Deploying AI at Meta Scale

BUILDING DEVELOPER-FRIENDLY HIGH-PERFORMANCE SCALABLE SYSTEMS FOR AI + PYTORCH

Alexis Bjorlin
VP, Infrastructure Engineering

Meta

Community Statistics

3.65B

people use at least one of our services **monthly**, approximately

2.91B

monthly active users on Facebook

700M

people use Augmented Reality across our apps and devices monthly

Meta

Snapshot of Data Center Footprint

>100

Data centers

~50M

Square feet

100%

Renewable energy

New Albany, Ohio

How we **use** AI…

AI AT META

Content Understanding

- Identify and eliminate inappropriate content before being viewed
- Leverages several model types, including computer vision, image classification, and natural language processing

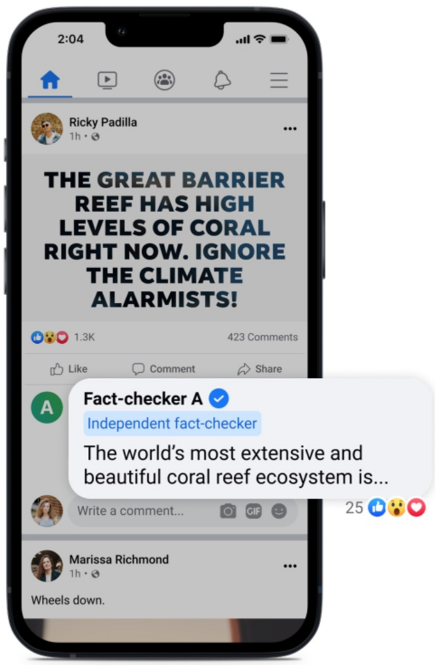

Ricky Padilla
1h

THE GREAT BARRIER REEF HAS HIGH LEVELS OF CORAL RIGHT NOW. IGNORE THE CLIMATE ALARMISTS!

Fact-checker A
Independent fact-checker
The world's most extensive and beautiful coral reef ecosystem is...

Marissa Richmond
1h
Wheels down.

~250M
Violating Content (Q2)

95%+
Actioned on by AI

AI AT META

Natural Language Processing

- Open Pretrained Transformer (OPT-175B)
- XMLR

Sources:
https://ai.facebook.com/research/no-language-left-behind/
https://ai.facebook.com/blog/democratizing-access-to-large-scale-language-models-with-opt-175b/

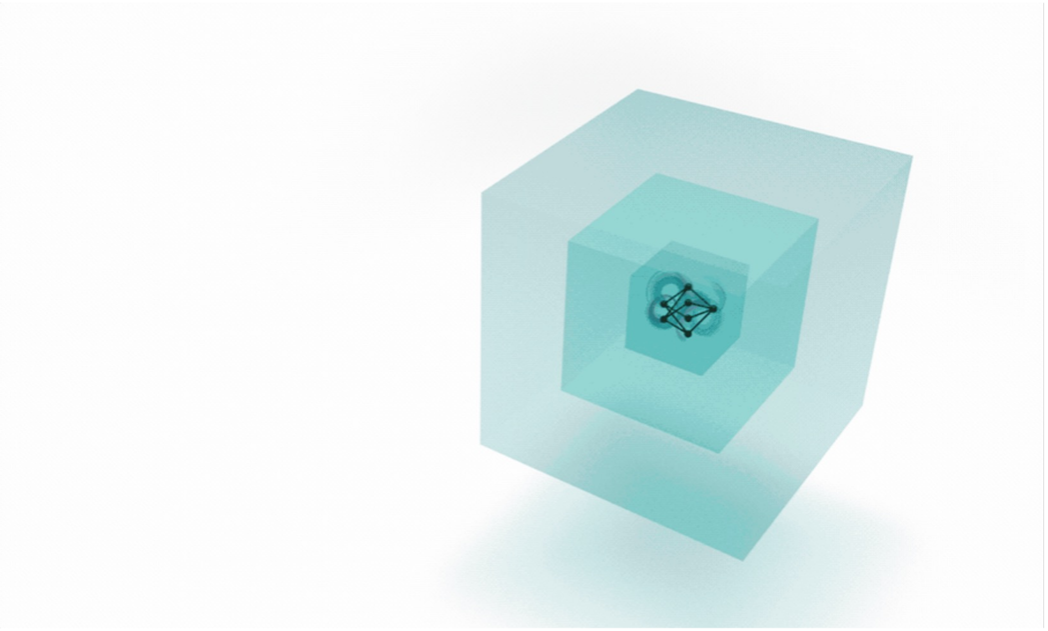

AI AT META

Recommendation & Personalization

- Deep Learning Recommendation Model (DLRM)
- TBSM, DHEN

Source: https://ai.facebook.com/blog/dlrm-an-advanced-open-source-deep-learning-recommendation-model/

