

# OCPJ Meet-up 2019

## Lightning Talks

Marvell

Delta

MiTAC

CREDO

Penguin Computing

OCPJ Meet-up

**M A R V E L L<sup>®</sup>**

2019年6月27日

山内正光

Business Development Manager

Marvell Japan

**M A R V E L L<sup>®</sup>**


2018年: MarvellによるCaviumの買収が完了

2016年: CaviumによるQLogicの買収が完了

# THUNDERX2<sup>®</sup>

- 非常に強力なArmプロセッサ
- 最多32コア（2ソケット構成では最多64コア）
- 8本のDDR4メモリーチャンネル（2ソケット構成では16本）
- 最新のTOP500で156位のスーパーコンピューター（Sandia National LaboratoriesのAstra）が使用（138,096コア）

# Ethernet Bunch of Flash (EBOF)



# Driving Open Networking to Production

## A Vendor Perspective



**OPEN**  
PLATINUM™

Kevin Chan  
Delta Electronics Japan  
Jun. 27<sup>th</sup>, 2019



# About DELTA

1971 Founded in Taiwan

World's Leader in Power Supplies

Networking Products business Since 1991

Delta Japan office Since 1991

## Business Categories

### Power Electronics

- Components
- Embedded Power
- Fans & Thermal Management
- Automotive Electronics
- Merchant & Mobile Power



### Automation

- Industrial Automation
- Building Automation

### Infrastructure


- ICT Infrastructure
- Energy Infrastructure

Delta Confidential

## Global Operations



## Worldwide Revenues






Micro Datacenter




Modular Datacenter



Containerized Datacenter



Power Container



Datacenter Infrastructure Management (DCIM)



AC Power



Eltek Converged Power Systems



Precision Cooling



Power Shelf



Racks





















Power Distribution Units (PDU)

Delta Confidential



# OCP Membership since 2016

AT&T (since 2015) 	Baidu (since 2019) 	Facebook (since 2011) 	<b>Delta Electronics (since 2016)</b> 	Intel (since 2011) 	IBM (since 2013) 
Deutsche Telekom (since 2016) 	LinkedIn (since 2018) 	Nokia (since 2015) 	Flex (since 2016) 	JD.com (since 2018) 	Lenovo (since 2016) 
Alibaba (since 2017) 	Google (since 2015) 	HPE (since 2015) 	Huawei (since 2018) 	Microsoft (since 2014) 	Tencent (since 2018) 

# Hyper-scale Data Center Power/Rack




12V Power Shelf + BBU




48V Power Shelf + BBU



ATS + PDU + PSUs  
→ 12V Power Shelf + BBU





12V Power Shelf + BBU + Fan Wall



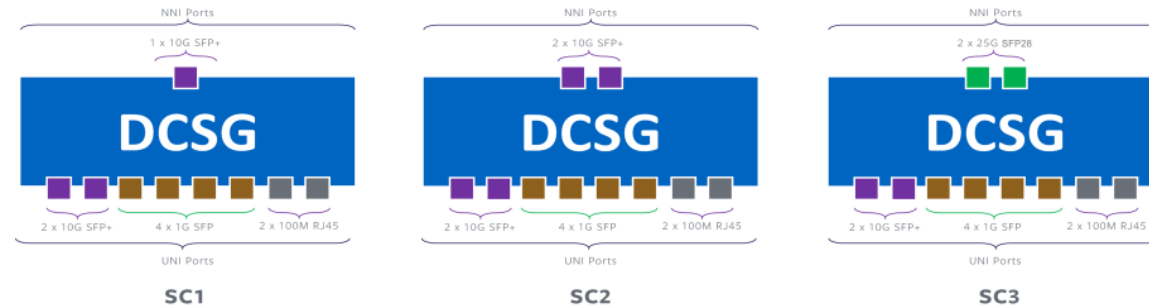
Combined PSU + PDU (ATS + BBU) → PSU

# OCP Inspired/Accepted White box Switches

<p><b>Delta AG7648C</b></p> <p>10GbE Leaf/TOR switch for Data Center</p> <p>Solution Part #: A Specifications</p>	<p><b>Delta AGC5648S</b></p> <p>25 GbE Deep Buffer Switch for Carrier/Service Provider</p> <p>Solution Part #: A Specifications</p>	<p><b>Delta AG5648v1</b></p> <p>25 GbE Leaf/TOR Switch for Data Center</p> <p>Solution Part #: A Specifications</p>	<p><b>Delta AGC7648A</b></p> <p>10 GbE Deep Buffer Switch for Carrier/Service Provider</p> <p>Solution Part #: A Specifications</p>	<p><b>Delta AG9032v1</b></p> <p>100 GbE Spine Switch for Data Center</p> <p>Solution Part #: A Specifications</p>	<p><b>Delta Agema 1G PoE Switch</b></p> <p>48 1G Copper PoE ports w/ first 12 ports as UPOE, last 2 as combo ports, 2 x 10G SFP+ uplinks, 2200 W max power capacity, AC Input, ONIE only</p> <p>Solution Provider: Delta Electronics Part #: AG6248C PoE-R Specifications</p> 	 <p><b>OCP</b> ACCEPTED</p>
---	---	---	---	---	---	--





