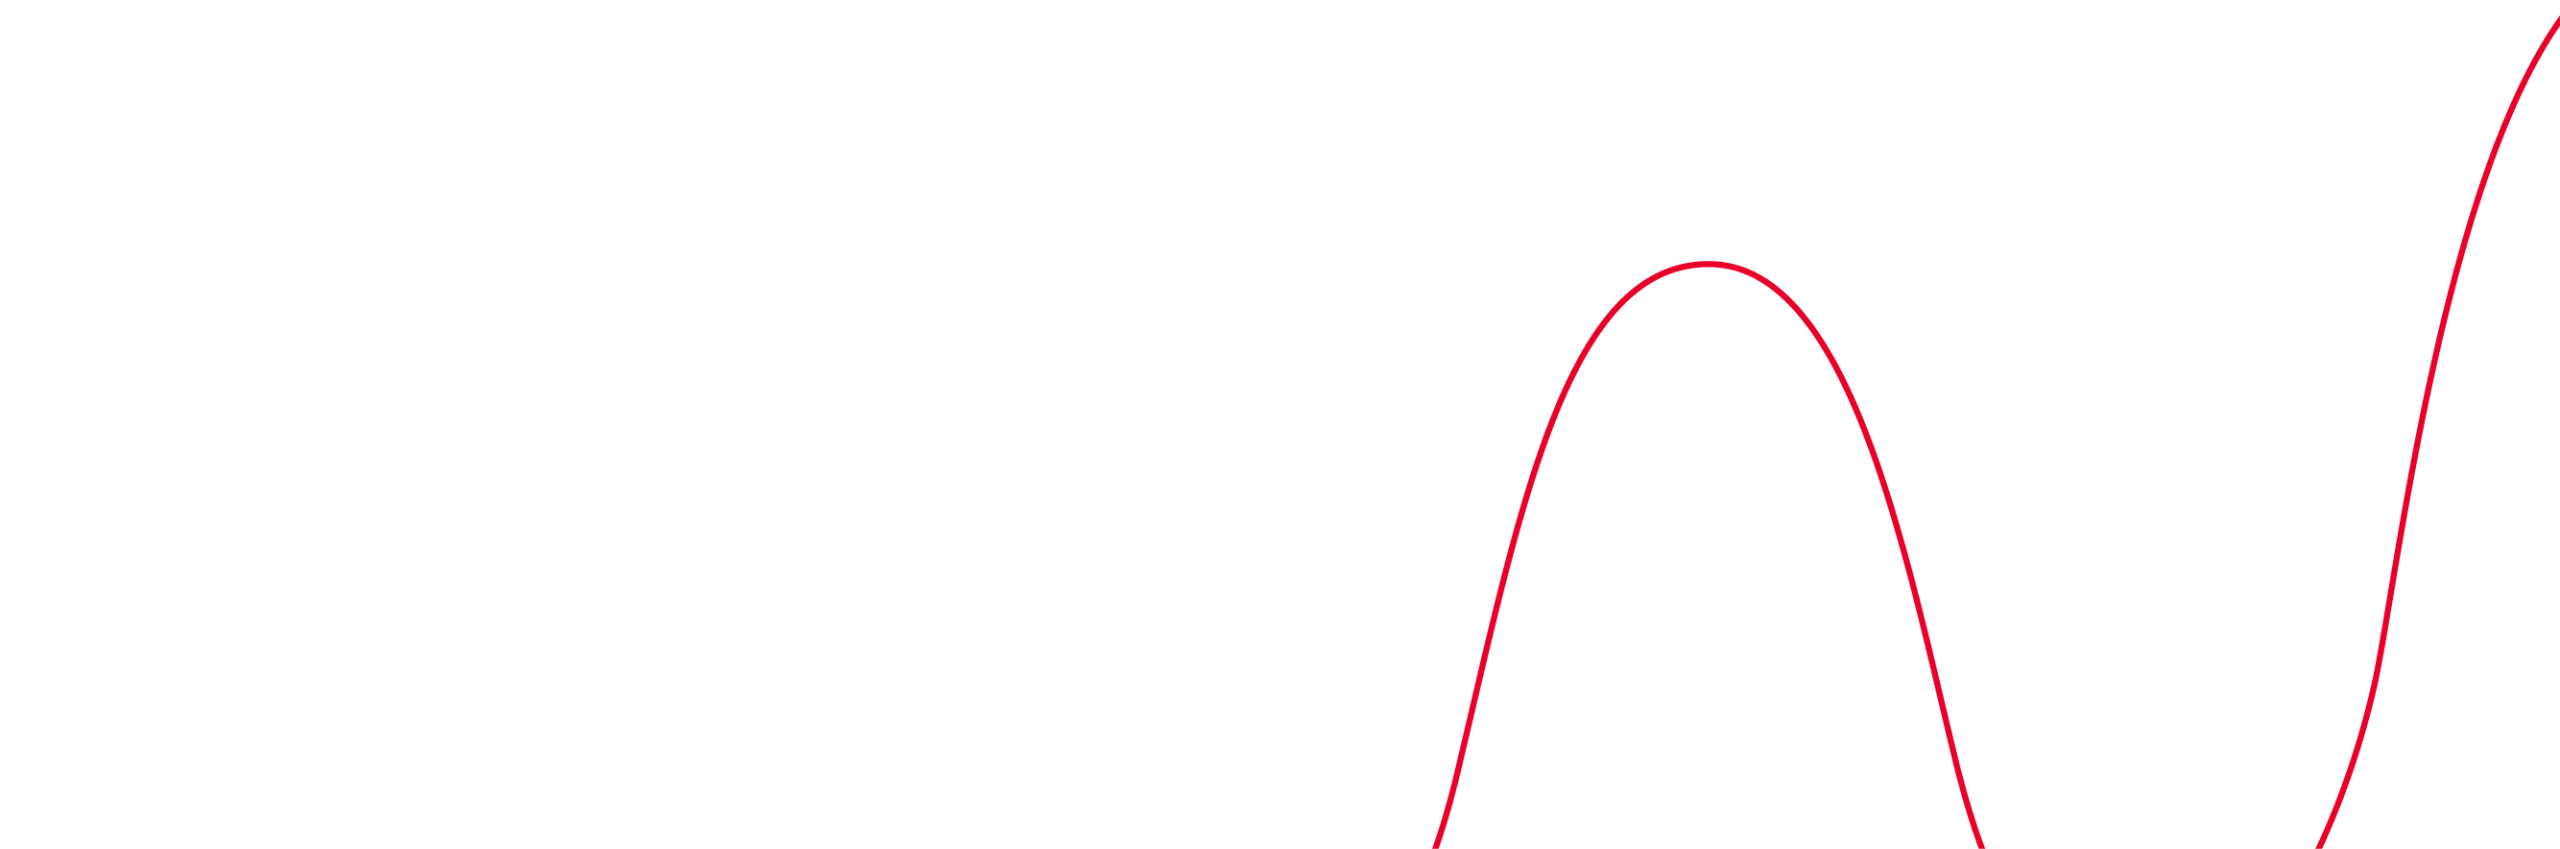
KEYSIGHT

Keysight AI Data Center Builder

October 2025

AI Infra Market Trends



AI: A Technical Inflection Point

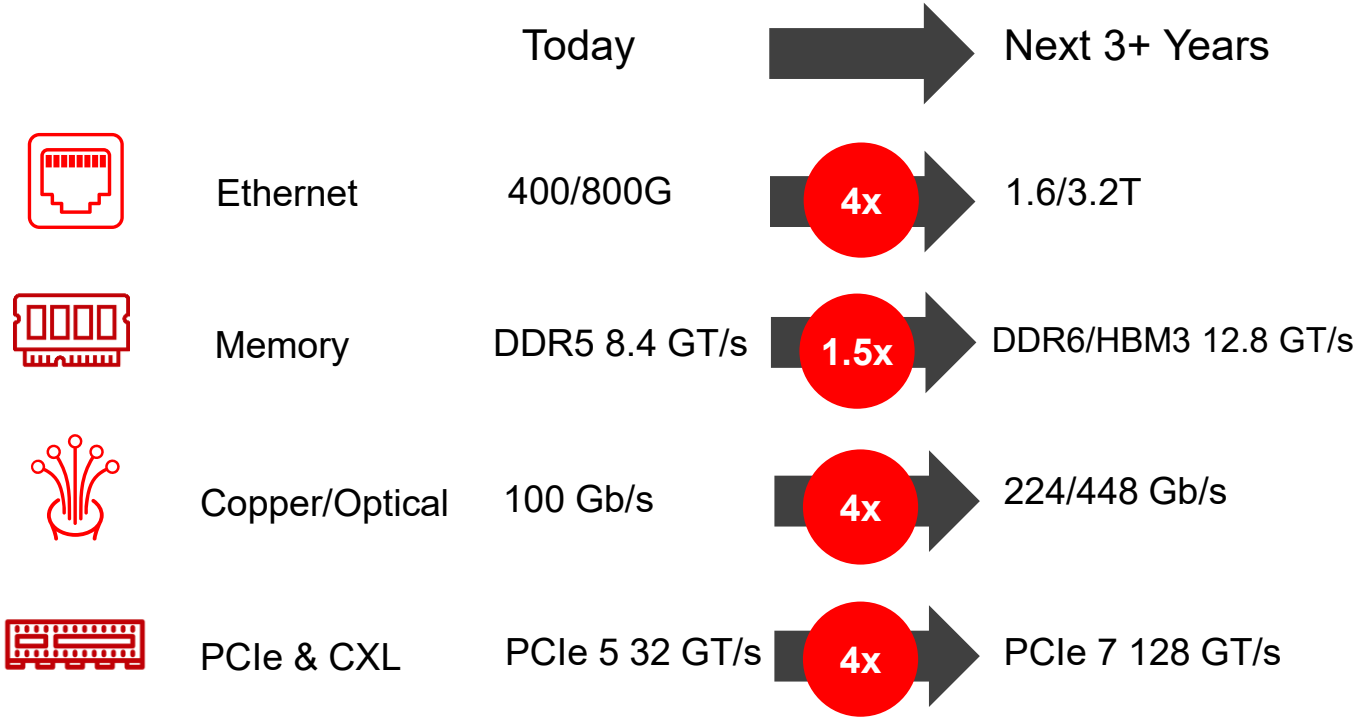
Bandwidth issues are often met by increasing bandwidth, but optimizing efficiency is equally crucial

AI market predicted to reach $1.3T by 2030 — up from 86.9 billion in 2022.

AI market will grow at a 35.7% annual rate from 2024 to 2030.



Market growth in AI will contribute a 21% net increase to US GDP by 2030.



Heterogeneous Designs Picking Up Steam

The next generation of AI data centers is not “one size fits all”

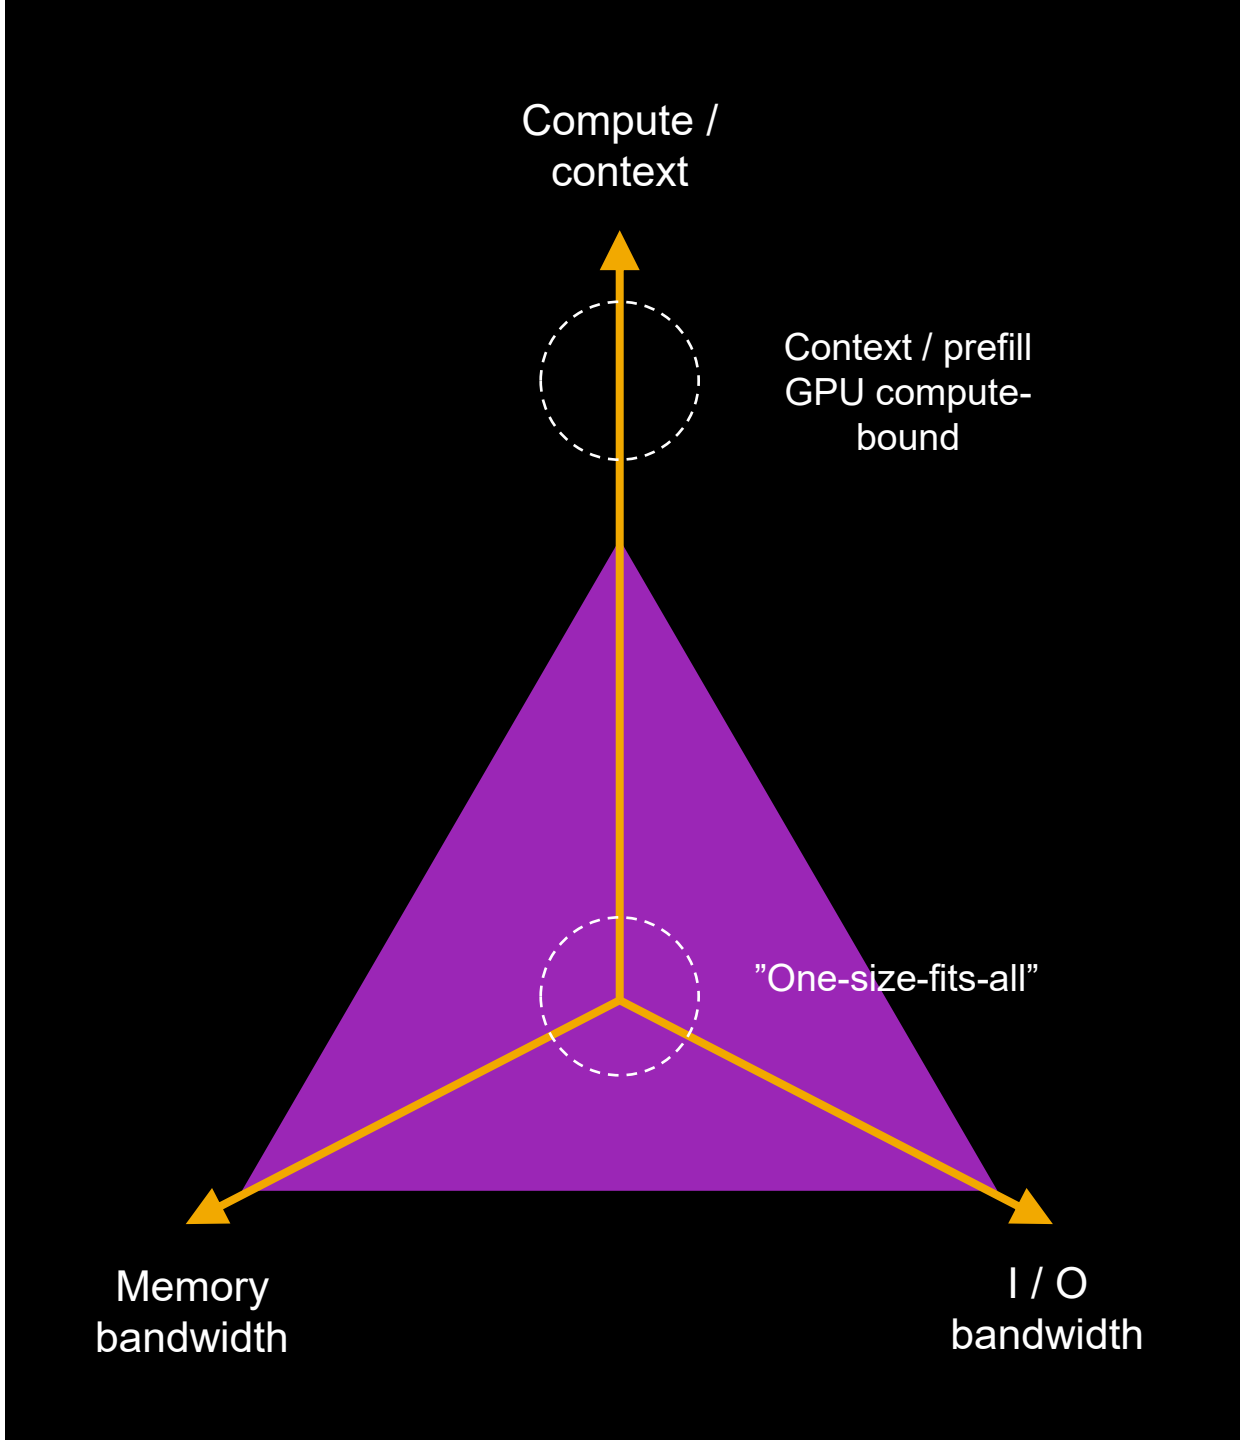
- AI has moved past **vertical integration** and **homogeneity**
- Variety is good, but it means we need to **standardize**
- New vendors make it easier to **customize systems** based on **individual requirements**



