SYMBEEO

PUBLIC

REFERENCE ARCHITECTURE

AI compute block on NVIDIA DGX B300, one scalable unit or a fraction of it, with the Rubin path

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

WRITTEN BY

Diksha Shah, Principal AI Infrastructure Architect

© 2026 Symbeeo. All rights reserved.

Document control

Document	SRA-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	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 practice lead signs this document before any build proceeds under it."		

Contents

1 · Scope 3
2 · Requirements register 7
3 · Architecture overview 8
4 · Component design 8
5 · Options and platform choices 9
6 · Failure conditions 11
7 · Validation 11
8 · Operations and lessons learned 11
9 · Appendix 12

© 2026 Symbeeo. All rights reserved. · Public

CONVENTIONS

Source tags. [S] read Sep 7, 2026: NVIDIA SuperPOD DGX B300 Systems, Spectrum-4 Ethernet and DC Busbar Power Reference Architecture, RA-11337-001 V01, publication May 29, 2025, release Aug 2025, 35 pp · NVIDIA DGX SuperPOD Rubin NVL8 Systems reference architecture, key components page (2026) · Isambard-AI arXiv 2410.11199 · Cisco IT BRKCOC-3005 for the Ethernet alternative. [S] read Sep 24, 2026: the sources listed under 2026 technology review in §1. [S] read Sep 27, 2026: the platform-vendor sources for §5 D1a, marked read Sep 27, 2026 in §1. [S*] a vendor figure taken on Sep 27, 2026 from the published summary of a vendor page that could not be opened. [A] typical-buyer assumption. **[8]** one of the eight delivery measures this design records. **The eight measures: 1 racks built · 2 first-time power-on rate · 3 burn-in hours per build · 4 on-time-to-go-live · 5 install time against unbuilt kit · 6 change orders per project · 7 first-reply time · 8 weekly throughput.** 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 · Scope

Who it is for. The head of research computing or the CTO buying an on-premises training or inference block in 2026, who has a data hall with a power number and needs the block to arrive built, burned in and cabled to a plan, on the platform NVIDIA ships in 2026 and with a path to the one it ships next.

Typical buyer [A]. 16 DGX B300 systems (a quarter SU), inference plus fine-tuning, an existing hall rated at 30 kW per rack on air, with budget for rear-door cooling.

What is in. The DGX B300 block, its compute and storage fabrics, management nodes, racks, power and cooling plan, staging and burn-in, acceptance, and the Rubin NVL8 upgrade path. **What is out.** Models and data pipelines; the enterprise network (SRA-NET-001); the second SU.

Background. The 2025 DGX SuperPOD reference architecture for B300 builds in scalable units of **64 DGX B300 systems**, up to **four per rack**, with rack power that "exceeds 50 kW" and a DC busbar power design [S]. The 2026 Rubin NVL8 architecture places **up to eight systems per rack** on **four 110 kW power shelves in N+N, up to 220 kW per rack**, with 800 Gb/s ConnectX-9 SuperNICs, BlueField-4 DPUs and a Quantum-X800 Q3400-RD 800 Gbps InfiniBand compute fabric, liquid-cooled, at 140 PFLOPS FP8 training and 400 PFLOPS FP4 inference per system [S]. Most enterprise halls are rated at 5–20 kW per rack. The design problem is power and cooling first, fabric second, compute third.

2026 technology review, read Sep 24, 2026 [S]. Every product, version, standard and date was checked against the vendor's or publisher's live page.

