OCP SUMMIT

March 20-21
2018
San Jose, CA

OPEN
Compute Project

AI HARDWARE INFRASTRUCTURE AT FACEBOOK

Xiaodong Wang
Kevin Lee

Today’s AI Hardware Infrastructure

New Hardware Announcement

What’s Next

Applied Machine Learning

AI Research

AI Infrastructure

OPEN
Compute Project

Facebook
Open Source

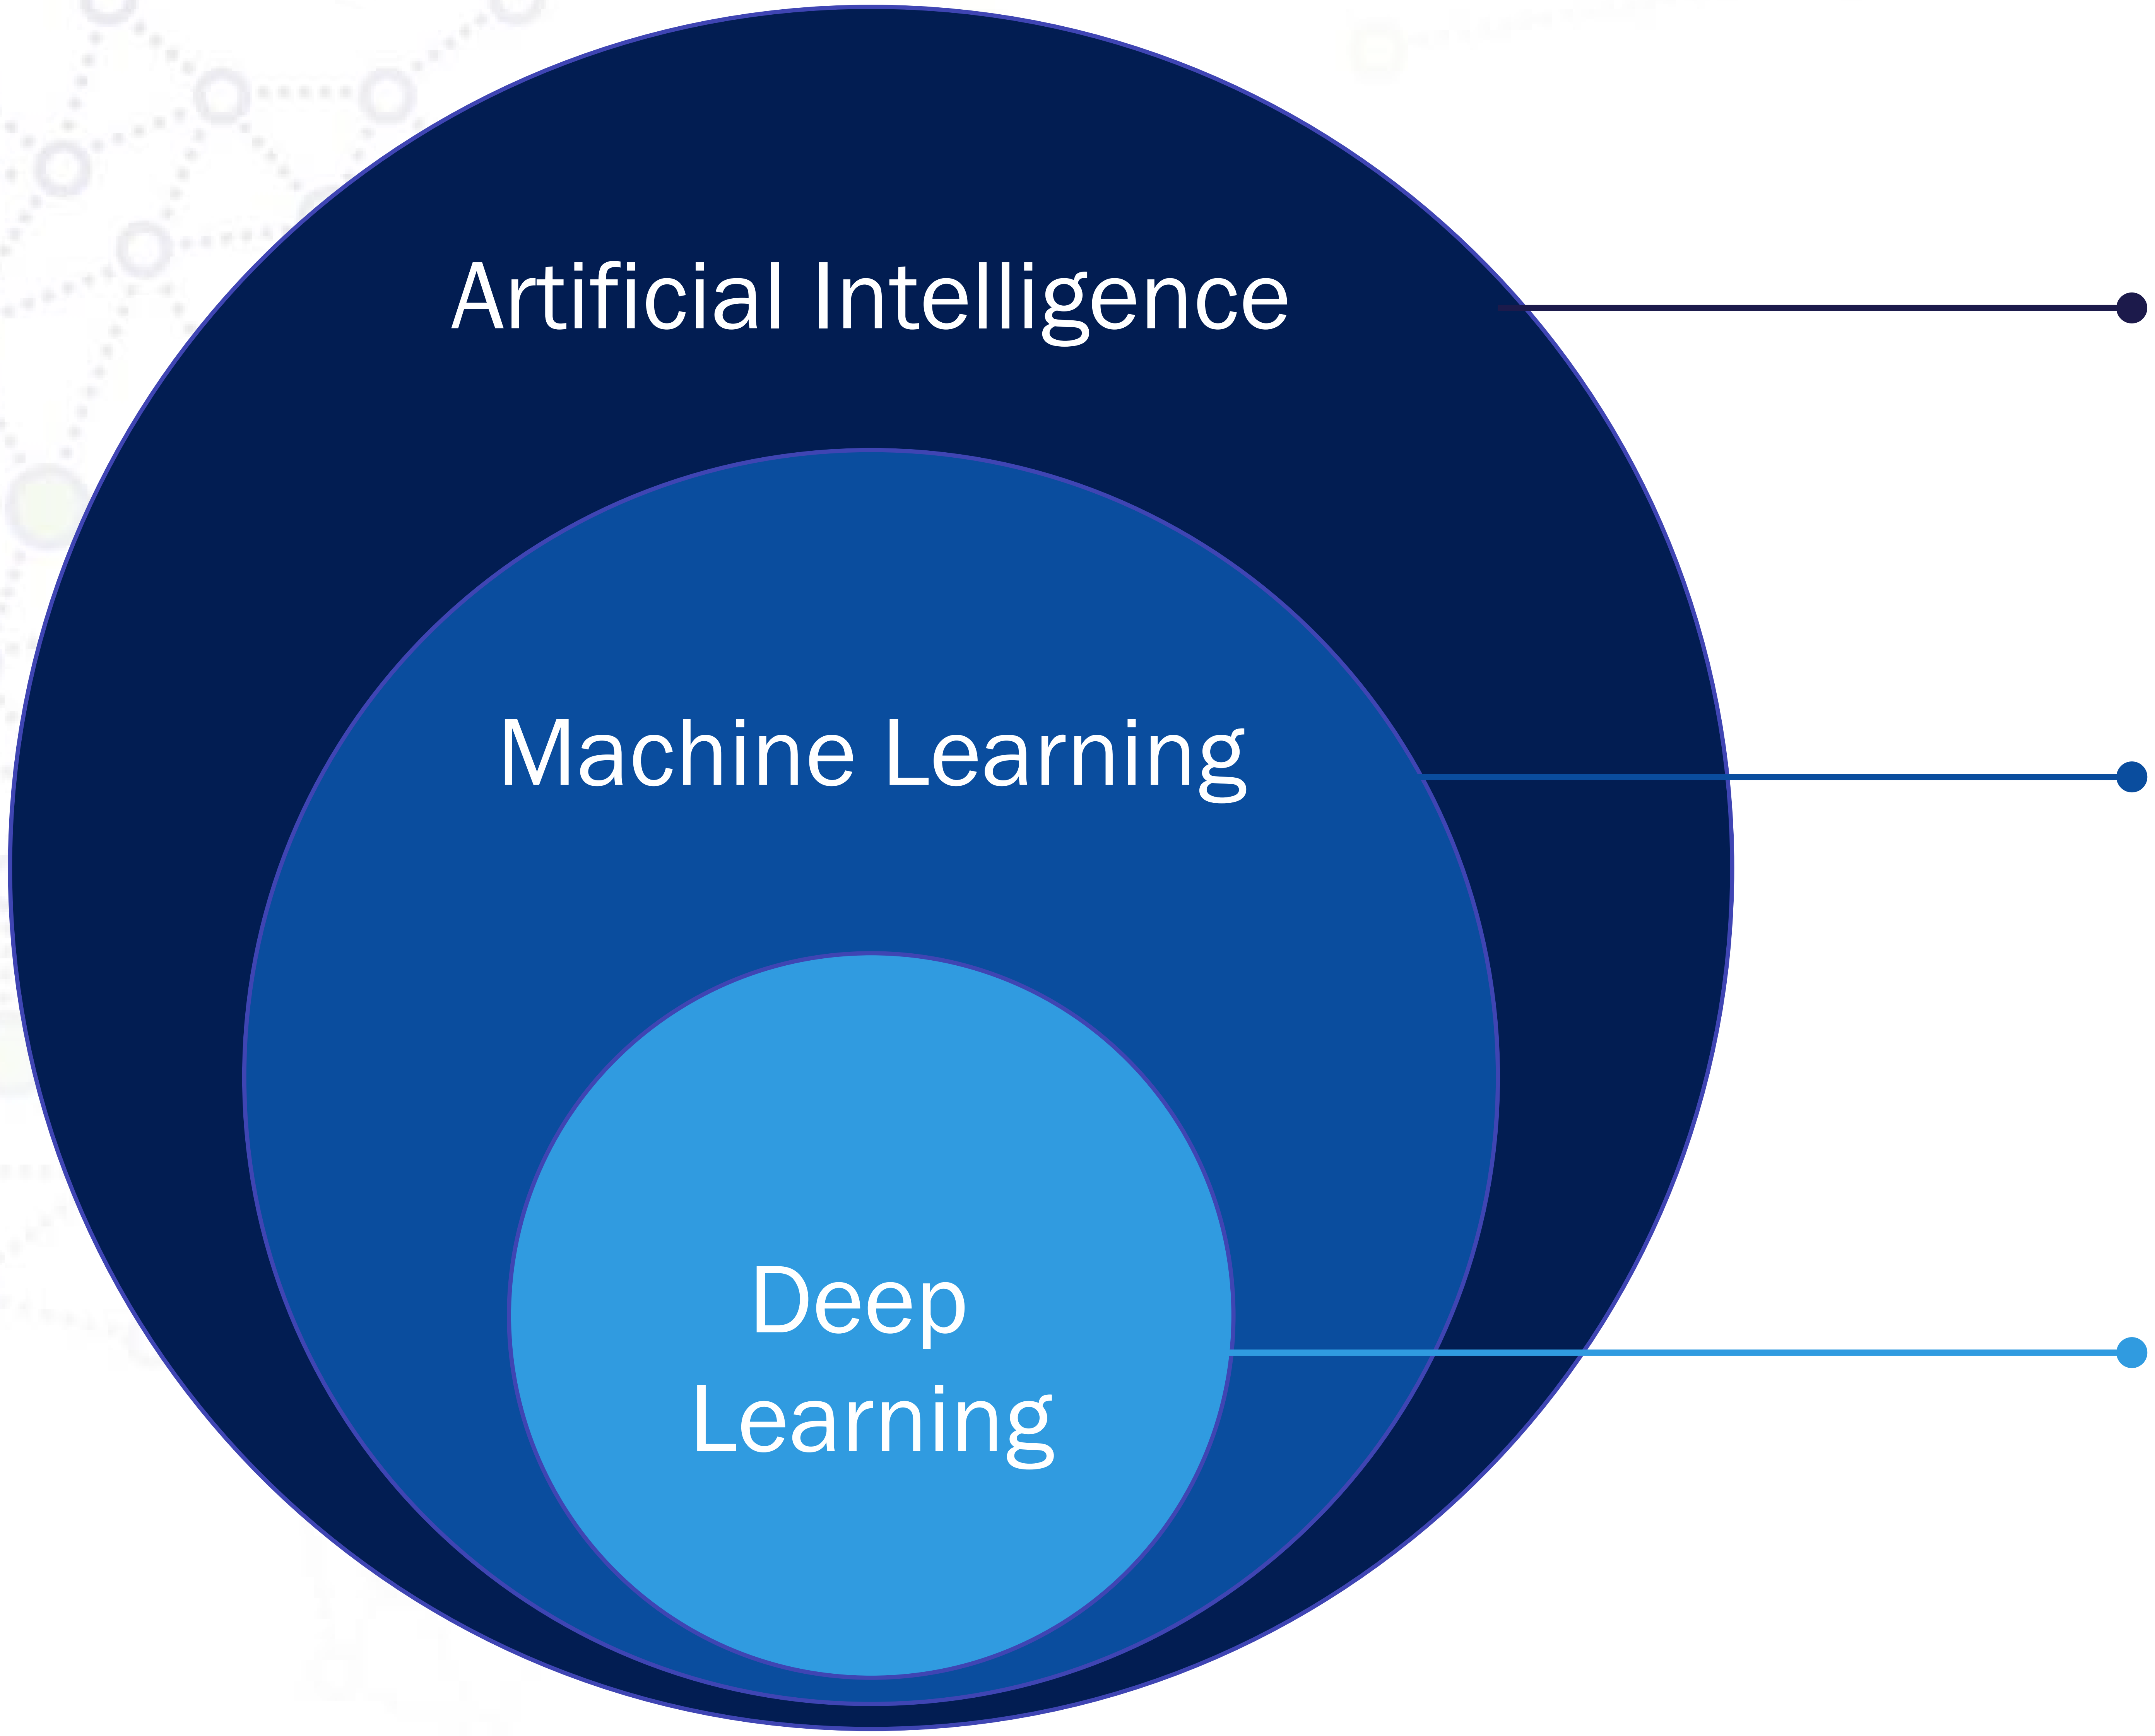

Program that can sense, reason, act, and adapt like humans

Program that can perform actions without explicitly being programmed

