



The bridge to possible

# Cisco Silicon for AI

Peter Jones

*Distinguished Engineer*

Dave Zacks

*Distinguished Engineer*

BRKARC-2095

CISCO *Live!*

#Cisco



**#HighBitRate**

# Cisco Webex App

## Questions?

Use Cisco Webex App to chat with the speaker after the session

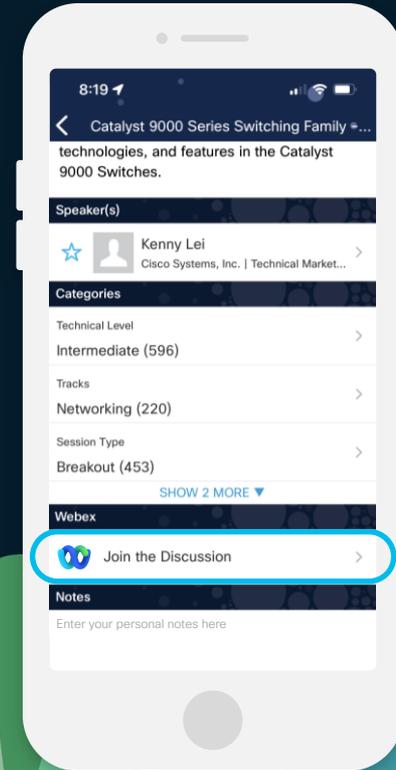
## How

- 1 Find this session in the Cisco Live Mobile App
- 2 Click “Join the Discussion”
- 3 Install the Webex App or go directly to the Webex space
- 4 Enter messages/questions in the Webex space

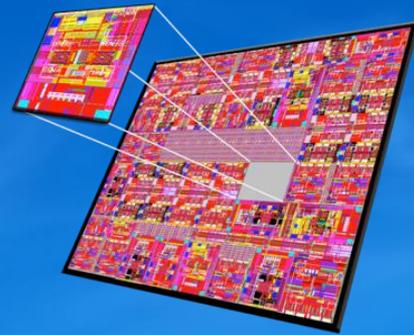
Webex spaces will be moderated by the speaker until June 7, 2024.

**CISCO** *Live!*

<https://ciscolive.ciscoevents.com/ciscolivebot/#BRKARC-2095>



# By Way of Introduction ...



I am a **Distinguished Engineer** in the Network Transformation team, and have been with Cisco since 1999.

I work primarily with large, high-performance Enterprise networks, designs, and systems. I have over 38 years of experience with designing, implementing, and supporting solutions with many diverse technologies.

I have a strong background in, and focus on, customer requirements, and integrating these into the products and solutions Cisco builds.

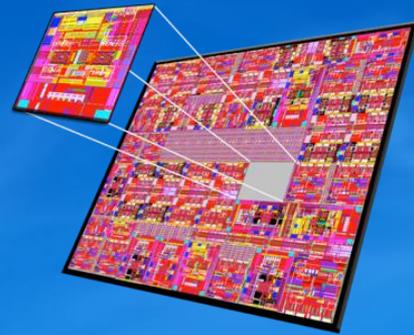
I have a special interest in **Flexible Hardware, Fabrics, Assurance and ML/AI**.

**Dave Zacks**  
Distinguished Engineer

Email: [zacks@cisco.com](mailto:zacks@cisco.com)  
Twitter/x: [@DaveZacks](https://twitter.com/DaveZacks)  
LinkedIn: [In/dave-zacks-43677474/](https://www.linkedin.com/in/dave-zacks-43677474/)



# By Way of Introduction ...



I am a **Distinguished Engineer** in the Cisco Networking Hardware team, and have been at Cisco since 2005.

I work on system architecture and standards strategy across the portfolio. I was a key figure in the development of the UADP switching ASIC architecture and the Catalyst switches that use it.

I work in defining and promoting new Ethernet standards through participating in IEEE 802.3 and as Ethernet Alliance Chairman.

I am passionate about **Network Evolution and Adoptable Technology**.

**Peter Jones**  
Distinguished Engineer

Email: [petejone@cisco.com](mailto:petejone@cisco.com)  
Twitter/X: [@petergjones](https://twitter.com/petergjones)  
LinkedIn: [in/petergjones/](https://www.linkedin.com/in/petergjones/)



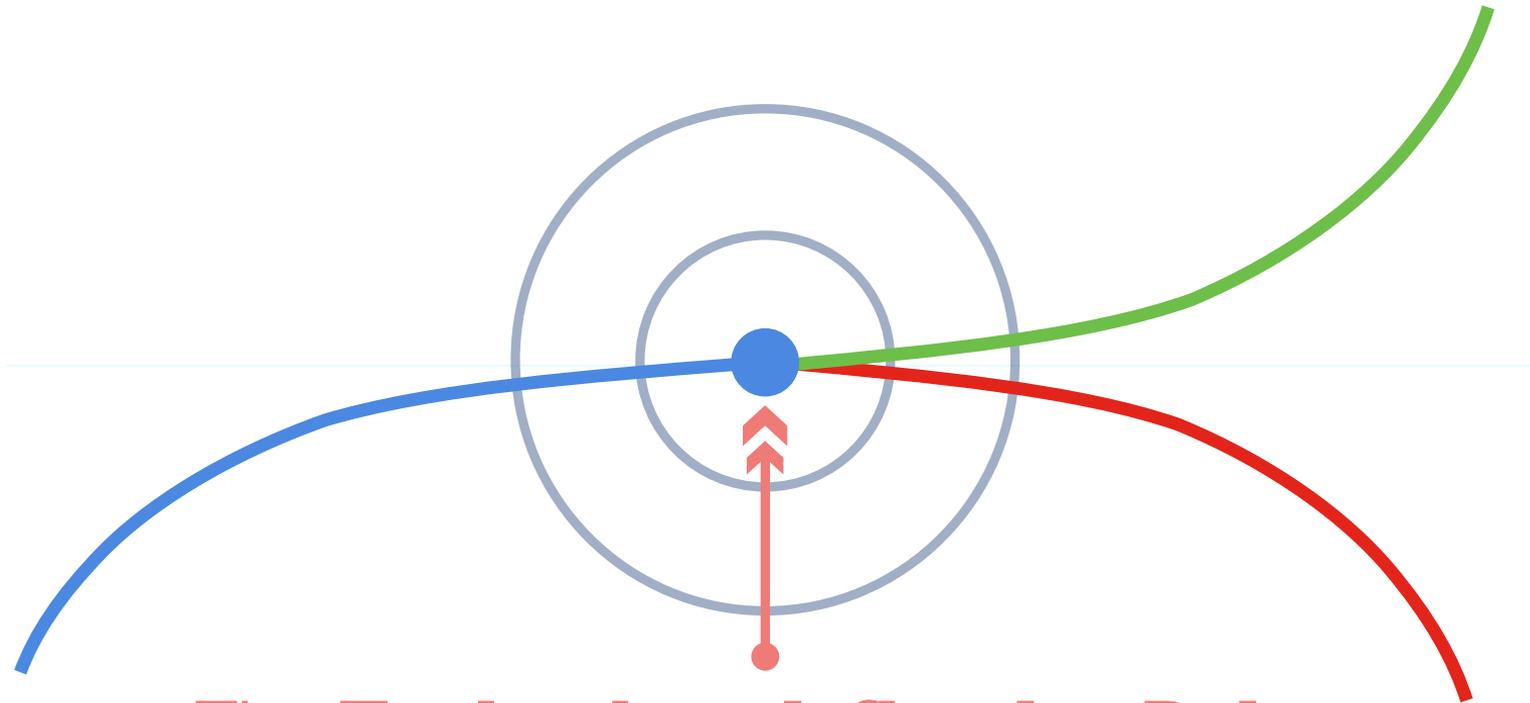
# AI - Overview

CISCO *Live!*

Dave

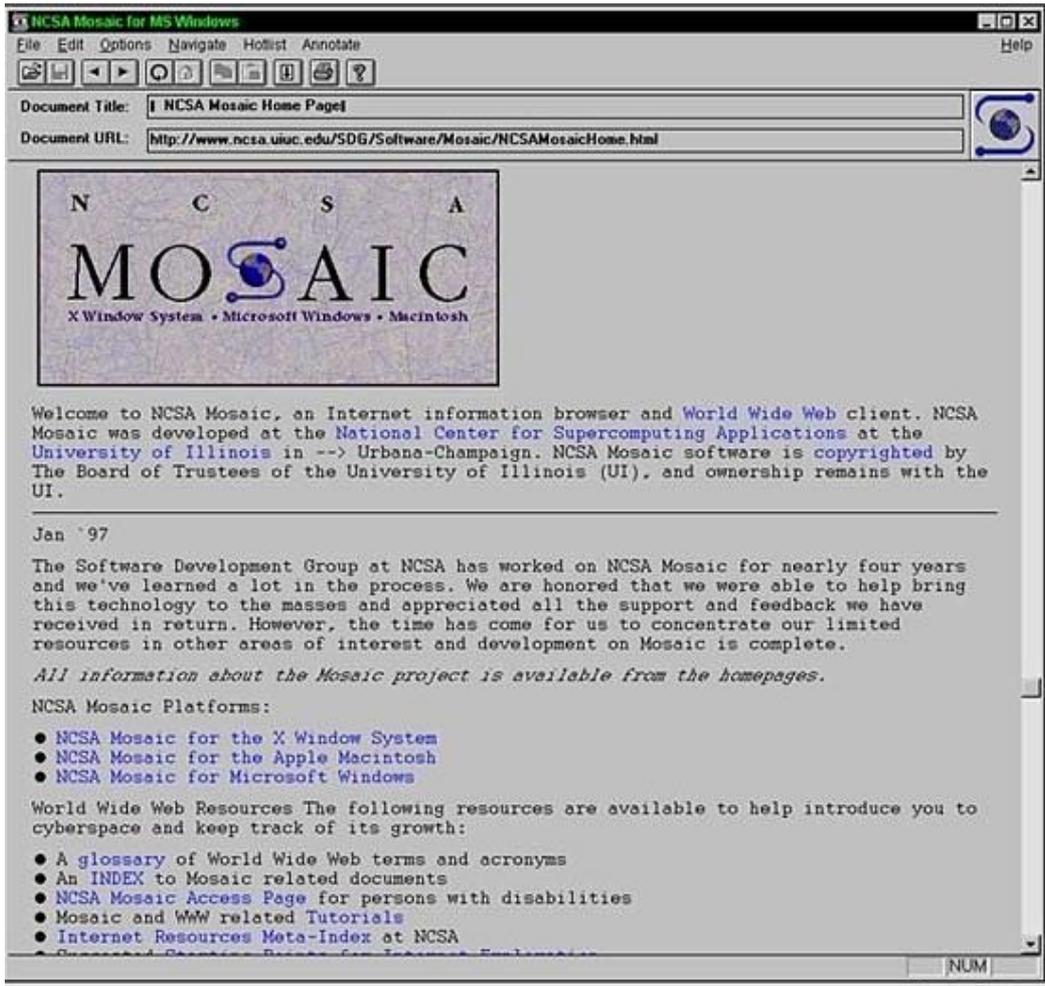


# Rapid change doesn't come in normal times

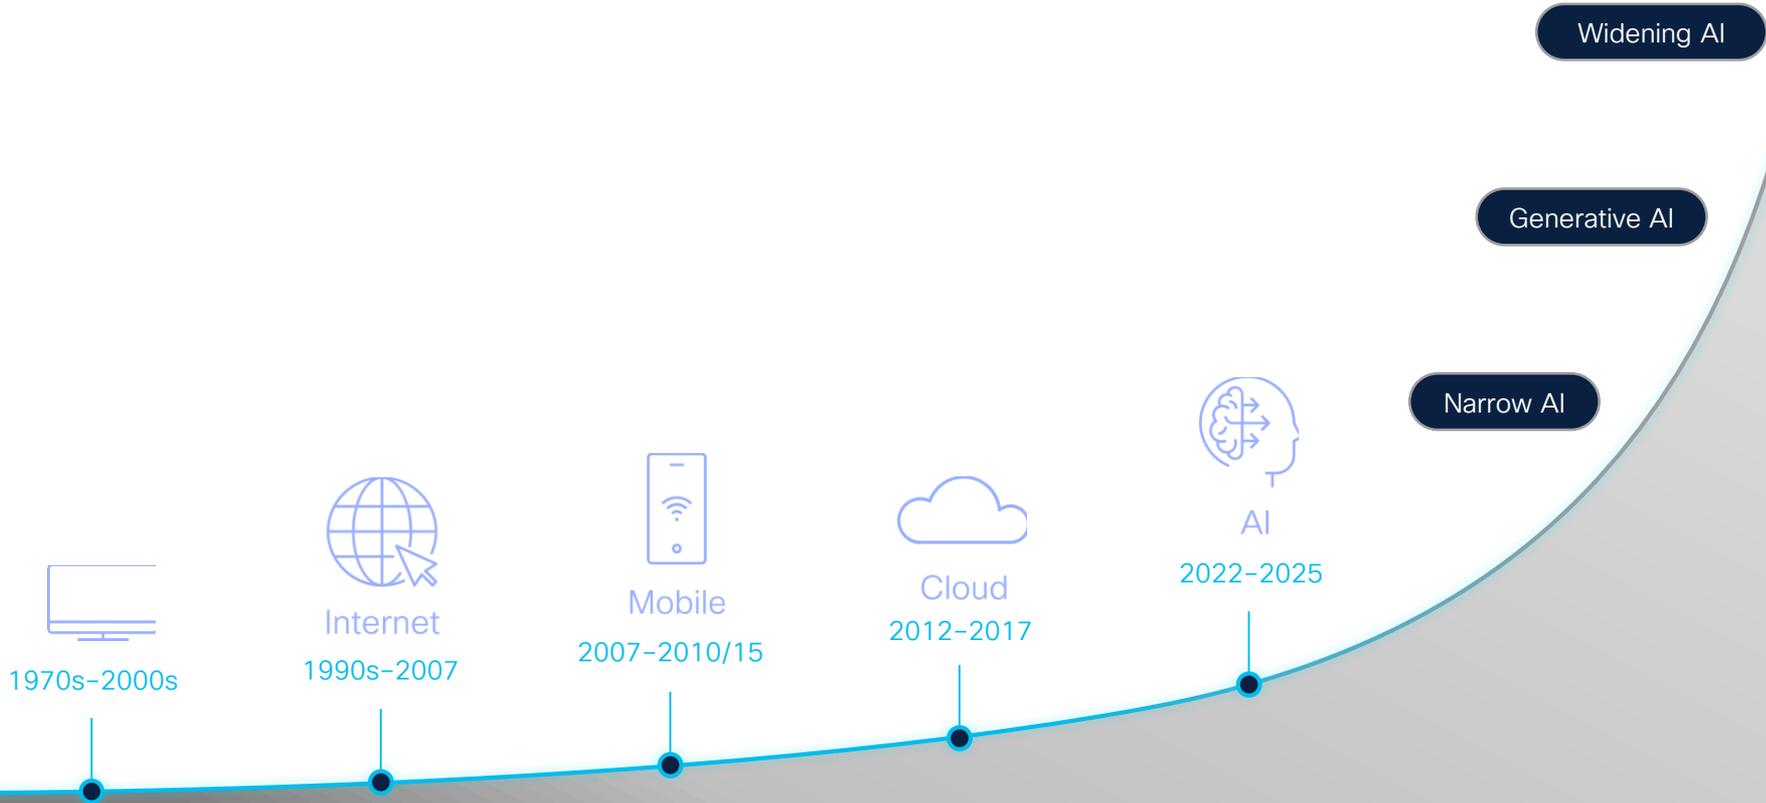


The **Technology Inflection Point**





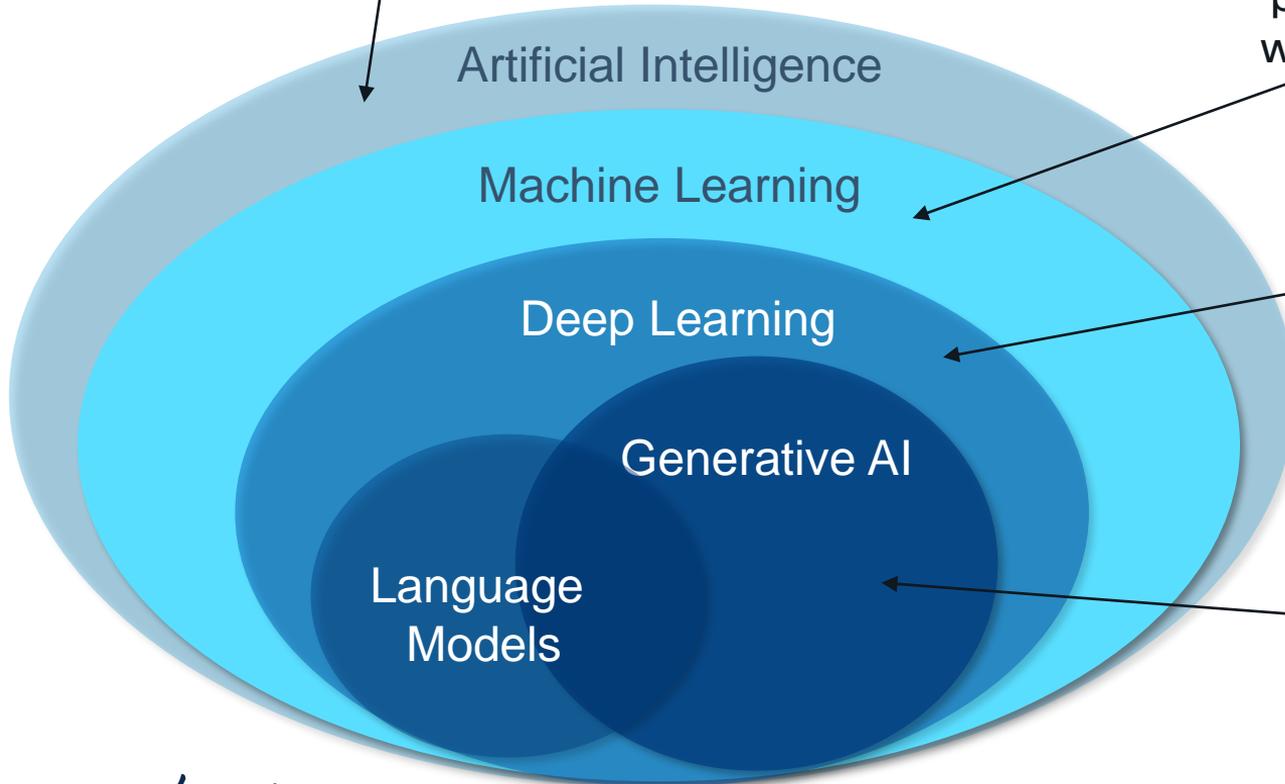
# Generations of Technology Inflection Points



INNOVATION AT SCALE

AI encompasses the whole world of ML and Deep Learning

An AI technology where the rules are not set in the program, but are **learned** while the program is used



A form of ML that uses **Neural Learning Networks**

A powerful mechanism that allow neural networks to **learn language and generate content**

# A Snapshot of GenAI Capabilities ...

Casual / Fun

R tell me a joke about the summer olympics in the style of Albert Einstein

R "Why did the photon refuse to compete in the Summer Olympics?"

Because it didn't want to be weighed down by mass and preferred to move at the speed of light, where time dilation makes every event feel like a lifetime!"

Prompt: A stylish woman walks down a Tokyo street filled with warm glowing neon and animated city signage. She wears a black leather jacket, a long red dress, and black boots, and carries a black purse. She wears sunglasses and red lipstick. She walks confidently and casually. The street is damp and reflective, creating a mirror effect of the colorful lights. Many pedestrians walk about.



Cisco Enterprise Chat AI - powered by Azure OpenAI API.

R Help me write an RFP response why a customer would choose DNA Center

Dear [Customer],

Thank you for the opportunity to respond to your Request for Proposal (RFP) regarding the selection of a network management solution. After carefully evaluating your requirements and considering various options, we believe that Cisco DNA Center is the ideal choice for your organization. Allow

1. Unified Network Management  
This unified approach overall efficiency.
2. Intent-Based Network  
Networking, enabling than manually configu

New Chat Ask me anything...

Cisco Enterprise Chat AI - powered by Azure OpenAI API.

R can you write a python script to add a VLAN and SVI to a catalyst switch using the python library paramiko?

Certainly! Here's an example Python script that uses the Paramiko library to add a VLAN and SVI (Switched Virtual Interface) to a Cisco Catalyst switch:

```
import paramiko

# Define the switch connection details
hostname = 'your_switch_hostname'
username = 'your_username'
password = 'your_password'

# Define the VLAN and SVI details
vlan_id = '10'
vlan_name = 'VLAN10'
svi_ip = '192.168.1.1'
```

New Chat Ask me anything...

BUSINESS -  
Answering RFPs,  
Writing code ...

MULTIMEDIA -  
Text, Pictures, Video, ...

