SYMBEEO

PUBLIC

DESIGN STUDY

Broadcast render and AI upscaling farm for a media group

Complete design on current platforms

CAPABILITY	INDUSTRY	SETTING
AI & Data	Media & Gaming	Broadcaster
DOCUMENT	**VERSION**	**DATE**
AI-005	1.3	Oct 10, 2026

WRITTEN BY

Diksha Shah, Principal AI Infrastructure Architect

© 2026 Symbeeo. All rights reserved.

Document control

Document	AI-005	Version	1.3
Issued	Oct 10, 2026	Status	Issued by Symbeeo
Classification	Public	Supersedes	1.2, Sep 27, 2026
Capability	AI & Data	Industry	Media & Gaming
Author	Diksha Shah	Role	Principal AI Infrastructure Architect
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. Design study: every figure is vendor-published, a design target or a marked assumption.		
Sign-off	"An engagement under this design is signed by the practice lead before any build proceeds."		

Contents

1 · Scope 3
1c · Public implementations this design follows 4
2 · Requirements register 4
3 · Architecture overview 5
4 · Design by subsystem 6
5 · Failure behavior 6
6 · Limits 7
7 · Lessons learned 7
8 · The author 7

© 2026 Symbeeo. All rights reserved. · Public

CONVENTIONS

Source tags. [S] figure read from the vendor's own current page on Sep 18, 2026 · [S*] figure from the vendor's page summary where the page itself could not be opened · [A] design assumption, to be confirmed by the signing engineer · [T] design target taken from the study record or derived from it. No figure comes from any prior engagement, file or asset. 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 18, 2026	1.0	First issue.
Sep 20, 2026	1.1	Public implementations added; figures reconciled with published sources.
Sep 27, 2026	1.2	Figures and platform details rechecked against published sources.
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 technology of a broadcaster whose overnight render and archive upscaling queue no longer finishes before the morning, who needs a farm that renders, upscales and transcodes on the same GPUs and reads the media store fast enough to keep them busy.

Typical buyer [A]. A broadcaster with two production centers: 2 sites, 1 primary farm at the larger site, 3 workloads (3D and motion-graphics render, AI upscaling of an SD and HD archive to UHD, transcode for distribution), a 6-hour overnight batch window, a media store of 4 PB active on an all-flash NAS.

What is in. A 48-GPU farm at the primary site; the 200 Gb/s path from each node to the media store; the render and upscaling job queue; the failure behavior. **What is out.** The second site's farm (a copy of this design at smaller scale, sized later); the media asset management system; the distribution network; the archive ingest path.

Assumptions. Dell is the server and storage vendor, NVIDIA the GPU vendor, AMD the CPU vendor [A]. Render and upscaling software supports the RTX PRO 6000 Blackwell Server Edition [A]. The media store is Dell PowerScale all-flash [A].

Related documents, read Sep 18, 2026. Dell PowerEdge XE7745 product page (https://www.dell.com/en-us/shop/ipovw/poweredge-xe7745) · Dell press release, Nov 17, 2025, enterprise AI with NVIDIA (https://www.dell.com/en-us/dt/corporate/newsroom/announcements/detailpage.press-releases~usa~2025~11~dell-technologies-and-nvidia-advance-enterprise-ai-innovation.htm) · NVIDIA RTX PRO 6000 Blackwell Server Edition (https://www.nvidia.com/en-us/data-center/rtx-pro-6000-blackwell-server-edition/) · AMD EPYC 9005 Series processors (https://www.amd.com/en/products/processors/server/epyc/9005-series.html) · Dell PowerScale F910 all-flash NAS platform (https://www.dell.com/en-us/shop/ipovw/powerscale-f910) · Dell PowerScale All-Flash specification sheet H15963.26, read Sep 20, 2026 (https://www.delltechnologies.com/assetlink/doc/en-us/h15963-ss-powerscale-all-flash-nodes-en-dl1m3kd-original.pdf) · Dell AI Factory for Government with NVIDIA, RTX PRO 6000 Blackwell Server Edition configuration on Dell PowerEdge servers, datasheet, read Sep 20, 2026 (https://www.delltechnologies.com/asset/en-us/products/storage/technical-support/dell-ai-factory-with-nvidia-rtx-pro-6000-on-dell-poweredge-servers-datasheet.pdf).

© 2026 Symbeeo. All rights reserved. · Public

1c · Public implementations this design follows

Deployment rows are real public deployments, each dated. Where no public deployment matches, the row is Vendor documentation: the vendor's document that this design, Symbeeo's own, is built on.

