

OCP Networking Project Meeting Notes

Date: 2025-11-05

Agenda/Presenters

Date/Time: Nov 5, 8.00-8.45am Pacific

Presenters

- Michael Schill (OCP)
- Manoj Wadekar (Meta)
- Rajesh Sankaran (Microsoft)

Agenda

At the 2025 OCP Global summit, OCP announced OCP ESUN (Ethernet for Scaleup Networking).

The agenda of this meeting is to kick-off the ESUN workstream under the OCP-Networking Project. Presenters will share the following:

- OCP ESUN Process
- ESUN Overview
 - Motivations for ESUN
 - Goals and non-goals for the OCP ESUN workstream
 - Framework for workstream participation
 - Workstream collaboration details (regular call schedule, mailer, wiki etc)

Attendees (175)



Ravindra Sunkad (OCP Networking) (me)



Michael Schill (OCP) (Host)



Rajesh Sankaran (MSFT)



Manoj Wadekar (Meta)



Mohan Kalkunte (Broadcom)



1408****256



Aviran Kadosh



Clarence Filsfils (Cisco)



Danny V.



David Ofelt (Juniper/HPE)



Gerald Schmidt (NVIDIA)



Greg Whelan (Teradio)



Jesin Kim (Samsung)



Jose Perez



JunyiZhang (Huawei)



Mark Gravel [Hewlett Packard Enterprise]



xingjun chu (Huawei Canada)



Ahmed Abdelsalam (Cisco)

W(Wenjun (CoMIRA)

W Wesley Richards (CoMIRA)

WL Winston Liu (Keysight)

XF Xiao Feng Wang [Qualcomm]

ys yuki saka

富 富士通 瓦井

PH Paul Hartke

 Sree Chadalavada (Cisco)

 Sridhar SR

 Srihari Vegesna (Upscale AI)

 Srini

 Sridhar SR

 Srinivas Gangam (Upscale AI)

 Subrata Banerjee (Upscale AI)

 Sumiran Shubhi (Ampere)

 Tian Fang

 Tina Tsou

 Tom Herbert (XDPnet)

 Uday Mini

 uelzur

 Varun Vij (Upscale AI)

 Veerasenareddy Burru

 Vitorio's - Samsung

 vivek kashyap (intel)

 Weiqiang Cheng (China Mobile)



Richard Shannon (Ampere)



rmahatme



Robert Hormuth (AMD)



Roberto Guerreiro



Robin Grindley



Ronen Hyatt (UnifabriX)



RUNLONG HU (CMCC)



Ryan Yu, TeraHop



Sai Rahul Ch(d-Matrix)



Sandeep Banerjee [UpscaleAI]



Santhosh(Upscale AI)



Scott Irwin (Adtran)



Sean Kao



shainer



Shankar Vasudevan



Shawn Chen (Carnegie Mellon University)



Shawn Dube (Dell)



Shrawan Chittoor Surender (Upscale AI)



Paul Congdon (Huawei)



Pete Del Vecchio [Broadcom]



Prakash Rashinkar



Pratik Marolia (Msft)



Priyank Shukla (Synopsys)



Purush Gupta L B @CLS



pvautee



Radha Chintakuntla



Rajiv Krishnamurthy (Meta)



Rakesh Cheerla



Rakesh Kumar(Ampere Computing)



Ramana Dasari(Upscale AI)



Ramanan Natarajan (UpscaleAI)



Ramesh Sivakolundu (Cisco)



Ravi (Marvell)



Ravi Kiran Gummaluri(Micron)



Reddy (Intel)



Richard Shannon (Ampere)

 Lisa Nguyen (Broadcom)

 Lowell Lamb

 Mayuresh(Upscale AI)

 mdemers

 Meg Lin (Marvell)

 Mesut Ali Ergin (Intel)

 mike

 Mike Andrewartha

 Mikel Jimenez Fernandez

 Moshe Voloshin

 Nerses

 Nick.Kriczky - Teledyne LeCroy

 Nilesh Kamdar (Microsoft)

 Nitin Nangare (Marvell)

 Nolan Morgan (Nereus)

 Omar Baldonado (Meta)

 Pankaj Mehra

 Pat Fleming



Jay Pottenger



Jennifer Anselmo



Jieyu Li [CMCC]



Jim Panian (Qualcomm)



JMC (Teledyne Xena)



John Cheung (Upscale AI)



John Leung (Intel)



John Marshall (AMD)



Jon Hermes (Arm)



Juan-Carlos Calderon (AlphaWave)



Jungmin Choi (SK hynix)



kannan (Celestica)



Kingshuk Mandal(Keysight)



Krishna Mallampati



Krithika's iPhone



Kumuthini Ratnas (HPE)



Kurtis Bowman (AMD)



Lakshmi Seshan



Don Cober (CoMira)



Dong Zhang (APD)



Douglas Lee (Teledyne LeCroy)



Durgesh Srivastava



Earl Parsons [CommScope]



Eduardo Mendes (AxeleraAI)



eisraeli



Gary Goldman (Tensordyne)



