[01]

REFERENCE DESIGN

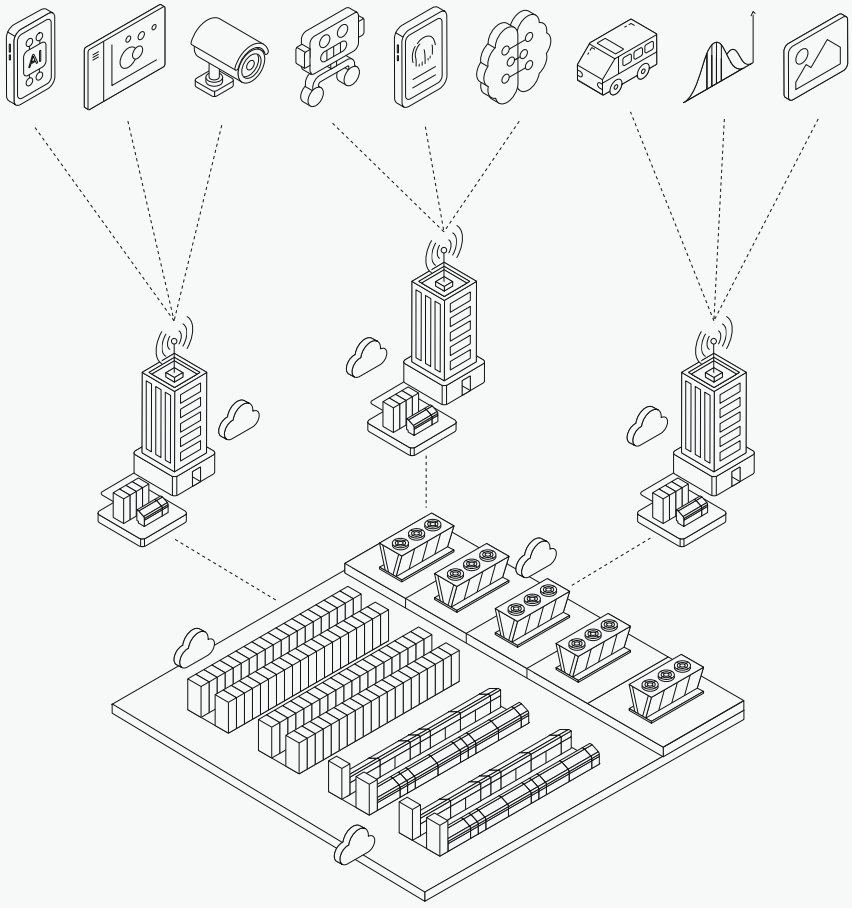
AI infrastructure ready to deploy:

Hybrid liquid cooling architecture for scalable AI inference

The global demand for GPUs has surged, fueled by the exponential growth of AI training and inference. Yet, the combination of high hardware costs and complex infrastructure often slows down scalability and time-to-deployment.

2CRSi and Submer have partnered to deliver a ready-to-deploy reference design that simplifies the path to AI infrastructure. Fully integrated and plug-and-play by design, this solution combines immersion-cooled GPU clusters with networking, storage, and management in a pre-validated architecture.

This approach enables customers to accelerate deployment, reduce integration risks, and scale AI workloads quickly — while maintaining efficiency, reliability, and sustainability at every step.



[02]

Technology set up

The selected nodes for this solution are a combination of highly efficient 2CRSi immersion servers and Submer's SmartPod EXO product. This solution will allow customers to maximize efficiency and extend the lifespan of IT, protecting it against dust, moisture and vibrations while keeping a long-stable temperature.

We can scale to any size clusters with NVIDIA's high-end GPU's such as H200 NVL, or RTX PRO 6000 Blackwell. This cluster also includes the same elements as a top tier market turnkey solution for DC Design (Power, Network and Fiber Cables, Patch Panels, PDUs, etc.) This reference design delivers a 2 MW AI-ready architecture built on a hybrid model, combining immersion cooling for compute with air-cooled racks for networking, storage, and management.

2MW AI-ready architecture

This reference design delivers a 2 MW AI-ready architecture built on a hybrid model, combining immersion cooling for compute with air-cooled racks for networking, storage, and management.