# TIP Disaggregation Cell Site Gateway

TIP DCSG Requirement

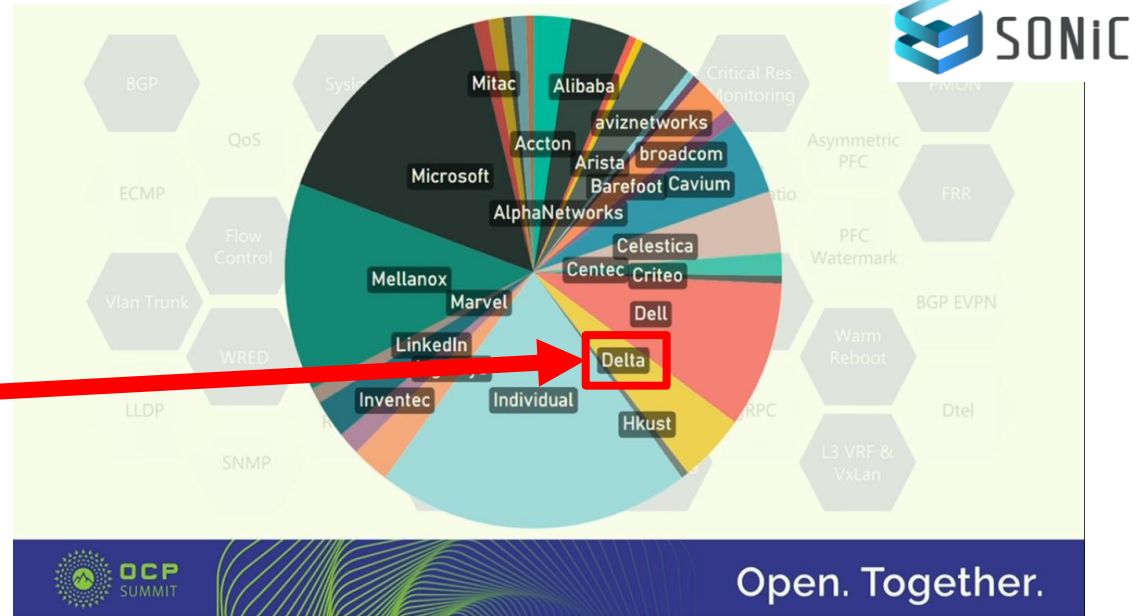



**AGCV208S (4\*RJ45 + 4\*1G SFP/10G SFP + 8\*25G SFP28 + 2\*100G QSFP28)**

Delta's Suggestion



# SONiC Contribution



Tomahawk 2



Tomahawk



Tomahawk+





Networking Vendors



NoviFlow



