TMTB: OCP Highlights - AMD's Lisa Su, NVIDIA on Rubin Ultra, META on 5GW AI Clusters, Broadcom on Ethernet, Credo on Optics, Astera Labs on Rack-Scale AI & More

TMTB：OCP 重点内容——AMD 的苏姿丰谈 Rubin Ultra，NVIDIA 谈 Rubin Ultra，META 谈 5GW AI 集群，Broadcom 谈以太网，Credo 谈光互连，Astera Labs 谈机架级 AI 及更多内容

TMT BREAKOUT
11 MINUTES AGO • PAID

Some good OCP talks 几场精彩的 OCP 演讲

- **AMD Keynote: CPU TAM, Helios shipments and the new agentic rack**

  AMD 主题演讲：CPU 总可寻址市场（TAM）、Helios 出货量与新型智能体机架

  Lisa Su (Chair & CEO, AMD).

  Lisa Su（AMD 董事长兼首席执行官）。

- **NVIDIA Keynote: Rubin deployment, dedicated CPU racks and optical NVLink**

  NVIDIA 主题演讲：Rubin 部署、专用 CPU 机架与光互连 NVLink

  Ian Buck (VP, Hyperscale & High-Performance Computing, NVIDIA).

  Ian Buck（NVIDIA 超大规模与高性能计算副总裁）。

- **Meta Networking: Ethernet scale-up, merchant switches and multi-building training**

  Meta 网络：以太网纵向扩展、商用交换机与跨楼宇训练

  Omar Baldonado (Director of Engineering, Meta).

  Omar Baldonado（Meta 工程总监）。

- **Credo: The optical roadmap beyond active electrical cables**

Credo：有源电缆之后的光互连路线图

Bill Brennan (President & CEO, Credo).

Bill Brennan（Credo 总裁兼首席执行官）。

- **Compute Meets Optics: OpenAI's architecture and the path into scale-up**

计算与光学的交汇：OpenAI 的架构及迈向规模化扩展之路

Binbin Guan (Member of Technical Staff, OpenAI); Manish Mehta (VP, Marketing & Operations, Broadcom); Tareq Bustami (VP, Market Technology, Celestica).

Binbin Guan（OpenAI 技术人员）；Manish Mehta（博通市场营销与运营副总裁）；Tareq Bustami（Celestica 市场技术副总裁）。

- **Astera Labs: Orion, rack-scale connectivity and a dedicated KV-cache tier**

Astera Labs：Orion、机架级互连与专用 KV 缓存层

Chris Petersen (VP & Technologist, Astera Labs); Dimitris Ziakas (CTO & VP, Architecture, Hyve Solutions).

Chris Petersen（Astera Labs 副总裁兼技术专家）；Dimitris Ziakas（Hyve Solutions 首席技术官兼架构副总裁）。

- **Google: 800V power, batteries, liquid cooling and flexible grid access**

Google：800V 供电、电池、液冷与灵活接入电网

Tom Garvens (VP, Hardware Platform Engineering, Google Cloud); Amber Huffman (Technologist, AI & Infrastructure, Google).

Tom Garvens（Google Cloud 硬件平台工程副总裁）；Amber Huffman（Google AI 与基础设施技术专家）。

- **Microsoft: Getting more compute from the power already available**

Microsoft：如何利用现有电力提升算力

Brijesh Warrier (VP, Software Engineering, Microsoft); Nithish Mahalingam (Partner Director of Engineering, Microsoft).

Brijesh Warrier（微软软件工程副总裁）；Nithish Mahalingam（微软工程合作伙伴总监）。

- **Meta Keynote: Building and operating a gigawatt-scale training cluster**

Meta 主题演讲：构建和运营吉瓦级训练集群

Santosh Janardhan (Head of Infrastructure, Meta).

