

# NASDAQ:AEHR Q1 2026 Earnings Call Transcript

Generated on 6/10/2026

## **Operator | Conference Operator:**

Greetings. Welcome to the Air Test Systems Fiscal 2026 First Quarter Financial Results Conference Call. At this time, all participants are in a listen-only mode. A question-and-answer session will follow the formal presentation. If anyone should require operator assistance during the conference, please press star zero on your telephone keypad. Please note, this conference is being recorded. I will now turn the conference over to your host, Jim Byers of Pondell Wilkinson Investor Relations. You may begin.

## **Jim Byers | Investor Relations (Host):**

Thank you, operator. Good afternoon. Welcome to Airtest Systems' first quarter fiscal 2026 financial results conference call. But with me on today's call are Airtest Systems President and Chief Executive Officer, Gane Erickson, and Chief Financial Officer, Chris Yu. Before I turn the call over to Gane and Chris, I'd like to cover a few quick items. This afternoon, right after market close, Airtest issued a press release announcing its first quarter fiscal 2026 results. That release is available on the company's website at [air.com](http://air.com). This call is being broadcast live over the Internet for all interested parties, and the webcast will be archived on the investor relations page of Airtest's website. I'd like to remind everyone that on today's call, management will be making forward-looking statements that are based on current information and estimates and are subject to a number of risks and uncertainties that could cause actual results to differ materially from those in the forward-looking statements. These factors are discussed in the company's most recent periodic and current reports filed with the SEC and are only valid as of this date, and Airtest Systems undertakes no obligation to update the forward-looking statements. And now with that said, I'd like to turn the conference call over to Gane Erickson, President and CEO. Thanks, Jim.

## **Gane Erickson | President and Chief Executive Officer:**

Good afternoon, everyone, and welcome to our first quarter fiscal 2026 earnings conference call. I'll begin with an update on the exciting markets areas targeting for semiconductor test and burn-in with an emphasis on how these markets seem to share a common thread of market growth related to the massive expansion of data center infrastructure and AI. After that, Chris will provide a detailed review of our financial performance. And finally, we'll open up the floor for your questions. Although we started with the typical low first quarter revenue consistent with the last few years, and actually higher on both top and bottom lines in Wall Street analyst consensus, we're pleased with our start to this fiscal year. We had revenue from several market segments and strong momentum in sales and customer engagement, both wafer level and package part test and burn-in of artificial intelligence or AI processors. Again, although we did not provide guidance for the quarter, our first quarter results surpassed analyst consensus estimates for both the top and bottom lines. We saw continued momentum in the qualification and production burden of packaged parts for AI processors, which is fueling sales growth in our new Sonoma ultra-high-power packaged part burden systems and consumables. During the quarter, our lead production customer, a leading hyperscaler, placed multiple follow-on volume production orders for Sonoma systems, requesting shorter lead times to support higher-than-expected volumes as they accelerate the development of their own advanced AI processors. This customer is one of the premier large-scale data center providers and has already outlined plans to expand capacity for this device and introduce new AI processors over the coming year to be tested and burned in on our Sonoma platform at one of the world's leading test houses. We're also collaborating with them on future generations of processors to ensure we can meet their long-term production needs for both package and

even wafer-level burn-in. Hyperscalers like Microsoft, Amazon, Google, and Meta are increasingly designing and deploying their own application-specific integrated circuits, or ASICs, for AI processing to meet the unique demands of their massive-scale workloads and gain a competitive advantage. AIR allows customers to perform production burn-in screening, qualification, and reliability testing for GPUs, AI processors CPUs and network processors directly in package form our Sonoma systems provide what we believe to be the industry's most cost effective solution enabling customers to smoothly move from early reliability testing to full production burn-in and early life failure screening which helps reduce costs improve quality and speed up time to market In the last year, AIR has implemented several enhancements to the Sonoma system to meet qualification and production test and burning requirements across a wide range of AI processor suppliers, test labs, and outsourced assembly and test houses, or OSATs. Major upgrades include increasing power per device to 2,000 watts, boosting parallelism, and adding full automation with a new fully integrated package device handler. Over the last quarter, including a very successful customer open house we held last week at our Fremont, California headquarters, 10 different companies visited AIR to see our next-generation Sonoma system and new features, including a fully automated device handler for completely hands-free operation, which we've installed here at our Fremont facility. Customer feedback regarding these enhancements has been very positive. And we expect these new features to open up new applications and generate additional orders this fiscal year. As I've mentioned before, one of the biggest benefits of our acquisition of in-cal technology one year ago is that it gives us a front row seat to the future needs of many top AI processor customers, providing us with close insight into the burning requirements. As the only company worldwide that offers both proven wafer-level and package-part burn-in systems for qualification and production burn-in of AI processors, AIR is ideally positioned to assist them regardless of their burn-in method. Consequently, we are experiencing increased interest in our Sonoma high-volume production solution for package-level burn-in, and some of these same customers, as well as other AI processor companies, are approaching us to learn about our production wafer-level burn-in capabilities. This past year, we delivered the world's first production wafer-level burn-in systems for AI processors. Importantly, these systems are installed at one of the largest OSATs worldwide, providing a highly visible showcase to other potential AI customers of our proven solution for high volume testing and burn-in of AI processors in wafer form, thereby strengthening our market position. We anticipate follow-on orders from this innovative AI customer as volumes increase. and other AI processor suppliers have already approached us about the feasibility of wafer-level burn-in of their devices. We're also developing a strategic partnership with this world leading OSAT to provide advanced wafer-level test and burn-in solutions for high-performance computing and AI processors. This joint solution, already in operation at their facility, marks a significant milestone for the industry. By combining AIR's technological leadership with this OSAT's global reach, we can provide unique capabilities to the market. This model offers a complete turnkey solution from design to high-volume production, and several customers have already begun discussions to learn more about our high-volume wafer-level test and burn-in solutions for AI processors. This OSAT and AIR have a long history of innovation together, including the first FOX-NP wafer-level burn-in system installed in an OSAT for high-power silicon photonics wafers. Now the world's first wafer-level test and burn-in of HPC AI products using AIR's FOX-XP systems. And they're also one of the largest installed bases of AIR's Sonoma system for high-power AI and high-performance computing processors. Additionally, this last quarter, we launched an evaluation program with a top-tier AI processor supplier for production wafer-level testing burn-in for one of their high-volume processors. This paid evaluation, which includes a custom high-power wafer pack and the development of a production wafer-level burn-in test program, will feature a comprehensive characterization and correlation plan to validate AIR's Fox XP production systems for wafer-level burn-in and functional testing of one of this supplier's high-performance, high-power processors on 300-millimeter wafers. We believe this represents a significant step toward adopting wafer-level burn-in as an alternative to later-stage burn-in and into future generations of their products. Our Fox XP multi-wafer test and burn-in system is the only production-proven solution for full wafer-level test and burn-in of high-power devices such as AI processors, silicon carbide, and gallium nitride power semiconductors, and silicon photonics integrated circuits. Beyond AI processors, we're seeing signs of increasing demand in other segments we serve, including silicon photonics, hard disk drives, gallium nitride, and silicon carbide semiconductors. We're experiencing ongoing growth in the silicon photonics market driven by the adoption of optical chip-to-chip communication and optical network switching. This quarter, we upgraded another one of