Table 2. Public implementations this design follows

Implementation	Who published it · date	What it settles for this design	URL
Vendor documentation. Dell AI Factory for Government with NVIDIA: endorsed RTX PRO 6000 Blackwell Server Edition configuration on PowerEdge XE7745, 8 GPUs per node, 4-node and 4–32-node configurations	Dell Technologies · Mar 24, 2026 (the PDF's document date; no date is printed on the page), read Sep 27, 2026	The XE7745 with 8 × RTX PRO 6000 as a shipped configuration (R5, R6); about 8 kW per node and 120–160 kW per rack for the hall check in §6	https://www.delltechnologies.com/asset/en-us/products/storage/technical-support/dell-ai-factory-with-nvidia-rtx-pro-6000-on-dell-poweredge-servers-datasheet.pdf
Vendor documentation. Dell PowerScale All-Flash specification sheet, F910	Dell Technologies · H15963.26, Jun 8, 2026 (the PDF's document date; © 2026 printed), read Sep 27, 2026	F910 node figures in R10 and the front-end NIC options in R11	https://www.delltechnologies.com/assetlink/doc/en-us/h15963-ss-powerscale-all-flash-nodes-en-dl1m3kd-original.pdf

2 · Requirements register

Table 3. Requirements register

#	Requirement	Value	Source
R1	GPU count	48 GPUs for render and upscaling	[T]
R2	GPU	NVIDIA RTX PRO 6000 Blackwell Server Edition, 96 GB GDDR7, 1,597 GB/s, up to 600 W	[S] NVIDIA
R3	GPU media engines	4 × NVENC (9th gen), 4 × NVDEC (6th gen) per GPU	[S] NVIDIA
R4	GPU ray tracing	188 RT cores, 355 TFLOPS RT	[S] NVIDIA
R5	Node	Dell PowerEdge XE7745, 4U, 8 × PCIe Gen5 double-width up to 600 W	[S] Dell XE7745
R6	Node count	6 nodes × 8 GPUs	[T] derived from R1 and R5
R7	Node CPU	2 × AMD EPYC 9005, up to 192 Zen 5 cores per socket	[S] Dell XE7745; [S] AMD
R8	Node memory	up to 3.072 TB DDR5, 24 DIMM slots, up to 6,000 MT/s	[S] Dell XE7745
R9	Node to media store	200 Gb/s per node	[T]
R10	Media store	PowerScale F910, 2U, 24 × NVMe SSD, 92 TB–2.9 PB raw per node, 3–252 nodes	[S] Dell F910

© 2026 Symbeeo. All rights reserved. · Public

#	Requirement	Value	Source
R11	Media store front-end	per F910 node, one dual-port NIC: 25G, 100G, 200G or 400G Ethernet, or 200G HDR InfiniBand; dual-port 100G chosen	[S] Dell PowerScale spec sheet H15963.26 (https://www.delltechnologies.com/assetlink/doc/en-us/h15963-ss-powerscale-all-flash-nodes-en-dl1m3kd-original.pdf); [A] choice
R12	Batch window	6 hours overnight; cut against the current farm, measured at first build	[T]
R13	Local scratch per node	up to 8 × E3.S Gen5 NVMe, 122.88 TB max	[S] Dell XE7745
R14	Availability	95 % of GPUs schedulable at the start of every window	[A]

Fourteen rows; nine carry a vendor figure, four are design targets, one is a design assumption.

3 · Architecture overview

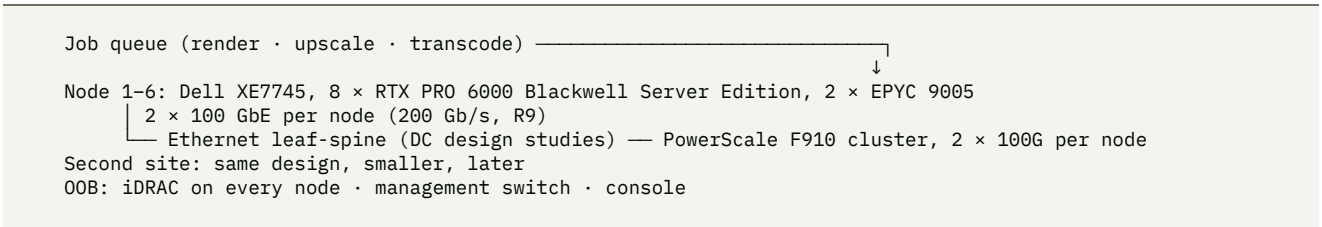


Figure 1. Render and upscaling farm

Six 8-GPU PCIe nodes; the same GPU renders, upscales and transcodes so the queue moves work between job types without moving hardware; two 100 GbE links per node meet the 200 Gb/s target.