Domain-Specific GPUs Are the Future

GPUs optimized for prefill and decode emerging

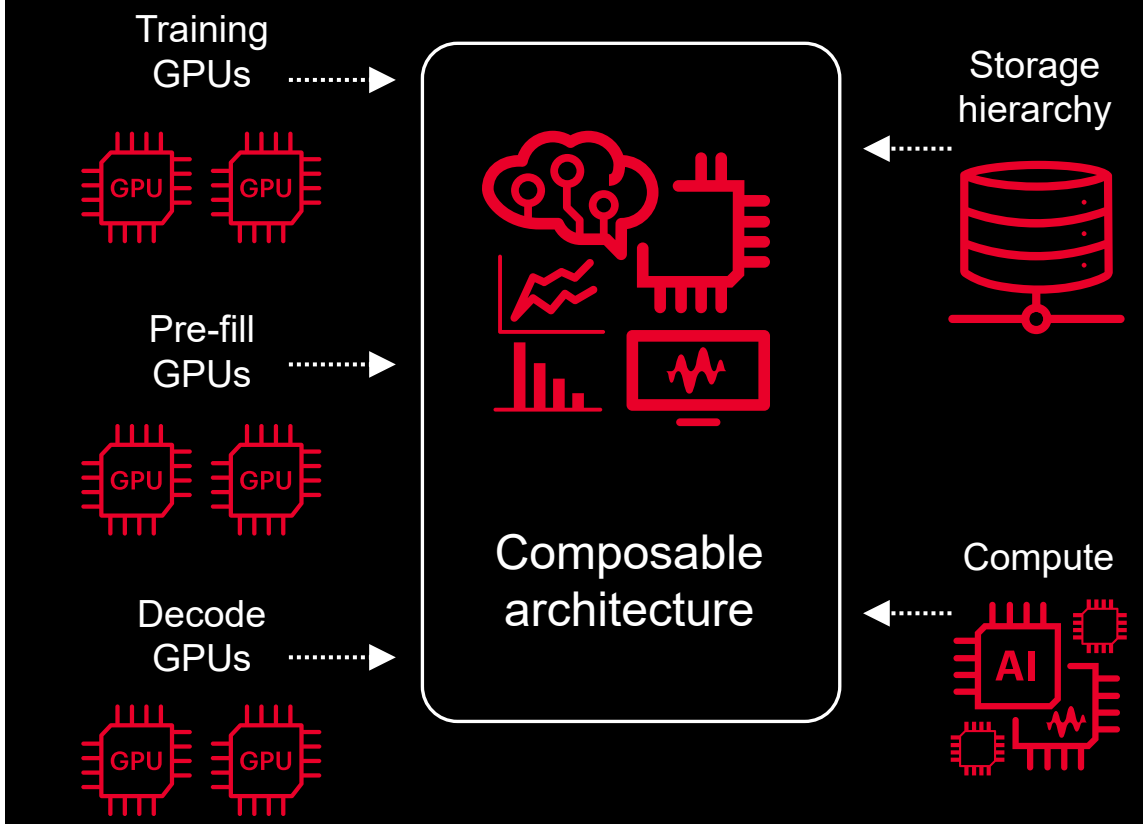
- **Training and inference were once the sole differentiator; now it's the tip of the iceberg**
- **GPUs optimized uniquely for each industry vertical**
- **As specialized infrastructure grows, how do we aggregate it all into a cohesive system?**



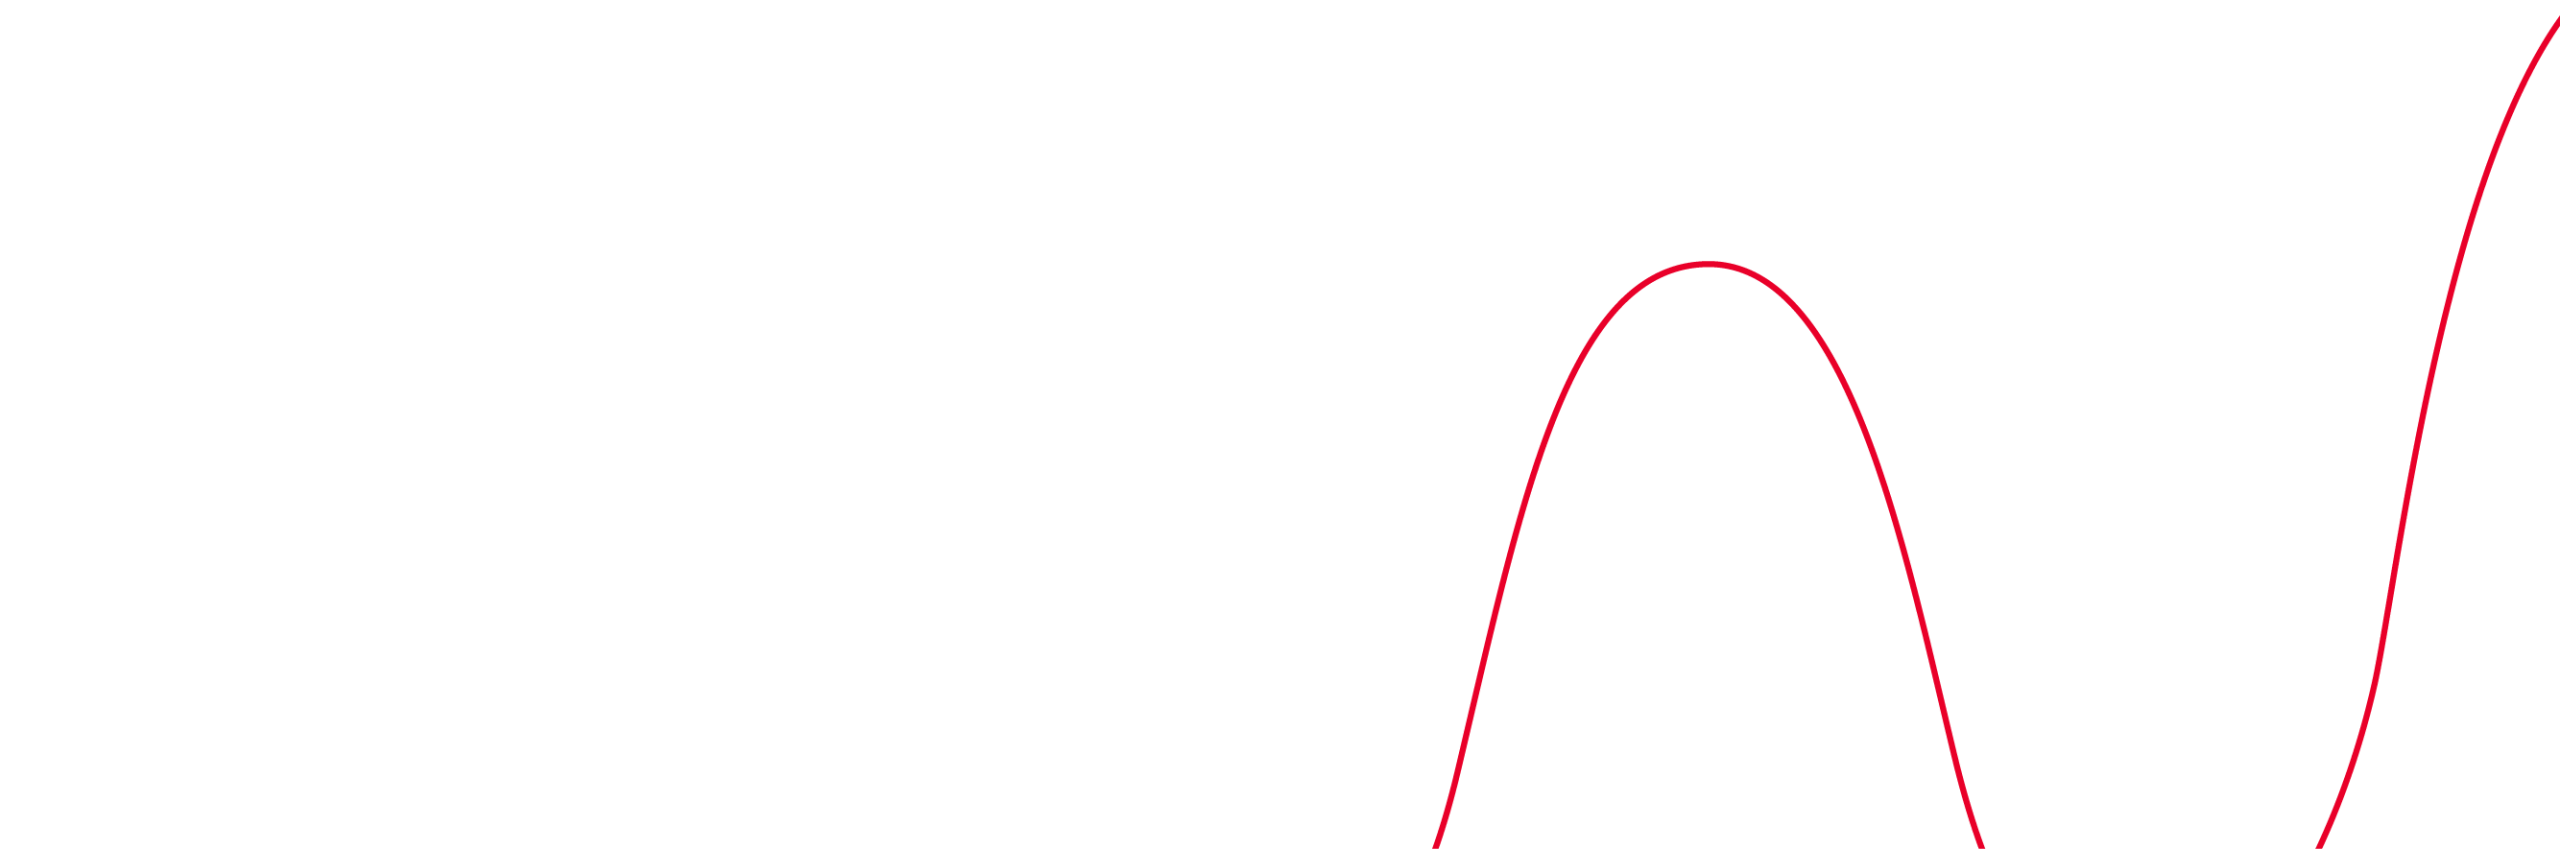
Flexibility for Applications

Ensures AI data centers evolve with requirements

- As AI moves deeper into the enterprise, it's **hard to predict** what **workloads** will look like
- Apps drive requirements, but they're **constantly changing**
- **Composable infrastructure** is a dynamic alternative, creating new systems as needed