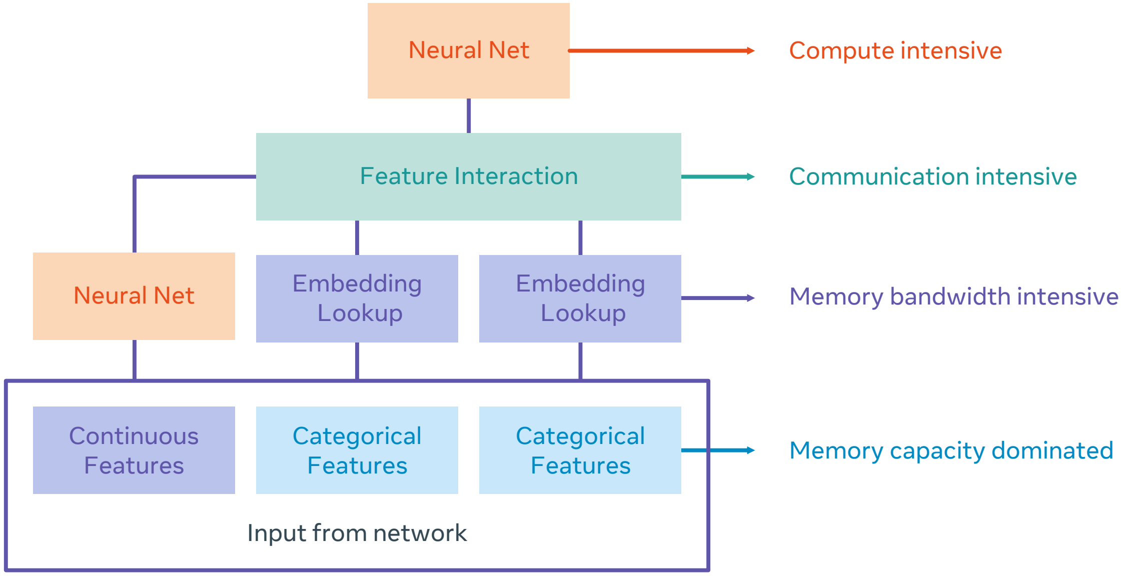

DEEP LEARNING WORKLOADS - CHARACTERISTICS

SIZE

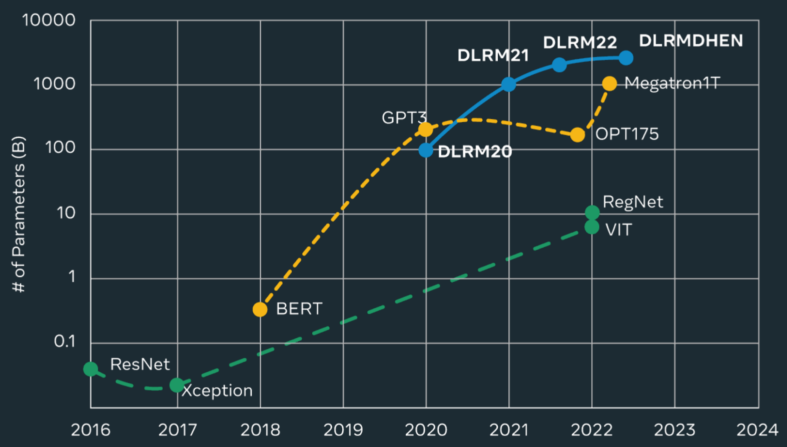

COMPUTE

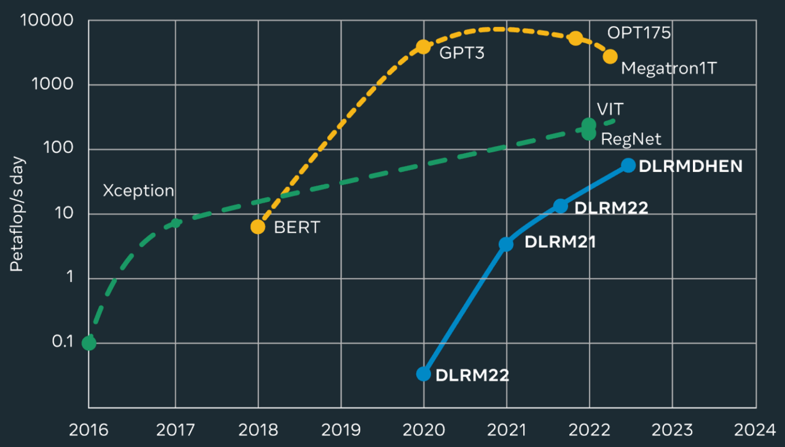

AI IS POWERING EVERYTHING WE DO: AI-related statistics

6B

Training images

20B+

Translations per day

200T+

Predictions per day

How we **develop** AI…

THE LANGUAGE OF AI

PyTorch

ALL IN ON AI WITH PYTORCH

Rapid research to production

- Benefit of deploying state-of-art models fast can be huge

Source:
https://paperswithcode.com/trends

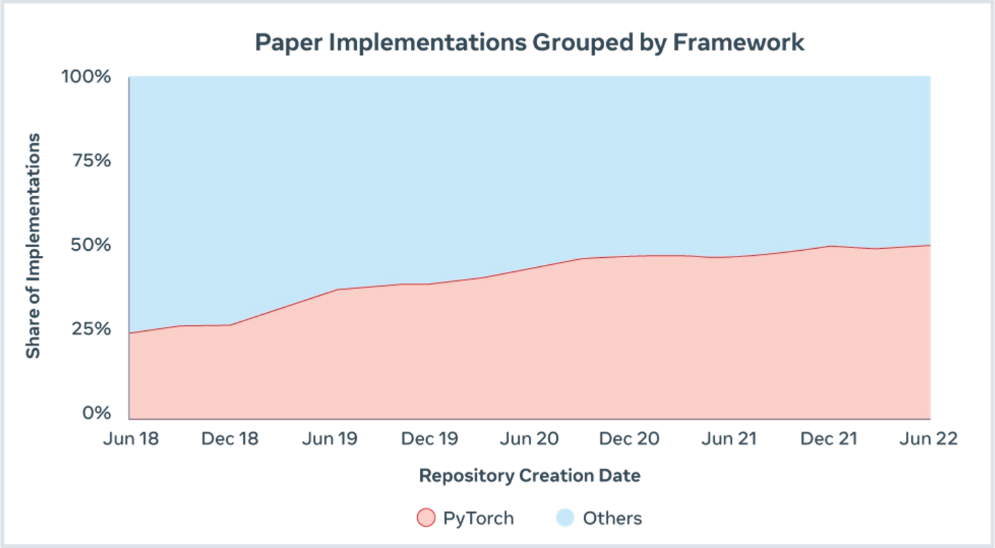

PYTORCH: A DEVELOPER-FIRST MINDSET

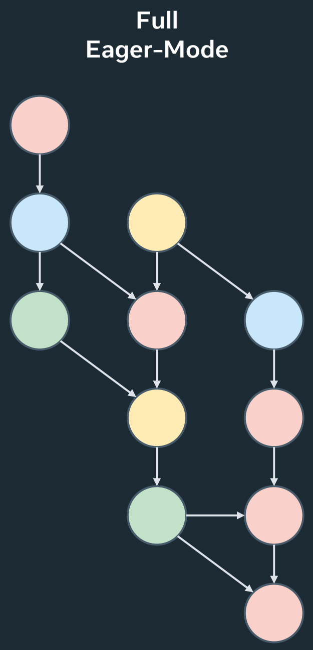

```
for epoch in range (max_epochs):
    for model_target, model_input in enumerate(training_data):
        if numpy.random.randint(100) > 90: # 10% noise
            model_target = torch.from_numpy(numpy.random.randint(2))
        model_output = dlrm(model_input)
        model_loss   = torch.nn.BCELoss(model_output, model_target)
        model_loss.backward()
        optimizer.step()
        print("BCE loss " + str(model_loss))
        matplotlib.pyplot.plot(...)
```

