SEMICON West 2025 Technology Breakfast

KEY POINTS | October 7, 2025

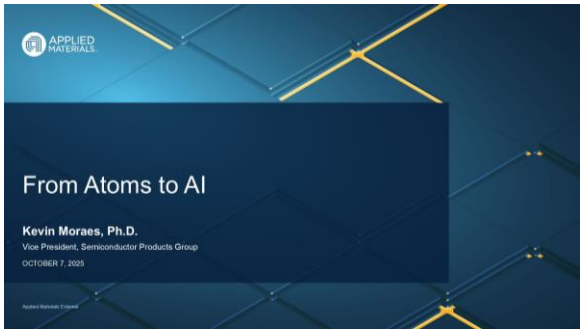

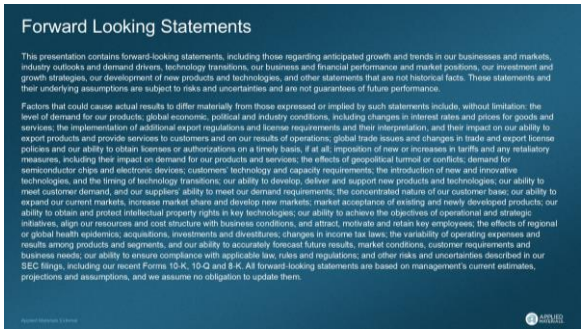

- Good morning, everyone, and welcome to the Technology Breakfast at SEMICON West!
- Before we dive in, I need to share our standard forward-looking statements disclaimer.
- This reminds you that today's discussion may include predictions or expectations about future performance, which are subject to risks and uncertainties. Please review the details on this slide.

- This year marks a special milestone. For the first time, we are gathering here in Phoenix. And I am excited to see so many familiar faces in the room today!

SEMICON West 2025 Technology Breakfast

KEY POINTS | October 7, 2025

- Last year in San Francisco, the room was buzzing as the industry debated how quickly AI would reshape everything – from chips to fabs.
- After last year's event, I had a hallway conversation with a longtime customer who said: "This isn't just about chips anymore, it's about how we work together to solve problems no one can tackle alone." That stuck with me – and it's never felt more true.

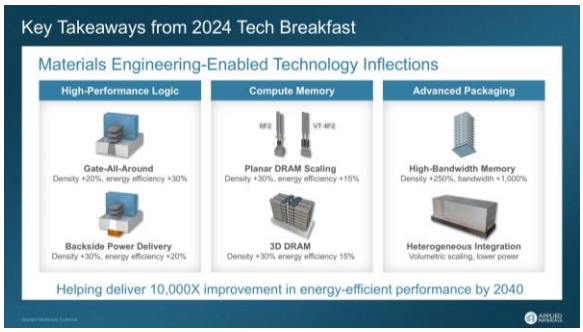

- Looking back at the Key Takeaways from last year there was a shared urgency around energy-efficient AI.
- New architectures like Gate-All-Around, and Advanced Packaging were key to energy-efficient performance.
- It was clear that the industry needed deeper ecosystem collaboration.

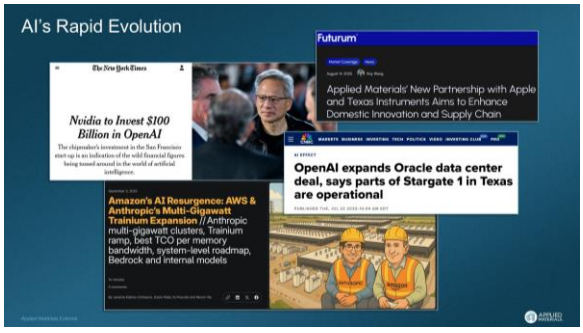

- So, what's different this year?
- The pace of change has accelerated – AI workloads, chip complexity, and energy demands are scaling faster than ever.

- Technologies we previewed last year are now in high-volume production.
- The ecosystem is connecting in new ways – integrated solutions, cross-company partnerships, and platforms like EPIC are gaining real momentum.
- The challenges are bigger, but so are the opportunities.
- We are not just talking about the future, we are building it, together.

- Today, we are raising the bar.
- You will hear from industry leaders and experts tackling the toughest AI bottlenecks – from system workloads, to manufacturing and device enablement.
- We will start with opening insights from Prabu Raja, framing the opportunities and challenges of the AI era.
- Then, we will hear perspectives from Jim Chambers at NVIDIA and Rose Castaneres at TSMC, sharing their views on AI workloads and manufacturing roadmaps.
- Finally, Mukund Srinivasan will lead a fireside Q&A with Applied technologists, diving deeper into how we're enabling the next wave of innovation.
- And with that quick introduction, please join me in welcoming Prabu Raja to the stage.

- We are in the early innings of AI deployment, most of it happening in the cloud.
- We are moving beyond narrow, task-specific models to systems that can reason, generate, and adapt in real time.
- It's already transforming how we create, how we design, and even how we innovate.

- But the future? It's going to be even more dramatic and far more exciting.
- AI will be everywhere – embedded in devices, powering new experiences, and reshaping how we live and work.
- And as it scales, it won't just change industries. It will redefine the global economy.

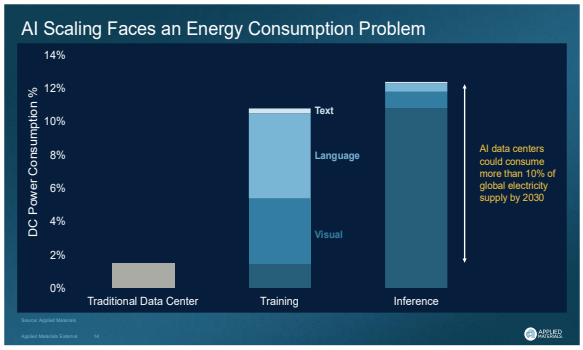

- AI data centers could consume more than 10% of global electricity supply by 2030. Without a sustainable approach, demand may outpace supply.

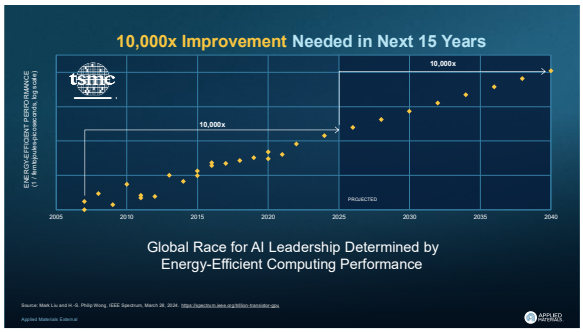

- To lead in AI, we must deliver 10,000x better performance and efficiency by 2040. That's the challenge – and Applied is built for it.

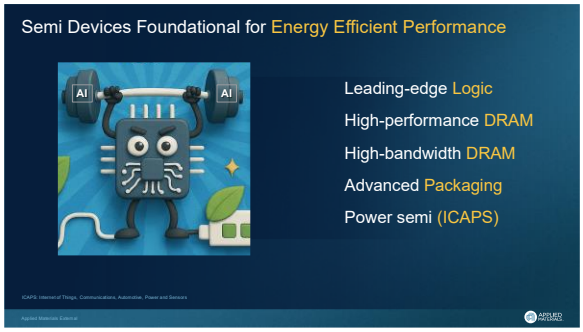

- The journey starts with foundational technologies: Leading-edge Logic, High-performance DRAM, High-bandwidth Memory, Advanced Packaging and Power semiconductors
- These define the architecture of AI systems.

- New 3D architecture inflections like gate-all-around, backside power, vertical transistors and advanced packaging are helping deliver big improvements in Energy-Efficient Performance
- These aren't tweaks – they are breakthroughs!
- They unlock performance and energy gains that scale.

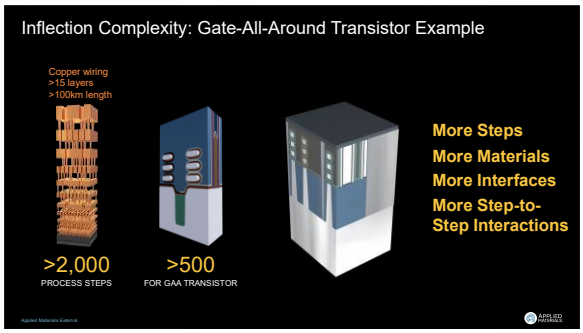

- But complexity is rising fast.
- Just one GAA device takes 2,000+ steps – 500 for the transistor alone.
- Every step matters and interacts with others.

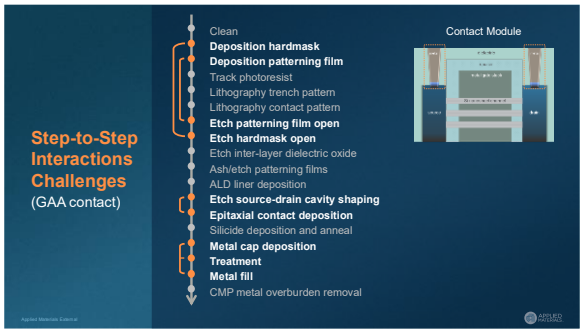

- One step can affect the next, or one 20 steps later.
- Managing these interactions is critical to yield and performance.