Subset of machine learning in which multilayered neural networks learn

EXIT

Services and Use Cases

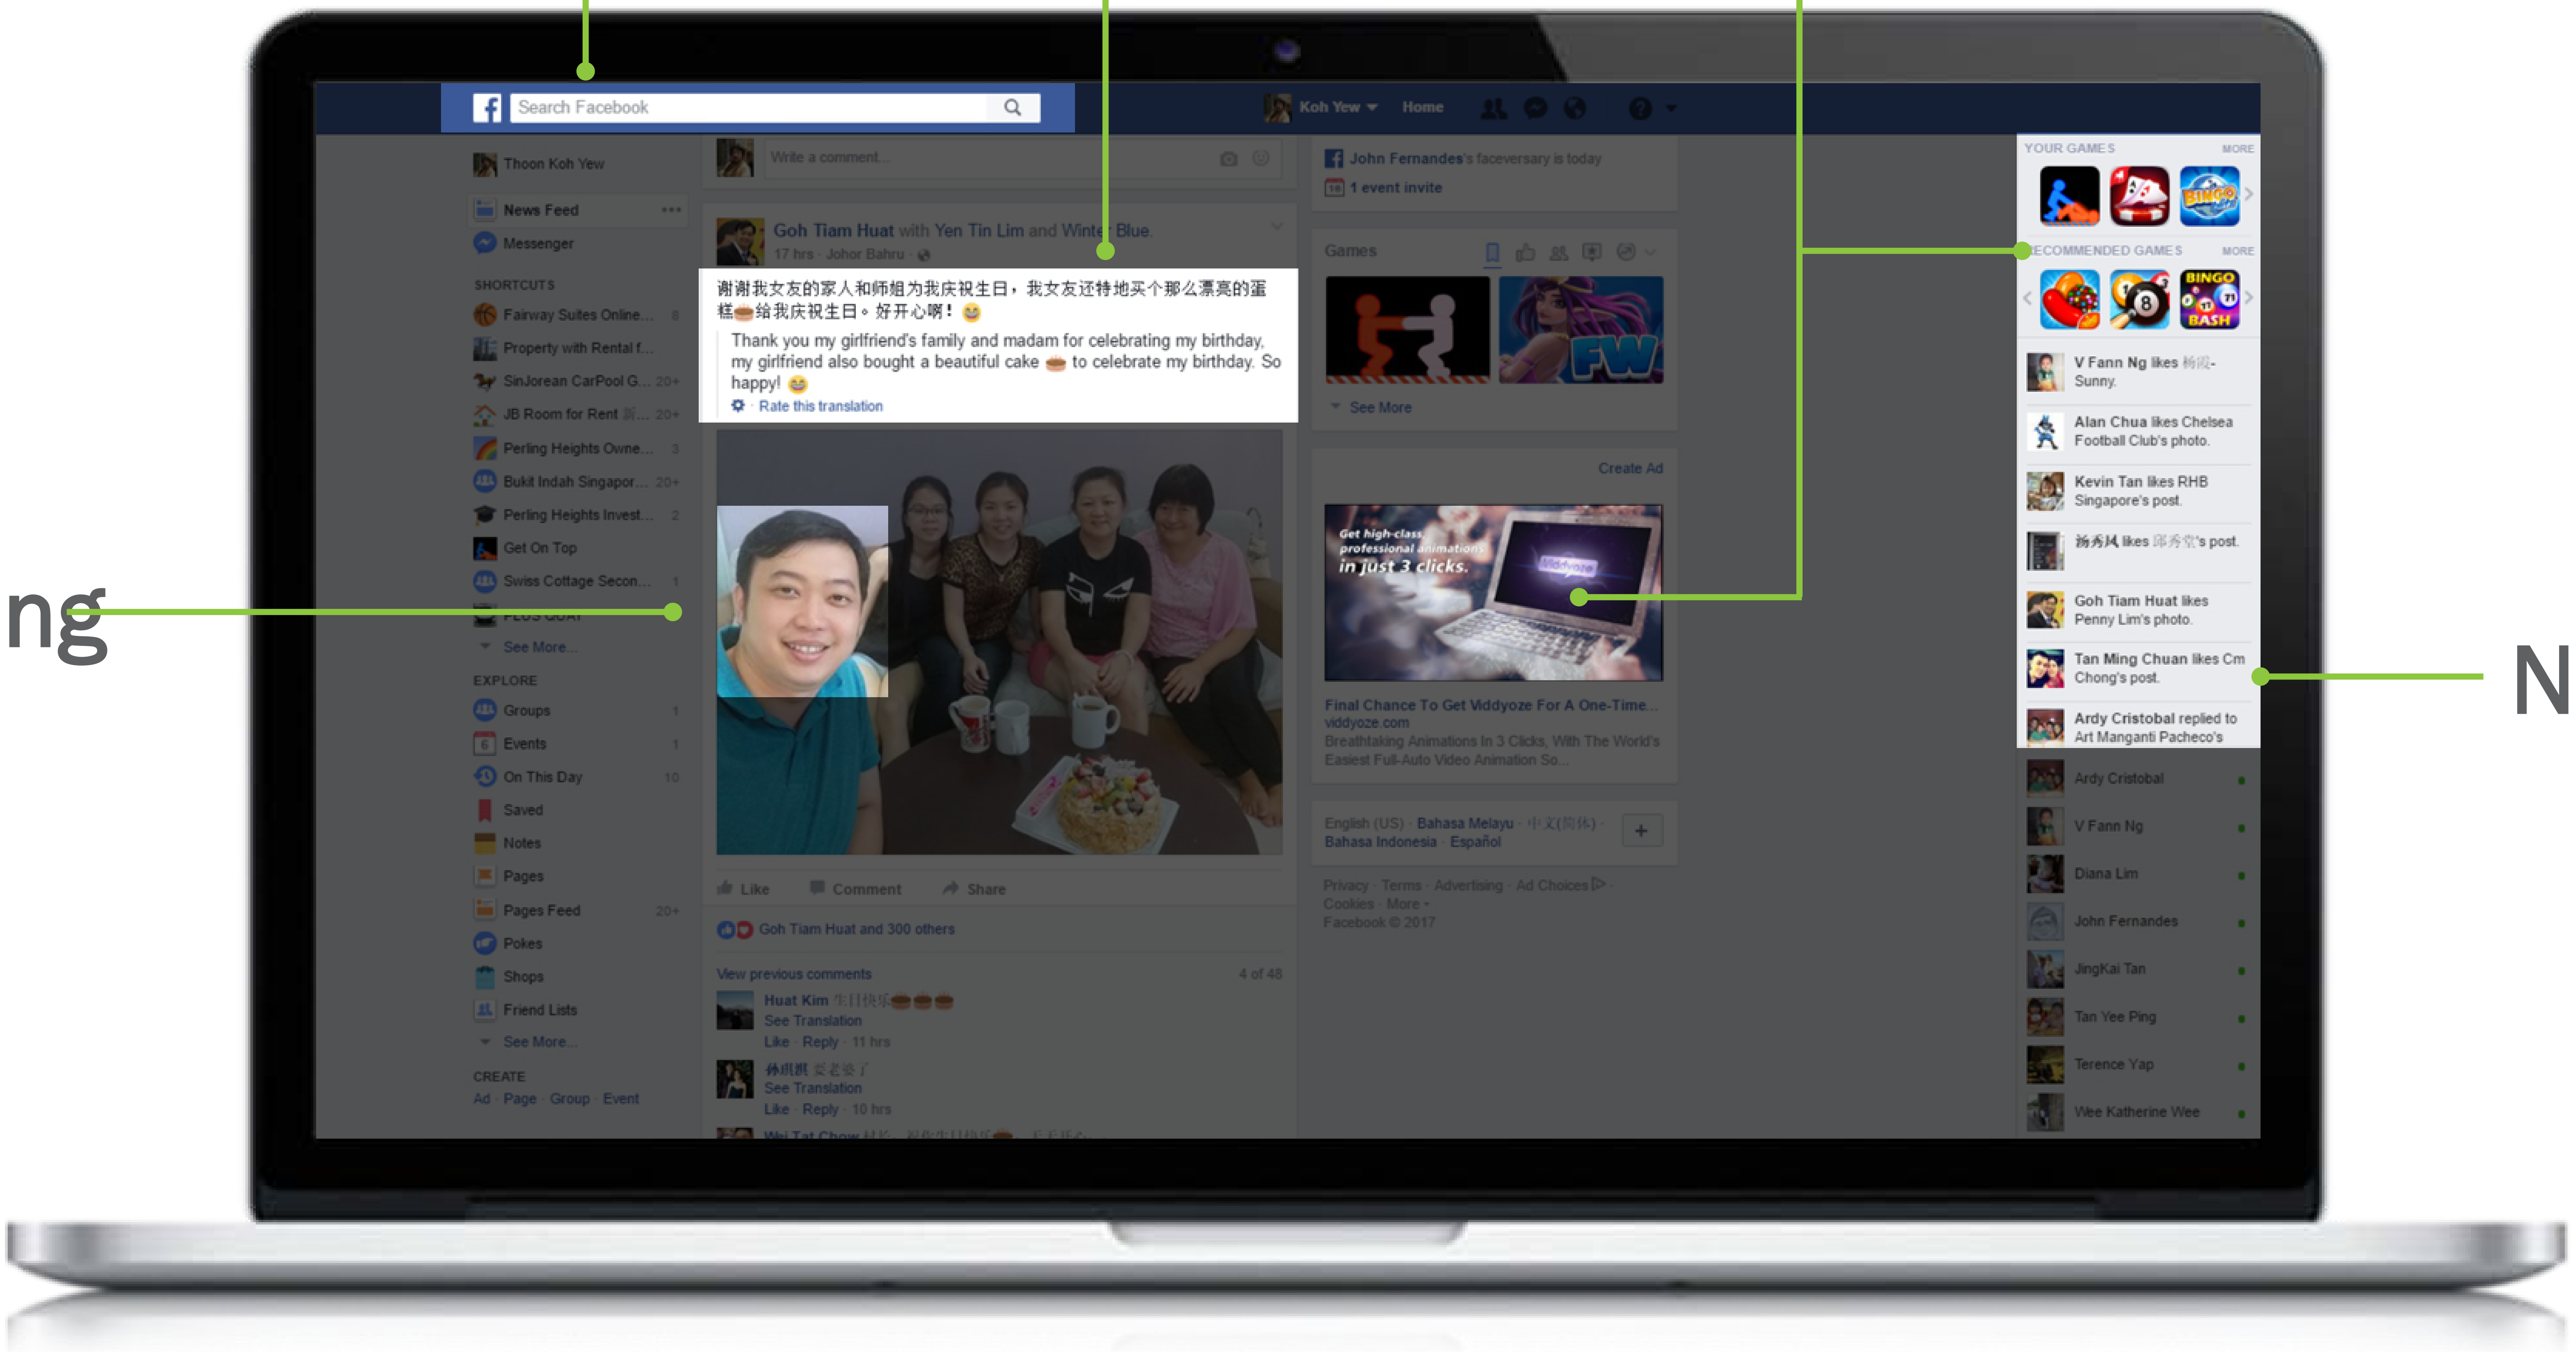

Search

Translation

Ads

Face tagging

News Feed

Goh Tiam Huat with Yen Tin Lim and Winter Blue.
17 hrs · Johor Bahru

谢谢我女友的家人和师姐为我庆祝生日，我女友还特地买个那么漂亮的蛋糕🎂给我庆祝生日。好开心唷！😄

Thank you my girlfriend's family and madam for celebrating my birthday, my girlfriend also bought a beautiful cake 🎂 to celebrate my birthday. So happy! 😄

