AMERICAS TECHNOLOGY: SOFTWARE

Highlights from GS Comm+Tech conversations

Gabriela Borges, CFA
+1(212)902-7839 | gabriela.borges@gs.com
Goldman Sachs & Co. LLC

Matthew Martino
+1(212)902-0695 | matt.martino@gs.com
Goldman Sachs & Co. LLC

Callie Valenti
+1(212)357-2790 | callie.valenti@gs.com
Goldman Sachs & Co. LLC

Max Gamperl
+1(415)249-7311 | max.gamperl@gs.com
Goldman Sachs & Co. LLC

Maura Hager
+1(212)9028724 | maura.hager@gs.com
Goldman Sachs & Co. LLC

Noah Naparst
+1(917)343-6395 | noah.x.naparst@gs.com
Goldman Sachs & Co. LLC

Greyson Sklba
+1(212)357-2534 | greyson.sklba@gs.com
Goldman Sachs & Co. LLC

Selina Zhang
+1(212)357-9979 | selina.zhang@gs.com
Goldman Sachs & Co. LLC

We hosted ~40 Software companies at our Communacopia + Technology Conference last week. Investor sentiment was better than at any time in the prior eight months and has moved to be significantly more micro/idiosyncratic focused, with questions and commentary becoming much more product and tech stack specific. Generally speaking, for companies with exposure to the inference economy, there are good reasons to believe that the best part of AI product cycles are still ahead, given how early enterprises are in scaling production inference use cases. At the same time, there is a fair amount of investor discomfort in a) the move higher in application software stocks, without a commensurate move higher in fundamentals or good visibility into how the next 12 months will unfold from an estimates revisions standpoint; b) the potential for a wildcard change in the security narrative or underlying architectures, such as from frontier models entering the security TAM. Most notable incremental positive datapoints in our view were for CRWD, DDOG, DOCN, DT, SHOP and SNOW; more mixed datapoints/low visibility in our view were for CHKP, HUBS, INTU and NAVN.

Infrastructure

- **AI infrastructure demand remains exceptionally strong, even when accounting for mix shift between frontier and open source/weight tokens. Enterprises are increasingly moving AI workloads into production and treating compute as a strategic resource.**
  - **CoreWeave** highlighted enterprise customer wins (Caterpillar, IBM), noting that customers increasingly view AI infrastructure as a strategic asset and are building dedicated clusters to support internal models and inference workloads.
  - **DigitalOcean** pointed to increasing production adoption across coding, customer experience, personal productivity, demand generation, and generative media use cases. Management argued that inference demand is increasingly tied to real business outcomes and monetization rather than venture funded experimentation, suggesting the market is progressing from proof of concept to production deployment.

Goldman Sachs does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. For Reg AC certification and other important disclosures, see the Disclosure Appendix, or go to www.gs.com/research/hedge.html. Analysts employed by non-US affiliates are not registered/qualified as research analysts with FINRA in the U.S.

- **Akamai** highlighted sustained customer demand extending into 2028, with opportunities spanning both GPU and CPU infrastructure. Management emphasized that demand is no longer limited to model training and increasingly includes the supporting infrastructure required to operationalize AI applications.
- **Microsoft's** comments on capex were balanced: Microsoft already has a framework for the right long-term investments in land and power, and a good understanding of the quantum of capacity it can bring on in the next year. Instead, its focus is on how quickly it can get new supply "revenue ready". The company reiterated its >50% reduction in dock-to-live times over the past year and pointed to efficiency gains across hardware, Azure infrastructure, and application layers that are expanding effective supply without a commensurate increase in capital intensity. With the amount of GW of new infrastructure able to come online over the NTM largely fixed, Microsoft's focus is on accelerating the pace at which deployed capacity becomes revenue generating, effectively extracting more Azure growth from every dollar of capital already invested.