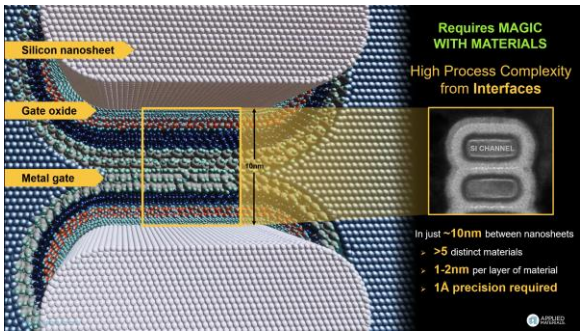

- We are engineering at the atomic level. Five 1nm layers must behave perfectly.
- This is materials magic and Applied's core strength.

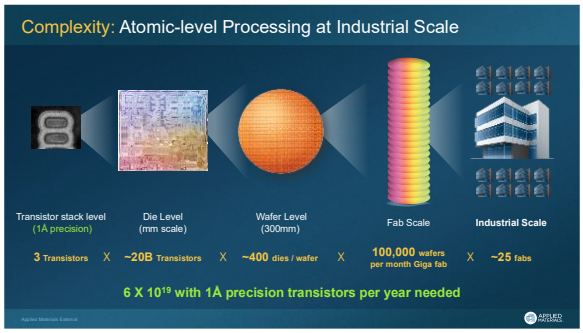

- And we do it at scale: 10^{20} transistors a year, with angstrom-level control.
- Precision meets volume.

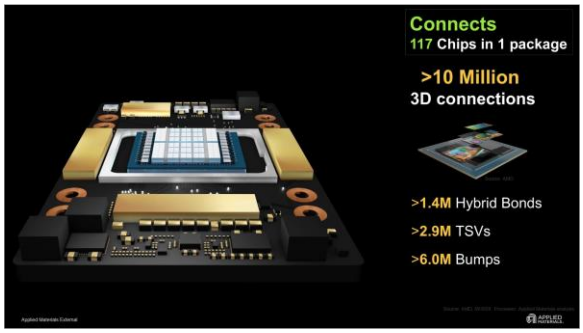

- Packaging is now a performance driver. Systems like MI300x have 10M+ connections, 117 chips, 153B transistors.
- Yield matters – and Applied delivers.

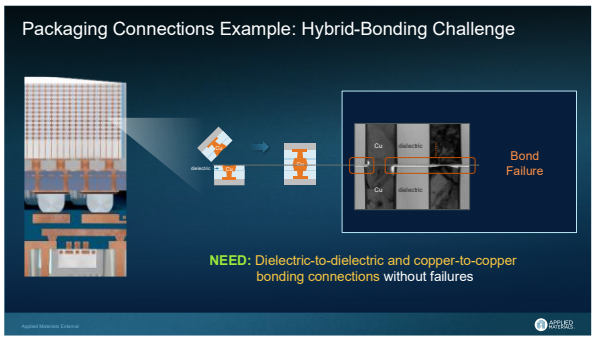

- Hybrid bonding demands millions of copper-to-copper connections.
- Dielectric fusion. Perfect alignment. No room for error.

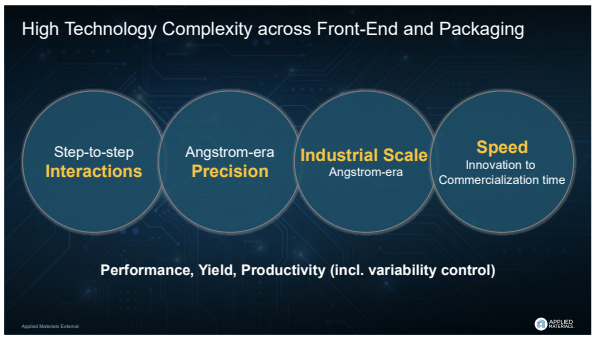

- We are innovating across front-end and packaging: angstrom precision, step control, speed, yield, scale.
- No one can do this alone.

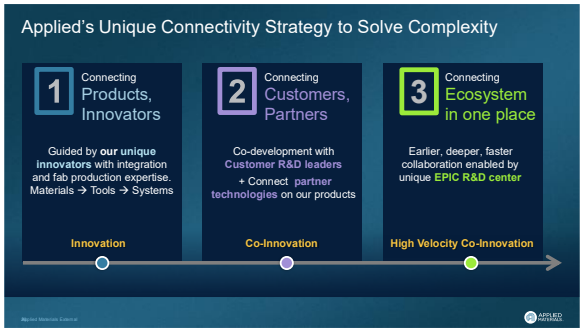

- That's why Applied is connecting everything: In 2013, we began linking our broad portfolio, in 2018, we connected our innovators, and in 2023, we launched EPIC – partnering with customers, equipment makers, and universities.
- We are building a connected innovation network.

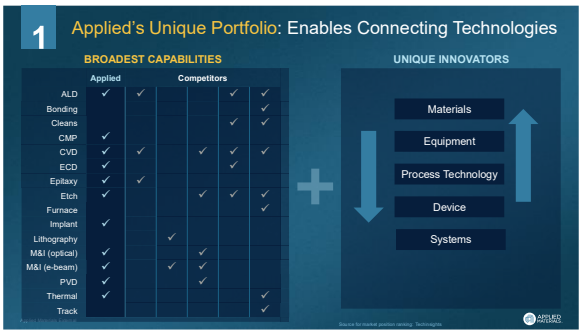

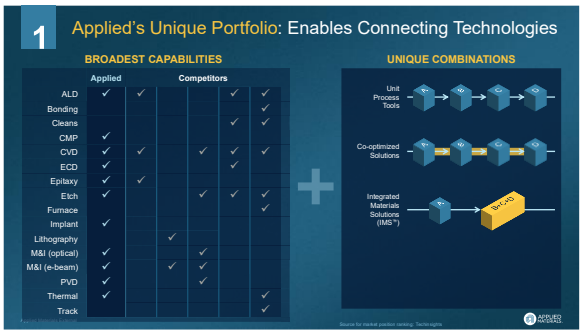

- Applied leads in breadth and depth.
- We combine technologies in ways no one else can. From systems down to materials, including device and process technologies, and manufacturing equipment.
- Sometimes this means co-optimized solutions, sometimes this results in Integrated Materials Solutions (IMS).
- That's how we solve the hardest problems.

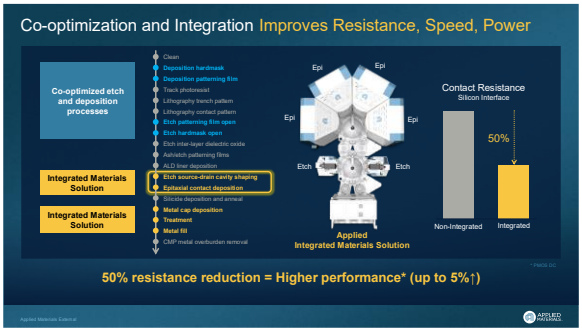

- For example, the contact is a bottleneck where power from a transistor must connect to the interconnect.
- It takes over 20 steps to form the contact, and some of these steps must be co-optimized and even integrated for best performance.
- By combining our unique etch capability to shape the contact and our leading epitaxy capability to create a low resistance interface, we can lower contact resistance by as much as 50%.
- That's Applied synergy.

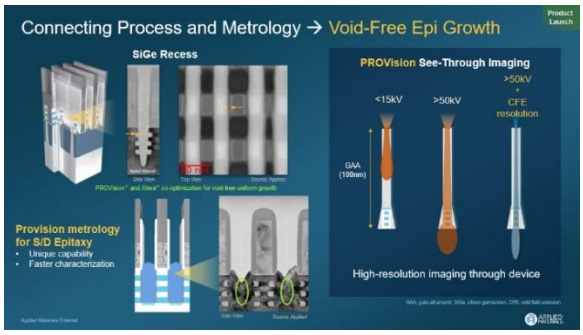

- Our eBeam metrology sees what others miss.
- With cold field emission (CFE) technology and a 2x throughput improvement, PROVision™ 10 finds buried voids in GAA – at scale.
- By increasing the landing energy, it can image deeply buried features, for example detecting buried voids, to accelerate process optimization and yield ramps.

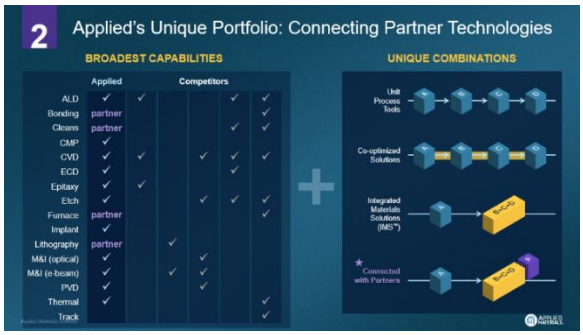

- But we don't just stop with technologies that Applied has invented.
- We co-optimize with partners, because no one wins alone.