our major silicon photonics customers, FoxXPs, to the new higher power configuration, doubling their device test parallelism with up to 3.5 kilowatts of power per wafer in a nine wafer configuration. This latest system shipment includes our fully integrated and automated wafer pack aligner configured for single touchdown test and burning of all devices on their 300-millimeter wafers. We anticipate additional orders and shipments this fiscal year to support their production capacity needs for their optical I.O. silicon photonics integrated circuits. In hard disk drives, AI-driven applications are generating unprecedented amounts of data, creating ever-increasing demand for data storage and driving new read-write technologies for higher-density drives, particularly for data center applications. We're ramping and have shipped multiple Fox CP wafer-level testing burn-in systems integrated with the high-power wafer probe and unique wafer pack high-power contactors to a world-leading supplier of hard disk drives to meet the test, burn-in, and stabilization needs of a new device used in their next-generation read-write heads. This customer is one of the top suppliers of hard disk drives worldwide and has indicated they're planning additional purchases in the near term as this product line grows. Gallium nitride devices are increasingly used for data center power efficiency, solar energy, automotive systems, and electrical infrastructure. Gallium nitride offers a much broader application range than silicon carbide and is set for significant growth in the next decade. Our lead production customer is a leading automotive semiconductor supplier and a key player in the GaN power semiconductor market. And we have multiple new engagements with other potential GaN customers in progress. We're currently in design and development of a large number of wafer packs for new device designs targeted for high volume manufacturing on our Fox XP systems. Although silicon carbide growth is expected to be weighted toward the second half of the year, we continue to see opportunities for upgrades, wafer packs, and capacity expansion as that market recovers. Demand for silicon carbide remains heavily driven by battery electric vehicles, but silicon carbide devices are also gaining traction in other markets, including power infrastructure, solar, and various industrial applications. Late in last fiscal year, we shipped our first 18 wafer high voltage Fox XP system, extending beyond our previous nine wafer capability to test and burden 100% of the EV inverter devices on six or eight inch wafers in a single path with up to plus or minus 2000 volt test and stress conditions at high temperature. We believe we're well positioned in this market with a large customer base and industry leading solutions for wafer level burden. I also want to give a quick update on the flash memory wafer level burden benchmark we've discussed earlier. This benchmark is ongoing, and we've now begun testing with our new fine-pitch wafer pack that can meet the finer pitches and higher pin count costs more cost-effectively for flash memory, but also can be applicable for DRAM and even AI processors if they require fine-pitch wafer probing. This is the first wafer pack full wafer contactor demonstrating this capability. The benchmark has gone slower than expected with some challenges with the test system bring-up, but appears to show positive results of the new wafer pack, our ability to do an 18-wafer test cell, and using our full automated wafer handler and wafer pack aligner for the 300-millimeter NAND flash wafers. Interestingly, the market for NAND flash is in a state of flux. with earlier announced transition to hybrid bonding technologies for higher-density NAND flash on 300-millimeter wafers, driving new requirements for higher parallelism and higher power, to now a push for high-bandwidth flash, or HBF, which drives very different requirements in terms of test system capabilities. This is exciting news for AIR, as both are driving power requirements up substantially, which is right in our wheelhouse. High-bandwidth flash, or HPF, is an emerging technology developed by two of the flash market leaders and aims to provide a massive capacity memory tier for AI workloads by combining the DRAM high-bandwidth memory, or HBM-like packaging, with 3D NAND flash. This innovation is said to offer 8 to 16 times the capacity of HBM DRAM at a similar cost. delivering comparable bandwidth to dramatically accelerate AI inference and process larger models more efficiently while using less power than traditional DRAM. We're working with one of these lead customers on the now newer tester requirements to provide them with a proposal to meet even these newer, higher performance and higher power requirements within our Fox XP 18 wafer test and burn-in system infrastructure. We expect to have yet another update out at next quarter's earnings call. The rapid advancement of generative artificial intelligence and the accelerating electrification of transportation and global infrastructure represent two of the most significant macro trends impacting the semiconductor industry today. These transformative forces are driving enormous growth in semiconductor demand while fundamentally increasing the performance, reliability, safety, and security requirements of these devices across computing and data infrastructure, telecommunications networks, hard disk drive and solid state storage solutions, electric vehicles, charging

systems, and renewable energy generation. As these applications operate at ever higher power levels and in increasingly mission-critical environments, the need for comprehensive tests in burn-in has become more essential than ever. Semiconductor manufacturers are turning to advanced wafer-level and package-level burn-in systems to screen for early life failures, validate long-term reliability, and ensure consistent performance under extreme electrical and thermal stress. This growing emphasis on reliability testing reflects a fundamental shift in the industry. From simply achieving functionality to guaranteeing dependable operation throughout a product's lifetime, a requirement that continues to expand alongside the scale and complexity of next-generation semiconductor devices. To conclude, we're excited about the year ahead and believe nearly all of our served markets will see order growth in the fiscal year, with silicon carbide growth expected to strengthen further into fiscal 2027. Although we remain cautious due to ongoing tariff-related uncertainty and are not yet reinstating formal guidance, we're confident in the broad-based growth opportunities ahead across AI and our other markets. With that, let me turn it over to Chris, and then we'll open up the lines for questions.