Rate this translation

facebook

10 YEAR ROADMAP

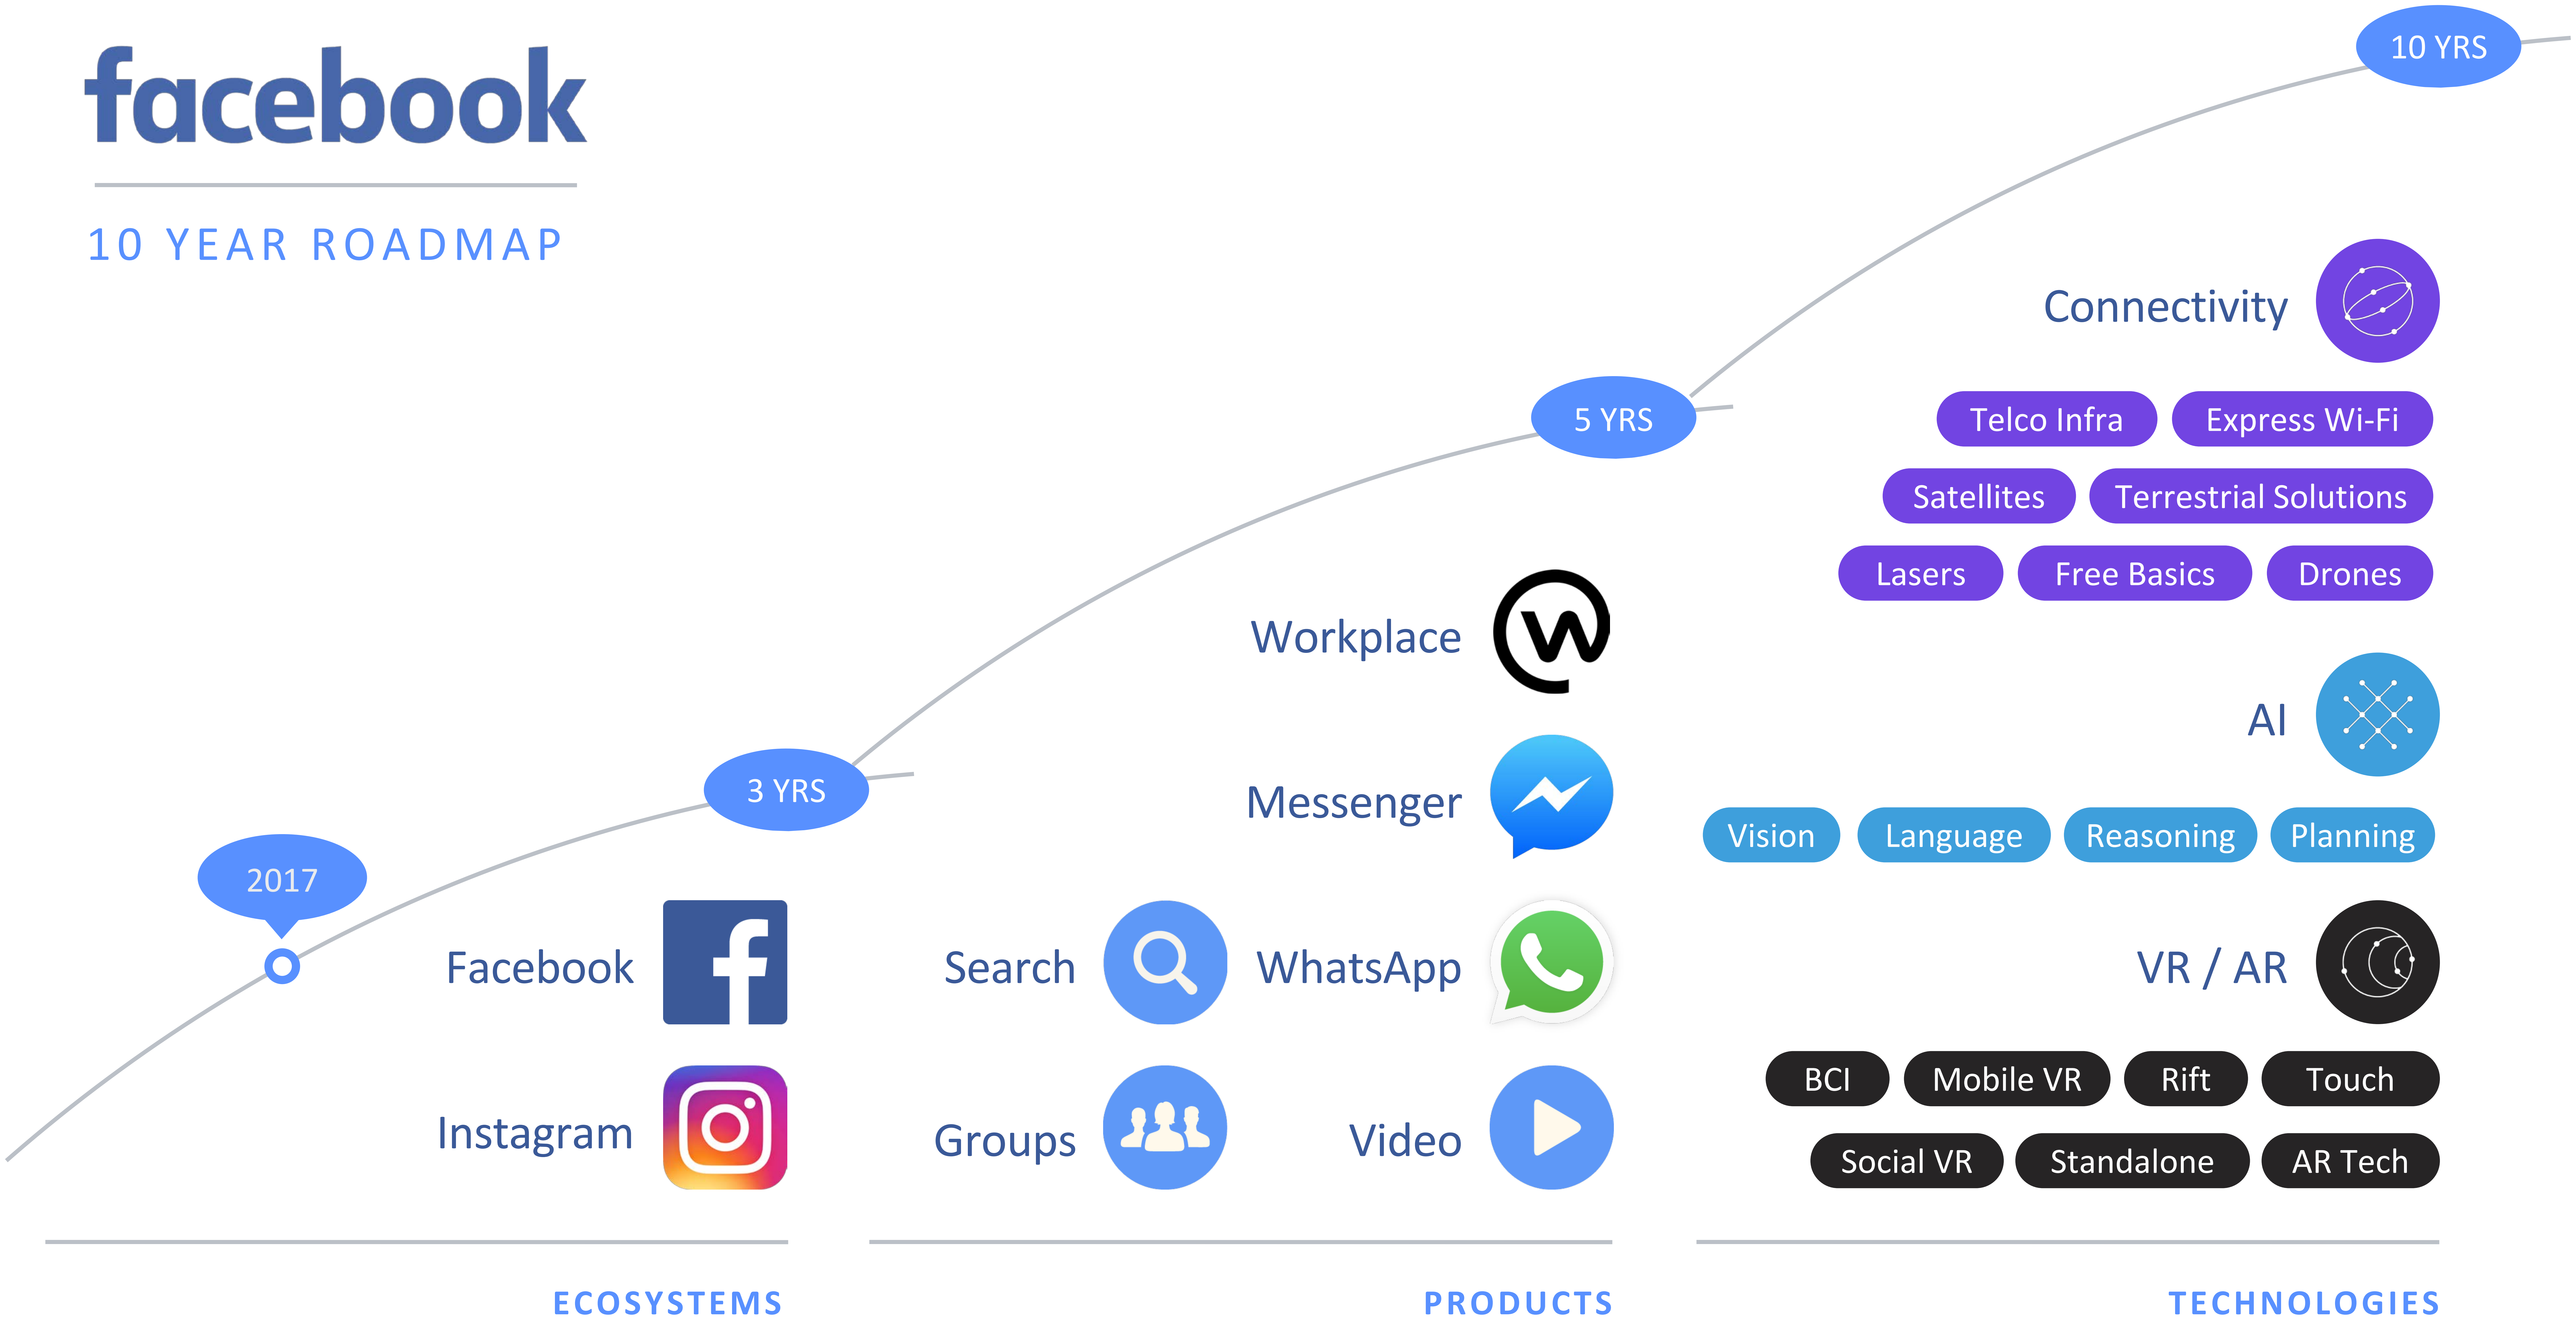

2017

3 YRS

5 YRS

10 YRS

Facebook

Instagram

ECOSYSTEMS

Search

Groups

Workplace

Messenger

WhatsApp

Video

PRODUCTS

Connectivity

Telco Infra

Express Wi-Fi

Satellites

Terrestrial Solutions

Lasers

Free Basics

Drones

AI

Vision

Language

Reasoning

Planning

VR / AR

BCI

Mobile VR

Rift

Touch

Social VR

Standalone

AR Tech

TECHNOLOGIES

Why Now?

Research

Data

Compute

Internal Platforms & Toolkits	Frameworks & Infrastructure
FBLearner Feature Store	Caffe2
FBLearner Flow	PYTORCH
FBLearner Predictor	ONNX

How Are We Leveraging Machine Learning at FB?