One Source Integrations



Connectivity of Tomorrow



White Box ODM Vendors

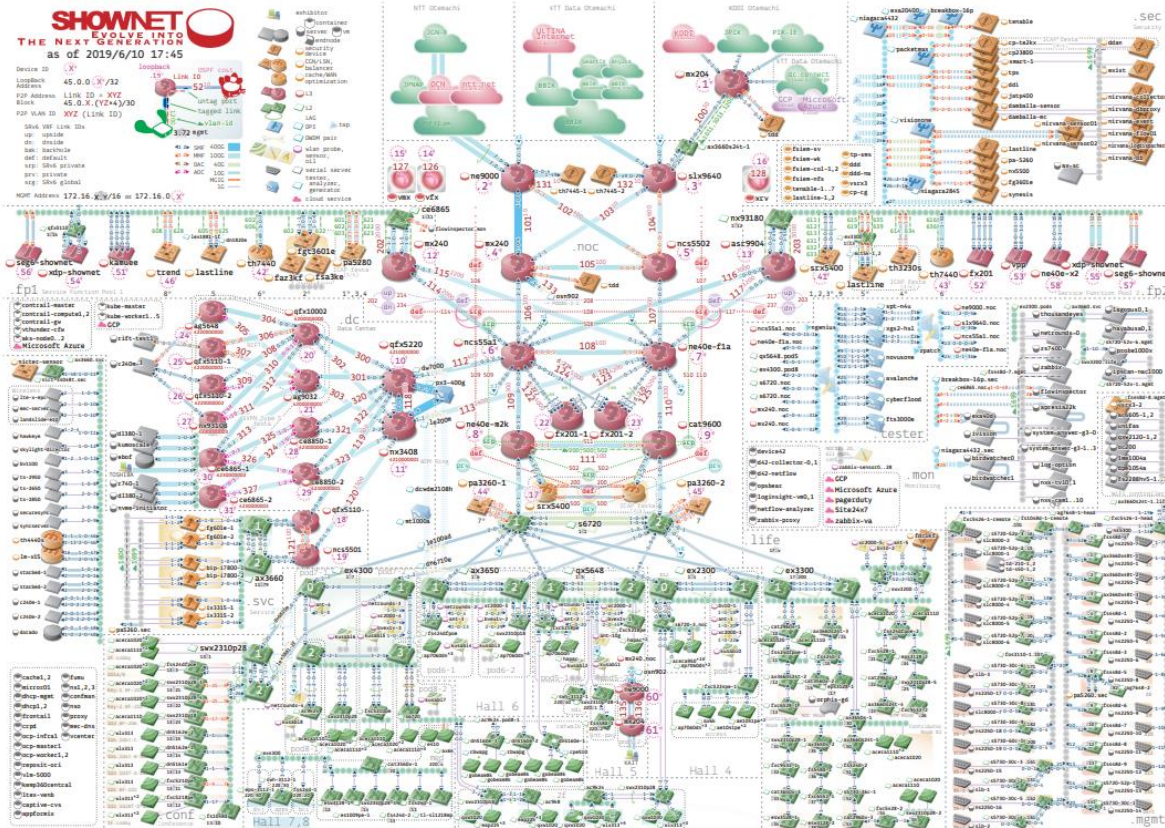


Silicon Vendors



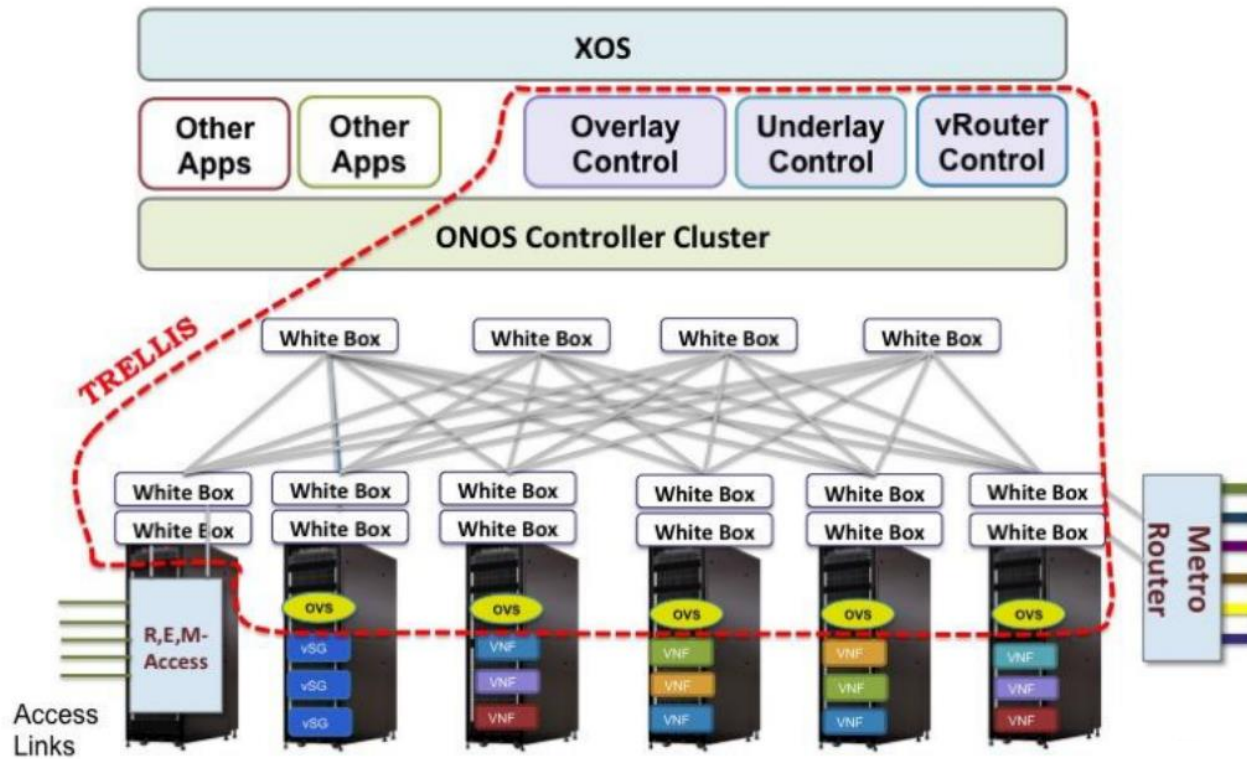
Other Open Source Projects





- AG9032V1      32 x 100G
- AG5648V1      48 x 25G + 6 x 100G
- AG7648        48 x 10G + 6 x 40G