© 2026 Symbeeo. All rights reserved. · Public

Table 2. Scope

Item	Status found	Use in this design	Source
DGX B300	shipping; "Shipping Now"; ~14 kW per system; 8 × ConnectX-8 VPI, 2 × BlueField-3	used	nvidia.com/en-us/data-center/dgx-b300/
RA-11337-001 V01 (B300, Spectrum-4, DC busbar)	current, last updated Nov 19, 2025	used	docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-b300/latest/
RA-11339-001 V01 (B300, Quantum-X800, AC power)	current, last updated Sep 2, 2026; ~56 kW per rack; SU of 72	the AC-power alternative (§5 D2)	docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-b300-xdr/latest/
Cluster management	the RA's standard is now NVIDIA Mission Control, which includes Base Command Manager and Run:ai; current release 2.3.1 (BCM 11.33.1, Run:ai 2.25)	**NVIDIA Mission Control 2.3.1**	…/scalable-infrastructure-b300/latest/dgx-software.html · docs.nvidia.com/mission-control/
UFM	RA names UFM 3.5 Appliance, Enterprise Edition	**UFM 3.5 Appliance, Enterprise Edition**	…/scalable-infrastructure-b300/latest/dgx-superpod-components.html
DGX OS	DGX OS 7.6.0, Sep 14, 2026, supports DGX B300	pinned **DGX OS 7.6.0** as baseline (STP §3)	docs.nvidia.com/dgx/dgx-os-7-user-guide/release_notes.html
Rubin compute and storage fabric	Rubin NVL8 RA names Quantum-X800 Q3400-RD InfiniBand compute fabric, BlueField-4 DPUs to Spectrum-4 leafs, SN5600D storage fabric; last updated May 5, 2026	the RA's components (§4.6)	docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-rubinx86/latest/dgx-superpod-components.html
Rubin availability	Rubin products from partners in H2 2026 (NVIDIA, Jan 5, 2026)	date in §8	nvidianews.nvidia.com/news/rubin-platform-ai-supercomputer

© 2026 Symbeeo. All rights reserved. · Public

Item	Status found	Use in this design	Source
Spectrum-4 SN5610 / SN5600D	still the storage fabric in both B300 RAs and the Rubin RA; Spectrum-6 announced Jan 5, 2026 for Rubin	used; Spectrum-6 considered (§5 D5)	same RA pages · NVIDIA newsroom above
Cisco Ethernet AI switch	current Cisco AI switch: N9364E-SG2, Silicon One G200, 64 × 800G, RoCEv2 with PFC and ECN; datasheet updated Jul 31, 2026, no end-of-sale notice	named **N9364E-SG2** as the Cisco option	cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/nexus-9364e-sg2-switch-ds.html
Dell H19686.1	superseded by H19686.2 (Mar 2024), per Dell's own listings; the PDF could not be read, so H19686.2 was not read	cited for its document structure only; no product or figure is taken from it	delltechnologies.com (h19686-gen-ai-inferencing-dg.pdf)
NIST SP 800-88	Rev. 2 final, published September 2025	used in SDD §10	csrc.nist.gov/pubs/sp/800/88/r2/final
DGX B300, as a platform-vendor option	read Sep 27, 2026: 8 GPUs, ~14 kW, 10U, 8 × ConnectX-8 VPI up to 800 Gb/s, DGX OS, Mission Control, "three-year business-standard hardware and software support", shipping now	recorded in §5 D1a	nvidia.com/en-us/data-center/dgx-b300/

© 2026 Symbeeo. All rights reserved. · Public

Item	Status found	Use in this design	Source
NVIDIA HGX AI Factory enterprise RA	read Sep 27, 2026: HGX B300 systems, "ConnectX-8 HCA (on baseboard)", Spectrum-X Ethernet compute and converged fabrics, 32, 64 and 128 nodes; last updated May 18, 2026	named in §5 D1a as the RA of the OEM HGX path	docs.nvidia.com/enterprise-reference-architectures/hgx-ai-factory/latest/index.html
Dell PowerEdge XE9780 / XE9785	read Sep 27, 2026: XE9780, 8 × HGX B300 NVL8 270 GB 1,100 W SXM6 (or HGX B200), air cooling, 17.30 in high, 12 × 3,200 W PSUs, sold through Dell sales [S]; XE9785, 2 × AMD EPYC 9005 with 8 × HGX B300 or 8 × Instinct MI355X [S*]	recorded in §5 D1a	dell.com/en-us/shop/ipovw/poweredge-xe9780 · dell.com/en-us/shop/ipovw/poweredge-xe9785
HPE ProLiant Compute XD685	read Sep 27, 2026: 8 × HGX B300, B200, H200 or AMD Instinct MI355X; 5U direct liquid cooling or 6U air; liquid cooling for every GPU option, air cooling for H200 only; 8 nodes per rack in the 5U form [S*]	recorded in §5 D1a	hpe.com/us/en/compute/proliant-compute-xd685.html · buy.hpe.com (XD685 Direct Liquid Cooling Server, S4R96A)

