US Semiconductors

SpaceX's Terafab - Putting Things In Perspective

Equities

Americas

Semiconductors

Timothy Arcuri
Analyst
timothy.arcuri@ubs.com
+1-415-352 5676

Gianmarco Vella
Associate Analyst
gianmarco.vella@ubs.com
+1-415-352 4555

Alex Kivali
Analyst
alex.kivali@ubs.com
+1-212-713 3945

Natalia Winkler, CFA
Analyst
natalia.winkler@ubs.com
+1-415-352 4626

Aaryan Wadhwa
Associate Analyst
aaryan.wadhwa@ubs.com
+1-212-821 6481

The What, Why, and When of Musk's Terafab plans

In conjunction with UBS' initiation of SpaceX (covered by John Hodulik), we outline the size, scale, and potential impact to the broader Semi Equipment supply chain. While it is fair to question the magnitude, Elon Musk has set forth a clear vision to build a vertically integrated semiconductor complex capable of producing terawatt-scale AI compute to resemble what has been demonstrated with Tesla's battery strategy (i.e., the Gigafactory). In that case, when the industry could not meet Tesla's battery capacity requirements in 2014, the company effectively built its own fully integrated battery production supply chain. Similarly, Terafab appears to target the full semiconductor complex including mask shop, leading foundry/logic & memory process flows, packaging and testing, all under one roof - enabling rapid "design/build/test/redesign" cycles (Figure 1). Importantly, we believe our supply chain datapoints appear more credible than many investors still seem willing to underwrite. Grimes County (TX) public hearing notices, recent executive hires and an extensive list of Terafab-related process engineering jobs are now being further substantiated by our discussion with equipment suppliers indicating orders having been placed for a pilot line that looks to us to be something in the ~$5B range in C2027. Taken at face value, Terafab will essentially grow to something in the range of TSMC's scale in terms of WFE within 4-5yrs - obviously, a huge game changer for how high WFE can go over the coming years and putting $300B + WFE in sight for C2029. We expect this to be a major theme over the next few earnings seasons as the project takes shape and tool orders are placed.

WFE math for Terafab - A new >$50B/yr WFE spender

UBS' analyst John Hodulik forecasts SpaceX's AI business alone (ex-Connectivity and ex-Launch segments) to spend a total of ~$1.1T in capex over the next 5yrs, with ~20% of this capex (or ~$225B) dedicated to Terafab. Assuming the typical 60% capex split to WFE, this would imply ~$135B in WFE spend over the next 5yrs - about the size of the entire global WFE TAM this year. Because this will grow over time and start with a pilot line that we believe to be in the ~$5B range for C2027 and grow to ~$10B in C2028 (already assumed in our ~$250B WFE for '28), the math would suggest that Terafab's WFE spending would have to grow above ~$50B/yr in the C2030/2031 timeframe - which is akin to adding spending in the same ballpark as TSMC's WFE spend over the next few years. Based on the magnitude of the investment and existing filings thus far (specifically, plans that outline an initial $55B semiconductor fab investment by SpaceX with the potential to expand to $119B if all planned phases are completed (all in Grimes County, TX)), we infer SpaceX could be initially planning for a very large fab cluster - potentially something in the range of ~80k wsm for memory and two ~20k wsm foundry/logic fabs, alongside a mask shop and back-end test/assembly capacity. Our assessment is based on Intel's Ohio project (leading-edge greenfield fabs without a pre-existing semiconductor ecosystem) which initially carried an estimated ~$28B investment and, similarly, Boise ID1 fab requiring a ~$15B investment (for ~80k wsm of capacity). We would note that Terafab has indicated an initial focus on memory prior to building out other capabilities - we discuss this more in light of the fact that it has already engaged with INTC (see below). This is all consistent with our recent checks suggesting that INTC may be helping it to secure some slots at toolmakers. Ultimately, even if part of the potential Terafab WFE spend was to materialize, it would reinforce our long-held conviction (here) that industry WFE could approach ~$300B by C2029.