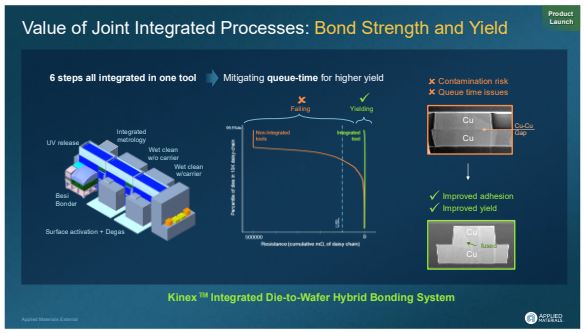

- Going back to the example of hybrid bonding, to be assured of the best performance, it requires many steps to be co-optimized, and in some cases even integrated.
- Let's zoom into the pre-treatment and bonding step. Applied has developed a solution with six bonding steps, co-optimized with Besi's industry-leading bonder, all integrated on a single system we call Kinex.
- You would have seen the press release earlier today, that the Kinex system launches the industry's first integrated die-to-wafer hybrid bonding system.

- It eliminates queue time challenges, improves adhesion and yield.
- That's collaboration in action.

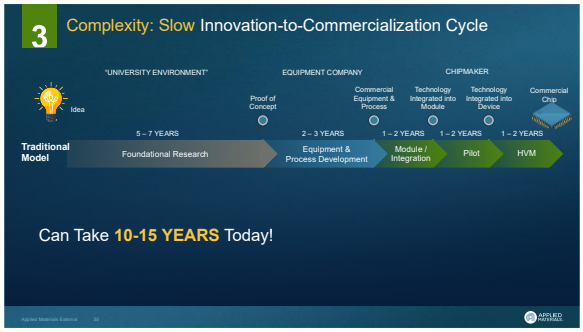

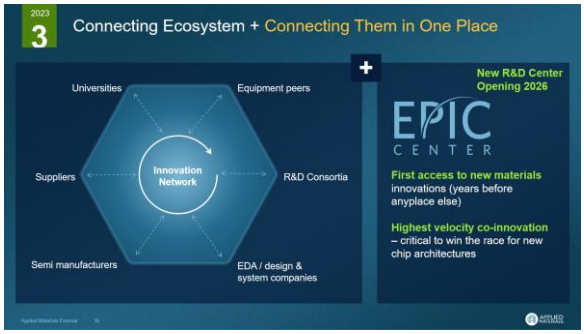

- Traditionally, it can take 10-15 years to create these new semiconductor manufacturing technologies.
- We need to solve integration challenges with new materials and new interfaces, new devices, and robustness of manufacturing technologies.
- This is too slow. We can and should do better.
- This is why Applied announced the EPIC Center, connecting the eco-system, for faster access to new materials and the highest velocity co-innovation.

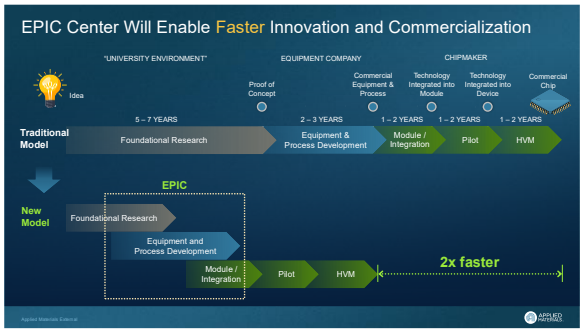

- As I said before, "Conventional approach no longer works. We must innovate the way we innovate to speed the cycle up"
- This new EPIC model will allow us to innovate in parallel, and accelerate the cycle of development and commercialization, potentially 2x faster.

- Let me recap.
- Device complexity is accelerating – driven by more process steps, tighter interactions, and the need for angstrom-level precision at industrial scale – all while speeding up innovation and commercialization to stay competitive.
- The answer to this rising complexity is deeper connectedness: bringing together innovators, products, device makers, and partners in one ecosystem to co-invent the solutions we need faster.

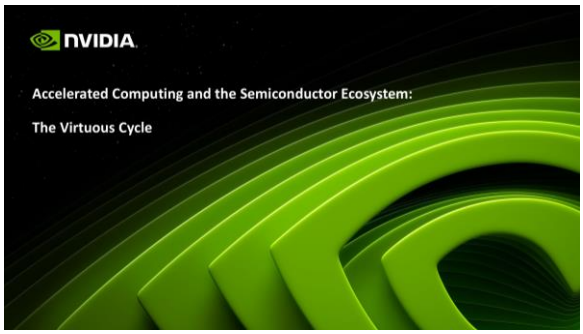

- Now, let's hear from the leaders shaping this future.
- Please welcome Jim Chambers from NVIDIA to the stage.

- Now, please welcome Rose, who will take us deeper into the device innovations driving performance per watt.

- Now, please welcome Mukund, who will take us into the next chapter – how we enable the devices that will power the future of AI.

- Thank you, Prabu, for the vision. Thank you, Jim, for the challenge. And thank you, Rose, for the roadmap.
- Together, they've shown us what's possible. and what's required, to scale AI to its full potential.

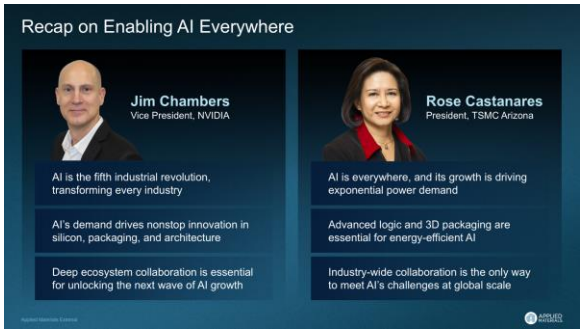

- Jim spotlighted the transformative power of AI and the system-level constraints we must overcome.
- Rose then showed us how TSMC is breaking through those limits with innovations in transistors, power delivery, and advanced packaging.

- AI is powered by a new generation of devices: advanced logic, low-power and high-speed DRAM, high-bandwidth memory, cutting-edge packaging, and efficient power devices.
- But as AI scales, so does its energy demand – and the challenge is immense.
- To keep pace, we need a 10,000X leap in energy-efficient computing, and we need it fast.
- This is our call to action. As an industry, we must deliver More Compute. More Bandwidth. Less Energy. Faster.

- As Prabu showed us, meeting the energy challenge requires bold architectural inflections – especially in 3D.
- We are entering a new era of 3D innovation with gate-all-around transistors, backside power delivery, advanced DRAM devices, hybrid bonding, and photonics.
- These aren’t incremental steps. They are foundational shifts that will unlock the next wave of AI performance.
- This is where Applied comes in.
- Our portfolio is built for this moment: 3D architectures are enabled by materials engineering solutions.
- We work side-by-side with customers accelerating innovation through deep R&D partnerships, AI-powered process control, and integration expertise across the stack.

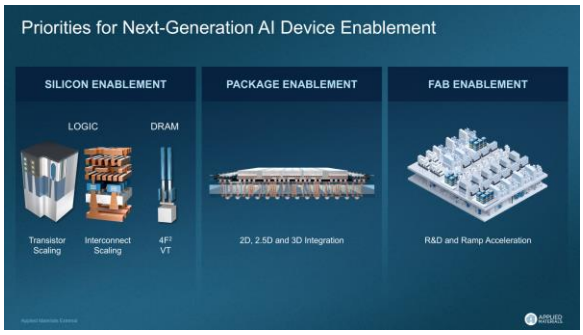

- To show you how we do this, let me introduce our enablement priorities:
- Silicon Enablement: Transistor and interconnect scaling with new materials and architectures.
- Package Enablement: 2.5D, and 3D integration, including photonics and HBM.
- Fab Enablement: Accelerating R&D cycles and ramping to high-volume manufacturing.

- You will hear from our technology leaders who will take you through each of these areas.
- You will also hear real-world success stories – how Applied helped customers accelerate both R&D and production ramps.
- Let me introduce our technology leaders:
- Bala Haran, who leads integrated solutions for logic and memory scaling for energy-efficient computing.
- Subi Kengeri, driving innovation in advanced packaging and integration.
- Mike Chudzik, who leads integrated solutions for ICAPS and Packaging, and focuses on accelerating R&D.
- Steve Frezon, who leads our North America Service Group and specializes in manufacturing ramp acceleration.
- Each brings deep expertise and hands-on experience from working with customers and partners around the world.
- To kick things off, I'd like to invite Bala to the stage as our first speaker.