## CORD Architecture



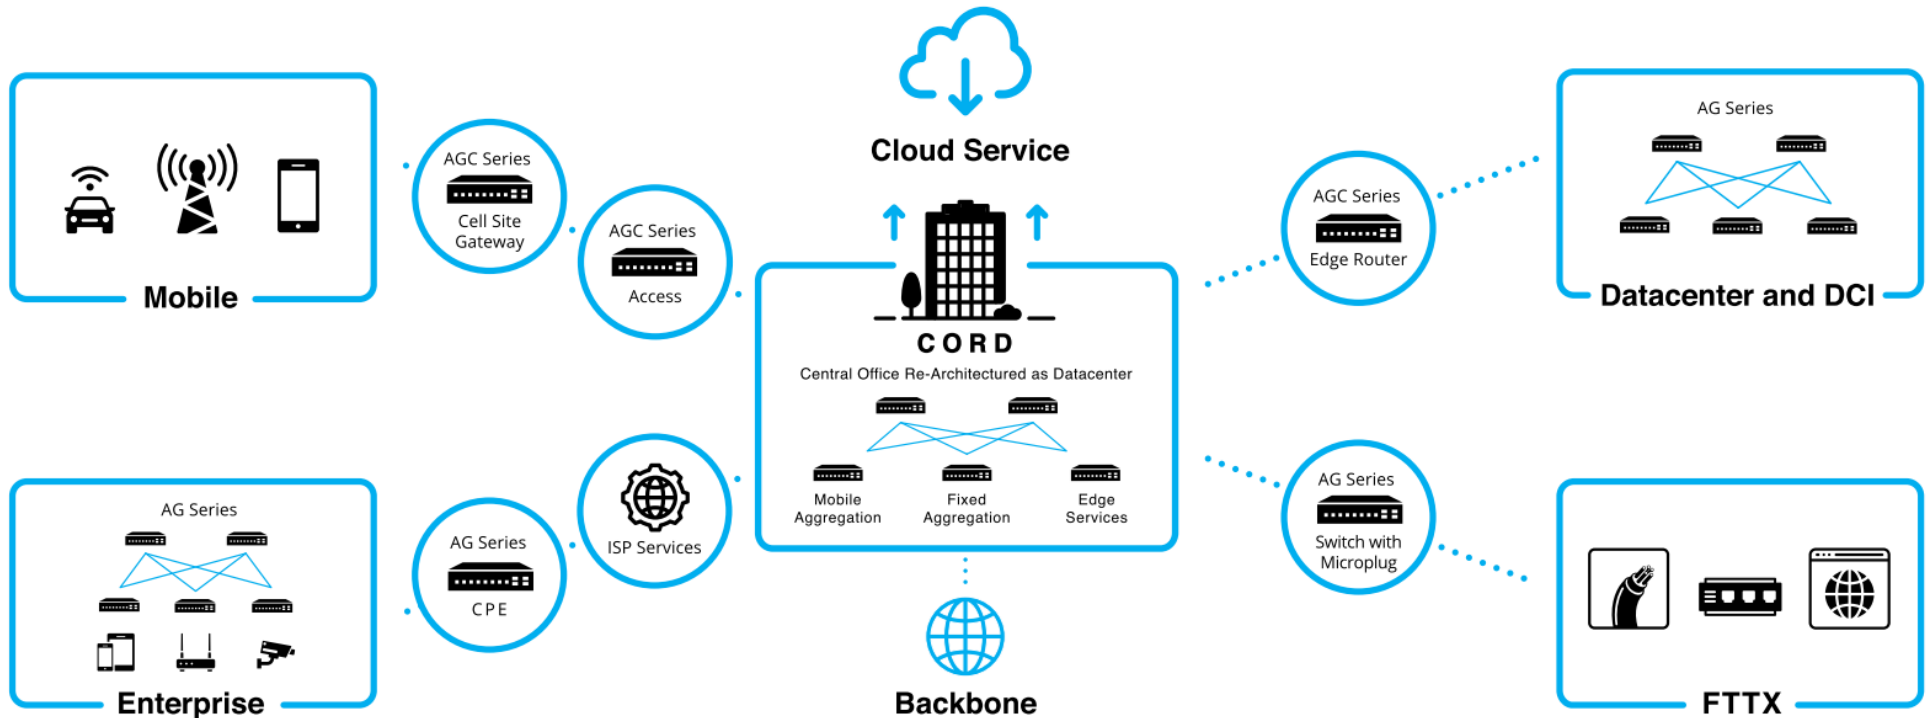


Platform	ASIC	Port Configuration	Fabric	Packet Buffer	CPU	Memory	Storage (SSD)	Power Budget
AGC6248C	Helix 4	48x1GbE BASE-T / 2xSFP+	68 Gbps	4 MB	Embedded-dual 1GHz Cortex ARM9	1 GB DDR3	NAND Flash 4 GB	460W
AGC6248C POE	Helix 4	48x1GbE BASE-T / 2xSFP+	68 Gbps	4 MB	Embedded-dual 1GHz Cortex ARM9	1 GB DDR3	NAND Flash 4 GB	1100W
AG7648	Trident 2	48xSFP / 6xQSFP28	720 Gbps	12 MB	Intel Rangley C2338	2 GB DDR3	8 GB	460W
AG7648C	Trident 2	48x10GbE BASE-T / 6xQSFP	720 Gbps	12 MB	Intel Rangley C2338	4 GB DDR3	8 GB	460W
AG8032	Trident 2	32xQSFP	1.28 Tbps	12 MB	Intel Rangley C2338	8 GB DDR3	128 GB	460W
AG5648V1	Tomahawk +	48xSFP28 / 6xQSFP28	1.8 Tbps	16 MB	Intel Rangley C2538	4 GB DDR3	8 GB	550W
AG9032V1	Tomahawk	32xQSFP28	3.2 Tbps	16 MB	Intel Rangley C2538	8 GB DDR3	128 GB	800W
AG9032V2	Trident 3	32xQSFP28 / 2xSFP+	3.2 Tbps	32 MB	Intel Denverton C3538	16 GB DDR4	64 GB	650W
AG9032V2A	Trident 3	32xQSFP28 / 1xSFP+	3.2 Tbps	32 MB	Intel Broadwell D1527	2 x 16 GB DDR4	128 GB	650W
AG9064	Tomahawk 2	64xQSFP28	6.4 Tbps	42 MB	Intel Broadwell D1548	32 GB DDR4	128 GB	1300W
AGC032	Tomahawk 3	32xQSFP56-DD	12.8 Tbps	42 MB	Intel Broadwell D1527	32 GB DDR4	128 GB	1600W
AGCV208S	Qumran-AX	4xSFP+ / 8x25G / 2xQSFP28	300 Gbps	3GB GDDR4	Intel Denverton C3538	8 GB DDR4	64 GB	400W