© 2026 Symbeeo. All rights reserved. · Public

Item	Status found	Use in this design	Source
AMD Instinct MI355X platform and Helios	read Sep 27, 2026: MI355X platform, 8 × MI355X on UBB 2.0, 2.3 TB HBM3E, launched Jun 12, 2025; MI355X 288 GB HBM3E, 1,400 W typical board power, ROCm; Helios rack, 72 × MI455X, liquid cooled, volume deployments expected 2H 2026	recorded in §5 D1a	amd.com/en/products/accelerators/instinct/mi350/mi355x/platform.html · amd.com/en/products/accelerators/instinct/mi350/mi355x.html · amd.com/en/products/rackscale-solutions/helios.html

2 · Requirements register

Table 3. Requirements register

#	Requirement	Value	Source
R1	Block size	16 DGX B300, a quarter SU; inference plus fine-tuning	[A]
R2	Unit of sale	one SU of 64 DGX B300, or a fraction in whole racks	[S] RA-11337-001
R3	Partial-SU cabling	a partial SU still cables the full SU's leaf and spine positions	[S] RA-11337-001
R4	Rack density and power	up to 4 DGX B300 per rack; > 50 kW per compute rack, A + B	[S] RA-11337-001
R5	Existing hall	30 kW per rack on air	[A]
R6	Cooling	rear-door heat exchangers on every compute rack	[A]
R7	Compute fabric	800 Gbps XDR InfiniBand, rail-optimized, non-blocking, twin-plane fat tree; ConnectX-8 on each B300	[S] RA-11337-001
R8	Storage fabric	Spectrum-4 SN5610 / SN5600D at 800 Gbps Ethernet; storage access independent of compute nodes	[S] RA-11337-001
R9	Storage tiers	flash for active sets, object for archive	[A]
R10	Cluster and fabric management	NVIDIA Mission Control; UFM 3.5 Appliance, Enterprise Edition	[S] RA-11337-001, read Sep 24, 2026

© 2026 Symbeeo. All rights reserved. · Public

#	Requirement	Value	Source
R11	Scheduler and identity	Slurm or Kubernetes per client; access through the client's identity provider (SRA-SEC-001 R2)	[A]
R12	Security	management plane on its own network; DPU-enforced tenant isolation; certificates as SRA-SEC-001 §4.3; RBAC per the client's identity provider (SRA-SEC-001 R2)	[A]
R13	Staging	every rack built, cabled and burned in before shipment	[A]
R14	Acceptance	NCCL all-reduce across the block, a reference fine-tuning run, storage throughput to the storage fabric	[A]
R15	Growth	full SU of 64 in the same fabric positions; Rubin NVL8 in a liquid-cooled hall at up to 220 kW per rack	[S] both RAs

Fifteen rows; seven carry a vendor figure or rule, eight are typical-buyer assumptions.

3 · Architecture overview

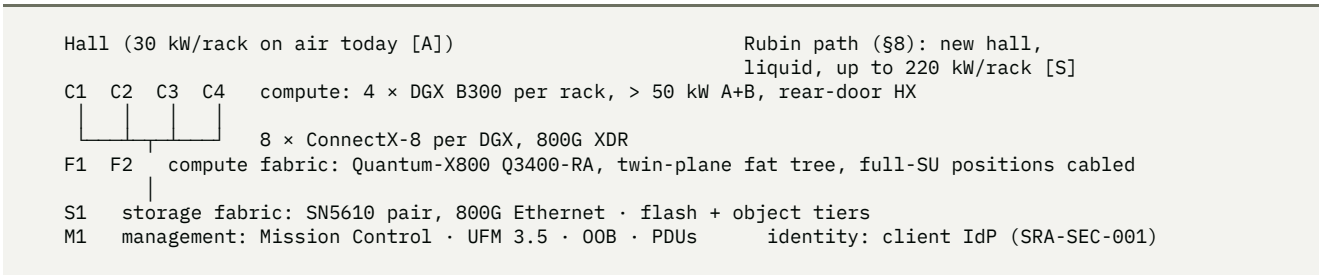


Figure 1. Block architecture

A quarter SU in eight racks (4 compute, 2 fabric, 1 storage, 1 management) [A], cabled for the full SU so growth fills positions rather than re-cables.