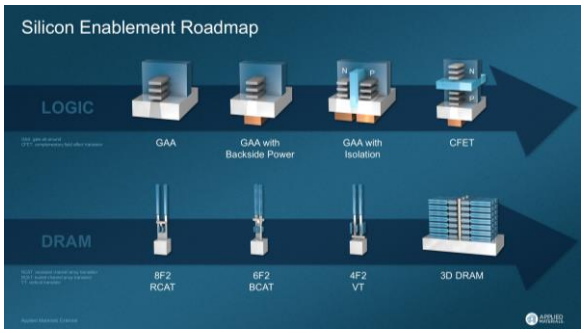

- *Mukund*: Quick roadmap for the audience – today we'll cover both logic devices and DRAM, because both matter for energy efficient AI.
- Bala, to set the stage, what are the key levers and why do they matter?

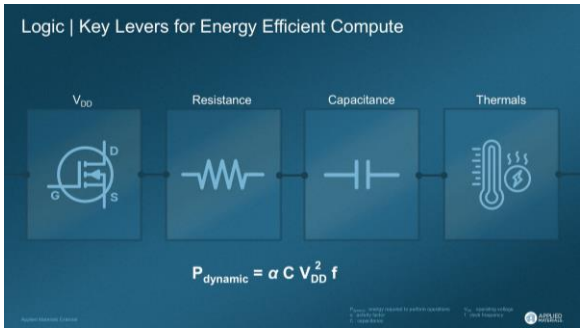

- Four levers, in this order:
- Lower VDD (transistor turn-on voltage), lower resistance (R), lower capacitance (C), and better thermals.
- Power scales roughly with V^2, so voltage is the biggest step; then cutting R and C boosts speed and bandwidth at lower power; and thermals keep performance sustained under load.
- I will show how each applies to logic and DRAM.

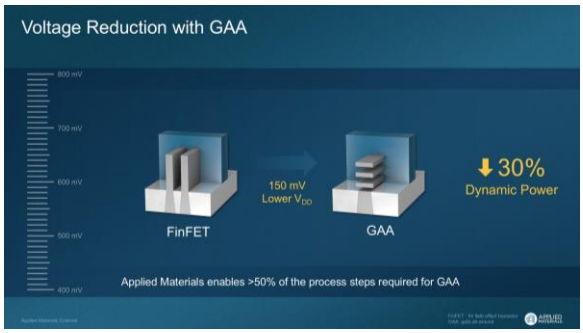

- Let's start with logic. Reducing voltage has become a critical scaling element. How is this accomplished?
- The fundamental building block for Logic is the transistor – in simple terms this is a switch that when turned on allows digital operations.

- Energy-efficient compute requires full control of the switch with the gate surrounding the active device on all sides.
- These Gate-All-Around (GAA) nanosheets are superior switches with tighter gate control and less leakage, so designers can drop operating voltage while maintaining performance.
- The FinFET→GAA inflection is enabling ~150 mV lower operating voltage, which translates to ~30% lower dynamic power at comparable performance.
- And Applied has developed many of the process and metrology tools enabling this inflection.

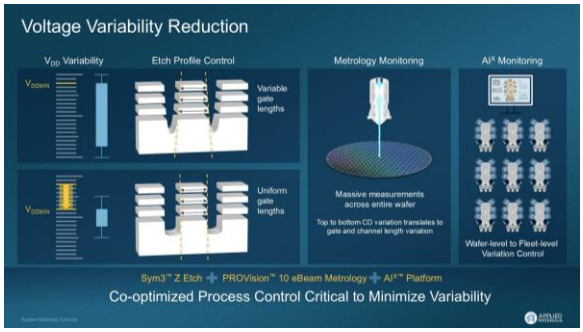

- Lowering V_{DD} is key for power reduction, but as process windows tighten, variability can limit how far we can go. How are we addressing this challenge?
- Variability sets the floor for V_{DD}. For example, if gate lengths are varying from top nanosheet to bottom nanosheet caused by a tapered etch profile, margin disappears.
- We uniquely solve this with Sym3™ Z with our unique Pulsed Voltage Technology or PVT, that delivers ultra-vertical profiles and flat bottoms, locking in electrostatics so voltage can scale.
- And with ProVision™ 10, we deliver massive eBeam coverage – 70 million CD measurements per hour – to measure within die, within wafer and and lot trends like no one else.
- Connecting process and metrology gives us unique insight into delivering transistor performance through materials engineering
- On top of that, AI-driven process control, in-tool sensors, and digital twins slash variation across wafers and tools, even at the edge.
- That's why we're the enabler for pushing V_{DD} lower without sacrificing yield.

- Reducing resistance is critical for performance. Where are the biggest challenges, and how is Applied addressing them?
- At today's dimensions, resistance bottleneck shows up in three places: near the transistor or FEOL, at the narrowest contact or MOL, and across the interconnect stack or BEOL.

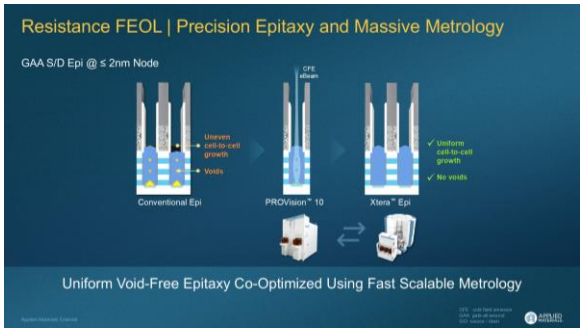

- First bottleneck, near the transistor.
- As electrons leave the nanosheet channel, the geometry is tight and high-aspect ratio.
- To pull full current from GAA, we use Xtera™ EPI with integrated high-aspect ratio cleans and selective deposition to grow highly-doped epi that cuts resistance and boosts drive.
- Buried voids and uneven growth quietly raise resistance. That's why we pair epi with ProVision™ 10, delivering 70 million eBeam measurements per hour to rapidly co-optimize and lock in void-free, uniform epi.

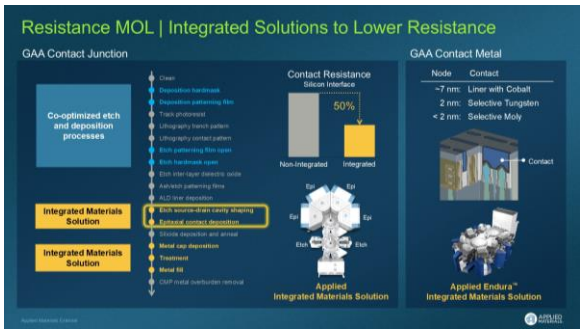

- Second bottleneck, at the contact between silicon and metal.
- As Prabu showed earlier, we connect anisotropic directional etch shaping and epitaxy technologies to engineer the silicon-metal interface for 50% lower resistance.

- And we need a low resistance contact metal fill.
- We deposit Tungsten bottom up - no barriers, no interfaces. Selective W scales to 2nm, and barrier-free Selective Mo drives resistance even lower.

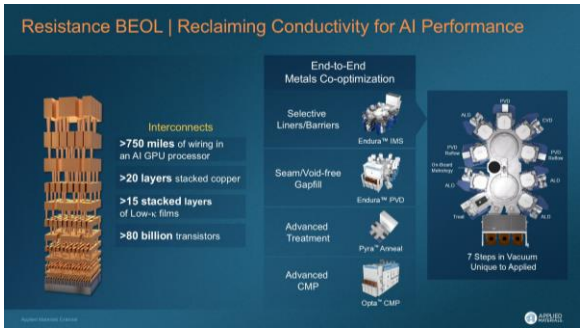

- Third bottleneck, across the interconnect stack.
- AI devices route signals through 20-plus metal layers and hundreds of miles of wiring.
- Long lines and tiny dimensions add wire resistance.
- We tackle that with end-to-end metals co-optimization.
- The most unique product here is Endura®, integrating seven advanced technologies to deposit ultra-thin, selective liners and barriers that reclaim conductor area and lower resistance.
- Together, these steps uniquely reduce resistance from device to interconnect, protecting timing margins and enabling higher performance at lower power.

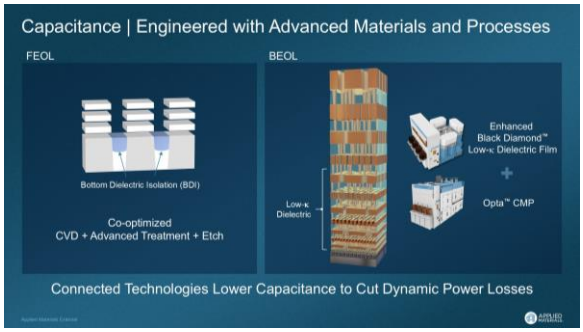

- In my customer conversations, I keep hearing about the need to lower capacitance for dynamic power. What's Applied doing to tackle this?
- Capacitance can hurt both the transistor and the interconnect.
- That is why we have developed a unique co-optimized solution that combines deposition, treatment, and etch to create a dielectric that isolates adjacent transistors – cutting parasitic coupling right at the device level.
- On the back end where AI chips pack 20-plus metal layers, we deploy ultra-low-k dielectrics like Enhanced Black Diamond™ and pair them with Opta™ CMP for clean, planar interfaces across dense and variable patterns.