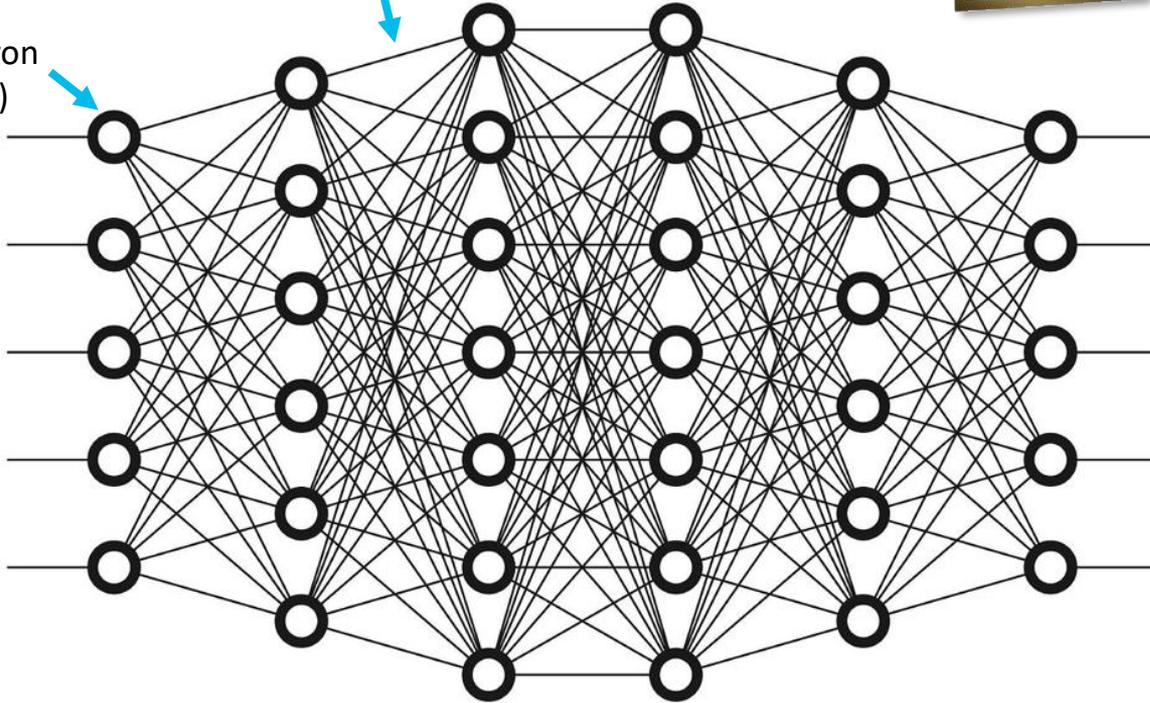
Cisco Live!

# Why is this happening now?

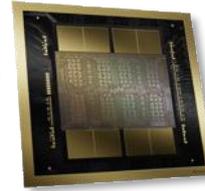
# Scale

Parameter  
(Synapses)

Perceptron  
(Neuron)



  
NVIDIA  
Blackwell



**Advances in Silicon –**  
High-density, High-performance GPUs

NVIDIA Flagship Accelerator Specification Comparison

	B200	H100	A100 (80GB)
FP32 CUDA Cores	A Whole Lot	16896	6912
Tensor Cores	As Many As Possible	528	432
Boost Clock	To The Moon	1.98GHz	1.41GHz
Memory Clock	8Gbps HBM3E	5.23Gbps HBM3	3.2Gbps HBM2e
Memory Bus Width	2x 4096-bit	5120-bit	5120-bit
Memory Bandwidth	8TB/sec	3.35TB/sec	2TB/sec
VRAM	192GB (2x 96GB)	80GB	80GB

Geoffrey Hinton - the “Godfather” of Deep Learning



*CISCO Live!*

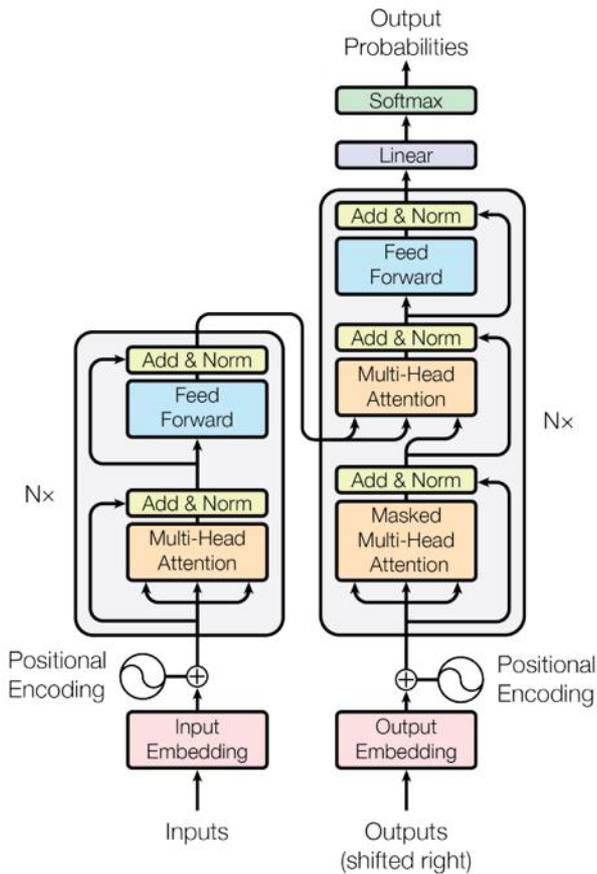


Figure 1: The Transformer - model architecture.

## Attention Is All You Need

---

<b>Ashish Vaswani*</b> Google Brain avaswani@google.com	<b>Noam Shazeer*</b> Google Brain noam@google.com	<b>Niki Parmar*</b> Google Research nikip@google.com	<b>Jakob Uszkoreit*</b> Google Research usz@google.com
<b>Llion Jones*</b> Google Research llion@google.com	<b>Aidan N. Gomez* †</b> University of Toronto aidan@cs.toronto.edu	<b>Łukasz Kaiser*</b> Google Brain lukaszkaizer@google.com	
<b>Illia Polosukhin* ‡</b> illia.polosukhin@gmail.com			

[arXiv:1706.03762](https://arxiv.org/abs/1706.03762) [cs.CL]

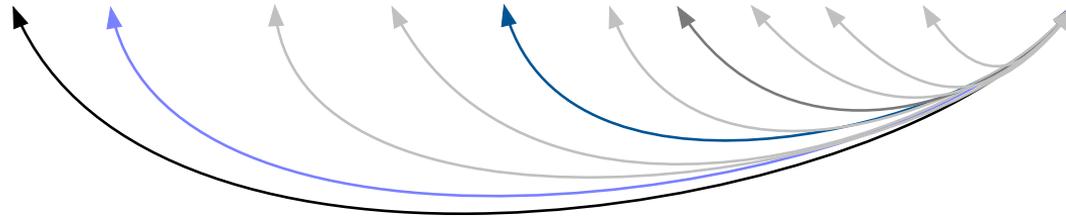
# Attention Mechanism – Overview

You have no problem interpreting “bank” in the following sentence:

“I swam across the river to get to the other bank.”

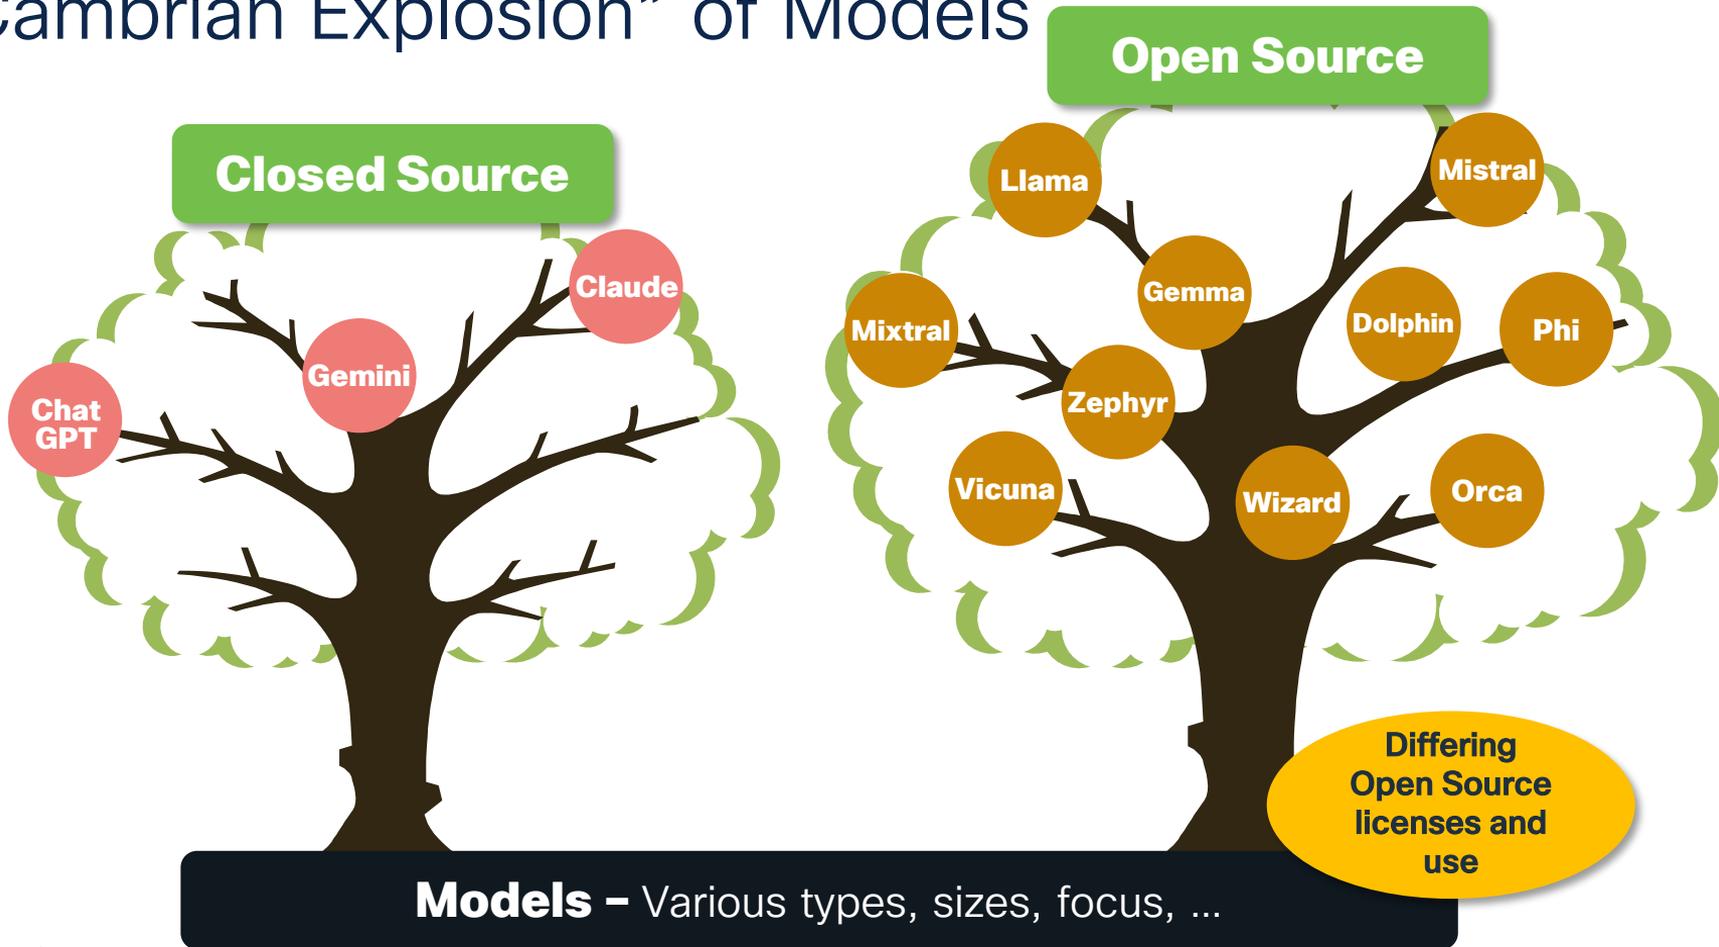
A machine needs some help...

I swam across the river to get to the other bank.

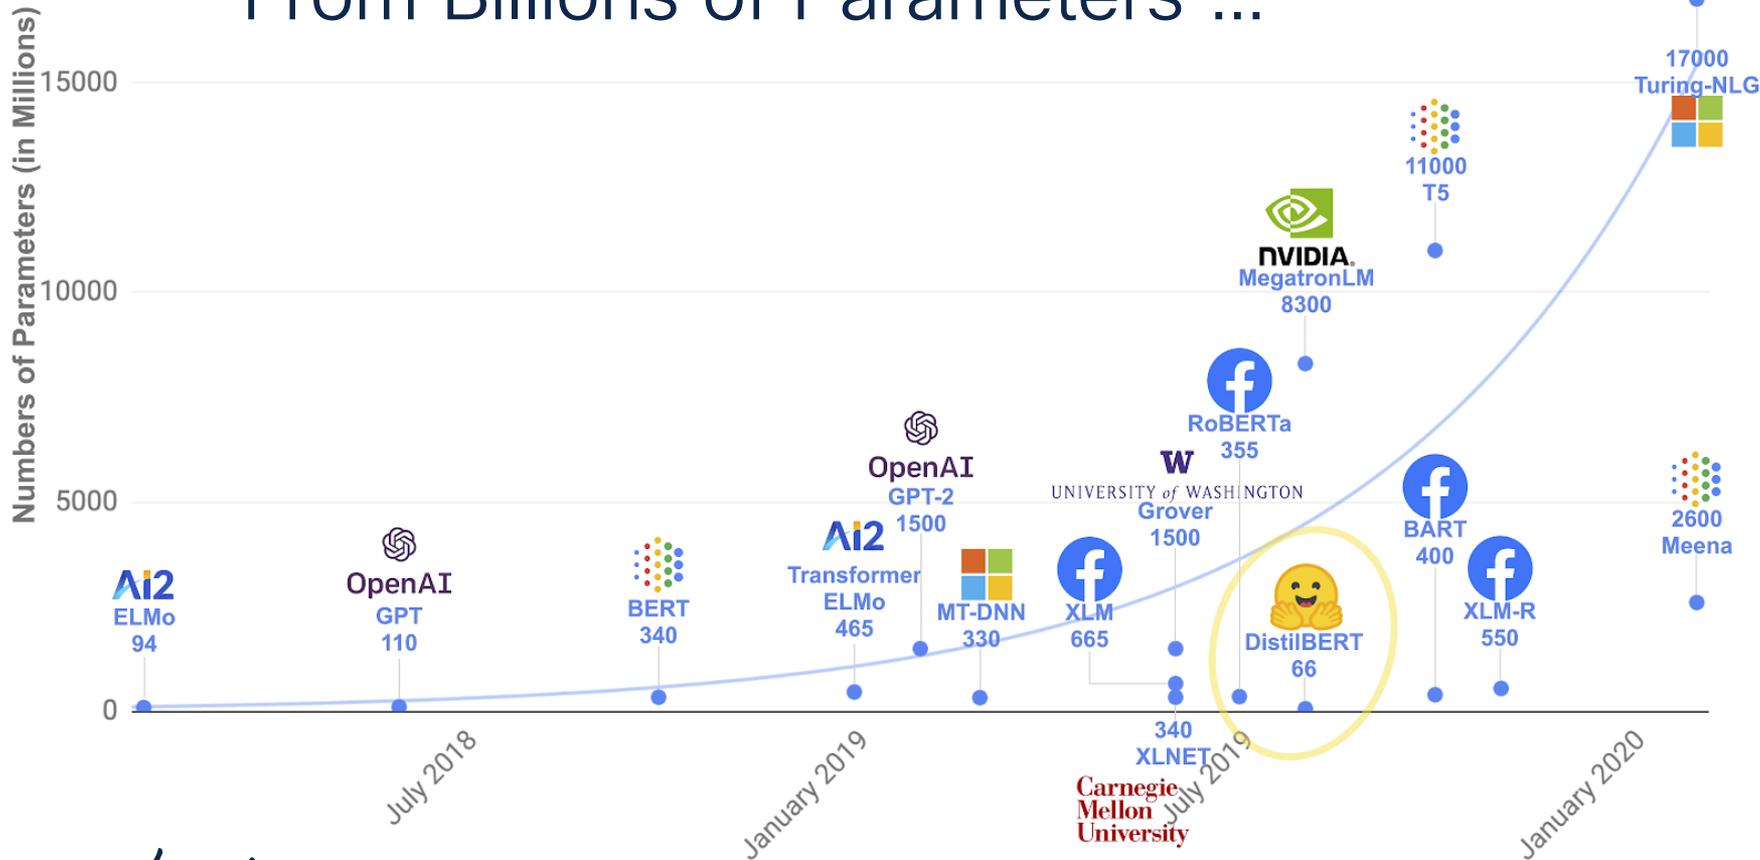


The goal of the attention mechanism is to add **contextual information** to words in a sentence.

# “Cambrian Explosion” of Models



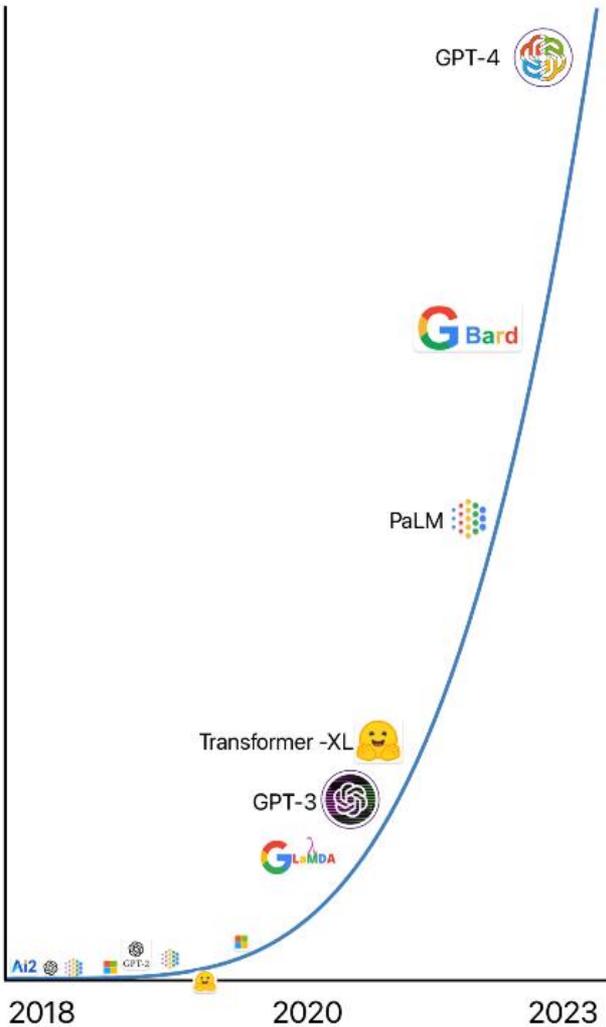
# From Billions of Parameters ...



# To Trillions ...

1,000,000,000,000

500,000,000,000



**CISCO** Live!

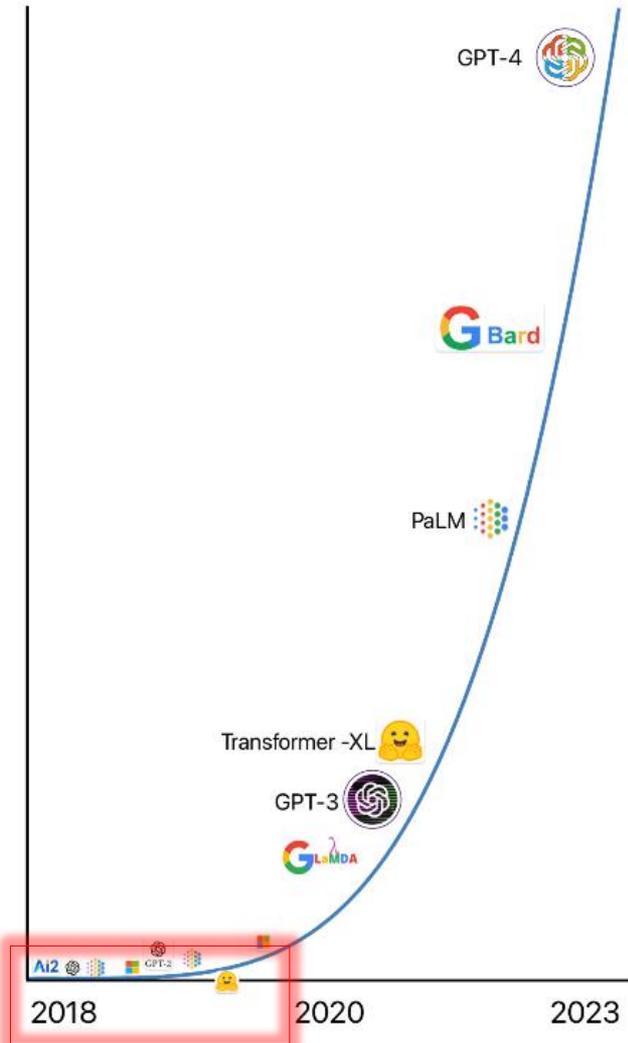
#CiscoLive

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

# To Trillions ...

1,000,000,000,000

500,000,000,000



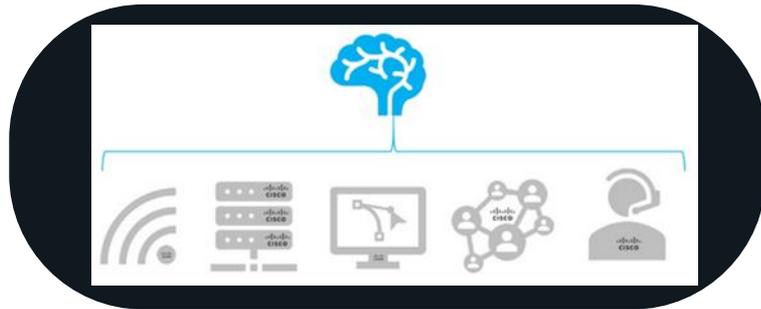
## FUN FACT!