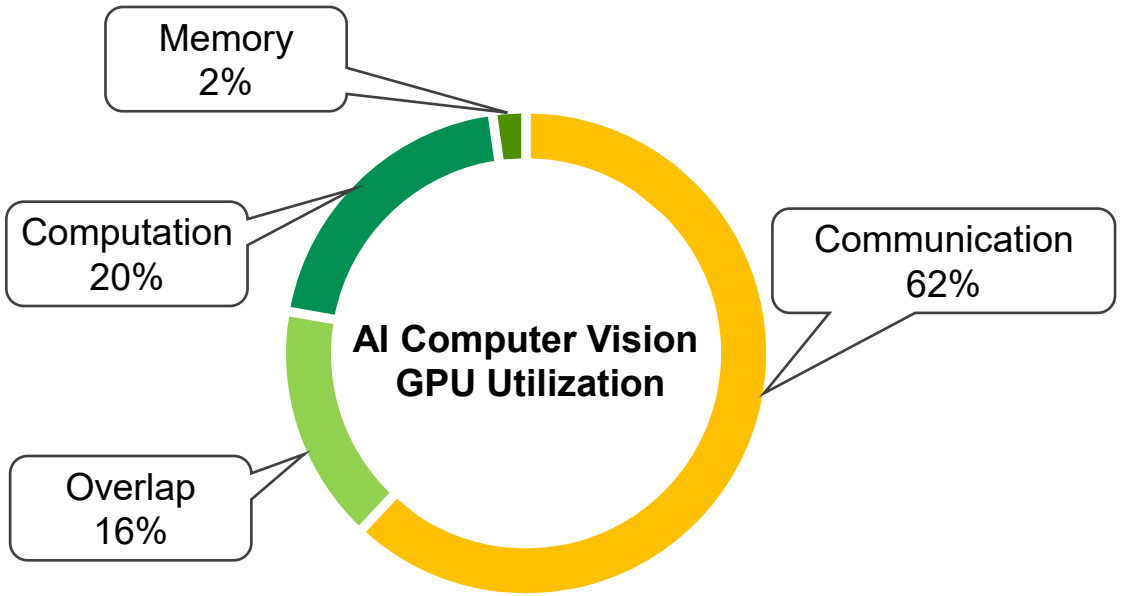


AI Cluster Networking in Focus



Why the Network & Components Matters in an AI Cluster

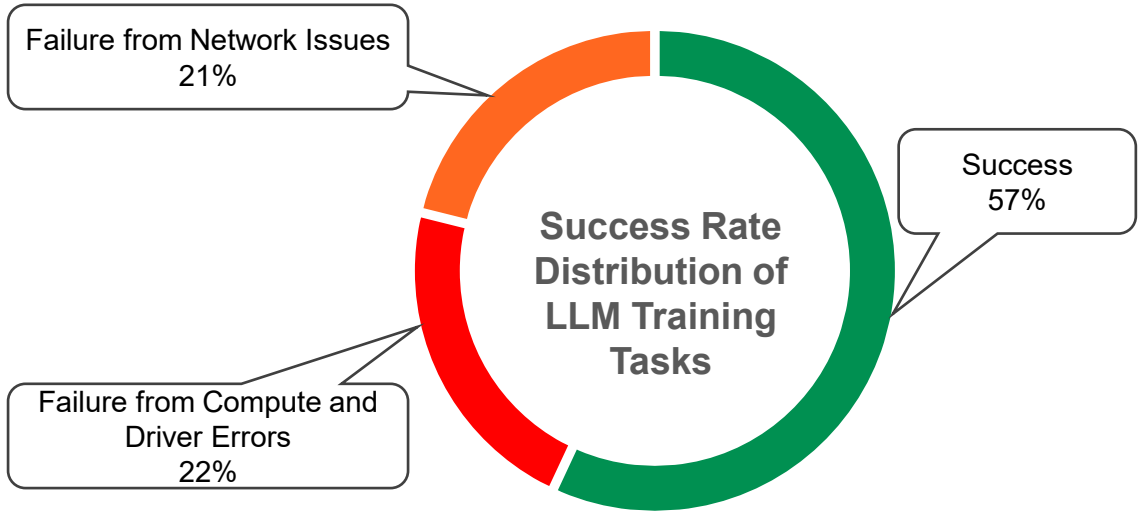
AI is Compute, Network & Data Intensive and requires validation at System Scale



GPUs waiting on data

>50%

Vision transformer (ViT) example. Source: https://github.com/facebookresearch/HolisticTraceAnalysis/



Training task failures

>43%

Source: Unicron: Economizing Self-Healing LLM Training at Scale, Tao He[1], Xue Li[1], Zhibin Wang[1,2], Kun Qian[1], Jingbo Xu[1], Wenyuan Yu[1], Jingren Zhou[1] [1]Alibaba Group, [2]Nanjing University

Validating Networks in AI clusters



Keysight
CyPerf

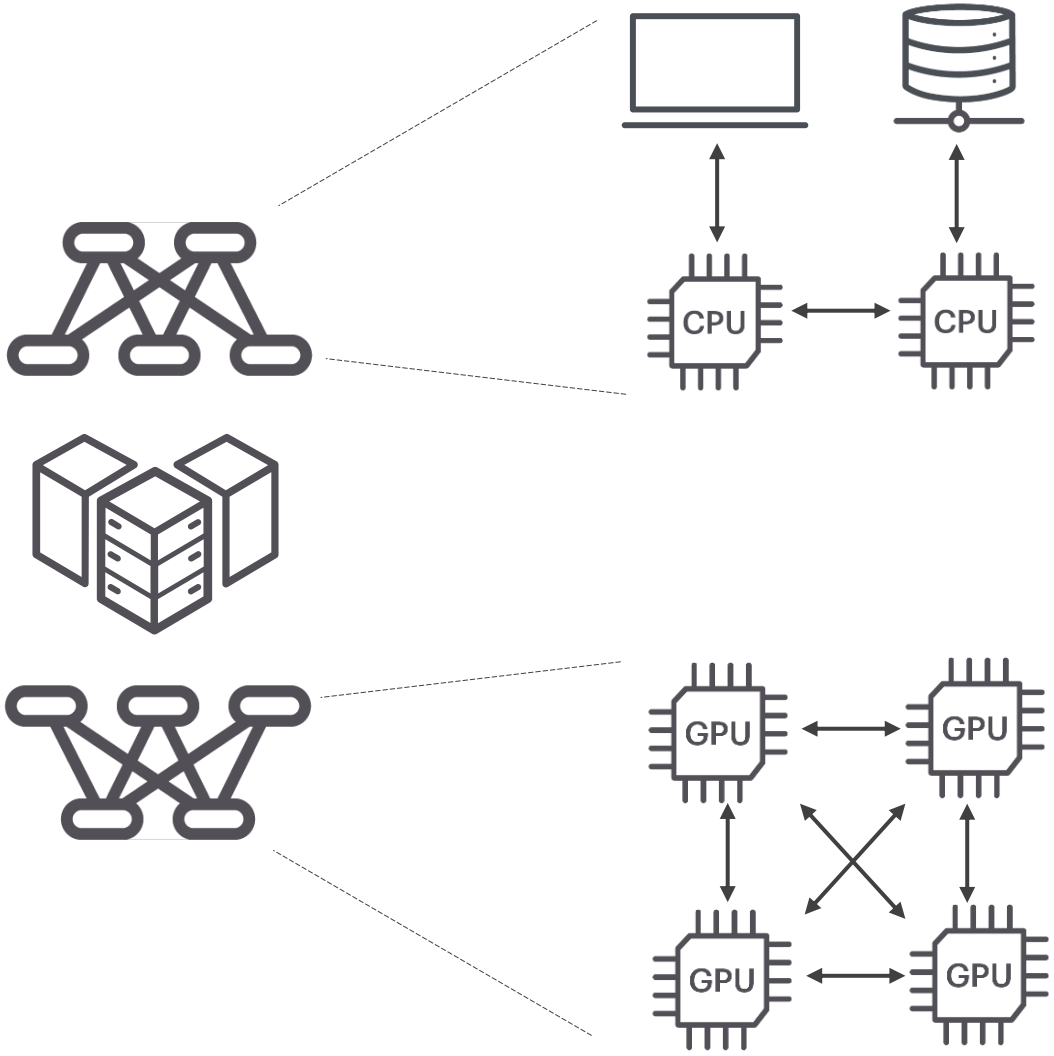
Frontend



KAI DC Builder

Backend

Focus of this presentation



KAI Data Center Builder

AI Infrastructure Design & Validation Solution

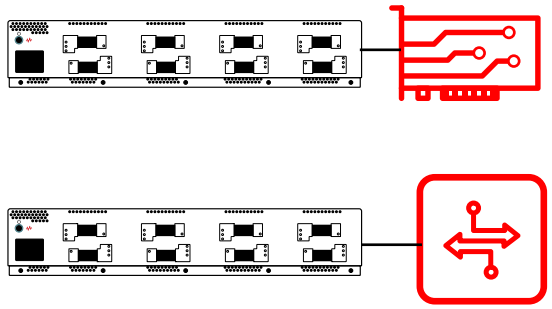
- Benchmark AI network infrastructure
- Co-tune AI cluster performance
- Reproduce issues seen in production
- Prove new designs at system scale



High fidelity emulation of GPU servers

KAI Data Center Builder Use Cases

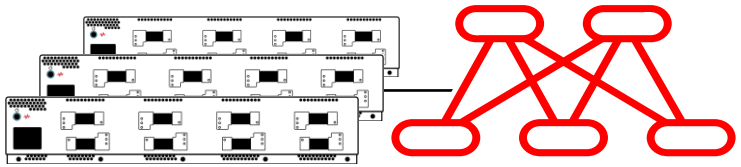
1 Device



RoCEv2 & UEC

- Compliance
- Throughput
- Connection rates
- Congestion handling
- Negative tests

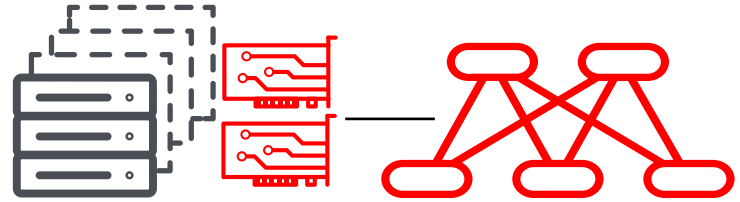
2 Network



Scale-up & -out

- Throughput
- Lossless queuing
- Congestion signaling
- Tail latencies

3 System

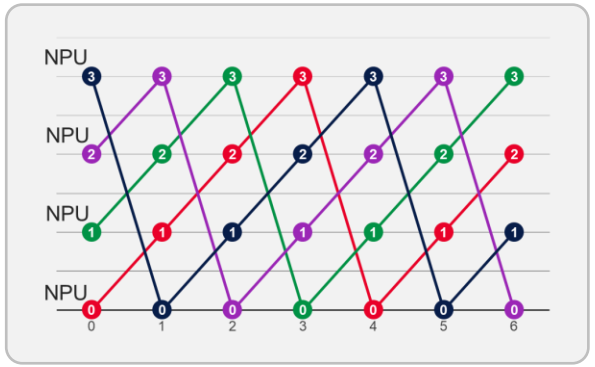


Workloads

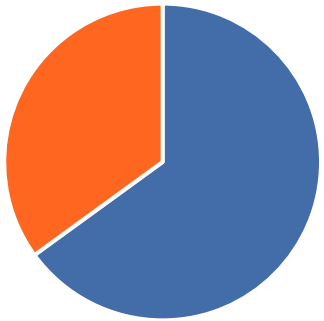
- JCT
- Benchmarks
- Performance Isolation
- Load balancing
- Failure recovery