Santosh Janardhan（Meta 基础设施负责人）。

- **Broadcom Keynote: Ethernet adoption and why copper lasts longer**

Broadcom 主题演讲：以太网的普及，以及铜缆为何能延续更久

Asad Khamisy (SVP & GM, Core Switching Group, Broadcom).

Asad Khamisy（Broadcom 核心交换事业部高级副总裁兼总经理）。

- **Arm Keynote: Three CPU markets inside the AI data center**

Arm 主题演讲：AI 数据中心中的三类 CPU 市场

Mohamed Awad (EVP & GM, Cloud AI Business Unit, Arm).

Mohamed Awad（Arm 云与 AI 业务部门执行副总裁兼总经理）。

- **AMD Helios Technical: Memory, integrated networking and rack-level co-design**

AMD Helios 技术专题：内存、集成网络与机架级协同设计

Andrew Dieckmann (Corporate VP & GM, AMD); Steve Scott (Corporate Fellow, AMD); Alan Smith (Corporate Fellow, AMD).

Andrew Dieckmann（企业副总裁兼总经理，AMD）；Steve Scott（企业研究员，AMD）；Alan Smith（企业研究员，AMD）。

- **Nokia: AI networking beyond the data center**

诺基亚：数据中心之外的 AI 网络

Clayton Wagar (Principal Consulting Engineer, Nokia).

Clayton Wagar（首席咨询工程师，诺基亚）。

- **Marvell / Microsoft / Oracle: Moving hardware security into every compute node**

Marvell / Microsoft / Oracle：将硬件安全扩展至每个计算节点

Cary Ussery (SVP & GM, Multimarket Business Group, Marvell); Mohan Kumar (Architecture & Ecosystems, Oracle Cloud Infrastructure); Jayant Gandhi (Security

Architect, Microsoft Azure).

Cary Ussery（高级副总裁兼总经理，多市场业务集团，Marvell）；Mohan Kumar（架构与生态系统，Oracle Cloud Infrastructure）；Jayant Gandhi（安全架构师，Microsoft Azure）。

---

1. AMD Keynote: A much larger CPU opportunity, Helios entering production and a separate rack for agents

1. AMD 主题演讲：规模大幅扩张的 CPU 机遇、Helios 投产，以及面向智能体的独立机架

Speaker: Lisa Su, Chair & CEO, AMD.

演讲者：Lisa Su，AMD 董事长兼首席执行官。

Summary: AMD's most consequential claims were that its CPU TAM estimate has been raised three times in nine months and now exceeds $200B by 2030, and that Helios production shipments started in September. The new Helion rack turns the agentic-CPU thesis into a product, while Verano extends the CPU roadmap toward lower-power, memory-intensive configurations.

摘要：AMD 最重要的表述是，其 CPU 总潜在市场（TAM）估算在九个月内上调了三次，目前预计到 2030 年将超过 2,000 亿美元；此外，Helios 已于 9 月开始量产出货。全新的 Helion 机架将智能体 CPU 这一论点转化为产品，而 Verano 则拓展了 CPU 路线图，面向功耗更低、内存密集型的配置。

1. The CPU market is being redefined by agent execution, not just GPU hosting

1. CPU 市场正因智能体执行而被重新定义，而不仅仅是 GPU 托管

Su put the data-center AI accelerator opportunity at approximately $1.5T by 2030, but the sharper change was CPUs: AMD now sees a market exceeding $200B by 2030, growing at roughly a 50% CAGR over the next several years. These are AMD's market forecasts, not Street estimates.

苏姿丰预计，到 2030 年，数据中心 AI 加速器市场规模约为 1.5 万亿美元；但更显著的变化在于 CPU：AMD 目前预计，到 2030 年，CPU 市场规模将超过 2,000 亿美元，未来几年复合年增长率（CAGR）约为 50%。这些是 AMD 对市场的预测，并非华尔街一致预期。

The underlying argument is that agents create substantial work outside model inference: executing code, querying applications, managing context and coordinating other agents. That adds a CPU workload alongside the processors already supporting GPUs. The keynote's older ~$20B market characterization should not be treated as a directly comparable starting point for the new CAGR.

