

Isambard AI

A National AI Research Resource

Matt Williams

University of Bristol



University of
BRISTOL



**SOCIETY OF RESEARCH
SOFTWARE ENGINEERING**



-  Isambard
-  Engineering and Physical Sciences Research Council
- GW4
-  UNIVERSITY OF BATH
-  University of BRISTOL
-  CARDIFF UNIVERSITY
PRIFYSGOL CAERDYDD
-  UNIVERSITY OF EXETER
-  Met Office
- arm
-  MARVELL

CRAY



CRAY





Isambard 2

- ~25,000 CPU cores
- Multiple architectures
 - Mostly ARM
- Infiniband interconnect
- Fujitsu, Cray, Arm and GNU compilers

- Coming to end of life

Isambard 3: a new class of Arm-based CPU supercomputer



>2 × Isambard 2

- **55,000 ARM cores**
NVIDIA Grace-Grace
- Slingshot 11
- 2 PiB Sonexion
- 300kW
- £9M including the
Modular Data Centre
- Updated MACS
- In production **Summer
2024**

REAR View

Rack#		κ3010		κ3009		κ3008		κ3007		κ3004		κ3003		MACS RACK κ3002		SERVICE NODE RACK κ3001		STORAGE RACK κ3000	
UID	800mm x1200mm	UID	800mm x1200mm	UID	800mm x1200mm	UID	800mm x1200mm	UID	800mm x1200mm	UID	800mm x1200mm	UID	800mm x1200mm	UID	600mm x1200mm	UID	800mm x1200mm	UID	800mm x1200mm
42	Slingshot Switch sw-hsn28	42	Slingshot Switch sw-hsn24	42	Slingshot Switch sw-hsn20	42	Slingshot Switch sw-hsn16	42	Slingshot Switch sw-hsn12	42	Slingshot Switch sw-hsn08	42	Slingshot Switch sw-hsn04	42	BLANK	42	JL625A/R9F65A sw-25g02	42	8360/R9G13A sw-lmn2
41	Slingshot Switch sw-hsn27	41	Slingshot Switch sw-hsn23	41	Slingshot Switch sw-hsn19	41	Slingshot Switch sw-hsn15	41	Slingshot Switch sw-hsn11	41	Slingshot Switch sw-hsn07	41	Slingshot Switch sw-hsn03	41	BLANK	41	JL625A/R9F65A sw-25g01	41	8360/R9G13A sw-lmn1
40	Slingshot Switch sw-hsn26	40	Slingshot Switch sw-hsn22	40	Slingshot Switch sw-hsn18	40	Slingshot Switch sw-hsn14	40	Slingshot Switch sw-hsn10	40	Slingshot Switch sw-hsn06	40	Slingshot Switch sw-hsn02	40	Slingshot Switch sw-hsn00	40	Slingshot Switch sw-hsn02	40	ew1 (repurposed FMN MAC)
39	Slingshot Switch sw-hsn25	39	Slingshot Switch sw-hsn21	39	Slingshot Switch sw-hsn17	39	Slingshot Switch sw-hsn13	39	Slingshot Switch sw-hsn09	39	Slingshot Switch sw-hsn05	39	Slingshot Switch sw-hsn01	39	Slingshot Switch sw-hsn03	39	Slingshot Switch sw-hsn01	39	CDS 3
38	BLANK	38	BLANK	38	BLANK	38	BLANK	38	BLANK	38	BLANK	38	BLANK	38	BLANK	38	BLANK	38	CDS 2