- **The value is moving up the stack from GPU access to AI platforms.** As AI applications mature, customers increasingly care about deploying, managing, optimizing, and scaling AI workloads rather than simply securing capacity.
  - **DigitalOcean** provided a clear articulation of this framework, arguing that durable value accrues through "token flow" and "agent flow" rather than scarcity driven GPU pricing. The company emphasized that higher level services such as agent runtimes, databases, storage, and orchestration increasingly resemble software businesses rather than infrastructure businesses, carrying significantly higher margins.
  - **CoreWeave** highlighted managed inference, observability, networking, storage, and software tooling as key pieces of its long term strategy. Most customers already consume multiple services across the platform, which it views as evidence that AI infrastructure is evolving toward a broader cloud architecture rather than remaining a pure capacity business.
  - **Akamai** noted GPU deployments frequently drive incremental demand for CPU infrastructure, storage, security, networking, and edge compute.
- **Sovereign AI and enterprise control of infrastructure.** AI is increasingly treated as strategic infrastructure, with demand from governments, regulated industries, and enterprises seeking greater control over compute resources, data residency, and deployment architectures. We heard arguments both ways on compute architectures: some vendors noted customers increasingly pushing to cloud modernization, but others noted that AI could reinforce hybrid architectures. For example, Fortinet noted customer interest in hybrid and sovereign SASE architectures as enterprises look for more local control over AI infrastructure.
  - **Datadog** noted that customers are increasingly bringing their AI stack in house: rather than simply observing API calls into third party frontier models, it now has the opportunity to capture full AI stacks at enterprise customers

that span open models, proprietary data and harnesses, and training/post-training/inference workloads. Federated Logs and BYOC help Datadog stay in the workflow without requiring every log or dataset to move into Datadog's environment. Production inference is still the cleaner near-term monetization path, but some AI-native and hyperscaler customers are already using Datadog around training/post-training workloads as the boundary between pre-production and production starts to blur.

- **CoreWeave** highlighted demand from customers building their own clusters rather than relying entirely on hyperscalers and pointed to its software platform as a way to support customers that prefer local ownership of infrastructure (via CoreWeave Omni). The company also noted it has contracted more than 1GW of power outside the U.S., reflecting growing international demand and a more geographically distributed AI infrastructure footprint.
- **Nutanix's** AI deployment offered a useful window into how enterprise architectures could evolve as usage grows. Management invested approximately $20M in internal clusters expected to support roughly 80% of the company's AI needs, while continuing to call frontier models for the remaining 20% when their capabilities are required. Most AI activity still runs in the cloud, and we would not treat one company's experience as a template for the broader market. Still, the logic is relevant beyond Nutanix: data locality, latency, security, sovereignty and control may all support running certain workloads on dedicated infrastructure, particularly as inference becomes more persistent.
- **MongoDB** pointed to renewed demand for self-managed deployments as some enterprises seek to build their AI layer in-house, maintain greater control over sensitive data or revisit the economics of public cloud infrastructure. Management sees its ability to run across cloud and self-managed environments gaining relevance as customers assemble AI architectures around the requirements of individual workloads, with this flexibility complementing Atlas instead of redirecting demand away from it.
- **Rubrik** sees sovereign and self-hosted offerings becoming a more meaningful contributor as customers seek greater control across operations, data and infrastructure locality, with geopolitical risk adding urgency to those requirements. Management emphasized that sovereignty can take several forms, spanning customer data centers, colocation facilities and locally hosted cloud environments, giving Rubrik multiple ways to participate as enterprises place greater value on control over where critical data and AI workloads run.

- **There are many levers for cost optimization,** with discussions centered on token generation economics, inference optimization, agent orchestration, utilization, and software tooling improving productivity per megawatt.
  - **DigitalOcean** discussed the importance of model routing, quantization, post training, reinforcement learning, agent harnesses, and token optimization, noting that not all tokens are monetized equally and that competitive

differentiation increasingly comes from improving throughput, quality, utilization, and cost rather than simply provisioning more hardware. The company also highlighted growing adoption of open-weight models as customers seek lower inference costs and greater control over their AI stacks.

- **CoreWeave** emphasized its software rails, observability, and infrastructure tooling that help customers extract more value from each compute dollar.
- **Akamai** noted that as inference workloads become more latency sensitive and distributed, infrastructure optimization becomes increasingly differentiated; with opportunities in edge and distributed inference architectures where performance and efficiency are paramount.

