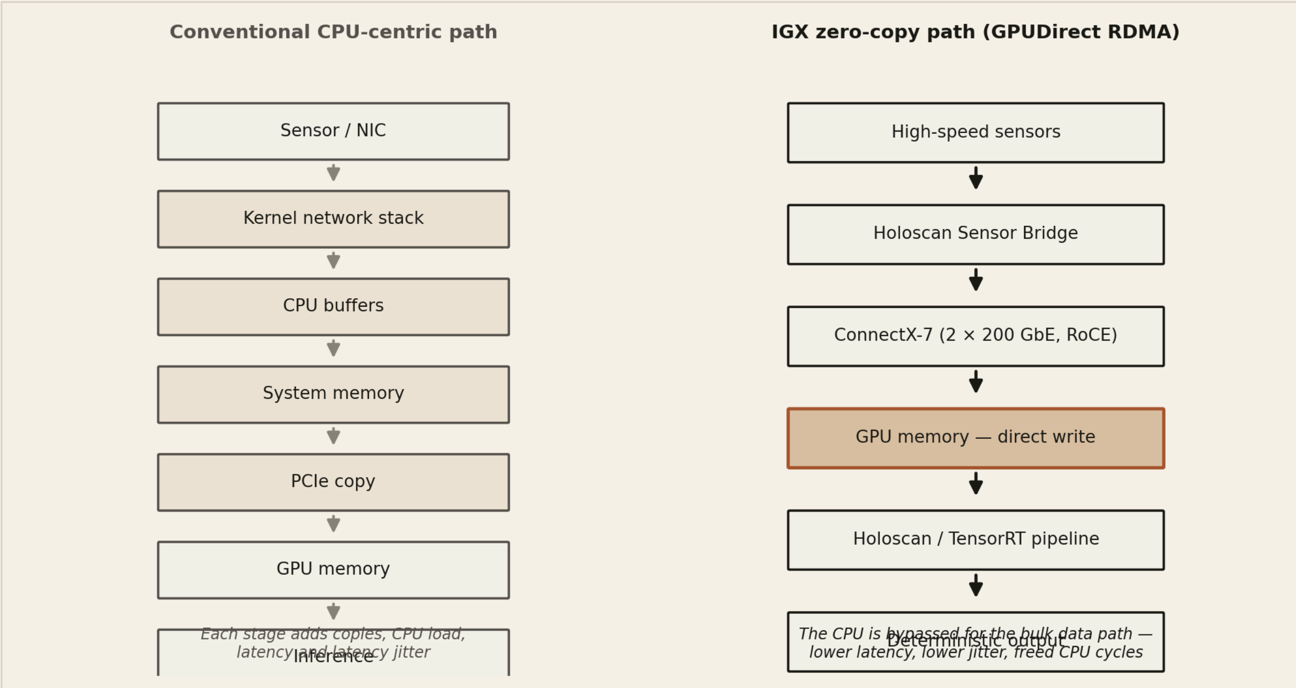

Figure 3 · Conventional CPU-centric ingest and the IGX ConnectX-7 / GPUDirect zero-copy path · Source: [2][4].

A worked example: five sensors producing 20 + 20 + 10 + 10 + 40 Gb/s yield roughly 100 Gb/s of continuous ingest. A T7000 carries this on a single 200 GbE port, leaving headroom for burst and synchronisation traffic and a second port for redundancy. A CPU-centric system sustaining the same rate through the kernel would consume a substantial share of CPU capacity before any inference begins.

4.1 Latency and determinism

DGX optimises primarily for throughput — work completed per unit time. IGX optimises for latency, throughput and determinism together. For real-time sensor systems the tail governs: a profile of 5 ms average and 12 ms worst-case supports closed-loop control; 3 ms average with 80 ms at the 99.9th percentile does not. IGX's real-time Linux support, hardware safety supervision and offloaded data path are directed at bounding that tail [2].