37	BLANK	37	BLANK	37	BLANK	37	BLANK	37	BLANK	37	BLANK	37	BLANK	37	6300M /R9F63A sw-smn04	37	6300M /R9F63A sw-smn02	37	CDS 1
36	6300M /R9F63A sw-smn22	36	6300M /R9F63A sw-smn19	36	6300M /R9F63A sw-smn16	36	6300M /R9F63A sw-smn13	36	6300M /R9F63A sw-smn10	36	6300M /R9F63A sw-smn07	36	6300M /R9F63A sw-smn03	36	6300M /R9F63A sw-smn03	36	6300M /R9F63A sw-smn01	36	smu 1
35	6300M /R9F63A sw-smn21	35	6300M /R9F63A sw-smn18	35	6300M /R9F63A sw-smn15	35	6300M /R9F63A sw-smn12	35	6300M /R9F63A sw-smn09	35	6300M /R9F63A sw-smn06	35	6300M /R9F63A sw-smn02	35	BLANK	35	BLANK	35	BLANK
34	6300M /R9F63A sw-smn20	34	6300M /R9F63A sw-smn17	34	6300M /R9F63A sw-smn14	34	6300M /R9F63A sw-smn11	34	6300M /R9F63A sw-smn08	34	6300M /R9F63A sw-smn05	34	6300M /R9F63A sw-smn01	34	BLANK	34	BLANK	34	mdu 1
33		33		33		33		33		33		33		33	BLANK	33	BLANK	33	BLANK
32	cn-0383	32	cn-0382	32	cn-0319	32	cn-0318	32	cn-0255	32	cn-0254	32	cn-0191	32	cn-0190	32	cn-0063	32	cn-0062
31	cn-0384	31	cn-0381	31	cn-0320	31	cn-0317	31	cn-0256	31	cn-0253	31	cn-0192	31	cn-0189	31	cn-0064	31	cn-0061
30	cn-0379	30	cn-0378	30	cn-0315	30	cn-0314	30	cn-0251	30	cn-0250	30	cn-0187	30	cn-0186	30	cn-0059	30	cn-0058
29	cn-0380	29	cn-0377	29	cn-0316	29	cn-0313	29	cn-0252	29	cn-0249	29	cn-0188	29	cn-0185	29	cn-0060	29	cn-0057
28	cn-0375	28	cn-0374	28	cn-0311	28	cn-0310	28	cn-0247	28	cn-0246	28	cn-0183	28	cn-0182	28	cn-0055	28	cn-0054
27	cn-0376	27	cn-0373	27	cn-0312	27	cn-0309	27	cn-0248	27	cn-0245	27	cn-0184	27	cn-0181	27	cn-0056	27	cn-0053
26	cn-0371	26	cn-0370	26	cn-0307	26	cn-0306	26	cn-0243	26	cn-0242	26	cn-0179	26	cn-0178	26	cn-0051	26	cn-0050
25	cn-0372	25	cn-0369	25	cn-0308	25	cn-0305	25	cn-0244	25	cn-0241	25	cn-0180	25	cn-0177	25	cn-0052	25	cn-0049
24	cn-0367	24	cn-0366	24	cn-0303	24	cn-0302	24	cn-0239	24	cn-0238	24	cn-0175	24	cn-0174	24	cn-0047	24	cn-0046
23	cn-0368	23	cn-0365	23	cn-0304	23	cn-0301	23	cn-0240	23	cn-0237	23	cn-0176	23	cn-0173	23	cn-0048	23	cn-0045
22	cn-0363	22	cn-0362	22	cn-0299	22	cn-0298	22	cn-0235	22	cn-0234	22	cn-0171	22	cn-0170	22	cn-0043	22	cn-0042
21	cn-0364	21	cn-0361	21	cn-0300	21	cn-0297	21	cn-0236	21	cn-0233	21	cn-0172	21	cn-0169	21	cn-0044	21	cn-0041
20	cn-0359	20	cn-0358	20	cn-0295	20	cn-0294	20	cn-0231	20	cn-0230	20	cn-0167	20	cn-0166	20	cn-0039	20	cn-0038