KAI DC Builder Applications

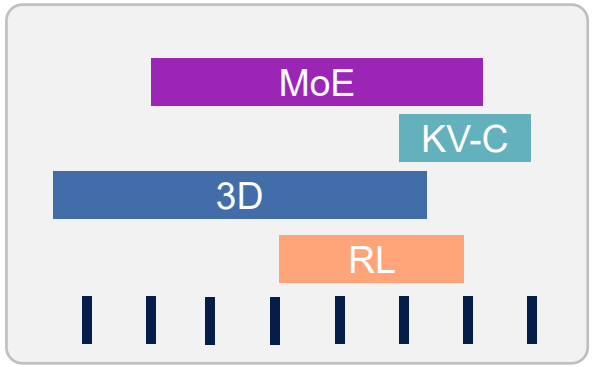
Collective Benchmarks



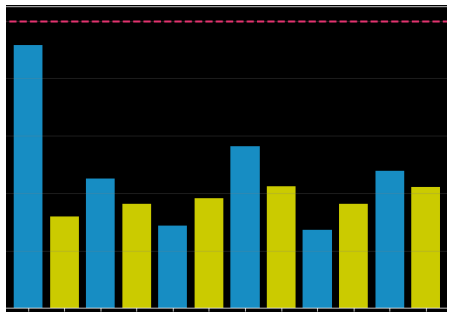
Network Utilization



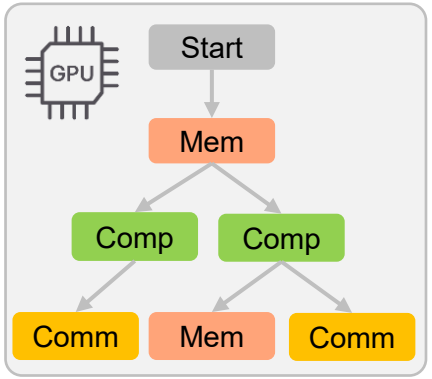
Mix of Collectives



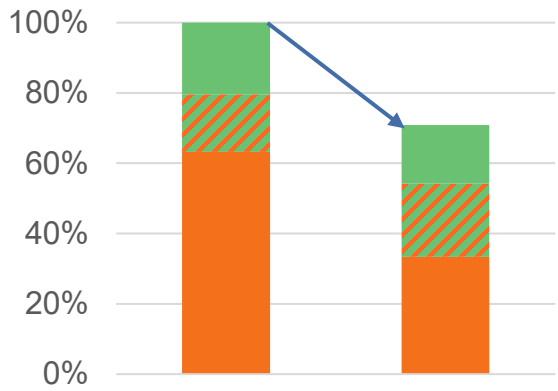
Fabric Validation



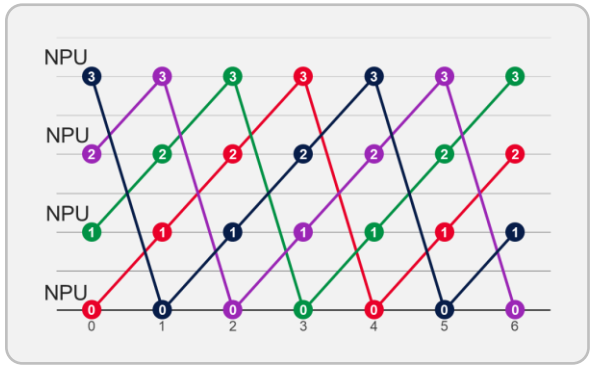
Workload Emulation



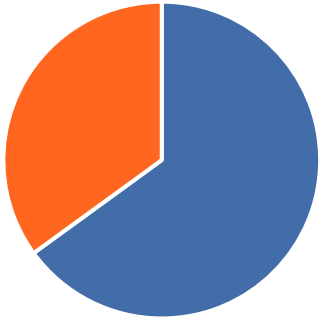
Training Time



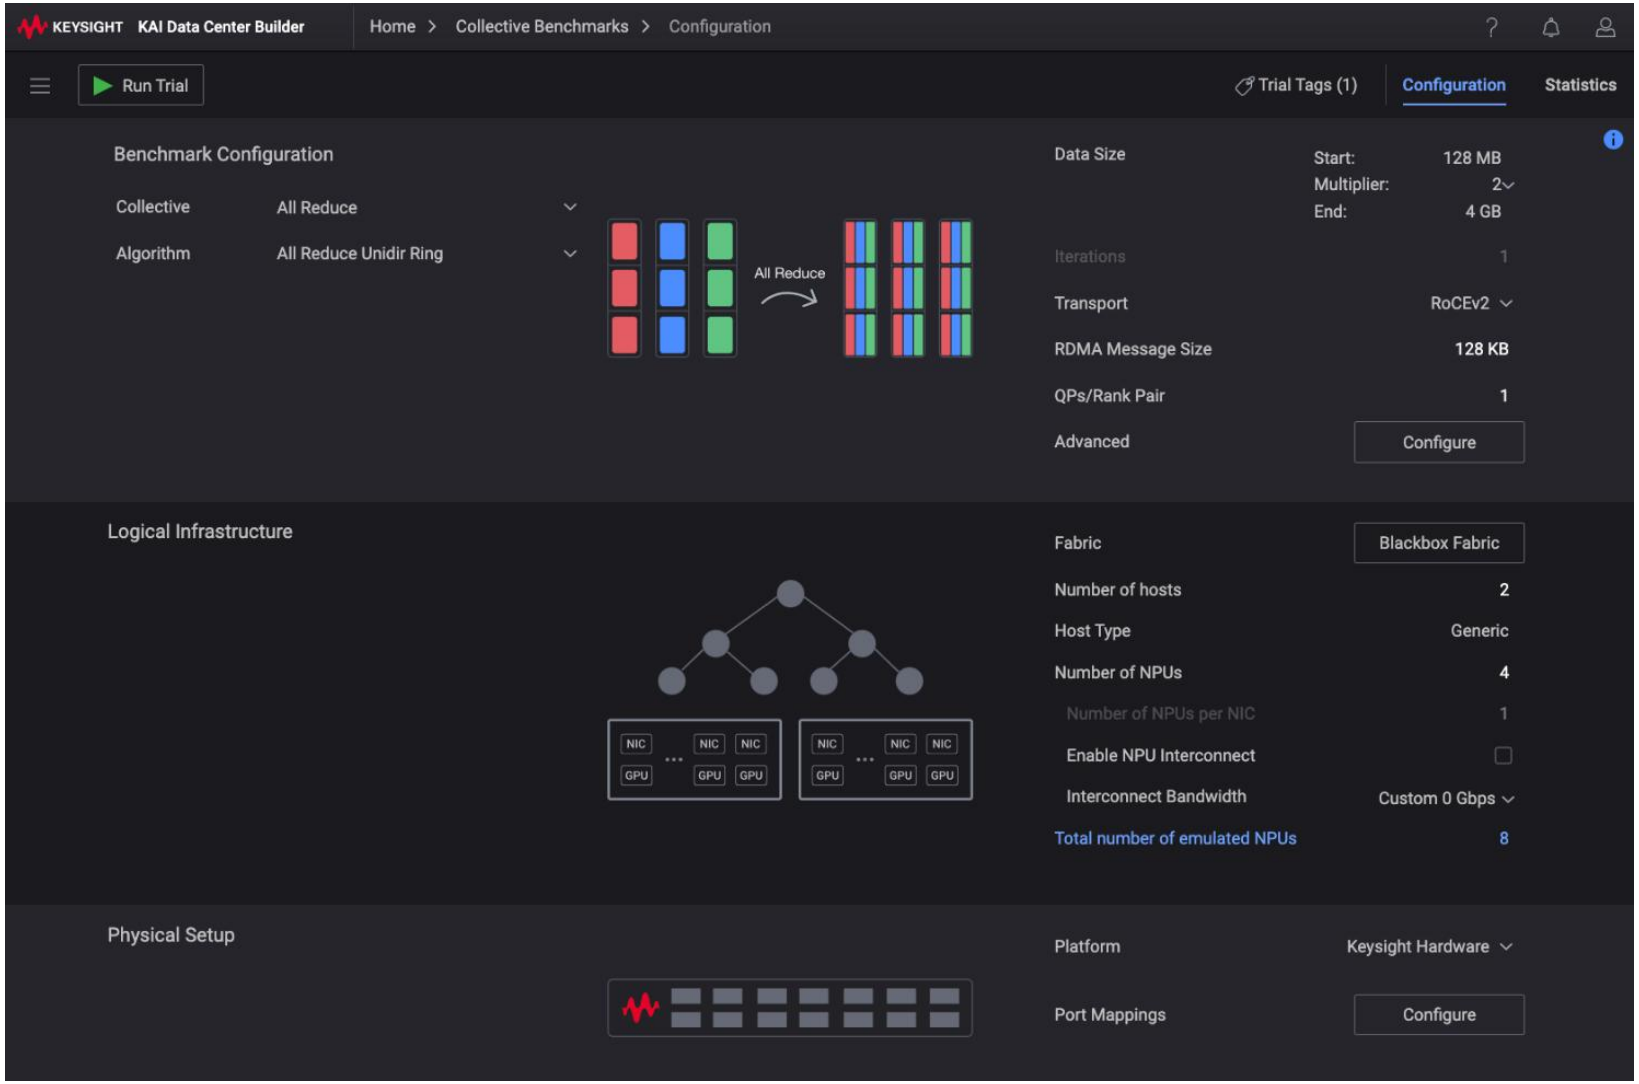
Collective Benchmarks



Network Utilization



■ Measured ■ Idle



Characterizing Efficiency

Example: Data Size sweep

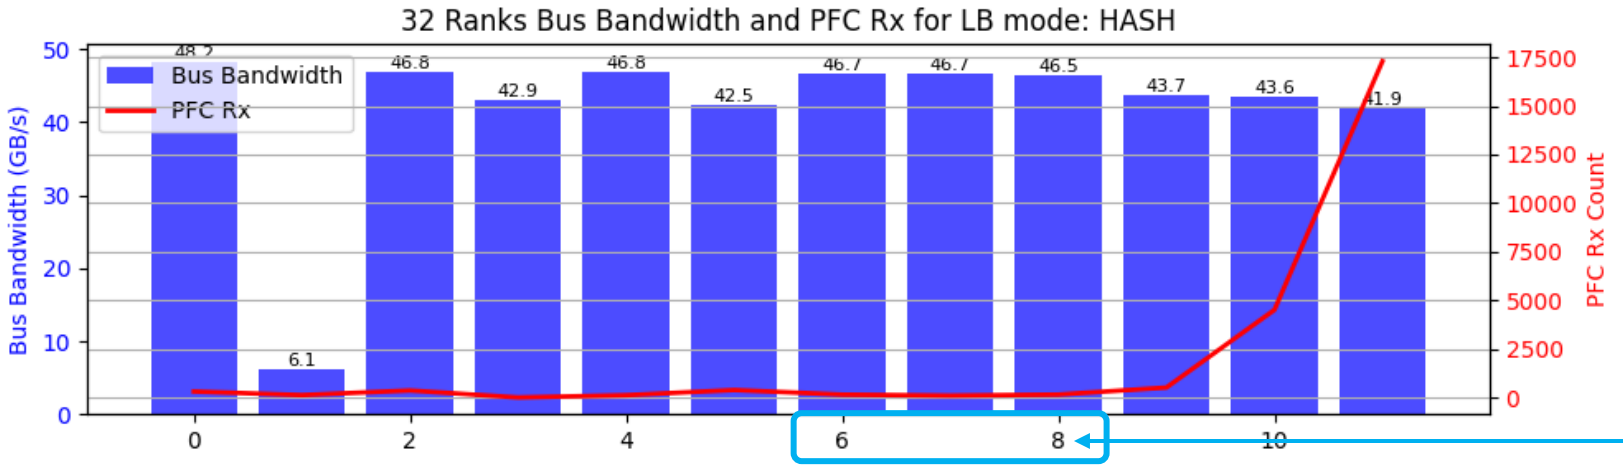
Trial Summary

Index	Data Size	Collective	Completion Time (ms)	Algbw (GB/s)	Busbw (GB/s)	Ideal (%)	PFC Rx	ECN CE Rx	Frames ReTx	Min FCT (ms)	P50 FCT (ms)	P95 FCT (ms)	Max FCT (ms)
1	128.00 MB	ALL_TO_ALL-1	7.936	16.91	14.80	30.05	0	1458	0	7.220	7.720	7.910	7.936
2	256.00 MB	ALL_TO_ALL-2	15.362	17.47	15.29	31.05	2	1982	0	14.858	15.150	15.300	15.362
3	512.00 MB	ALL_TO_ALL-3	30.103	17.83	15.61	31.69	0	1787	0	28.978	29.582	29.799	30.103
4	1,024.00 MB	ALL_TO_ALL-4	58.697	18.29	16.01	32.50	0	2154	0	55.448	57.516	58.650	58.697
5	2,048.00 MB	ALL_TO_ALL-5	113.709	18.89	16.53	33.55	2	2720	0	105.642	109.996	112.258	113.708
6	4,096.00 MB	ALL_TO_ALL-6	221.121	19.42	17.00	34.51	0	3371	0	189.554	208.750	218.249	221.121

Performance

Trouble signals

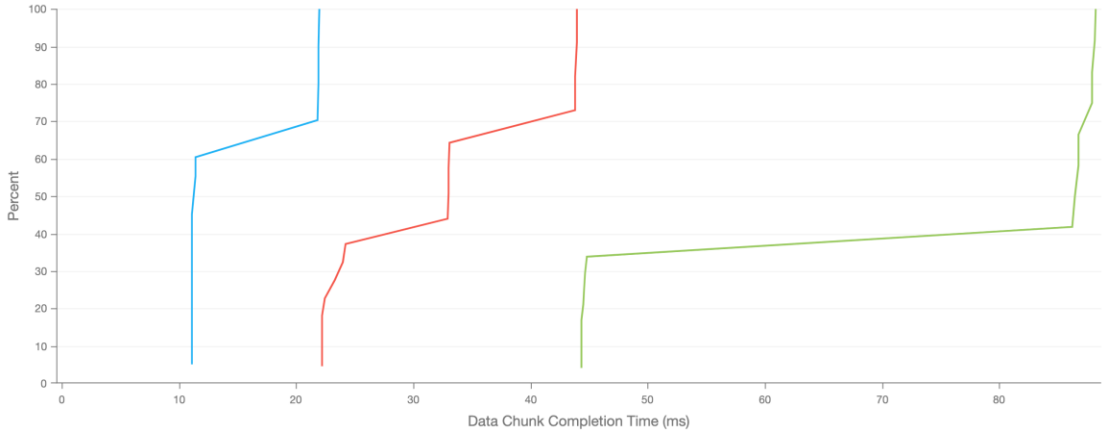
Stragglers



Example: DCQCN parameter sweep

Set of trials with optimal parameters

Visualizing Performance Observations

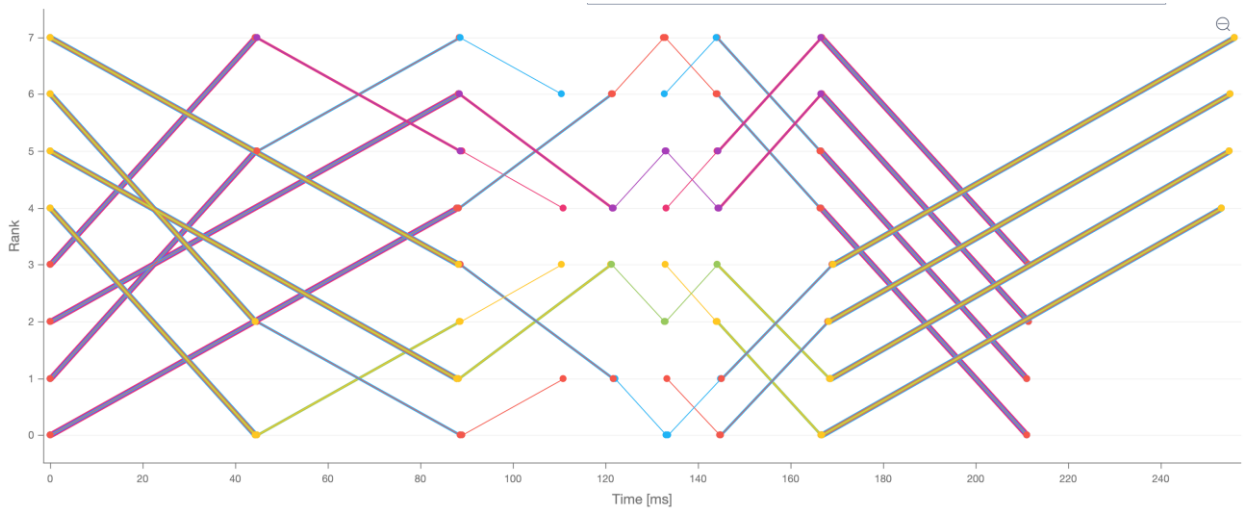


Cumulative Distribution Function (CDF)

