SYMBEEO

PUBLIC

SOLUTION DESIGN DOCUMENT

Sixteen-node DGX B300 block design for a typical buyer, applying SRA-AI-001

Design document (HLD + LLD)

CAPABILITY	DOCUMENT	VERSION	DATE
AI & Data	SDD-AI-001	1.3	Oct 10, 2026

WRITTEN BY

Diksha Shah, Principal AI Infrastructure Architect

© 2026 Symbeeo. All rights reserved.

Document control

Document	SDD-AI-001	Version	1.3
Issued	Oct 10, 2026	Status	Issued by Symbeeo
Classification	Public	Supersedes	1.2, Sep 27, 2026
Capability	AI & Data	Author	Diksha Shah
Role	Principal AI Infrastructure Architect	Applies to	SRA-AI-001 v1.2
Tested by	STP-AI-001		
Credentials	NVIDIA Certified Professional: AI Operations, NVIDIA Certified Associate: AI Infrastructure and Operations, Google Cloud Professional Cloud Architect, AWS Certified Data Engineer		
Approval	Approved by Edwin Ozer, Founder, Oct 10, 2026. Issued by Symbeeo, Oct 10, 2026.		
Sign-off	"The customer's named representative signs this document before build; the practice lead signs it after."		

Contents

1 · Introduction 3
2 · Prerequisites 3
3 · Solution overview 4
4 · Logical design 4
5 · Physical design 5
6 · Staging and delivery 5
7 · Migration and growth 5
8 · Roles and deliverables 5
9 · Change control and reporting 5
10 · Security and lifecycle clauses 6

© 2026 Symbeeo. All rights reserved. · Public

CONVENTIONS

The reference architecture applied to the typical buyer: a quarter scalable unit, sixteen DGX B300 systems, in a hall rated at 30 kW per rack on air. Figures marked [S] are from NVIDIA RA-11337-001 and the Rubin NVL8 reference architecture, and, where marked read Sep 24, 2026, from the sources in SRA-AI-001 §1's technology review; [A] are the typical buyer's until the site survey replaces them. A row with more than one tag is counted once: as a published figure if any tag is [S] or [S*], otherwise as an assumption if any tag is [A], otherwise as a design target.

Table 1. Revision history

Date	Version	Change
Sep 15, 2026	1.0	First issue.
Sep 21, 2026	1.1	Corrections to wording, figure numbering and citations.
Sep 27, 2026	1.2	Revised to current platform versions and the current document format.
Oct 10, 2026	1.3	Approval line now reads: approved by the Founder, issued by Symbeeo. Diksha Shah remains the author.

1 · Introduction

1.1 Purpose.

Design for one AI compute block: 16 DGX B300, compute fabric, storage fabric, management, racks, power and cooling, delivered staged and burned in. **1.2 Goals.** All 16 nodes visible to one scheduler; NCCL all-reduce across the block at the RA's expected bandwidth; a fine-tuning job runs end to end on day one. **1.3 Approach.** Site works first (power, cooling); fabric racks second; compute racks staged at Symbeeo and shipped built; acceptance runs per STP-AI-001.

2 · Prerequisites

Table 2. Prerequisites

#	Prerequisite	Owner	Status
P1	Hall survey: power per rack available, floor loading, containment, chilled water for rear doors; DC busbar path or AC PDUs (SRA-AI-001 §5 D2)	Symbeeo	open
P2	Power upgrade to > 50 kW A+B on 4 compute racks (RA: B300 at 4 per rack exceeds 50 kW) [S]	client	open
P3	Rear-door heat exchangers ordered for 4 racks [A]	client	open
P4	Storage platform chosen from the DGX SuperPOD certified list	client	open
P5	Scheduler choice (Slurm or Kubernetes) and identity provider	client	open
P6	Fiber plant between compute, fabric and storage racks measured	Symbeeo	open
P7	Staging space, power and burn-in window at Symbeeo	Symbeeo	open

© 2026 Symbeeo. All rights reserved. · Public

3 · Solution overview

3.1 Racks [A].

Table 3. Racks [A]

Rack	Contents	Power
C1–C4	4 × DGX B300 each (16 total)	> 50 kW each, A+B [S]
F1–F2	compute fabric: Quantum-X800 Q3400-RA leaf and spine positions for the full SU [S]	per switch spec
S1	storage fabric: Spectrum-4 SN5610 pair; certified storage system	per spec
M1	management nodes, NVIDIA Mission Control, UFM 3.5 Appliance, OOB, PDUs	per spec

3.2 Compute fabric.

800 Gbps XDR InfiniBand, rail-optimized, non-blocking, twin-plane fat tree; eight XDR connections per DGX via ConnectX-8 [S]. Leaf and spine cabled for the full SU of 64 even though 16 are populated, per the RA's partial-SU rule [S] (SRA-AI-001 §4.1, §4.6).