AGC7648A	Qumran-MX	48xSFP+ / 6xQSFP28	800 Gbps	4 GB GDDR5	Intel Rangeley C2538	8 GB DDR3	32 GB	800W
AGC7648SV1A	Qumran-MX	48xSFP+ / 6xQSFP28	800 Gbps	4 GB GDDR5	Intel Broadwell D1527	32 GB DDR4	128 GB	800W
AGC5648S	Jericho +	48xSFP28 / 6xQSFP28	1.8 Tbps	2 x 16 GB GDDR5	Intel Broadwell D1548	32 GB DDR4	128 GB	1600W
AGC9020S	Jericho +	20xQSFP28	1.8 Tbps	2 x 16 GB GDDR5	Intel Broadwell D1548	32 GB DDR4	128 GB	1300W
AGC9080	Jericho 2	80xQSFP28	9.6 Tbps	8 GB	Intel Broadwell D1528	2 x 16 GB DDR4	128 GB	2200W
AGCX848V1S	Jericho 2	48xQSFP28 / 8xQSFP-DD	9.6 Tbps	8 GB	Intel Broadwell D1528	2 x 16 GB DDR4	128 GB	2200W

DC power aviable with all models

# Targeted Applications

Delta focuses on few applications base on IP Networking so that the design expertise could be accumulated and exceeded



Smarter. Greener. Together.

To learn more about Delta, please visit [www.deltaww.com](http://www.deltaww.com)





# MiTAC OCP Products for OCPJ Meet up

June-27, 2019

# MiTAC OCP Solutions



**OPEN PLATINUM**

**Compute Node**


- Mezzanine Card
- Storage/JBOF
- 2 OU Shelf
- Busbar
- 8 OU-Height Rail

Visit us @ **MiTAC Portal:** <http://www.mitac.com/Product/Open-Compute-Project.html>  
**Market Place:** <https://www.opencompute.org/products?query=mitac&page=1>