DEVELOPER EFFICIENCY VS. PERFORMANCE

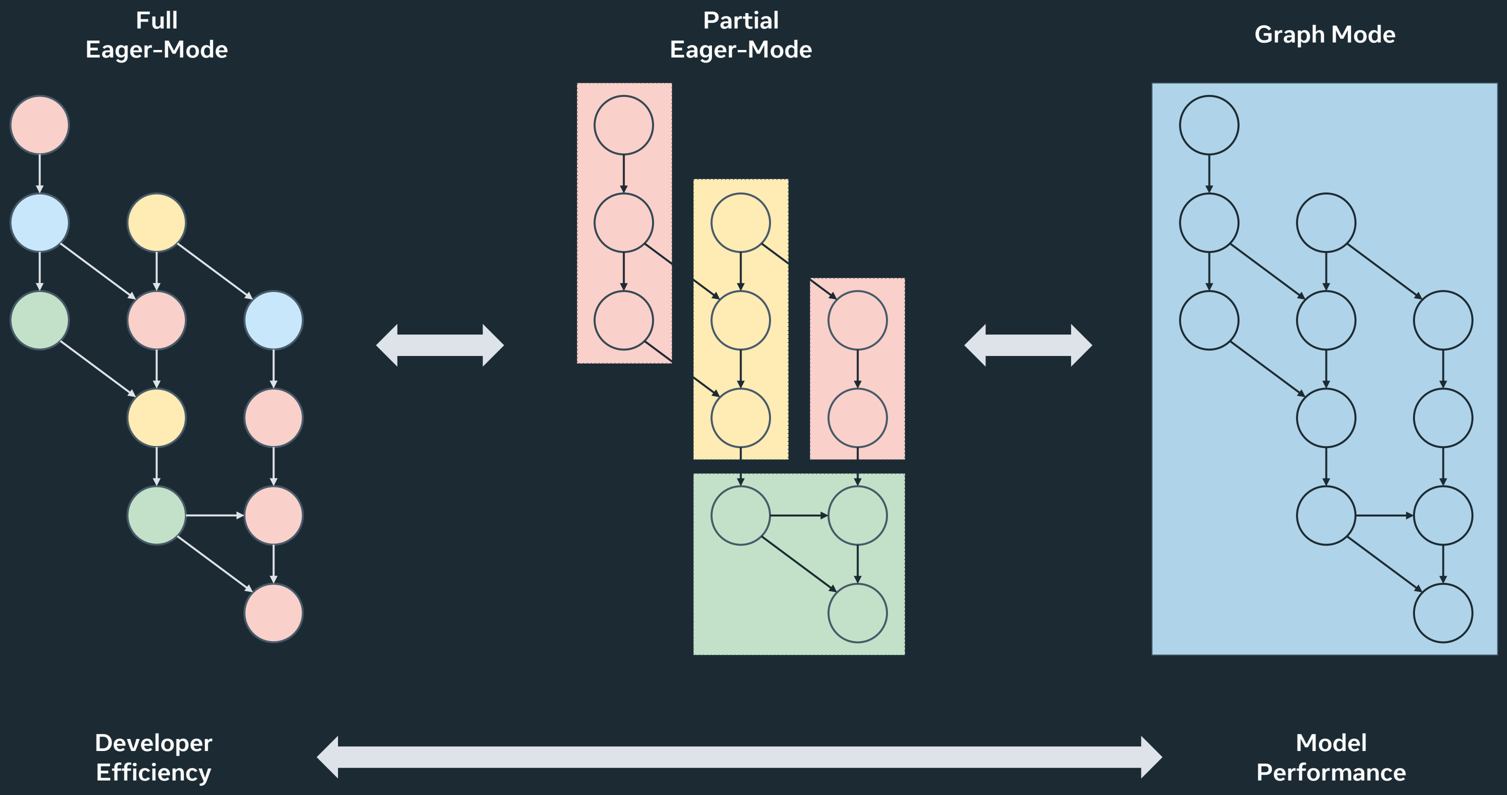

How we **enable** AI…

OPTIMIZING THE AI SYSTEM FOR PYTORCH

Programmable

Easy operator authoring → new compute primitives for model innovation.

Dynamic

Fast operator launch.
Dynamic tensor shapes, memory allocation, easy to prototype

Scalable & Tunable

Future-proof.

Balance between compute, memory, network.