其核心逻辑是，智能体会在模型推理之外产生大量工作负载，包括运行代码、查询应用程序、管理上下文以及协调其他智能体。这使 CPU 工作负载有所增加，与现有支持 GPU 的处理器形成互补。演讲中此前将市场规模描述为 ~$20B，这一数字不应被视为新 CAGR 的直接可比基准。

2. Helios has moved from launch to initial customer deployment

2. Helios 已从发布阶段迈入初步客户部署

Su said production shipments began in September, with the ramp continuing through Q4 and 1H27. AMD said initial systems are installed and running at OpenAI, Meta, Microsoft and Oracle, with deliveries, deployments and workload validation underway.

苏姿丰表示，Helios 于 9 月开始量产出货，预计在第四季度和 2027 年上半年继续提升出货规模。AMD 表示，初始系统已在 OpenAI、Meta、Microsoft 和 Oracle 安装并运行；相关交付、部署和工作负载验证正在进行。

That is a meaningful execution milestone, but initial installations are not equivalent to completing the previously announced multi-gigawatt commitments. The next question is the pace at which those installations convert into repeatable, large-scale deployment.

这是一个意义重大的执行里程碑，但初始安装并不等同于此前宣布的数吉瓦级承诺已全部兑现。接下来的关键问题是，这些初始安装将以多快的速度转化为可重复、规模化的部署。

3. Helion targets a different workload from Helios

3. Helion 面向的工作负载与 Helios 不同

Helion is designed around agentic execution, with CPU-heavy configurations for hyperscalers and integrated CPU/GPU configurations for enterprises. AMD's published specification supports up to 32,768 CPU cores, while an integrated model router sends requests to the appropriate model and compute resources.

Helion 的设计围绕智能体式执行展开：面向超大规模云服务商的配置以 CPU 为主，面向企业的配置则集成 CPU/GPU。AMD 公布的规格显示，其最多支持 32,768 个 CPU 核心；集成式模型路由器则会将请求发送至合适的模型和计算资源。

The distinction is deliberate: hyperscalers may operate separate inference and agent-execution racks, while an enterprise running smaller models on premises may prefer both functions together. Partner configurations are planned for November 2026. This is a second deployment architecture, not simply a smaller Helios.

这种区分是有意为之：超大规模云服务商可能会分别部署推理机架和智能体执行机架；而在本地运行较小模型的企业，可能更倾向于将两种功能集成在一起。合作伙伴配置计划于 2026 年 11 月推出。这是第二种部署架构，并非只是缩小版的 Helios。

4. Verano makes memory and system power part of the CPU debate

4. Verano 将内存和系统功耗纳入 CPU 讨论

AMD showed Verano, a Venice-family EPYC variant optimized for lower system power, with an LPDDR-based memory approach and updated CPU-to-GPU connectivity. Production is planned for 2027.

AMD 展示了 Verano，这是 Venice 系列 EPYC 处理器的一个变体，针对降低系统功耗进行了优化，采用基于 LPDDR 的内存方案，并更新了 CPU 与 GPU 之间的互连。该产品计划于 2027 年投产。

The important distinction is that AMD is not treating every agent-related CPU workload as a contest for maximum core count. Host processors, dense agent execution and technically demanding agents require different balances of memory bandwidth, connectivity, per-core performance and power.

关键区别在于，AMD 并未将所有与智能体相关的 CPU 工作负载[illegible]心数量最大化的竞赛。主机处理器、高密度智能体执行以及技术[illegible]能体，需要在内存带宽、互连、单核性能和功耗之间作出不同权



5. Rack economics depend on the surrounding system

5. 机架经济性取决于整个配套系统

AMD claimed a 30% tokens-per-dollar advantage for Helios, but the keynote did not provide enough comparison detail to extrapolate that advantage across workloads. Its broader engineering argument was more durable: compute, memory capacity, memory bandwidth, networking and software all need to improve together.