### **Chris Yu | Chief Financial Officer:**

Thank you, Gabe, and good afternoon, everyone. Looking at our Q1 performance results, we exceeded analysts' expectations for both revenue and profit. First quarter revenue was \$11 million, a \$16 million decrease from \$13.1 million in the same period last year. It is important to note that last year Q1 benefited from a very strong consumables revenue quarter, which makes direct comparisons challenging. This quarter's revenue was primarily driven by demand for our FOX CP and XP products. In Q1, we shipped multiple Fox CP single wafer production tests and burning systems, featuring an integrated high-power wafer probe for new high-volume application involving burning and stabilization of new devices for our lead customer in the hard disk drive industry. Consumable revenues, which include wafer packs for wafer-level burning business and BIMs and BIPs for our packaged part burning business, totaled \$2.6 million and made up 24% of our total revenue in the first quarter, significantly lower than \$12.1 million or 92% of the previous year's first quarter revenue. As we have discussed in the past, this consumable business is ongoing even when customers are not purchasing capital equipment for expansion. We feel that this revenue will continue to grow both in terms of absolute value but also as a percentage of our overall revenue over time. Non-gap gross margin for the first quarter was 37.5%, down from the 54.7% year-over-year. The decline in non-gap gross margin was mainly due to lower sales volume and a less favorable product mix compared to a previous year, which included a higher volume of higher margin wafer packs. Also, our product ship this quarter included lower margin probes and an automated aligner, both manufactured by third parties and sold as part of our overall product offerings. Non-GAAP operating expenses in the first quarter were \$5.9 million, an 8% increase from \$4.5 million in last year. Our operating expenses increased due to high research and development expenses for our ongoing project, and we continue to invest resources and efforts to support AI engineering initiatives and the memory project. As we previously announced, we successfully closed the in-cow facility on May 30, 2025, and completed the consolidation of personnel and manufacturing into a three-month facility at the end of fiscal year 2025. In connection with the facility consolidation, we eliminated a small number of headcount due to redundancy in our global supply chain and incurred a one-time restructuring charge of \$219,000 in our fiscal first quarter. In the first fiscal quarter of 2026, we received \$1.3 million of employee retention credit from the press for eligible businesses affected by the COVID-19 pandemic. We reported this cash credit minus the professional fee to process the refund in other income on our income statement. In Q1, we recorded an income tax benefit of \$0.8 million and a productive tax rate of 26.5%. Non-cash income for the first quarter, which includes the impact of stock-based compensation, action-related investment and return charges, was \$22 million, or \$0.01 per share, compared to \$2.3 million, or \$0.07 per share, first quarter of fiscal 2025. The consensus on the income for the first quarter of fiscal 2025 was even. Our patent lock at the end of the year was \$15.5 million. With \$2 million in bookings received in the first five weeks of the second quarter of fiscal year 2026, our effective ad lock now totals \$17.5 million. Turning to our cash flows and balance sheet. During fiscal year, we used \$0.3 million in operating cash flows. We ended the quarter with \$24.7 million in cash, cash equivalents and restricted cash, compared to \$26.5 million at the end of Q4, mainly due to the final 1.4 million payment for facility renovation. In total, we have spent 6.3 million on

remodeling our manufacturing facility. With the renovation now complete, we have significantly upgraded our manufacturing floor, customer and application test labs, and clean room space for wafer pack, full wafer contactors. Improvements increase our power and water cooling capacity enabling us to manufacture all of our FoxWave-level burning products and packaged part burning products, including Sonoma, Tahoe, and Backhoe products on a single floor. We are very excited about this renovation as it was specifically designed to enable us to manufacture more high-power systems for AI configuration. We believe investment in this facility renovation has increased our overall manufacturing capacity by at least five times, depending on the current product configuration and we are more ready than ever to support the growth of our customers. We celebrated these upgrades with a customer open house that was well attended and received very positively. Over the past quarter, we hosted many package-targeting, buffer-level running customers who had the opportunity to see our expanded capabilities firsthand. Importantly, we do not expect and anticipate additional capital expenditures for facility expansion in the near future. We have no doubt and continue to invest our excess cash in money market funds. As Kay mentioned, we started the year by withholding formal guidance due to the ongoing tariff related uncertainty. Since we remain cautious, we'll continue with that approach for now. However, looking ahead, we're confident in our base growth opportunities across AI and our other markets. Lastly, look at the investor relations calendar. EdTech will meet with investors at the 17th annual CEO Summit in Phoenix tomorrow, Tuesday, October 7th. This following month, we will participate in the Big Hal on the 16th annual Alpha Conference in New York on Tuesday, November 18th. And on Tuesday, December 16th, we will return to New York City to attend the NYC CEO Summit. We hope to see some of the EdTech conferences. This concludes our prepared remarks. We're now ready to take your questions. Operator, please go ahead.

### **Operator | Conference Operator:**

Thank you. At this time, we'll be conducting our question and answer session. If you would like to ask a question, please press the star 1 on your telephone keypad. Confirmation tone will indicate your lives in the question you give. You may press star two if you would like to remove your action from the give. For participants using secret equipment, it may be necessary to pick up their handset before pressing the stop keys. One moment, please, while we poll for questions. Once again, please press star one if you have a question or a comment. Please continue to hold while we adjust some sound tech issues. One moment.

### **Operator | Conference Operator:**

Thank you for standing by. This is the operator once again. Christian Schwab, your line is live. Please go ahead. Great.

### **Christian Schwab | Analyst:**

That sounds like a much better connection. So, Gain, you know, as we kind of get into the second half of the year and kind of these more open-ended growth opportunities in AI that you've talked about in particular, you know, when do you think we'll see a material improvement in bookings to drive revenue down the road?