The human brain contains **86 billion neurons**, and over **100 trillion synaptic connections**

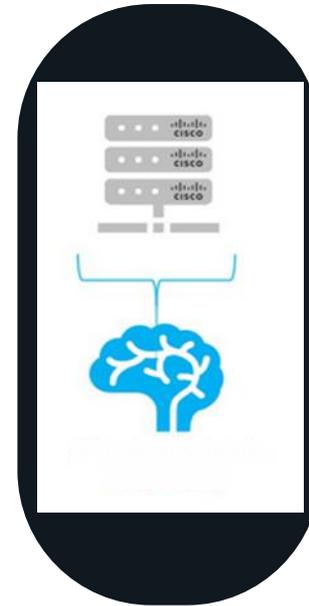


# Artificial Intelligence and Cisco

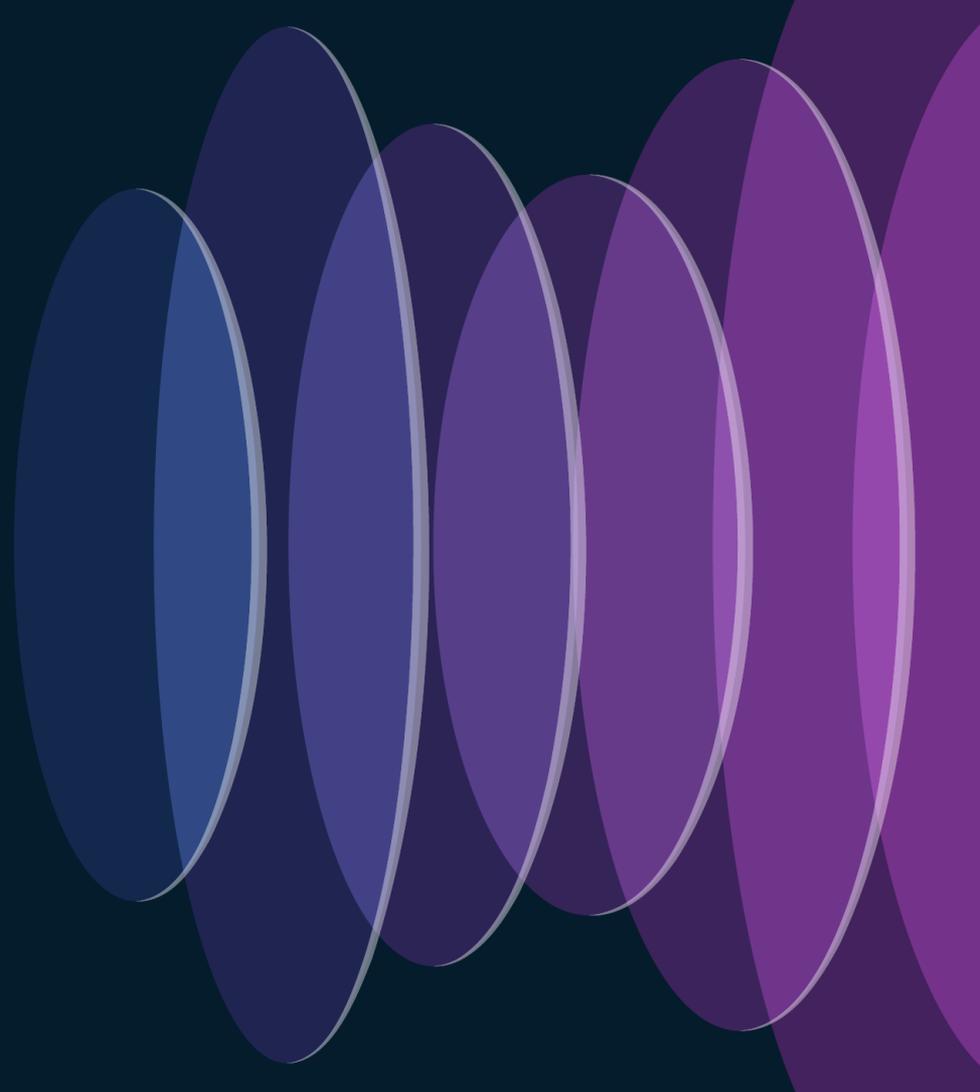
## AI in Cisco – AI to improve products



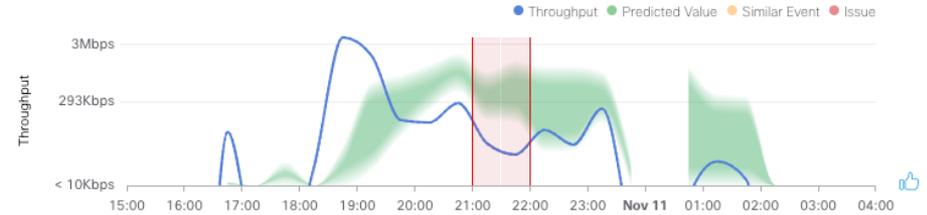
## AI on Cisco – Products to improve AI



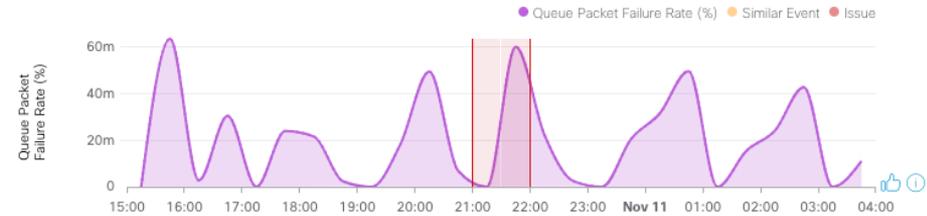
# AI **in** Cisco – Overview



# AI-Powered Wireless Analytics



## Probable network causes

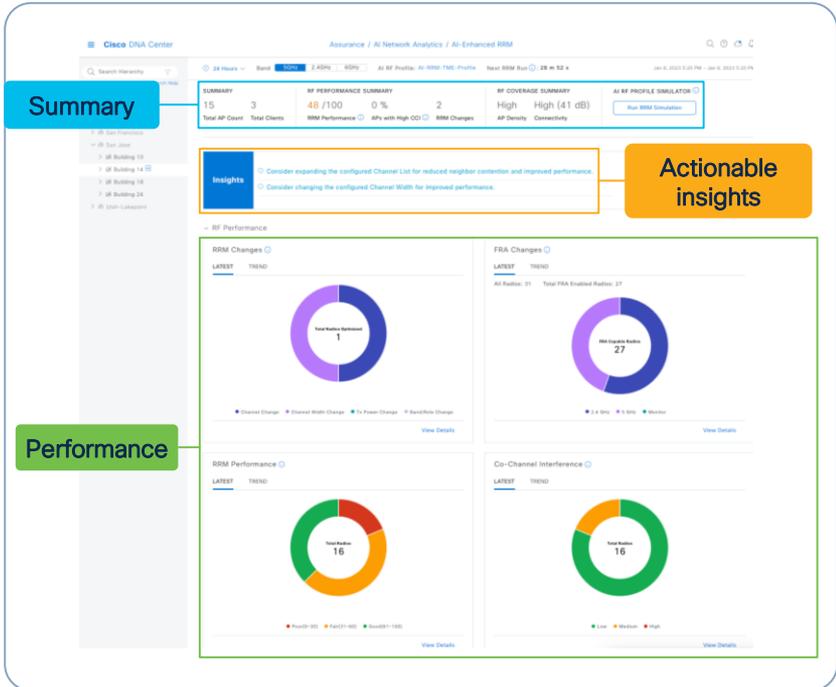


**AI-Driven Anomaly Detection**  
Find + Root Cause Complex Issues

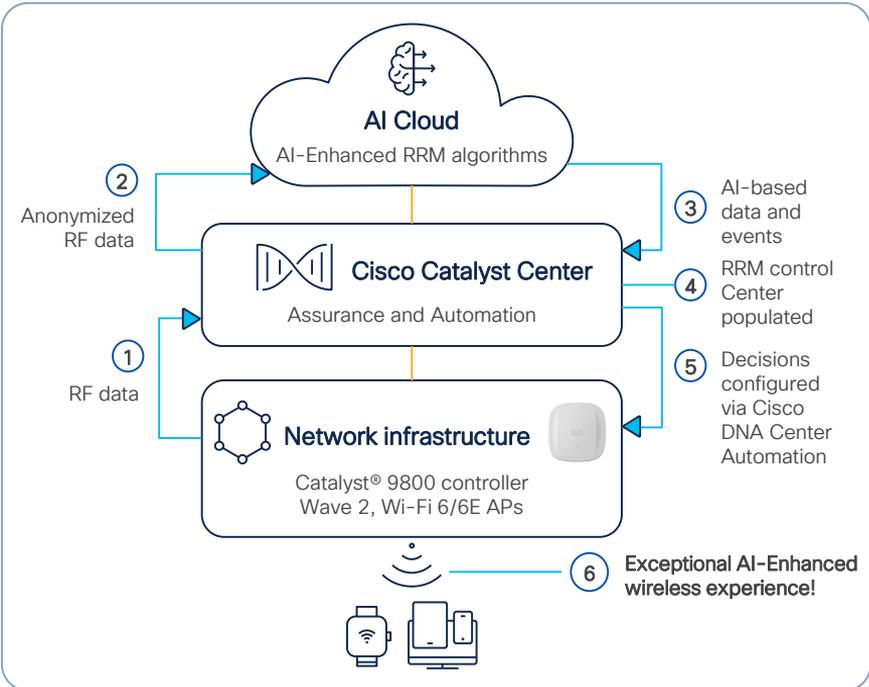
# AI-Enhanced RRM

## Catalyst AI-Driven RRM solution

Deep RF visibility & advanced control



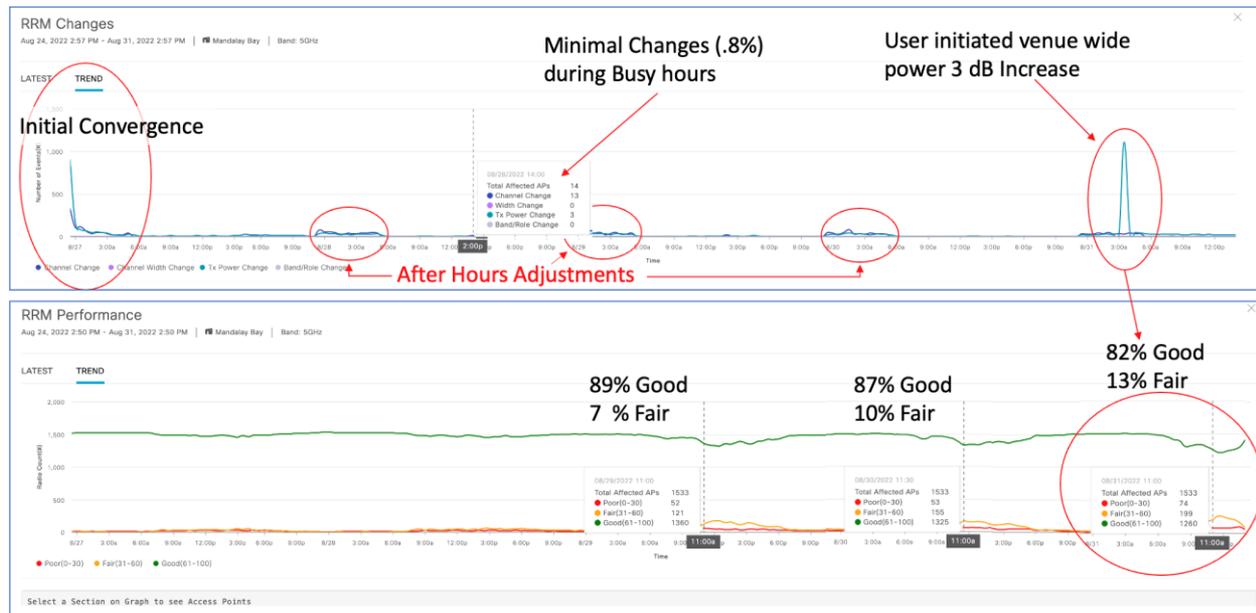
Proactive optimizations for all deployment sizes



# AI-Enhanced RRM in Action

## Cisco Impact 20022

- Initial Convergence  
~3 Hours
- Changes made at Night
- Health stayed above 85%  
(very good with load)
- Manual Changes made last day, easy to spot the decrease in efficiency



# AI Policy Assistant

## Cisco Live Melbourne (Dec 2023)

## Policy Assistant

- Reduces security policy complexity
- Optimizes policy for efficacy & efficiency

The screenshot shows the Cisco Defense Orchestrator (DO) interface for the 'ACP - Production' configuration page. The page displays a table of access control rules with columns for Name, Action, Source (Zones, Networks, Ports, Dynamic Attributes), and Destination (Zones, Networks, Ports, Applications). A 'Policy Assistant' chat window is overlaid on the right side of the screen, containing a user query and an AI response.

**AI Assistant Chat:**

**User:** I'm trying to determine what rules are poorly written by analyzing the number of hits. Can you show me the rules with the least number of hits in the last 30 days?

**AI Assistant 8:12 AM:** There are five access control rules with low hit counts in the last 30 days.

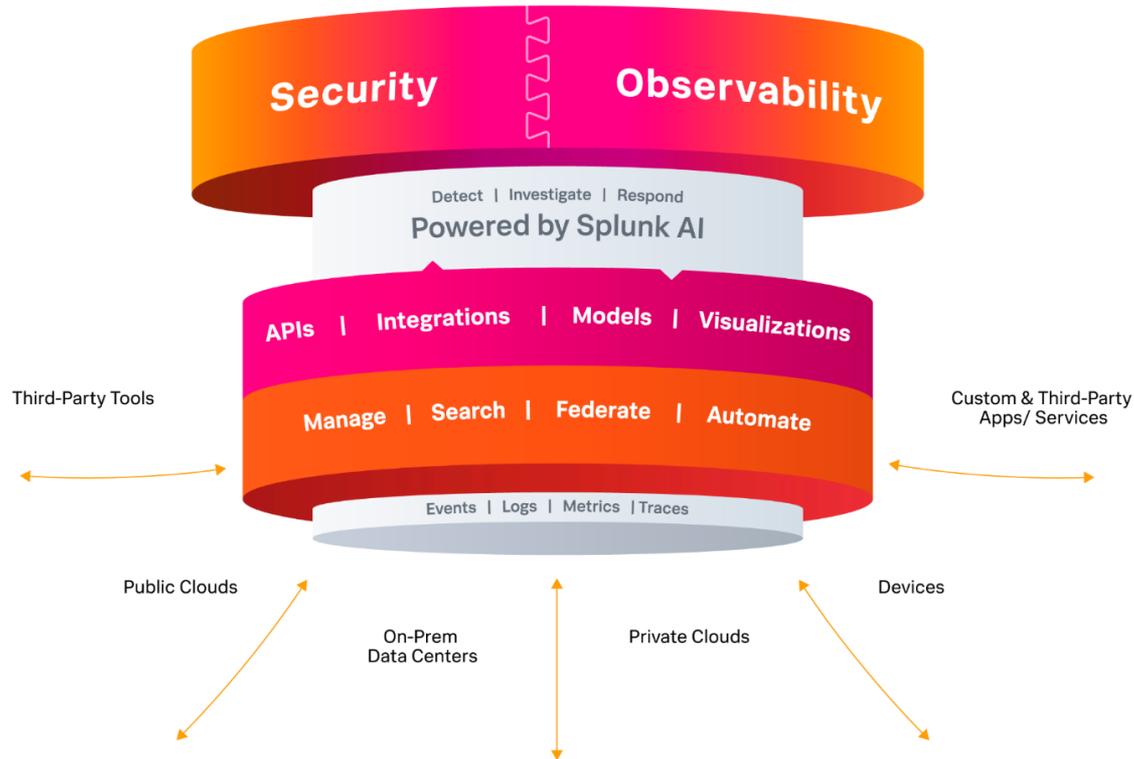
[View details](#)

[Regenerate](#)

Ask a question or request, or type "?" for suggestions

# Unified Security & Observability Platform

*Powered by Splunk AI*



# AI/ML Capabilities across the Splunk Portfolio

## AI FOR SECURITY

Enterprise Security, UBA

ML-Powered Detections, Anomalous User Actions

## AI FOR OBSERVABILITY

ITSI, APM, IM, On-Call

Predictive Analytics, Anomaly Detection, Adaptive Thresholding, Incident Correlation, Alert Noise Reduction, Alert Autodetect, Suggested Responders

### Assistive Intelligence Experiences

Splunk AI Assistant, Splunk App for Anomaly Detection

### Customizable ML

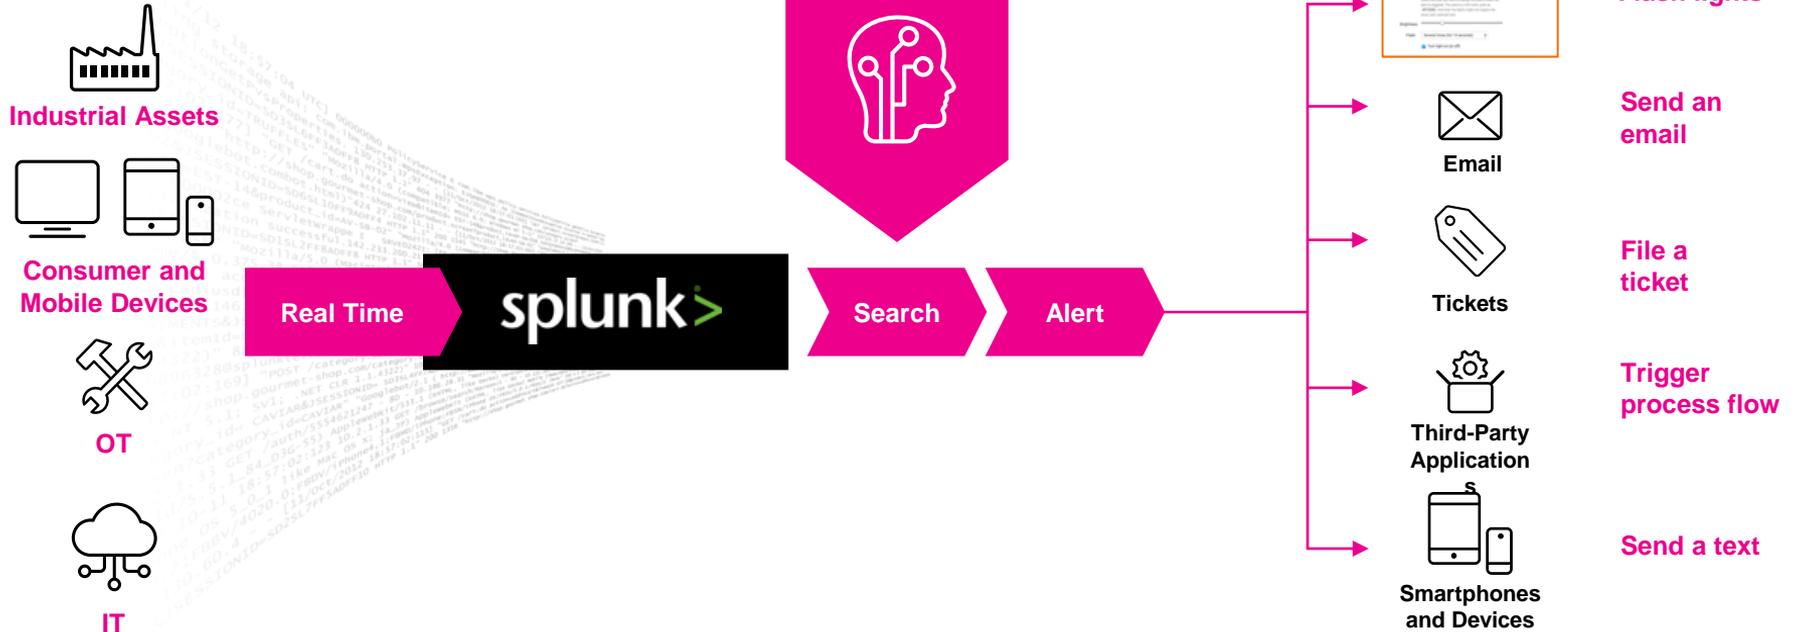
Splunk Machine Learning Toolkit, Splunk App for Data Science and Deep Learning, Python for Scientific Computing

## THE SPLUNK PLATFORM

Powerful search commands for statistical analysis, predictive analytics and clustering

# Splunk Machine Learning Tool Kit (MLTK)

Every Search Can Use Machine Learning



# Splunk App for Data Science and Deep Learning (DSDL)

Extension for MLTK to operationalize advanced custom AI / ML use cases



- **Built for Data Scientists**
- **Code Examples:** Guided model building, testing, and deployment of deep learning frameworks
- **Container Management:** Containers & Models can be productionized for scalability & optimization of resources, e.g. CPU & GPU
- **Frameworks:** PyTorch, Tensorflow, SpaCy, DASK, Rapids, Spark, Jupyter Notebooks & Tensorboard, ...
- **Extensible to operationalize any use case**
- **Open Source:**  
New frameworks & Python Libraries freely available for integration into DSDL via github
- **Acceleration:** Support for GPU accelerated machine learning and inference pipelines

Content Overview

Architecture

Learn about the architecture how the Deep Learning Toolkit works

Model Development Workflow

Learn about the predefined rapid model development workflow using Jupyter Lab Notebooks

Multi GPU Computing

Learn how you can leverage multi GPU computing for high computational workloads

File Edit View Run Kernel Task Settings Help

Python 3

Notebook for Deep Learning App for Splunk

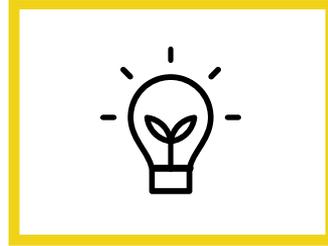
Logistic Regression in PyTorch

MACHINE LEARNING TOOLKIT + SPLUNK APP FOR DATA SCIENCE AND DEEP LEARNING

splunk>enterprise splunk>cloud

# Splunk – Investing in Foundational + Generative AI

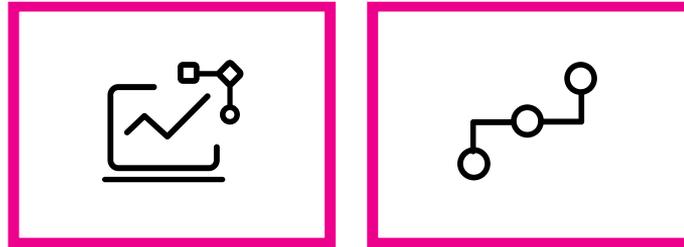
*Combining predictive analytics, accelerated investigation, and workflow enhancements*



Generative AI Capabilities

## Make Everyone an Expert

Reduce need for environment and tool expertise by simplifying content creation and investigation workflows