Compute power in immersion twenty-four

Submer SmartPod EXO tanks integrate all the GPU compute power, hosting up to 2688 GPUs. Immersion cooling ensures maximum efficiency, stable fluid operating temperatures, and extended hardware lifespan, while enabling intensive AI and HPC workloads.



High-bandwidth networking

Dedicated network rack houses Spectrum-X switches, providing low latency and high throughput. This ensures seamless connectivity across the entire cluster, fully aligned with NVIDIA's reference architecture for AI Factories.



Storage & management

One combined storage and management rack supports high-performance storage platforms and cluster orchestration nodes, ensuring resilience, redundancy, and smooth operations at scale.

By leveraging immersion for compute and air-cooled racks for network and storage, this 2 MW / 2688 GPU design delivers the best of both worlds: exceptional energy efficiency, optimized cooling, and robust scalability for next-generation AI data centers.

[02]

Tech-mix

GPU

NVIDIA H200 NVL

- 141GB Memory
- 4,8 TB/s GPU-to-GPU Bandwidth
- FP64 Tensor Core - 60 TFLOPS
- 600W TDP



SERVER

Atlas 1.8 SP5-D 3OU 21”

- Dual AMD EPYC 9005/9004 Turin
- x12 - channel per CPU, 12 RDIMM per CPU
- x8 -7/15mm NVMe PCIe 5.0 x4 hot-swappable bays [4 from each CPU PLX]
- 1x 1GbE (Realtek RTL8211E) for dedicated management GLAN
- 2x 1GbE LAN ports (Intel i350-AM2), support NCSI function
- 2x PCIe x16 (Gen5 x link), HH-HL [PLX CPU0 and CPU1] + 1x PCIe x8 (Gen5 x link), HH-HL [CPU1]



TANK

SmartPod EXO TwinTank

- Optimized Immersion Tank for smooth & predictable daily IT Operations
- 99,999% SLA Availability (Full-Redundant | HA)
- FWS Inlet of 35°C and Fluid Setpoint up to 45°C for high-efficiency
- Improved Cable Management for structured layouts



FLUID

ENEOS Type H

- Meets the OCP (Open Compute Project) standard of over 300°C for spontaneous auto-ignition
- Low viscosity fluid type for improved circulation efficiency and reduced power pump consumption
- Excellent oxidation stability and robust dielectric strength properties

[03]

Datacenter set up

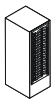
FIG. 1
DC aerial view



Pipes
Facility Water Supply (95°F/35°C inlet)