4 · Component design

4.1 Self-contained block.

Unit of sale: one SU of 64 DGX B300, or a fraction in whole racks [S]. Typical buyer: 16 systems in 4 racks at 4 per rack [A]. The RA's rule: a partial SU must still cable the full SU's leaf and spine positions [S].

4.2 Site.

Per compute rack: > 50 kW A + B (B300 at 4 per rack) [S]; DC busbar power distribution per the RA's title design [S], or AC PDUs per RA-11339-001 (~56 kW per rack) where the hall has no busbar path (§5 D2) [S]; floor loading and

© 2026 Symbeeo. All rights reserved. · Public

containment per site survey. Rubin path: 220 kW per rack with N+N shelves [S] — a different hall, not an upgrade in place.

4.3 Physical space.

4 compute racks + 2 fabric racks + 1 storage rack + 1 management rack for a typical buyer [A]; cabinet elevations drawn per rack before shipment (Cisco IT's practice, BRKCOC-3005 slides 26–27) [S].

4.4 Cooling.

B300 at 4 per rack exceeds air-only at 30 kW: rear-door heat exchangers on all compute racks [A]; the Rubin generation is liquid-cooled at rack scale [S].

4.5 Storage and data platforms.

Storage fabric on NVIDIA Spectrum-4 SN5610 / SN5600D at 800 Gbps Ethernet [S]; two tiers, flash for active sets and object for archive [A]; management nodes give storage access independent of compute nodes [S].

4.6 Fabric.

Compute fabric: NVIDIA Quantum-X800 Q3400-RA (Quantum-3), 800 Gbps XDR InfiniBand, rail-optimized, non-blocking, twin-plane fat tree [S]; ConnectX-8 on the B300, eight 800 Gb/s ports per system [S]. Ethernet alternative: Spectrum-X with Spectrum-4 and ConnectX-8 [S], or for a Cisco-standardized client, Cisco Nexus N9364E-SG2 (Silicon One G200, 64 × 800G) with PFC and ECN [S], as Cisco IT runs its own two clusters [S] BRKCOC-3005. Rubin path: ConnectX-9 SuperNICs at 800 Gb/s, Quantum-X800 Q3400-RD InfiniBand compute fabric, BlueField-4 DPUs to Spectrum-4 leaf switches [S].

4.7 Software, access and federation.

DGX OS 7; NVIDIA Mission Control, the RA's standard, which includes Base Command Manager and Run:ai [S]; Slurm or Kubernetes per client; UFM 3.5 Appliance, Enterprise Edition for the InfiniBand fabric [S]; access through the client's identity provider (SRA-SEC-001 R2) [A].

4.8 Cybersecurity.

Management plane on its own network; DPU-enforced tenant isolation (BlueField-3 on the B300 [S]); certificates as SRA-SEC-001 §4.3; RBAC per the client's identity provider (SRA-SEC-001 R2) [A].