NETWORK IS LAGGING ON EVOLUTION CURVE

Scaling of Peak hardware FLOPS, and Memory/Interconnect Bandwidth

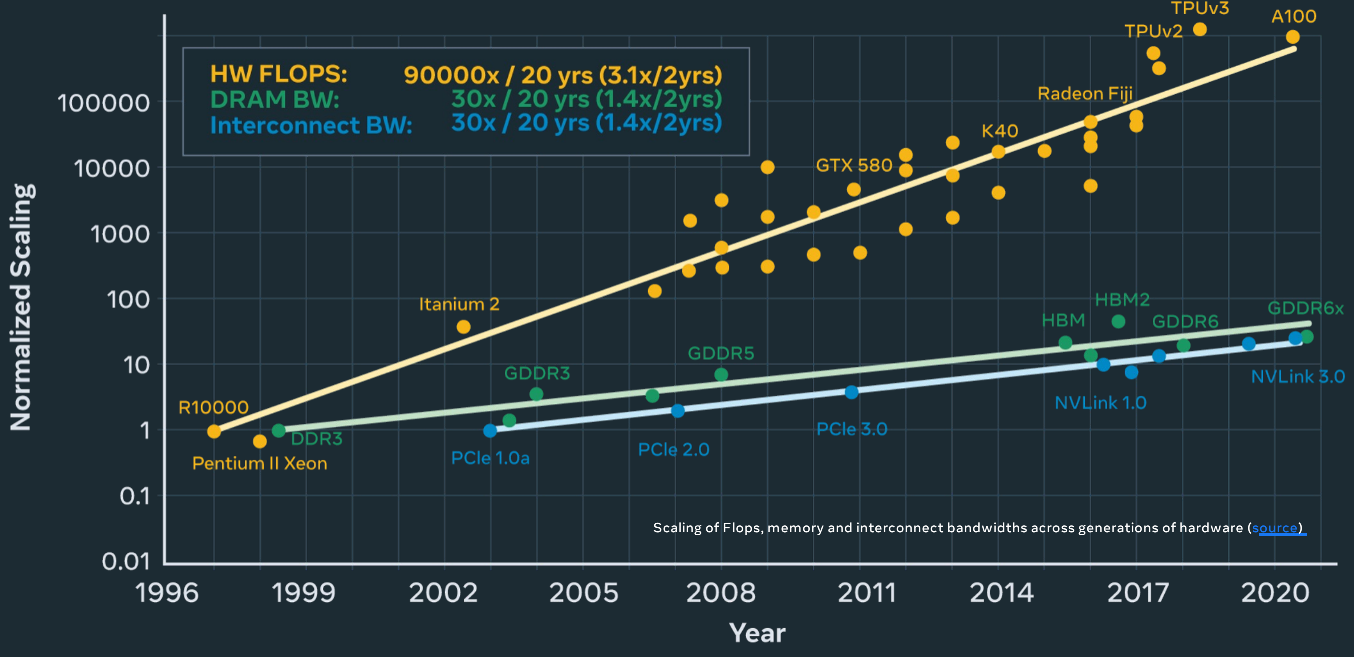

Scaling of Flops, memory and interconnect bandwidths across generations of hardware (source)

NETWORK I/O IS KEY FOR RECOMMENDATION WORKLOADS

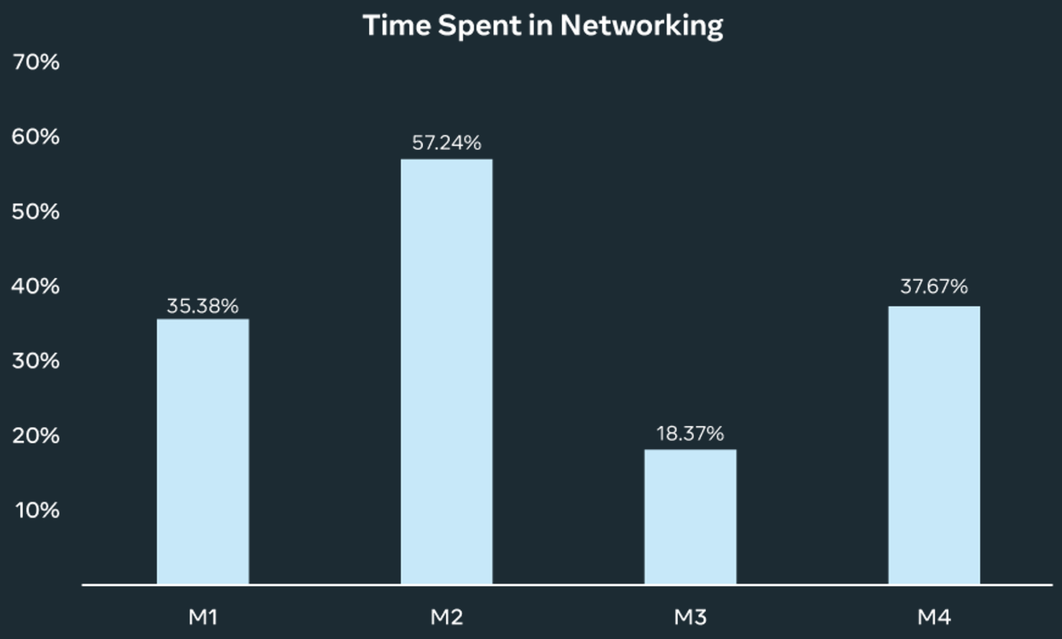
Time Spent in Networking

Ranking requires high **injection** & **bisection** bandwidth

WHAT MAKES AN AI TRAINING CLUSTER?