Read to INTC + Memory suppliers (MU/SK Hynix/Samsung)

We think of INTC's engagement with SpaceX as potentially serving as the "knowledge owner" in the Terafab - in other words, being potentially structured as a technology licensing arrangement, similar to historical models such as the IBM-AMD technology

This report has been prepared by UBS Securities LLC. **ANALYST CERTIFICATION AND REQUIRED DISCLOSURES, INCLUDING INFORMATION ON THE QUANTITATIVE RESEARCH REVIEW PUBLISHED BY UBS, BEGIN ON PAGE 4.**

transfer framework that ultimately supported GFS. Under such a structure, INTC could provide access to process flows, manufacturing IP, PDKs, design rules, tool recipes, and potentially operational know-how, while retaining ownership of the underlying technology and collecting a license/royalty payments. Another potential scenario (upon successful pilot) could be INTC dedicating an entire site to Terafab in a JV format. This could be its "Ohio One" site, which is large enough to support two (2) leading-edge fabs - as one of Terafab's campuses. We could also foresee a scenario in which SpaceX invests in INTC further down the road, with the likelihood increasing if Terafab ultimately proved unsuccessful. On the memory front, Musk's comments (Figure 1) explicitly call out memory as part of the effort, although a lot of questions remain unanswered here - most notably, we are unclear where Terafab would source memory IP and see limited incentives for existing memory suppliers to license their IP. However, the emergence of another potential supplier capable of producing memory at scale - if successful - could ultimately prompt Korean companies to invest in US front-end manufacturing for memory (more likely for DRAM – more here), also noting that, for example, Samsung has ample land available to add capacity in Taylor, TX. Overall, we continue to view this development as supportive of SPE suppliers.

Figure 1: Notable comments from the Terafab presentation (Mar-2026)

Topic	Commentary
Today's compute capacity	*The current output of AI compute is roughly 20 GW per year [...] we need to build the terafab because all of the rest of the output from Earth is about 2% of what we need. So* ***if you add up all the fabs on Earth combined, they're only about 2% of what we need for the terawatt*** *project or Terafab project.*
Capacity expansions by TSMC, Samsung, Micron	*[...] to be clear, we're very grateful to our existing supply chain — to* ***Samsung, TSMC, Micron and others. And we would like them to expand as quickly as they can.*** *[...]* ***But there's a maximum rate at which they're comfortable expanding****, but that rate is much less than we would like. And* ***so we either build the terafab or we don't have the chips*** *and we need the chips so we're going to build the terafab. And we're starting off with an advanced technology fab here in Austin.*
100-200GW of Terrestrial compute, rest in Space	*[...] for the space compute my guess is that is the vast majority of the compute because you're power constrained on Earth.* ***That's why I think it's probably 100-200 GW a year of terrestrial chips and probably on the order of a terawatt of chips in space.*** *Just because of power constraints on the ground is probably how it ends up.*
Silicon for Inference first, and then Space	*[...] we expect to make two kinds of chips.* ***One will be optimized for edge inference.*** *So that'll be used primarily in Optimus [...] because I expect the robots — humanoid robots — to be made 10 to 100 times more than the volume of cars. [...]* ***And then we need a high power chip that is designed for space.*** *[...] It's a hostile environment in space. So you want to design the chip — you want to optimize it for space. And* ***you also want to generally run it a little hotter than you would normally run a chip on Earth to minimize the radiator mass.***
Logic, Memory, Packaging, Test, Masks	*In the advanced technology fab, we* ***will have all of the equipment necessary to make a chip of any kind, logic or memory, and we will also have all of the equipment necessary to make the lithography masks****. So in a single building we can create a lithography mask, make the chip, test the chip, make another mask and have an incredibly fast recursive loop for improving the chip design [...] you've got* ***everything necessary [...] and just keep looping it***

Source: Company Presentation

Valuation Method and Risk Statement

We use various valuation techniques such as P/E, EV/FCF for valuing the companies in this report. Risk factors include but are not limited to macroeconomic factors such as a downturn in the economy, a disruption of international trade, technological disruption due to new inventions, or business model innovation whereby structural changes in the industry alter the future course of unit sales, ASPs, and revenues.

Intel Corp.:
Our PT is based on a SOTP analysis of the five contributing segments (Intel Products, Intel Foundry, Altera, Mobileye and All Other) each one valued using a tailored P/S multiple applied to F2026E revenues. NVDA has built a formidable moat for new compute-intensive workloads in the data center and could ultimately leverage its GPU architecture to more broadly displace INTC. New client and server CPUs from AMD also present a threat that we could be underestimating.

SpaceX:
Our price target is SOTP and multiples based. Risks include dependence on Elon Musk as a key decision-maker, delays in achieving technological milestones, failure to scale and monetize AI initiatives, regulatory, geopolitical, and competitive pressures, and high capital intensity with execution risks around supply chains and unproven space-based compute models.