Platform

- **AI is changing the economics of data modernization, shifting historically lengthy, services-intensive projects into automated, outcome-driven, faster processes.**
  - **Snowflake** discussed AI coding agents accelerating both migration execution and problem resolution e.g. a large Teradata migration now moves to <3 quarters from 2-3 years prior. Partners are increasingly willing to deliver migrations using fixed-price, outcome-based models rather than traditional time-and-materials engagements, all of which is enabling customers to ramp consumption faster, reach production more quickly, and generate value sooner than in prior cohorts. While faster migrations theoretically lower switching costs in both directions, platform vendors noted that easier onboarding outweighs any incremental churn risk, with the understanding that continued innovation is key to customer lifetime value.
  - **MongoDB** is increasing its investment in AI-assisted tooling with the ambition of compressing modernization projects from years to months and eventually weeks. Enterprises already have ample reason to retire aging applications and data estates; the practical hurdle has been the cost, disruption and scarce expertise required to do it. Automating more of the conversion, testing and remediation work could turn modernization into a more immediate source of workload growth across databases, data platforms and infrastructure.
- **As data gravity weakens and models become interchangeable, value is shifting higher up the stack toward governance, semantics, and enterprise context.**
  - **Databricks** emphasized that governance is becoming a central buying criterion as enterprises face three converging trends: increasingly rapid model proliferation, growing scrutiny around AI spending, and rising security concerns around autonomous agents. In this context, Unity Catalog is one of its most important products and frequently the primary differentiator in competitive evaluations.
  - **Snowflake** acknowledged that open formats and AI-assisted migrations are reducing traditional switching costs at the storage layer, making openness increasingly standard; as a result, differentiation is shifting toward governance, disaster recovery, observability, agent orchestration, and application services that sit above the underlying data platform. Snowflake's view is that if customers can move data more freely, long-term value creation must come from providing the control, visibility, and operational capabilities needed to run AI at enterprise scale.
  - **Dynatrace** argued that confidence remains one of the primary bottlenecks to broader enterprise AI adoption. As organizations move AI workloads into production, they need visibility not only into whether systems are running, but whether models are generating reliable outputs and agents are behaving as intended. Management believes this increases the value of platforms that can bring together logs, metrics, traces and business events into a shared

operational context, creating a stronger foundation for governance, automation and autonomous operations.

- **Coding agents are beginning to influence infrastructure selection at the point where new applications are created.** This was a notable theme in our discussion with MongoDB, where Management acknowledged that agents often default to Postgres because decades of documentation, code and community content make it easier for models to understand and recommend. Supabase has been an early beneficiary, emerging as a preferred backend across several vibe-coding platforms and AI startups. MongoDB is responding by making Atlas easier for agents to discover, provision and use, including through deeper integration with Vercel v0 and continued investment in MCP. We increasingly see agent mindshare becoming as important as developer mindshare in determining where new workloads begin. Vendors that fall outside the agent's consideration set risk losing those workloads at the top of the funnel, with the disadvantage compounding as early selections shape the usage, content and familiarity that influence future recommendations.
- **AI is accelerating the convergence between the traditional infra and app layers**. Historically, platforms primarily sold infrastructure for analytics, storage, and data engineering. Increasingly, vendors are pushing into workflow, business intelligence, security, marketing, and operational software layers built on top of their data foundation.
  - **Databricks** noted that AI adoption in enterprises remains far below its technical potential and views reducing that gap as the largest market opportunity ahead; products such as Genie are intended to move beyond traditional BI by allowing users to reason, compute, and interact directly with enterprise data. Management described ontology as the equivalent of an indexing layer for enterprise AI, helping translate raw data into business context that agents and models can reason over. Rather than positioning itself solely as a data platform, Databricks increasingly sees itself addressing knowledge workers, analysts, and operational users. This broader opportunity is also driving expansion into adjacent product categories. Databricks highlighted two large markets where it believes ownership of the data layer provides a structural advantage: customer data platforms (CDP) and SIEM/security analytics (with security a particularly large opportunity as agent activity drives significant log growth and increases monitoring requirements).
  - **Snowflake** also noted that the line between applications and data platforms is blurring, with future applications increasingly resembling collections of AI-powered skills operating directly on enterprise data. This opens opportunities for custom applications, workflow automation, orchestration, and AI-native software built directly on top of the platform layer.
  - **Twilio** described communications infrastructure increasingly evolving into a higher-level intelligence and orchestration layer. Management highlighted new products (ex. Conversation Memory and Conversation Intelligence) that help preserve context across channels, agents, and customer interactions. Twilio's view is that long-term value accrues not simply from connecting

communications channels, but from maintaining the contextual layer that enables AI agents and humans to operate consistently across voice, messaging, email, and other engagement surfaces.