### **Gane Erickson | President and Chief Executive Officer:**

Well, that sounds an awful lot like guidance again here, but so what we believe and what we've tried to communicate in our previous calls as well is that, you know, our lead, our first AI way for low-earning production customer, we anticipate that they will need additional capacity that would be both bookings and revenue for this year, and That could be more than last year, and we won't put a top on that. So, you know, the question is timing of that. We're not sitting on an order. We didn't get it yet and just put it in our pocket. But as that order comes in, we typically will announce those within, you know, a couple of business days or

so. What we are seeing is additional wafer-level customer engagements. It's pretty interesting that kind of span from, processors and asics and I'm sorry okay hold on that was a person can you hear me okay oh wow all right I can hear you again okay All right. So I'll assume that Christian is on mute or something that he can hear me as well. So we're seeing it across several different groups from hyperscalers, AI processors, kind of across the board. And it's interesting. We had direct people that have come in saying that's what they're interested in. We have people that are talking to us about Sonoma because their current customer is already doing qualifications and are looking to do burn-in for the first time and are looking at either package and also now exploring the wafer-level side of things. So these generally do take some time, and, you know, so I would – Probably guess these tend to be more second half, this being the second fiscal quarter of fiscal 26 for us. But, you know, at this point, we're just scrambling as fast as we can to address all the requests and requirements and keeping our head down to focus on them. On the package part, same thing, both additional quals and additional processes that are being put on our system and its enhancements to the Sonoma. as well as we've got customer interest to do additional production customers, you know, with and without the fully automated integration of the pick-and-place handler that bolts right onto the front of Sonoma. So I think it's ongoing and very interesting, and we're just really happy to have this number of engaged and active customers. Operator, can you hear us?

**Operator | Conference Operator:**

Yes, I can hear you. And are you ready for the next question?

**Gane Erickson | President and Chief Executive Officer:**

Yeah. Christian, do you have any other questions?

**Operator | Conference Operator:**

It seems a little rough this time. We have a question coming from Christian Schwab. Christian, your line is live. Go ahead, please.

**Christian Schwab | Analyst:**

Sorry about that, Gane. I was telling you I could hear you, but it wasn't working either. So we have, you know, a few customers here currently. You talked about, you know, a bunch of more customers coming in there. As we look to the end of your fiscal year, do you have a target number of customers that you think you'll be in the process of shipping to by then or shipping to fairly shortly afterwards?

**Gane Erickson | President and Chief Executive Officer:**

That's a good question in terms of targets. Actually, we do have some discrete quantity targets. In fact, some of the KBOs, which are the bonus structures for our officers, are based upon not only numbers but specific targeted AI customers. We're really given a lot of insight, but I would say in plural for additional package part and also for wafer level. At this point, we're not really limiting ourselves, but we're just trying to be cautious about oversetting expectations either in terms of the timeline of it. One of the things that was interesting that really came to fruition, and I apologize if I said this before on the last call, is I'm starting to also understand a couple of things going on. One of them that was kind of new is there are many of the ASIC suppliers in particular, and there's some evidence within the you know, the GPU or just the processor suppliers themselves, they don't do a production burn-in like you think about it, like using one of our tools. They're doing it at system level, like as in the rack. So these processors are getting all the way to the end, and then

they're simply running them in rack form, sometimes at elevated temperatures and sometimes not, to try and get, you know, the first seven days of failures out of them, which is so inefficient and uses a ton of power. And, you know, there's only so many processors per rack, if you will. And so I was sort of surprised at some of this. You know, some of the test vectors that we're getting from customers are not, this is on a production tool today. This is just an HTAL, which is like a qualification vector instead of a production vector. And that's because they weren't doing production yet. so you know you know you're really at the leading edge of this um but one thing is really clear from the data we've seen so far the devices are failing we do see the failures in the burn-in so they're absolutely able to screen them using our tools at wafer and production and so you know that creates the leading edge of this market and why we're so excited about it's really easy i mean obviously Every single call you get on, your CEOs are talking about how they're using AI one way or the other. But this is really happening to us. I mean, it was 40% of our business last year from zero. We think it's going to grow, both package and wafer level, this year. And, you know, we're still seeing the other businesses grow as well. So we're really glad to have gotten the facility upgrade, you know, behind us. There's a lot of work to get that there. Now we have the capacity to be able to ship so many more systems, particularly the high-power ones. And if you come on our floor right now, you'll see AI wafer-level burning systems right next to Sonoma systems being built today. So, you know, I think that we believe that we have the opportunity to capture, you know, multiple customers in both package and wafer levels.

### **Christian Schwab | Analyst:**

And then my last question, Gain, is, you know, last call you were quite enthusiastic about the TAM for, you know, AI-driven products for you to be three to five times bigger than silicon carbide. And is there a timeframe that we should be thinking about that, you know, becomes evident? Again, I kind of asked it on the backlog question, but I'll ask it again more directly. You know, are we going to see material orders from, you know, one or two customers, you know, this fiscal year? Or is that something that it's just too early to know? But, yeah, you can feel confident it's going to come. How should we be thinking about that?

### **Gane Erickson | President and Chief Executive Officer:**

I feel the latter is the easy out to say that I'm confident they'll come. I think timing it can both be, you know, a lot more guidance than we're providing right now. But there's also just some of these evaluations as we prove it, the customers can actually start contemplating how many and when they would want to install them. You know, the new evaluation, I think we already alluded to it, it's for a processor that is expected to go into volume production at the end of next year or in the second half of next year. So, you know, tools would be needed to be going in in that timeline. So, you know, if you just – we do fiscal years through, in this case, fiscal 26 is through May of 26. If you talk about calendar 26 – there's a lot of opportunities in play that need to play out that would be production for both wafer level as well as package. So it's not that far away. I mean, even something that seems like is a one year away in our space, there's a lot of work that needs to be done to actually ramp a customer to be one year out. And so, um, we'll keep, uh, you know, we'll keep focused on this thing as we get a little closer, we'd hope to give you answers. Um, To be candid, this will probably feel like, you know, you'll hear enthusiasm and we think we're winning and, you know, the customer has gotten good results. Those will be early indicators. And then, you know, we're going to surprise everyone with a large production order, not unlike what happened with the first wafer level system, except for some of these customers are just significantly bigger.