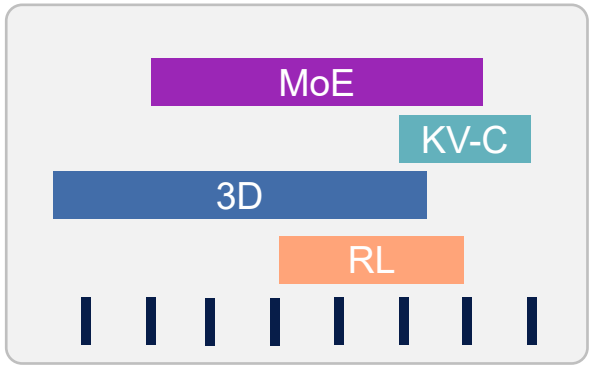
Severity and uniformity of issues

Data Movement Timeline

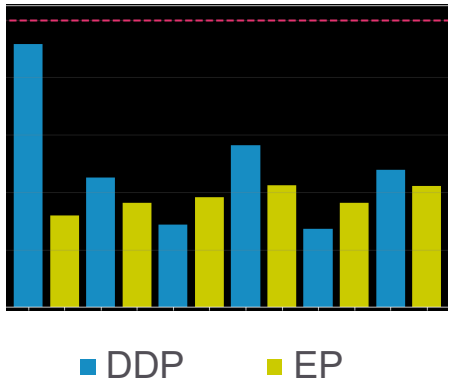
Localize place and time of issues



Mix of Collectives



Fabric Validation



Curated Mixes

- Training-Mix-2023 3D Parallel
- Training-Mix-2024 Mixture of Experts
- Training-Mix-2025 MoE + RL
- Inference-Mix-2025 MoE, KV Cache

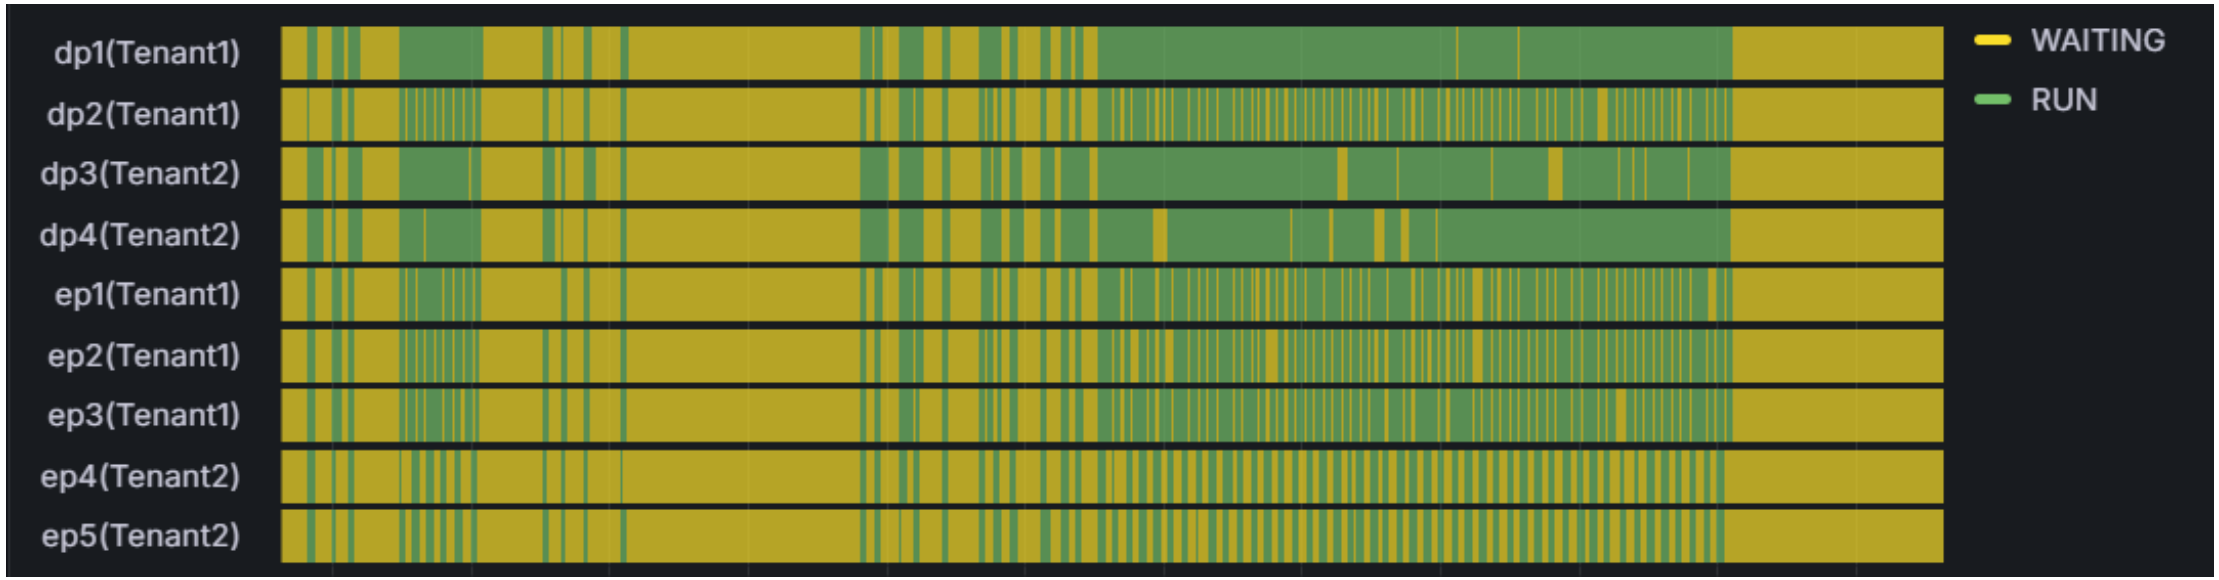
Use Cases

- Switch Regression
- Vendor Qualification
- Multi-Tenant Performance Isolation

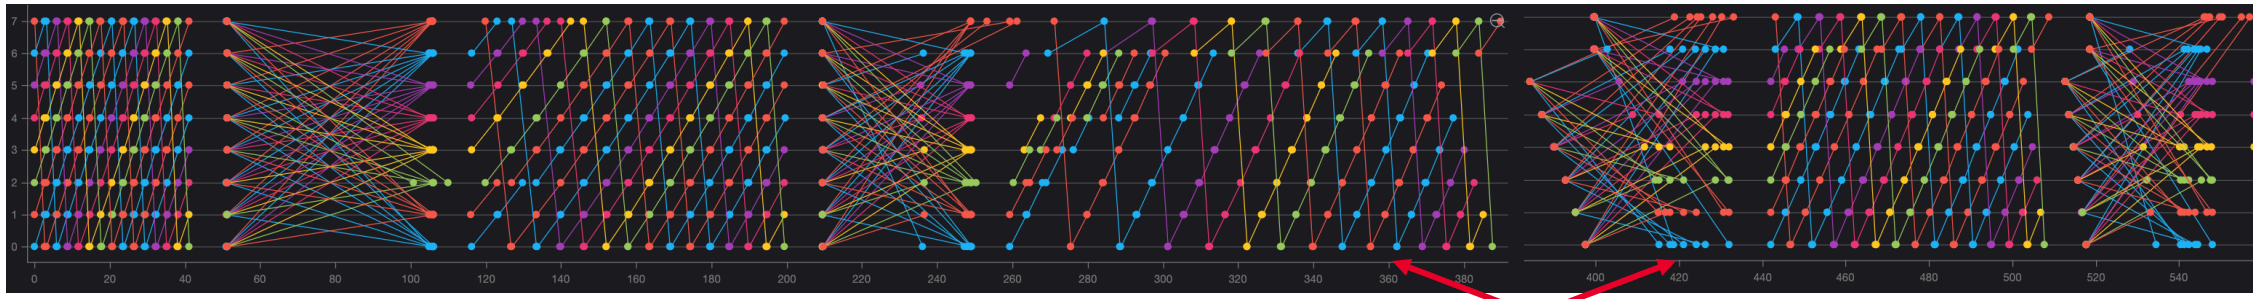
Training & Inference Mix 2025 Example

Tenant	Comm Group	Job Type	Data Size	Iterations	Ranks	Algorithm	
Tenant 1 64 ranks	A	DP1.1	1GB	>= 1000	16	AllReduce Ring	Job 1: 3D Parallel Training
	B	PP1.1	128MB	>= 100	8	Point-2-Point	
	C	DP1.2	1GB	>= 1000	16	AllReduce Ring	Job 2: MoE Model Training
	D	PP1.2	128MB	>= 100	8	Point-2-Point	
	E	EP1.1	512MB	>= 100	16	AlltoAll	
Tenant 2 24 ranks	F	DP2.1	1GB	>= 1000	16	AllReduce Ring	Job 3: 3D Parallel Training Job
	G	PP2.2	64MB	>= 100	8	Point-2-Point	
Tenant 3 16 ranks	H	KV3.1	4GB	>= 100	8	AlltoAll	Inference KV-cache access
	I	KV3.2	16GB	>= 100	8	AlltoAll	
Storage	J	7:1	Continuous	---	8	Incast	Storage
Total					**128**		