4.2 Holoscan

For this requirement the Holoscan SDK is nearly as consequential as the hardware. It provides the end-to-end pipeline — acquire, synchronise, pre-process, filter, infer, fuse, analyse, output — and integrates GPUDirect RDMA, Rivermax, TensorRT, UCX, PVA offload and NVDEC/NVENC [4]. NVIDIA's Holoscan guidance recommends an IGX system with ConnectX SmartNICs for best performance, and Holoscan Sensor Bridge is supported on IGX Orin, IGX T7000 and DGX Spark — the basis for the development-workflow recommendation in Section 6 [4].

4.3 Functional safety

IGX Thor includes a dedicated Functional Safety Island — an independent safety processor that isolates safety-critical workloads — and a Safety MCU, with design targets including IEC 61508 and ISO 26262 [2]. If a sensor feed drops or an anomaly occurs while the system is controlling machinery or a medical device, the safety domain can enforce deterministic fail-safe behaviour independently of the main AI domain. DGX provides no machine-level safety architecture, because data centres do not require one. In regulated deployments this is frequently the deciding factor, ahead of raw performance.

5. Comparative Assessment

5.1 DGX Spark and IGX Thor T7000

Because DGX Spark is the one DGX product with genuine sensor-pipeline credentials, the Spark–T7000 comparison is the most decision-relevant.

ATTRIBUTE	DGX SPARK	IGX THOR T7000
Architecture	GB10 Grace Blackwell	Thor (Blackwell iGPU) + optional dGPU
AI compute	~1 PFLOP FP4	2.07 PFLOPS iGPU; up to 5.58 PFLOPS with RTX PRO 6000
Memory	128 GB unified, 273 GB/s	128 GB LPDDR5X, 273 GB/s + dGPU GDDR7
Sensor NIC	ConnectX-7, 200 Gb/s class	ConnectX-7, 2 × 200 GbE QSFP112
GPUDirect / Sensor Bridge	Yes / yes	Yes / yes — primary platform
PCIe expansion	Limited	Gen5 x16 + x8
PVA · Safety Island · sMCU · BMC	— · — · — · —	Yes · yes · yes · yes
Industrial lifecycle, 10-yr support	No	Yes
Power	~240 W system	40–130 W SoC + up to ~300 W dGPU
Best role	Developer desktop; pipeline prototyping	Production edge sensor fusion and inference

Table 3 · DGX Spark and IGX Thor T7000 · Source: [2][4][6].

Spark approximates the IGX sensor environment on a developer's desk — the same NIC family, Sensor Bridge support and CUDA/Holoscan stack — at a fraction of a full IGX build-out. It lacks the safety domain, BMC, industrial lifecycle, PCIe expansion and dual-port 200 GbE aggregation of the T7000, and is therefore not the production edge node.

5.2 Weighted assessment

The scoring below, carried forward from the preliminary investigation and consistent with the specifications above, rates each candidate against criteria weighted for edge AI with high-speed sensor processing. Scores are an architectural assessment, not vendor benchmarks.

CRITERION	WEIGHT	DGX SERVER	DGX SPARK	IGX T5000	IGX T7000
Sensor bandwidth	20%	9	8	7	10
Low latency	20%	6	8	9	10
GPUDirect sensor path	15%	8	9	6	10
AI compute	15%	10	7	8	9
Edge integration	10%	2	6	10	10
Functional safety	5%	2	2	9	10
Industrial lifecycle	5%	5	4	10	10
Expansion	5%	9	4	7	10
Power efficiency	5%	1	8	10	7
Weighted total	100%	6.6	7.4	8.2	9.8

Table 4 · Weighted assessment for the stated requirement · Source: author's assessment over [1]–[12].

The T7000's margin comes not from any single line item but from being strong wherever the requirement is weighted — ingest bandwidth, deterministic latency, the zero-copy path, edge integration and safety — while remaining close to far larger systems on usable AI compute once the dGPU is fitted.