Applications

- **The conversation has moved from "existential risk" to a more nuanced discussion of NTM product cycles, tech stack maturity, and value capture vs. value abstraction.** Companies continue to describe a future where software serves as the orchestration layer connecting agents and delivering outcomes, rather than simply acting as the system of record. Management teams generally argued that durable differentiation will come from owning workflow context and vast amounts of historical data, not simply access to frontier models. Shopify, HubSpot and Intuit each discussed how they are adapting to marketing channel fragmentation (e.g. AI-driven impact to SEO). Intuit noted it is not yet seeing impacts from AI native tax competition, but investing heavily in AI features in TurboTax and sharpening their price points to ensure they remain competitive.
- **Monetization continues to be a work in progress.** Most vendors continue to prioritize adoption and usage, while experimenting with new forms of pricing/packaging. At the same time, there are headwinds from elongating sales cycles, app budget pressure to fund token projects, and more uneven competition. As companies tie their products to quantifiable outcomes, there does seem to be some willingness to source some portion of apps budgets from labor budgets.
  - **Microsoft** described AI as accelerating a shift from seat-based licensing toward pricing aligned with units of work, outcomes, and productivity delivered. The company pushed back on the notion that AI simply reallocates existing software budgets, arguing instead that AI enables new workflows, new business processes, and entirely new revenue opportunities that were previously uneconomic or impossible to deliver. In this framework, AI is not just an efficiency tool but a mechanism for expanding software spend by creating incremental demand for new outcomes. At the same time, Microsoft emphasized that enterprise adoption increasingly depends on the ability to measure, govern, and optimize AI usage. Agent 365 is the control plane that helps customers monitor consumption, manage costs, and validate ROI across models and workflows, reflecting a view that observability and spend management may become as important to enterprise deployment as model performance itself.
  - **Salesforce** sees a sizable upgrade opportunity within the installed base, with only ~5% of customers currently on the AE1 premium edition which carries a ~60%-80% uplift. Rather than positioning AI as a standalone product sale, Salesforce framed the strategy around consolidating sales enablement, performance management, analytics, collaboration, Data 360, Agentforce, and ClaudeForce into a single higher value offering (AE1). Salesforce further quantified the value proposition, noting that sales organizations often spend ~$1,300 per seller across point solutions vs. $550 price point for Salesforce's bundled premium edition. This monetization framework could extend well beyond traditional seat licenses; with agents, consumption, workflow execution, and business outcomes becoming incremental pricing vectors as agents represent a larger share of enterprise activity. The combination of

premium SKU migration, vendor consolidation, and emerging agent based pricing models points to a broader expansion in monetization opportunities over time.

- **Hubspot** noted adoption of agentic use cases across the customer base, with both reach (the percentage of customers using agents) and usage depth (the number of agentic actions performed) continuing to improve. Adoption is uneven as customers work through data consolidation efforts, evaluate AI-generated outcomes within their own environments, and seek greater visibility into usage-based costs. HubSpot views these as normal adoption barriers during a platform shift and is focused on reducing friction through longer trials, customer education, and outcome based pricing.
- **Figma** reduced credit pricing by up to 50% after determining its prior pricing was too high, with the near term focus shifting toward driving adoption and volume rather than maximizing unit economics. Price was not always the primary constraint on adoption, with customer feedback pointing more consistently to quality as the limiting factor. As Make quality improved, adoption friction has eased, while declining frontier model costs provide a meaningful offset to lower pricing and support broader deployment across the user base.
- **Tyler** gave an example of a customer who was paying $1mn annually for its core software, purchased an AI-add on for another ~$1mn, and realized $2mn in labor savings; they noted this trend of funding AI products from labor rather than IT budgets could expand the TAM.
- **ServiceNow** provided several quantitative examples that suggest AI deployments are moving beyond experimentation: 1) Robinhood, which has automated ITSM, with agents now managing 70% of cases previously handled by humans; 2) Rossmann, which operates 5,200 drugstores, reduced customer case resolution time from 5-9 minutes to five seconds; and 3) an airline automating millions of calls. ServiceNow reiterated that AI can materially expand contract value over time, with the average seat expected to generate 4.5x its original contract value as customers add AI use cases and redesign workflows across the enterprise. ServiceNow acknowledged that enterprise adoption remains uneven, with customers navigating significant confusion, experimentation, and competing point-solution claims. As a result, the company is increasingly relying on C-suite engagement and FDE motion to accelerate deployment and enterprise-wide standardization. ServiceNow is also beta testing a simpler AI help desk offering that can be deployed independently while remaining connected to the core platform.
- **Klaviyo** noted that agent monetization remains secondary to adoption today, with management focused on driving usage, repeat engagement, and product-market fit before optimizing pricing. Early data points were encouraging, including 95k Composer users within one month of launch and Customer Agent usage growing ~40% q/q, supporting management's view that agents can become a future growth accelerator as adoption scales.