Power / Cabling



Air-Rack



Tank

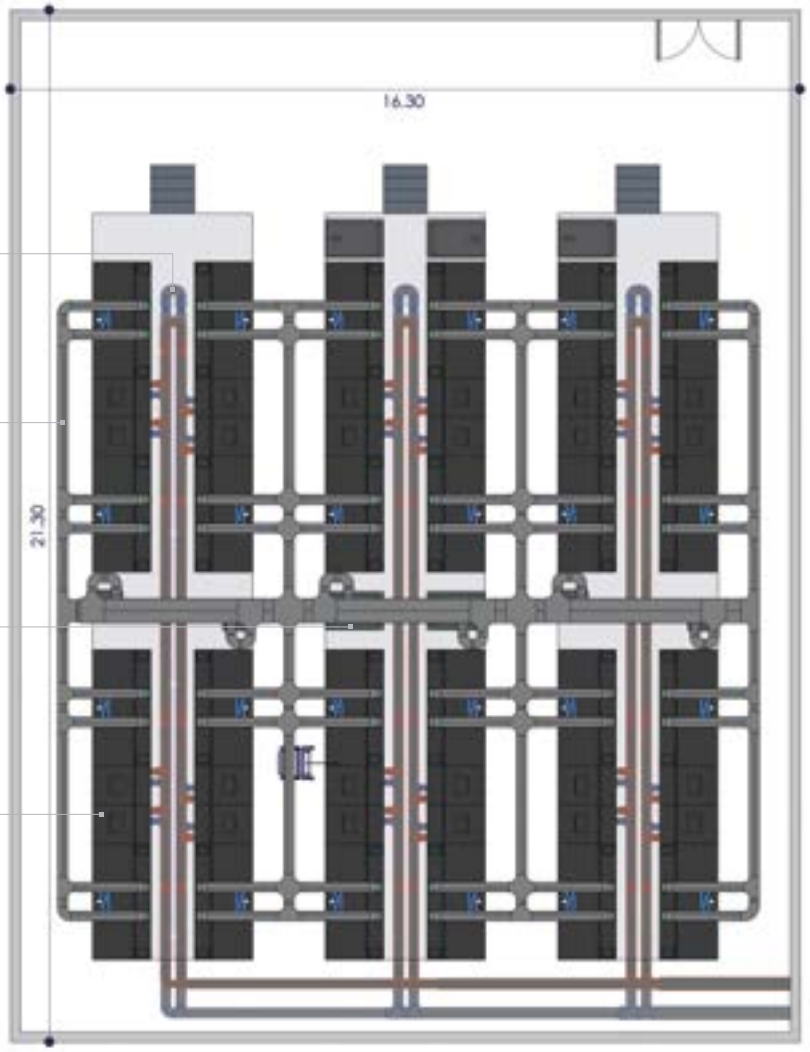
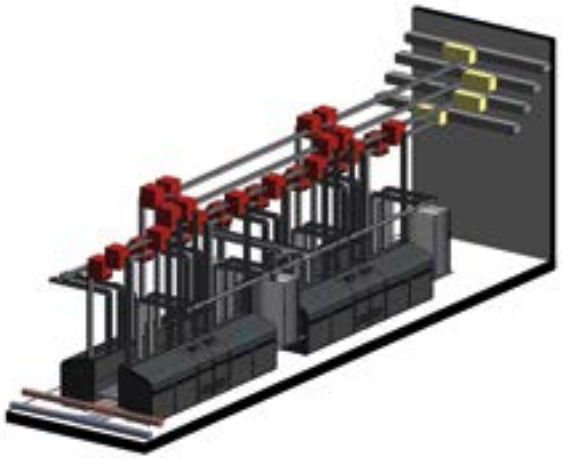


FIG. 2
Full clustered data hall

- Tap off from the main busbar to the secondary busbar
- Tap off from secondary busbar to the SmartPods



[03]

Datacenter set up

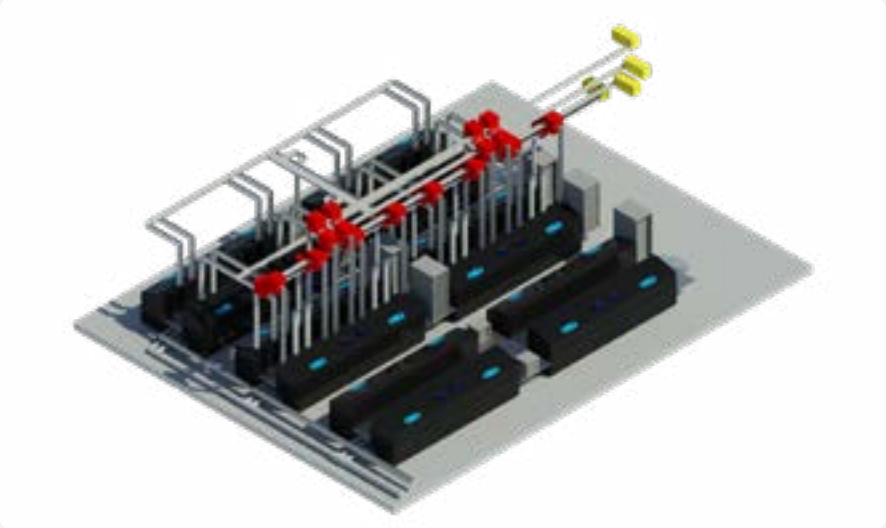


FIG. 3
Aerial cluster view



FIG. 4
Row cluster connectivity



FIG. 5
Aerial cluster view

[04]

Datasheets / SKUs

COOLING SKUS
(Tanks, racks and fluid)

PART#	QTY	MANUFACTURER	DESCRIPTION
SP_51S1TWI2HA (Tanks)	12	Submer	SmartPod EXO TwinTank S100
Type H (Liters)	40000	ENEOS	ENEOS Type H Fluid
AR3300B2 (Racks)	5	Schneider Electric	APC NetShelter SX Server Rack, 42U, 1991H x 600W x 1200D

COMPUTE SKUS
(Servers, switches and storage)