- OCP Mezzanine
- OCP Compute Node
- OCP ESA Kit
- OCP JBOF

<https://www.youtube.com/watch?v=bJAJ4gutbYo>

# Continuous Commitment on OCP

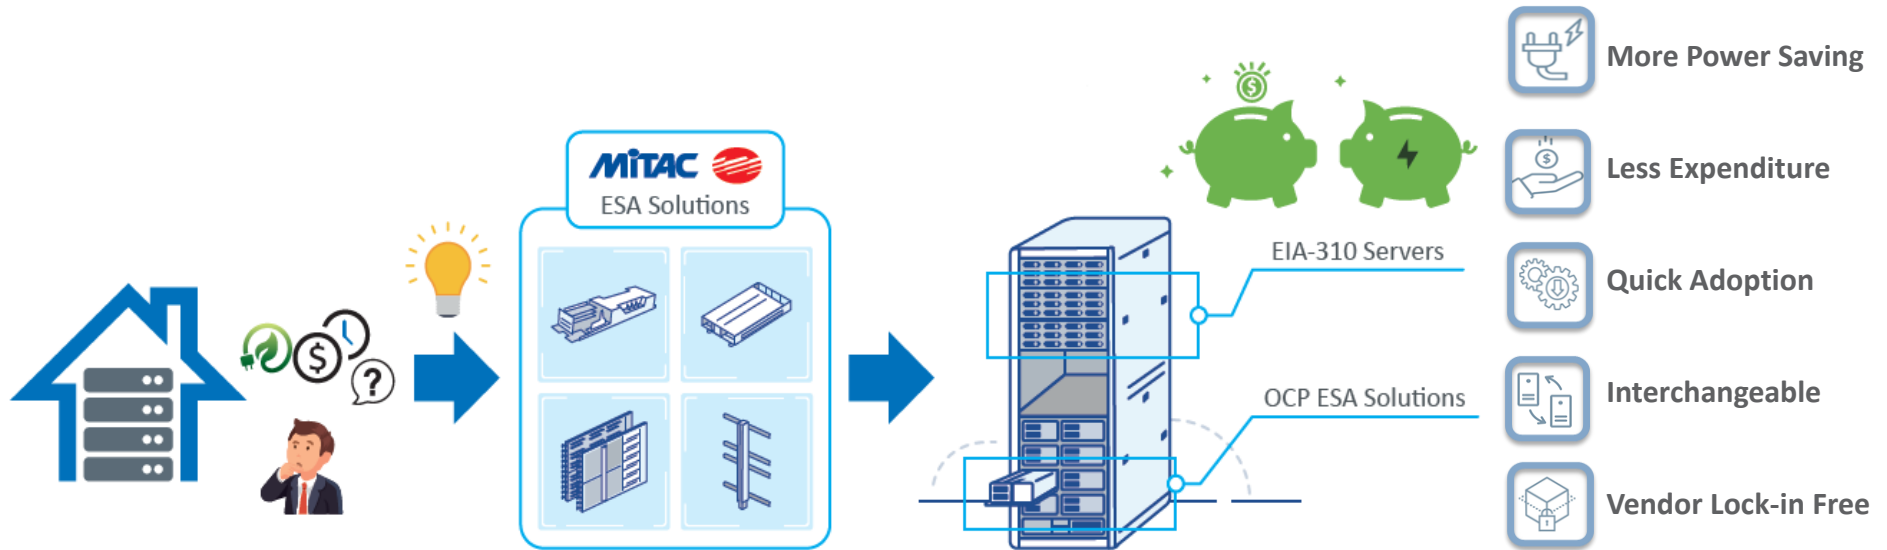


# What Is ESA - Enclosure Sub Ass.



★ MiTAC ESA kit helps convert traditional EIA-310 racks to be adoptable with OCP solution;

It's no need to rebuild the whole data center facility to deploy OCP solutions!



★ **Spec has been contributed in 2018 July**

**Product Details:** <https://www.youtube.com/watch?v=bJAJ4gutbYo>



OCP ESA Kit



19" Power Shelf





OCP Compute Node





OCP JBOF



# Tioga Pass Spec



	Standard (E7278-S)	Advanced (E7278-A)	Ultra (E7278-U)
<b>Processor</b>	(2) Intel® Xeon® Scalable processor, TDP 165W ➢ Support both SkyLake SP and Cascade Lake SP, Intel® C621-AT chipset		
<b>Memory</b>	(16) DIMM slots, 6 channels ➢ Up to 1024 GB R-DIMM, DDR4 2666MT/s		
<b>PCIe Extension</b> 	(2) FHHL PCIe Gen3 x16	(2) HHHL PCIe Gen3 x16	(1) FHHL PCIe Gen3 x16
<b>Storage Extension</b> 	(1) 3.5" hard drive bay <sup>[1]</sup>	(6) 2.5" hard drive bay <sup>[1][2]</sup>	(4) 2.5" hard drive bay <sup>[1][3]</sup> (1) 2.5" hard drive bay <sup>[1][2]</sup>
<b>Mezzanine</b>	(1) OCP NIC 2.0		
<b>On Board Storage</b>	(1) Up to 110mm M.2 (SATA3 / PCIe Gen3 x4)		
<b>Management IO</b>	(1) 1GbE RJ45, support NC-SI		
<b>BMC</b>	ASPEED AST2500 ➢ Support IPMI 2.0 and DMTF Redfish® 1.0		
<b>Dimension &amp; Weight</b>	D28.5" x W6.9" x H3.5" (724 x 175 x 89mm) / 17.6~18.7lbs (8~8.5kg)		
<b>Power Supply</b>	Centralized OCP power shelf (12V DC)		