19	cn-0360	19	cn-0357	19	cn-0296	19	cn-0293	19	cn-0232	19	cn-0229	19	cn-0168	19	cn-0165	19	cn-0040	19	cn-0037
18	cn-0355	18	cn-0354	18	cn-0291	18	cn-0290	18	cn-0228	18	cn-0225	18	cn-0169	18	cn-0166	18	cn-0035	18	cn-0034
17	cn-0356	17	cn-0353	17	cn-0292	17	cn-0289	17	cn-0227	17	cn-0224	17	cn-0170	17	cn-0167	17	cn-0036	17	cn-0033
16	cn-0351	16	cn-0350	16	cn-0287	16	cn-0286	16	cn-0226	16	cn-0223	16	cn-0171	16	cn-0168	16	cn-0031	16	cn-0030
15	cn-0352	15	cn-0349	15	cn-0288	15	cn-0285	15	cn-0224	15	cn-0221	15	cn-0160	15	cn-0157	15	cn-0032	15	cn-0029
14	cn-0347	14	cn-0346	14	cn-0283	14	cn-0282	14	cn-0219	14	cn-0218	14	cn-0155	14	cn-0154	14	cn-0027	14	cn-0026
13	cn-0348	13	cn-0345	13	cn-0284	13	cn-0281	13	cn-0220	13	cn-0217	13	cn-0156	13	cn-0153	13	cn-0028	13	cn-0025
12	cn-0343	12	cn-0342	12	cn-0279	12	cn-0278	12	cn-0215	12	cn-0214	12	cn-0151	12	cn-0150	12	cn-0023	12	cn-0022
11	cn-0344	11	cn-0341	11	cn-0280	11	cn-0277	11	cn-0216	11	cn-0213	11	cn-0152	11	cn-0149	11	cn-0024	11	cn-0021
10	cn-0339	10	cn-0338	10	cn-0275	10	cn-0274	10	cn-0211	10	cn-0210	10	cn-0147	10	cn-0146	10	cn-0019	10	cn-0018
9	cn-0340	9	cn-0337	9	cn-0276	9	cn-0273	9	cn-0212	9	cn-0209	9	cn-0148	9	cn-0145	9	cn-0020	9	cn-0017
8	cn-0335	8	cn-0334	8	cn-0271	8	cn-0270	8	cn-0207	8	cn-0206	8	cn-0143	8	cn-0142	8	cn-0015	8	cn-0014
7	cn-0336	7	cn-0333	7	cn-0272	7	cn-0269	7	cn-0208	7	cn-0205	7	cn-0144	7	cn-0141	7	cn-0016	7	cn-0013
6	cn-0331	6	cn-0330	6	cn-0267	6	cn-0266	6	cn-0203	6	cn-0202	6	cn-0139	6	cn-0138	6	cn-0011	6	cn-0010
5	cn-0332	5	cn-0329	5	cn-0268	5	cn-0265	5	cn-0204	5	cn-0201	5	cn-0140	5	cn-0137	5	cn-0012	5	cn-0009
4	cn-0327	4	cn-0326	4	cn-0263	4	cn-0262	4	cn-0199	4	cn-0198	4	cn-0135	4	cn-0134	4	cn-0007	4	cn-0006
3	cn-0328	3	cn-0325	3	cn-0264	3	cn-0261	3	cn-0200	3	cn-0197	3	cn-0136	3	cn-0133	3	cn-0008	3	cn-0005
2	cn-0323	2	cn-0322	2	cn-0259	2	cn-0258	2	cn-0195	2	cn-0194	2	cn-0131	2	cn-0130	2	cn-0003	2	cn-0002
1	cn-0324	1	cn-0321	1	cn-0260	1	cn-0257	1	cn-0196	1	cn-0193	1	cn-0132	1	cn-0129	1	cn-0004	1	cn-0001
	880463-B21		880463-B21		880463-B21		880463-B21		880463-B21		880463-B21		880463-B21		P9Q65A		P9R85A		P9S25A
	Power Bottom		Power Bottom		Power Bottom		Power Bottom		Power Bottom		Power Bottom		Power Bottom		Power Bottom		Power Bottom		Power Bottom