Bryce Canyon

Big Basin

Tioga Pass

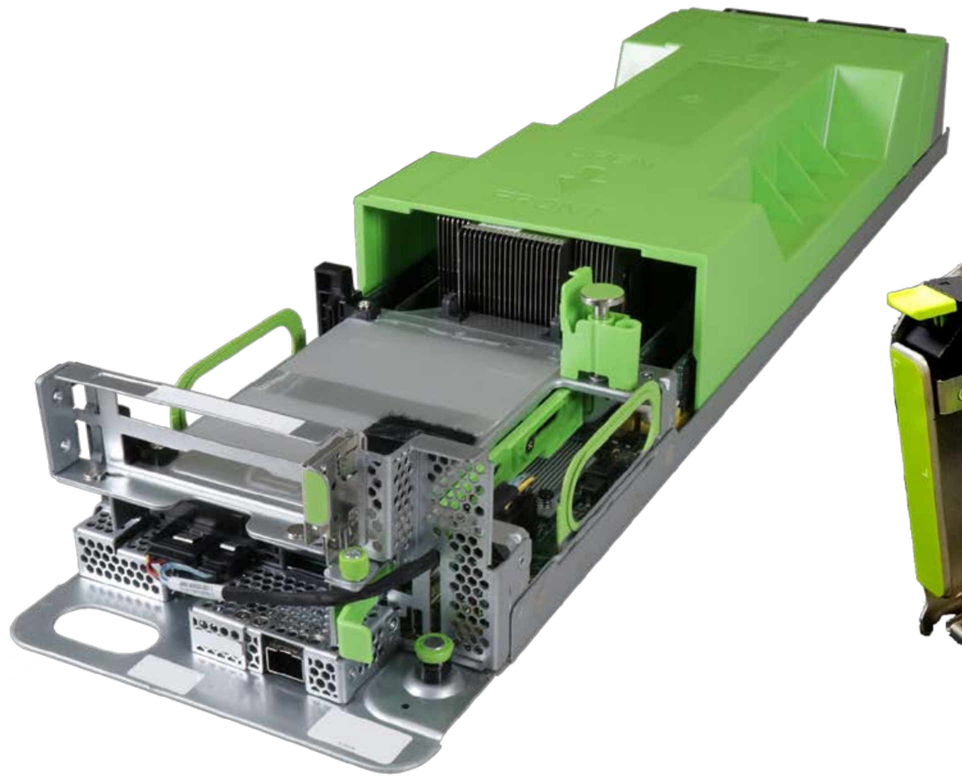

Twin Lakes

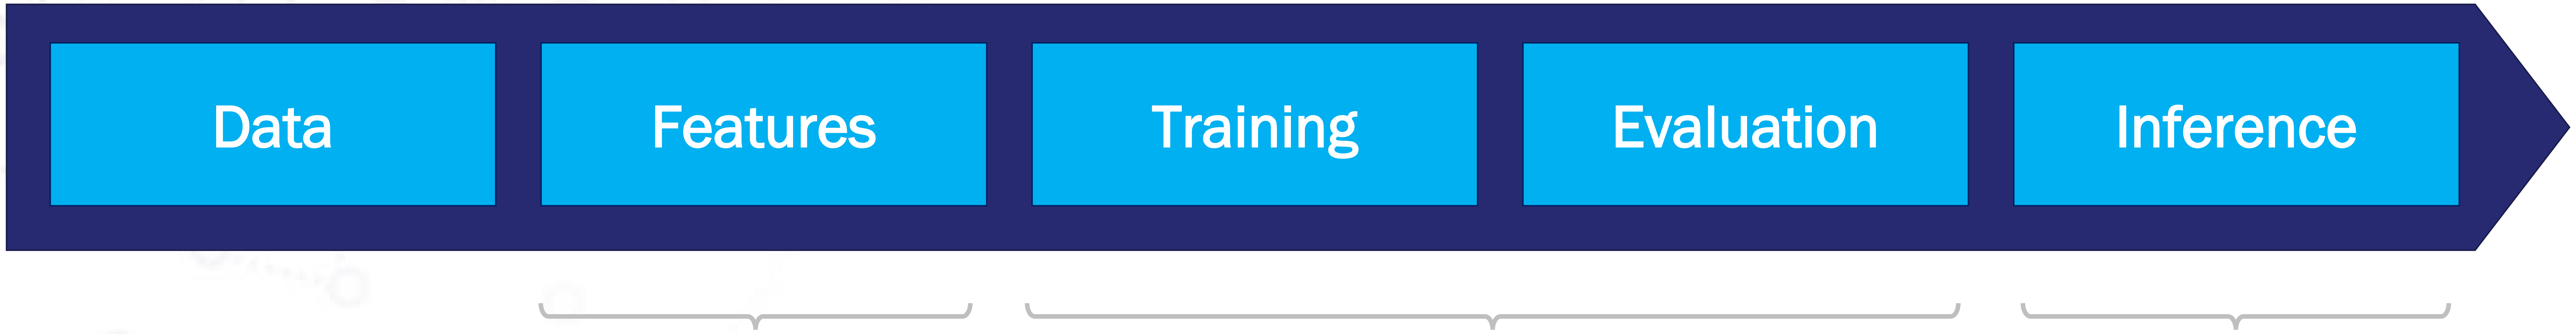

Data

Features

Training

Evaluation

Inference

FBLearner
Feature Store

FBLearner
Flow

FBLearner
Predictor

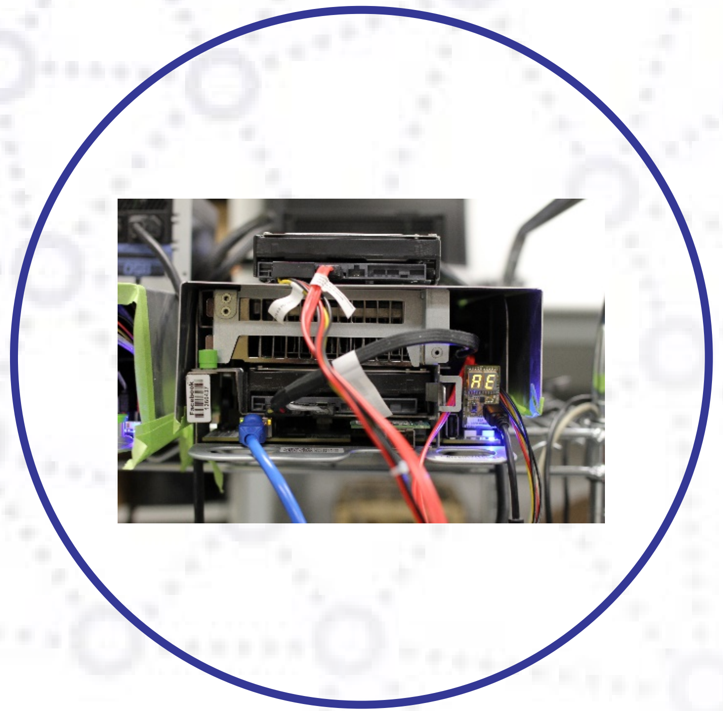

2012

Experimentation

2013

HP SL270s G8

2015

Big Sur

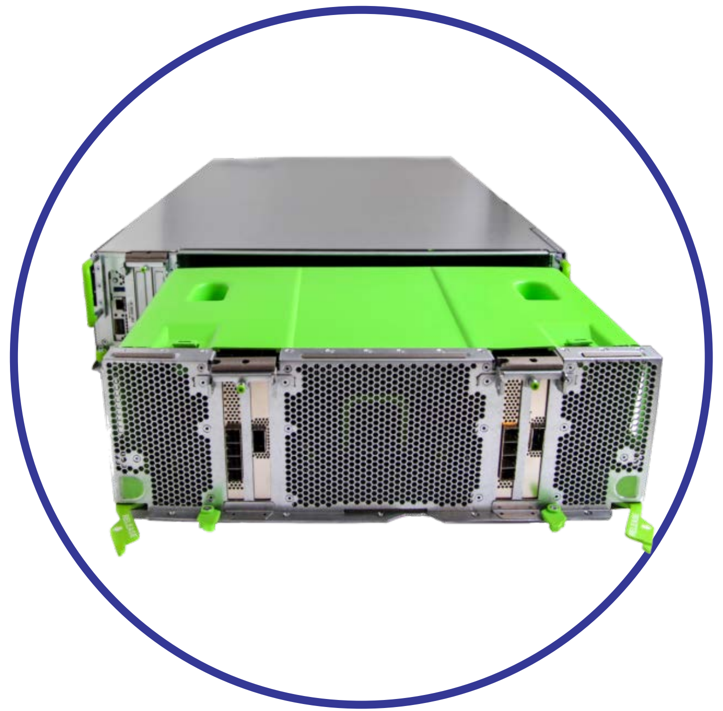

2016

Big Basin Pascal

2018

Announcing Today

Today’s AI Hardware
Infrastructure

Announcing
Big Basin Volta

What’s Next

Big Basin Volta at OCP 2018

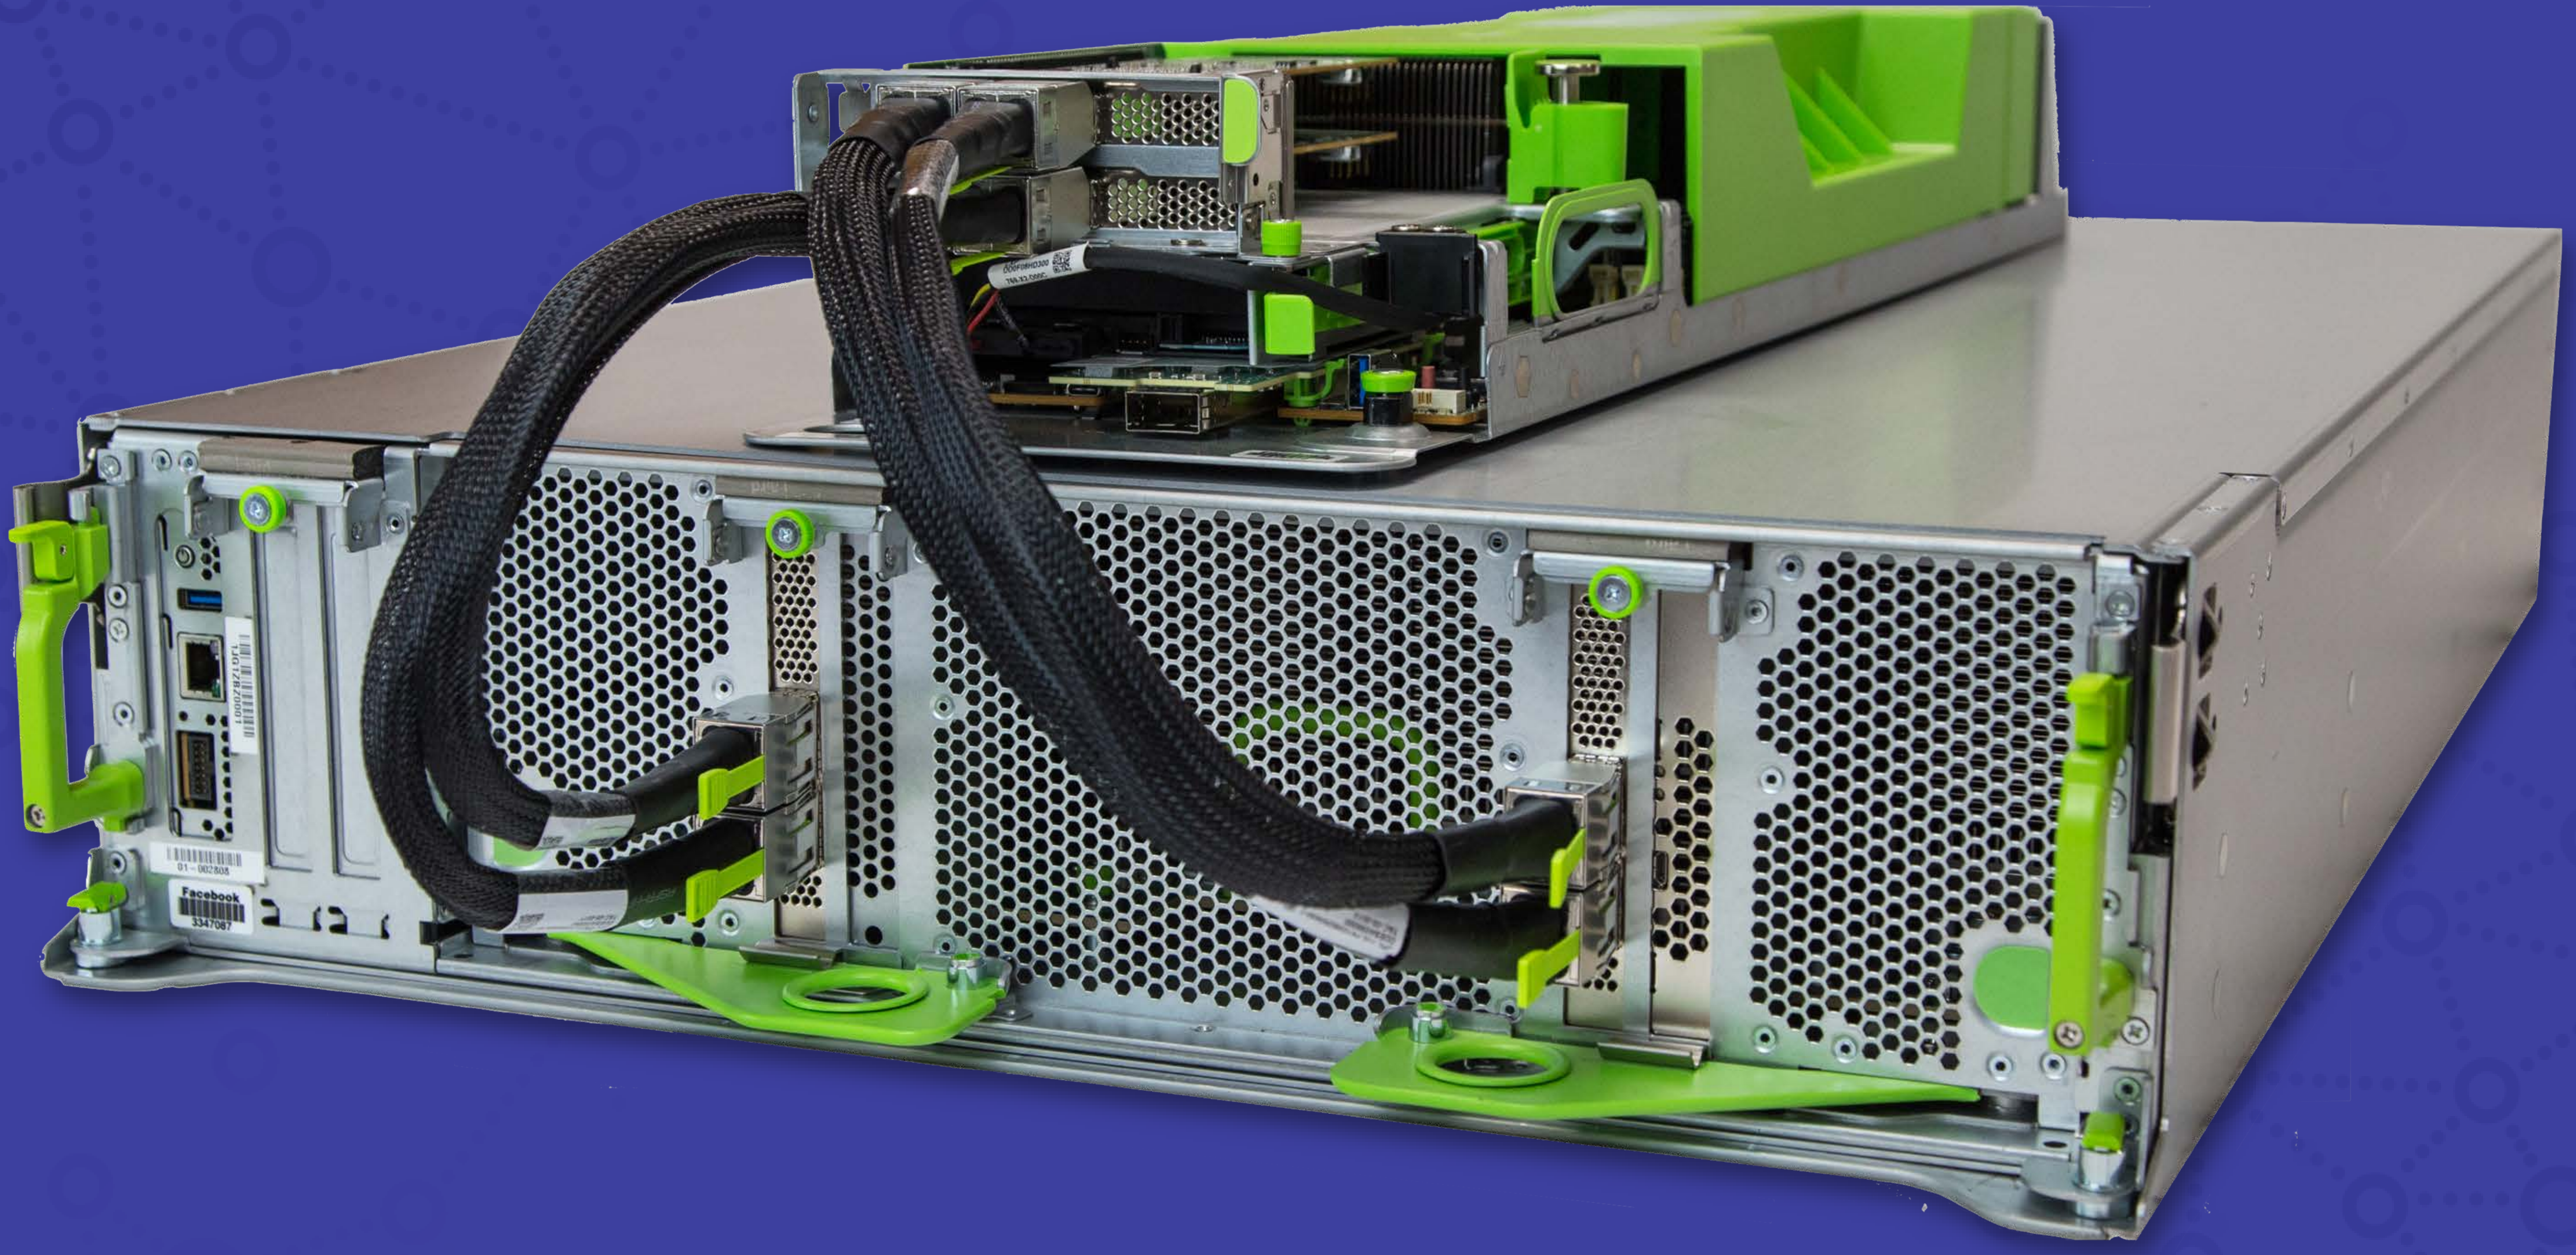

Big Basin Volta at OCP 2018

- Paired with OCP Tioga Pass platform
- Increased GPU performance
- Improved airflow and serviceability