[1] Shipment does not include hard drive

[2] Support 7mm SSD only

[3] Support SATA / SAS / U.2 hard drive








## Credo Introduction

OCP Japan Meet Up  
June 27, 2019


# クレド概要

- 2008年: 創業、高速 SerDes (Serializer-Deserializer) IP 開発開始
- 2014年: SerDes IP ライセンスと高速イーサネットシステム用製品の販売開始
- 高速SerDesの低電力低コスト化を、成熟した半導体プロセス(28nm)で実現し、市場シェアを拡大中



## クレドのビジネスモデルと成功例の一部



# FBOSS 初の 3rd Party Platform = Arista 7368X4 に Credo 400G ギャボックス 搭載



**OCP**  
GLOBAL  
SUMMIT



Facebook Minipack	Arista 7368X4
OCP MiniLake	Arista uServer
Brcm TH3	Brcm TH3
Brcm Gearbox	<b>Credo Gearbox</b>
Edgecore Board (FPGA/PIMs/PSUs/Fans)	Arista Board (FPGA/PIMs/PSUs/Fans)

Open. Together.

San Jose Convention Center | San Jose, California

March 14-15, 2019

# 400G AEC (Active Electrical Cable)



400G to 400G  
(QSFP-DD/OSFP)



400G (QSFP-DD/OSFP) to 4 x 100G (QSFP)  
ブレイクアウトケーブル

- 最新 12.8T Switch 400G Port (50G/lane) システムの課題
  - DAC (ダイレクトアタッチケーブル) ではデータ転送が困難
  - AOC (アクティブ光ケーブル) は高コスト高電力
- Credo Retimer/GB 搭載の銅線ケーブル (AEC) が課題を解決
  - AOCより低コスト低電力
  - 50G/lane から 25G/lane への変換がケーブル内で可能 (ブレイクアウト)
- 2019年後半から量産開始
  - 台湾ODM複数社で認定評価進行中
  - ハイパースケールデータセンターとの商談進行中
  - 25G/lane システムユーザからの引合い多数のため 100G AEC も量産予定