6. Reference Architecture

The stronger overall system treats DGX and IGX as complementary layers of one closed loop. DGX-class infrastructure in the data centre acts as the AI factory: it trains and fine-tunes models, generates synthetic data, runs fleet-wide analytics, and validates and optimises models for deployment — TensorRT compilation and quantisation to FP4. DGX Spark units on developer desks provide a faithful, lower-cost approximation of the edge sensor environment for Holoscan pipeline development and integration testing. IGX Thor T7000 nodes at the physical edge receive the compiled models and execute real-time sensor fusion and inference against live streams within a bounded latency, under independent safety supervision, returning curated field data and telemetry for the next training cycle.

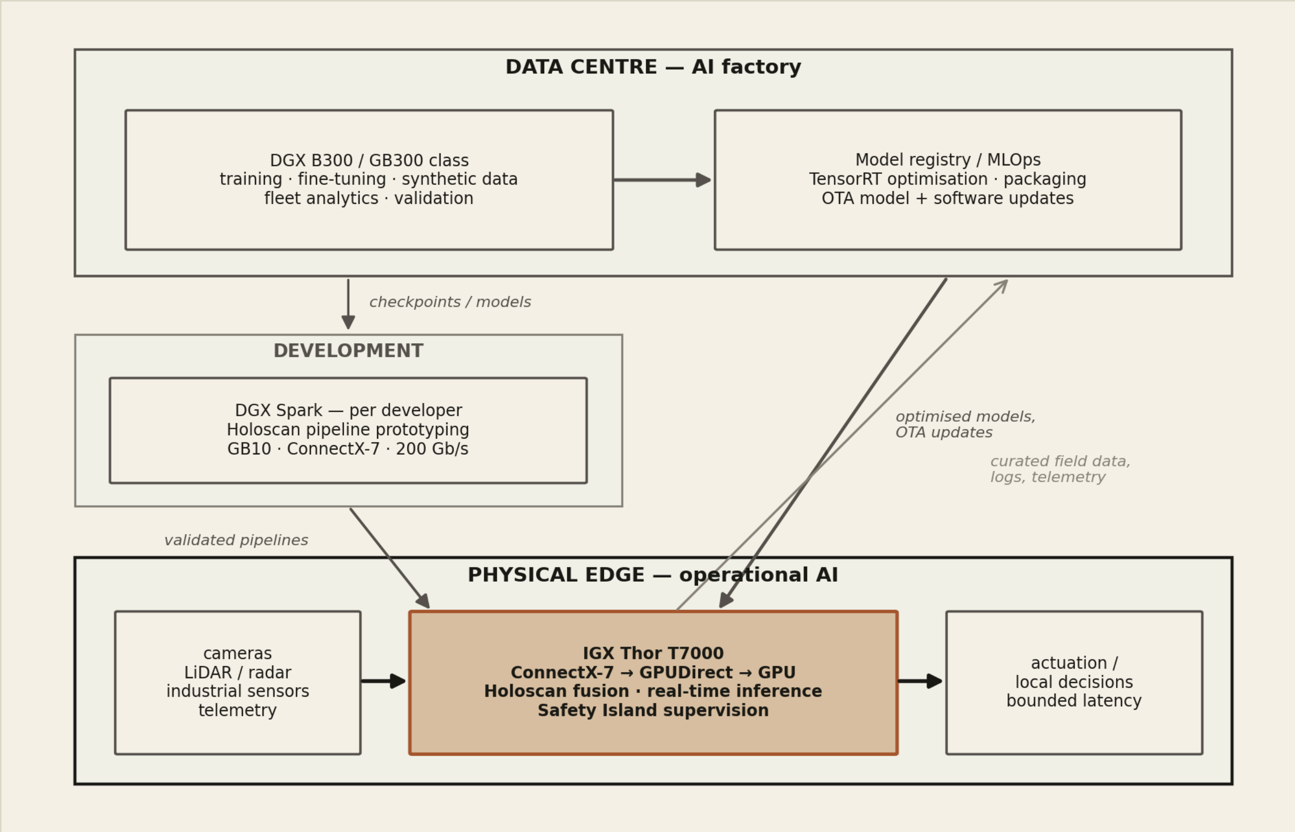

Figure 4 · Closed-loop deployment — DGX trains, DGX Spark develops, IGX T7000 operates · Source: [2][4][6].