- The result? Lower capacitance, less dynamic power, and more on-chip bandwidth for AI workloads.

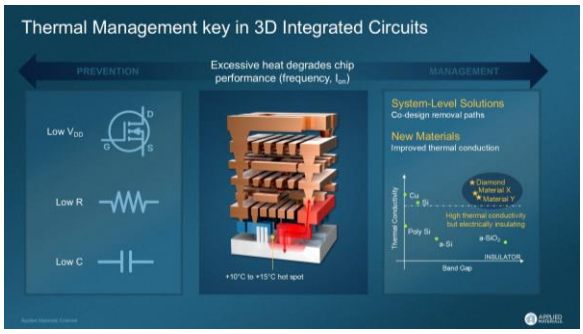

- Finally, thermals. Power density keeps rising. What's your philosophy on managing heat, and what are we doing about it?
- We start with prevention: generate the least heat possible by pushing lower V_{DD}, lower R, and lower C. That's switching and conduction loss you never have to remove.
- What we cannot avoid, we manage by incorporating thermal interface materials (TIM) that reduce hot spots and spread the heat more effectively.
- However, these materials must be engineered to work in different parts of the device.
- The net result: higher sustained frequency under AI workloads, less throttling, and better long-term reliability. Prevention first, removal second.

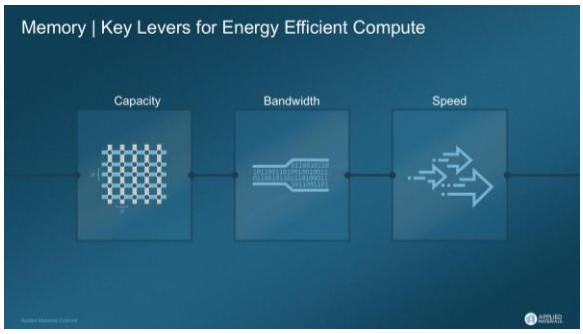

- Let's turn to memory. Why does DRAM remain central for AI workloads?
- AI is a bandwidth and capacity monster.
- DRAM is critical for providing fast access to vast amounts of data required for AI tasks with high speed LPDDR and High Bandwidth Memory being the key technologies.
- What must DRAM deliver next to keep up with AI?
- Three things: capacity, bandwidth, speed.
- Capacity: Compute intensive AI training workloads require access to large datasets instantaneously requiring high DRAM capacity
- Bandwidth: Memory-to-compute data speed is the bottleneck. We engineer the stack and co-optimize with advanced packaging to break through.

- Finally, Speed: Critical for error correction in large memory arrays. Fast periphery transistors with low-RC paths are key to accelerating deep learning and inference.

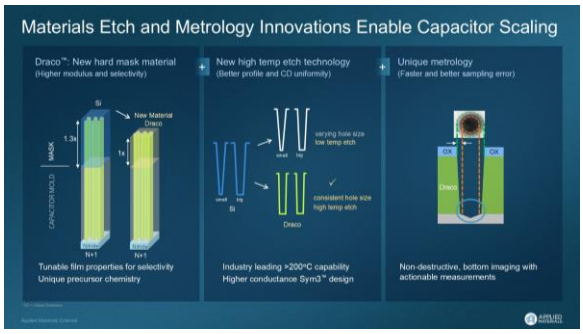

- As DRAM moves toward 3D architectures to improve density, what are the big process inflections, and how is Applied enabling them?
- First, etching tall, high-aspect-ratio vertical channels at tight pitch for the DRAM capacitor without bowing or collapse.
- Here Applied leads with unique hard mask materials and etch technologies, co-optimized with the power of our advanced metrology.
- Our CVD Draco™ and Sym3™ Z etch continue to deliver generation after generation of tighter pitch and highly vertical profiles with high etch selectivity – exactly what is required to scale the DRAM capacitor from node to node.

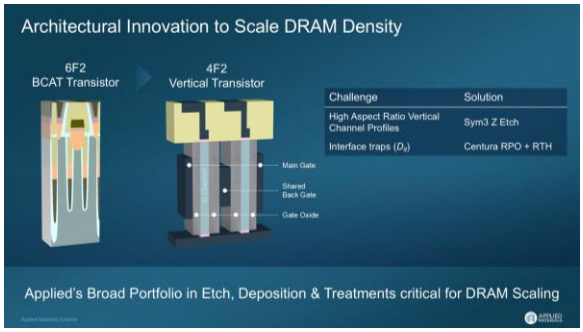

- Next, the industry is moving from a conventional $6F^2$ transistor architecture to a compact $4F^2$ vertical channel transistor – the Vertical Gate Transistor.
- A key challenge here is to improve the transistor performance with vertical channel profiles & better gate oxide with high interface quality.
- Interface traps degrade refresh and raise standby power. Applied Centura™ RPO with integrated treatments engineers the gate oxide to reduce interface trap density, improving refresh and overall power efficiency.
- That's how Applied is uniquely enabling the next generation of DRAM density – solving vertical device and interface quality challenges at production scale.

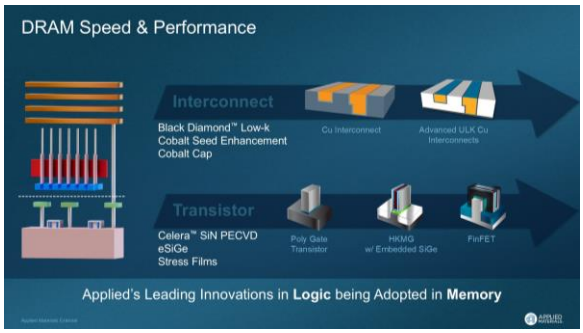

- What about speed? How are we helping DRAM achieve faster I/O and peripheral performance?
- We're transferring logic-proven speed boosters into the DRAM periphery: embedded SiGe and stress films for mobility, and interconnect upgrades – Enhanced Black Diamond™ low-k plus advanced Cu and capping – to lower RC.
- Net: faster I/O, better margins, and more cycles per joule.
- And on the near horizon, a shift in periphery transistor architecture from planar to FinFET promises another step in speed – a story for another day.
- Thanks for all the insights, Bala.
- It was great to be here, Mukund.

- Before we move to our next topic, let's hear from EVP Jong Myeong Lee, of Samsung Electronics, on the future of DRAM innovation and our partnership.
  [Samsung video plays]
- Thank you, EVP Lee, for sharing Samsung's perspective.
- Now, let's continue the conversation on integration and packaging.
- Let's invite Subi to join me on stage.

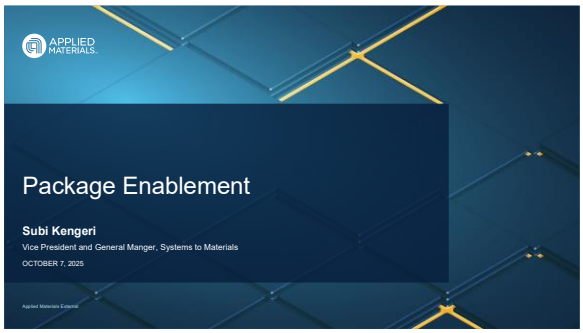

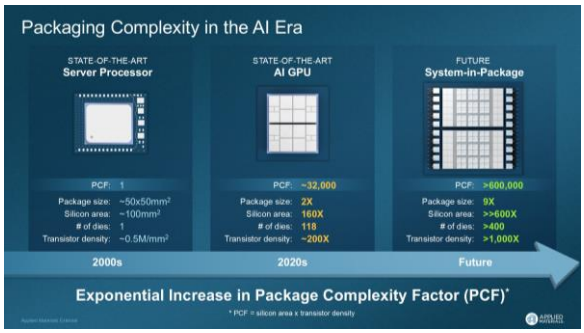

- Subi, let's shift gears and talk about packaging. It's often said that packaging is becoming the new frontier in enabling AI. Can you walk us through what's changing?
- Absolutely. Packaging has become far more complex than it used to be. A modern device holds 160 times more silicon area than one from 20 years ago, yet the physical size has barely doubled.
- The Packaging Complexity Factor, which multiplies transistor count by silicon area, has increased 600,000 times. That density is essential to meet AI's compute demands.

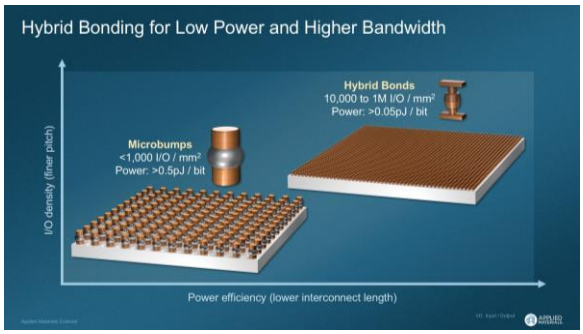

- That's a huge leap. What's driving this complexity?
- It's all about speed and power efficiency.
- AI workloads demand massive bandwidth and ultra-low latency.
- To achieve that, we need to increase interconnect density while shortening connections between chips.
- That's why the industry is moving from micro-bumps to hybrid bonding, bringing chiplets closer together to boost bandwidth and cut power.