AMD 声称 Helios 的每美元令牌数优势达 30%，但主题演讲没有提供足够的对比细节，无法据此将这一优势外推至各种工作负载。其更具持续性的整体工程论点是：计算能力、内存容量、内存带宽、网络和软件都需要协同提升。

The investment implication is twofold: AMD is addressing a larger CPU opportunity while moving its GPU proposition toward a complete deployable system. Neither the TAM forecast nor a vendor benchmark establishes the eventual revenue share or margin contribution.

这一投资逻辑包含两个层面：AMD 正在把握更大的 CPU 市场机遇，同时推动其 GPU 方案向完整、可部署的系统演进。TAM 预测和供应商基准测试均无法确定最终的收入份额或利润率贡献。

---

2. NVIDIA Keynote: Shipping infrastructure today, a larger optical NVLink domain next, and CPUs as a separate product

2. NVIDIA 主题演讲：当前交付基础设施，下一步扩展光互连 NVLink 域，并将 CPU 作为独立产品

Speaker: Ian Buck, VP, Hyperscale & High-Performance Computing, NVIDIA.

演讲嘉宾：Ian Buck，NVIDIA 超大规模与高性能计算副总裁。

Summary: NVIDIA combined current Vera Rubin deployment with a preview of a 576-GPU Rubin Ultra NVLink domain using co-packaged optics. Equally important was the dedicated Vera CPU rack for agent execution, plus software and power-management work intended to increase usable compute inside constrained facilities.

摘要：NVIDIA 将当前 Vera Rubin 部署与 Rubin Ultra 的预览结合起来。Rubin Ultra 将采用共封装光学技术，构建一个由 576 颗 GPU 组成的 NVLink 域。同样重要的是，NVIDIA 推出了用于智能体执行的专用 Vera CPU 机架，并开展软件和电源管理方面的工作，以提升受场地条件限制的数据中心内的可用算力。

1. Separate Vera Rubin deployment from the Rubin Ultra preview

1. 将 Vera Rubin 的部署与 Rubin Ultra 的预览区分开来

Buck described Vera Rubin as being in production and reaching customer data centers. The subsequent Rubin Ultra preview was a different milestone: a planned 576-GPU NVLink-connected system spanning multiple racks, using a next-generation co-packaged-optics switch tray.

Buck 表示，Vera Rubin 已进入量产，并正在进驻客户数据中心。随后预览的 Rubin Ultra 则代表另一项进展：计划打造一个由 576 颗 GPU 通过 NVLink 互连、跨多个机架部署的系统，并采用新一代共封装光学交换机托盘。

The read-through is that NVIDIA intends to expand the tightly coupled accelerator domain beyond today's copper-connected rack. The preview supports a future optical-content opportunity, but it should not be counted as a currently shipping 576-GPU configuration.

由此可见，NVIDIA 有意将紧密耦合的加速器域扩展至当前铜缆互连机架之外。这一预览表明未来存在光学组件的市场机遇，但不应将其视为目前已交付的 576-GPU 配置。

2. The Vera CPU rack is not merely a GPU accessory

2. Vera CPU 机架并非只是 GPU 的配套产品

Buck described a dedicated rack containing 256 Vera CPUs for agent execution and related workloads. NVIDIA's architecture separates its role from the Vera Rubin NVL72 training/inference system: the CPU rack supports tool calling, orchestration and sandboxed execution.

Buck 介绍了一个专用机架，其中配备 256 颗 Vera CPU，用于执行智能体及相关工作负载。NVIDIA 的架构将其与 Vera Rubin NVL72 训练/推理系统的功能区分开来：CPU 机架支持工具调用、编排和沙箱化执行。

The architectural preference is strong sustained per-core performance, memory

access and predictable throughput, rather than optimizing solely for the lowest cost per core. That gives NVIDIA another route into data-center spending, including tasks that do not run directly on its GPUs.