### **Christian Schwab | Analyst:**

Great. Thank you. No other questions. Thank you.

## **Operator | Conference Operator:**

Thanks, Christian. Thank you. Your next question is coming from Jed Dorsheimer. Jed, your line is live. Please go ahead.

## **Mark Shooter | Analyst:**

Hey, Gabe. You have Mark Shooter on for Jed Dorsheimer. Hey, Mark. Congrats on the success for this quarter and the announcements for the AI customers, and that's great. Can you give us a little color? How should we think about the engagement in the qualification cycle for these customers? Do you need a new product cycle to occur? Do you need to slide in between Blackwell and Rubin? And if you can give us a little bit of what's it like in the room with the customers. Is the tenor of these as risk aversion, or is the overwhelming demand spur some willingness to try a new equipment like AIR?

## **Gane Erickson | President and Chief Executive Officer:**

Oh, that's actually, there's a lot in there. Those are good ones. All right, so let me talk about sort of the qualification process. So, so far in the engagements that we've had so far, we don't need a new product. So we are doing some things depending on the pitch of their probe cards, which we call our wafer packs. We may need to do some things specifically for that. We have some design for testability features that we have been touting to our customer base that allow them very short lead time, high volume, low cost wafer packs. We can also supply them at higher cost and a little bit longer lead time. if they don't hit those DFT targets. We've got some of both. And so, like, one of the engagements, we made a conversation related to them about their pitch of their devices. And we're like, wow, you know, you happen to choose a pitch on these so many pins. That's driving the cost of your wafer pack-up. And they're like, well, why didn't you tell me before? And they kind of joke because they hadn't talked to us before. And they're like, well, this will be no problem to cut in for our next generation, but we're just going to have to live with it on the current one. so you know they're they're engaged with us in kind of a roll up the sleeves working the qualification in some cases is just validating that we can do the same type of dft and power delivery as we've done with the other processors on their devices i think customers i get it they're they're kind of like it's hard to imagine that we can really pull this off you know if they haven't seen it with their own eyes and so we're just showing it and demonstrating it to them somewhat like what we ended up doing with the first silicon carbide customers. And then at some point, people get it. Now, one thing that also seems to be going on is, you know, these are pretty visible. I already said that these systems are sitting at an OSAT, and there aren't that many of them, okay? So especially not that many of the biggest, right? There's a lot of people out there that are aware of the success of this. And even though the analysts and all are still trying to figure out everything, there's a lot of people that have pretty intimate knowledge and seem to know what's happening. And so they're like, can you do – can I do it that way too? So they're leaning in. So it's a little less of, you know, complete disbelief, can you do it, but more of can you prove it for me. Now, from a timing perspective, it's, you know, just typical of the industry – Normally, when people are buying test equipment, like semiconductor test equipment like ours, you do it at some disconnect. Either you're putting a new fab in, if you're an IBM, or it's with some new product. Or if just simply the volume is growing so fast that you want to buy a tool that has more output per dollar or so. So in this case, you know, outside of one supplier, everybody is using TSMC today, and eventually Tesla will be using the Samsung stuff. But it's not like there's a new fab, although there are new fabs coming online. People are just getting access to those TSMC wafers and then want to be able to test them. And they either do it in a package for Vernon on something like Sonoma or a system-level test or off-lead all the way back at the RAC. So customers are engaging because they need to buy capacity for these new products and for new things coming out. So it is a fair way of looking at it to look at the intercept between product A to product B. That's at least what's been communicated to us with this latest one we just announced. And similarly, our first customer intercepted us with their transition to a newer device. We announced that a year ago. So that's pretty typical, and sometimes that's the gating item of their timing, and sometimes that's fast or slow, but it's

sort of you need to time it with that. Just the tenor, the tone. So, you know, if you guys, people that have followed us understand that, you know, our value proposition, our pitch, if you will, is that semiconductors are growing, you know, extremely high. So, you know, within... It took 40 years to get to 500 billion. It's going to take less than 10 to double that. Much of that is driven by either directly AI or all of the pieces surrounding all of the explosive data center growth. What's happening is these devices are not more reliable for multiple reasons. The smaller and smaller geometries and the fact that they're putting multiple devices into one package because they can't make the devices any bigger are driving the requirements for reliability and burn-in test. And if you look at the roadmaps from all of the players, every single one of them, from all of the NVIDIA products to everyone else, from the ASIC suppliers, all of their products going forward are pulling multiple compute processors to make it generic in a single package, along with many, many stacks of HBM and ultimately optical I.O. chipsets. They put these on these complex advanced packaging substrates, and they're extremely expensive. And I always remind people, the reason you burn them in is because they fail. And when they fail, you take out all the other devices. So the value proposition, if someone could ever do wafer-level burn-in, is overwhelming because the cost of the wafer-level burn-in is cheaper than the yield loss. I actually alluded to it in my prepared remarks that our lead customer for Package Part Burn-In is going to do a couple few generations in Package Part and then wants to switch to wafer level. So, you know, what are they going to do with all those Sonomas? It doesn't matter. The yield advantage of moving it to wafer level pays for it all. So, you know, that's a thing that's a macro trend heading our way. And it's not just AI. It happened to us in the silicon carbide side of things. We see it in stacked memories in both DRAM and flash. We see it in other complex devices in GAN that are going to automotive that are mixing different devices together and why it's driving for wafer level. And these large trends are good for both reliability as a tide that's rising for all, and really good for us, but also for our unique products, particularly the Sonoma and the high-power wafer-level burning systems we have with our Fox products.

### **Mark Shooter | Analyst:**

Dane, all that color is very helpful. Thank you. To dig in a bit around that last part of the Sonoma versus the Fox products, what's the gating factor of why customers are going first with the Sonoma and not right to wafer-level burning? What needs to be proven out? for wafer-level burn-in for those customers? And how – I'm assuming there's a sales cycle there of you'd like to start with Sonoma and then push people to wafer-level burn-in. So how does that transition go?

