Broadcom Inc

Mgmt Meeting: TPU “v9” Roadmap on Track, AAPL Extension/Expansion Includes New ASICs + RF, Next-Gen OpenAI AI XPU Taping Out “Very Soon”, Tomahawk 6 Demand Outstripping Supply

Overweight

AVGO, AVGO US
Price (07 Jul 26):$370.78

Semiconductors & Semiconductor Capital Equipment / IT Hardware

Harlan Sur [AC]
(1-415) 315-6700
harlan.sur@jpmorgan.com

Mayur Ramdhani
(1-212) 622-1664
mayur.ramdhani@jpmorgan.com
J.P. Morgan Securities LLC

Yesterday, we hosted an investor group meeting with Broadcom's President and CEO, Hock Tan, CFO, Amie Thuener, and head of IR, Ji Yoo. The discussions centered around accelerating custom AI XPU demand driven by rapid growth in inference workloads, the expansion of Broadcom’s decade-plus Apple relationship through 2031, continued strength in AI networking demand, and Broadcom’s differentiated ASIC design, IP, and advanced packaging capabilities. Key takeaways in our view were: 1) inference demand is driving a faster custom XPU adoption curve, with management suggesting that among the largest frontier model builders, XPU/GPU unit shipments could approach a ~50/50 mix next year; 2) customers are attempting to pull in next-generation XPU tape-out schedules and broaden chip design variants as inference workloads evolve from one-shot to reasoning and agentic use cases; 3) Broadcom/OpenAI’s first-gen XPU program, Jalapeño, further validates Broadcom’s ASIC design and execution strength (initial design to tape-out in just 9 months), with management expecting the follow-on/next-gen chip to tape out very soon; 4) competitive concerns around hyperscaler customer-owned tooling (COT) appear overstated given the complexity of multi-die XPU design, advanced packaging, and proprietary IP; 5) Broadcom’s next-gen TPU (v9) roadmap remains on track to ramp in CY28, with working 400G SerDes silicon already developed; 6) AI networking demand remains extremely strong, with Tomahawk 6 sold out and Broadcom prioritizing XPU customers for switch allocation; 7) the Apple agreement extends and expands Broadcom’s design/supply relationship through 2031, and includes both RF and ASICs (which we believe will include an AI XPU ASIC accelerator; see here); 8) VMware remains a longer-dated but strategically relevant AI-on-prem opportunity; and 9) management remains committed to maintaining ASIC margins relatively stable and within target despite input cost inflation and supply constraints across wafers, substrates, and HBM. Overall, we come away with increased confidence in Broadcom’s AI compute and networking growth runway, supported by accelerating inference demand, expanding strategic customer engagements, and Broadcom’s strong technology leadership across ASIC design, SerDes, packaging and Ethernet switching. We see accelerating AI fundamentals combined with an improving non-AI semiconductor business and further VMware revenue unlock driving strong revenue and earnings growth. Broadcom is the #2 global AI semiconductor supplier, #1 custom chip ASIC supplier, and #1 networking semiconductor supplier, and we reiterate our Overweight (OW) rating.

- **Apple ASIC partnership expansion reinforces Broadcom’s long-duration ASIC franchise — potentially expanding into AI/datacenter ASICs; RF supply agreement extended through 2031 as well…** The team emphasized that Broadcom has been engaged with Apple for multiple generations of ASIC products over 10+ years, including touch controller ASICs, wireless charging ASICs and other design sockets. The recently announced agreement extends and expands the ASIC design relationship through 2031, covers new design

See page 4 for analyst certification and important disclosures.

sockets going forward (including AI/datacenter ASICs, in our view), and includes the prior RF supply agreement as part of the broader extension/expansion through 2031. Management also pointed to the Fort Collins fab expansion as being indicative of Broadcom's confidence in the long-term opportunity and return profile.

- **Broadcom's next-gen TPU (v9) roadmap remains on track; 400G SerDes already developed and running.** On the specific question of whether the next-generation TPU program is delayed or cancelled, management indicated that the roadmap is "totally on track" - consistent with our own primary research but notable given recent concerns around competitive TPU programs and potential share shifts. The team also indicated that Broadcom has already developed working 400G-per-lane SerDes capable of driving copper over one meter (and believes Broadcom is currently the only company with this capability).