7. Indicative Pricing (Canadian Dollars)

Pricing is indicative and intended for budgeting comparison only. USD figures are drawn from published list prices, marketplace listings and reseller or OEM quotations as of August 2026 [5]–[9]; Canadian-dollar conversion assumes 1.37 CAD per USD and moves with the exchange rate (Appendix A). All figures exclude taxes, duties, shipping and support subscriptions. IGX hardware is largely quote-based through authorised distributors — Arrow Electronics and NVIDIA partners such as PNY, Leadtek, Aetina and ADLINK; in Canada typically via Arrow, CDW Canada or Insight Canada — so rows marked “est.” require a formal quotation.

ITEM	TYPICAL USD PRICE	APPROX. CAD	NOTES
Jetson AGX Thor Developer Kit (T5000 SoM)	$3,499 list; $5,499 bundle listings [5]	C$4,800 – C$7,550	Closest low-cost proxy for Thor-class compute; limited-release kit
NVIDIA DGX Spark	$3,999 at launch; ~$4,699 current [6]	C$5,500 – C$6,450	Per-developer workstation; ConnectX-7; Sensor Bridge supported
IGX Orin Developer Kit	Quote-based; est. $6,000–$10,000 [3]	C$8,000 – C$14,000 (est.)	Via Arrow / partners; ConnectX-7 on dev kit
RTX PRO 6000 Blackwell dGPU	$8,200 – $10,000 [9]	C$11,200 – C$13,700	Optional dGPU for T7000; 96 GB GDDR7

ITEM	TYPICAL USD PRICE	APPROX. CAD	NOTES
IGX Thor dev kit + RTX PRO 6000	Quote-based; est. $13,000–$22,000	C$18,000 – C$30,000 (est.)	Production-representative edge configuration
DGX Station (GB300)	$90,000 – $115,000 [7]	C$123,000 – C$158,000	OEM-built — ASUS, Dell, HP, MSI, Supermicro, GIGABYTE
DGX B300 system	$300,000 – $350,000 [8]	C$411,000 – C$480,000	Rack system; excludes fabric, storage, facility
GB300 NVL72 / SuperPOD	Multi-million per rack [8]	Multi-million CAD	Cloud or colocation often preferable to capital purchase

Table 5 · Indicative pricing by platform tier · Source: [5][6][7][8][9]; conversion basis in Appendix A.