Example of Multi-Tenant Collective Mix Test

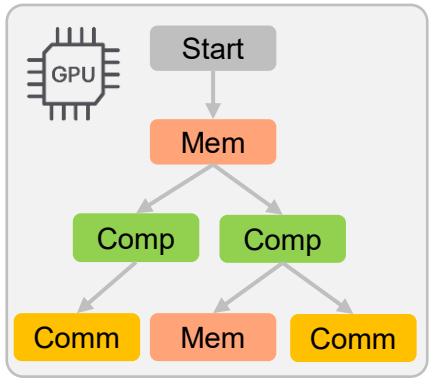


Duty Cycle Distribution

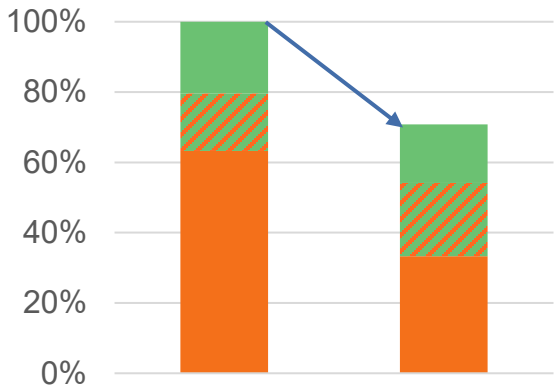


Performance Degradation

Workload Emulation



Training Time



KEYSIGHT AI Data Center Builder | Home > Browse Workloads

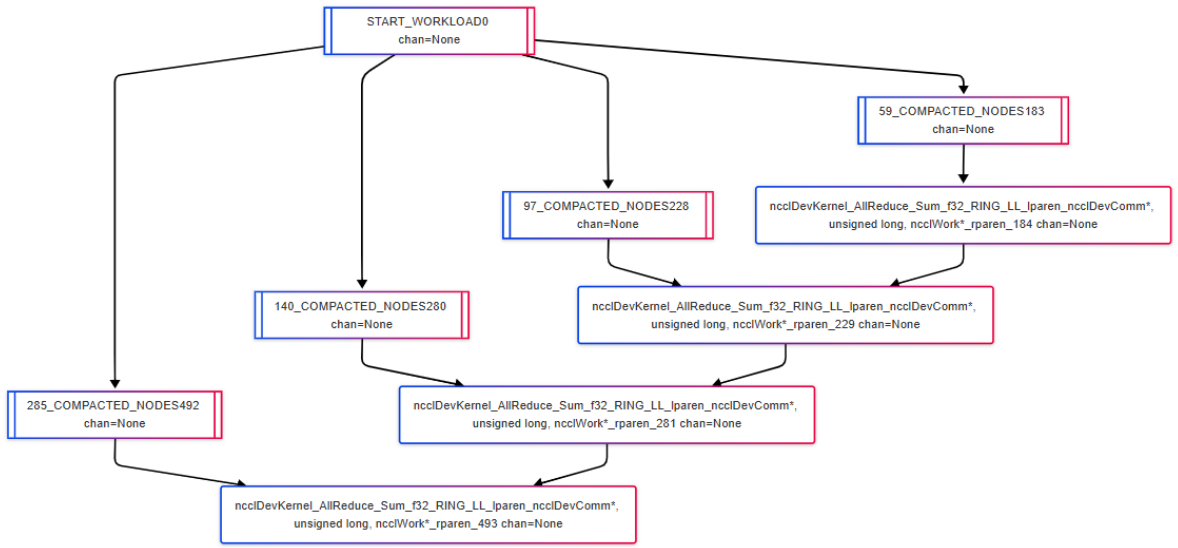
Browse Workloads

Filter

- All Workloads
- ddp
  - 8ranks
  - 16ranks

Num Ranks Info ddp / 16ranks

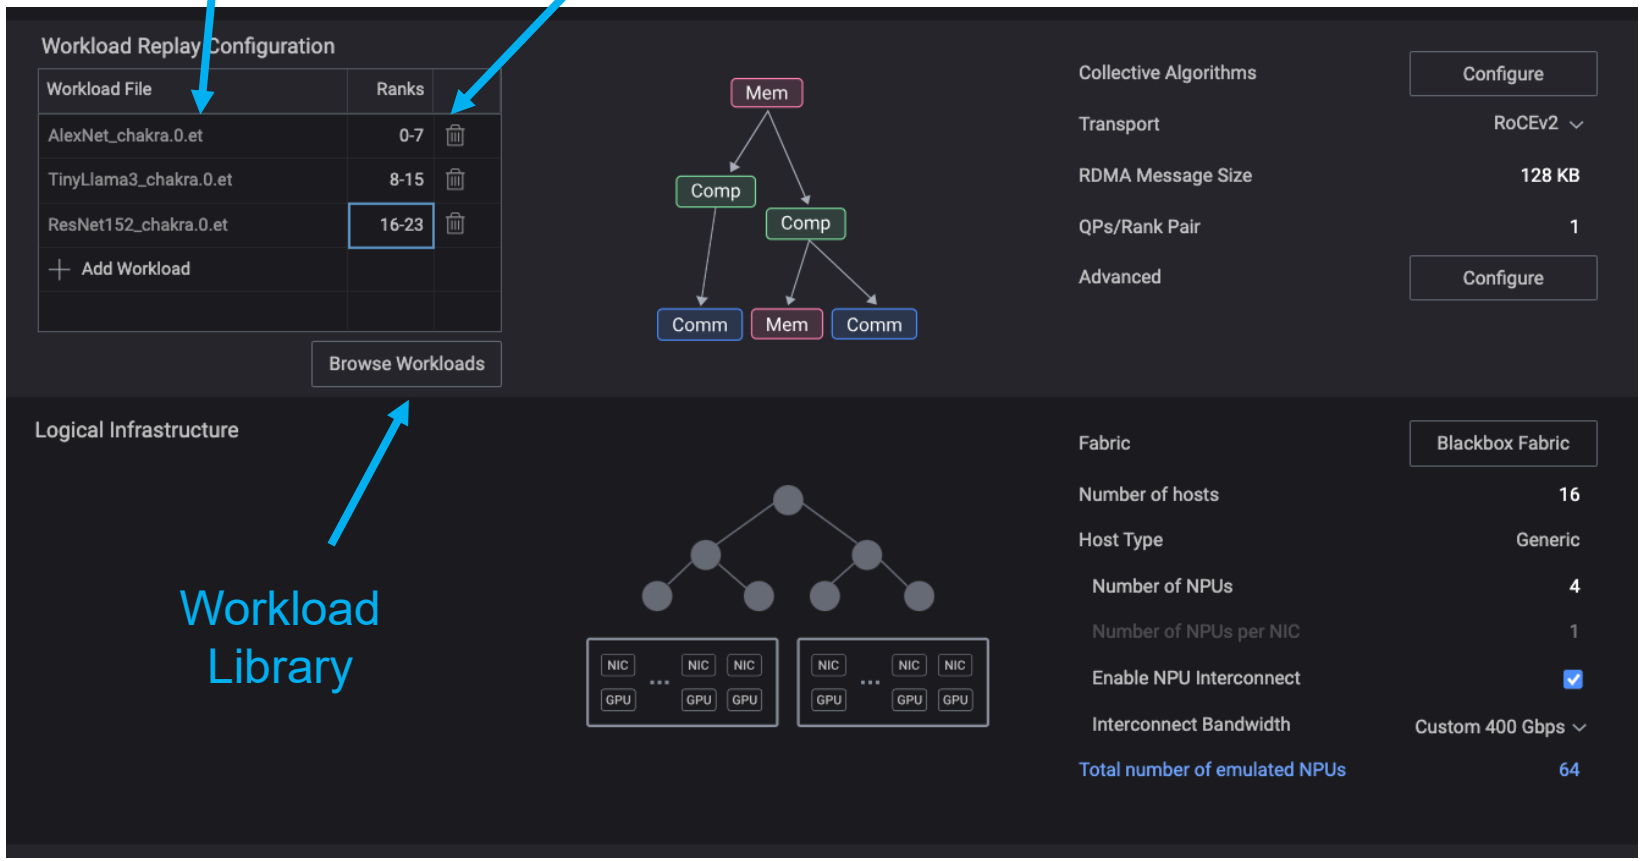
Name	Capture Date	PyTorch Version	NCCL Version	CUDA Version	Chakra Version	Partitioning Schema	Rank Count	Process Group Count	Batch Size	Epochs	Params
AlexNet	2/28/2025, 11:11:07 AM	2.4.0+cu121	2.20.5	11.5.119	0.0.4	DDP	16	1	8	1	61,100,840
ResNet152	2/28/2025, 11:29:54 AM	2.4.0+cu121	2.20.5	11.5.119	0.0.4	DDP	16	1	8	1	60,192,808
TinyLlam...	2/28/2025, 12:05:33 PM	2.4.0+cu121	2.20.5	11.5.119	0.0.4	DDP	16	1	1	1	1,204,881,408
TinyLlam...	2/28/2025, 3:44:16 PM	2.4.0+cu121	2.20.5	11.5.119	0.0.4	DDP	16	1	1	1	396,657,664
TinyMistral	2/28/2025, 4:24:19 PM	2.4.0+cu121	2.20.5	11.5.119	0.0.4	DDP	16	1	1	1	416,319,488
TinyPhi3	2/28/2025, 8:36:07 PM	2.4.0+cu121	2.20.5	11.5.119	0.0.4	DDP	16	1	1	1	433,358,848



