

光进入机架之后，怎样让任务继续跑

OCP 第二天下午 · 独立现场记

宽并行光学 → 封装与连接 → 链路韧性

2026.10.14 · System Boundary

先把四种扩展放回各自的位置

01

Scale-in : 封装内、芯片与存储之间

02

Scale-up : 多个加速器形成紧耦合计算域

03

Scale-out / across : 更大网络与跨数据中心



Avicena : 把高速串行改成宽并行

01


μ LED 阵列 + PD 阵列 + ASIC
+ 微透镜 + 纤路

02

Scale-up 与 scale-in 共用
并行架构思路

03

主机若仍是高速 SerDes , 适配开销仍要计算



1 Tb/s 已进入评估；更高密度仍是目标

已披露评估平台

1 Tb/s 级 eKit
最多 335 通道
每路最高 3 Gb/s

现场方向

<1 pJ/bit
>200° C
>10 Tb/s/mm 或 mm²

待定义条件

功耗参考面
温度 / 距离 / 寿命
同时工作通道数

连接器正在进入宽并行光路

01


光纤束、多芯光纤与波导是不同实现

02

MPO 外形 ≠ 与传统 MPO 光学互通

03

核心是逐通道映射、损耗、串扰和维修



一排机架成为计算单元

计算机架

计算托盘 / 加速器
光 I/O 与冷却
可维护单元

光互连

跨机架纤路
前面板可换光源
结构化布线

交换机架

统一 scale-up 域
1024+ 参考目标
系统级联合验证

托盘里，光纤也是机械结构的一部分

系统约束

工作负载 / 拓扑
可维护单元
供电与冷却

托盘整合

光纤回路 / 弯曲
冷板与水管
装配和返修顺序

光电接口

光引擎 / 光源
耦合与热应力
逐通道测试

热协同设计仍要接受动态负载验证

静态模型

给定功率
材料与几何
冷却边界

动态验证

负载跃迁
温度时间轨迹
光学与电气响应

系统结果

BER / 重训练
任务进度
维护与循环寿命

Corning : 安装工时也会成为规模瓶颈

01


现场模型：约 1k → 45k
纤 / 机架

02

按 16F 跳线及每端作业时间计算

03

420 小时是模型工时，不是普遍部署实测



扩束连接：用光斑面积降低污染敏感性

01

PRIZM TMT：透镜、插芯与被动定位

02

现场比较：反复插拔且不清洁的 ΔIL

03

低污染敏感性不等于免清洁、免检查



从插得上，到制造得出来

01

16F 现有设计，32F 为路线图

02

现场：90 天制造导入、2 天既有产线转换

03

0.5 dB 是目标路径，非已保证规格



TeraHop : 把激光留在 NPO 里面

01

Diablo-1 : 6.4T NPO , 集成激光

02

专用冷板与独立热区域

03

现场宣称约 5.5 pJ/bit , 含内部激光



初步眼图与一公里演示说明了什么

01


Preliminary Test Data : 初步测试

02

约 12 dB 指 ASIC 至 NPO 电通道插损

03

1 km 照片不等于全温、全通道、长期验收



同一张功耗图，可能混着不同边界

01


器件 → 光引擎 → 端点 → 完整链路

02

反向 gearbox 与主机接口不能漏计

03

同方向、同吞吐、同距离再比较



Open CPX : 开放接口与路线图分开读

01


1.0 : 32/36 × 200G ;
6.4/7.2T

02

内置 / 外置光源、铜 / 光与模块管理

03

64/72 × 200G 及 400G/lane
属于后续计划



永久失效与短暂掉线，是两类统计

01


器件 FIT：永久失效与寿命数据

02

Flap-FIT：链路短暂 down/up 事件

03

Job impact：恢复是否赶得上任务超时



Motivation

- Optical Scale-Up allows 1-tier Pod size expanded up to the radix of the switch. 1024+ GPUs/XPUs per Pod are possible.
- Optics has its benefits and challenges. One challenge is Flaps.
- Flaps are intermittent link-down events, uS – minutes in duration.
- Some lanes might flap all the time and be persistently “down”.

OK, why do we care?

- Network Switches know how to deal with flaps for 30+ years. GPUs/XPUs don’t...
- If every flap causes a job crash,
we could see tens to hundreds of job crashes a day... not acceptable!

 OCP FUTURE TECHNOLOGIES SYMPOSIUM | 2026

规模放大的是暴露面，不是每颗器件突然变差

01

总事件率约等于各独立通道事件率之和

02

示例：1024 GPU、10k Flap-FIT

03

约 0.76 小时 MTBF；约 0.5 小时为 50% 首次事件时间

Some example stats...MTBF and Time (hrs) to 1st Flap

- Sweep over Flap-FIT rate (10-1M) and Pod size (16-1024 GPUs/XPUs)

MTBF (hrs)	16 GPUs	32 GPUs	64 GPUs	128 GPUs	256 GPUs	512 GPUs	1024 GPUs
10 Flap FIT	48828.13	24414.06	12207.03	6103.52	3051.76	1525.88	762.94
100 Flap FIT	4882.81	2441.41	1220.70	610.35	305.18	152.59	76.29
1000 Flap FIT	488.28	244.14	122.07	61.04	30.52	15.26	7.63
10000 Flap FIT	48.83	24.41	12.21	6.10	3.05	1.53	0.76
100000 Flap FIT	4.88	2.44	1.22	0.61	0.31	0.15	0.08
1000000 Flap FIT	0.49	0.24	0.12	0.06	0.03	0.02	0.01

Hours to 1st failure with P-fail = 0.5 (50%)	16 GPUs	32 GPUs	64 GPUs	128 GPUs	256 GPUs	512 GPUs	1024 GPUs
10 Flap FIT	33845.1	16922.5	8461.3	4230.6	2115.3	1057.7	528.8
100 Flap FIT	3384.5	1692.3	846.1	423.1	211.5	105.8	52.9
1000 Flap FIT	338.5	169.2	84.6	42.3	21.2	10.6	5.3
10000 Flap FIT	33.8	16.9	8.5	4.2	2.1	1.1	0.5
100000 Flap FIT	3.4	1.7	0.8	0.4	0.2	0.1	0.1
1000000 Flap FIT	0.3	0.2	0.1	0.0	0.0	0.0	0.0

Example highlighted (**NOT** indicative of field data):

- With a Flap-FIT rate of 10K, 1K GPU pod might hit 1st flap within first 30min!