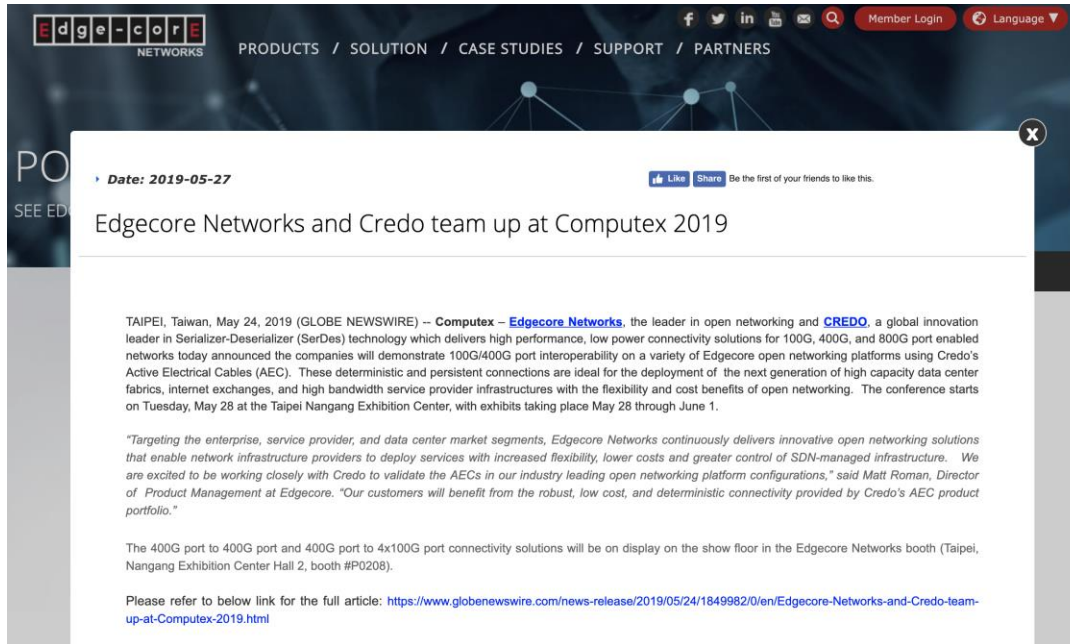


400G  
Retimer



400G  
Gearbox

# Computex 2019 での 400G AEC 共同出展 (Edgecore 社様)



The screenshot shows a web page for Edgecore Networks. At the top, there is a navigation bar with the Edgecore Networks logo and links for PRODUCTS, SOLUTION, CASE STUDIES, SUPPORT, and PARTNERS. Below the navigation bar, there is a social media sharing section with 'Like' and 'Share' buttons. The main headline reads 'Edgecore Networks and Credo team up at Computex 2019'. The article text describes the partnership between Edgecore Networks and Credo, highlighting their joint demonstration of 100G/400G port interoperability using Credo's Active Electrical Cables (AEC). It mentions the event taking place at the Taipei Nangang Exhibition Center from May 28 to June 1, 2019. A quote from Matt Roman, Director of Product Management at Edgecore, is included, along with information about the display of 400G port connectivity solutions at the booth. A link to the full article is provided at the bottom.

Edgecore Networks and Credo team up at Computex 2019

TAIPEI, Taiwan, May 24, 2019 (GLOBE NEWSWIRE) -- Computex – **Edgecore Networks**, the leader in open networking and **CREDO**, a global innovation leader in Serializer-Deserializer (SerDes) technology which delivers high performance, low power connectivity solutions for 100G, 400G, and 800G port enabled networks today announced the companies will demonstrate 100G/400G port interoperability on a variety of Edgecore open networking platforms using Credo's Active Electrical Cables (AEC). These deterministic and persistent connections are ideal for the deployment of the next generation of high capacity data center fabrics, internet exchanges, and high bandwidth service provider infrastructures with the flexibility and cost benefits of open networking. The conference starts on Tuesday, May 28 at the Taipei Nangang Exhibition Center, with exhibits taking place May 28 through June 1.

"Targeting the enterprise, service provider, and data center market segments, Edgecore Networks continuously delivers innovative open networking solutions that enable network infrastructure providers to deploy services with increased flexibility, lower costs and greater control of SDN-managed infrastructure. We are excited to be working closely with Credo to validate the AECs in our industry leading open networking platform configurations," said Matt Roman, Director of Product Management at Edgecore. "Our customers will benefit from the robust, low cost, and deterministic connectivity provided by Credo's AEC product portfolio."

The 400G port to 400G port and 400G port to 4x100G port connectivity solutions will be on display on the show floor in the Edgecore Networks booth (Taipei, Nangang Exhibition Center Hall 2, booth #P0208).

Please refer to below link for the full article: <https://www.globenewswire.com/news-release/2019/05/24/1849982/0/en/Edgecore-Networks-and-Credo-team-up-at-Computex-2019.html>

Edgecore 社様展示ブースにおける  
同社製 400G (32 port) Switch AS9716-32D と  
Credo 400G AEC (QSFP-DD to QSFP-DD)の展示






# Credo プライベートスイートでの各種AEC展示とライブデモ



各種 400G AEC、200G AEC、100G AEC の展示  
並びに 800G AEC の参考出展





**CREDO**  
we connect.



# OCP Solutions by Penguin Computing

2019 OCP-J Meet-up  
June 27<sup>th</sup>, 2019

# Penguin Computing

Developing Open, Linux-Based Solutions for Enterprise Data Centers, HPC and Cloud.


Tux – Linux's mascot



[lewing@isc.tamu.edu](mailto:lewing@isc.tamu.edu) Larry Ewing and The GIMP



**Linux Expertise and Innovation**



- What We Do – Help startups, Fortune 500, government, and academic organizations around the globe solve the most complex technology and business challenges
- How We Do It – AI, engineering, and computer science experts use deep technical knowledge and out-of-the-box thinking to reimagine what is possible

# 2012 OCP named Solution Vendor

Early primary member of Open Compute Project (OCP)



2018

<https://www.youtube.com/watch?v=p4D9UNXKGV>

<https://youtu.be/SXXReqf9uWQ>



*Phillip Pokorny, CTO, Penguin Computing*

## Penguin Computing CTO Selected as HPC Representative of OCP Incubation Committee

In February 2018, Penguin Chief Technology Officer Phillip Pokorny was selected as the HPC representative of the OCP Incubation Committee.

The Incubation Committee of OCP works with the OCP Foundation to review all specifications and designs that are submitted. Its members help set goals and direction for growing the community. Because of their diverse technical backgrounds, as well as their insight into technology verticals and trends, subject members advise on contribution guidelines and project scope and reach.

[Learn More](#)

# Tundra Extreme Scale

**16 OCP-based Tundra supercomputers – running complex nuclear testing simulations, among other tasks – were ranked top 500**

- Supports a diverse computing architectures, including GPU, while providing the benefits of an OCP form factor.
- 10U–40U with a capacity for over 100 nodes per rack for unmatched power and density.
- Built with the latest processors high-speed SDN, and localized performance SDS.
- Comes with Tundra Extreme Scale Rack for greater flexibility, operability and efficiency.