该架构优先考虑强劲且持续的单核性能、内存访问能力和可预测的吞吐量，而非一味追求最低的单核成本。这为 NVIDIA 进一步参与数据中心支出提供了另一条路径，包括承接不直接运行在其 GPU 上的任务。

3. Co-packaged optics addresses scale-out first, then larger scale-up domains

3. 共封装光学首先解决横向扩展问题，随后拓展至更大的纵向扩展域

Spectrum-X Photonics was presented as a shipping-year scale-out product, while optical NVLink was attached to the next-generation scale-up roadmap. Those are different addressable markets and adoption schedules.

Spectrum-X Photonics 被定位为将在当年出货的横向扩展产品，而光学 NVLink 则被纳入下一代纵向扩展路线图。两者面向不同的目标市场，采用时间表也各不相同。

Buck also emphasized multiplane networking and rapid hardware-level failure detection. The economic objective is not simply higher bandwidth: it is avoiding situations where a bad link slows or interrupts a much larger pool of expensive accelerators. His power-efficiency and switch-count improvements were vendor claims tied to particular architectural comparisons.

Buck 还强调了多平面网络以及在硬件层面快速检测故障的能力。其经济效益目标并非单纯提升带宽，而是避免链路故障导致规模大得多、成本高昂的加速器集群降速或中断运行。他提到的能效提升和交换机数量减少，是针对特定架构对比提出的厂商主张。

4. Power management can unlock capacity without a new utility connection

4. 电力管理可在无需接入新的公用电网电源的情况下释放容量

NVIDIA’s MaxLPS discussion focused on the gap between nameplate power and actual workload consumption. If software continuously monitors demand and enforces electrical limits, operators can potentially install or use more compute within the same provisioned power envelope.

NVIDIA 围绕 MaxLPS 的讨论重点关注铭牌功率与实际工作负载功耗之间的差距。如果软件能够持续监控需求并严格控制用电上限，运营商便有望在既有供电容量范围内部署或使用更多算力。

This is a utilization opportunity, not the creation of additional electricity. The amount recoverable depends on workload diversity, existing provisioning conservatism and control-system responsiveness. The transcript's example contained inconsistent figures, so the mechanism is more reliable than a precise fleet-wide uplift.

这是提升利用率的机会，而不是新增发电量。可回收的电力取决于工作负载的多样性、现有配置中的保守余量以及控制系统的响应速度。访谈中的示例数据前后不一致，因此其机制比精确估算整个机群的提升幅度更可信。

5. Manufacturing and deployment are becoming part of platform differentiation

5. 制造与部署正成为平台差异化的一部分

Buck contrasted roughly 90 minutes to assemble an earlier compute tray with about one minute for the newer cableless tray design. The comparison concerned a compute tray, not an entire completed rack or commissioned cluster.

Buck 将早期计算托盘约 90 分钟的组装时间，与新款无缆线托盘设计约 1 分钟的组装时间进行了对比。这一比较针对的是计算托盘，而非整台已完成组装的机架或已调试投运的集群。

The broader MGX contribution standardizes rack, cooling, power and service interfaces. NVIDIA is trying to make manufacturing, installation and servicing repeatable across partners while retaining its system-level co-design advantage. The 800V discussion extends that effort into facilities, rather than assuming every existing data center can accept the next rack unchanged.

MGX 的更广泛贡献在于标准化机架、冷却、电力和维护接口。NVIDIA 正努力让合作伙伴能够以可重复的方式开展制造、安装和维护，同时保留其系统级协同设计优势。800V 相关讨论将这项工作延伸至数据中心基础设施，而不是假设所有现有数据中心都能直接接纳下一代机架。

---

3. Meta Networking: Ethernet scale-up is already deployed, and merchant silicon extends