Workload Emulation Configuration

Single or multiple parallel workloads

Workload placement

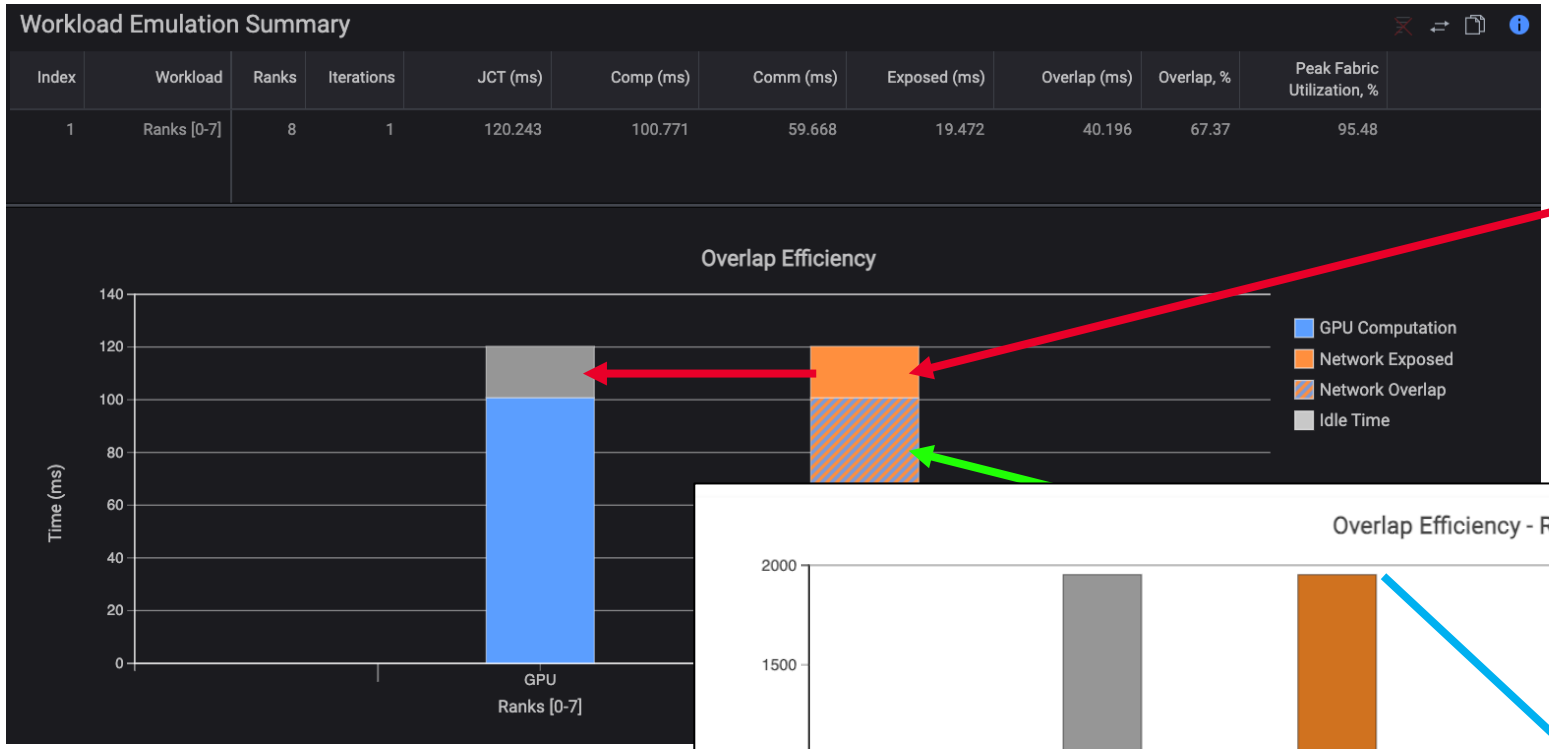


Workload Library

CCL and RDMA configuration

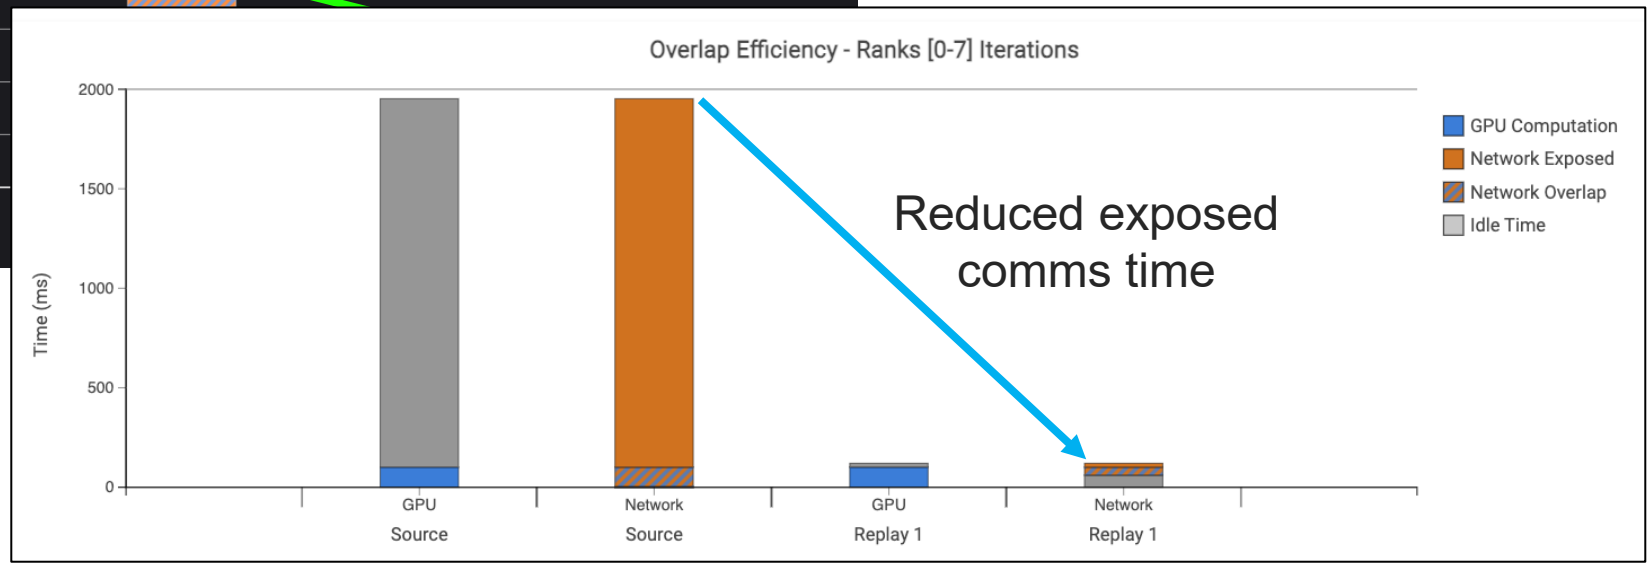
Emulated AI host configuration

Workload Emulation Reports



Exposed communications

GPU Activity

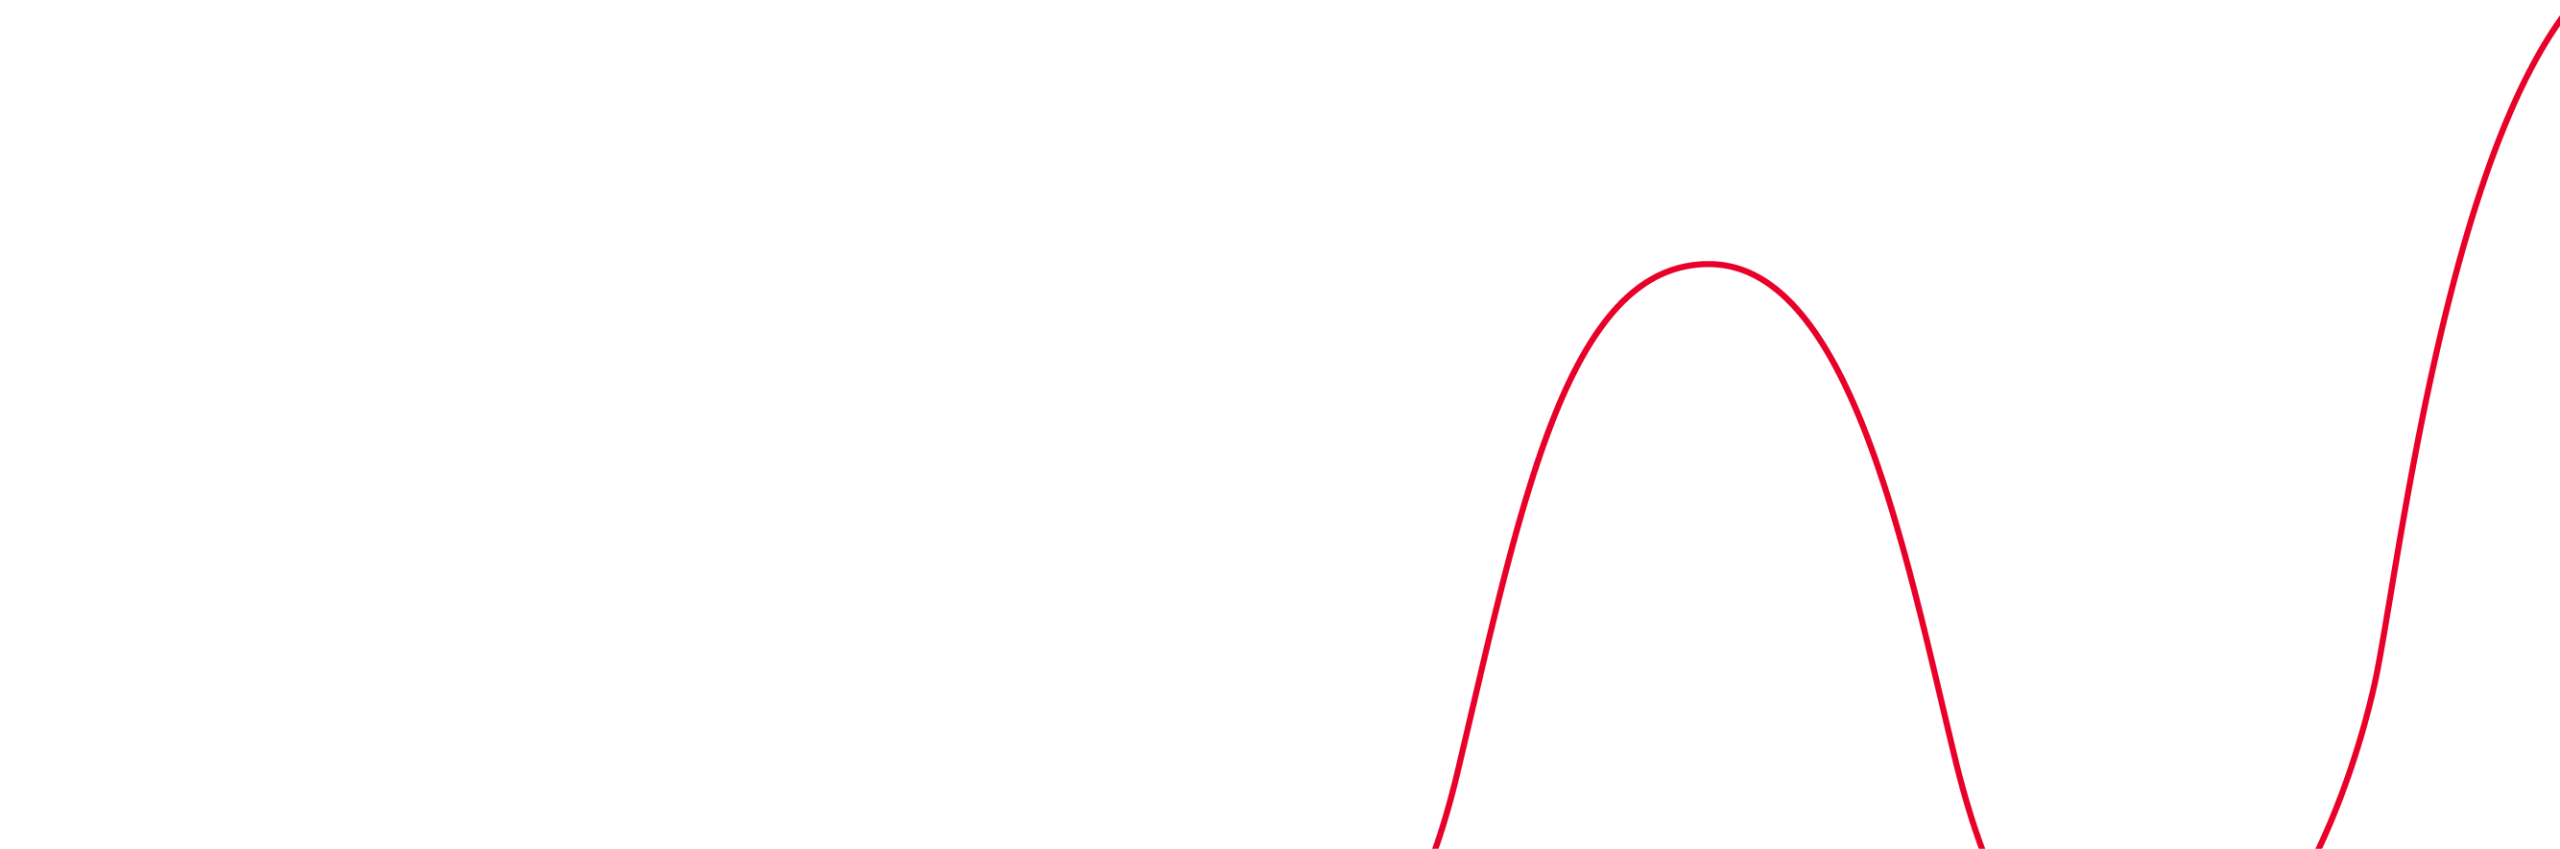


Reduced exposed comms time

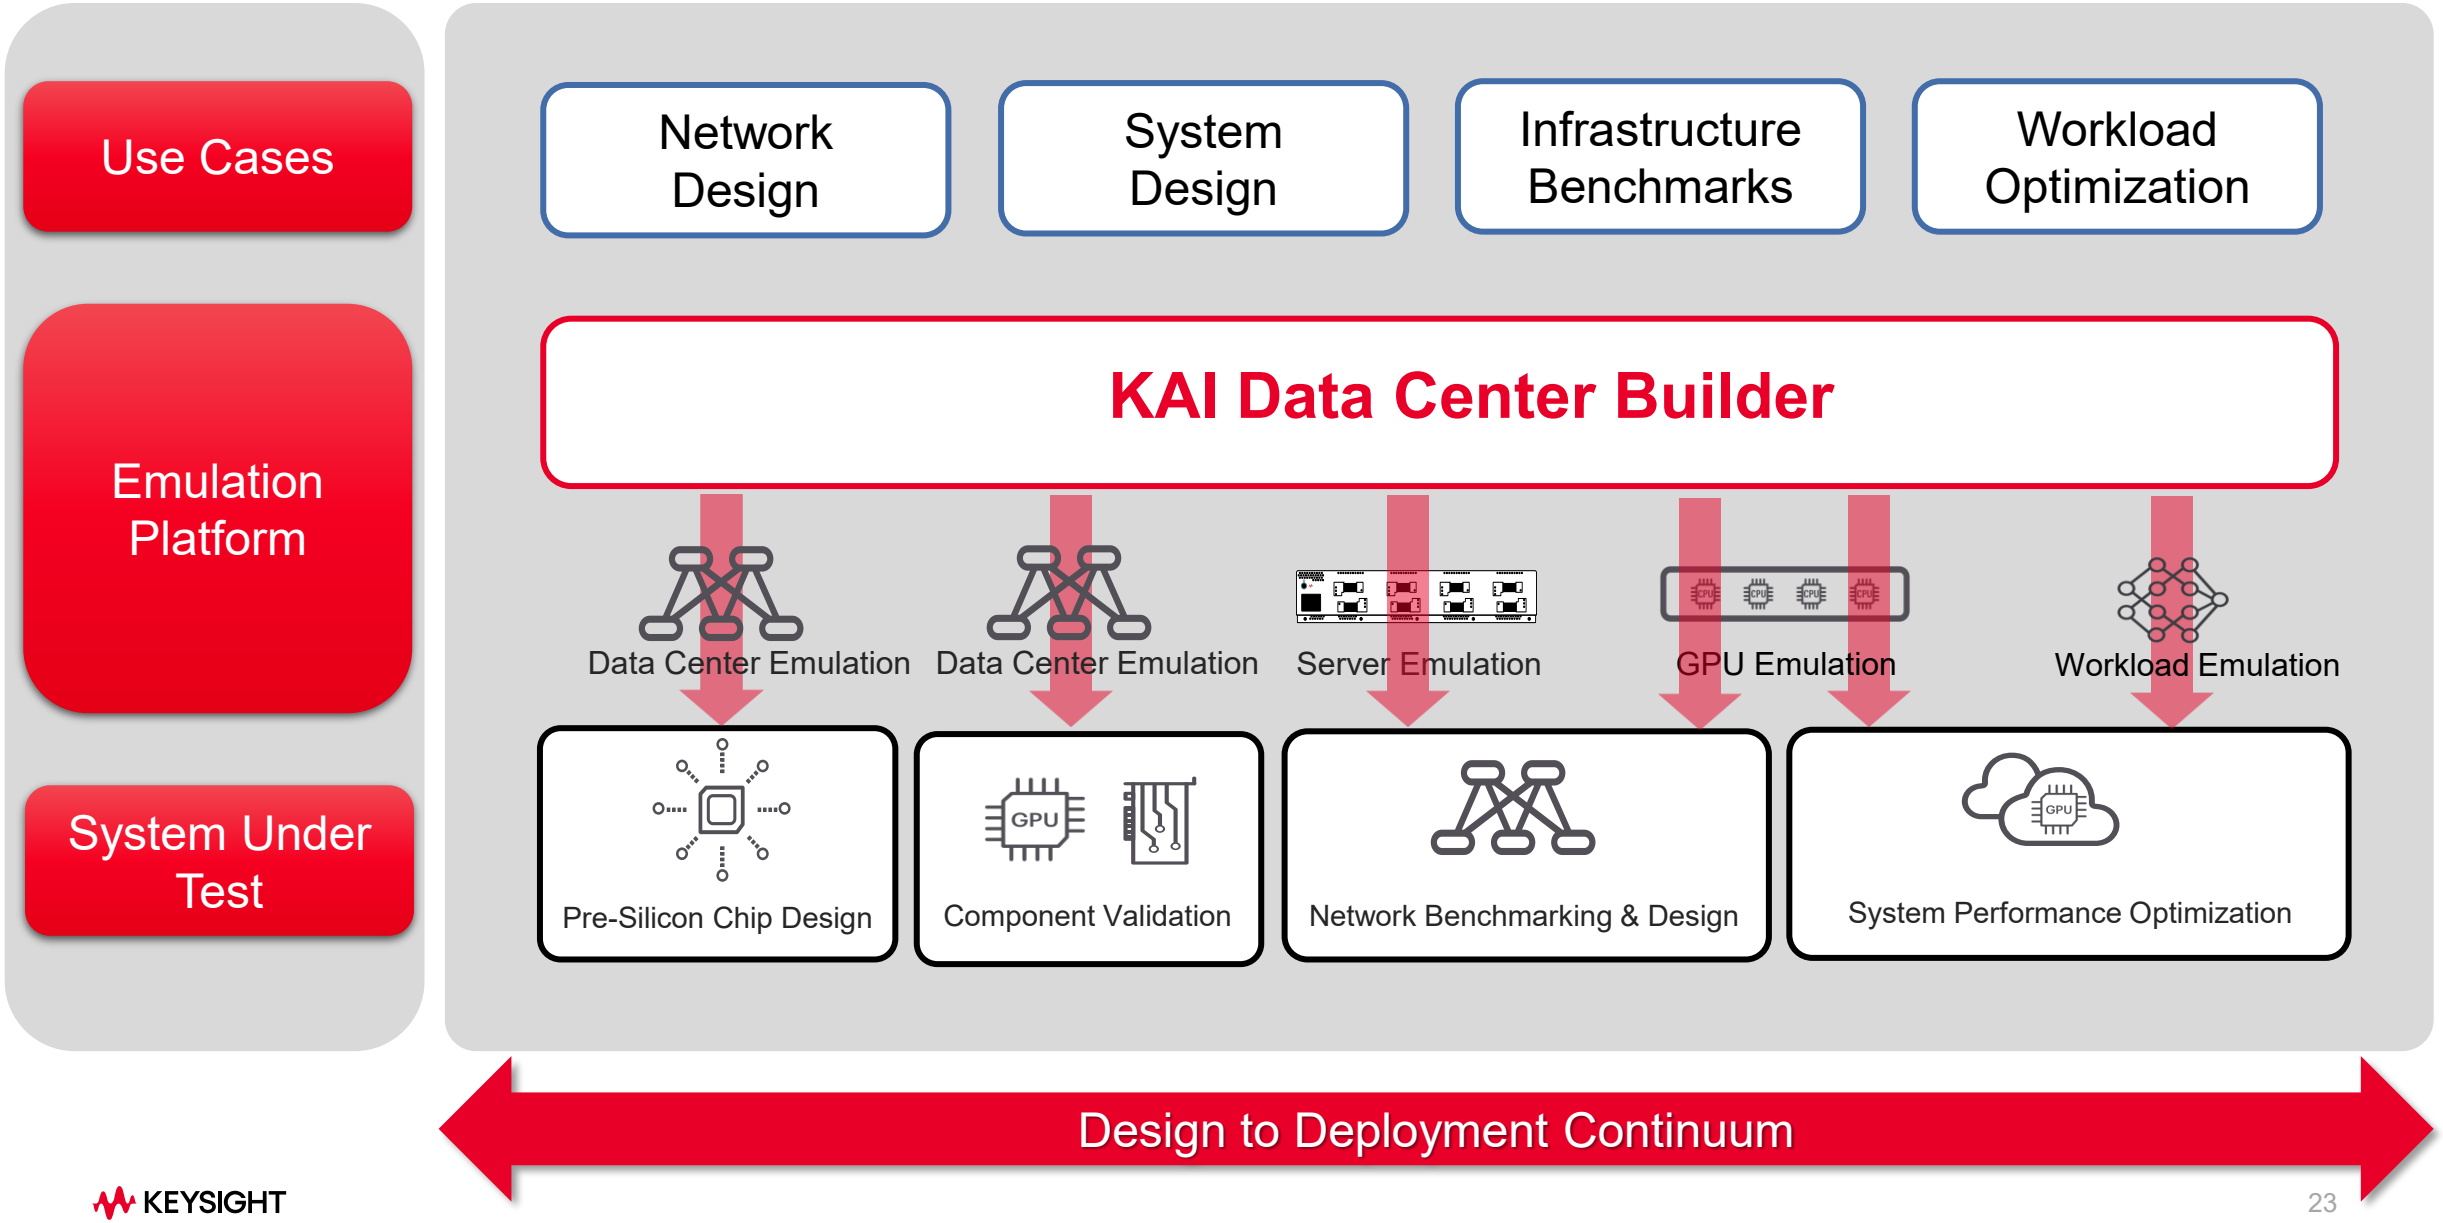
Original workload

Replay on faster network

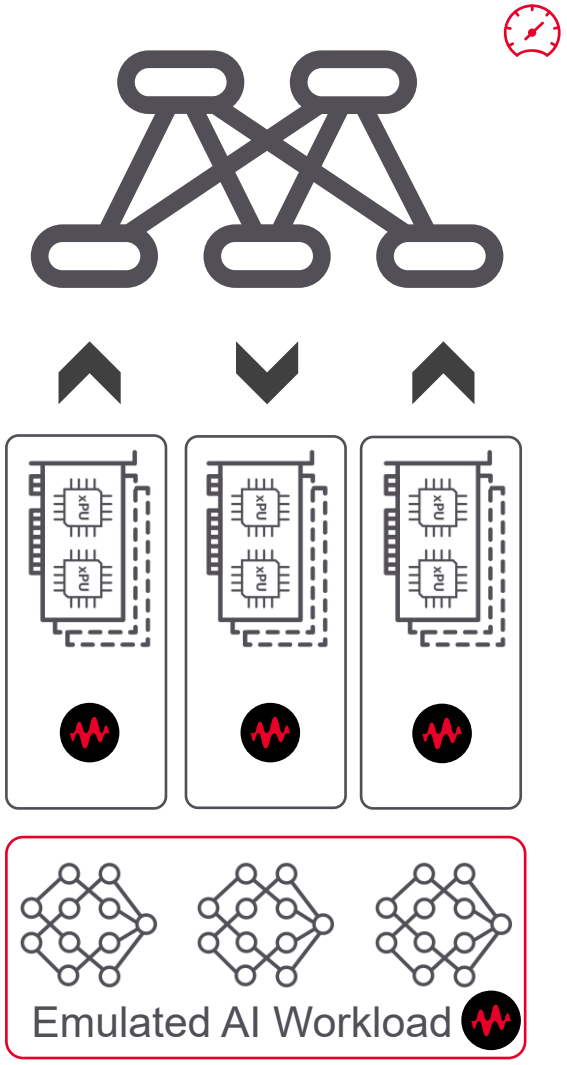
KAI DC Builder Platform



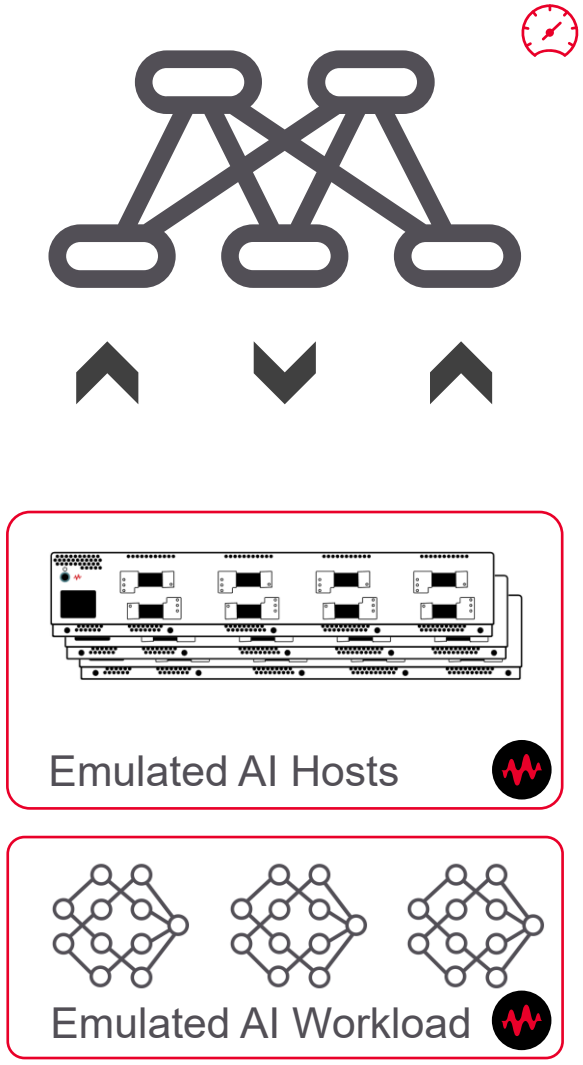
Product lifecycle & Use case Coverage



1 Keysight Software Solution



2 Keysight Hardware Solution



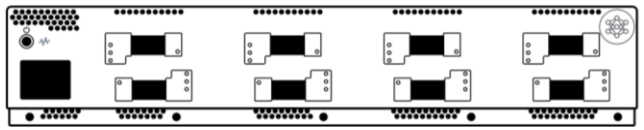
Software Benefits

- ✓ NIC+Fabric co-tuning
- ✓ Cost
- ✓ New transports
- ✓ Production & Cloud

Hardware Benefits

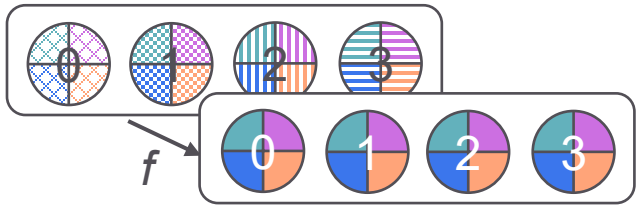
- Isolated fabric validation
- 800G throughput
- Deep network insights

Key Capabilities



AresONE NIC & RoCEv2 Emulation

- RoCEv2 RC mode
- 100/200/400/800GE at 56/112
- Per-QP performance metrics
- 4K QPs per port
- PFC, ECN, DCQCN



KAI CCL Emulation

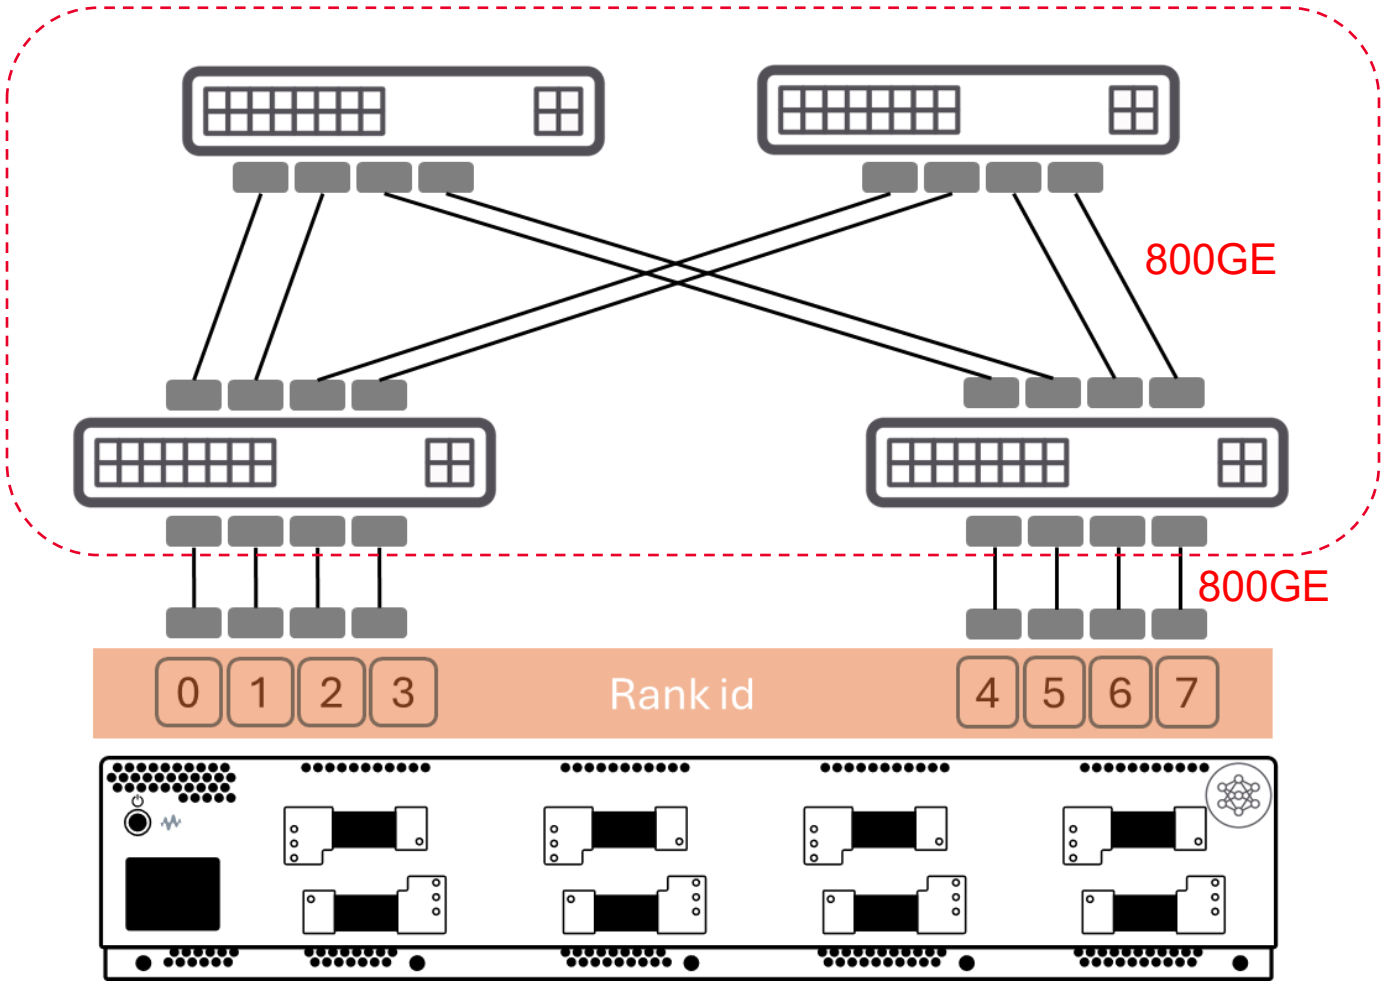
- Microbenchmarks
- Workloads