3.3 Storage fabric.

Spectrum-4 SN5610 pair at 800 Gbps Ethernet; management nodes provide storage access independent of compute [S] (SRA-AI-001 §4.5).

3.4 Management.

UFM 3.5 Appliance, Enterprise Edition for the fabric [S]; NVIDIA Mission Control, which includes Base Command Manager and Run:ai, as the RA's standard [S] read Sep 24, 2026; OOB network for BMCs; the client's identity provider for users (SRA-SEC-001).

4 · Logical design

Table 4. Logical design

Network	Purpose	Notes
compute fabric	NCCL, collective operations	XDR, twin-plane [S]
storage fabric	dataset and checkpoint I/O	800G Ethernet [S]
in-band management	scheduler, node management	SN5610 / SN5600D at 200 Gbps bonded, per the RA [S] read Sep 24, 2026
OOB	BMC, PDU, switch management	1G, isolated [A]; SN2201 per the RA [S] read Sep 24, 2026

Scheduler: Slurm with GPU partitions, or Kubernetes with the GPU operator, per P5, both under Mission Control. Users and projects mapped from the identity provider; quotas per project [A].

© 2026 Symbeeo. All rights reserved. · Public

5 · Physical design

Cabinet elevation per rack drawn before staging (Cisco IT's practice, BRKCOC-3005). Cooling: rear-door heat exchangers on C1–C4 sized to the measured load; fabric and storage racks on air [A]. Power: DC busbar or A/B PDUs at the RA's rack figure [S]; if P1 finds no busbar path, the AC form of RA-11339-001 (~56 kW per rack, SU of 72) applies and the full-SU cabling in §3.2 is re-drawn to it [S]. Fiber: XDR transceivers and cables per the RA's cable tables; lengths from P6.

6 · Staging and delivery

Each compute rack built, cabled and burned in at Symbeeo: burn-in for the hours set at the site survey and recorded as measure 3 of the eight, power-on test, per-node GPU diagnostics, NCCL test within the rack; then shipped assembled. On site: rack placement, fabric cabling, power-on, STP-AI-001.

7 · Migration and growth

No migration; new capacity. Growth to 32 and 64 nodes fills the fabric positions already cabled [S]. The Rubin NVL8 generation, from partners in H2 2026 [S], needs a liquid-cooled hall at 220 kW per rack [S] — a new-hall decision, recorded here so nobody expects an in-place upgrade.

8 · Roles and deliverables

Design authority: practice lead and the client's head of research computing (both sign). Deliverables: this document · STP-AI-001 with results · as-built (nodes, serials, fabric map, versions) · operations runbook · acceptance certificate · the handover pack (§10).

9 · Change control and reporting

Change orders. After this document is signed, any change is a written change order stating what changes, why, and its cost and schedule effect. Both parties sign it before it is built. Change orders per project are measure 6 of the eight.

Weekly status report. Each week Symbeeo sends the client a report giving % complete per deliverable (§8) and slippage against plan.

Escalation contacts by role [A].

Table 5. Change control and reporting

Level	Symbeeo	Client
1	Principal AI Infrastructure Architect	head of research computing
2	practice lead	CTO
3	the Founder	the client's executive sponsor

© 2026 Symbeeo. All rights reserved. · Public

10 · Security and lifecycle clauses

- A security incident affecting the client's data or systems is **reported to the client within 24 hours** of discovery.
- Storage media leaving the client's control (failed drives, returned nodes) are **sanitized to NIST SP 800-88 Rev. 2** (published September 2025 [S] read Sep 24, 2026), with a certificate per device; or the media stay with the client under a non-return service. Media holding PHI or other regulated personal data are never taken back unsanitized.
- **No equipment is quoted past its vendor's end-of-sale date**, and the end-of-support date of every item is shown on the quote. At Sep 24, 2026 the DGX B300 is shipping and the switches named here appear in NVIDIA's current reference architectures [S]; end-of-sale and end-of-support dates were not checked item by item and are taken from the vendor at quote time.
- **Every quote states the vendor support level by severity**: the response time for severity 1 to 4, as the vendor's support contract gives it.
- **Handover pack** at the end: as-built drawings, configurations, credentials transferred, licenses (Mission Control, UFM), support contracts, the signed STP-AI-001, and open items.

© 2026 Symbeeo. All rights reserved. · Public

SYMBEEO

Each vendor's tool shows you their option. We show you all of them against one requirement.

Compared. Then Delivered.

hello@symbeeo.com · symbeeo.com

© 2026 Symbeeo. All rights reserved.

Product and company names and logos are their owners' trademarks, shown only to identify the lines Symbeeo sources. Naming a vendor or showing its logo does not mean the vendor sponsors or approves Symbeeo's comparisons or any recommendation Symbeeo makes.

Figures are as published by their sources on the dates given in this document; figures marked [A] are assumptions.

SDD-AI-001 · Version 1.3 · Oct 10, 2026