PART#	QTY (per tank)	MANUFACTURER	DESCRIPTION
SRV-A21-308-SP5D-U573 (Server)	5	2CRSi	High-End GPU Server (H200 NVL)
SRV-A21-308-SP5D-U574 (Server)	9	2CRSi	Mid-Range GPU Server (RTX6000 PRO)
920-9N110-00R1-0C0 (Switch)	1	NVIDIA	OOB Network Switch

PART#	QTY (per rack)	MANUFACTURER	DESCRIPTION
SRV-CE1-112G0-R101 (Server)	12	2CRSi	High-Performance Storage Server for WEKA
SRV-CE1-112G0-R142 (Server)	2	2CRSi	Compute Server
SRV-CE1-112G0-R112 (Server)	5	2CRSi	Management Server
SRV-CE1-112G0-R132 (Server)	4	2CRSi	Storage Server
AFF-A20-20639667 (Server)	1	2CRSi	NFS Storage
920-9N110-00R1-0C0 (Switch)	9	NVIDIA	OOB Network Switch
920-9N42F-00RI-3C1 (Switch)	28	NVIDIA	Compute Fabric Switch Lead / Leaf / Spine
920-9N42F-00RI-7C0 (Switch)	2	NVIDIA	Common Network Switch S-Leaf

NETWORK SKUS
(Fiber cables, copper cables and patch pannels)

PART#	QTY (per tank)	MANUFACTURER	DESCRIPTION
09 48 474 774 3 005	18	Harting	1G RJ45 cable
TQS-P167C-P8M01	28	FormericaOE	200G QSFP112 Immersion AOP
TQS-P167C-P8M01	56	FormericaOE	400G QSFP112 Immersion AOP
KAXCSV-02408-C001	1	Corning	Patch Panel RJ45
FHD-1U-CMP400 (SKU:70419)	3	FS.com	FHD High Density 1U Rack Mount Modular Enclosure
FHD-FAP12MTPHA (SKU:35510)	12	FS.com	FHD Fiber Adapter Panel
09 48 474 774 3 030	19	Harting	Cat6A Ethernet cable, 3m
TES-K18HF-8CZ	42	FormericaOE	800G OSFP Transceiver (for SRV links)
980-9I510-F4NS00	14	NVIDIA	800G OSFP Transceiver (for SW uplink ports)

[04]

Datasheets / SKUs

NETWORK SKUS
(Fiber cables, copper cables and patch pannels)

PART#	QTY (per rack)	MANUFACTURER	DESCRIPTION
09 48 474 774 3 005	84	Harting	1G RJ45 cable
QSFP-SR4-100G (SKU:71009)	120	FS.com	100G QSFP28 Transceiver (for SW uplink ports)
980-9I570-00N030	440	NVIDIA	PFC MPO-12 30M
QSFP-400G-AO02 (SKU:344599)	256	FS.com	400G QSFP112 AOC 2M
980-9I510-F4NS00	880	NVIDIA	800G OSFP Transceiver
OSFP-400G-4Q4xNAO02 (SKU:342395)	2	FS.com	400G OSFP to 4xQSFP28 AOC

POWER SKUS
(PDUs and power cables)

PART#	QTY (per tank)	MANUFACTURER	DESCRIPTION
PX3-5048JV-E2V2	3	Legrand	PDU 400V, 3ΦWye, 63A, 43.6kVA, (x24) IEC60320 C13 Locking, 10A, (x12) IEC60320 C19 Locking, 16A
PC5009	15	Scolmore	Power cord Blue (2M)
PC5003	29	Scolmore	Power cord Red (2M)

PART#	QTY (per rack)	MANUFACTURER	DESCRIPTION
PX3-1703-V2	18	Legrand	PDU - 400V, 3ϕ Wye, 32A, 22kVA, (x36) IEC320 C13 10A
PC5007	9	Scolmore	Power cord Blue (1M)
PC5001	9	Scolmore	Power cord Red (1M)
PC5009	70	Scolmore	Power cord Blue (2M)
PC5003	70	Scolmore	Power cord Red (2M)

OPERATIONS SKUS
(Operations cart or crane lift)

PART#	QTY (per tank)	MANUFACTURER	DESCRIPTION
AC_SPLIFTBC-159B	1	Submer	Server Lift

2crsi | Submer | ENEOS

Learn more at 2crsi.com