Foundational AI Capabilities

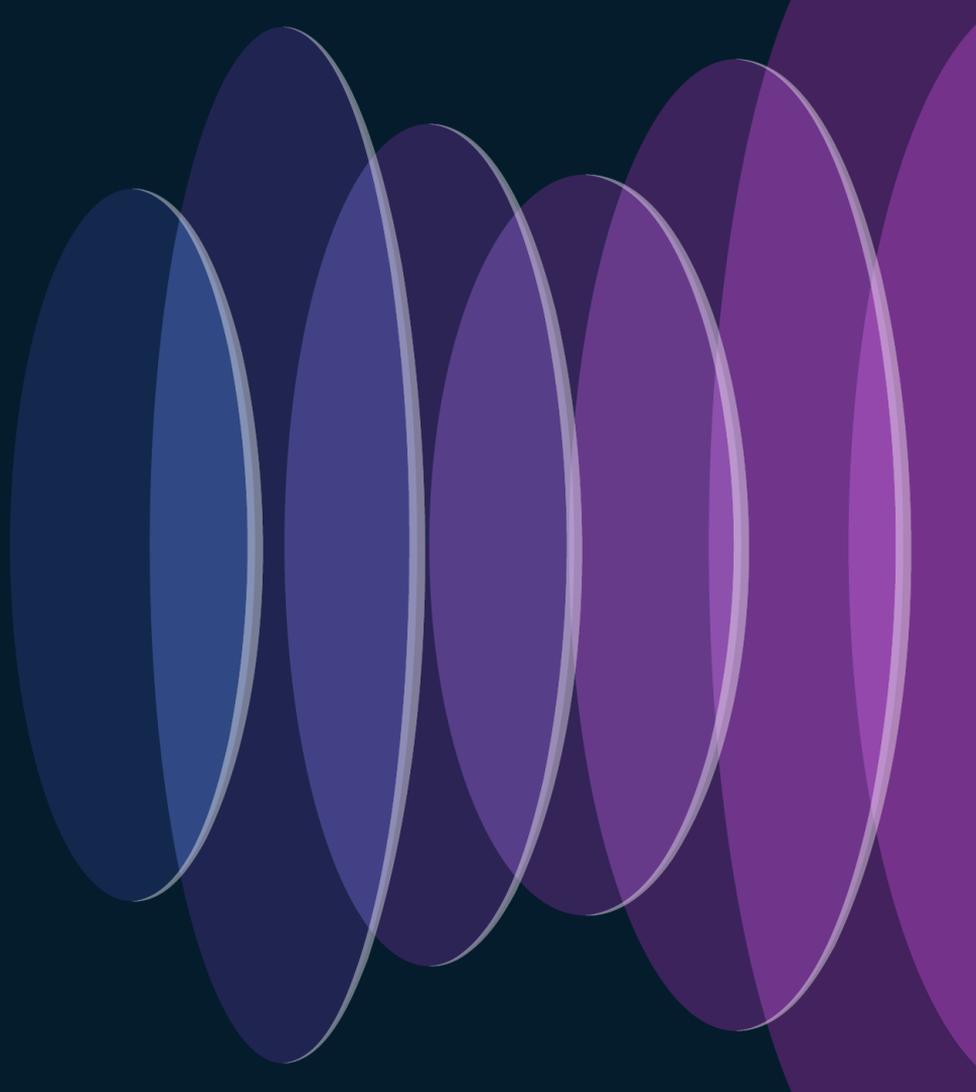
## Detect and Predict

Real-time, streaming analysis to detect anomalies and forecast trends

## Correlate and Diagnose

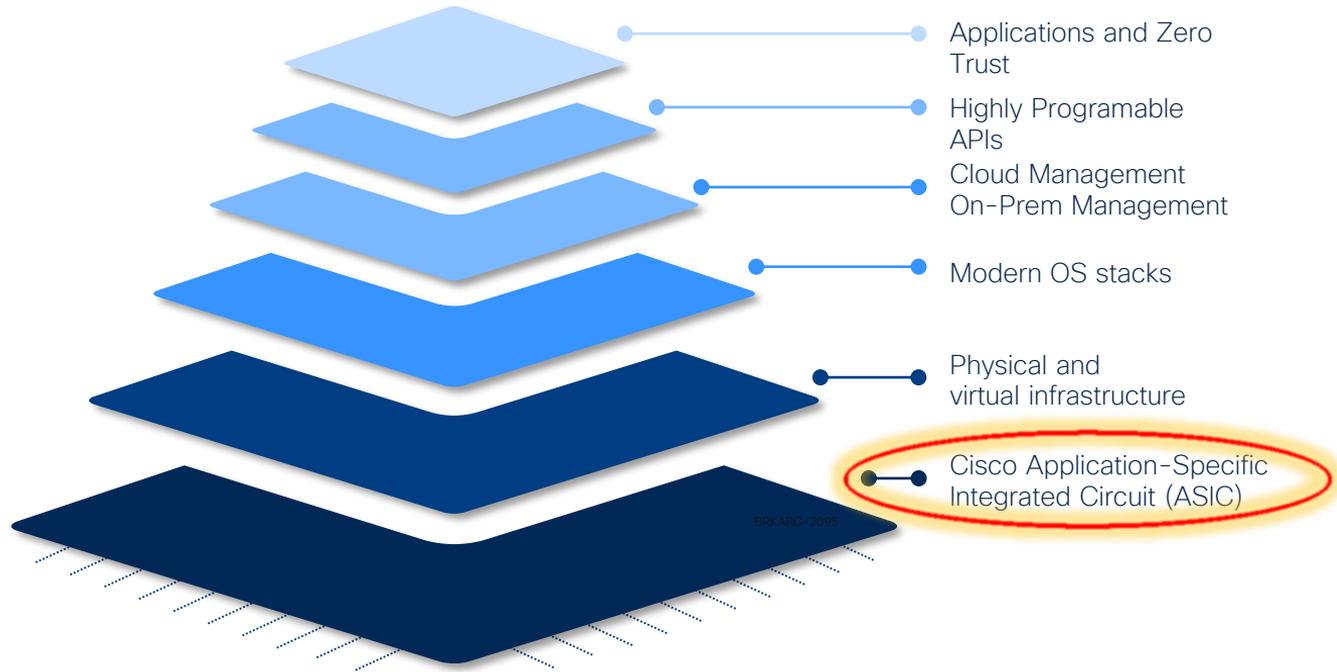
Aggregate and analyze all data to investigate and identify root causes

# AI **on** Cisco – Networking Silicon Overview



# Cisco Silicon

## Foundational Elements



# ASICs

## From Definition to Deployment



Marketing



Engineering



Then, it starts with coding...

Verilog  
VHDL

```
library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.numeric_std.all; -- for the unsigned type

entity COUNTER is
  generic (
    WIDTH : in natural := 32);
  port (
    RST   : in std_logic;
    CLK   : in std_logic;
    LOAD  : in std_logic;
    DATA : in std_logic_vector(WIDTH-1 downto 0);
    Q     : out std_logic_vector(WIDTH-1 downto 0));
end entity COUNTER;

architecture RTL of COUNTER is
  signal CNT : unsigned(WIDTH-1 downto 0);
begin
  process(RST, CLK) is
    begin
      if RST = '1' then
        CNT <= (others => '0');
      elsif rising_edge(CLK) then
        if LOAD = '1' then
          CNT <= unsigned(DATA); -- type is converted to
        unsigned
        else
          CNT <= CNT + 1;
        end if;
      end if;
    end process;

    Q <= std_logic_vector(CNT); -- type is converted back to
  std_logic_vector
end architecture RTL;
```

# ASICs

## From Definition to Deployment



Marketing

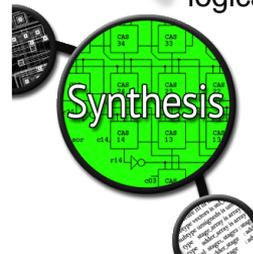


Engineering

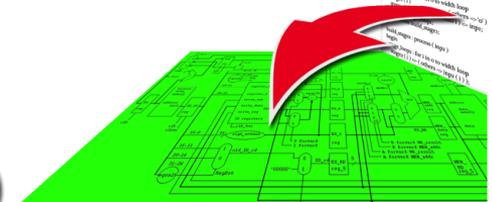


Then, it starts with coding...

Verilog  
VHDL



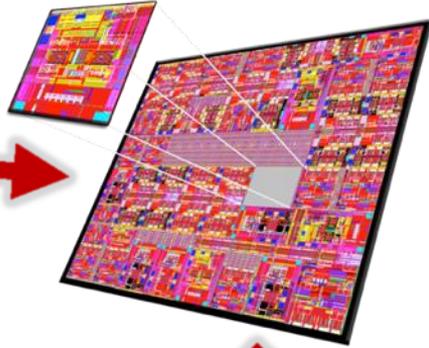
Synthesis Process  
Converts code into  
logical gate constructs (Netlist)



# ASICs

## From Definition to Deployment

Arrange and interconnect constructs, connect power, minimize crosstalk, etc...



**Floor  
Planning and  
Placement**



**Etch Design onto Silicon Wafer ...**

Light Source

Photo  
Masking

Silicon  
Wafer

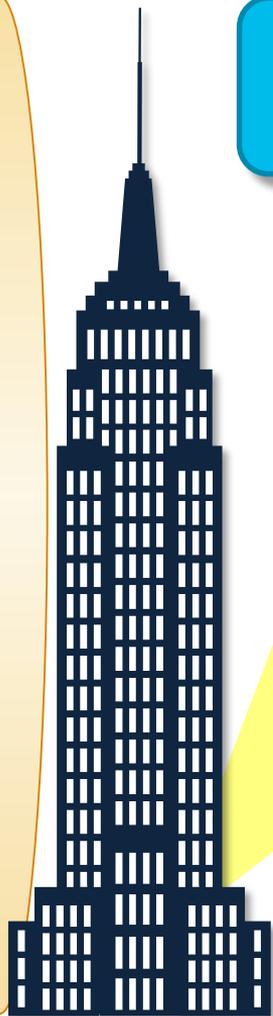
Photo  
Mask

How **SMALL**  
is **SMALL**?

**ONE NANOMETER –**  
**less than 1/4<sup>th</sup> of an inch!**  
*... about the same thickness  
as three pennies  
on this scale ...*

Single human hair

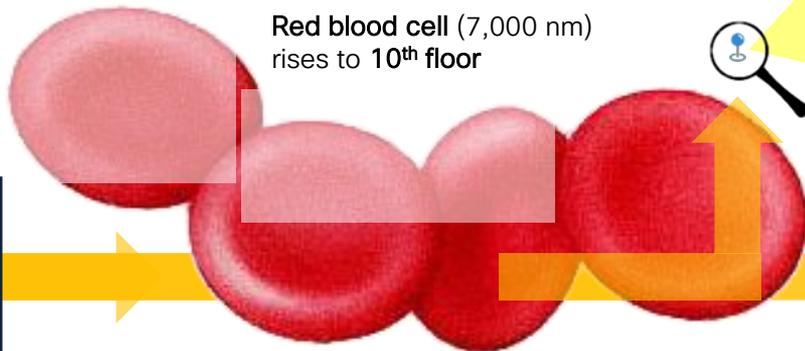
100,000  
nm



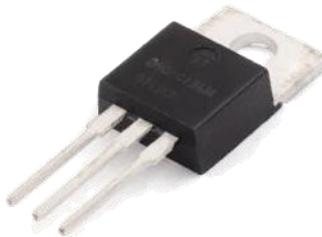
Empire State  
Building =  
1454 feet  
to tip =  
443 meters

... and then we come to  
**this little pinprick** over here ...

Red blood cell (7,000 nm)  
rises to 10<sup>th</sup> floor



*... and we **build**  
**transistors** measured in  
**nanometers** ...*



# How do Humans Build Something This Small?

ASML

## Driving affordable scaling

Using EUV makes scaling more affordable for chipmakers and allows the semiconductor industry to continue its pursuit of Moore's Law. The EUV systems are used to print the most intricate layers on a chip. For example, an advanced Logic chip at the 3 nm node is made up of about 80 layers. Around twenty six of these are printed using EUV, with the rest of the layers printed using various DUV systems. Both types of technology will be required in parallel for many years to come, and we're continuing to advance both technologies.

NXE:3400B



The size of a bus

Weight: 180 tons

Cost: \$150M+ USD

It takes something **BIG**  
to make something small 😊

# Fun Fact!

## Apollo Guidance Computer



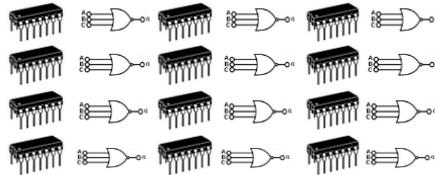
We put a man here ...

... using this ...



... which was built from nothing but **that** ...

4100 ICs,  
each of which  
contained a  
single 3-input NOR gate



In other words ...

we put a man on the moon with less than 10,000 transistors ...

It takes 19.2 billion transistors to route your packets!

With the appropriate security, segmentation, QoS, encryption, fragmentation, etc, etc ...

*cisco Live!*

# ASICs

## From Definition to Deployment

### How much exactly?

Numbers vary, but SemiAnalysis puts it at **\$50M-\$75M** for TSMC 7nm and mask costs have increased from **~\$10M at 7nm to ~\$40M for 3nm**.

Source: <https://www.semianalysis.com/p/the-dark-side-of-the-semiconductor>



Marketing requirements



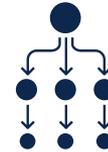
Architecture



RTL design



Synthesis



Floor planning



Fabrication



2 to 5 years

Building a new ASIC takes a lot of time and money



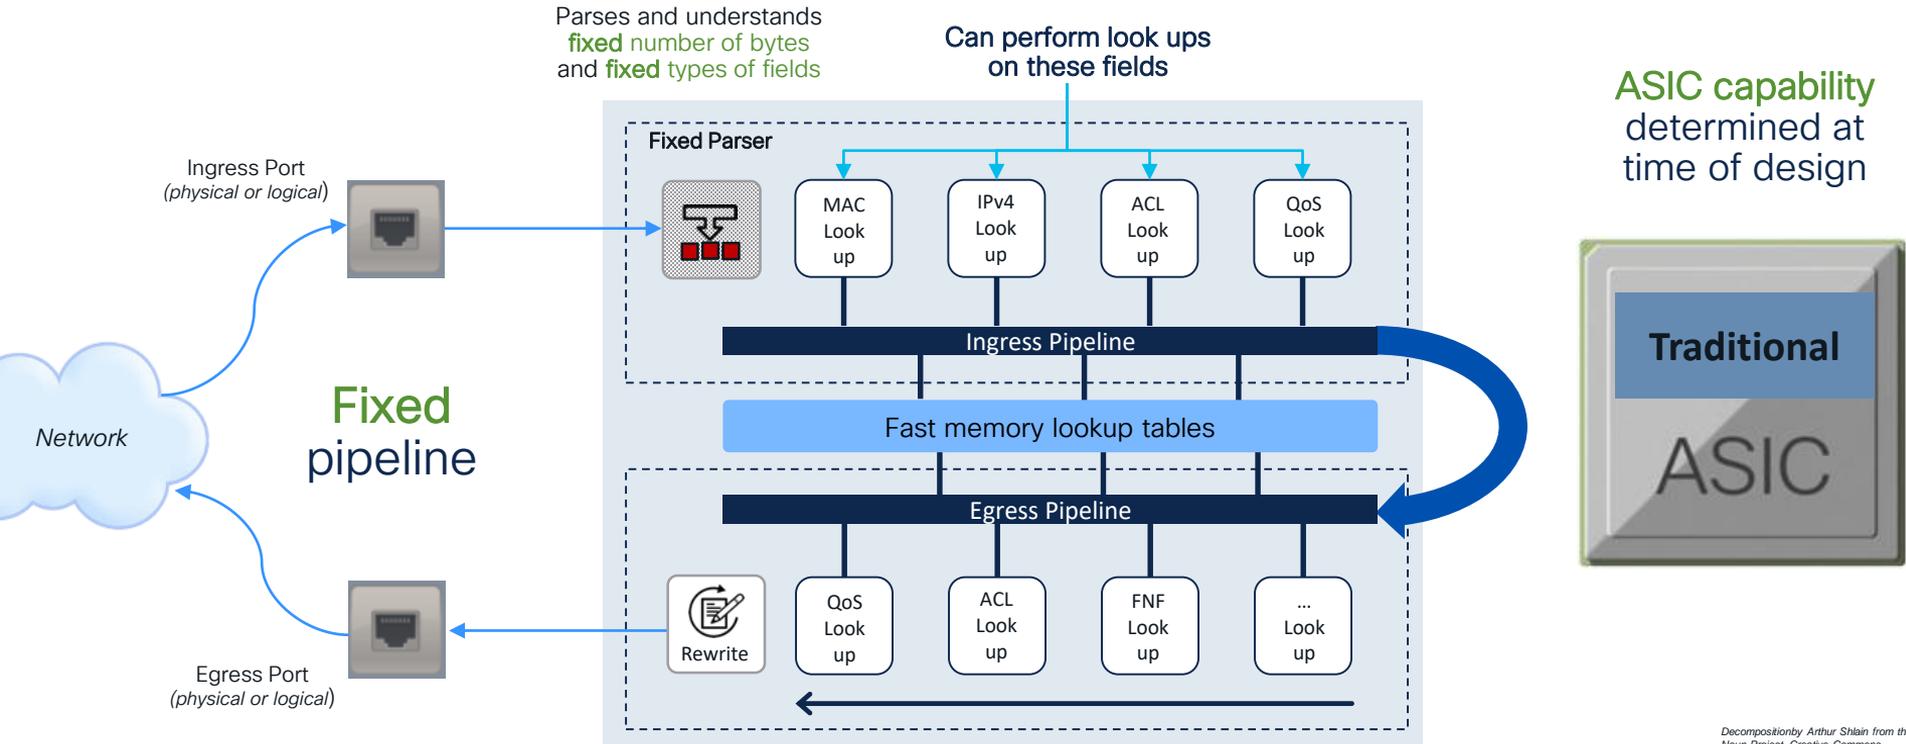
Traditional  
ASIC Silicon

Functional – Fast – but **FIXED** in nature



WHAT'S IN A CHIP?

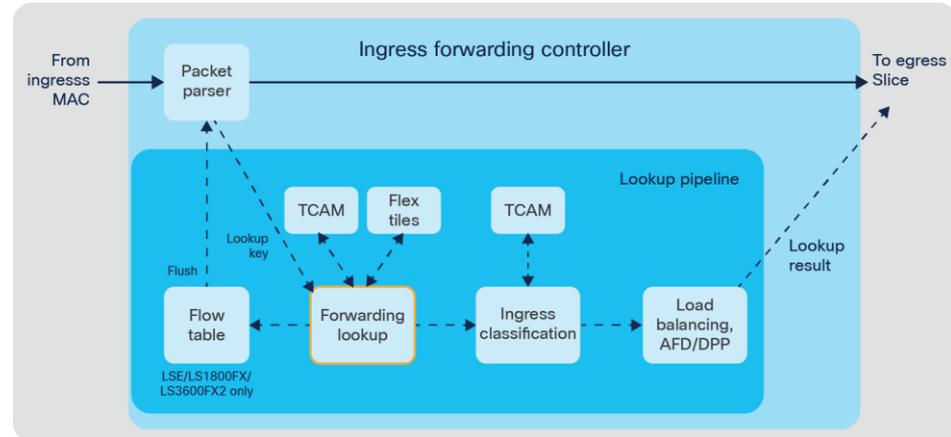
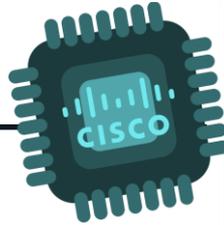
# Networking ASICs – Fixed Pipelines



# Cloud Scale ASIC – Themes

## Purpose Built for DCN

- SSX – Streaming Statistics
- Smart Buffers – DPP + AFD
- Flex Scale –  
Route, Policy, Host Tables
- Integrated Line-rate Encryption
- Align HW and SW Innovations
- No 3rd-party SDK Limitations