Arm CPUs

Intel

NVIDIA

AMD

NVIDIA



<https://www.techrepublic.com/article/uk-ai-supercomputer/>

Isambard-AI — National AI Research Resource

- **5,448 NVIDIA Hopper GPUs**
- **392,256 ARM-based CPU cores** $7 \times$ *Isambard 3*
- HPE system design
- ~25 PiByte all-flash storage

- >21 ExaFLOP/s of 8-bit floating point for AI
- >250 PetaFLOP/s 64-bit for hybrid workflows
- In top 5-10 systems worldwide
- PUE under 1.1




Timeline:

- Phase 1 online in May for early users, **open from June**
- Phase 2 online from October for early users, **open from Nov**

<https://www.techrepublic.com/article/uk-ai-supercomputer/>



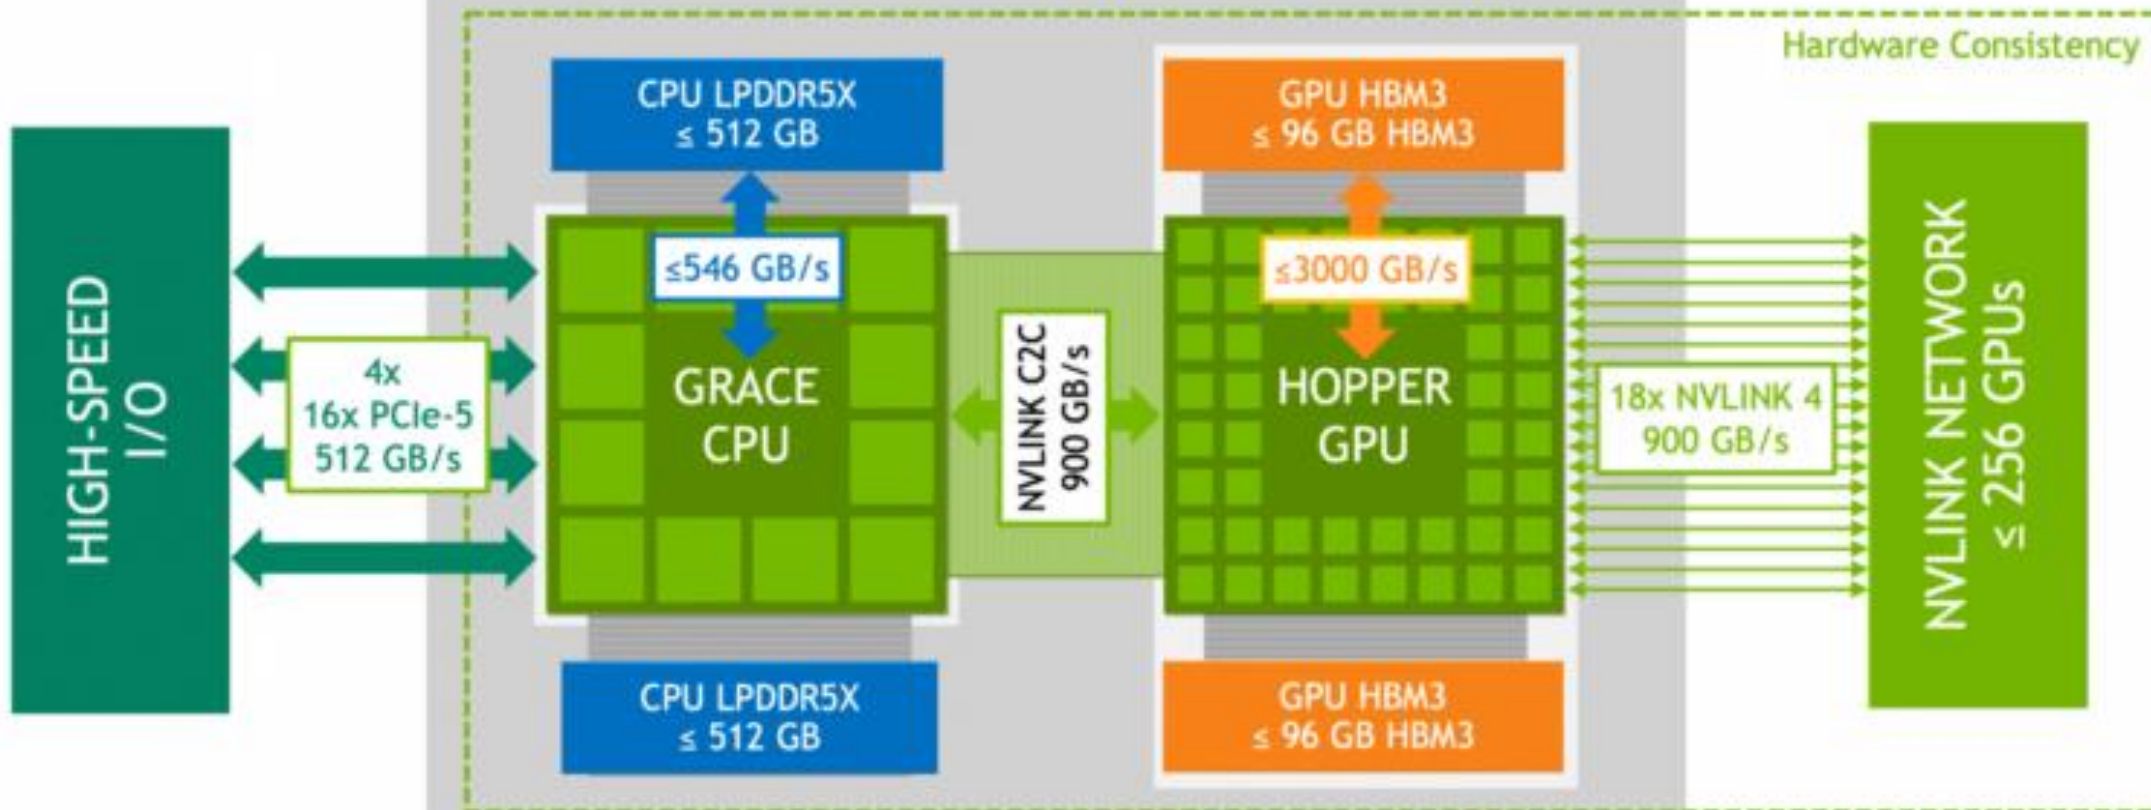




Grace-Hopper Superchip

<https://resources.nvidia.com/en-us-grace-cpu/grace-hopper-superchip>

NVIDIA Grace Hopper Superchip



72 Neoverse V2 Armv9 CPU cores

128 GB Memory

1 Hopper GPU

96 GB High Bandwidth Memory

224 GB

Memory Total

Grace-Hopper Superchip

<https://resources.nvidia.com/en-us-grace-cpu/grace-hopper-superchip>

Isambard AI Phase 1

Three shelves of 7 blades (21 total)