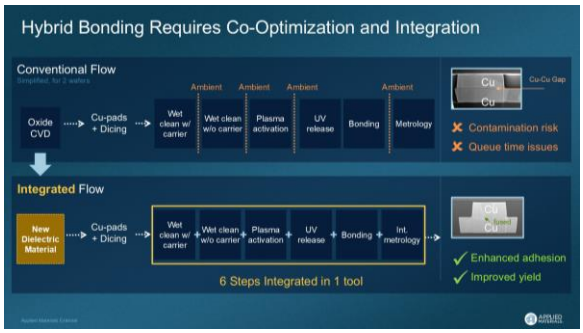

- And what makes hybrid bonding so challenging?
- Prabu described this earlier in his talk.
- Hybrid bonding isn't just about putting chips together; it's about doing it with nanometer precision and high yield.
- That requires new materials, advanced surface treatments, and integrated bonding processes.
- Every step – cleaning, alignment, bonding, and inspection – must be co-optimized to prevent particles, manage queue times, and maintain overlay accuracy.
- To maximize bond strength and yield, these critical steps need to be integrated and co-optimized on a single platform, and that's exactly what Applied delivers.

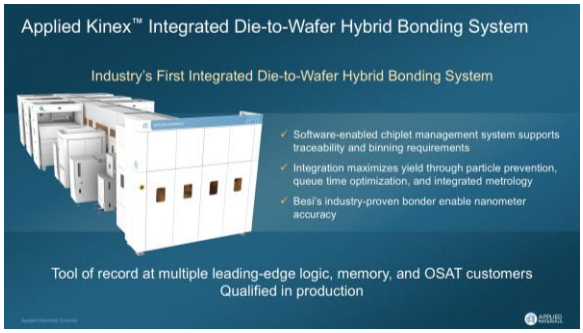

- So how is Applied addressing these challenges?
- This morning, we launched Kinex™ – the world's first integrated die-to-wafer hybrid bonding system.
- Kinex combines bonding, cleaning, and metrology on a single platform, streamlining the process and improving throughput.
- What makes Kinex unique is its software-enabled chiplet management system, which supports traceability and binning – critical for managing chiplets from multiple sources.
- It also maximizes yield through particle prevention, queue time optimization, and integrated metrology.
- And at the heart of Kinex is Besi's industry-proven bonder, delivering nanometer-level alignment accuracy.

- Sounds like yield management is a big part of the story.
- Exactly. Integration isn't just about putting chips together. It's also about the economics.
- High yield is critical as the material cost of the package is very high at this stage. Kinex ensures that by maintaining a super-clean mini environment, controlling queue times, and enabling precise overlay control. These capabilities are essential for high-volume manufacturing of complex AI packages.
- At the heart of it is our software-enabled chiplet management system. It optimizes queue time – the window between surface activation and bonding – and balances throughput across multiple chiplets. This simplifies how we manage different chiplets, with different performance bins, across multiple configurations.
- It's smart, scalable, and essential for high-volume AI package manufacturing.

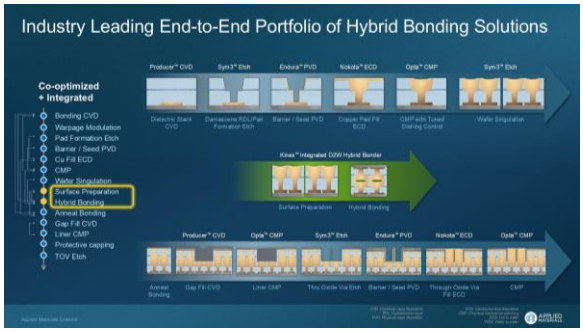

- Beyond Kinex, how else is Applied supporting the ecosystem?
- Applied is the only company with a complete end-to-end portfolio for hybrid bonding.
- We leverage our strengths in interconnect for advanced logic to co-optimize every critical step before and after bonding – CVD, Etch, PVD Plating, CMP, and inspection – so they work seamlessly with bonding.
- This integration ensures higher yield, faster ramps, and scalable production for the most advanced AI packages.

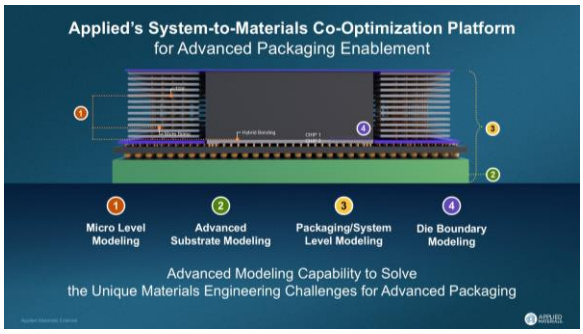

- You mentioned earlier that packaging is becoming a system-level challenge. How is Applied addressing that?
- That's where our Modeling and Simulation Platform comes in – a unique capability in the industry.

- It enables Systems-to-Materials co-optimization, giving us early visibility into application-specific design and technology challenges.
- This helps us prioritize the most critical solutions and develop process, metrology & inspection capabilities for faster performance, lower thermal, better reliability, and yield performance.

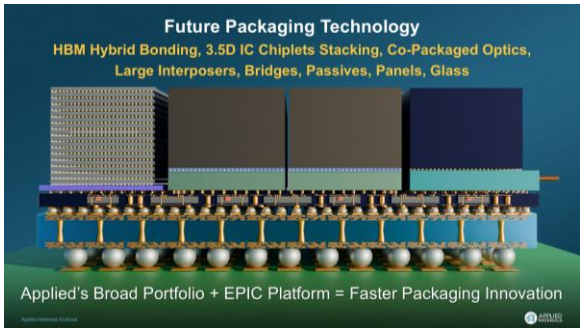

- Looking ahead, what's next?
- Complexity is only going up.
- HBM stacks are getting taller, and the industry is moving from micro-bumps to hybrid bonding.
- We're seeing more embedded devices, advanced power delivery, co-packaged optics, and new substrate formats – panels of different sizes and new materials.
- With the largest portfolio for advanced packaging, Applied is leading innovation across all these formats.
- For example, we're developing maskless digital lithography for sub-2-micron patterning, and panel eBeam test technology adapted from the LCD industry.
- So packaging is no longer just about protecting the chip, it's central to performance.
- Exactly. It's becoming a strategic enabler for AI. It's where materials engineering meets system-level innovation. The choices we make in packaging will shape how fast and efficient future AI systems can be.
- Great! Thank you, Subi, for that deep dive into advanced packaging and integration.
- It was great to be here, Mukund.

- Next, let's hear from VP Hoyoung Son, Head of Advanced Package Development at SK hynix, on the challenges and opportunities in HBM stacking and yield acceleration.
  [SK hynix video plays]
- Thank you, VP Son, for sharing SK Hynix's perspective.
- Now, let's shift to accelerating R&D in the fab.
- Mike, please join me on stage.

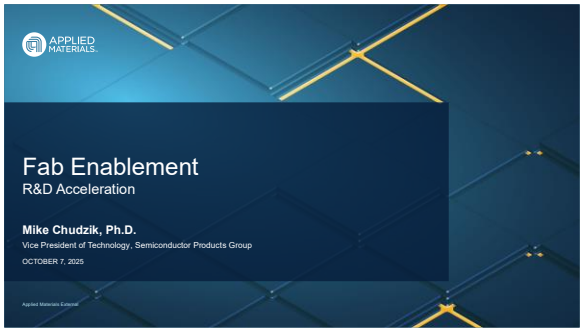

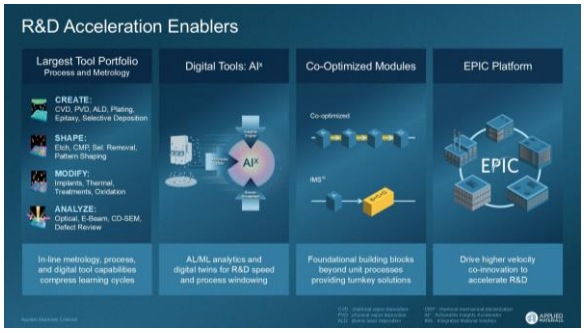

- As technologies get more complex and development timelines stretch out, how can Applied Materials help customers accelerate R&D?
- Applied is uniquely positioned to accelerate R&D through 4 pillars:
  - Interconnected Process + Metrology – Co-located steps with in-flow feedback that compress learning cycles.
  - Digital Tools in our AI^{X} platform – Sensors, AI/ML analytics, and digital twins accelerate R&D and enable process windowing.
  - Co-optimized Modules – Larger building blocks beyond unit processes that customers can drop into their flows.
  - EPIC Infrastructure – A global network of fab-like environments for realistic development with strict IP controls.