### **Gane Erickson | President and Chief Executive Officer:**

Yeah, you know, the way we look at it is we say we're just neutral. If you want to do package part or you want to do wafer-level, we love you both, okay? Okay. it's not easy to just go, you know, talk someone out or whatever it is they're used to. So in this case, we don't have to. We just say, listen, we think we make the best machine for qualification reliability of your complex packages with Sonoma. They can test all the processors, HBM, and all the chipsets inside of it in a single path during your quals. If you want, we'll do it in production as well, and we're now adding automation to it. But if you'd like to kind of go to the next step, you could take the high-failing devices out of there and do a wafer-level burn-in of them before you put them in those packages. And our data would suggest you don't need to burn them in again. But, you know, if you still need a little burn-in, that may be fine, but you don't want to have the massive yield loss. Some of these processors have four and eight CPU chips in them, right, compute chips, and have another, you know, six or eight HBM stacks on it. Just the co-loss substrate is extremely expensive and rare. And so, you know, it makes sense to go to wafer-level, but, you know, to be candid, one year ago, 12 months ago, we didn't even have the first order. There was not one machine in the world that could do a wafer-level burn-in of an AI processor. None. We're the only ones, and we've now just, you know, And, you know, we're at the front end of this thing. I understand people are sort of in a doubting mode. Let us prove it to them. And for those that are on the call, if you have a processor, you can sit down with us under non-disclosure. We can tell you which exact specific files we need, and we can do a paper benchmark and give you an answer within a couple of days. as to the feasibility of

your devices. And so far, we have not found one that we haven't been able to test that we've been given that detailed data on. So, I'm sure there are some out there. But for now, we're on a roll.

### **Mark Shooter | Analyst:**

Okay.

### **Operator | Conference Operator:**

Much appreciated. Thank you. Thank you. Your next question is coming from Bradford Ferguson. Bradford, your line is live. Please go ahead.

### **Bradford Ferguson | Analyst:**

Hello, Gane. I'm curious about the cost to wait until you get to the motherboard or the package part or the final part. When we were talking about silicon carbide, you could have 24 or 48 sick devices in one inverter, and then the whole inverter is bad, and maybe that's a \$1,000 or \$2,000, but the retail price on these NVIDIAs is, what, \$40,000?

### **Gane Erickson | President and Chief Executive Officer:**

Well, the rumor is they have really high margins, and I'd love it if the customers would give me credit for their sales price. They really only give me credit for their cost, but fair enough. But their cost is significantly higher than any silicon carbide module ever would be. Fair enough. Yeah, I mean, it's, you know, and by the way, to me, the craziest thing is how many people are doing it at the rack level. Like, you're talking about all the way at the computer level side of things and burning it in. And, you know, obviously, a failure there is, you know, a lot more expensive than it would be all the way back at wafer level. So you want to move, in our industry, we refer to shift left. You want it to go as far left in the process as possible because it's way more cost effective. In this case, we have the first two steps in the left side, wafer level and when it's just the module level. Before that module is then put actually into the system level where you'd start to see all of the power supplies and everything else on it, you know, like the GB200 module itself. And then you, certainly before it goes over super micro or \$2 or something in some mainframe rack. So, you know, one thing to put in perspective, and I don't think this is the value proposition yet, but it is interesting. We know that people are doing this burn-in at the rack level or the computer level, right? When you're in the computer level, basically what burn-in does is you're basically applying stress condition of, power via voltages or current, and temperature. And what it does is it accelerates the life of the part without killing it. So I can take a device and in 24 hours make it look like it's one year old, and if it hasn't died by then, it's going to last 20 years. There's all kinds of books on it. You can read it, Google it, or something, and you can find out about the basic process of burn-in and why you do it. Thanks. The key here is you want to do it in, you know, 24 hours or four hours or two hours or something along those lines to get the infant mortality rate out so it doesn't shift to the customer or take down your large language model compilation, okay? Now, when you're at system level, you can't run that rack at 125 degrees C. Everything will burn up. In fact, those racks are running cold water through them. they're probably running 30 degrees C temperature maximum. I know of a company that was trying to do some things to try and get an isolation of the GPU or the processors to 60 degrees C, and their burn-in time was measured in days at the system level. That's what they were doing. Now, by moving it to wafer level, we can actually run the devices at a junction temperature at 125 degrees C which is an accelerant that's more than 10x. We can also run the voltages extremely closely to their edge, and we can get the burning times to come down. So when we do that, we're actually applying only power to the processor, not the HBM, not all the inefficiencies everywhere else, not the rack, et cetera, just to the processor, and we can do it for a significantly less amount of time. The long and short of it is I can burn it in to the same level of quality at a fraction of the power. Now, I don't

think anyone's going to buy our system because of that per se, although there's some argument for it. But you know what's hard? Getting a permit for a megawatt burn-in floor for your racks. So people may buy our systems because they can actually get the power infrastructure to burn in hundreds of wafers at a time in parallel and in a regular 480-volt, maybe 1,000 or multi-thousand-amp circuit like we have in our building. You wouldn't be able to do that. If you had to burn in a bunch of racks in our building, you wouldn't be able to do it. But I could have 10 systems running with nine wafers apiece and test 100 wafers at a time with the power that I have in my facility, which is not that atypical of a facility – in the Bay Area in Silicon Valley. So there is a value proposition there. In addition to the real cost savings, it might just be feasibility of power.

### **Bradford Ferguson | Analyst:**

And so you mentioned the high bandwidth flash. I'm hearing from some systems makers that they're focused on burn-in more. just because of how expensive it is to, you know, scrap the whole motherboard or whatever. Do you have any kind of end to high bandwidth memory, or is it mainly the high bandwidth flash?