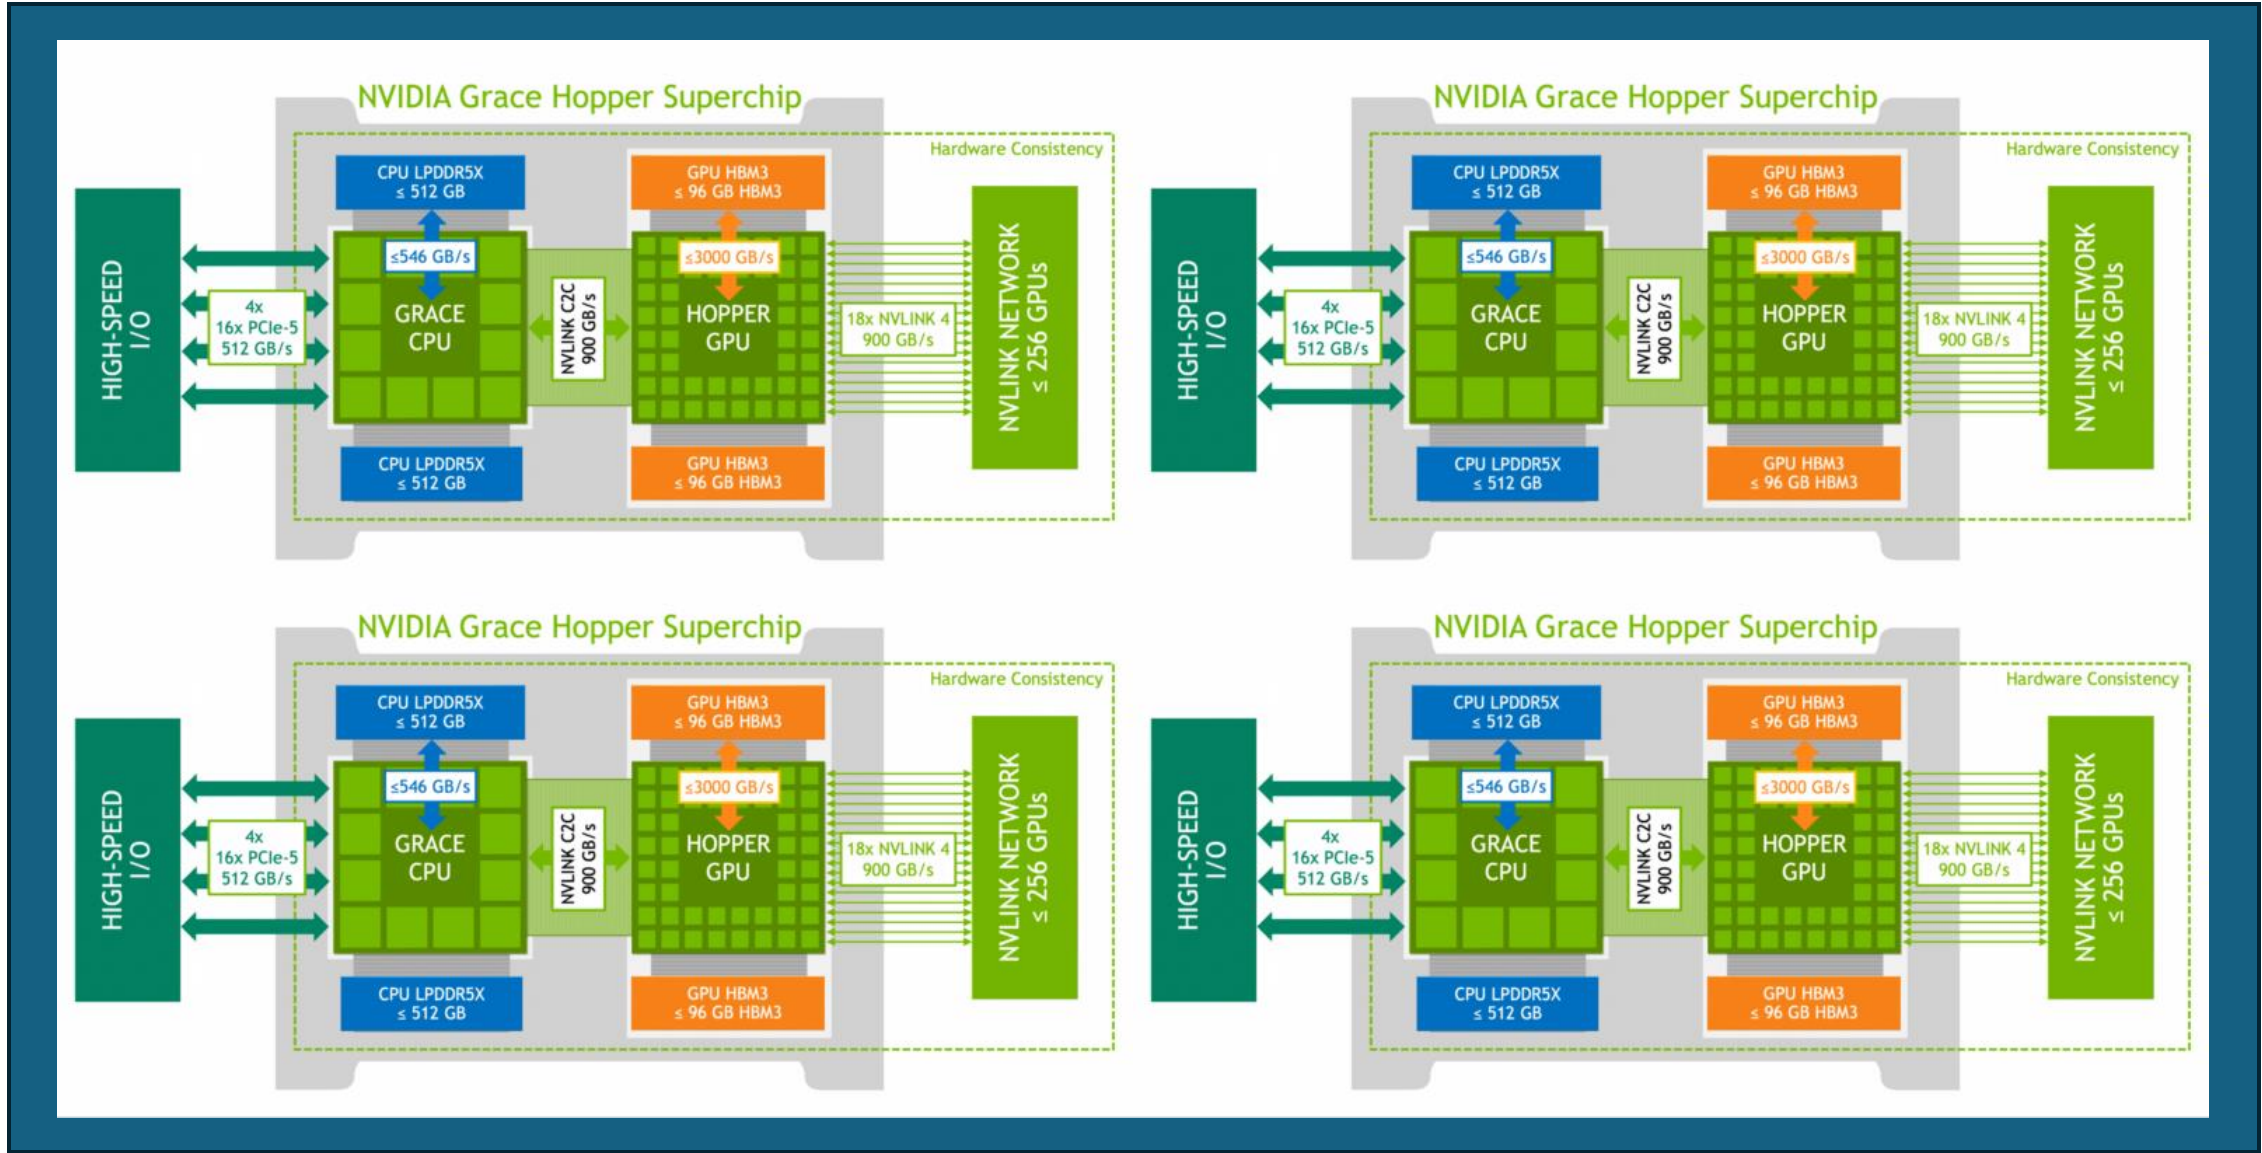
each with

two “super-nodes”
(42 total)

each with

four Grace-Hopper
“superchips”
(168 total)