- **Zeta** highlighted Athena as a potentially meaningful monetization driver, with 40% of super-scale customers already monthly active users within four months of launch. Management noted that Athena users are creating significantly more audiences and campaigns, driving higher platform consumption and creating a flywheel where increased AI adoption translates directly into greater usage of the core platform.

- **We view Shopify as a good case study of leveraging proprietary data into proprietary AI product, creating a moat vs. new entrants**. We believe these comments are particularly relevant in a period of time when investors are debating the impact of new personal assistants like Meta Muse. For example, Sidekick benefits from Shopify's visibility into merchant behavior and transaction workflows, enabling increasingly proactive recommendations based on "next best action" prediction models. The company views this intelligence layer, combined with network-scale payments and conversion data, as a durable competitive differentiator that should drive incremental GMV and ultimately gross profit dollars over time. While Sidekick incurs incremental inference costs, Shopify is actively driving efficiencies through model distillation, caching, and broader use of lower-cost models where quality remains acceptable. Rather than focusing on direct AI monetization, the company expects value creation to occur indirectly through higher merchant success rates, increased GMV, and greater adoption of merchant solutions such as payments, tax, and related services. Successful AI adoption ultimately translates into gross profit expansion across the broader platform, reinforcing Shopify's historical preference for ecosystem monetization rather than standalone AI pricing.

Security

Several industry experts described AI as compressing time: for the first time, CrowdStrike noted that security is being viewed as an accelerant to business outcomes rather than a break; while Palo Alto noted that new AI headline attacks have been a more useful marketing tool than the marketing efforts it drove internally for the last eight years. At the same time, Palo Alto acknowledged that AI value-added services are being built as we speak, and no vendor today has a complete stack. At speed, it takes 3-6 months to build the right security hooks into new agentic stacks, such that a complete agentic stack with a robust accompanying security stack is still very much a moving target today.

Against this backdrop and the Security stocks trading at elevated multiples, investors are spending an outsized amount of time looking to understand where there could be disruption, such as with a frontier model entering the security TAM.

- **CrowdStrike** further detailed its introduction of SafeMind at Fal.con last week as a cyber frontier model that can spin up continuous penetration testing, built on Nvidia's Nemotron paired with CrowdStrike's 15+ years of proprietary kill chain data. CrowdStrike believes that frontier models have a role in understanding vulnerabilities and patching - but do not have the benefit of being in line and are not taking action real time, as is the case with CrowdStrike.
- **Palo Alto** noted the following enduring moats: 180mn sensors deployed globally (endpoint and network), which can expand both organically and inorganically and 19PB of data run through Google Cloud every day. Observability (Chronosphere) and operational IT (Console) are good new examples of Palo Alto further inserting itself into token flow, and subsequently compounding on Palo Alto's traditional 2-5% take rate on cloud spend.

Across security subcategories, network security screened more positively as the AI debate translated into a more concrete traffic and throughput argument. There is also an underlying push to modernization, which we believe has the greatest impact in SIEM. The consensus view in Identity is that a rising tide lifts all boats; we believe the more nuanced answer to whether there is a step function change in architectures because of the ephemeral nature of agentic identities will become clearer in 2027.

- **Fortinet** shared that customers are buying larger capacity firewalls, as AI-driven east-west traffic, machine-to-machine activity and internal segmentation create higher capacity needs. In other words, AI could help the existing firewall refresh and SASE cycle, not only create a separate set of AI-security products. Sovereignty came up in the same context, with Fortinet tying demand for on-premise and sovereign SASE to customers that want more local control over AI infrastructure and traffic.
- **Okta** and **SailPoint** highlighted identity as an increasingly critical control plane for AI adoption, with enterprises becoming more focused on agent discovery, authorization, and auditability as deployments scale. While views differed on the extent to which agentic identities require new architectures, both companies emphasized the growing importance of understanding what agents exist, what they

can access, and what actions they can take. As enterprises move from pilots into production deployments, we expect the debate around ephemeral identities, policy inheritance, and governance architectures to become increasingly important.

