# 

Organized by



#### **Speaker**

Topic Open source hardware for AI and its impact on the data center

#### Steve Helvie | VP of Channel

steve@opencompute.org





## Today's Topics







**Change in Hardware** 

#### AI Accelerators

#### Data Center Density

Advanced Cooling

#### **Predictive Server Modelling**

Data Center Software



11011010 10101001 101011010

#### **Circular Economy**

Reuse at Scale

# Cost per Human Genome



# "Full and free exchange of all sequence data"

https://www.genome.gov/about-genomics/fact-sheets/Sequencing-Human-Genome-cost

# "Publicly accessible database within 24 hours "

### Architects of Intelligence

The truth about AI from the people building it

Yoshua Bengio Stuart Russell Ceoffrey Hinton Nick Bostrom Yann LeCun Fei-Fei Li Demis Hassabis Andrew Ng Rana el Kallouby Ray Kurzweil Daniela Rus James Manyika Gary Marcus Barbara Grosz Judea Pearl James Manyika Gary Marcus Barbara Grosz Judea Pearl Joff Dean Daphne Koller David Ferrucci Rodney Brooks Cynthia Breazeal Josh Tenenbaum Oren Etzioni Bryan Johnson

Interviewed by: New York Times Bestselling Author of Rise of the Robots

Martin Ford

# "Architects of Intelligence"

By Martin Ford



### The AI J-Curve



Time

Performance

 $\infty$ 

Productivity

**Compute Project** 



.......

What stakeholders (mistakenly) expect

What stakeholders can expect with proper execution

What actually happens in most cases







### January 2010

**COMPANIES > FACEBOOK** 

# Facebook to Build Its Own Data Centers

Facebook has decided to begin building its own data centers, and may announce its first facility as soon as tomorrow. The fast-growing social network has previously leased server space from wholesale data center providers.

Rich Miller | Jan 20, 2010

Rich Miller | Jan 20, 2010

tomorrow. The rast-growing social network has p leased server space from wholesale data center providers.



#### What can we remove from the system?

Can we raise operating temperatures and have the servers survive?

Can we increase relative humidity operational ranges to make the system much more efficient?

Do we need a centralized power supplies?





# Open Compute Project

A collaborative community focused on redesigning hardware technology to efficiently support the growing demands on compute infrastructure.

![](_page_8_Picture_4.jpeg)

![](_page_8_Picture_5.jpeg)

#### **Goldman Sachs**

![](_page_8_Picture_7.jpeg)

![](_page_8_Picture_8.jpeg)

![](_page_8_Picture_9.jpeg)

![](_page_9_Picture_0.jpeg)

![](_page_9_Picture_1.jpeg)

# 150+ companies189 contributions6K engineers

![](_page_9_Picture_3.jpeg)

![](_page_9_Picture_4.jpeg)

![](_page_10_Picture_0.jpeg)

# Our Projects

![](_page_10_Figure_2.jpeg)

https://www.opencompute.org/projects

![](_page_10_Picture_4.jpeg)

#### Advanced Cooling

![](_page_10_Picture_8.jpeg)

Security

![](_page_10_Picture_10.jpeg)

![](_page_10_Picture_11.jpeg)

![](_page_11_Picture_0.jpeg)

![](_page_11_Figure_1.jpeg)

Efficiency

Scale

Impact

Openness

![](_page_11_Picture_7.jpeg)

![](_page_11_Picture_8.jpeg)

![](_page_11_Picture_9.jpeg)

![](_page_12_Picture_0.jpeg)

....gratuitous differentiation

![](_page_13_Picture_0.jpeg)

### Tool less OCP vanity free open source cubby servers

![](_page_13_Picture_2.jpeg)

![](_page_13_Picture_3.jpeg)

![](_page_13_Picture_4.jpeg)

![](_page_13_Picture_5.jpeg)

![](_page_13_Picture_6.jpeg)

![](_page_13_Picture_7.jpeg)

# Open Rack Unit Innovations Server Fans "Cube Law"

40 mm Fans

![](_page_14_Picture_2.jpeg)

#### <sup>1</sup>/<sub>2</sub> the speed 7/8<sup>th</sup> energy reduction

![](_page_14_Picture_4.jpeg)