2026

Marvell 的 3R : 活下来、恢复、减少重演

01


Redundancy : 独立路径与故障域隔离

02

Resiliency : 本地检测、快速切换

03

Reliability : 遥测、关联分析与群体修复



遥测的终点，是可执行的恢复动作

01


物理层 → 链路重试 → 链路层
→ 应用 → 系统

02

需要知道故障会持续多久、何时可恢复使用

03

告警、状态、隔离与回归策略
应统一



Ranovus : 平均 BER 掩盖不了突发错误

01

低延迟 FEC 意味着更低的原始误码要求

02

突发误码可能很短，却触发重试或掉线

03

温度、机械、污染、电源、固件都要入测试

Potential causes of burst errors

RANOVUS

	Possible Burst Error Contributors	Susceptible Configurations
Environmental (OE)	Mechanical shock, vibration, & warpage Humidity, thermal variation, & EMI Thermal & power supply transients Transient high optical power effects on Rx	Potentially all OEs
Pattern Sensitivity (Baseline Wander)	Laser biasing operating point dynamics Modulator operating point dynamics	For direct modulation implementations For external modulation implementations
Optical Interface & Fiber	Loss variation (optical connectors) Optical reflections (Tx & Rx I/O) Short fiber MPI Fiber movement (mechanical vibration & shock)	Due to mechanical or dust Connectors and I/O ports Multi-path interference for short fibers
Electrical Interface (High-speed)	Reflection interference for short electrical traces Humidity, thermal variation, & EMI Stability of electrical contacts / reliability changes	High-speed channels, (not D2D) For socketed implementations
Laser Input Stability	Impact of reflection variability into laser cavity Impact of stress, thermal variation for laser die Fiber (mechanical vibration, shock, & MPI)	Pre-isolator, or non-isolator implementations For external laser implementations
Firmware/Management	Varied causes	
Other (Higher level)	Timing issues / synchronization failures / auto negotiation Interoperability, incompatibility, & configuration issues Stability, degradation of PCBA components	Especially systems with different retimer/SerDes implementations For example, power supply components

- Brainstormed list of potential causes of burst errors which can lead to link flaps
- Short periods of burst errors can yet lead to retries yet have a minor impact on the random BER (depending on the measurement period)
- Some of these causes are
 - Detectable well before failure
 - Other causes may appear suddenly

Each implementation should use design techniques to minimize the potential for burst errors

14 Oct 2026 FTI Workshop 23

800G 端口：一条大管道，还是四条可切换路径

01

示例：1×800 模式与 4×200 模式

02

单 lane 失效后能否降级继续传输

03

关键在 NIC 内部的数据能否及时换路

GPU/XPU resiliency to flaps

- Key Q – when lane(s) flap, how we avoid a job abort?

Example

800Gbps 4-lane EP

800Gbps 4-lane EP

What happens if 1 lane flaps?

- 1x800 mode → Job abort
- 4x200 mode → degrade to 3x200 or 2x200

OCF FUTURE TECHNOLOGIES SYMPOSIUM | 2026

把共因故障画出来，再谈冗余收益

01

同光源、连接器、电源或交换芯片可能同失效

02

p^4 只适用于相应独立假设


03

分散故障域也会增加布线与管理复杂度

GPU/XPU resiliency to flaps

- Key Q – when lane(s) flap, how we avoid a job abort?

Example



Timely?

800Gbps 4-lane EP

800Gbps 4-lane EP

What happens if 1 lane flaps?

- 1x800 mode → Job abort
- 4x200 mode → degrade to 3x200 or 2x200

What happens if all 4 lanes flap together?

- NIC becomes blocked – no valid path
- Is MRC going to help?
- If traffic can be moved in **timely** fashion to another NIC, job survives, otherwise job crashes.

What is a likelihood of all 4 lanes flapping together?

- It depends on whether all lanes of a given EP belong to the same flap fault domain
- Multi-lane flaps do exist and can be quite common **within the same flap fault domain**.
- If all lanes per port are in **different flap fault domains**, with P as prob of 1-lane flapping, prob of all 4-lanes flapping per port is $\sim P^4$ – a much smaller value.

OCP FUTURE TECHNOLOGIES SYMPOSIUM | 2026

把三条宽并行路线放在同一套问题里

Credo

完整线缆系统
Gearbox / 监测
备用与管理通道

LightXcelerate

μ VCSEL / PD / EIC
19 芯 MCF 耦合
单 lane 与全阵列分开

Avicena

μ LED 宽并行
评估平台已披露
原生集成继续演进

四类供应商，可以从同一份验证表开始

光纤 / 连接器

映射与定位
损耗 / 串扰
弯曲与污染

材料 / 器件

热膨胀与应力
带宽与偏置
逐通道分布

模组 / 系统

完整功耗
遥测与故障域
任务恢复

结尾：带宽能交付，故障能收敛

能耗要算完整；指标要写边界

连接和散热要进入共同设计

每一次链路恢复，都要回到任务结果