- Storage Cache
- Data ingestion
- Compute Nodes
- Compute fabric

PyTorch AI Training Cluster

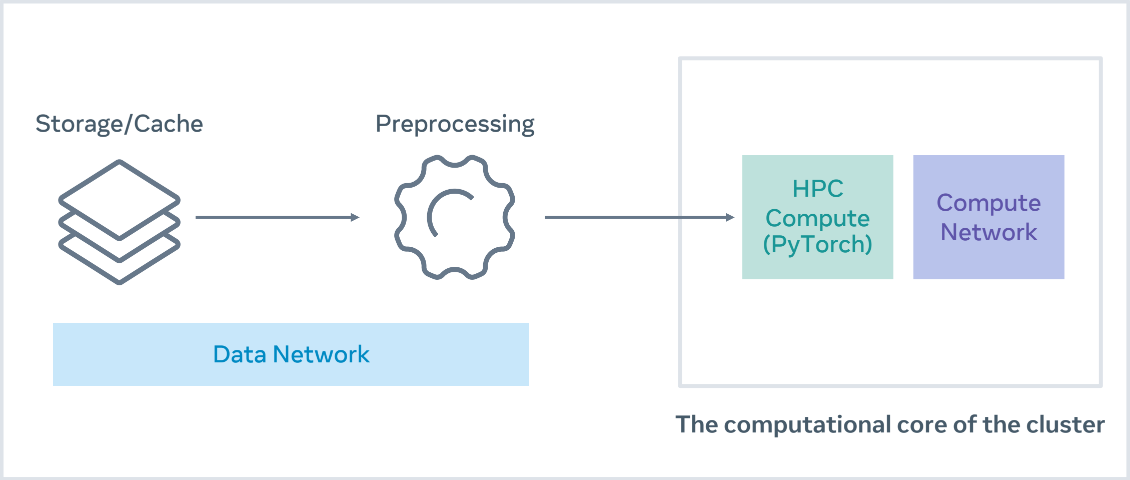

THE PYTORCH AI TRAINING CLUSTER

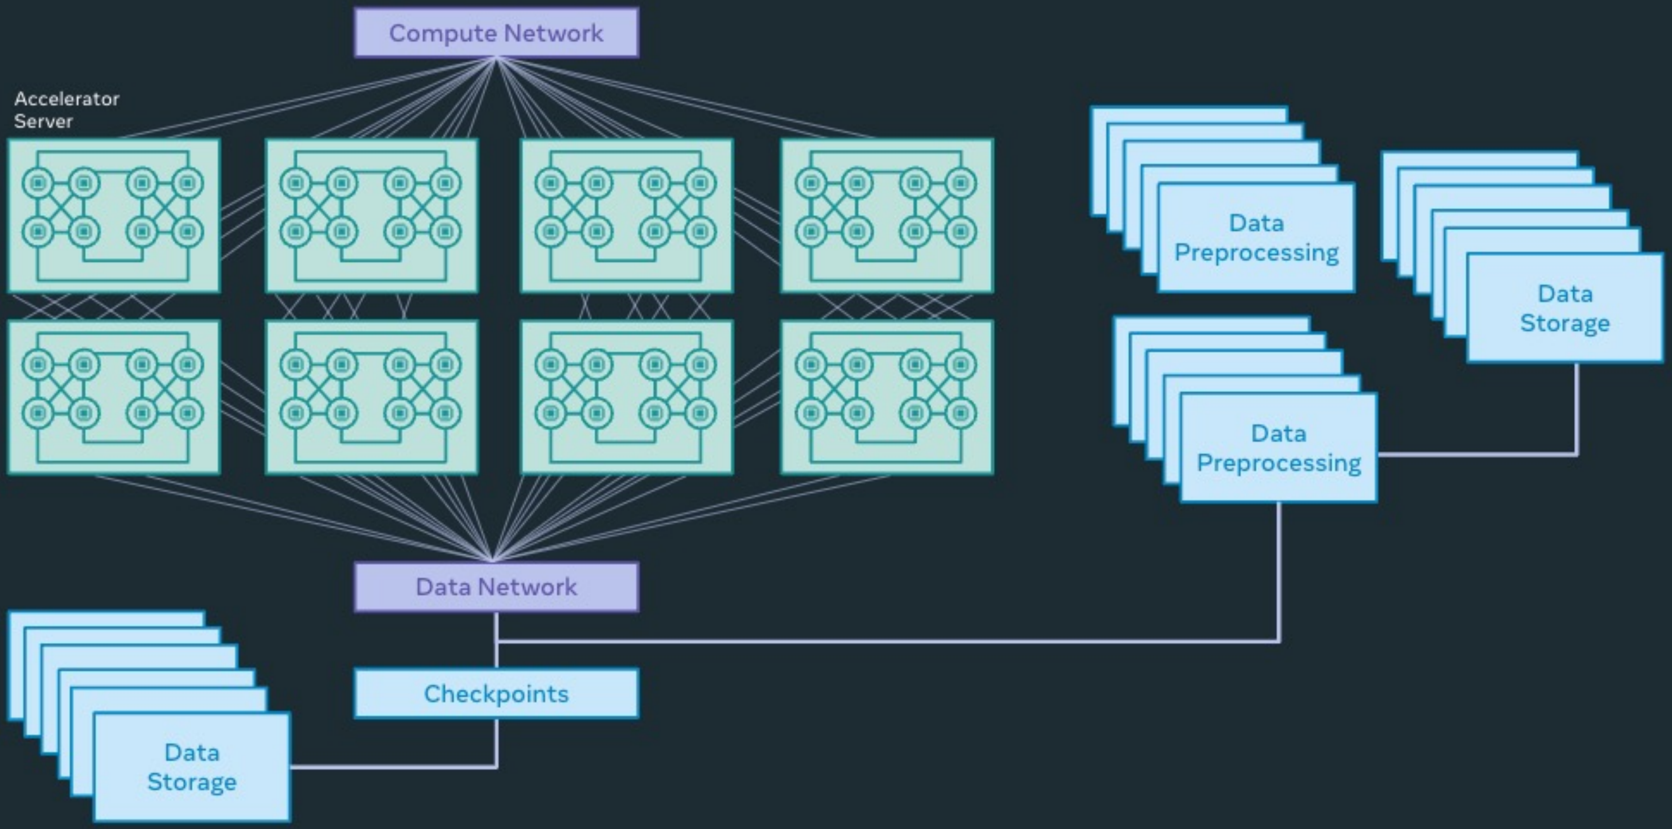

Compute Network

Accelerator Server

Data Preprocessing

Data Storage

Data Preprocessing

Data Network

Checkpoints

Data Storage

THE PYTORCH AI TRAINING CLUSTER OF THE FUTURE (2025+)