- Why does integrating metrology and digital tools with process matter so much, and can you share an example?
- Because speed is everything. A process is only as good as its ability to meet spec, and the fastest way to know that is massive data feedback in real time.
- Take moly contacts in advanced logic: As contacts shrink, tungsten hits limits.
- Selective, bottom-up Molybdenum delivers lower resistivity and better scaling – but tuning deposition and CMP for perfect fill and planarity is tough. There are a trillion contacts on a wafer.

- Applied solves this by pairing process with ProVision™ 10 eBeam metrology – tens of thousands of on-device measurements in under an hour – and AI^{X} digital tools for instant modeling feedback.
- Instead of waiting days for end-of-line electrical test, engineers see dishing, protrusion, and uniformity in hours, then adjust recipes on the fly.
- This approach compresses learning cycles dramatically, and it's why Applied won a major Mo metallization selection, covering deposition, CMP, and metrology.

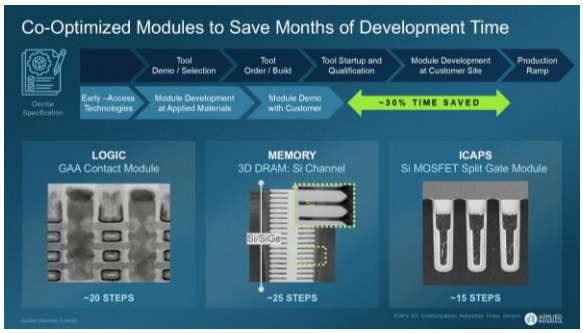

- Moving to the modules pillar – what exactly is a Co-Optimized Module, and how does it help customers accelerate development? Could you share an example?
- Here's the value: instead of waiting for fab readiness, we design, integrate, and validate the module upfront.
- We run demos with customer wafers, so when their fab is online, the module drops straight into production – saving months of development time.
- We've delivered modules for logic, memory, patterning, specialty devices, and heterogeneous integration, helping customers accelerate ramps and reduce risk.

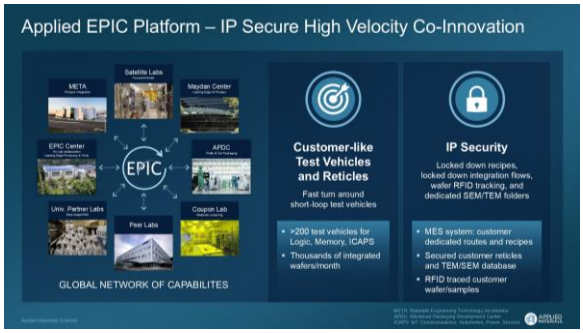

- And the last pillar. What kind of infrastructure do customers need, and how do you address confidentiality?
- Customers need realistic environments that replicate production fabs, and that's exactly what Applied delivers with our EPIC platform.
- It's a global network of co-innovation labs offering early access to breakthrough materials and tool innovations, years in advance, all in secure environments with uncompromising IP protection.
- We invest heavily in custom test vehicles that mimic real integration flows, helping us understand challenges and improve our materials and equipment for future inflections.
- EPIC connects fab-like sites worldwide:
    - Anchored on FEOL technologies in California
    - Advanced packaging in Singapore
    - Scales from coupon-level work to full 300mm wafers

    - Linked to major ecosystems like NY-Creates in Albany and IME in Singapore – giving us access to EUV lithography and advanced packaging flows.
- We combine industry-standard lithography and non-competitive steps with Applied's tools to deliver complete modules.
- Customers get:
    - Dedicated wafer corridors
    - Custom mask tape-outs
    - In-house TEM/SEM
    - Locked-down IP protocols
- So, they can innovate fast – and securely.

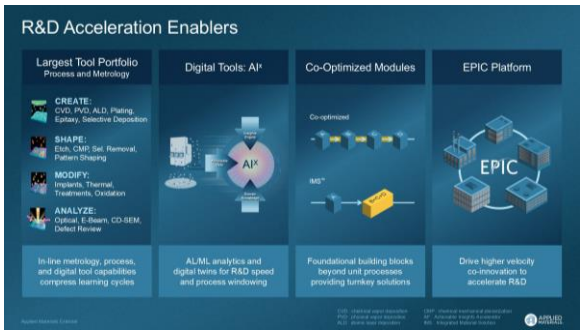

- We have covered a lot: integrated metrology, digital tools, co-optimized modules, and fab-like infrastructure. If you had to leave the audience with one thought on why these pillars matter now, what would it be?
- For our Partners, it all comes down to speed and strategic advantage.
- Every node and every new architecture brings more complexity and the cost of delay is enormous. Applied accelerates R&D by giving customers:
    - First access to breakthrough materials and tool innovations, years in advance,
    - The ability to shape and influence early technologies, and
    - Faster validation on their own devices.
- By combining fast feedback loops, ready-to-drop-in modules, and realistic development environments, we help customers learn faster, transfer faster, and ramp faster – all while protecting their IP.
- That's how Applied turns R&D acceleration into competitive advantage.
- Great! Thank you, Mike, for showing us how Applied is speeding up R&D and enabling innovation.
- It was great to be here, Mukund.

- Let's hear from David Bloss, CVP at Intel, on the importance of collaboration and yield acceleration in manufacturing.
  [Intel quote plays]
- Thank you, David, for highlighting the power of partnership.
- Now, let's talk about ramping to high-volume manufacturing. Steve, please join me on stage.

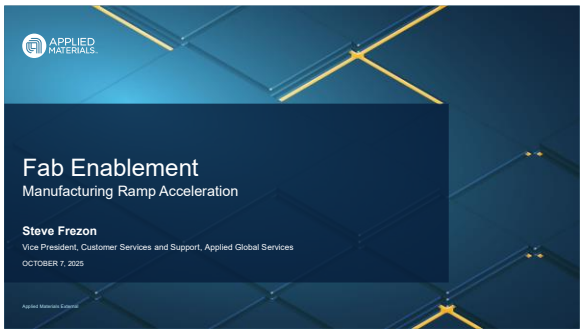

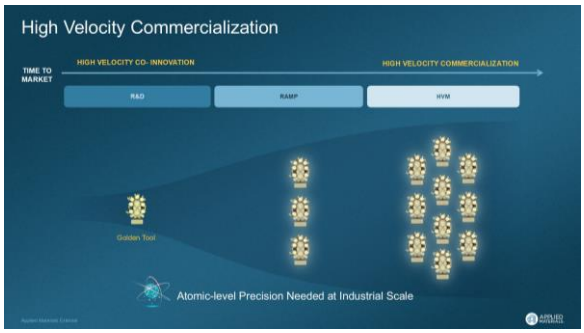

- Today, we're diving into how Applied Global Services is helping customers accelerate commercialization in the AI era.
- Let's start with the big picture. What do customers really need to win in the AI era?
- In one word – speed.
- Customers are racing to ramp faster, boost productivity, and hit yield targets.
- And as Prabu said, precision at scale is now non-negotiable.
- Applied is helping customers move from R&D to high-volume manufacturing with greater speed and precision.

- That's a bold promise. How exactly is Applied helping customers ramp faster and smarter?
- It starts with tackling complexity.
- The path from R&D to high-volume is complex, risky, and costly.

- That's why we built AI^X – Actionable Insight Accelerator.
- It's a powerful platform that blends AI, deep domain expertise, and data from our systems.
- With over 55 years of process knowledge, AI^X helps customers ramp new technologies with angstrom-level precision.

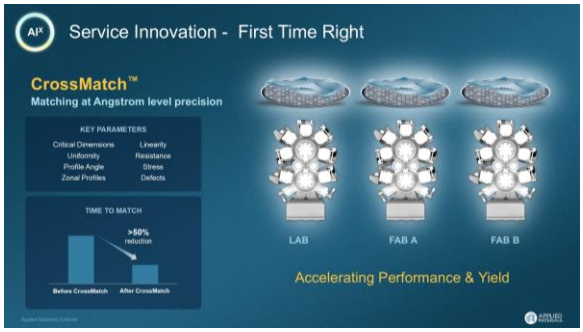

- Can you give us a real-world example of how AI^X is making a difference today?
- Absolutely. One of AI^X's standout capabilities is CrossMatch™.
- As Bala mentioned, reducing variability is critical.

- CrossMatch uses AI to match tools from Lab to Fab, and Fab to Fab – ensuring on-wafer results match the first time.
- It's automated, it's fast, and it works.
- In fact, it's helped customers cut time-to-match by over 50%. That's real acceleration.