Enter

# PROGRAMMABLE Silicon

Much faster than a



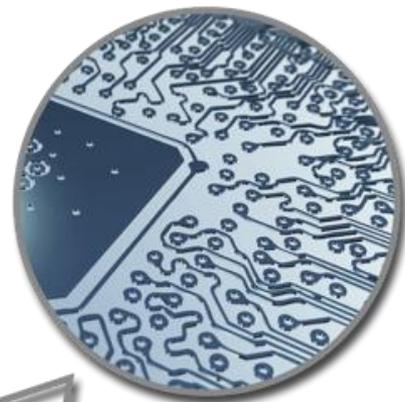
Much more capable than an



ASIC Engineer



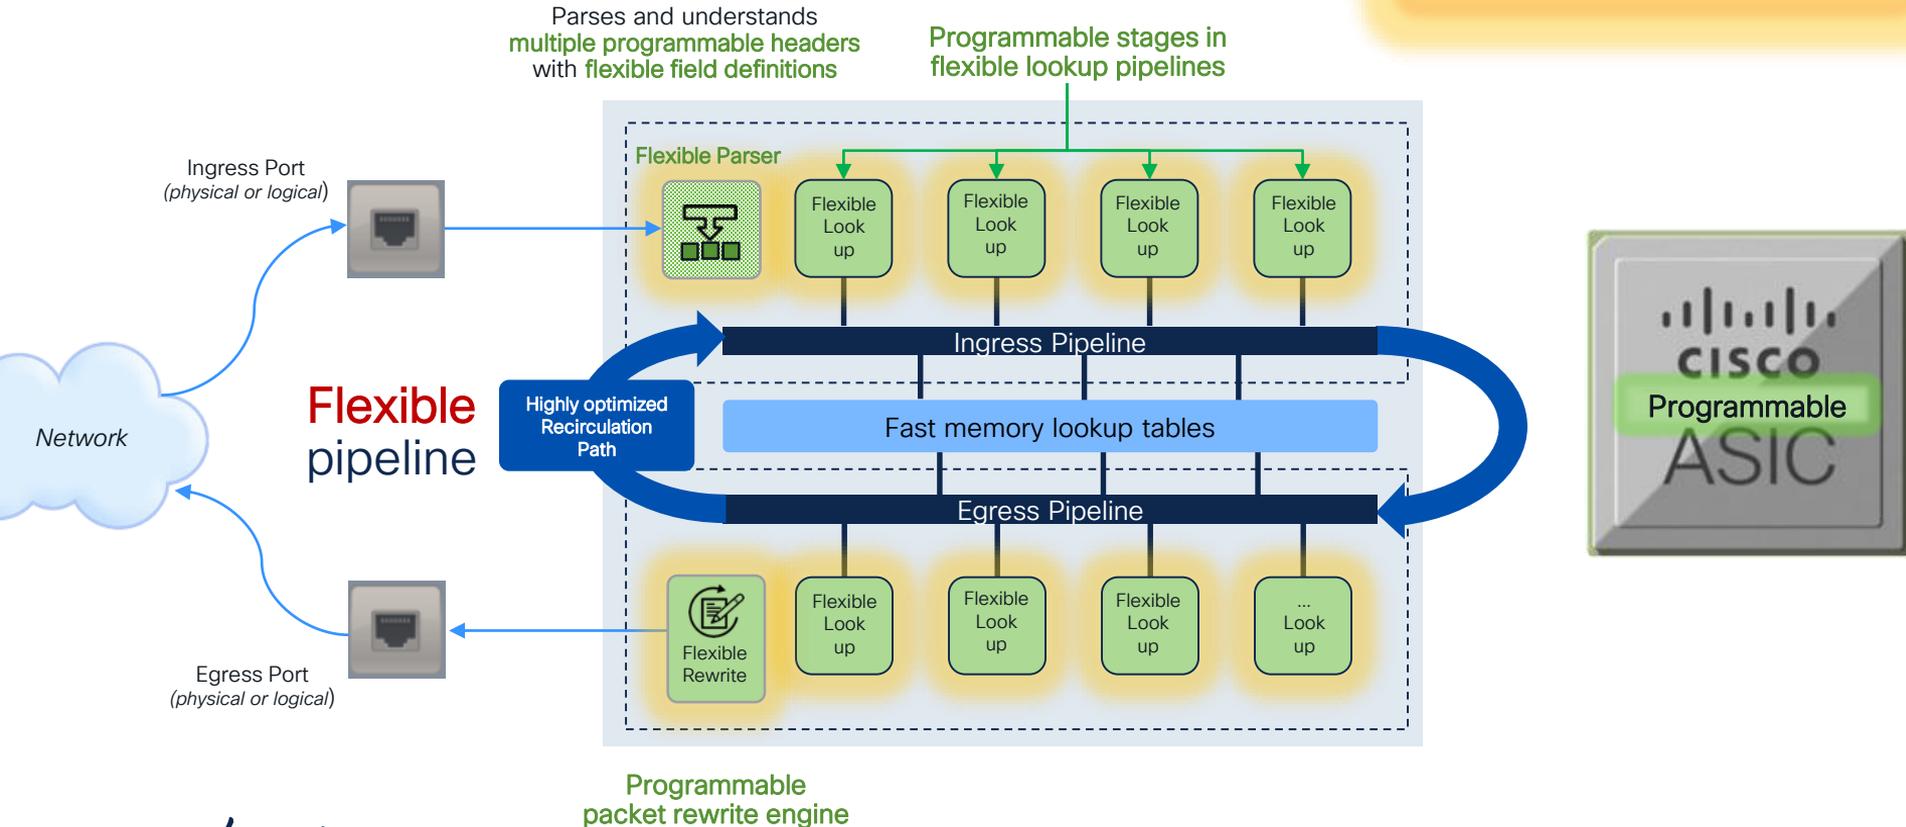
Microcode Update



New ASIC Functionality

# Cisco Programmable ASIC

ASIC Capability can be updated via software, operating at hardware speeds



# Catalyst 9000 Series



## UADP

## Programmable Switching

Up to 1.6Tbps



# Catalyst 9000 Series



## UADP

### Programmable Switching

Up to 1.6Tbps



## Cisco Silicon One™

### Programmable Switching & Routing

Up to 25.6Tbps  
Global Routing Scale  
Deep Buffering



# Catalyst 9000 Series



## UADP Programmable Switching

Up to 1.6Tbps



## Catalyst 9600X/9500X



## Cisco 8000 Series



## Nexus 9800

# Cisco Silicon One™ Programmable Switching & Routing

Up to 25.6Tbps  
Global Routing Scale  
Deep Buffering



# Cisco Silicon One™

CISCO *Live!*

Peter



# Introducing Cisco Silicon One™

## Industry-leading switching and routing silicon



### Switching silicon

- **High throughput**  
Extremely fast hardware-based Layer 2 through 4 forwarding and services (measured in terabits per second)
- **Optimized scale**  
Optimized for campus LAN environments with moderate IP and MAC scale (tens to hundreds of thousands)
- **Low latency**  
Extremely low hardware-based system latency (measured in nanoseconds and microseconds)
- **Streamlined buffering**  
Shallow buffering systems to reduce latency, with very high throughput



### Routing silicon

- **Flexible features**  
Complex Layer 3 through 7 forwarding and services (measured in 10s-100s of gigabits per second)
- **Massive scale**  
Optimized for WAN and service provider environments with very high IP scale (hundreds of thousands to millions)
- **Mixed interfaces**  
Support for Ethernet, serial, cellular, and other types and speeds in a single system
- **Deeper buffering**  
Deep buffers to accommodate different speeds, bursts, and different flow patterns

Cisco Silicon One™: Bringing switching and routing convergence

# Cisco Silicon One™ Q200



12.8 Tbps



8.1 Bpps



8GB HBM  
for  
deep buffers



2M IPv4  
or 1M IPv6  
routes



P4 NPL  
Programmable  
Pipeline



50G  
PAM4  
Serdes

Cisco Silicon One™  
Q200  
Industry Leading  
12.8T System on Chip



**First 7nm ASIC**  
lowest Watts/Gbit



**Fully P4 programmable**  
enabling feature velocity



**Multi-Slice architecture**  
for flexibility and scale

Routing Capabilities with Switching Power and Performance

# One Architecture – Many Devices



**Q2xx/Q2xxL Family**  
Q02, Q201, Q101, Q200  
3.2T, 6.4T, 10.8T, 12.8T

**G100/G20x Family**  
G100, G202, G200  
25.6T, 25.6T, 51.2T

- First network silicon to break the 10-Tbps barrier
- Comprehensive routing, with switching efficiency
- Flexible P4 NPL programmable packet processing

- **Multiple functions:** system-on-chip, line-card or fabric
- **Multiple form factors:** fixed or modular
- **Multiple networks:** Enterprise, Data Center and SP

# Multi-Slice Approach

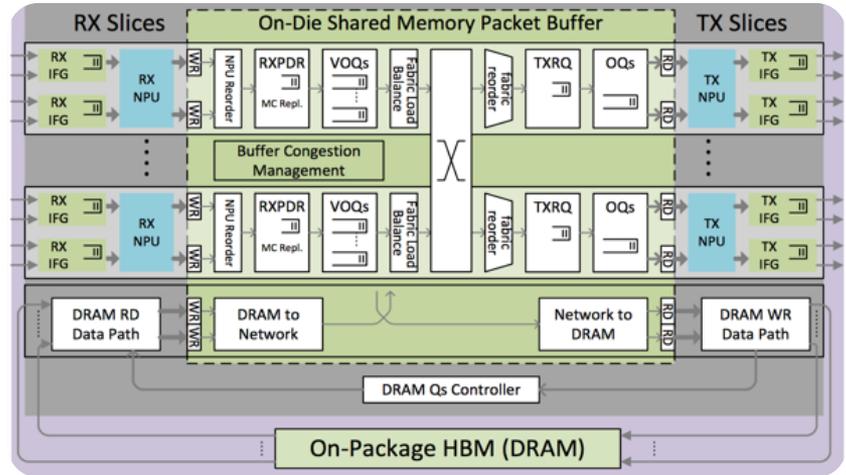
Multi-Slice systems use **multiple full Network Processing Unit (NPU) pipelines** on a single ASIC die.

Each NPU pipeline (or ‘Slice’) **operates independently** and are connected via an integrated crossbar “fabric”, using an integrated Virtual Output Queue (VoQ) buffer architecture, to manage traffic between Slices.

Each Slice has **dedicated RX/TX resources** for parsing, QoS, replication and other ASIC “forwarding” components.



Cisco Silicon One Q200 is an example of a Multi-Slice ASIC design



[www.cisco.com/c/en/us/solutions/service-provider/innovation/silicon-one.html](http://www.cisco.com/c/en/us/solutions/service-provider/innovation/silicon-one.html)

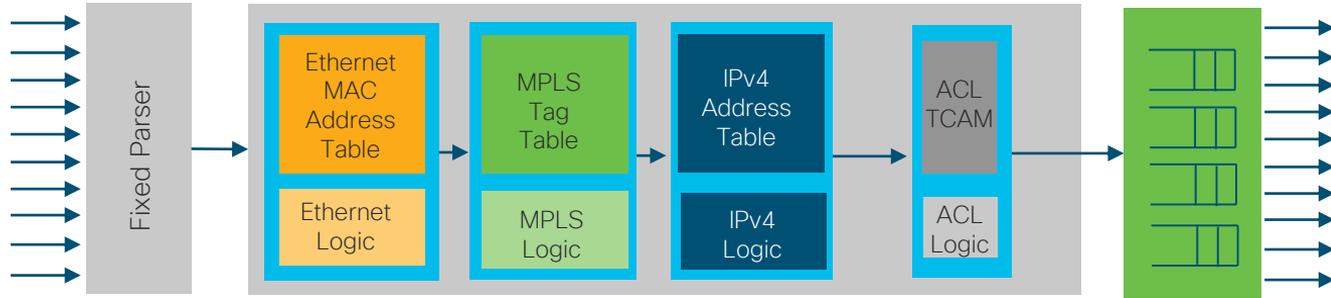
# What is P4?



- **P**rogramming **P**rotocol-Independent **P**acket **P**rocessors
- Programming language designed to allow the definition of data planes
- Open-source, permissively-licensed language
- Designed to be protocol-independent, implementation-independent
- Protocol independence and the abstract language model allow for re-configurability, target-independence

<https://p4.org>

# Fixed vs. Programmable Packet Processing

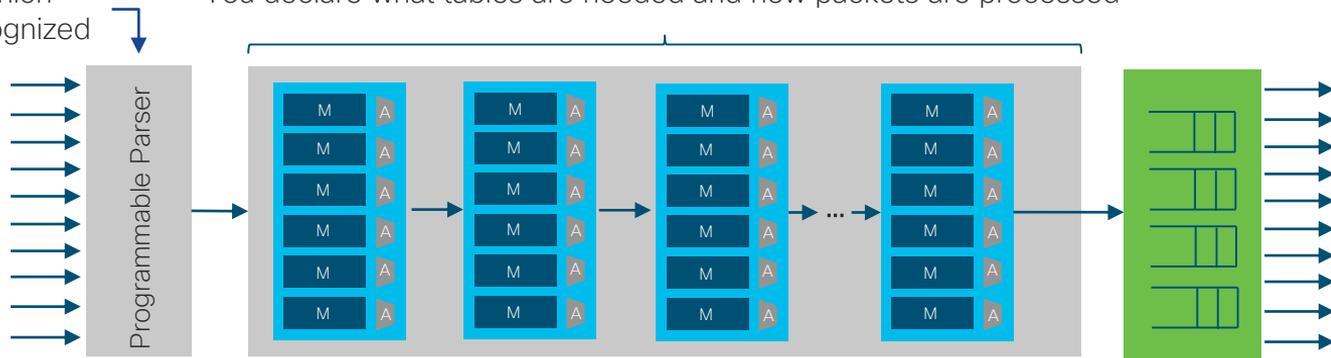


**Fixed Pipeline:** features and functionality are baked-in at design time



You declare which headers are recognized

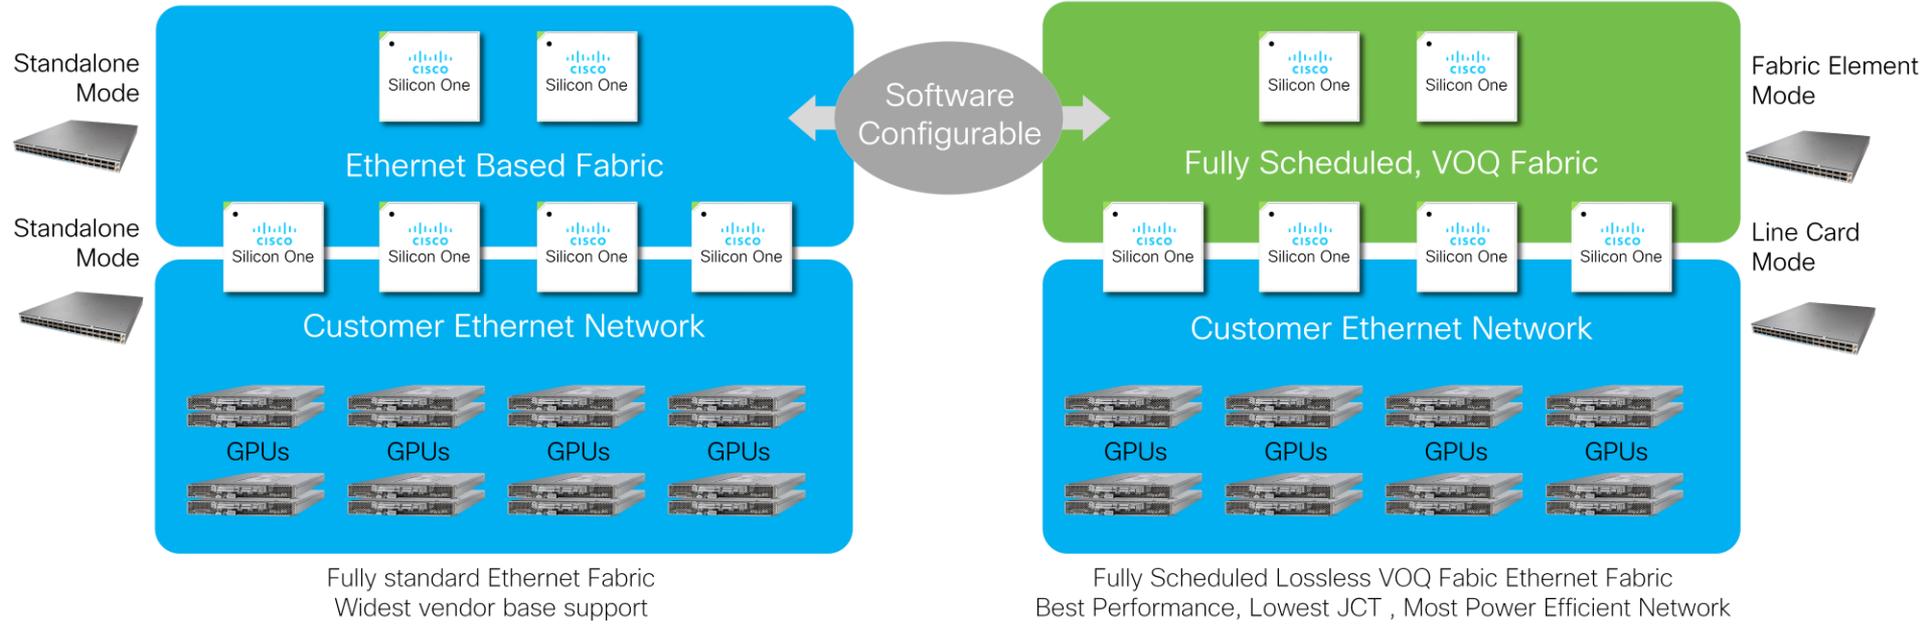
You declare what tables are needed and how packets are processed



**Programmable Pipeline:** all stages identical, customer-defined match-action logic

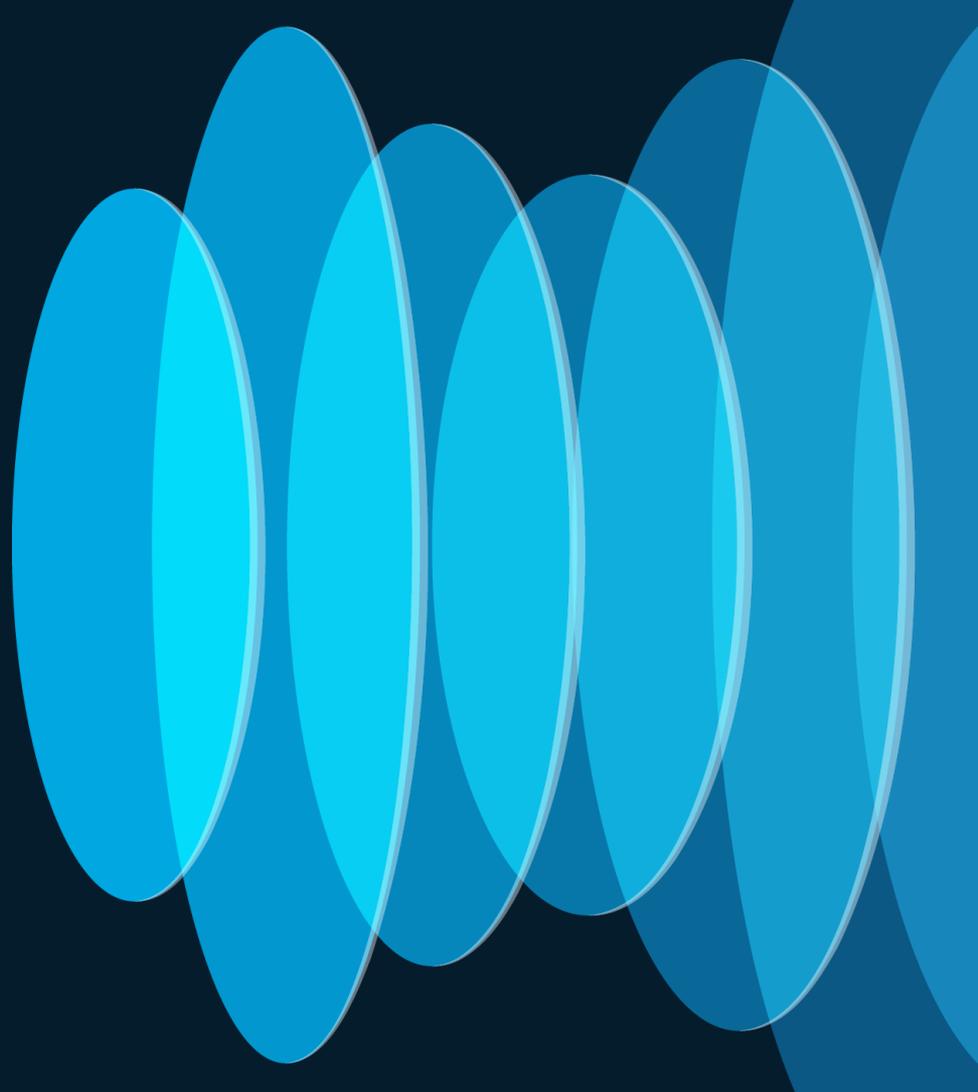


# Cisco Silicon One's Programmability

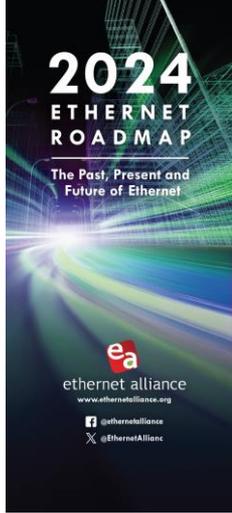
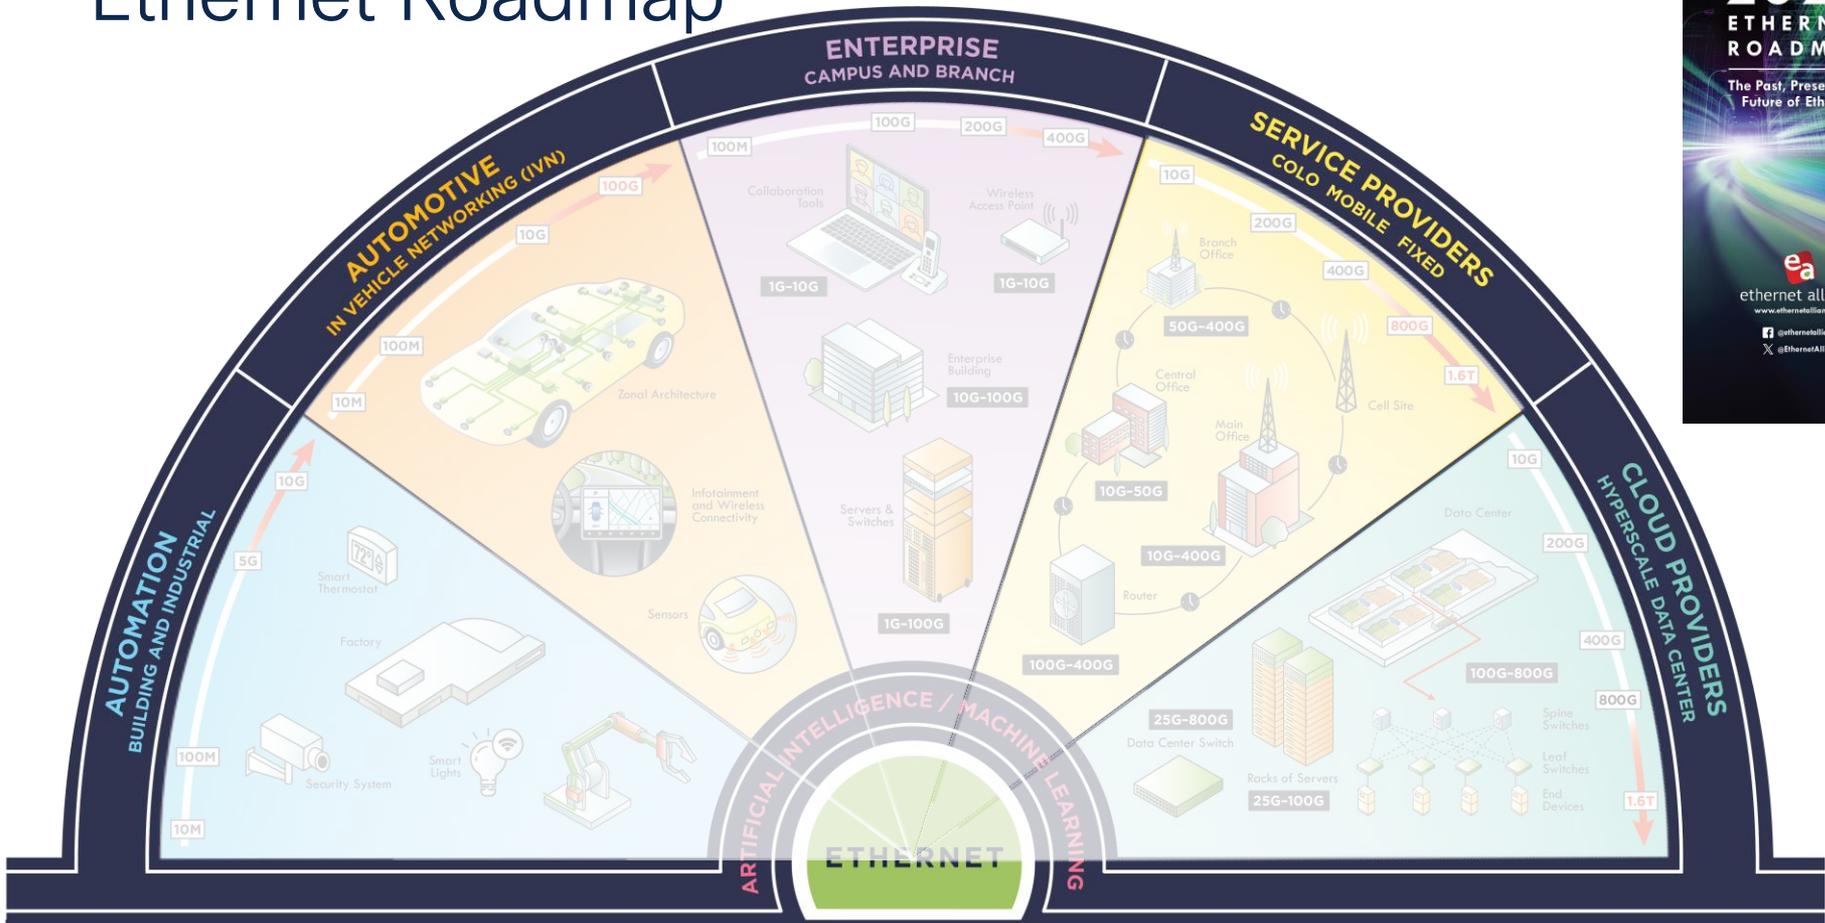