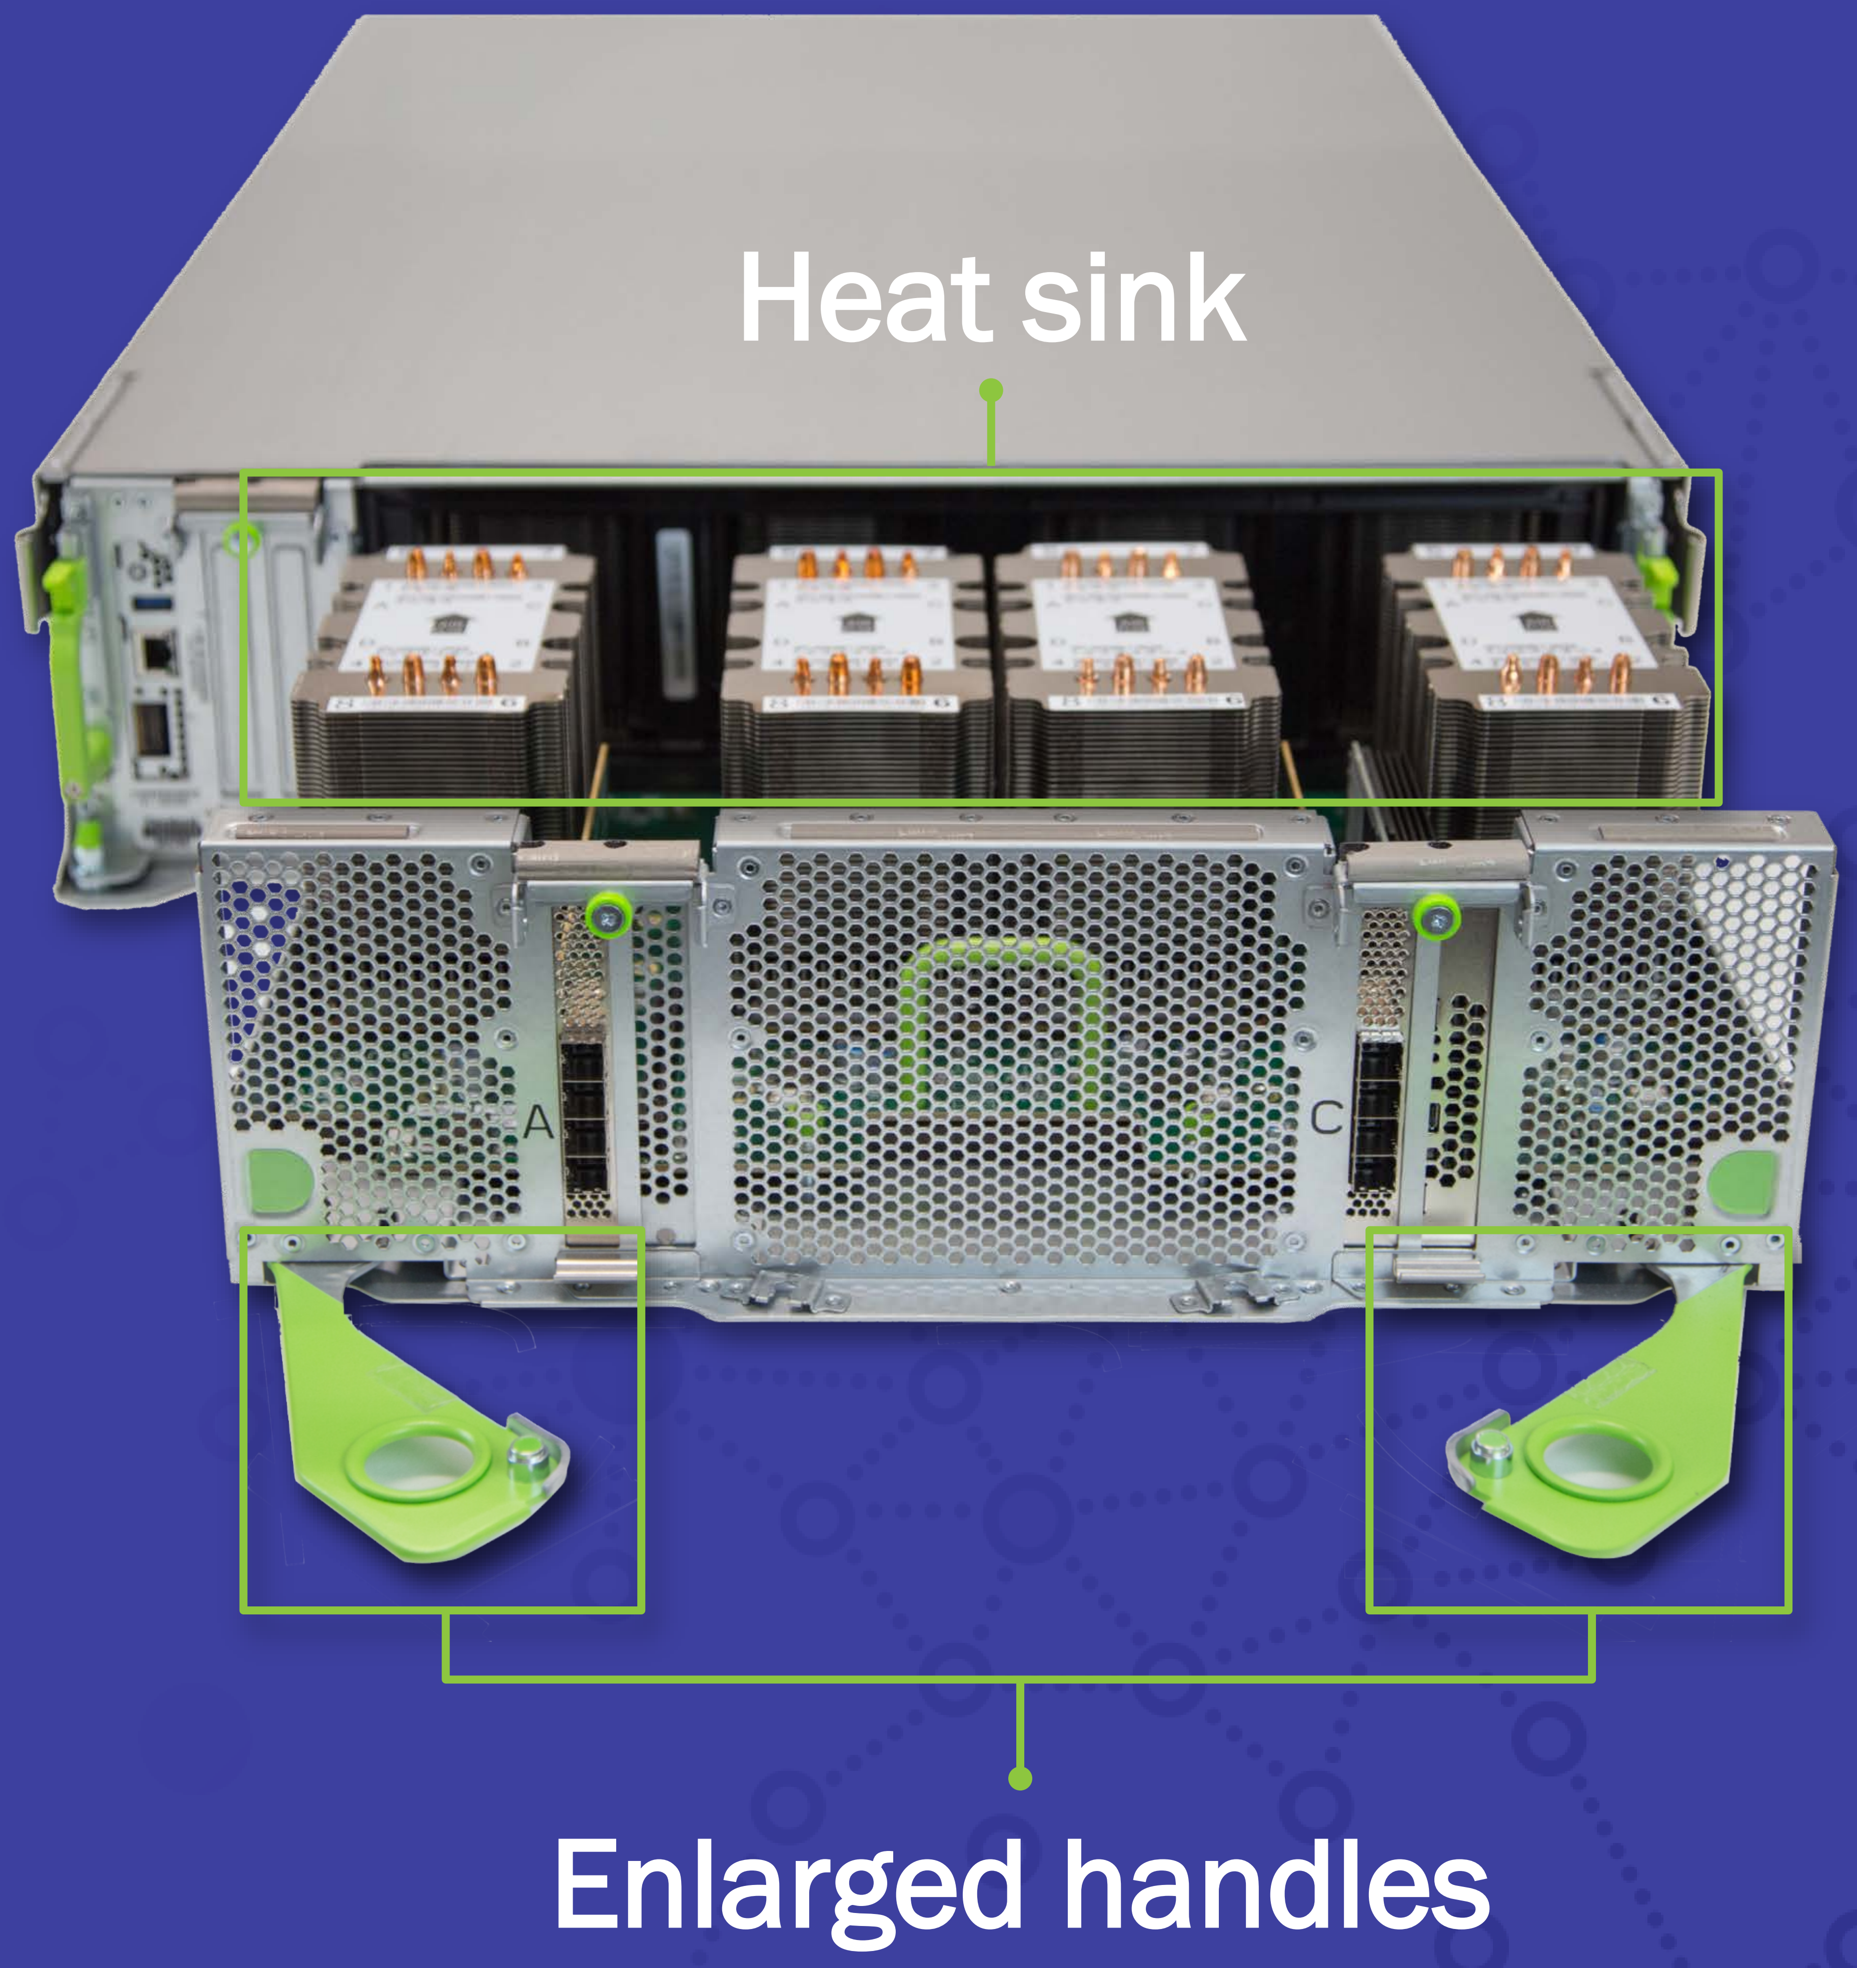

Computer Vision at FB

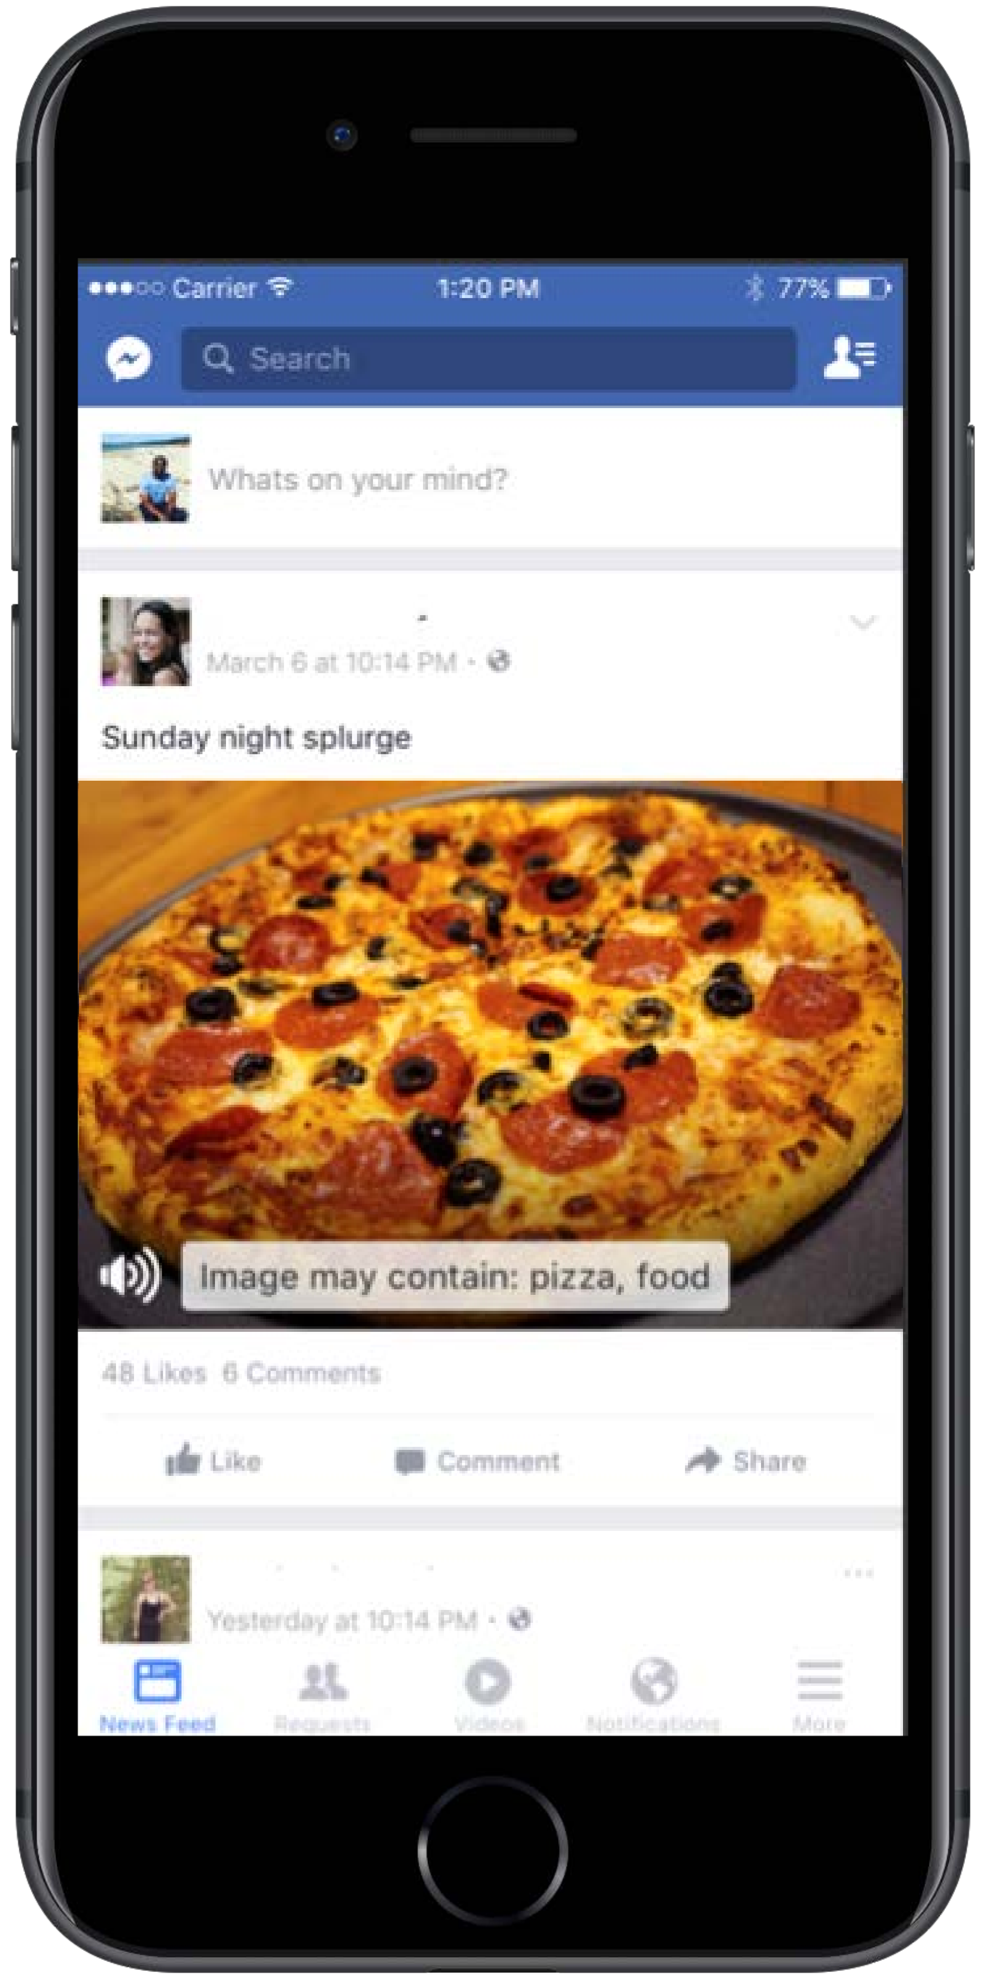

Sunday night splurge

Image may contain: pizza, food

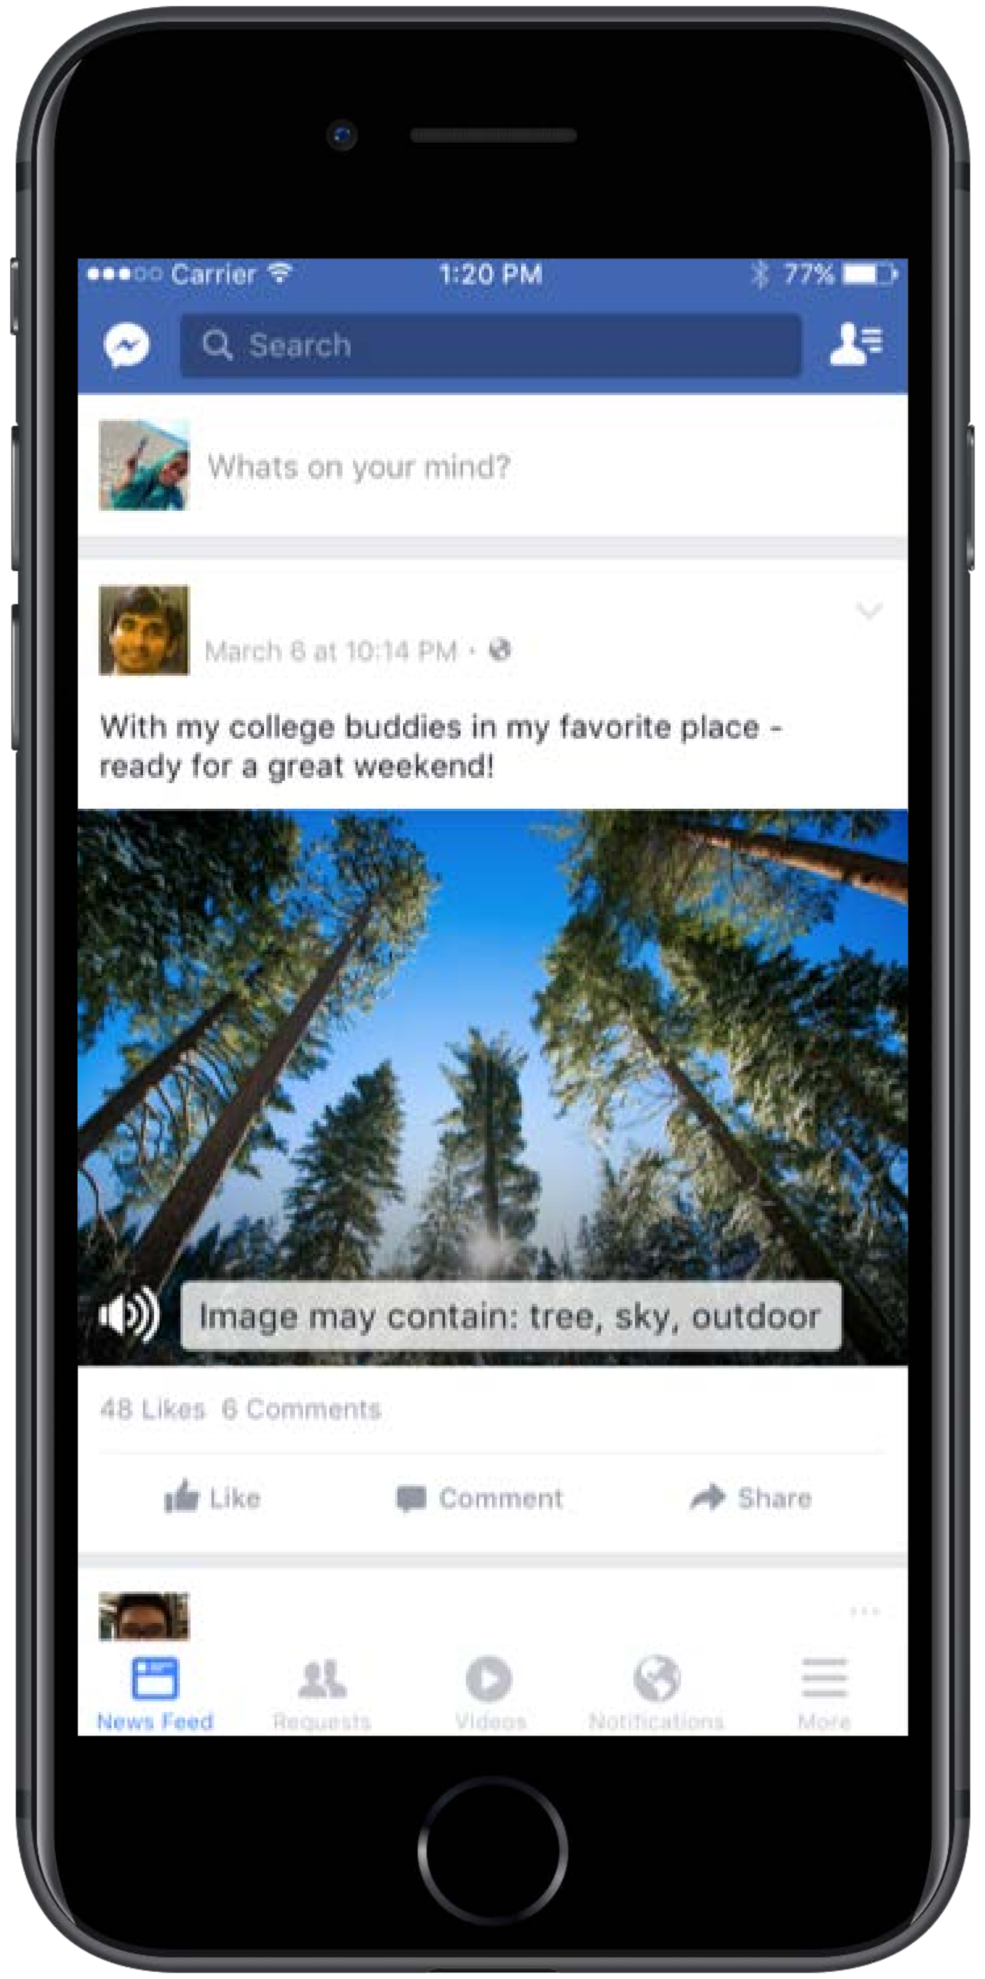