5 · Options and platform choices

Table 4. Options and platform choices

ID	Decision	Requirement	Alternatives considered	Decision taken	Rationale
D1	Platform generation for a 2026 build	R1, R5, R15	DGX B300 (shipping) · DGX Rubin NVL8 (partner availability H2 2026, liquid-cooled, up to 220 kW per rack)	**DGX B300**; Rubin NVL8 recorded as the path (§8)	B300 fits an air hall with rear doors at > 50 kW per rack; Rubin needs a liquid-cooled hall at 220 kW per rack — a new-hall decision [S]

© 2026 Symbeeo. All rights reserved. · Public

ID	Decision	Requirement	Alternatives considered	Decision taken	Rationale
D1a	Platform vendor for the Blackwell Ultra build	R1, R2, R3, R4, R5, R6, R7, R10	NVIDIA DGX B300: 8 GPUs, ~14 kW, 10U, 8 × ConnectX-8 VPI up to 800 Gb/s, DGX OS, Mission Control, three-year NVIDIA hardware and software support [S] · Dell PowerEdge XE9780: 8-GPU HGX B300, air-cooled, 12 × 3,200 W PSUs [S]; XE9785 with AMD EPYC hosts, HGX B300 or MI355X [S*] · HPE ProLiant Compute XD685: HGX B300 in the 5U direct-liquid-cooled chassis only, air cooling for H200 only [S*] · AMD Instinct MI355X platform: 8 × MI355X, 288 GB HBM3E and 1,400 W each, ROCm [S]; AMD Helios: 72 × MI455X per liquid-cooled rack, volume 2H 2026 [S]	**DGX B300**	R2, R3, R7 and R10 come from the DGX SuperPOD RA, which is written for DGX systems, and the DGX B300 carries one vendor's support for the whole node, hardware and software [S]. The XD685 fails R5 and R6: its B300 build is liquid-cooled only [S*]. AMD fails R7 and R10 (ConnectX-8, Quantum-X800, UFM and Mission Control are NVIDIA's) and moves the software stack to ROCm; Helios needs a liquid-cooled hall, as Rubin does [S]. The Dell XE9780 meets R1's GPU count and generation and R5 and R6 as far as Dell's page shows (the same 8-GPU HGX B300, air-cooled) [S]; its per-system power is not on the page, so R4 is checked at the survey [A]. Its NVIDIA reference is the HGX AI Factory enterprise RA (HGX B300, ConnectX-8 on the baseboard, Spectrum-X Ethernet, 32/64/128 nodes) [S], not the SuperPOD RA. It becomes the choice where the client standardizes on Dell servers and Dell support, or takes the Ethernet option under D4; R2, R3, R7 and R10 are then re-drawn from that RA [A]
D2	Rack power delivery	R4, R5	DC busbar, MGX racks (RA-11337-001, > 50 kW) · AC PDUs in EIA racks (RA-11339-001, ~56 kW, SU of 72) · legacy PDUs inside RA-11337-001 ("still possible")	**DC busbar per RA-11337-001**; AC PDUs where the site survey finds no busbar path	RA-11337-001 is the RA the design follows [S]; the AC form changes the SU to 72, so the partial-SU cabling in R3 is re-drawn if it is chosen [S]; survey decides [A]
D3	Cooling	R5, R6	air only · rear-door heat exchangers · direct liquid	**Rear-door heat exchangers** on compute racks; fabric and storage on air	air only fails > 50 kW in a 30 kW hall [S]/[A]; direct liquid is the Rubin generation's hall [S]; rear doors are budgeted by the client [A]
D4	Compute fabric	R7	Quantum-X800 Q3400-RA XDR InfiniBand · Spectrum-X Ethernet (Spectrum-4 + ConnectX-8) · Cisco N9364E-SG2 with PFC and ECN	**Quantum-X800 InfiniBand** per the RA; Spectrum-X or Cisco for an Ethernet-standardized client	the RA's validated design [S]; Ethernet is supported by NVIDIA and run by Cisco IT on its own clusters [S] BRKCOC-3005; choice follows the client's operations team [A]
D5	Storage fabric	R8	Spectrum-4 SN5610 / SN5600D 800G Ethernet · Quantum QM9700 NDR 400G InfiniBand (RA option) · Spectrum-6 (announced Jan 5, 2026, not in any B300 RA)	**SN5610 pair**	named in RA-11337-001 and RA-11339-001 [S]; Spectrum-6 is not yet in a DGX RA [S]
D6	Storage tiers	R9	single flash tier · flash + object	**Flash + object**	active sets on flash, archive on object [A]; platform from the DGX SuperPOD certified list (SDD P4)