One Architecture, One SDK, One Experience

# Evolution of Networks for ML/AI Workloads



# Ethernet Roadmap



# GenAI Data Center Networks Massive Growth

Single Building (360,000 ft<sup>2</sup>)



Silicon One™  
Q200  
©Cisco 2020

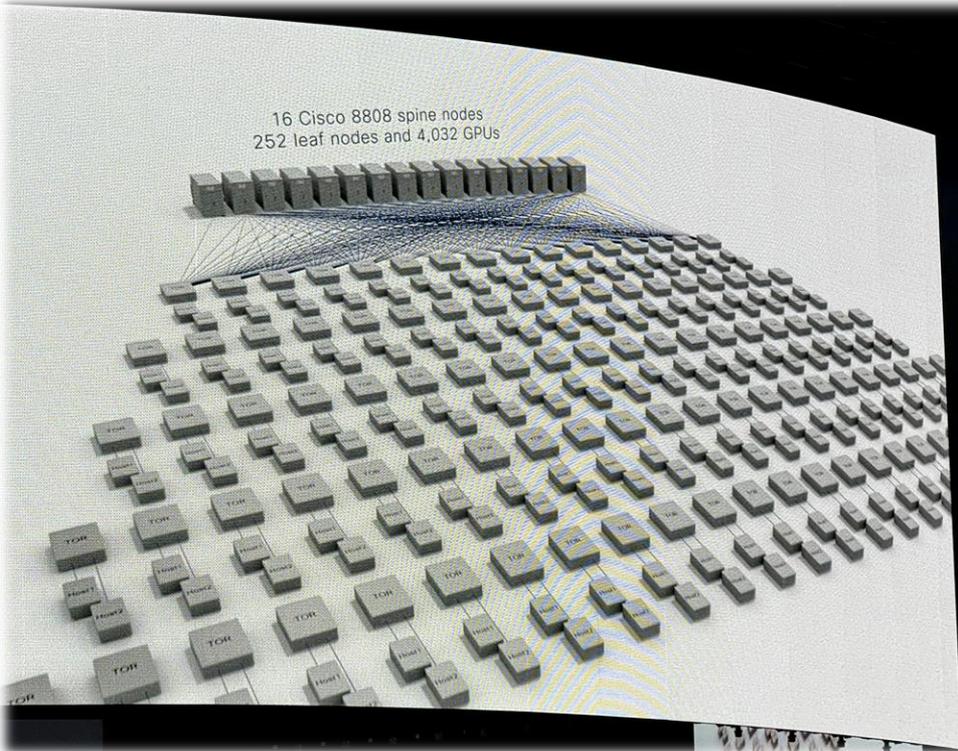
039



DC Location with 5 buildings (6,095,616 ft<sup>2</sup>)

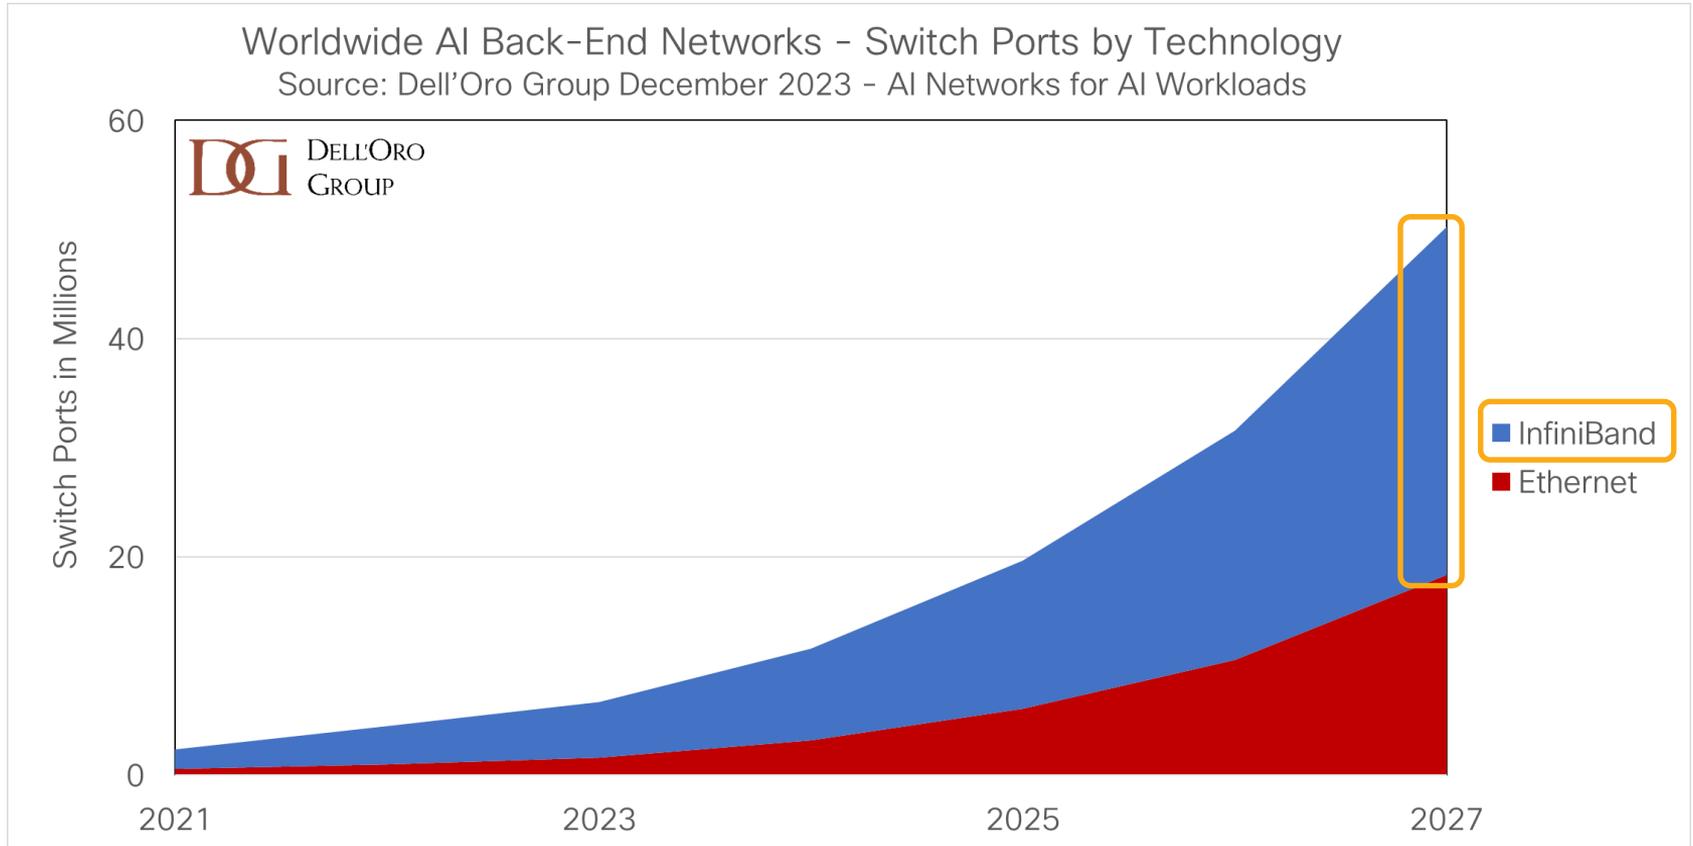


16 Cisco 8808 spine nodes  
252 leaf nodes and 4,032 GPUs



cisco *Live!*

# AI Network Market Trends



# Why InfiniBand?

*For AI training and inference infrastructure among the hyperscalers and cloud builders, Nvidia will tell you plain and simple that the network represents **20 percent of the cluster cost.***

*InfiniBand, explains Nvidia co-founder and chief executive officer Jensen Huang, delivers 20 percent better performance at scale at the same bandwidth than Ethernet, however, so **“InfiniBand is effectively free.”***

The Next Platform: META PLATFORMS IS DETERMINED TO MAKE ETHERNET WORK FOR AI - September 26, 2023  
<https://www.nextplatform.com/2023/09/26/meta-platforms-is-determined-to-make-ethernet-work-for-ai/>

# Ethernet vs. InfiniBand?

ChatGPT

You:

What are the advantages of Ethernet vs. InfiniBand?

ChatGPT

Certainly! Ethernet and InfiniBand are two different types of networking technologies, each with its own advantages and use cases. Here's a comparison of the advantages of Ethernet over InfiniBand:

- 1. Cost:** Ethernet is generally more cost-effective compared to InfiniBand. Ethernet hardware, such as switches and network interface cards (NICs), tend to be more widely available and affordable due to the widespread adoption of Ethernet in various industries. In contrast, InfiniBand equipment can be more specialized and costly.
- 2. Compatibility:** Ethernet is a widely adopted networking standard and is supported by a wide range of devices and operating systems. This broad compatibility makes it easier to integrate Ethernet networks into existing infrastructure and work with different types of hardware and software. InfiniBand, while powerful, may require specific drivers and configurations, which can limit compatibility in certain environments.
- 3. Scalability:** Ethernet networks are highly scalable and can support large-scale deployments with thousands of nodes. Modern Ethernet standards, such as 10 Gigabit Ethernet (10GbE), 25 Gigabit Ethernet (25GbE), and 100 Gigabit Ethernet (100GbE), offer high bandwidth and low latency, making them suitable for demanding workloads in data centers and high-performance computing (HPC) environments. InfiniBand also offers high scalability, but Ethernet's widespread adoption and compatibility make it a more flexible choice for many applications.
- 4. Ease of Management:** Ethernet networks typically have mature management tools and protocols, making them easier to manage and troubleshoot compared to InfiniBand. Ethernet switches often support standard management protocols such as SNMP (Simple Network Management Protocol) and CLI (Command-Line Interface), which facilitate network monitoring and configuration. InfiniBand management may require specialized tools and expertise, which can add complexity to network administration.
- 5. Flexibility:** Ethernet supports a wide range of network topologies, including star, tree, and mesh configurations, as well as point-to-point and point-to-multipoint connections. This flexibility allows Ethernet networks to adapt to different deployment scenarios and application requirements. InfiniBand also offers flexibility in terms of network topology, but Ethernet's widespread adoption and compatibility make it a more versatile choice for many organizations.

Overall, Ethernet offers advantages in terms of cost-effectiveness, compatibility, scalability, ease of management, and flexibility compared to InfiniBand. However, the choice between Ethernet and InfiniBand depends on specific requirements, performance considerations, and existing infrastructure in a given environment.

# Ethernet vs. InfiniBand?

## Gemini



What are the advantages of Ethernet vs. InfiniBand?



Ethernet holds the upper hand in several areas when compared to InfiniBand:

- **Cost:** Ethernet is significantly cheaper than InfiniBand. This covers the cost of the cables, switches, and NICs (Network Interface Cards) needed to set up the network.
- **Compatibility:** Ethernet is universally compatible with almost any device you can plug into a network. This makes it far easier to integrate new devices into an existing Ethernet network. InfiniBand, on the other hand, may require specific hardware and software configurations that limit its compatibility.
- **Security & Management:** Ethernet has a wider range of robust security protocols and comprehensive management capabilities that have been developed over decades. This makes it more favorable for organizations that prioritize data security and efficient network administration.

While InfiniBand boasts superior performance, Ethernet's advantages often make it the more practical choice for many situations.



tell me more



Certainly! Here's a deeper dive into the advantages of Ethernet over InfiniBand:

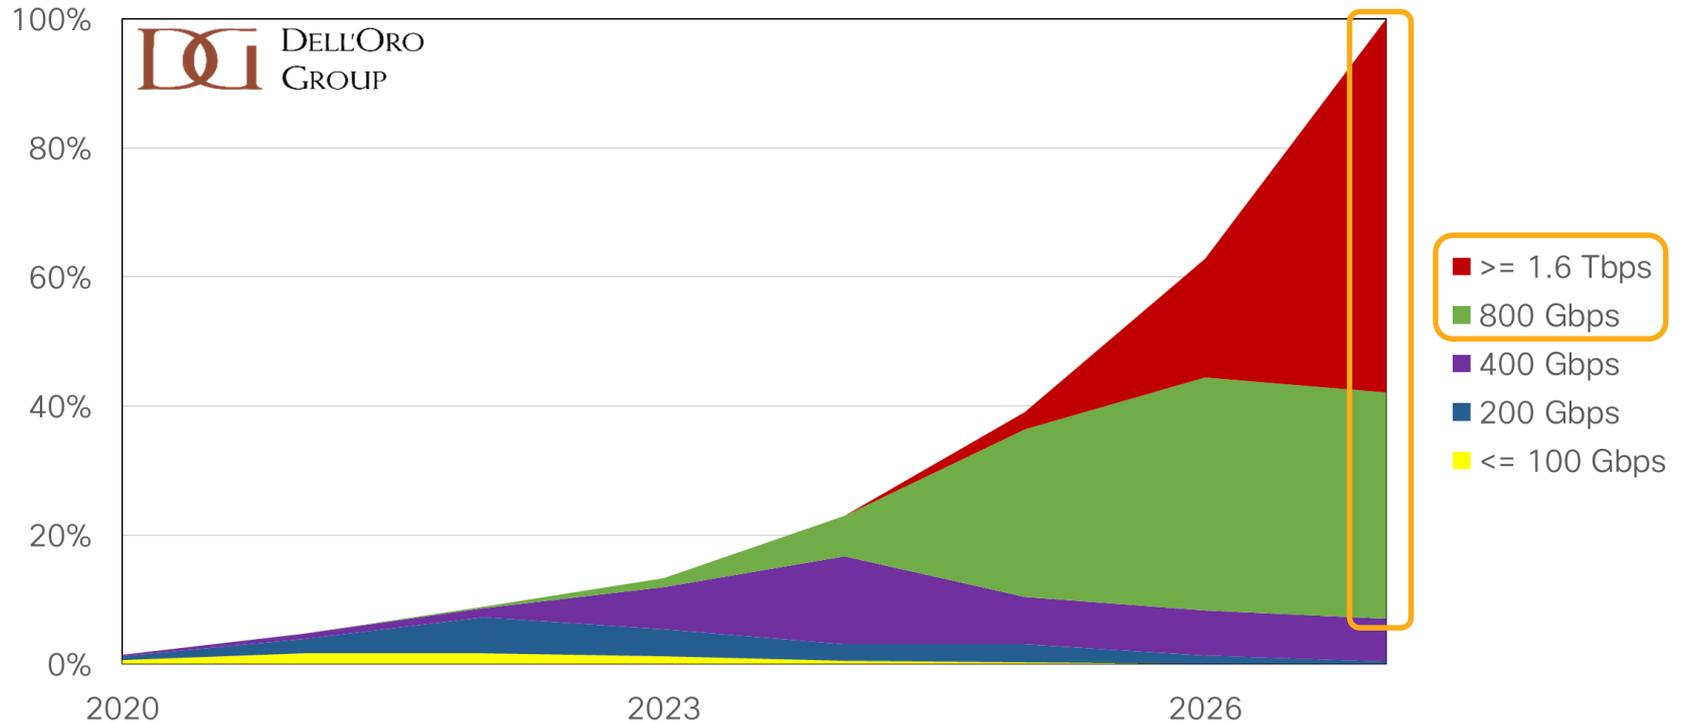
- **Widespread Adoption:** Ethernet is everywhere. Homes, businesses, and data centers all rely on Ethernet for their networking needs. This ubiquity translates to a vast ecosystem of manufacturers producing compatible equipment at competitive prices. You're almost guaranteed to find what you need readily available and at a reasonable cost.
- **Simplicity:** Ethernet is a mature technology with well-established standards and protocols. This makes it relatively easy to set up, manage, and troubleshoot an Ethernet network. There's a wealth of documentation and a large pool of networking professionals familiar with Ethernet.
- **Scalability:** Don't be fooled by Ethernet's reputation for being for simpler networks. Modern Ethernet can handle massive data centers with thousands of devices. Speeds continue to climb, with the latest versions reaching speeds of 400 Gbps and beyond. While InfiniBand might boast higher theoretical maximums, Ethernet offers excellent scalability for most applications.
- **Security:** Over the years, robust security protocols have been developed for Ethernet. These include measures like access control lists, VLANs (Virtual Local Area Networks), and encryption standards. This makes it suitable for organizations that prioritize data security.
- **Management Tools:** A wide range of software tools are available for managing Ethernet networks. These tools can help with tasks like network monitoring, configuration, and troubleshooting. This simplifies network administration and helps ensure smooth operation.

In conclusion, while InfiniBand offers unmatched speed and low latency, Ethernet's affordability, ease of use, vast compatibility, and robust security features make it the go-to choice for most networking needs. If your application demands the absolute pinnacle of performance, InfiniBand might be the way to go, but for most users, the advantages of Ethernet make it the clear winner.

# AI Network Market Trends

Worldwide AI Back-End Networks - Ethernet Switch Port Growth by Speed

Source: Dell'Oro Group December 2023 - AI Networks for AI Workloads

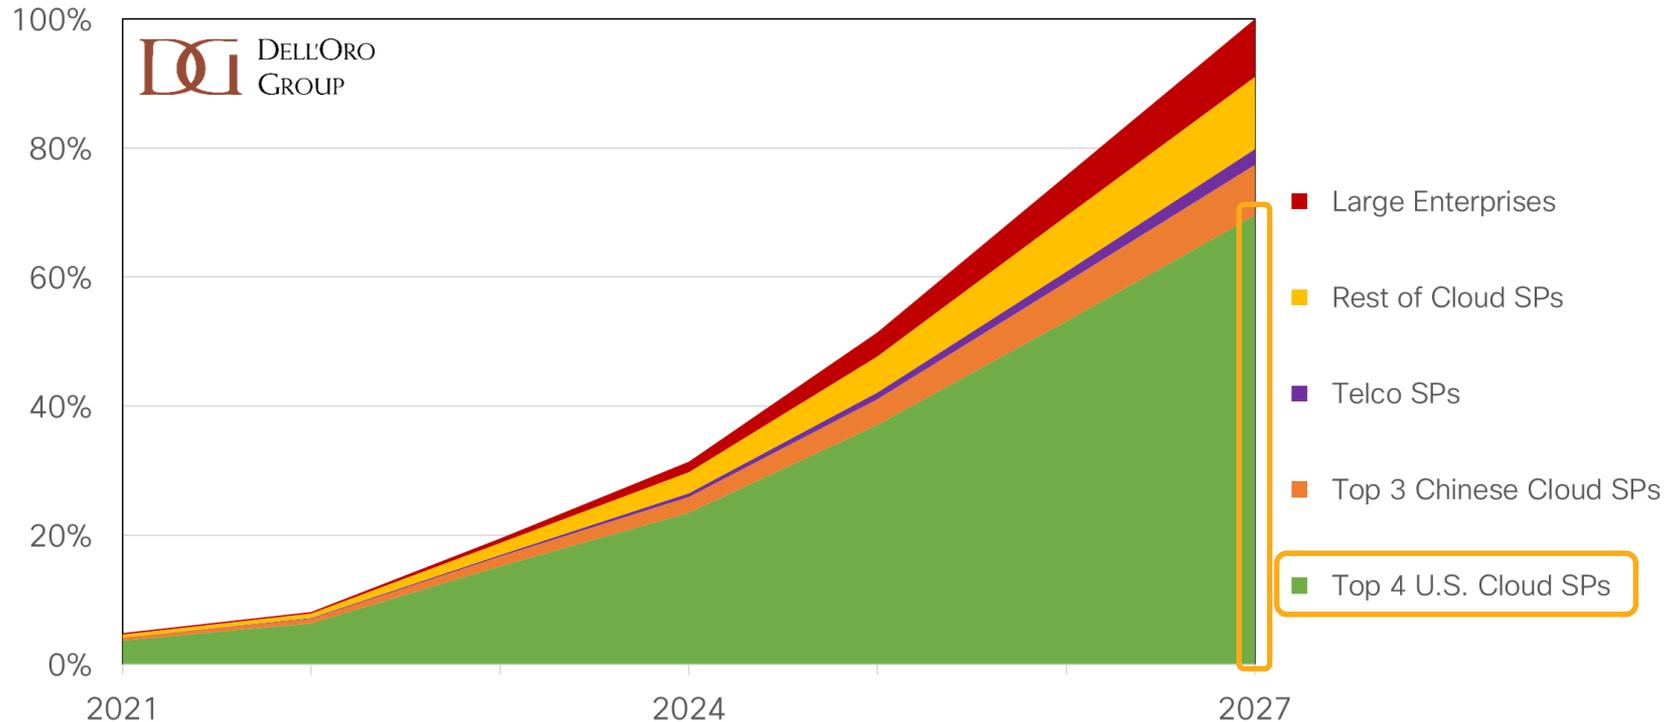


NOTE: The graph above shows percentage of 2027 total ports.