With my college buddies in my favorite place - ready for a great weekend!

Image may contain: tree, sky, outdoor

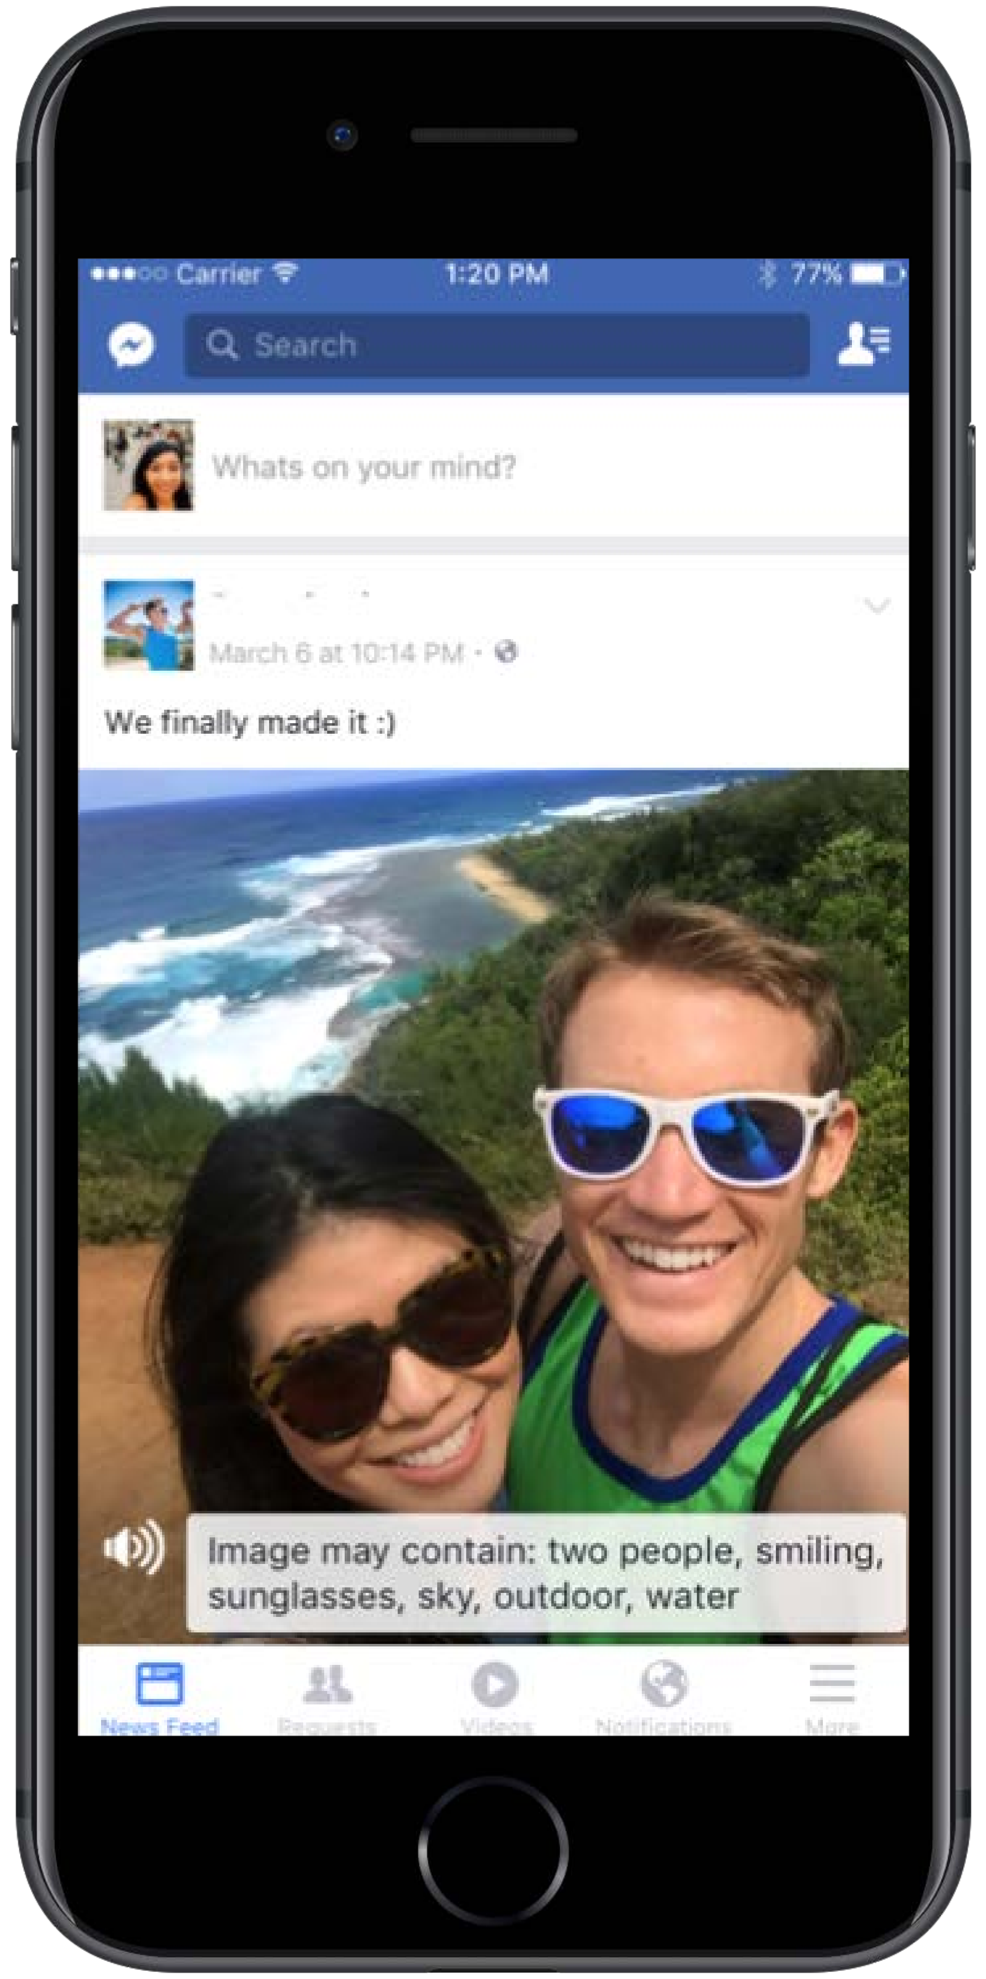

We finally made it :)