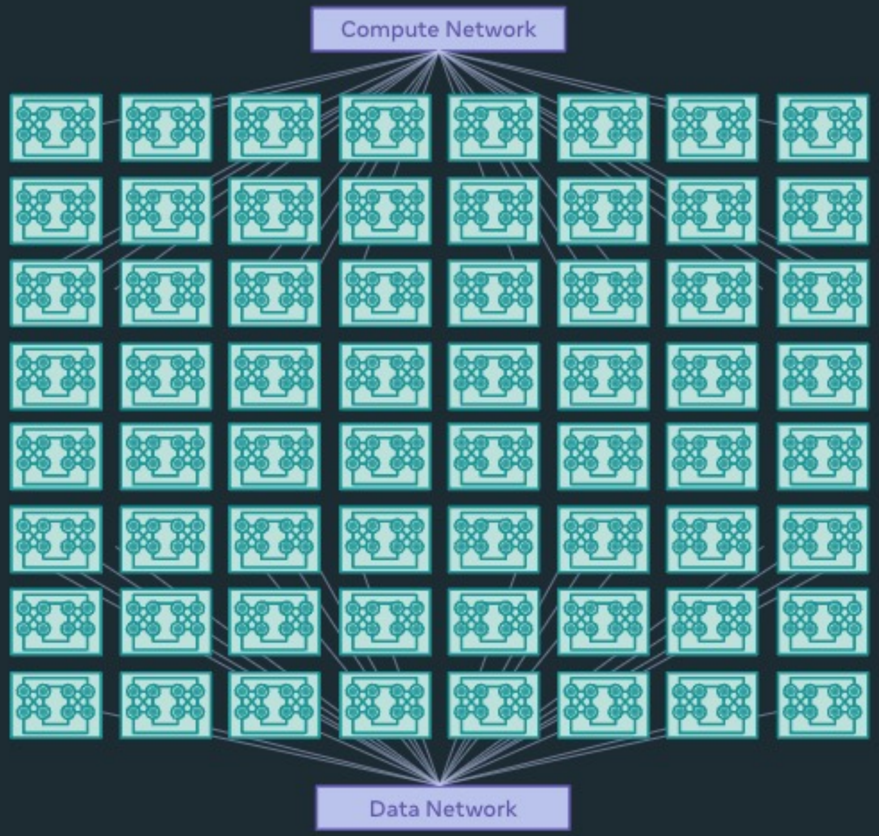

~4K

Accelerators

~1TB/s

of compute
network per accelerator

The cluster fabric is **Non-blocking**

THE PROMISE OF OPTICAL IO

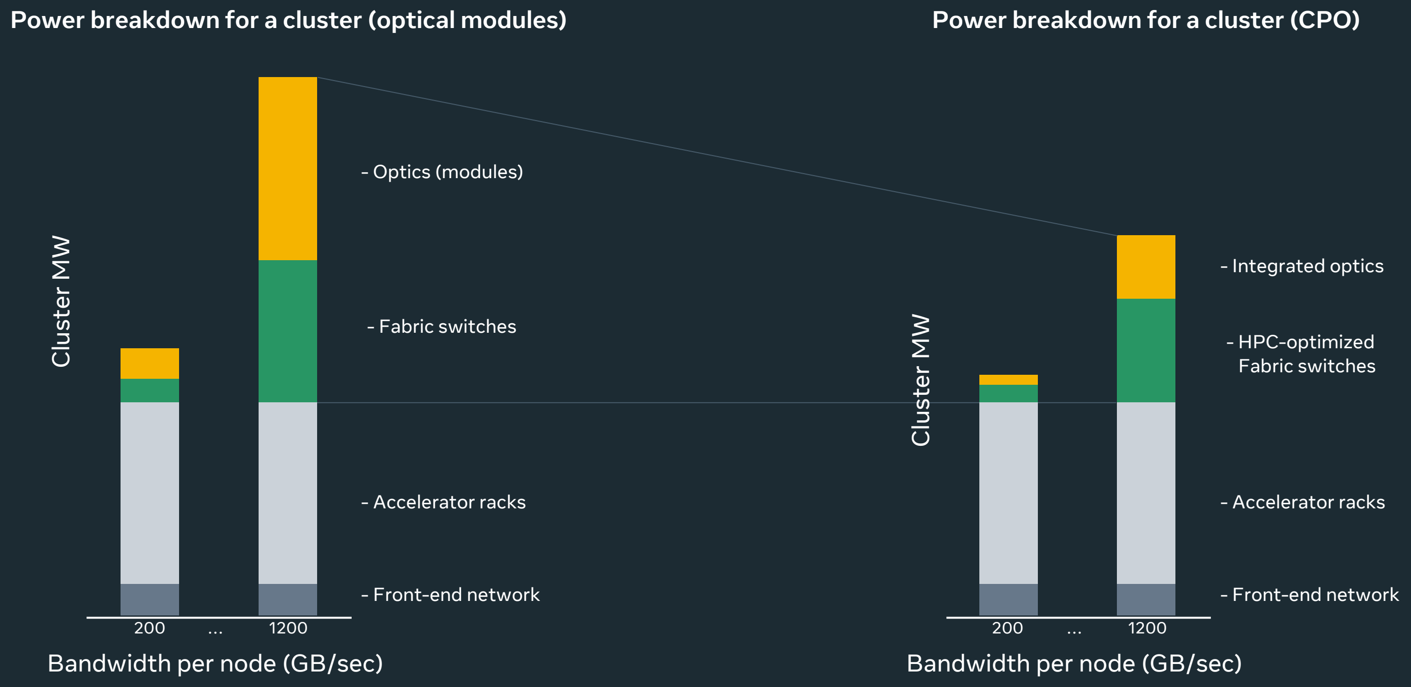

THE BENEFIT OF A PYTORCH AI TRAINING CLUSTER

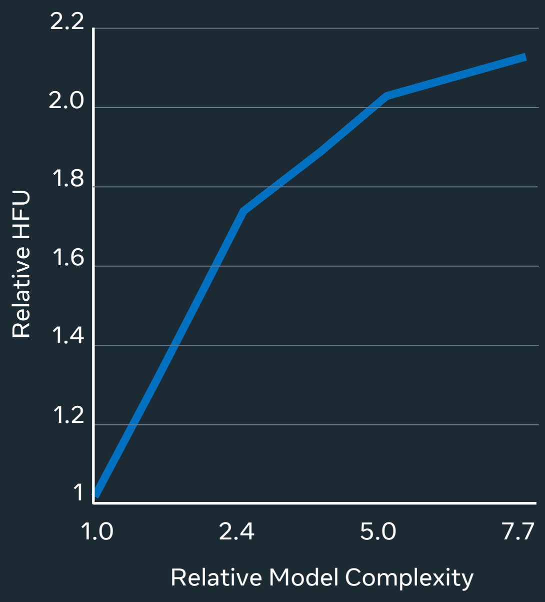