Gobinath Krishnamoorthy (Upscale AI)



Haider Salman



Helia



Hemal Shah



Hongjian Fan (Seagate)



Houxiang Ji



Hugh Holbrook (Arista Networks)



Ian Cox (Broadcom)



Jai Kumar (Broadcom)



James Kelly (OCP)



Azadeh Farzin



Balaji Srinivasan



Barak Gafni



Bijan Nowroozi (Lightmatter)



Bruck Girmay



Carmine Scarpitta (Cisco)



casper.yoon



Chad Hintz (AMD)



Christian Kuhtz (Nokia)



Chris's iPhone 13



Cliff.T [Keysight]



Craig's iPhone (Teledyne LeCroy)



Dan Biederman (Eridu AI)



Dan Wilkinson



Daniel Morato (UPNA/Naudit)



David Riddoch



Deepti (Upscale AI)



DJ (Meta)

Aibing (UpscaleAI)

Alberto Cicalini (Qualcomm)

Alex Umansky "Huawei"

Ali

Alicia Schap

Amnon Izhar (Dell)

Andrew J Herdrich

Anil's iPhone Intel

Anish Mathew [Cadence]

Ankur Sheth

Anu Ramamurthy [Microchip]

Arif Khan -Cadence

Arihant Jain (Arista)

Arto

Arun Saha

Avdhesh

Avik Bhattacharya

Avinash Natarajan (Celestica)

Notes/Minutes: [Video Recording Link](#)

Notes

- Michel Schill's (OCP) presentation on [OCP ESUN Process](#)
- Manoj Wadekar (Meta) and Rajesh Sankaran (Microsoft) presentation on [ESUN Overview](#)

Process:-

- ESUN mailing list to be joined as part of the process
 - Here is the link to the OCP ESUN Workstream:
<https://www.opencompute.org/wiki/Networking/ESUN>
 - (Check back later today for a link to access the Feedback Agreement)

Zoom Chat QA/Comments

- Do we need to be a paying member or contributor to join the ESUN list?
 - OCP/Michael Schill to answer
- Is there a plan for a WG to cover the transport level (with SUE-T being the first example)
 - [Manoj Wadekar](#) - ESUN will focus on Fabric functionality and interfaces.
- +1 on focusing the scope at the fabric. From an operator view, our asks are simple: preserve end-to-end ECN semantics, define a compact 'IP-optional' header that still carries class and entropy, pick one lossless mechanism that avoids pause storms, and publish buffer and timing budgets so silicon and NIC/GPU teams can converge. I'm ready to contribute traces and workload profiles to ground the tradeoffs.
 - From Tina Tsou (any idea which company?)
- Any target timeline when we want a draft ESUN spec done?
 - [Omar Baldonado](#)
 - Goal is to drive multi Ethernet vendor support
 - Open discussion progressively evolving to agreement.
 - Drawing parallel to SAI, with Microsoft driving it early on. Quickly converged to base L2/L3 and slowly evolved to cover broad areas over the years
- Is there an early draft specification for ESUN that the member companies can review?
 - Seed proposal will be contributed and the group can take it further.
- It seems likely that 'Efficient headers' would be tightly coupled with the particular transport. Any initial ideas on how to make 'efficient headers' generic and in support of different transports?
 - [Manoj Wadekar](#) Goal is to make AI headers appropriate for different transports.
- Would In Network Collectives be focused as a part of ESUN ?

- [Manoj Wadekar](#) Collectives and INC are tightly coupled with transport, which is currently out of scope. Can discuss as we go along
- Is there any thought about getting into the PHY to optimize for smaller scale networks, lower latency, etc.?
 - [Manoj Wadekar](#) If there are good proposals to be discussed - workstream can consider those. At this time there is no specific proposal.
- Is there already a preliminary spec available to review?
 - POPULAR QUESTION: May want to double click on this once more.
 - [Omar Baldonado](#)
- Is there any thought about getting into the PHY to optimize for smaller scale networks, lower latency, etc.?
 - Manoj: That is open for feedback
- The initial version of the SUE spec includes AFH header design. Does this get moved into the ESUN workstream and removed from the next SUE spec?
 - [Manoj Wadekar](#) Since packet header is part of interface between end node and the fabric - this needs to be defined in ESUN.
- Can we share analysis as other (non-Ethernet) technologies were considered and ruled out?
 - Rajesh/Manoj: Ethernet choice primarily drive due to the large expertise/knowledge-base across vendors, operators
 - Omar: Earlier proposals have been vertically integrated. Our goal was to disaggregate transport and networking to lend flexibility for independent evolution of the two. e.g Mohan chimed in that as feedback for SUE
- Comment: I think any L2/L3 header extensions/changes and functionality needs to be explicitly codified in a spec
- Not sure if this already answered, but is SAI going to be the API of choice, with additions as needed?
 - Any new functionality to Ethernet Fabric via ESUN, can be supported as enhancement to SAI
- Is the goal to make Ethernet usable for Scale Up or to actually compete with NVLink.
 - From Chad @ AMD!
 - ESUN will focus on offering a solid option for Scale-Up network that participating GPUs can use.
- Any target max. size (number of Accelerators) of one ESUN domain?
 - No specific limit was discussed.
- Elaborate on Hierarchical topologies?
 - As scale up domain crosses 1-rack topologies of connections will come in picture
- Comment: To make it easier, let's bucket questions into: 1) headers/encapsulation, 2) congestion and flow control, 3) link reliability/PHY, and 4) interoperability/testing. Please prefix your question with the bucket.
- Exactly. By standardizing headers, ECN/DSCP, LLR, and PFC/CBFC at the fabric, vendors can expose a clean Ethernet port while preserving transport flexibility.