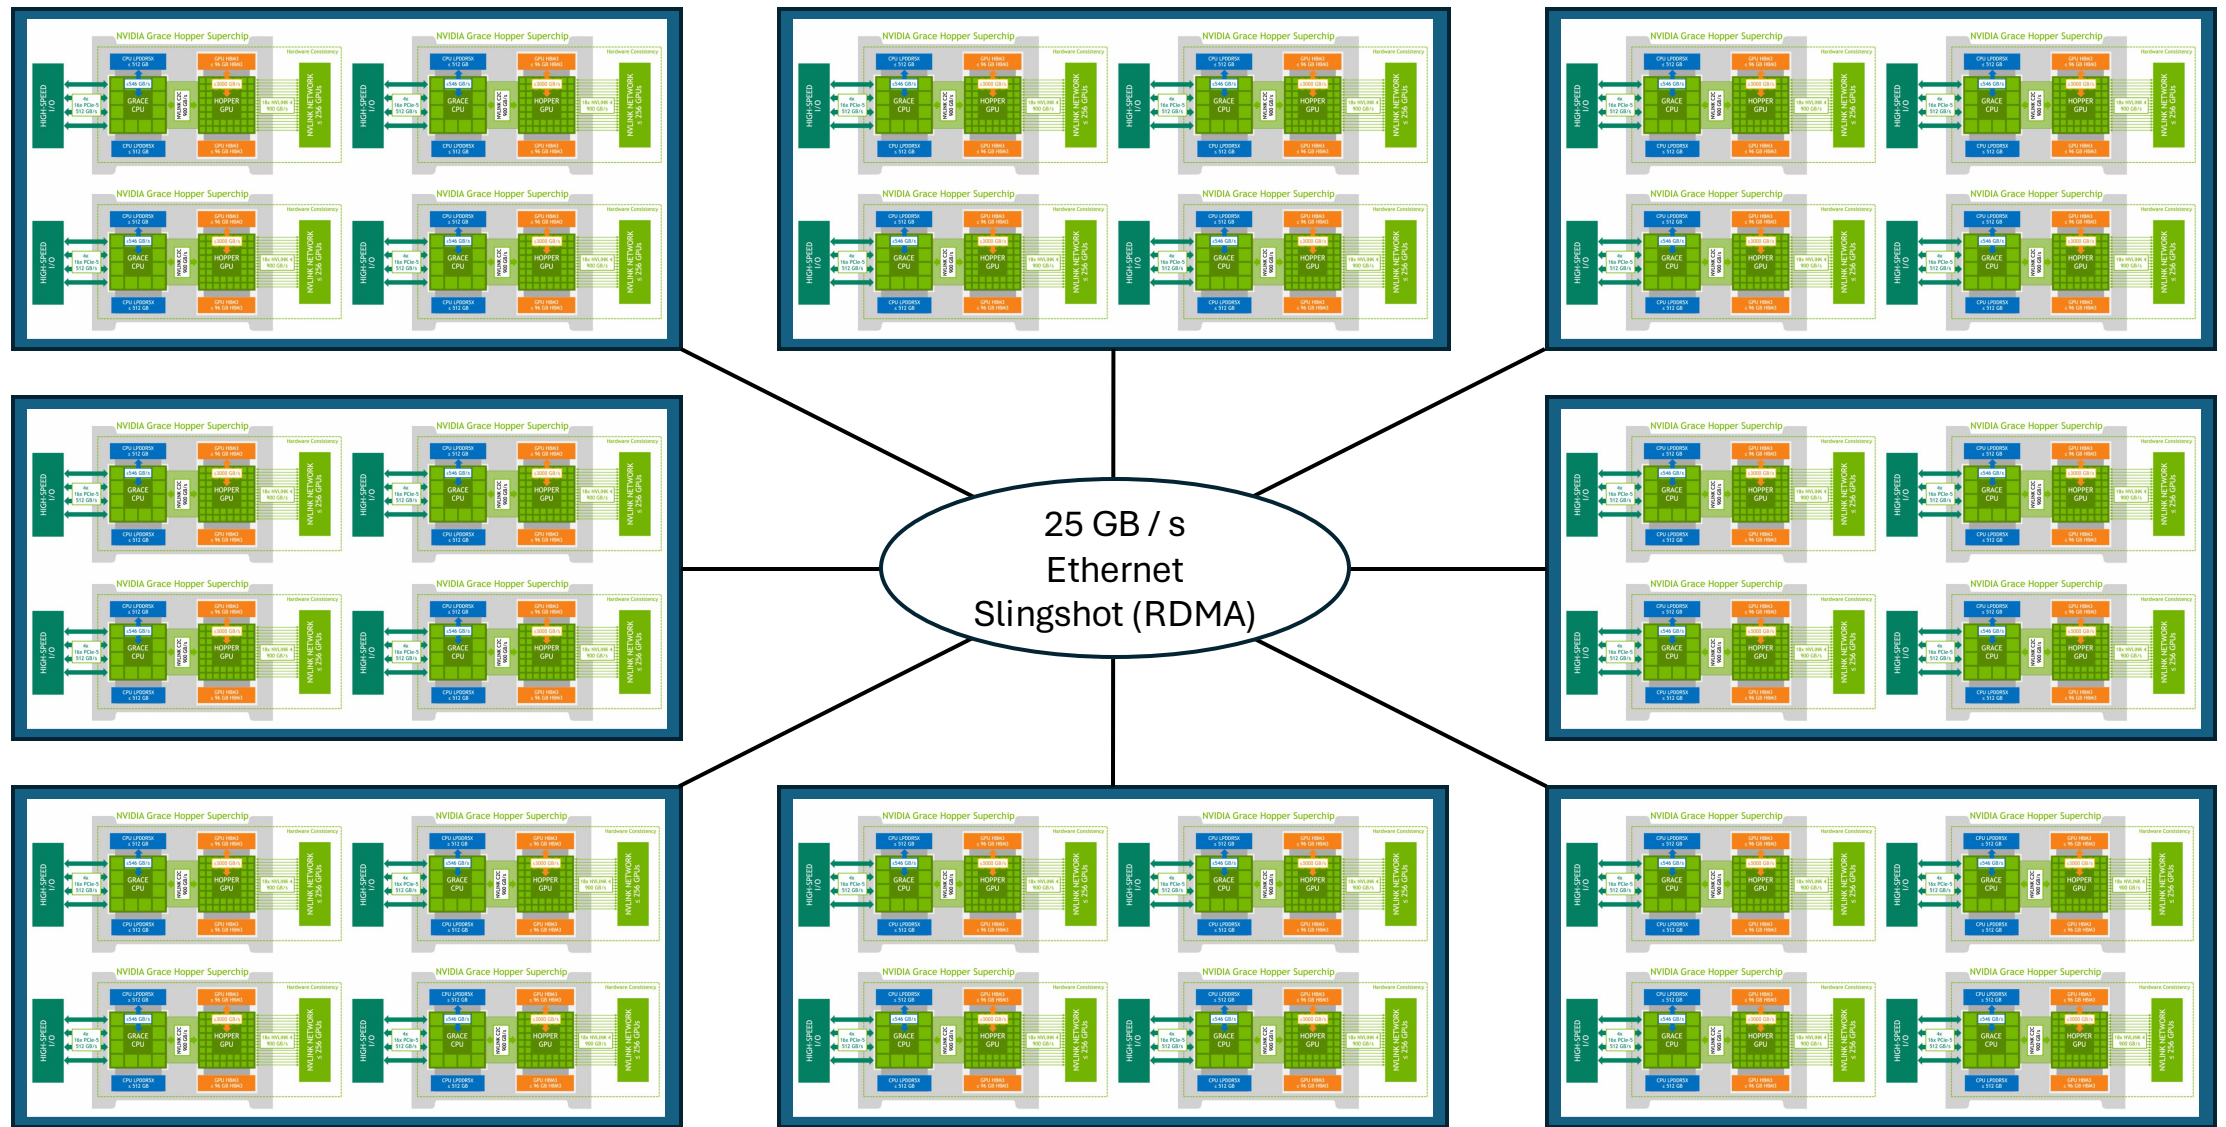


1 "supernode" =

288 CPU cores
512 GB Memory

4 Hopper GPUs
384 GB High Bandwidth Memory

896 GB
Memory Total



Each Supernode connected via 200 Gb/s HPE slingshot high speed network fabric
 Supports both ethernet and Slingshot RDMA packets simultaneously
 Optimised for ultra-low latency and minimizing switch hops
 GPU = 3000 GB/s, Supernode = 900 GB/s, Fabric = 25 GB/s

Isambard-AI Phase 1 *versus* Phase 2

- Arrived last week
- 21 blades
- 42 supernodes
- 168 GH superchips
 - **12,096** Neoverse V2 Armv9 CPU cores
 - **168** Hopper GPUs
- 21.5 TB CPU memory
- 16.1 TB high bandwidth GPU memory
- 37.6 TB total memory

x30!

- ...plus ~1 PB all-flash storage

- Arriving Summer – Isambard AI MDC
- 660 blades
- 1320 supernodes
- 5280 GH superchips
 - **380,160** Neoverse V2 Armv9 CPU cores
 - **5280** Hopper GPUs
- 675 TB CPU memory
- 506 TB high bandwidth GPU memory
- 1.18 PB total memory

- ...plus ~27 PB all-flash storage!
 - (~20 PB posix, ~7 PB object)



Traditional batch interface

Interactive interface



Isambard 3

- General HPC workloads
- **55k ARM CPU cores**
- **20 GPUs (mixed)**
- Summer 2024



Isambard AI

Phase 1

- AI/ML only
- **12k ARM CPU cores**
- **168 NVIDIA GPUs**
- Summer 2024
- NVIDIA GPUs

Phase 2

- AI/ML only
- **380k ARM CPU cores**
- **5280 NVIDIA GPUs**
- Autumn 2024
- NVIDIA GPUs