DLRM - DHEN

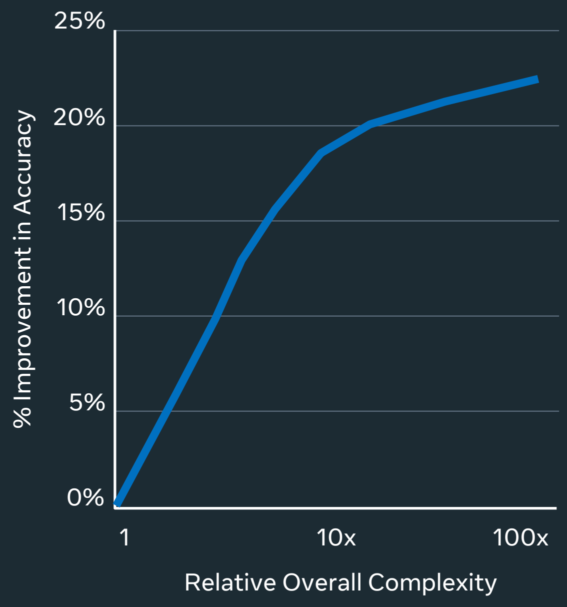

Improving Content Relevance

+24%

improvement in content relevance over 2021-22

CONCLUSION / TAKEAWAYS

Co-design
for PyTorch

Flexibly balance compute,
memory, and network

Plan for
Innovation as
AI evolves

Meta