- **5-year agreement with Google includes increasing TPU revenues per year and a majority volume share.** While it was widely known that AVGO's revenues would scale each year under its 5-year contract with Google, management incrementally confirmed that it expects to retain a majority volume share through the course of the agreement and that this was a key reason AVGO entered into the partnership. In our view, this confirms that AVGO is anchored as the primary supplier across the 5-year term, which should ease widespread investor concerns around longer-term share-loss.

- **Inference demand is accelerating the custom XPU adoption curve.** Management indicated that true compute demand (ex-China) is concentrated among roughly five frontier model builders that account for the vast majority of demand. While hyperscaler infrastructure can appear GPU-skewed at the aggregate level, management suggested that, when honing in specifically on leading frontier model builders, the mix could be roughly 50/50 between XPUs and GPUs next year (in line with our view that the shift from training-dominated workloads to inference monetization is pulling forward custom accelerator/XPU adoption).

- **Customers are attempting to pull in XPU schedules and pursue more design variants.** The team confirmed that customers are attempting to accelerate next-generation ASIC XPU tape-out schedules and are pursuing a wider variety of chip designs. This reflects the rapid evolution of inference workloads from one-shot inference to reasoning and agentic workloads, as well as customer focus on lowering cost per token (which in our view increases the strategic importance of workload-optimized custom silicon).

- **OpenAI/Jalapeño highlights Broadcom's ASIC execution capabilities.** Management discussed OpenAI's first-generation XPU custom accelerator, Jalapeno, co-designed by Broadcom, and emphasized that the complexity of developing these cutting-edge XPUs is often underestimated. These products "push the envelope" across multiple design and technology dimensions, including bandwidth, SerDes, power and packaging. Broadcom expects to tape out the follow-on/next-gen OAI chip very soon.

- **COT competitive concerns appear overstated; management believes COT TPU v8t is still not optimized and ready for production deployment…** Management pushed back on concerns that hyperscaler customer-owned tooling (COT) efforts will materially displace Broadcom's XPU/TPU share. The key point was that modern XPUs are no longer simple single-die chips; they increasingly require multi-die configurations, advanced packaging, high-speed SerDes and deep productization expertise. Management referenced 6-die and 8-die configurations as examples of the complexity involved and suggested that internal hyperscaler efforts are likely to face significant technical hurdles.



- **Tomahawk 6 demand remains extremely strong; XPU customers are prioritized for AI networking allocation.** Management indicated that all of its major frontier model builder customers are adopting Tomahawk 6 (which is effectively sold out), with unconstrained demand characterized as being meaningfully higher than available supply. The team also noted that its XPU customers do receive priority when allocating Tomahawk and Jericho switching products. In our view, this reinforces the strategic value of Broadcom's full-stack AI silicon portfolio: success in custom compute can strengthen networking share, while scarcity in networking can further deepen customer relationships around XPUs.

- **No near-term oversupply concerns for AI compute.** In response to questions around hyperscalers potentially leasing or reselling excess compute, management said it does not see an oversupply of compute for 2027 or 2028. Management also noted there is no reason hyperscalers could not rent out custom XPUs to third parties, but this would not materially alter Broadcom's P&L as Broadcom sells the chips and customers determine how to deploy or monetize them.

- **Management intends to maintain stable ASIC GMs despite input cost inflation.** On rising component costs, including memory, the team is aiming to keep ASIC margins intact "one way or another," emphasizing that technology leadership is the key source of pricing power. Management also pushed back on the idea of customers consigning HBM, noting that HBM is an integral part of the XPU design and must be tightly connected to the core logic die for the product to perform properly.

- **VMware AI-on-prem opportunity remains early but relevant.** Management indicated that it is still early for enterprises to run AI workloads on-prem, but believes that as this opportunity develops, VMware is well positioned as a software stack for running AI workloads in enterprise data centers. Management has seen some early signs of this use case and is leveraging this capability in VMware customer conversations.