# AI Network Market Trends

## Worldwide AI Back-End Networks - Revenue Growth by Customer Segment

Source: Dell'Oro Group December 2023 AI Networks for AI Workloads



NOTE: The graph above shows percentage of 2027 total ports.

# Cisco Silicon One

Flexibility  
via Programmability

Cisco **Silicon One**

---



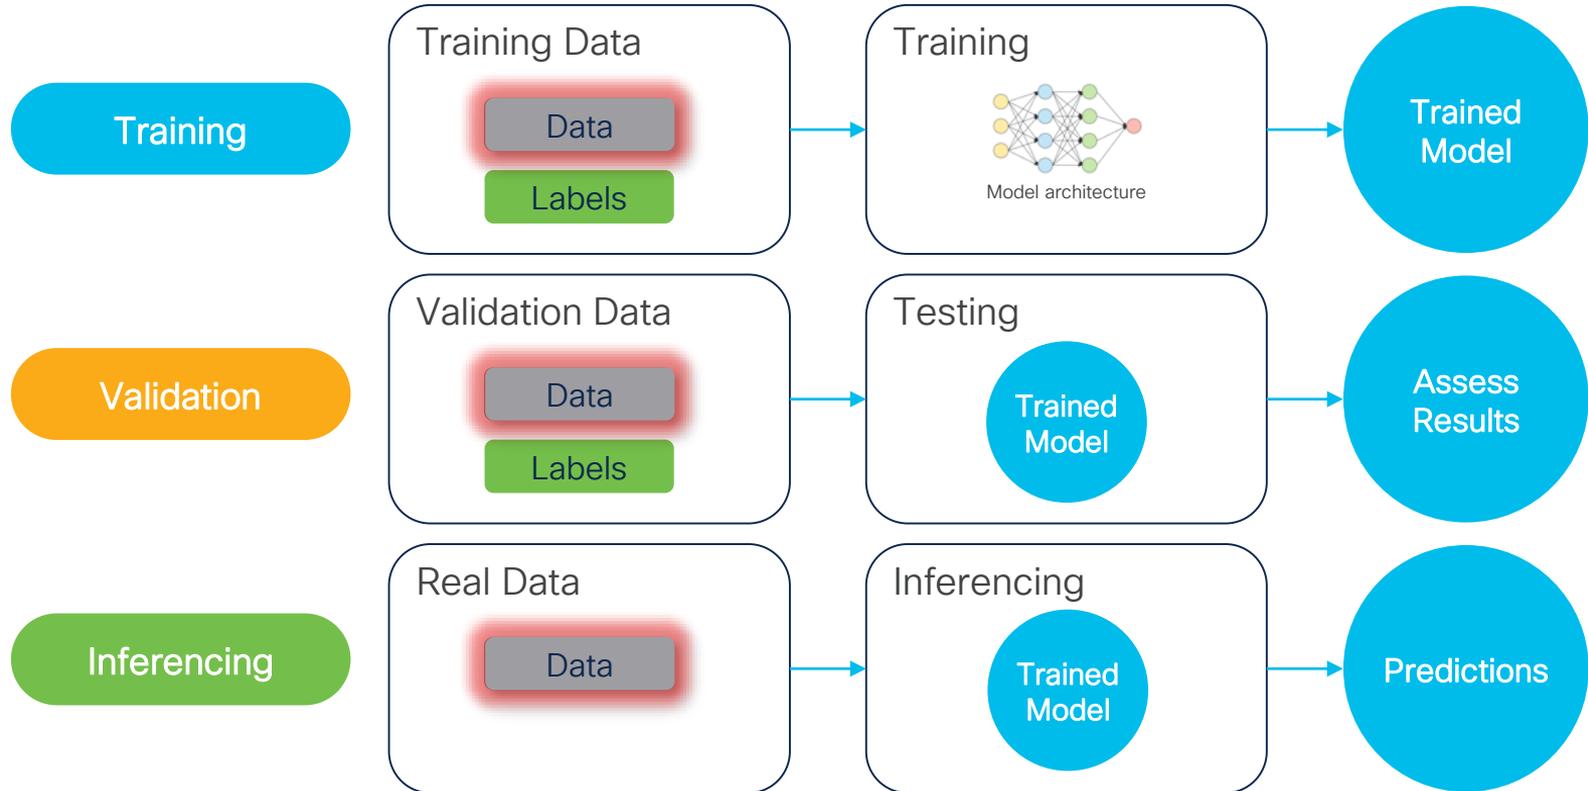
---

One Architecture, One SDK, One Experience

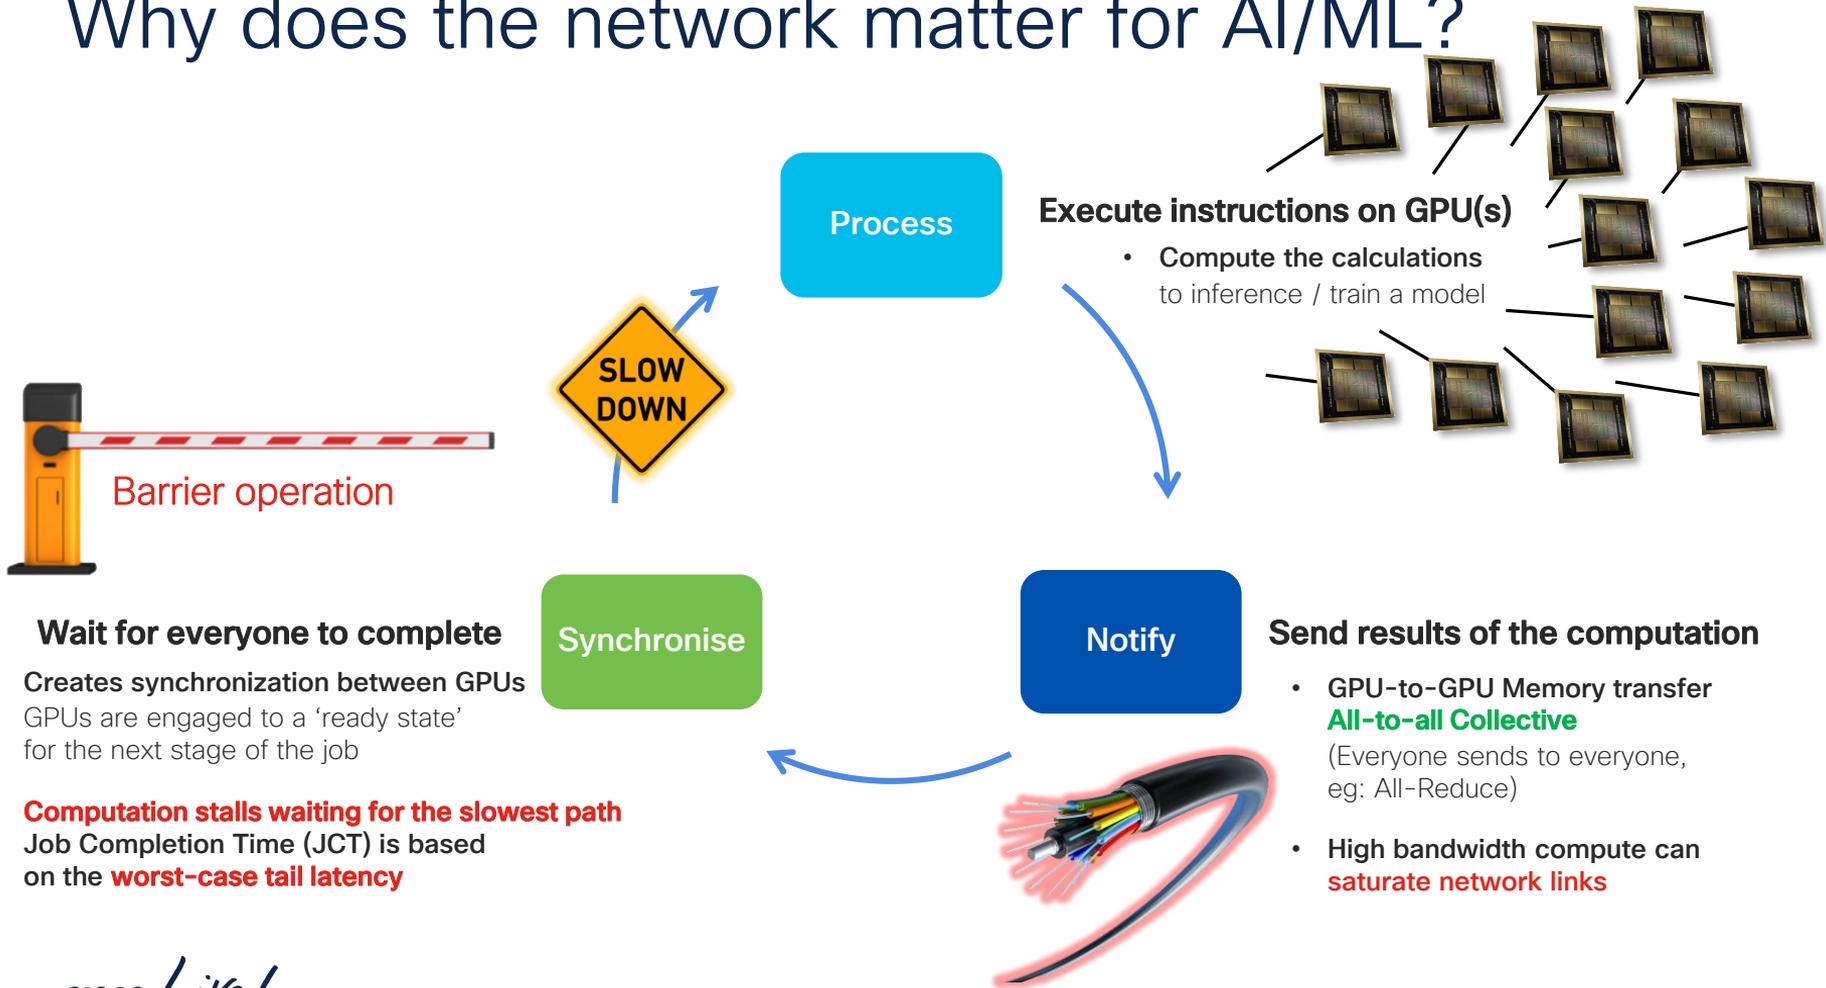
**cisco** *Live!*

# Bringing It All Together – *Networking for ML/AI*

# ML/AI- Stages



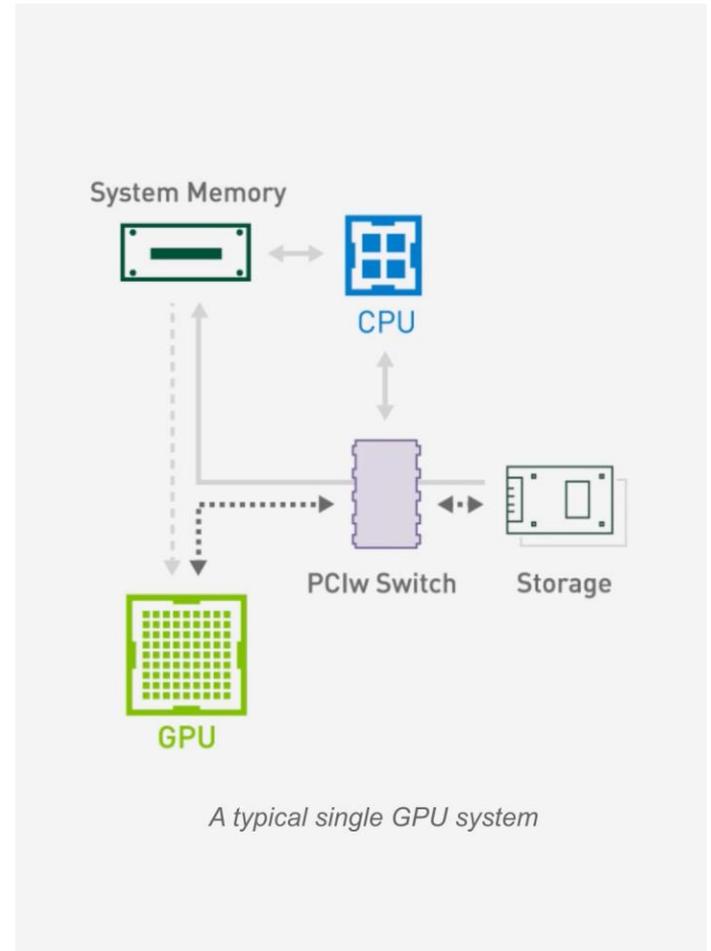
# Why does the network matter for AI/ML?



# Keeping the GPU Busy

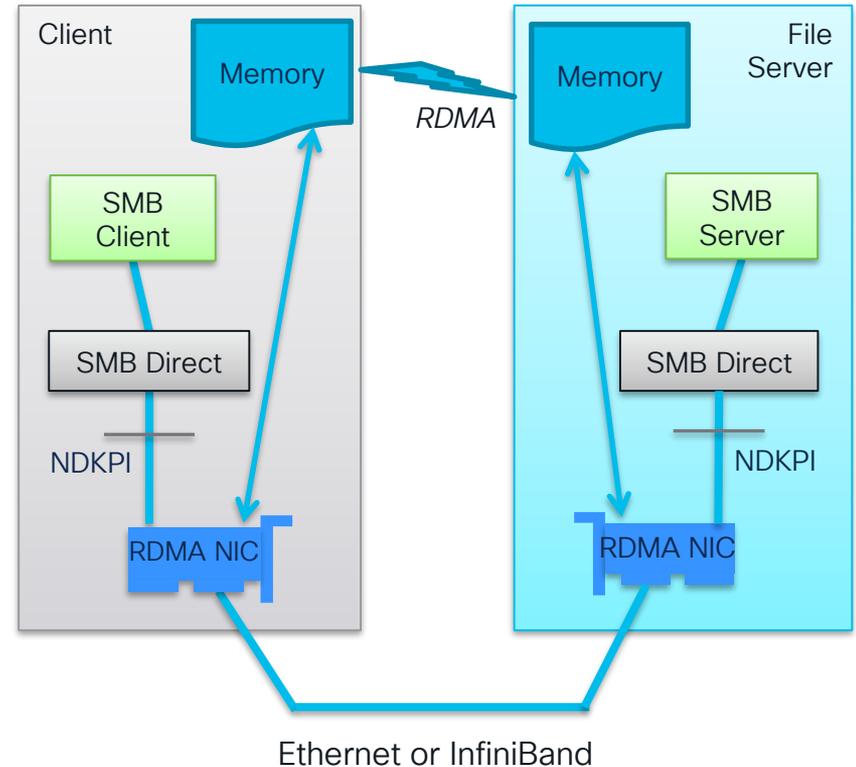
## The Network is Crucial for Maximizing GPU Acceleration

- Data sets continue to grow, increasing the I/O bottleneck on the data path between storage and GPU memory
- To maximize the acceleration that a GPU can provide, it is imperative that data is always available for the GPU to operate upon
- The system's components, CPU, memory and storage must support bandwidth, speeds and latency that do not cause significant GPU idle time



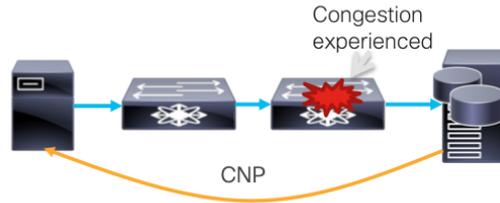
# RDMA – Remote Direct Memory Access

- Allows application software to communicate directly with the hardware (RDMA NIC)
- Bypasses OS stack
- RDMA delivers, low latency, high throughput, zero copy capabilities
- RDMA Hardware Technologies
  - **RoCEv2: RDMA over Converged Ethernet**
  - iWARP: RDMA over TCP/IP
  - Infiniband



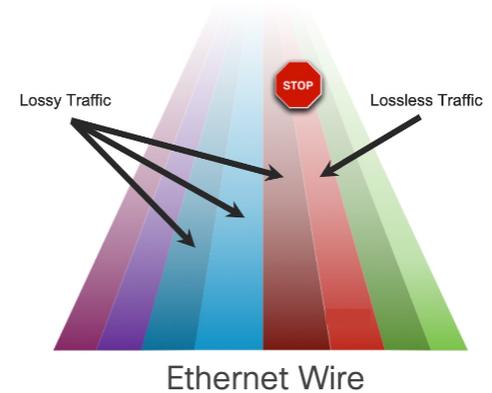
# ECN and PFC

- IP Explicit Congestion Notification (ECN) is used for congestion notification.
- ECN enables end-to-end congestion notification between two endpoints on IP network
- ECN uses 2 LSB of Type of Service field in IP header
- In case of congestion, ECN gets transmitting device to reduce transmission rate using Congestion Notification Packet (CNP) without pausing traffic.



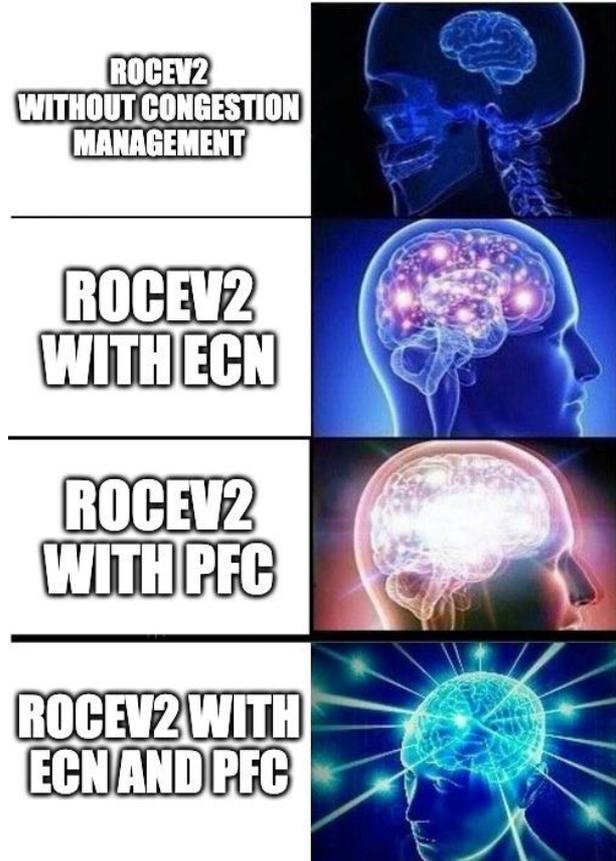
## Flow Control Mechanism – 802.1Qbb

- A.k.a "Lossless Ethernet"
- PFC enables Flow Control on a Per-Priority basis
- PFC is also called Per-Priority-Pause
- Therefore, we have the ability to have lossless and lossy priorities at the same time on the same wire
- Allows traffic to operate over a lossless priority independent of other priorities
- Other traffic assigned to other priority will continue to transmit and rely on upper layer protocols for retransmission



# Data Center Quantized Congestion Notification

- IP ECN or PFC cannot alone provide a valid Congestion Management framework
  - The two of them together provide the desired result of having lossless RDMA communications across Ethernet networks (this is called DCQCN)
  - The requirements are:
    - Ethernet devices compatible with both techniques
    - Proper configurations applied
- 
- Elephant v. Mice flows
  - AFD and Smart Buffers



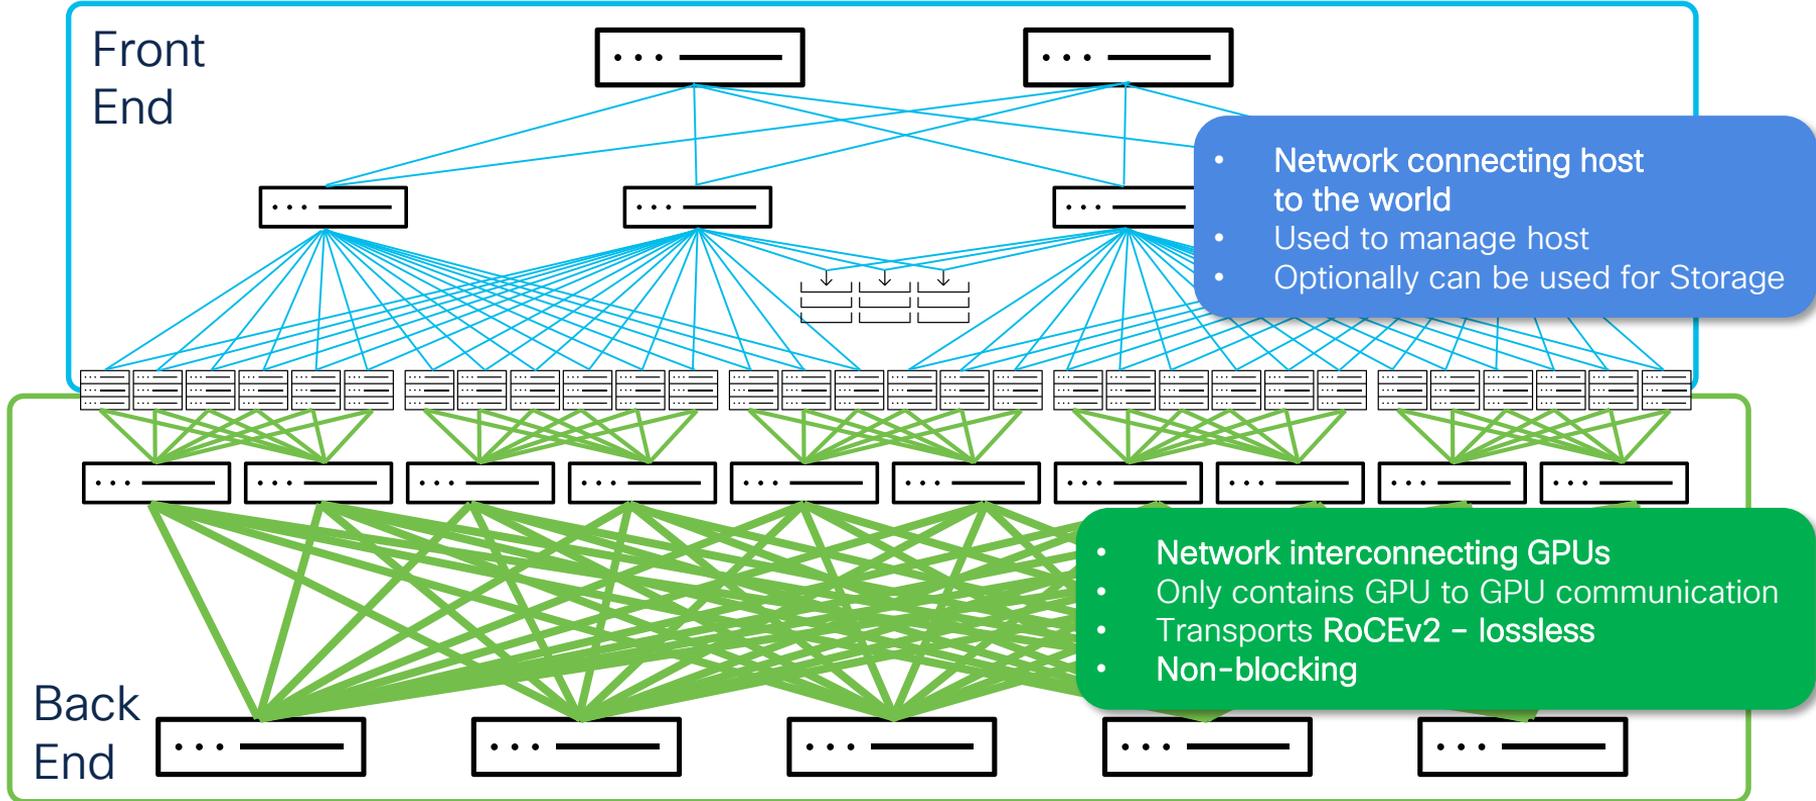
# Non-blocking Network

**TODAY!**

3:00 – 4:30pm

MORE INFO AT ...