Image may contain: two people, smiling, sunglasses, sky, outdoor, water

Computer Vision Performance

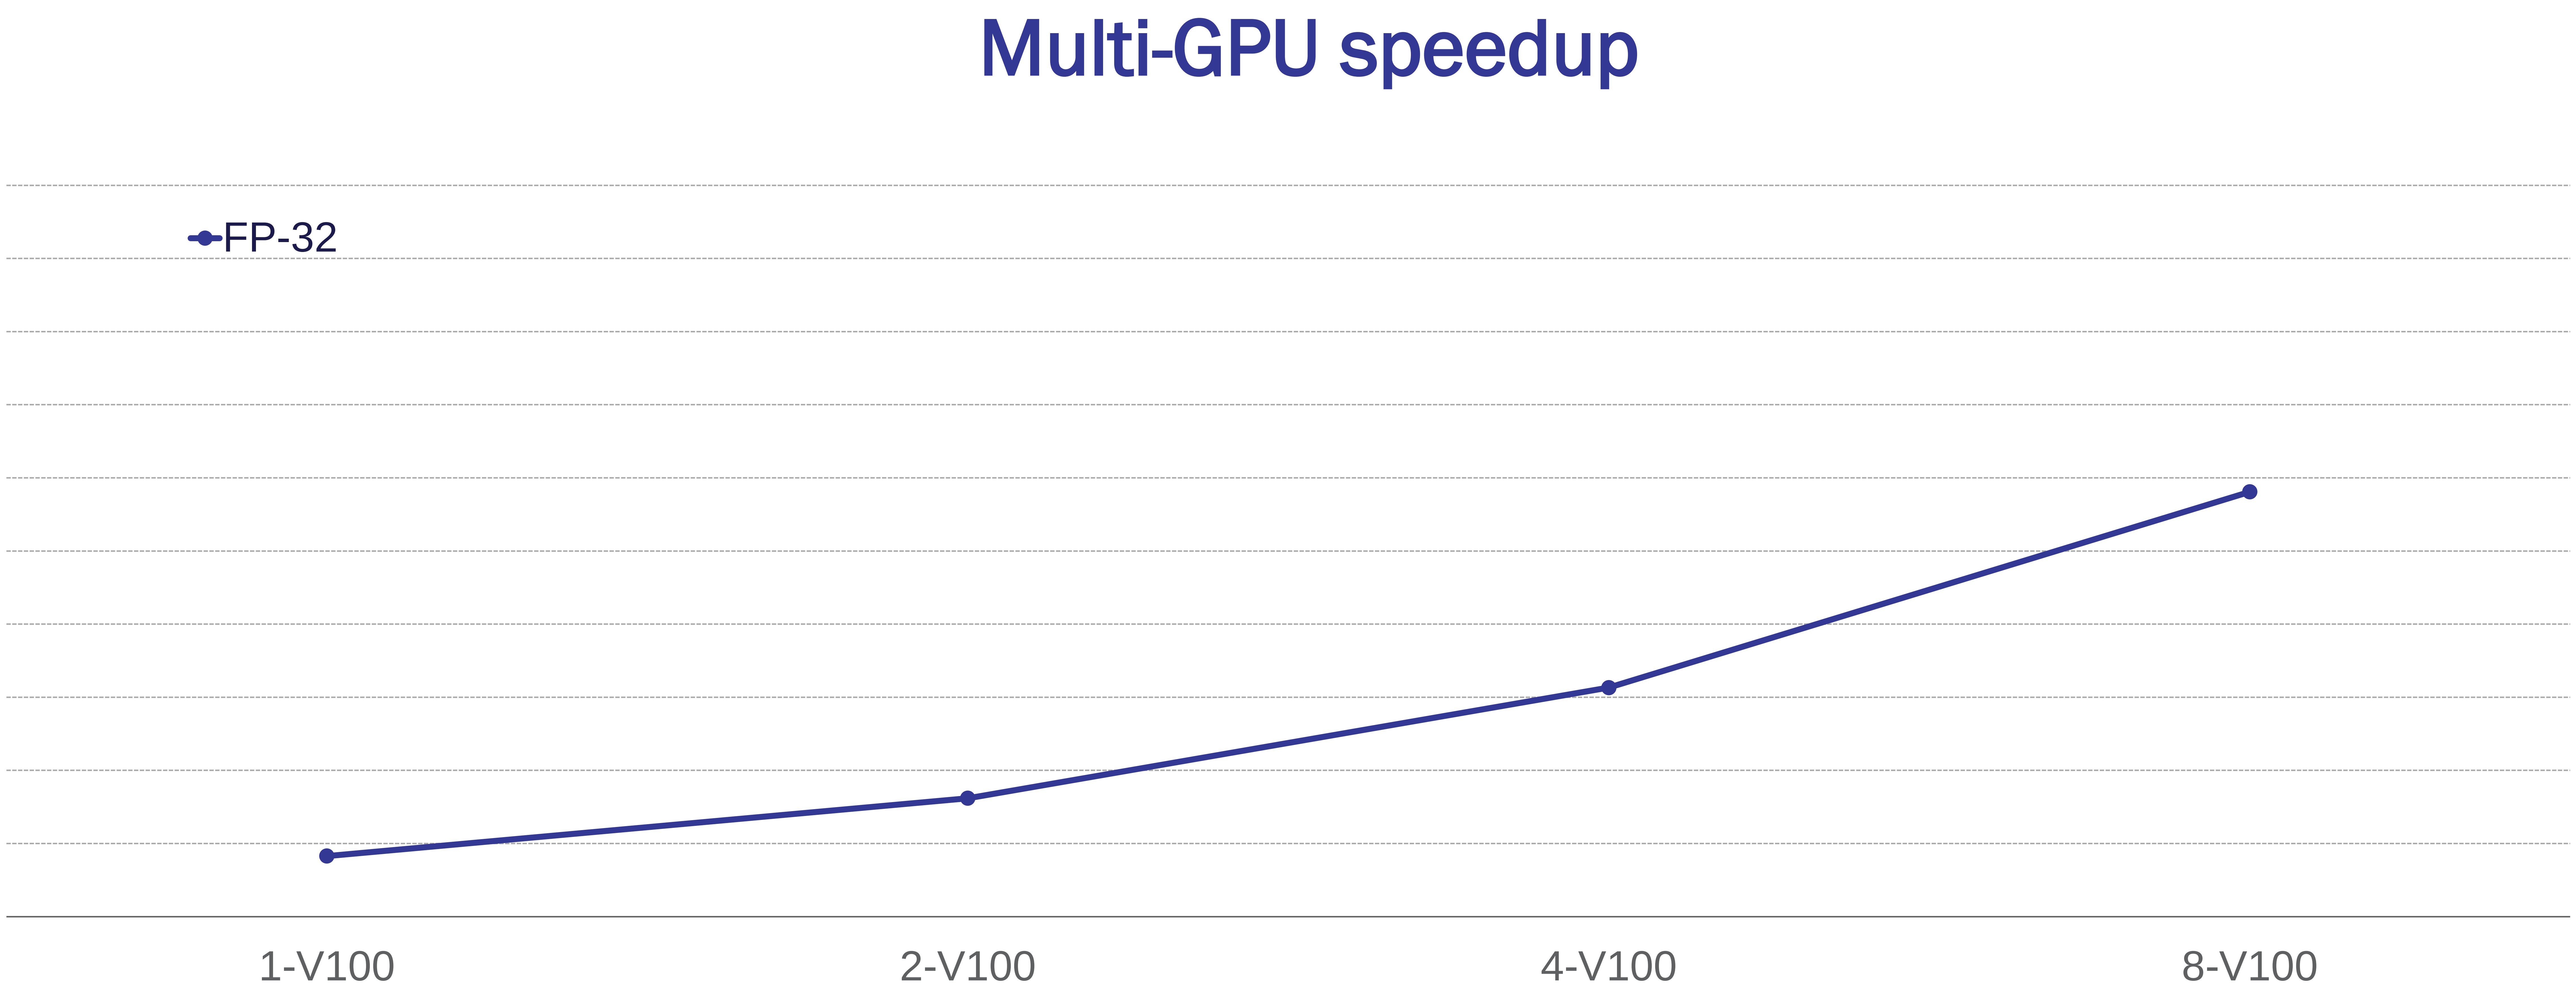

Computer Vision Performance

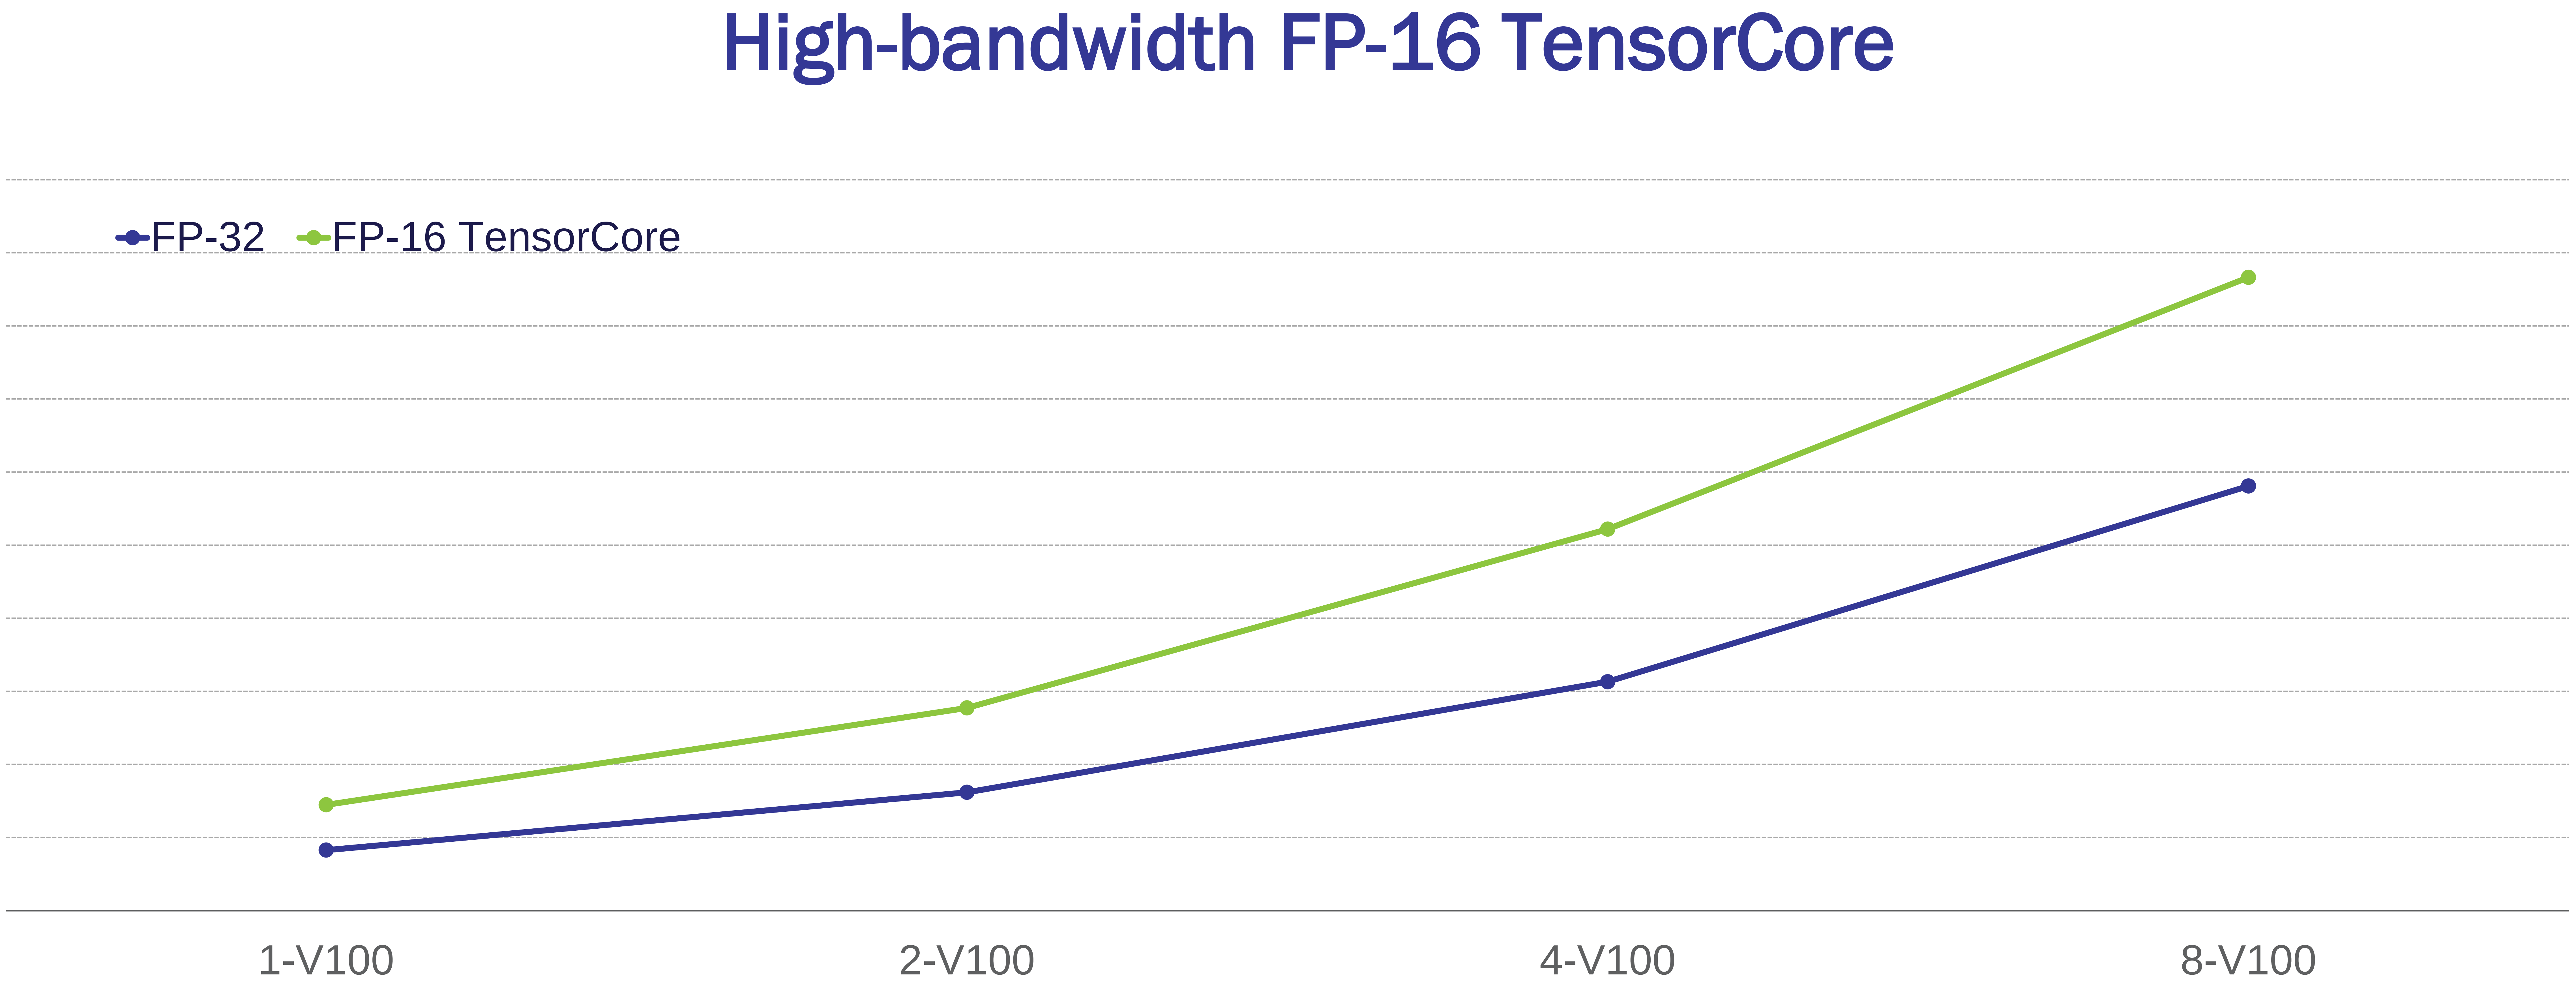

Machine Translation

Better Translation Quality

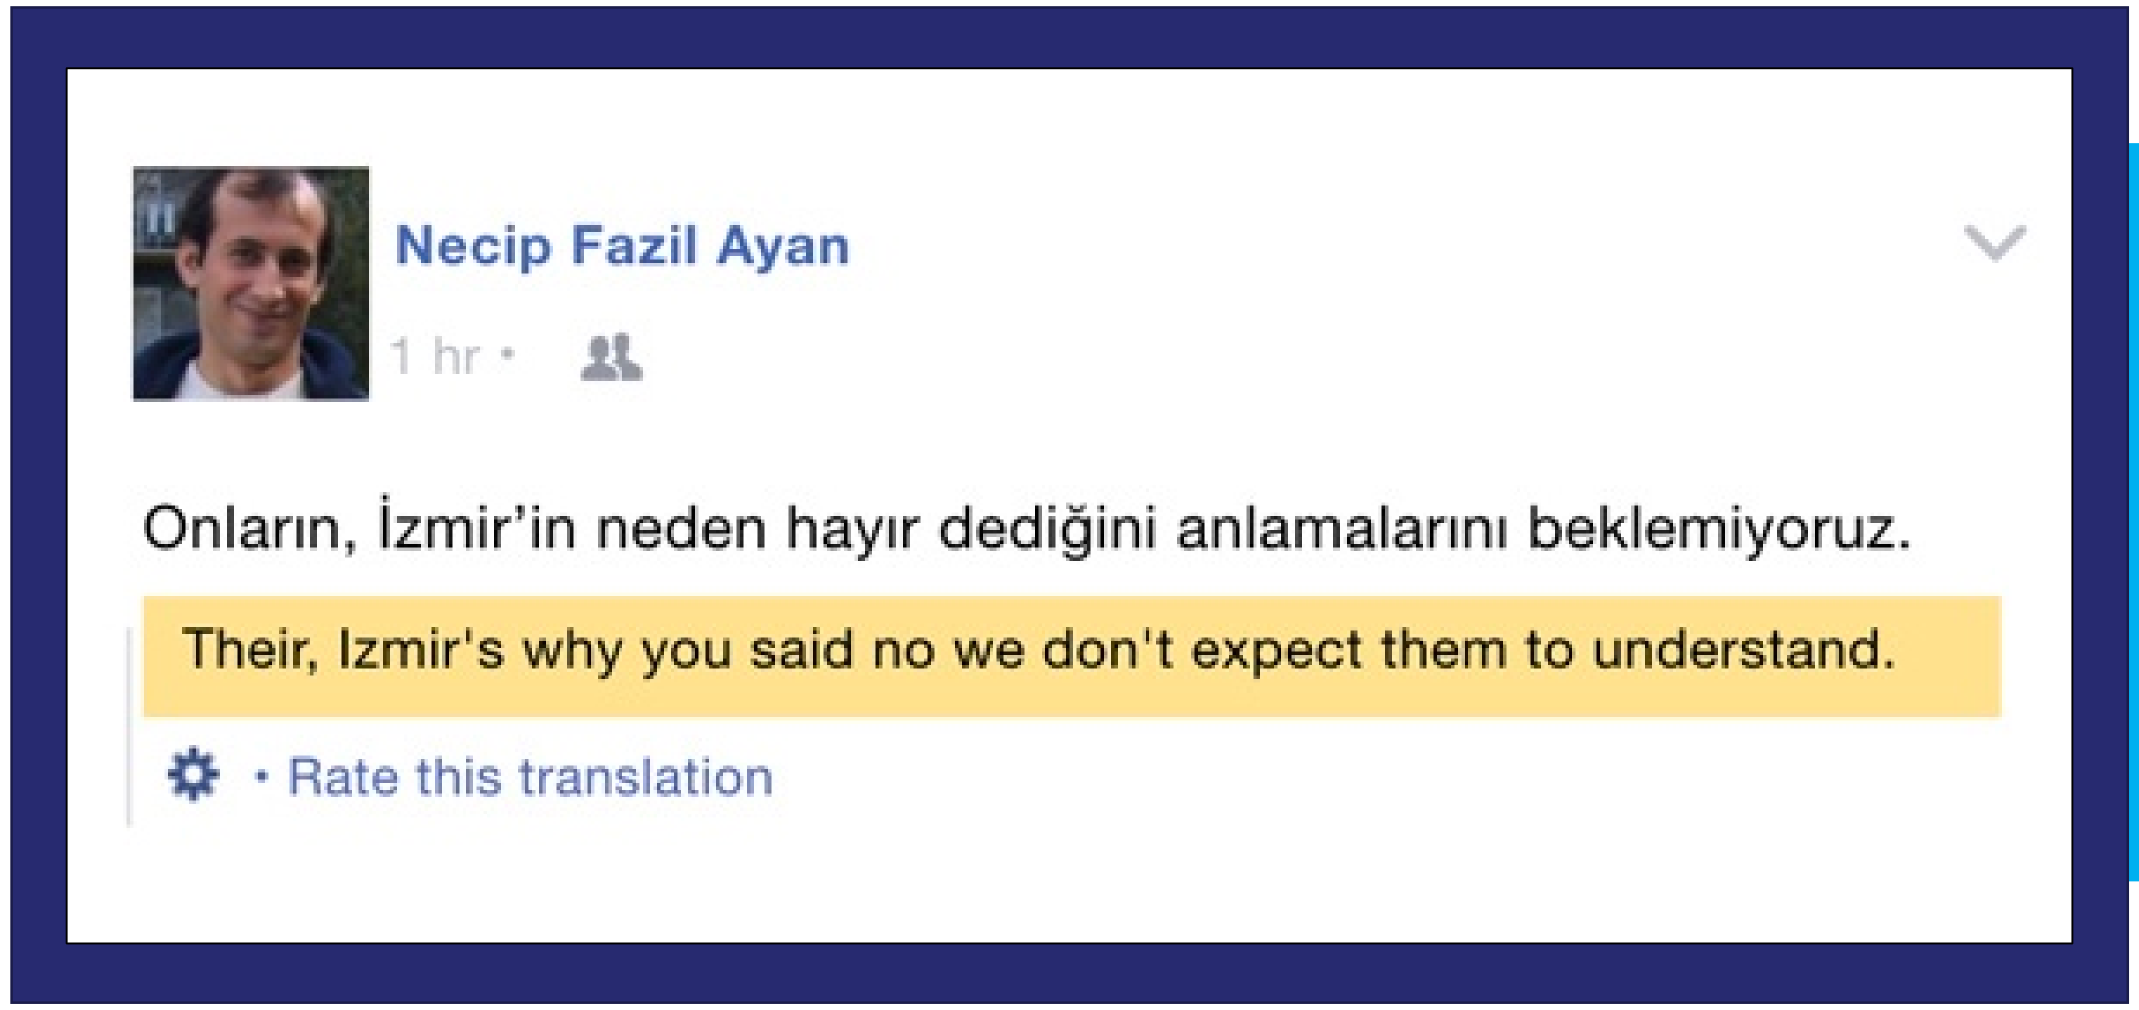

Necip Fazil Ayan

1 hr ·

Onların, İzmir’in neden hayır dediğini anlamalarını beklemiyoruz.

Their, Izmir's why you said no we don't expect them to understand.

· Rate this translation

Phrase-based statistical approach