Table 4. Architecture overview

Role	Product (exact model)	Qty	Capacity used	Why this unit
GPU node	Dell PowerEdge XE7745, 4U	6	8 of 8 double-width slots each [S]	Dell names XE7745 with RTX PRO 6000 Blackwell Server Edition [S] press release
GPU	NVIDIA RTX PRO 6000 Blackwell Server Edition, 96 GB	48	48 of 48 [T]	RT cores for render, 4 NVENC + 4 NVDEC for media, FP4 for upscaling models on one part [S]
CPU	AMD EPYC 9005 series, 2 per node	12	model chosen at first build [A]	Up to 192 cores per socket for CPU-side render stages [S]
Media store	Dell PowerScale F910, all-flash	6 nodes [A]	4 PB active [A] of 92 TB–2.9 PB raw per node [S]	Scale-out NAS; dual-port 100G chosen [A] from the 25G to 400G NIC options [S]
Node NIC	2 × 100 GbE per node	12	200 Gb/s per node [T]	Meets R9 on two links
Scratch	E3.S Gen5 NVMe in each XE7745	6 sets	per-job frames [A]	Keeps intermediate frames off the store [A]

© 2026 Symbeeo. All rights reserved. · Public

4 · Design by subsystem

4.1 GPU nodes.

Dell PowerEdge XE7745: 4U, two AMD EPYC 9005 processors at up to 192 Zen 5 cores each, up to 3.072 TB DDR5, eight PCIe Gen5 double-width accelerators at up to 600 W, up to 8 × E3.S Gen5 NVMe, 3,200 W Titanium PSUs [S]. Eight RTX PRO 6000 Blackwell Server Edition per node [S] press release pairing, [A] eight fitted. Rejected: 2U nodes with 2 double-width GPUs — 24 nodes for 48 GPUs, more switch ports and more rack space for the same count [A].

4.2 GPU.

RTX PRO 6000 Blackwell Server Edition: 96 GB GDDR7, 1,597 GB/s, 24,064 CUDA cores, 188 RT cores at 355 TFLOPS, 4 PFLOPS FP4, four NVENC and four NVDEC, up to 600 W, dual-slot air or single-slot liquid [S]. One part covers all three workloads (R2–R4). Rejected: a compute-only data-center GPU — no RT cores and no media engines for render and transcode [A].

4.3 Media store.

Dell PowerScale F910: 2U, 24 NVMe SSDs of 3.84–30.72 TB, 92 TB to 2.9 PB raw per node, 3 to 252 nodes per cluster, 512 GB per node, one dual-port front-end NIC at 25G, 100G, 200G or 400G Ethernet [S]. Six nodes for 4 PB active, dual-port 100G fitted [A]; NFS to the render nodes [A].

4.4 Node-to-store network.

Two 100 GbE ports per XE7745 into the leaf-spine, 200 Gb/s per node (R9) [T]; the F910's dual-port 100G per node [A] gives 1.2 Tb/s aggregate across six store nodes [A] derived; the 200G and 400G NIC options [S] are the upgrade path. Rejected: 1 × 200 GbE per node — one link, one failure takes the node off the store [A].

4.5 Job queue.

One queue with three job classes; GPUs are not partitioned by class, so the overnight window fills with whatever is pending [A]. The window-time cut (R12) is measured at first build against the current farm and recorded here; no figure is stated now [T].

4.6 Scratch.

Intermediate frames on the node's local E3.S NVMe (up to 122.88 TB [S]); only finished output returns to the store [A].

4.7 Management and out-of-band.

iDRAC on every node, management switch and console on their own VLAN [A].

5 · Failure behavior

Table 5. Failure behavior

Failure	Behavior	Source
One GPU fails	node keeps 7 GPUs schedulable; queue routes around it	[A]
One node fails	8 of 48 GPUs lost; window runs at 40 GPUs; jobs requeue	[A]
One 100 GbE link on a node fails	node continues at 100 Gb/s to the store	[A]

© 2026 Symbeeo. All rights reserved. · Public

Failure	Behavior	Source
One F910 node fails	cluster continues; capacity and front-end bandwidth reduced by one node	[S] 3-node minimum; [A] behavior
Job queue controller fails	running jobs finish; no new dispatch until restart	[A]
Site power lost	farm dark until power; second site unaffected	[A]

6 · Limits

- The batch-window reduction is a target measured at first build; this document states none.
- Render and upscaling software support for the Blackwell Server Edition GPU is assumed, not verified against any software vendor's page.
- The F910 product page lists 100G front-end NICs; the specification sheet H15963.26 also lists 200G and 400G Ethernet. This design fits dual-port 100G; store-side aggregate depends on the node count and the NIC chosen.
- The second production center is out of scope; a shared queue across sites is not designed.
- Air-cooled at 600 W per GPU, 8 per node: hall airflow and rack power are checked at first build against the XE7745 fan and PSU figures and Dell's about 8 kW per node (§1c).

7 · Lessons learned

None yet — this is a design study; populated after the first build against it.

8 · The author

WRITTEN BY

Diksha Shah

Principal AI Infrastructure Architect

NVIDIA Certified Professional: AI Operations, NVIDIA Certified Associate: AI Infrastructure and Operations, Google Cloud Professional Cloud Architect, AWS Certified Data Engineer

Signed off by Edwin Ozer, Founder.

© 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.

AI-005 · Version 1.3 · Oct 10, 2026