Zoom AI Summary

Quick recap

The OCP ESUN Workstream kickoff meeting introduced a new process (access restrictions for specialized IP content) and outlined the process for signing the OCP feedback agreement (Michael). Manoj and Rajesh presented an overview of Ethernet's role in AI system networking, discussing the need to optimize Ethernet for AI systems while maintaining essential IP functionality. The group established regular biweekly meetings to discuss Ethernet solutions for scale-up environments, focusing on Ethernet fabric capabilities for scale-up communication needs.

Next steps

- [Michael: Post the meeting recording on the wiki](#)
- [Michael: Add a link to the OCP feedback agreement form on the ESUN Wiki this afternoon](#)
- [Michael: Include link to ESUN Wiki in the chat](#)
- [Michael: Drop email in the chat for questions about feedback agreement](#)
- [Michael: Add folks to ESUN mailing list as their organizations sign the feedback agreement](#)
- [Manoj and Rajesh: Send out a poll with meeting time options for regular biweekly meetings](#)
- [Manoj and Rajesh: Work on codifying/capturing starting point documentation on targeted improvements for feedback](#)
- [Michael: Send out first email with logistics information](#)
- [Michael/Team: Post presentation slides onto the wiki](#)

Summary

- Michel Schill (OCP) presentation on OCP ESUN Process ([slides](#))
- Manoj Wadekar (Meta) and Rajesh Sankaran (Microsoft) presentation on [ESUN Overview](#)

OCP ESUN Workstream Kickoff

The OCP ESUN Workstream kickoff meeting attracted 125 participants, with Michael explaining that future calls containing specialized IP content would be restricted to those who signed the OCP feedback agreement. He clarified that the agreement, which must be signed once per company or individual, allows for the use of provided feedback in any way by contributors while protecting participants' rights. Michael encouraged attendees to visit the ESUN Wiki to access the agreement

form and join the mailing list for future restricted calls, with the process following the OCP closed CLA model.

Ethernet for AI System Networking

Manoj and Rajesh presented an overview of Ethernet's role in AI system networking, focusing on scale-up networks where GPUs collaborate with unified memory views. They discussed the need to optimize Ethernet for AI systems, including improvements to bandwidth, packet handling, and reliability, while maintaining essential IP functionality. The presentation outlined driving principles for the ESUN workstream, emphasizing interoperability, collaboration with standards bodies, and leveraging Ethernet's mature ecosystem. The session concluded with logistics for joining the ESUN workstream and participating in future discussions.

Ethernet Solutions for Scale-Up Environments

The meeting focused on the establishment of regular biweekly meetings to discuss Ethernet solutions for scale-up environments, with plans to record open calls for later review. Rajesh and Manoj highlighted the importance of focusing on Ethernet capabilities while leveraging existing standards to ensure fast adoption. Omar provided context on the development process, emphasizing the collaboration with Microsoft and other companies to assess momentum and align efforts. Mohan stressed the need for speed in adopting solutions and aligning on standards, while Michael clarified that while individual OCP memberships are free, corporate memberships require a fee.

Ethernet Fabric Transport Requirements Discussion

The group discussed the focus on Ethernet fabric capabilities for transport communications, with Rajesh emphasizing Ethernet's strength in supporting various upper-layer protocols. Omar explained that while multiple transport implementations are expected, the goal is to agree on common Ethernet switch capabilities, similar to previous networking projects. The team agreed to iterate on requirements through open discussions before formalizing them in a specification, with Manoj and Rajesh already working on capturing initial thoughts. The group also addressed questions about timeline, architecture comparisons, and lower-level optimizations, with Omar and Manoj clarifying that Ethernet was chosen as a common denominator across transports and that PHY-level optimizations would be considered within existing standards work.

ESUN Fabric Functionality Planning

The group discussed the scope and goals of ESUN, focusing on separating fabric functionality from higher-level integration. They agreed to focus on independent fabric functionality before addressing higher level integration functions like in-network collectives. The discussion covered hierarchical

topologies for multi-rack connectivity and the potential to support multiple transports over Ethernet. The team acknowledged the need to define Ethernet types and liaise with IEEE for future capabilities as the need arises. They concluded by planning to continue discussions via email and schedule bi-weekly meetings for which time will be picked based on a survey.