- Per-chunk, per-flow tracking
- Realistic LB and QoS pattens
- Generic NPU interconnect simulation



Impairments

- Background traffic
- Selective loss
- Selective reordering
- PFC backpressure

Starter Setup



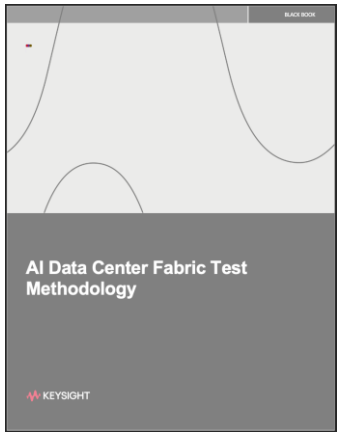
Top capabilities

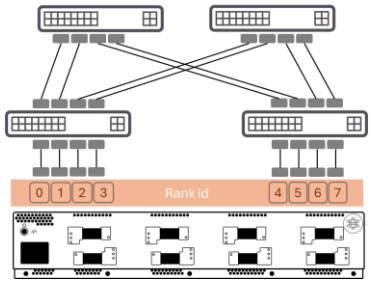
- 100/200/400/800GE
- 4K QP per port
- PFC, ECN & DCQCN
- 6 Collectives, 4 Algorithms
- NVLink Emulation
- Single, parallel collectives
- AI workloads
- Impairments

Keysight AI DC Builder Solution

Use case coverage today

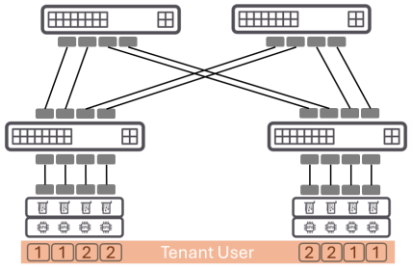
https://www.keysight.com/us/en/assets/3124-1729/application-notes/AI-Fabric-Test-Methodology.pdf





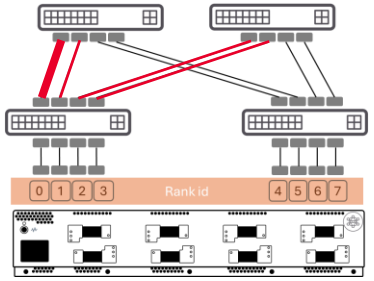
Benchmarking

- Topologies
- Algorithms
- Data sizes
- RDMA message sizes



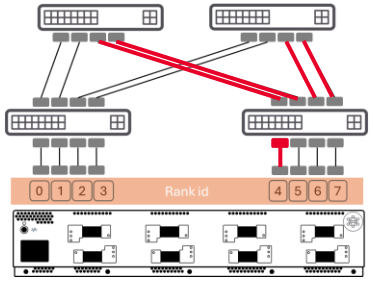
Performance Isolation

- Noisy neighbors
- Parallel collectives



Load Balancing

- ECMP hashing
- Traffic Engineering
- Q-Pair awareness
- Parallel Q-Pairs
- Dynamic Load Balancing



Congestion Control

- PFC
- ECN/DCQCN
- Fine-tuning

Learn More About the Keysight AI Data Center Builder



PDL-KAI-DC-Builder@keysight.com

Product Page

Methodology

Solution Brief

Documentation


KEYSIGHT