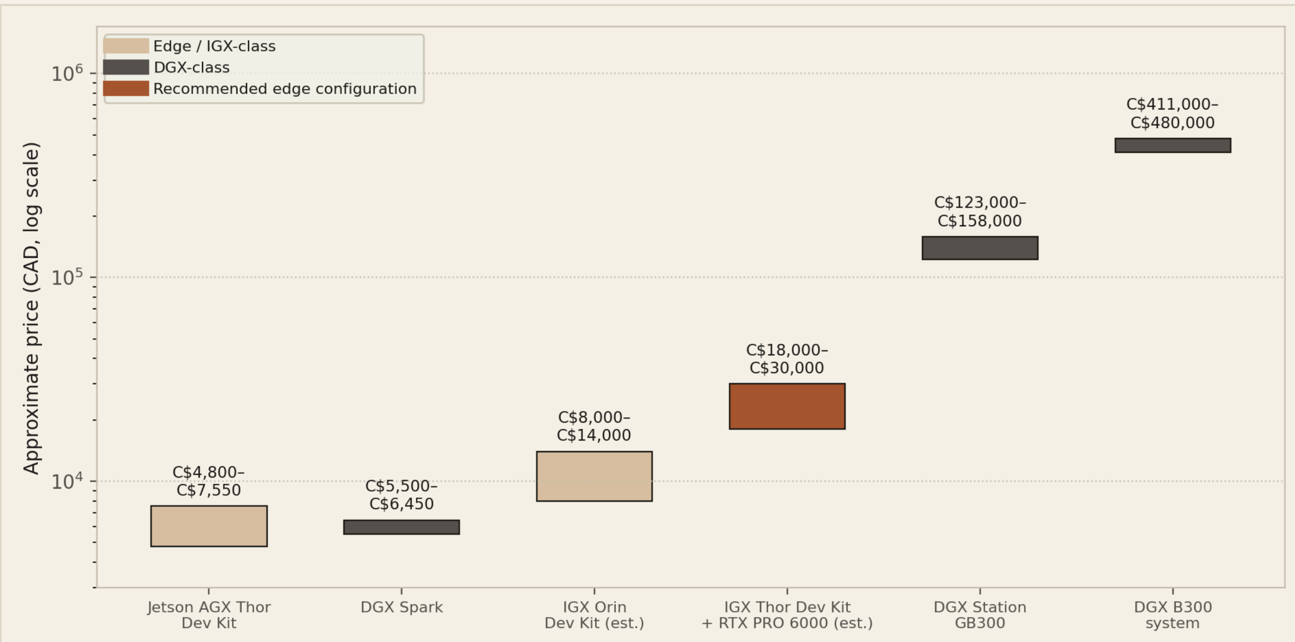

Figure 5 · Indicative CAD price ranges by platform tier, log scale · Source: Table 5.

Budgeting observations. A complete development environment — one DGX Spark per developer plus one shared IGX Thor T7000 system with dGPU — assembles for roughly C$25,000–C$45,000 depending on quotations, under a third of a single DGX Station. Training capacity is better rented — DGX Cloud or cloud GPU instances, with B300-class GPU-hours in the US$7.10–$17.80 range as of August 2026 [8] — until sustained utilisation justifies owned DGX hardware. Production IGX unit economics are dominated by the T7000-plus-dGPU bill of materials and should be quoted per volume tier.

8. Limitations and Scope

- Pricing for IGX Thor kits and modules is not published; the estimates here are anchored to the Jetson AGX Thor kit and RTX PRO 6000 prices [5][9] and must be replaced by formal distributor quotations before any budget commitment.

- The CAD conversion basis (1.37 CAD/USD) is an assumption as of August 2026; figures move with the exchange rate and exclude taxes, duties, shipping and support.
- Quoted AI-performance figures are vendor sparse-math peaks (FP4) and are not directly comparable to measured application throughput; no benchmark of the candidate systems was performed for this document.
- The weighted assessment in Table 4 is an architectural judgement, not a measurement; weights reflect the stated requirement and should be revisited if the requirement changes.
- Sensor counts, bandwidths and latency budgets for the target deployment were not yet available; the recommendation is at platform level and requires the sizing exercise in Section 9 before procurement.

9. Conclusion and Next Steps

If the question were which platform provides the most AI compute, DGX would carry it. Against the actual requirement — edge AI with high-speed sensor processing — IGX Thor T7000 is the stronger architecture, and the deciding factor is the complete sensor-processing design rather than the GPU alone: ConnectX-7 dual 200 GbE ingest, GPUDirect RDMA into GPU memory, Holoscan as the pipeline framework, PVA offload, the Functional Safety Island and Safety MCU, BMC-based manageability, and a 10-year industrial lifecycle [1][2][4].