### **Gane Erickson | President and Chief Executive Officer:**

Yeah, I mean, we talked at kind of our first – our belief was that the engagements and the interest was first on the HP – yeah, on the flash side of things. There is some things – there's discussions on the DRAM side of things. I mean, people are really scrambling to try and solve that through all kinds of mechanisms, and I won't get into all the technological things that we understand. You know, there's very different implications when you talk about Micron, Samsung, and Hynix. and what they do and how they stack their memories and how they test them and burn them in that have, you know, kind of key differentiating features amongst themselves that make tests interesting. We have a pretty good insight to that. I'm certainly not going to talk about it publicly, but that makes that interesting. Bottom line is, you know, high bandwidth memory and then eventually high bandwidth flash needs to be burnt in. and needs to have a cycle and stress to remove that somehow, or it's going to show up as it has been in the processors, in the AI stacks. You know, and that's widely known and understood. And, you know, NVIDIA came out last, what, six months ago, yelled at everybody and said, you need to figure out how to burn these things in before you ship them to me. We're sick and tired of it. So, you know, I'm not creating rumors. Those are widely understood reports. And so right now what we're seeing in the test community is sort of, you know, people overuse, you know, the Wild West, but there's just people scrambling for good ideas on how to address this and running as fast as they can. And, you know, it makes it exciting every day when you show up to work and you've got people that are like, how can, you know, how can you help us? So I love our hand. I love the cards we're dealt right now. I love our position. I love our visibility that we have within. Pretty much all, I think we can now say we have communicated with every single one of the AI players. And, you know, we have a line into them and some thread, either package or wafer level related, that gives us some great insights. And I think we may be completely unique in that realm. So I think the HP app, it looks pretty interesting. Again, you know, that stuff takes time. But more and more things are breaking the infrastructure of tests because of power at wafer level, and that's a good thing for us. We're really good at that. Our system, you know, I just throw out 3.5 kilowatts per wafer, and, you know, most people would not know what that means. That's crazy. I mean, you know, the world has wafer probers, you know, thousands of those installed, that has 300 watts of power capability. If you try to go get a prober that has 1,500 to 2,000 watts, it's a specialized half-a-million-dollar prober. It's what we ship with the CP to the hard disk drive guys. That's one wafer's capacity. Our systems can do 3,500 watts on each of nine wafers in one machine. Nobody can do 3,500 watts on one machine. I'm sorry, on one wafer on one machine. And so people are coming to us because of the thermal capabilities that are unique. Many, if not most of them, are patented around the whole wafer pack concept and the blade where we deliver thermal power without a wafer prober to create uniformity across a 3,000-plus-watt wafer is really awesome. And it's fun to talk about with the technical people. And they're, you know, I'd say that people are quite impressed with what they hear. And so, it's great to rotate people through here. And by the way, they see it. We can show them in operation when they come. This is not a story. I think, you know, the more and more of these

things, the rising tide, you know, the better shape we're in. And we're not abandoning our silicon carbide customers that are listening. I know they have ramps. They have opportunities. There's new fabs. There's new capacity coming on. They have new technologies. We're not abandoning the OEMs, the electric vehicle suppliers that we have met with personally and helped them to develop the burden structures and the burden plans that they drive their vendors towards. We're fully committed to those guys, and we'll be there as they rep, and we have more capacity than we ever had to be able to address their needs at a lower price point. I think we got that covered. We're not pivoting the company. We're just adding to it with this AI stuff.

### **Bradford Ferguson | Analyst:**

On silicon carbide, this will be my last one. Thank you for your generosity. On CIMI, I think one reason for their success is how aggressively they adopted air test systems, Fox XP systems. And we have a pretty large bankruptcy that happened with one of their competitors. Is there some kind of risk for the other chip makers if they don't take if they don't take burn-in more seriously, that it could spell issues for them?

### **Gane Erickson | President and Chief Executive Officer:**

So let me answer it this way. I have been invited to be a keynote speaker. I've spoken at multiple technical conferences around the world, and Silicon Carbide and Galvanitride conferences. I've sat on several panels, and I have been very almost emotional in some of those discussions. because we have seen the test and burden data of more, almost all of the wafers in the world, okay? That's pretty bold, okay? Certainly more than anyone by far, okay? Everybody would like to think that they are special and their devices are just so much better than everybody else's. The reality is that these devices fail immediately. during burn-in that represent the actual duty cycle or what's called the mission profile of electric vehicles. What that means is if you do not burn them in, it is our belief in the data that we have, they will fail during the life of the car, period. We've talked about that. I think I've quoted several times. Whatever you do, it is my opinion, never buy an electric vehicle that didn't have burn-in for something in the six to 18 hours, depending on the size of the engine and things like that. And there are OEM suppliers that have the data. They have failed customers who try to qualify without doing an extensive burn-in and kick them out. And there have been very large suppliers that have lost in the industry because of quality and reliability. So my call to arms for everybody is there's no reason – not to do wafer level burn-in or package part if, you know, if you don't want to go with us. But whatever you do, don't skip it. And we now with our 18-wafer system, even at high voltage, okay, so we've extended the capability with more capabilities. The cost of test at high voltage on our system with a capital depreciation of five years, et cetera, is about .5 cents per die today. on an 8-inch silicon carbide inverter wafer per hour. Per hour. You can do 24 hours of burn-in for 12 cents a die. And we have been very clear with that to all the OEMs, and they understand it. And so they drive for a level of quality that they can measure directly on our tools from their suppliers. And I think there is a difference between the people that have adopted a high level of quality and reliability in their market share. And all I'll say is, you know, I think On Semiconductor has done an incredible job. You know, in 2019, I think the year before, they had done \$10 million in silicon carbide, and they're now, you know, kind of neck and neck for market leadership, and they have won well more than their fair share of the industry across the, and I'm just repeating what they have said, across Europe, the U.S., Japan, and even China. They have done really, really well, and I commend them for that.

### **Operator | Conference Operator:**

Thank you. Your next question is coming from Larry Chapina. Larry, your line is live. Please go ahead.

**Larry Chapina | Analyst:**

Hi, Gane. The news today on the AMD – hook up with OpenAI. Does that accelerate your evaluation process that you have with that second process, or does that put more pressure on getting that done?