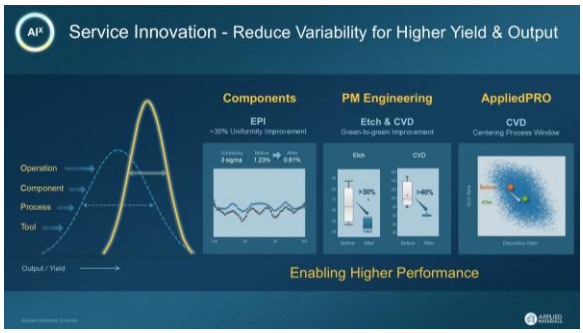

- Variability is clearly a big deal. How else is Applied helping customers tackle it?
- It's all about control.
- We reduce variability across operations, components, and processes.
- That means tighter specs, optimized process windows, and better tool performance.
- The result? Faster ramps, higher yields, and consistent output. Exactly what AI chipmakers need.

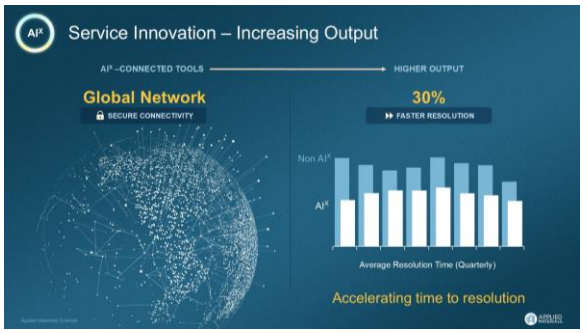

- Let's shift to connectivity. How does remote support change the game for customers?
- It's a game-changer.
- With remote connectivity, our AI^{X} experts can link into any fab securely and instantly.
- No flights. No delays.
- We're solving problems in real time, getting parts where they're needed, and cutting resolution time by 30%.
- And with predictive models, we're spotting issues a month in advance with up to 90% accuracy.
- That's impressive. Now, about talent. With the industry facing a talent shortage, how is Applied helping fabs do more with less?

- We're addressing the talent gap in two ways.
- First, we're scaling expertise – connecting global experts to every fab through AIx.
- Second, we're introducing cobots – collaborative robots that work alongside our teams.
- They boost precision, repeatability, & quality, helping fabs be more efficient.
- Let's go deeper on automation. How do cobots fit into the future of manufacturing?
- We're evolving, from onboard automation to human-assist systems.
- Cobots are key to enabling the human-assist strategy.
- They will handle repetitive, high-precision tasks like maintenance, calibration, and inspection.
- That frees up skilled workers to focus on higher-value tasks.
- It's smarter, faster, and helps fabs scale despite workforce challenges.

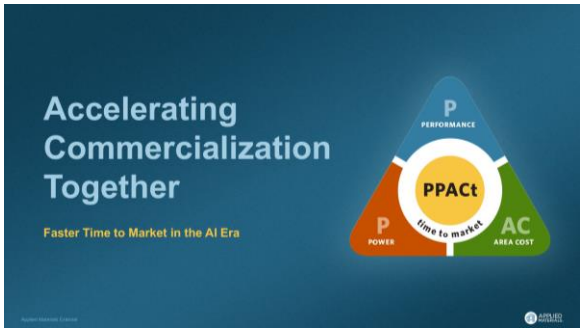

- Before we wrap, zoom out for us. How is Applied accelerating commercialization overall?
- We're building end-to-end service capabilities.
- From R&D, to ramp, to high-volume output, we are there.
- We bring materials engineering, advanced technologies, and critical components to help customers move faster and smarter.
- It's all about shrinking cycle times, optimizing yield, and maximizing value.
- Thank you, Steve, for sharing how Applied is accelerating commercialization and manufacturing ramp.
- Thanks for having me here.

- And thank you to everyone here for your attention and engagement throughout these discussions.
- Prabu will be back with us shortly to wrap up the session.
- But before that, let's hear from Olivier Hinsinger, Director of Operations at STMicroelectronics, who will share how digital technologies and AI^{X} are driving faster time-to-market.
  [STMicroelectronics video plays]

- Thank you, Mukund for laying out the path to next-generation AI devices.
- And a big thank you to our partners for sharing how Applied is working with them to enable the breakthroughs needed to power the future of AI.
- I hope you enjoyed seeing how Applied is co-innovating the inflection roadmap with our customers and the real progress we're making together.
- Now, let's shift gears. I'll summarize what this new roadmap means for our business.

- We are now tracking more than 100 fabs worldwide across all the devices. Some are in construction, some are ready, and some are in planning.

- Today, we talked about device architecture inflections enabled by materials engineering.
- These semiconductor inflections are driven by the industry's need for energy-efficient computing.

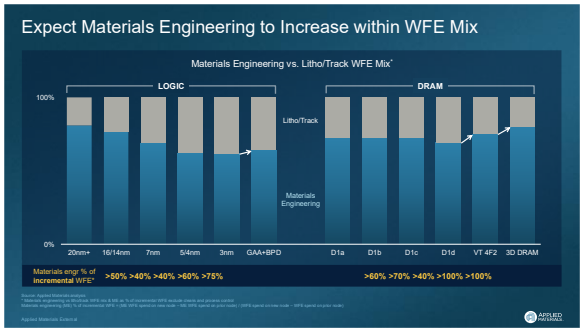

- Going in the third dimension is all about materials engineering.
- The chart shows materials engineering versus litho WFE mix over the nodes.
- In logic, this trend starts to reverse, with more materials engineering with these inflections.
- It is the same case in DRAM, with 4F2 and 3D DRAM. There is a big shift in the number of layers of materials engineering. Steps are going up, and the number of materials engineering content is going up.
- The row on the bottom shows the materials engineering percentage of incremental WFE.

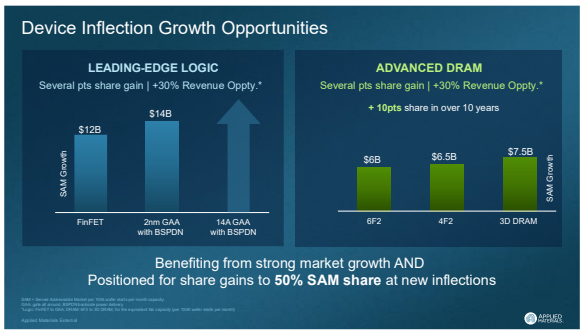

- We believe these inflections will increase materials engineering as a percentage of WFE, and this will increase Applied's opportunity.
- With Gate-All-Around and backside power, we expect our SAM to increase from $12 billion to $14 billion for 100,000 wafer starts per month of capacity. The SAM growth plus share gains increase our revenue opportunity by 30%.
- We expect logic SAM growth to continue as customers ramp the 14A node.
- We expect our DRAM SAM opportunity to grow from $6 billion to $6.5 billion for vertical transistor DRAM and by another $1 billion for 3D DRAM. The SAM growth plus share gains increase our revenue opportunity by 30%.

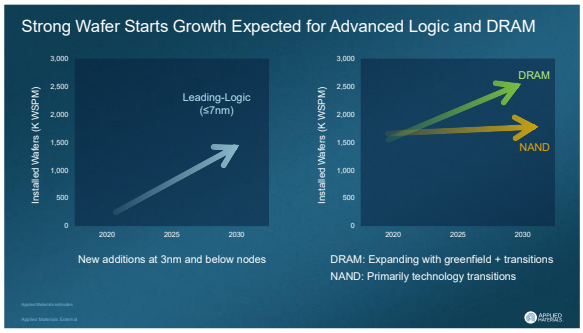

- In leading logic, we expect customers to continue to add capacity for the most advanced nodes ramping in the coming years.
- In DRAM, we expect continued growth in capacity, driven by high-performance DRAM and HBM.
- NAND is primarily technology transitions, and 20% bit growth is enough to supply the demand right now.

- Advanced packaging is a big story for Applied. There has been significant revenue growth in the past few years. It will continue to grow with all the inflections that are happening right now.
- The learnings from our leadership position in on-chip wiring expands to the packaging side for off-chip wiring.
- And we are building a leadership portfolio for all the advanced packaging inflections happening today.

- Our AGS business has had strong growth over many years, and the business is positioned to benefit from the industry's AI-driven growth.
- As tools become more complex, service becomes more important. There are a lot of digital tools, knowhow, and variabilities. This is a great opportunity for our AGS business, and we expect this business can continue to grow.
- Over two-thirds of our core AGS revenue is from subscriptions. We have over 18-thousand tools under service agreements, and those agreements have >90% renewal rate.

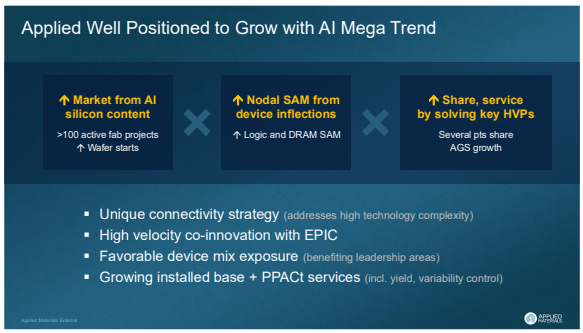

- To summarize – market growth, times SAM growth, times share and service increases – positions Applied well for growth.
- Thank you everyone for joining us today!