EIA "U"

![](_page_14_Picture_6.jpeg)

#### 80 mm Fans

![](_page_15_Figure_1.jpeg)

https://www.opencompute.org/events/past-summits

![](_page_16_Figure_0.jpeg)

https://www.opencompute.org/events/past-summits

![](_page_16_Picture_7.jpeg)

![](_page_16_Picture_8.jpeg)

Open Accelerator Infrastructure Project (OAI)

![](_page_17_Figure_1.jpeg)

![](_page_17_Picture_2.jpeg)

![](_page_18_Figure_0.jpeg)

https://www.opencompute.org/blog/new-open-accelerator-infrastructure-oai-sub-project-to-launch-within-the-ocp-server-project

![](_page_19_Figure_0.jpeg)

![](_page_20_Figure_0.jpeg)

![](_page_21_Picture_0.jpeg)

Open Accelerator Infrastructure (OAI)

()| OAM: Open Accelerator Module

**PSB: PCIe Switch Board** 03

02

()4

Tray

https://www.opencompute.org/blog/new-open-accelerator-infrastructure-oai-sub-project-to-launch-within-the-ocp-server-project

**UBB: Universal Baseboard** 

#### SCM: Secure Control Module

Chassis

### Open Accelerator Module (OAM)

![](_page_22_Picture_1.jpeg)

# UBB: Universal Baseboard

![](_page_23_Picture_1.jpeg)

![](_page_23_Picture_2.jpeg)

![](_page_24_Picture_0.jpeg)

https://onlinexperiences.com/Launch/Event.htm?ShowKey=86206&DisplayItem=E369184

![](_page_24_Picture_3.jpeg)

#### Stacked System Disaggregate from Host

![](_page_24_Picture_5.jpeg)

# Rack Level Solution

#### **48VDC Open Rack**

1 pairs 48V Bus Bar

1 shelf per Rack

#### **Power Shelf**

33KW(12xPSU)

40V-58V

93mm (H, 2OU) x 537mm (W, 21") x 586 (D) mm

#### System Devices

Inspur 30U OAI systems x4

Inspur 20U compute node x4

![](_page_25_Figure_11.jpeg)

![](_page_25_Figure_12.jpeg)

26	50	https://www.o	pencom	pute.org	/solutions/4	4/ope	en-cloud-	ai-rack

System	<u>Subtotal</u>		
Inspur OAI system	20kW (5kW x4)		
Inspur Compute Node	2.8kW (700W x4)	22.8KW	
Total Power/Ra	33kW		

![](_page_25_Figure_15.jpeg)

![](_page_25_Figure_16.jpeg)

OAI system

Compute node x2

OAI system

1	Cable symbol	Cable Type
]		Whisper Cable (PCIe)
		Eth cable
_		QSFP-DD (Serdes)
6		

![](_page_25_Picture_21.jpeg)

![](_page_26_Picture_0.jpeg)

2016Q2 X-MAN 1.0

Highest Density

![](_page_26_Picture_1.jpeg)

![](_page_26_Picture_2.jpeg)

AlBox is good show case for OAI Fully Connected topology

https://www.opencompute.org/summit/virtual-summit

https://onlinexperiences.com/Launch/Event.htm?ShowKey=86206&DisplayItem=E369188

Baidu is actively participating in OAI project and will build X-Man4.0 based on OAI

Baidu AlBox is able to train 10TB sparse models with single X-Man, and also scale to more nodes for bigger models and more samples. **Papers @ MLSys 2020** 

https://proceedings.mlsys.org/papers/2020/173

# Rack & Power Project

![](_page_27_Picture_1.jpeg)

![](_page_27_Picture_2.jpeg)

![](_page_27_Picture_3.jpeg)

Cloud & Networking News News

#### Facebook and Microsoft announce Open Rack V3 to address the power demands from artificial intelligence and networking

By **Bhagyashree R** - March 18, 2019 - 10:44 am • 1441

![](_page_28_Picture_5.jpeg)

https://engineering.fb.com/data-center-engineering/open-rack/

Converged rack frame

Flexible power shelf

Pluggable DC power shelf output interconnect

Battery backup systems

![](_page_28_Picture_12.jpeg)