© 2026 Symbeeo. All rights reserved. · Public

ID	Decision	Requirement	Alternatives considered	Decision taken	Rationale
D7	Cluster management and scheduler	R10, R11	Mission Control (BCM + Run:ai) with Slurm · Mission Control with Kubernetes	**Mission Control**; Slurm or Kubernetes per client (SDD P5)	Mission Control is the RA's standard [S]; scheduler follows the client's workload [A]
D8	Growth route	R3, R15	re-cable at growth · cable the full SU now · upgrade in place to Rubin	**Cable the full SU now**; Rubin in a new hall	the RA's partial-SU rule [S]; 220 kW liquid-cooled racks cannot go into a 30 kW air hall [S]

6 · Failure conditions

Conditions and expected behavior come from the RA's design (twin planes, A + B feeds) and the pass criteria already in STP-AI-001; the recovery times are filled in when the cases run. The engineer completes this table at acceptance.

Table 5. Failure conditions

#	Condition	What the user sees	Recovers by itself	Needs a hand	Time	Tested by
F1	One DGX node down under a job	job stops on that node	scheduler drains the node; job restarts from checkpoint	replace or repair node		STP T10
F2	One compute-fabric leaf down	none expected	traffic on the other plane [A]	replace switch		STP T11
F3	One power feed lost on a compute rack	none	rack continues on the other feed (A + B)	restore feed		STP T12
F4	One rear-door heat exchanger stopped	throttling alarm	job survives for the agreed window [A]	restore cooling		STP T13
F5	Partial-SU growth	none at growth	full-SU leaf and spine positions already cabled [S]	populate positions		inspection, STP S1
F6	Site power lost	block dark	no	—	until power	waived: site power is outside the block's scope

7 · Validation

Staging. Every rack built, cabled and burned in before shipment. Burn-in hours per build **[8]** measure 3; first-time power-on rate **[8]** measure 2; on-time to go-live **[8]** measure 4.

Acceptance. NCCL all-reduce across the block, a reference fine-tuning run, storage throughput to the storage fabric; results to be recorded once a build exists; until then this document has no validation.

8 · Operations and lessons learned

Lessons. Lessons every peer reports: NCCL variability, kernel panics, "fighting bugs that have already been solved" (CSCS, CUG 2025). The staging period exists to find them before the client does. Lessons from a Symbeeo build: to be

© 2026 Symbeeo. All rights reserved. · Public

written by the signing engineer after the first build under this document.

Growth and outlook. Growth to a full SU of 64 in the same fabric positions [S]; the Rubin NVL8 generation, from partners in H2 2026, in a liquid-cooled hall at 220 kW per rack [S]; Ethernet fabrics closing on InfiniBand (Spectrum-X, and Spectrum-6 with co-packaged optics announced for Rubin [S]; Cisco Silicon One G200 [S]).

---

9 · Appendix

Bill of materials, typical buyer [S]/[A]: 16 × DGX B300 · 4 × compute racks · compute fabric: Quantum-X800 Q3400-RA leaf and spine per the RA's SU tables, XDR cables and transceivers · storage fabric: SN5610 pair, 800G optics · management nodes · NVIDIA Mission Control · UFM 3.5 Appliance, Enterprise Edition · rear-door heat exchangers ×4 · DC busbar or PDUs at > 50 kW per rack, A/B. **References:** RA-11337-001 · RA-11339-001 · Rubin NVL8 RA · Isambard-AI 2410.11199 · BRKCOC-3005 · Cisco N9364E-SG2 datasheet · NVIDIA HGX AI Factory enterprise RA, Dell PowerEdge XE9780 and XE9785, HPE ProLiant Compute XD685, AMD Instinct MI355X platform and Helios pages (§5 D1a) · Dell H19686.1 (a 52-page validated design; superseded by H19686.2; cited for its document structure only; no product or figure is taken from it) · the sources in §1's technology review.

9.1 · Industry rules this architecture helps a buyer meet.

No industry rule in Symbeeo's industry research of September 2026 is linked to this document's requirements, so none is named here.

© 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.

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