- Production edge — IGX Thor T7000, with the RTX PRO 6000 Blackwell dGPU where model size or stream count warrants it; the T5000 SoM only for deeply embedded, lower-bandwidth designs.
- Development — one DGX Spark per developer for Holoscan pipeline prototyping, Sensor Bridge testing and small-scale fine-tuning.
- Central AI infrastructure — DGX B300/GB300 class, owned, colocated or via DGX Cloud, for training, fine-tuning, synthetic data, simulation and fleet analytics.
- Procurement — request formal CAD quotations for the IGX Thor developer kit, T7000 board kit and volume module pricing from an authorised IGX partner; validate the RTX PRO 6000 configuration and thermal envelope for the target enclosure.
- Sizing (recommended immediate next step) — enumerate sensor count and type, per-sensor bandwidth, required average and 99.9th-percentile latency, and target models; from these derive aggregate ingest rate, iGPU-versus-dGPU requirement, network topology and Holoscan/GPUDirect pipeline placement, and confirm DGX Spark versus IGX T5000 versus T7000 per deployment site.

With sensor inventory and latency targets in hand, this document extends to a concrete sizing architecture with estimated data rates, GPU counts and network design per site.

References

[1] NVIDIA, “NVIDIA IGX” platform overview, nvidia.com/en-eu/edge-computing/products/igx — accessed August 2026.

[2] NVIDIA Technical Blog, “NVIDIA IGX Thor Powers Industrial, Medical, and Robotics Edge AI Applications,” April 2026.

[3] NVIDIA documentation, “IGX Orin Developer Kit — Product Brief”; Arrow Electronics IGX Orin distribution page — accessed August 2026.

[4] NVIDIA Holoscan SDK documentation — sensor-processing performance guidance and Holoscan Sensor Bridge supported platforms — accessed August 2026.

[5] Hackster.io and CNX-Software, Jetson AGX Thor Developer Kit pricing (US$3,499 list), August 2025; NVIDIA marketplace bundle listing (US$5,499) — accessed August 2026.

[6] Engadget, “NVIDIA starts selling its $3,999 DGX Spark,” October 2025; IntuitionLabs DGX Spark review noting US$4,699 current pricing, 2026.

[7] Tom’s Hardware, DGX Station GB300 listings US$94,930–$108,350, August 2026; aiHola, MSI XpertStation WS300 CDW listing US$96,996, March 2026; VRLA Tech, DGX Station GB300 US$90,000–$115,000, August 2026.

[8] Tech-Insider, “NVIDIA Blackwell GPU Pricing,” August 2026 — DGX B300 systems US$300,000–$350,000; B300 cloud rates US$7.10–$17.80 per GPU-hour.

[9] TechRadar/PNY, RTX PRO 6000 Blackwell Workstation Edition at US$8,200; retail listings to approximately US$10,000 — accessed August 2026.

[10] NVIDIA Newsroom, “NVIDIA Launches IGX Edge AI Computing Platform for Safe, Secure Autonomous Systems,” September 2022.

[11] Aetina, AIE-MIX640 IGX Orin system announcement — 248 TOPS, expandable to 1,705 TOPS with RTX 6000 Ada, COMPUTEX 2024.

[12] ServeTheHome and VRLA Tech, DGX Station GB300 system analyses — 748 GB coherent memory, up to ~20 PFLOPS FP4, ~1.6 kW, 2026.

Appendix A – Pricing Assumptions

Conversion basis. Canadian-dollar figures apply a single assumed rate of 1.37 CAD per USD, rounded to the nearest C$50–C$1,000 depending on magnitude. The rate is an assumption for comparability, not a forecast; procurement budgets should apply the prevailing rate and the distributor's quoted currency.

Estimate construction. Where no list price exists — IGX Orin and IGX Thor kits — the estimate brackets are anchored as follows: the lower bound to the closest published NVIDIA hardware analogue (Jetson AGX Thor Developer Kit at US$3,499 [5], plus the IGX platform's industrial hardware, safety components and ConnectX-7 content), and the upper bound to community-reported and partner-system pricing for IGX-class hardware [3][11]. The combined "IGX Thor dev kit + RTX PRO 6000" row sums the kit estimate and the published dGPU range [9].

Exclusions. All figures exclude sales taxes, import duties, brokerage, shipping, extended support and NVIDIA AI Enterprise subscription components, enclosures, cabling, optics and installation. Rack-system figures exclude network fabric, storage and facility costs.