![](_page_28_Picture_14.jpeg)

![](_page_28_Picture_15.jpeg)

![](_page_28_Picture_16.jpeg)

plugin

Universal AC power interconnect

![](_page_29_Picture_0.jpeg)

# The Impact of AI on the average Density per Rack

![](_page_29_Figure_2.jpeg)

![](_page_29_Picture_4.jpeg)

![](_page_30_Picture_0.jpeg)

# operationally or economically viable by 2020.

![](_page_30_Picture_2.jpeg)

Gartner, Inc., predicts that more than 30 % of data centers that fail to sufficiently prepare for AI will no longer be

Advanced Cooling Solutions(ACS) Sub-Project

![](_page_31_Picture_1.jpeg)

![](_page_31_Picture_2.jpeg)

![](_page_31_Picture_3.jpeg)

![](_page_32_Picture_0.jpeg)

# Advanced Cooling (ACS) Sub-Project

01 Fully Immersed Systems

02 Cold Plate Direct Contact Liquid Loop

# 03 Door-based Heat Exchangers

![](_page_33_Picture_0.jpeg)

# Stretching Limits

#### Extend life of air-cooled facilities & permit smooth transition

![](_page_33_Figure_3.jpeg)

#### Increase in component power / density

![](_page_33_Picture_6.jpeg)

![](_page_33_Figure_8.jpeg)

![](_page_34_Picture_0.jpeg)

https://onlinexperiences.com/Launch/Event.htm?ShowKey=86206&DisplayItem=E369065

Compute sleds + blanking pizza boxes not shown

#### Promising: +50% power compared to original air-cooled solution

![](_page_35_Figure_1.jpeg)

![](_page_35_Picture_2.jpeg)

#### Hyper...Aggregated...Empowered...Edge

![](_page_36_Picture_1.jpeg)

https://techwireasia.com/2020/01/walmart-takes-on-amazon-with-edge-computing-and-5g/

![](_page_36_Picture_5.jpeg)

![](_page_37_Picture_0.jpeg)

![](_page_37_Picture_1.jpeg)

https://submer.com/

![](_page_37_Picture_3.jpeg)

![](_page_37_Picture_4.jpeg)

![](_page_37_Picture_5.jpeg)

https://www.opencompute.org/projects/acs-immersion-cooling-sub-project

![](_page_37_Picture_8.jpeg)

![](_page_37_Picture_9.jpeg)

https://www.asperitas.com/

![](_page_37_Picture_11.jpeg)

# Predictive Server Modelling

![](_page_38_Picture_1.jpeg)

https://www.opencompute.org/circular-economy

![](_page_38_Picture_3.jpeg)

![](_page_38_Picture_5.jpeg)

# Workloads normally placed .....

Energy & Cost

![](_page_39_Picture_2.jpeg)

#### Efficiency & Throughput

# Intelligent ML optimisation reduces energy by > 20%

![](_page_40_Picture_1.jpeg)

Energy & Cost

#### Efficiency & Throughput

![](_page_40_Picture_5.jpeg)

# Actual & Predictive Server Modelling

Actual readings (in blue)

Predicted values (in orange)

For a given software loading, one can predict over 100 attributes to 97% accuracy

This includes heat output & power consumption

![](_page_41_Figure_5.jpeg)

![](_page_42_Picture_0.jpeg)

https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/

# Circular Economy

![](_page_43_Picture_1.jpeg)

https://www.opencompute.org/circular-economy

![](_page_43_Picture_3.jpeg)

![](_page_43_Picture_4.jpeg)

![](_page_43_Picture_5.jpeg)

There is no such thing as "away"....when we throw things away it must go somewhere.

Annie Leonard

![](_page_44_Picture_4.jpeg)

#### The vision is a data center industry where no heat is wasted....

DATA ECONOMY 🏫 NEWS - ECONOMY - BUSINESS - MARKETS - LEADERSHIP - INDUSTRY - LIFE & ART

#### Around \$10bn is set to be invested in India's data centre market over the next six years, with reports predicting that \$3bn will be spent building data centres in the country.

A recent research report predicts that the data centre market in India will grow at a CAGR of around 5% during the period of 2019 to 2025, according to Arizton.