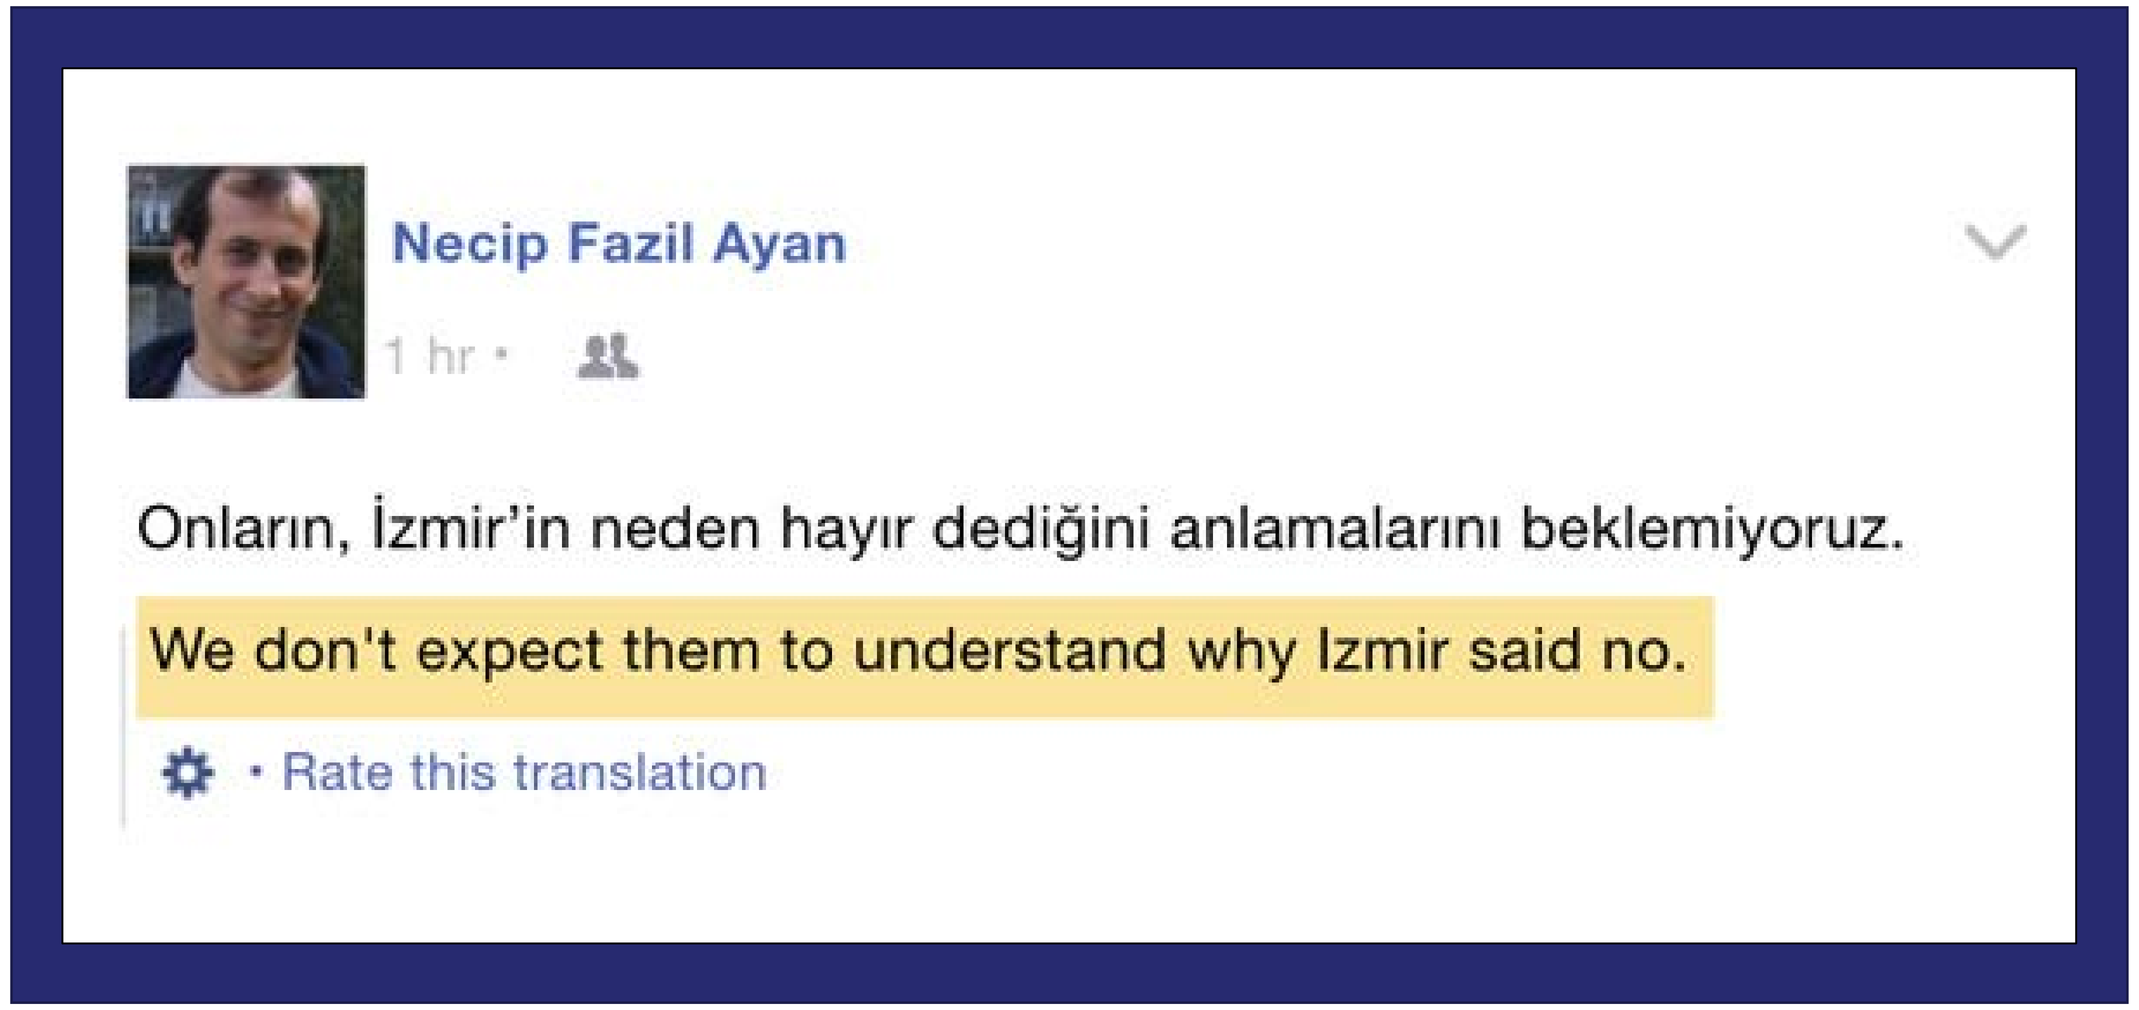

Necip Fazil Ayan

1 hr ·

Onların, İzmir’in neden hayır dediğini anlamalarını beklemiyoruz.

We don't expect them to understand why Izmir said no.

· Rate this translation

Neural network approach

Today’s AI Hardware Infrastructure

Announcing Big Basin Volta

What’s Next

Consolidated Hardware Design

Max Operating Efficiency

Disaggregated Design

Vendor Agnostic Platform

THIS JOURNEY 1% FINISHED

Please join us Wednesday at 11:30am for Big Basin Volta workshop!

OCP
SUMMIT