*Network Best Practices for Artificial Intelligence  
(AI/ML) Data Centre - BRKDCN-2921*



# Ultra Ethernet Consortia (UEC)

Deliver an Ethernet based open, interoperable, high performance, full-communications stack architecture to meet the growing network demands of AI & HPC at scale

THE NEW ERA  
NEEDS A  
NEW NETWORK

*Ultra* ~~E~~thernet

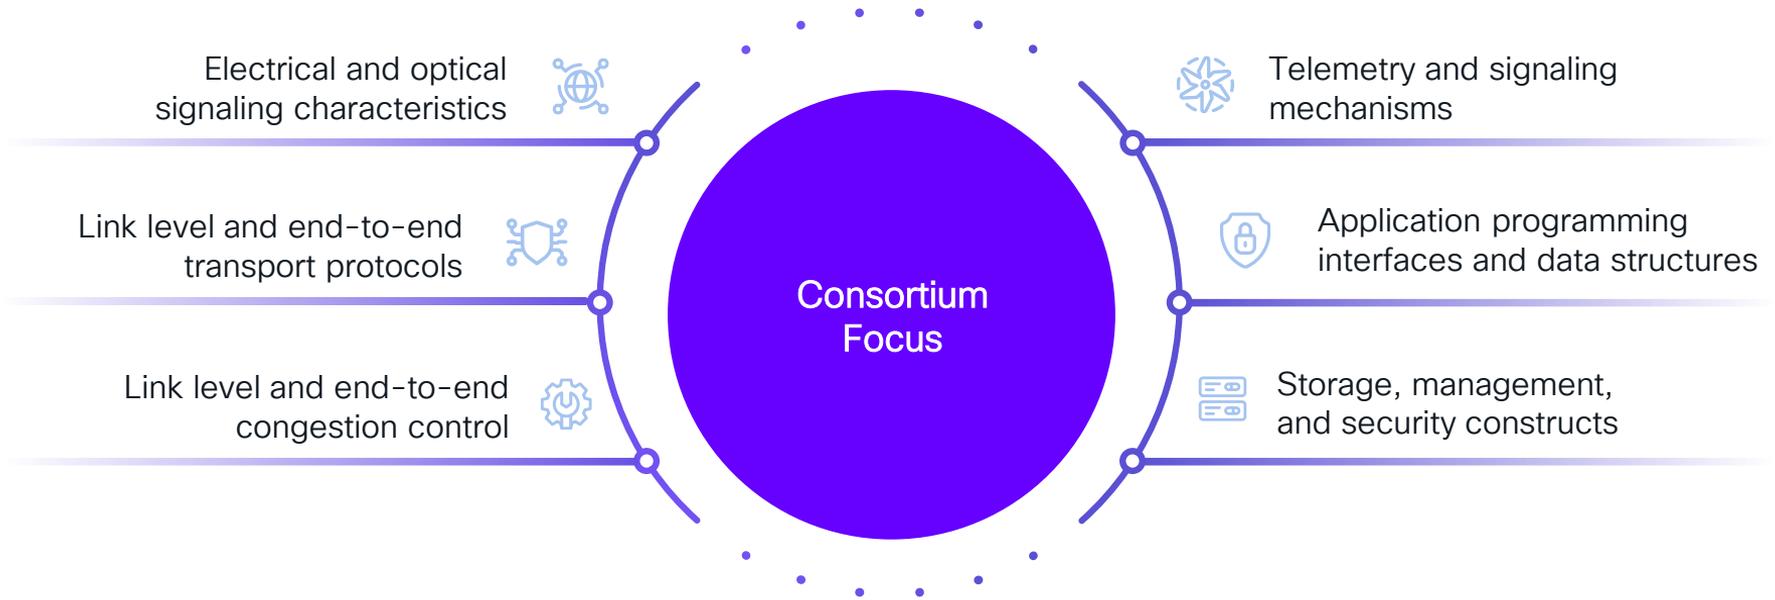
As *performant* as a  
supercomputing  
interconnect

As *ubiquitous* and *cost-  
effective* as Ethernet

As *scalable* as a cloud  
data center

# UEC Focus

Open specifications, APIs, source code for optimal performance of AI and HPC workloads at scale.



# UEC Elements

## Working Groups

Physical Layer

Link Layer

Transport Layer

Software Layer

Storage

Compliance

Management

Performance & Debug

## Key Features

L2 & L3 Header  
compression

Link Level Reliability

Multi-path packet  
spraying

Comprehensive  
congestion control

Multiple packet  
ordering options

Multiple transport  
services

Data processing  
offload to network

Security for huge  
scale and extreme  
speeds

Want to read more? - Visit <https://ultraethernet.org/>

# Building Networks for ML/AI Workloads

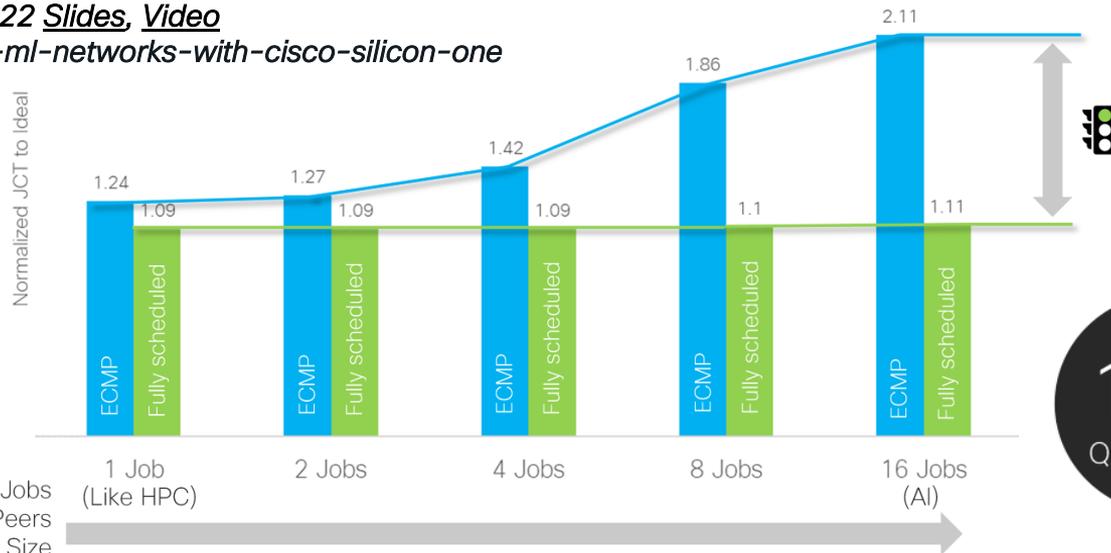
## Fully Scheduled Fabric Benefits

MORE INFO AT ...

*“Evolved Networking: The AI/ML Challenge”;*  
OCP Global Summit October 2022 [Slides](#), [Video](#)  
<https://blogs.cisco.com/sp/building-ai-ml-networks-with-cisco-silicon-one>

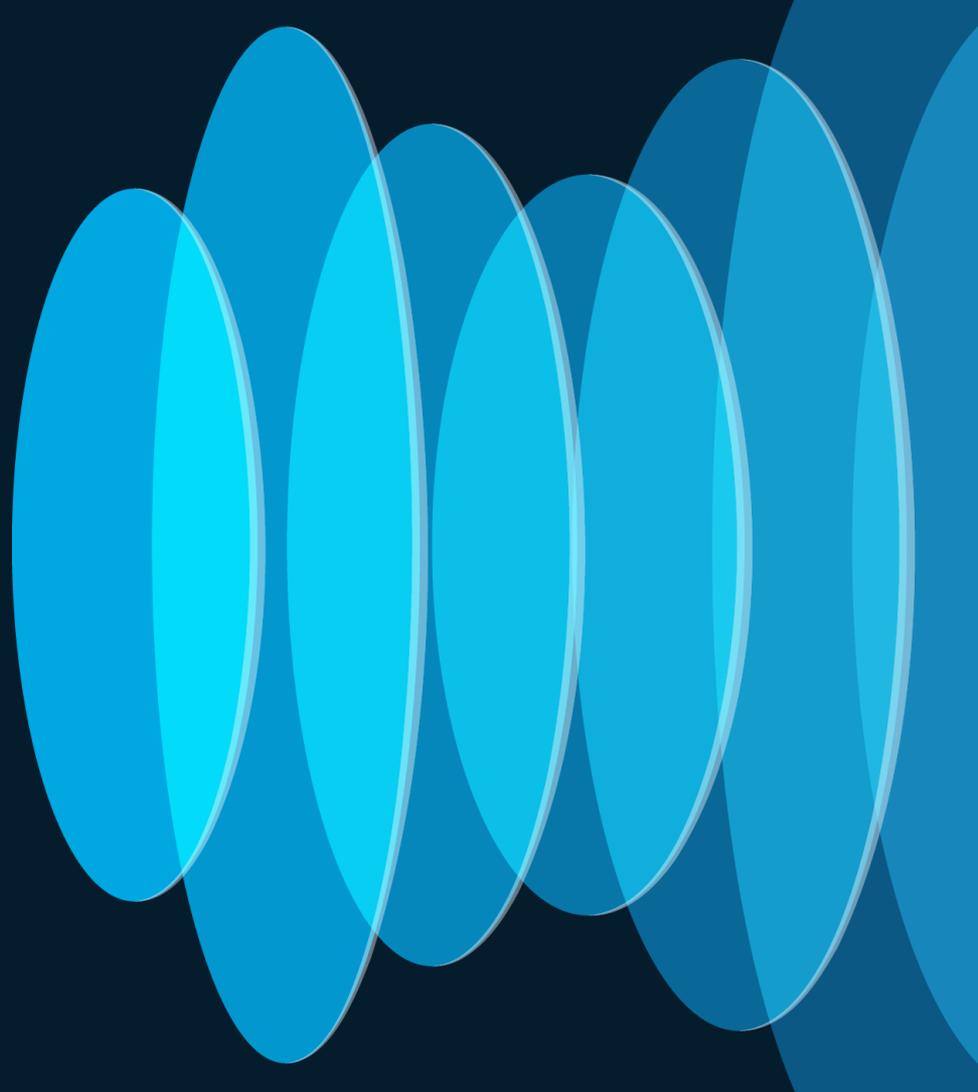


Impact on JCT of Increasing Number of Jobs

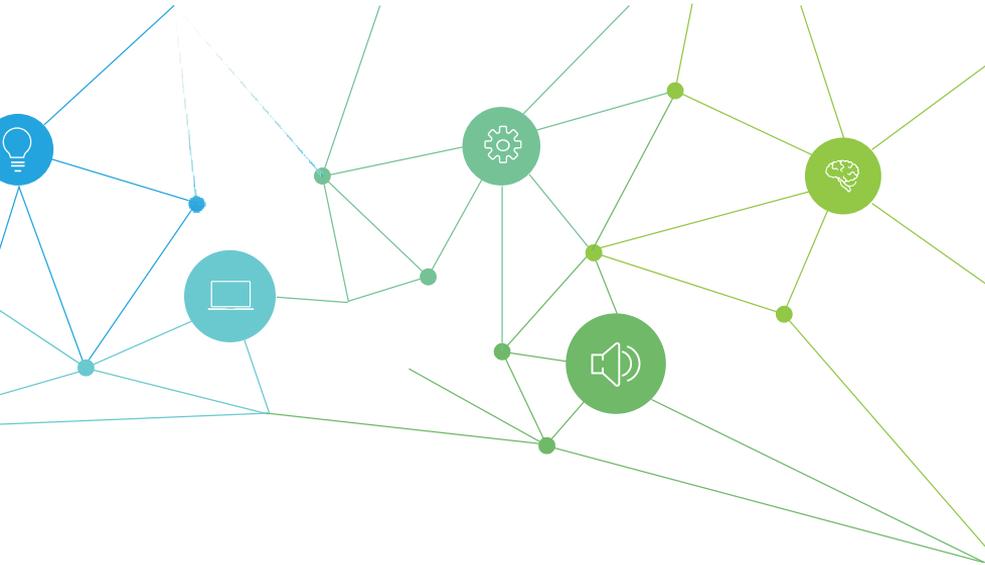


Fully scheduled fabric provides exceptional performance, providing lower job completion time

# Summary – Cisco Silicon for AI



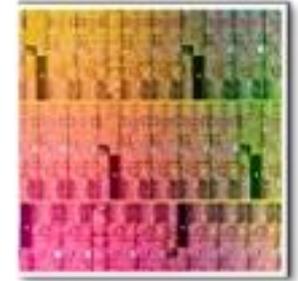
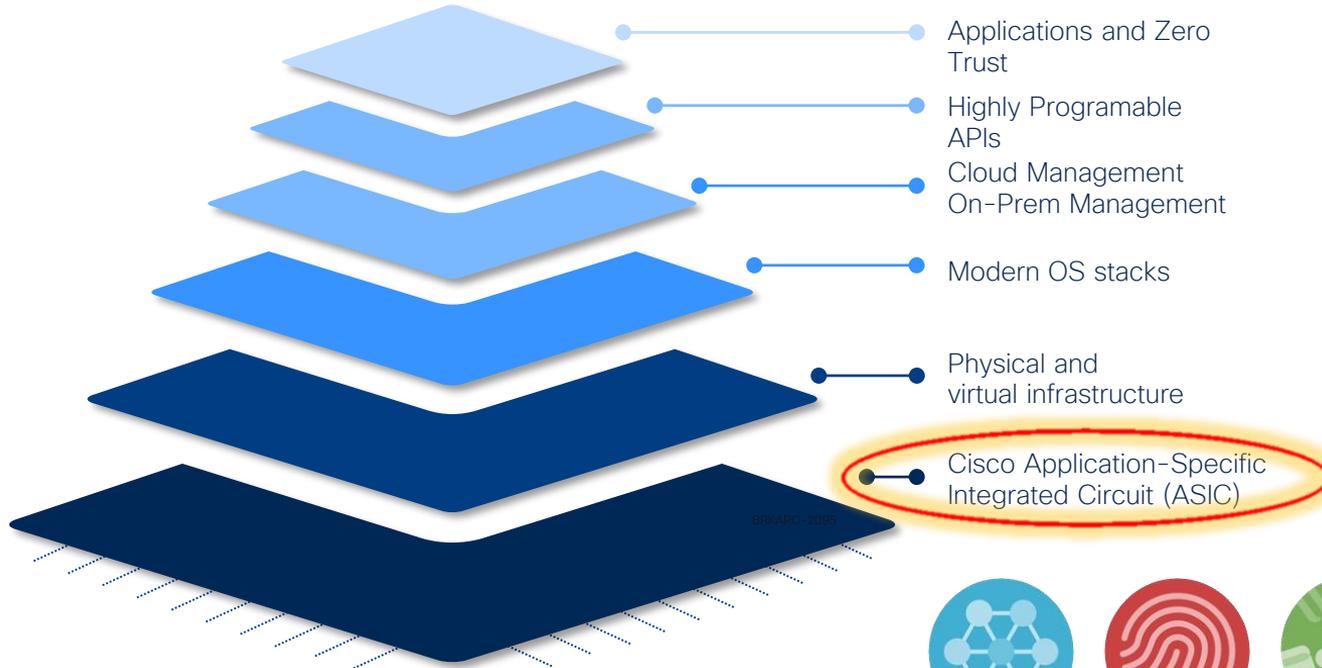
# Cisco and AI



- Cisco is **investing** in AI capabilities
- We have a focus on **creating AI solutions for use by customers**
- We have a focus on **creating solutions that support AI workloads**

# Cisco Silicon for AI

## Foundational Elements to Support AI Growth



Best-In-Class  
Hardware



Secure  
Networks



Sustainable



Highly  
Programmable



Cloud Ready

CISCO *Live!*

Cisco Networking Hardware



Networking the World

Cisco Networking Hardware



AI in the Network

Cisco Networking Hardware



Access to the Network

Cisco Networking Hardware



Heart of the Network

Cisco Networking Hardware



Securing the Network

Cisco Networking Hardware



Brains of the Network

Cisco Networking Hardware



Shielding the Network

Cisco Networking Hardware



Powering the Network

# YOUR NETWORK IS OUR LIFE'S WORK



**Silicon One™ G200**  
©Cisco 2023

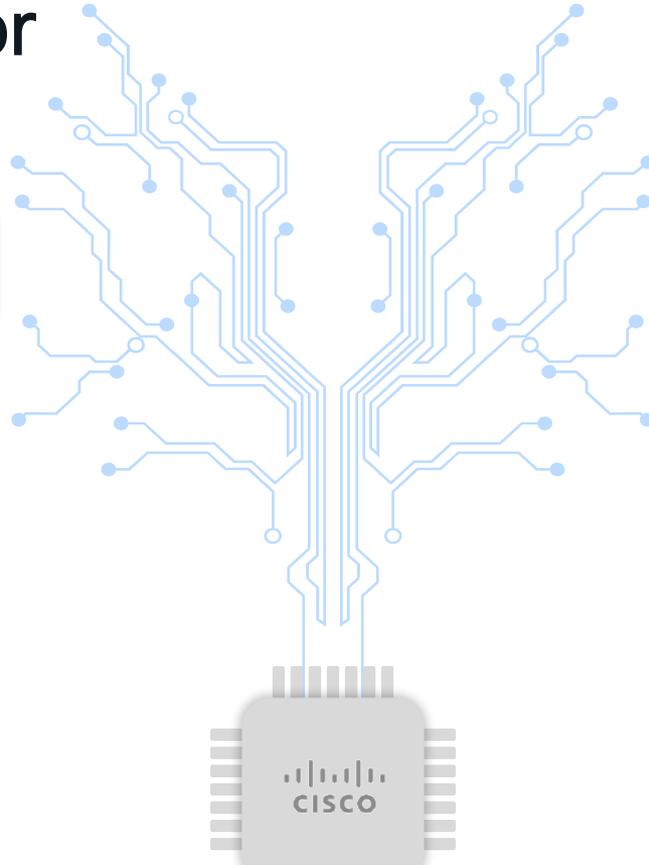
**Silicon One™ P100**  
©Cisco 2021

**Silicon One™ Q201**  
©Cisco 2020



# How Did We Do? Cisco Silicon for AI

Do You Have a Better  
Understanding ...



... of how **ASICs and  
Network Hardware** are  
Designed and Built ...

... of why **Hardware  
Functionality and Flexibility**  
are Key for AI Solutions ...

... and how You can  
**Leverage Cisco's Latest  
Flexible Hardware and  
Advanced Capabilities** in  
Your Own Network Designs?



# Complete Your Session Evaluations



Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to **win 1 of 5 full conference passes** to Cisco Live 2025.

---



**Earn 100 points** per survey completed and compete on the Cisco Live Challenge leaderboard.

---



Level up and earn **exclusive prizes!**

---



Complete your surveys in the **Cisco Live mobile app.**

# Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at [www.CiscoLive.com/on-demand](https://www.CiscoLive.com/on-demand)

See the “By Way of Introduction ...” slides for contact details.



The bridge to possible

# Thank you

CISCO *Live!*

#CiscoLive