The report also highlighted that data explosion and favourable government policies will aid the establishment of data centre parks across India.

As of October 2019, there were over ten data centre projects, which are expected to be operational between June 2020 and December 2022.

Adani Group, Bridge Data Centers, Colt DCS, and Yotta Infrastructure are among the new entrants to the market, targeting hyperscale investments with multiple facilities covering over 40 MW of power capacity.

Adani Group recently signed a Memorandum of Understanding with Digital Realty with the aim of enabling both companies to jointly evaluate developing and operating data centres, data centre parks and cultivating undersea cable provider communities of interest across India.

![](_page_45_Picture_8.jpeg)

### 46M servers to be "End of Life" (EOL'ed)" between 2019 and 2023

![](_page_46_Picture_1.jpeg)

https://www.ellenmacarthurfoundation.org/

#### **CEDaCI - Circular Economy for the Data Centre** Industry

https://www.nweurope.eu/projects/project-search/cedacicircular-economy-for-the-data-centre-industry/

![](_page_46_Picture_5.jpeg)

Ali Fenn – President of ITRenew

https://www.itrenew.com/

https://www.youtube.com/watch?v=PuzEMYCGeVk

![](_page_46_Picture_9.jpeg)

![](_page_46_Picture_11.jpeg)

![](_page_46_Picture_12.jpeg)

<image/> <text><text><text><text><text><text></text></text></text></text></text></text>						Sea	rch Q Co	ontact
<text><text><text></text></text></text>	OPEN Compute Project ®	About V Marketplace V Products	Contributions ~	Projects ∽	Events ~	SP/Reseller ∨	Membership ~	Blog
Open. For Business. The Open Compute Project (OCP) is reimagining hardware, making it more efficient, flexible, and scalable. Join our global community of technology leaders working together to break open the black box of proprietary IT infrastructure to achieve greater choice, customization, and cost savings.		Integrated Solutions Circular Economy		NR.			E SA	
	Open. For Busin	<b>ESS.</b> magining hardware, makin oin our global community o break open the black box greater choice, customiza	ng it of ¢ of ation,		25		26	

Sesame for AI/ML provides the high bandwidth, low latency performance you need to manage the huge volumes of data in your demanding learning environment.

![](_page_47_Picture_3.jpeg)

![](_page_47_Picture_4.jpeg)

https://www.opencompute.org/circular-economy

![](_page_47_Picture_6.jpeg)

### Open Cloud Al Rack

Customize your solution with a variety of compute nodes, JBOG, storage, FPGA, and networking components for applications ranging from high-density deep learning training to data mining and data analysis.

![](_page_48_Picture_2.jpeg)

49 | 50 <u>https://www.opencompute.org/solutions</u>

![](_page_48_Picture_4.jpeg)

![](_page_49_Picture_0.jpeg)

![](_page_49_Picture_4.jpeg)

# How to get

#### steve@opencompute.org

![](_page_50_Picture_2.jpeg)

Compute Project<sup>®</sup>

OCP Marketplace https://www.opencompute.org/products

Past Events (recordings and slides) https://www.opencompute.org/events/past-summits

https://www.opencompute.org/events/past-events

Social

in

f

https://www.linkedin.com/groups/4152886/

@OpenComputePrj

https://www.facebook.com/groups/opencompute/

**OCP** Projects https://www.opencompute.org/projects

OAI Project Mailing List https://ocp-all.groups.io/g/OCP-OAI

OAI Project Wiki with latest specification https://www.opencompute.org/wiki/Server/OAL

Advanced Cooling https://www.opencompute.org/projects/acs-immersion-cooling-subproject

https://www.youtube.com/user/OpenComputeProject

### The AI J-Curve

![](_page_51_Figure_1.jpeg)

Time

Productivity & Performance

**Compute Project** 

![](_page_51_Picture_5.jpeg)

.......

What stakeholders (mistakenly) expect

What stakeholders can expect with proper execution

What actually happens in most cases

![](_page_51_Figure_9.jpeg)

# Thank you

![](_page_52_Picture_2.jpeg)

# Questions?

# **piugin**

![](_page_53_Picture_2.jpeg)

![](_page_53_Picture_3.jpeg)