**Gane Erickson | President and Chief Executive Officer:**

We have not talked to the level of detail to determine who it is. We've given enough hints that it's amongst the top, suppliers of ai they it's not one of the asic guys so i i'm going to try and avoid being more specific i will restate we are in conversation with every one of the suppliers and i will then say including those guys okay um so my interpretation of that is you know it honestly just sort of warms my heart to see the different people's commitment to the different types of processors. I mean, without going into whether they are or could or might already be a customer or not, one thing about AMDs, and we've used that, again, not as an endorsement to them. We've used them as one of the examples because their MI325 has eight processor chips in addition to, I think, at least that many HBM stacks plus a chip set. in one substrate if there's anyone that ought to be doing wafer level burning they would be amongst them okay um but you know for example you know right now we provide uh opportunities for our customers uh including the likes of those guys to buy our tools for their their burning requirements for qualifications either themselves or to use it at one of the many test houses that have our systems to use our systems for package pump burn-in for the lowest cost alternative to things like system-level test systems that are being used out there, and if the most advanced process would be to do wafer-level burn-in over time. So, you know, I won't comment on anything more than that. Sorry, Larry. No, sir. You know, I think in general, you know, I think good news for the processor market is generally good for us right now.

**Larry Chapina | Analyst:**

The optical IO opportunity, is that going to involve actually new machines instead of upgrading existing machines? Is that transition going to happen here shortly, or do they have more machines that they're going to?

**Gane Erickson | President and Chief Executive Officer:**

The forecast includes both. So more upgrades and more new machines.

**Larry Chapina | Analyst:**

They've got to be running out of machines to upgrade, don't they?

**Gane Erickson | President and Chief Executive Officer:**

Yeah, but there's also a scenario where they also have a bunch of products on the current machines that haven't gone away. And so, you know, it's sort of, you know, while you're upgrading these systems, they're backwards compatible, so you can still use the old wafer packs and everything on them. But nevertheless, it's both. And then the other thing, and it's subtle, and those that don't know it, so we introduced a couple years ago, a front end to the Fox systems that allow you for fully hands-free operation with a wafer pack aligner. So you can come up to that with FOOPS, in this case with 300 millimeter, with both overhead or AGV, automatic ground vehicles, with an E86 compliant port that allows you to not even come and touch the machine. And the wafers can run around on the fab and they can run a burden cycle and then move on and go to the next step of test.

**Larry Chapina | Analyst:**

And you can upgrade them with the automation as well.

**Gane Erickson | President and Chief Executive Officer:**

Exactly. So we took what we actually took their tools that they had bought in the past with our older wafer pack aligners. And they are now upgrading to the new wafer pack aligner. But instead of it being offline, it's integrated with the system. So you know, that's kind of a that's kind of a good way. That's the advanced way of doing it. And particularly when you think about 300 millimeter fabs of like memory, big AI processors, you know, even the silicon photonics, you kind of want to do it. You know, that's the best way of doing it, full automation. But if they want offline, they can do that too with us.

**Larry Chapina | Analyst:**

On this HPF opportunity, is this a different company other than who you've been working with for two, three and a half?

**Gane Erickson | President and Chief Executive Officer:**

What's that? Same company. Just evolving requirements.

**Larry Chapina | Analyst:**

Okay. Yeah. I mean, do you expect anything to break loose on the original Enterprise Flash application, or is this going to continue on?

**Gane Erickson | President and Chief Executive Officer:**

It kind of feels like this is, let's just say trumping it, but that word means something different these days. It feels like this is such an enormous opportunity to the Flash guys. that it's sort of like, you know, the shiny bright light that may actually be better for us. I'm not sure it's better in terms of near term, like, you know, the opportunity is fast. We'll see. But they could, you know, they could configure a system. The new system configuration is a superset of the old requirements. And so we had already worked on the previous one. and we're working on an updated proposal to show them how they could build blades in our system that could do both their old devices and the new ones. So maybe that will help it be better. I think it is, but, you know, it's always interesting when things change. But the one thing, none of their old tools will work with this HP Flash.

**Larry Chapina | Analyst:**

No, I wouldn't think so.

**Gane Erickson | President and Chief Executive Officer:**

So that's, you know, maybe that's a good thing for us, right?

**Larry Chapina | Analyst:**

All right. That's all I had. I'll see you tomorrow, I guess.

**Gane Erickson | President and Chief Executive Officer:**

Thanks, Larry. And Larry's just alluding to, we're going to be over, we're here at Semicon in Arizona, Semicon West, and there's this CEO summit that Chris alluded to. Although, Chris, I don't know if you knew this, you were breaking up. And it sounds like we had operator problems with the operator connection. The new one has been a lot better. So sorry about that to folks that are on the line. Operator, any other questions?

**Operator | Conference Operator:**

I'm showing there are no further questions in queue at this time, and I'd now like to hand the floor back to management for closing remarks.

**Gane Erickson | President and Chief Executive Officer:**

Okay, thank you. You know, I meant to try and work this in. I'm going to do one little other thing. So the other one we haven't talked about, and maybe next call we'll spend a little bit more time on, we did a deep dive last time on the AI side of things. This time was more of an update on things. But there's other products that we have, and one of the things I want to highlight is the activities that we have within package part outside of AI. It turns out that with the NCAL acquisition, they have a low-power and a medium-power system called Echo and Tahoe that we've been shipping a lot of systems kind of quietly in the background. And recently we've had some customers, I think, egged on by some competitors that were saying, oh, there isn't even doing that stuff anymore, and that's just not true. These products are beloved by the customers for their software, their flexibility, and they did a really good job. In fact, those products were the products that honestly took air out of the Pacify burning market because the products were just better than ours. And, you know, we still love those. And if you come on our floor, you'll see them being built right alongside of the Sonoma systems and our Fox systems as well. So just a message out to our customers, we still love you. We're still committed to supporting those products, and we have way more manufacturing capacity than Intel ever did. So don't be timid. We're happy to continue to ship as we have, and we'll give the investors a little bit more insight on some of the systems we're building right now, some of the interesting applications that they're going into. that are also another part of this overall shift of all semiconductors needing more and more reliability tests from qualifications to burn-in. So with that, I thank everybody, and we appreciate your time. And putting up with a little bit of the stuff going on with the call, we'll work on that and make sure we do better next time. And we appreciate it, too. Thank you, now.

**Operator | Conference Operator:**

Goodbye. Thank you. This does conclude today's conference call. You may disconnect your phone lines at this time and have a wonderful day. Thank you once again for your participation.