- **Rubrik** added another dimension to the identity debate, arguing that resilience becomes increasingly important as agents inherit permissions and act on sensitive data. Management noted that identity systems are already a primary attack surface, with attackers frequently entering through valid credentials and leaving behind persistent changes that can survive a conventional recovery. The growth of agentic identities compounds that challenge by increasing both the number of actors and the speed at which they can alter systems, creating a need to recover identity environments to the correct state and unwind harmful actions across increasingly complex workflows.

Resilience is moving closer to the center of the security discussion as machine-speed attacks compress the time available to respond and recover. Boards and CEOs are increasingly evaluating cyber risk through downtime, financial exposure and the ability to keep the business running.

- **Rubrik** is building a broader resilience platform across data, identity and agents as recovery becomes increasingly intertwined with security operations. Management expects security and IT operations to converge around continuous patching and recovery, with machines implementing recovery plans and humans retaining authority over critical decisions. Rubrik's position in recovery is well established, and its progress in identity resilience extends that role into restoring identity systems to the correct state following an attack. Agent Cloud carries the same logic into AI governance by combining visibility into sanctioned and unsanctioned agents, identity-based controls, runtime understanding and the ability to unwind harmful actions. This gives Rubrik a credible entry point as enterprises begin to secure autonomous activity, although the broader AI governance category, its product boundaries and its eventual leaders are still taking shape.

Physical AI

Physical AI could become one of the more consequential shifts in vertical software as digital intelligence begins to connect with real-world execution. These vendors already sit inside critical workflows and increasingly control the operating data generated across those workflows. AI creates an opportunity to move beyond systems of record and into systems of action, expanding the role of software from documenting work to helping execute it. While adoption remains early and will vary by workflow, customer readiness and accuracy requirements, we came away increasingly focused on how these platforms are expanding the scope of data, context and operational responsibility they control.

- **Procore** highlighted a pragmatic adoption path through tiered Digital Coworker packages, allowing customers to start with narrowly defined workflows before expanding into more sophisticated automation. Management's approach suggests the market may develop incrementally rather than through a single step-function adoption curve.
- **Samsara** pointed to roughly 4x usage growth across certain AI capabilities over the past two months and highlighted measurable use cases in maintenance, safety and asset utilization. The common thread across these examples is that success can often be tied directly to labor productivity, downtime reduction or improved asset efficiency.
- **Autodesk** sees AI connecting design, construction and operations into a more automated project workflow, helping customers address persistent capacity constraints across AEC and manufacturing. Management argued that using fewer resources per project can free customers to work through backlogs and pursue additional bids, supporting seat demand across a larger volume of work. Given the accuracy requirements and long project cycles in these industries, monetization should develop gradually, with task automation increasing seat value and broader automation creating incremental consumption over time.
- **Across the broader ecosystem, companies are also expanding the context they control.** DroneDeploy gives Procore greater visibility into jobsite conditions, MaintainX connects Autodesk's design and construction workflows to operational performance, and Samsara continues to deepen its understanding of physical operations through its connected network. We see this expansion of context as an important part of the long-term opportunity, potentially allowing vertical software vendors to extend beyond task automation and into larger portions of the operating process.

Going into the week, investors were still working through what the latest AI developments could mean for Autodesk. The discussions gave us useful context, though few definitive answers. Autodesk's response to Astra was that generating geometry is only one part of professional design. The output still needs to reflect engineering constraints, project history and local standards, then be reviewed, modified and carried through to something that can actually be built. Management sees Autodesk Assistant as the layer that brings together its own Neural CAD models with open-weight

and frontier models depending on the task. There was also separate interest in whether Prometheus would offer a clearer view of its ambitions in engineering and design. Prometheus added less clarity, with its initial focus appearing closer to high-end manufacturing and simulation and too little product detail to draw firm competitive conclusions. For now, we see limited evidence that Autodesk's near-term position has changed. Longer term, the debate increasingly centers on how value will divide between models that make design creation easier and the data, context and workflow required to turn that initial output into a usable design. We expect to hear more on this when we attend AU26 next week.