Grid Dynamics

GDYN Q2 2026 Earnings Call
Recording Transcript

Company Participants
• Leonard Livschitz, Chief Executive Officer and Director
• Vasily Sizov, Chief Revenue Officer
• Yury Gryzlov, Chief Operating Officer
• Eugene Steinberg, Chief Technology Officer
• Anil Doradla, Chief Financial Officer
• Cary Savas, Director, Branding & Communications

Other Participants
• Mayank Tandon, Analyst, Needham
• Bryan Bergin, Analyst, TD Cowen
• Puneet Jain, Analyst, JP Morgan
• Matt Dezort, Analyst, William Blair
• Surinder Thind, Analyst, Jefferies

Cary Savas, Director, Branding & Communications:
Good afternoon, everyone. Welcome to Grid Dynamics' second quarter 2026 earnings conference call. I'm Cary Savas, Director of Branding and Communications. At this time, our participants are in listen-only mode. Joining us on the call today are CEO, Leonard Livschitz; CRO, Vasily Sizov; COO, Yury Gryzlov; CTO, Eugene Steinberg; and CFO, Anil Duradla. Following the prepared remarks we will open the call to your questions.

Please note that today's conference call is being recorded. Before we begin, I would like to remind everyone that today's discussion will contain forward-looking statements. This includes our business in a financial outlook and the answers to some of your questions. Such statements are subject to the risks and uncertainty as described in the company's earnings release and other filings with the SEC. During this call, we will discuss certain non-Gap measures of our performance. Gap to non-Gap financial reconciliations and supplemental financial information are provided in the earnings press release and the 8K filed with the SEC. You can find all the information I just described in the investor relations section of our website.

I now turn the call over to Leonard, our CEO.

####

Leonard Livschitz, Chief Executive Officer:
Thank you, Cary. Good afternoon, everyone, and thank you for joining us today.

We delivered a solid second quarter — consolidated revenue of $108.2 million, above the high end of our guidance range and ahead of Wall Street expectations, with non-GAAP earnings of $14.7 million also which is beating consensus.

As you may recall from my last quarter's commentary, there were three areas I highlighted: first, improving revenue trends, especially with key accounts in the areas of Technology and Financial Services; second, our AI adoption and growth; and third, improving profitability trends. I am happy to

report that on all three fronts, our execution is solid and we are seeing the benefits: growing top account relationships, continued AI momentum with expanded capabilities in robotics and Physical AI, and solid progress toward our 300 basis point margin expansion commitment.

For the second consecutive quarter, our top accounts are in technology and financial services. Technology and financial services now define our most strategic customer relationships, and those are precisely the sectors where AI adoption is moving fastest and where our capabilities are the most differentiated.

Our top accounts continue to drive our growth. Several delivered double-digit quarter-over-quarter growth, with standout performances. There are no incremental gains — they reflect expanding programs, deeper platform adoption, and the compounding effect of our capabilities that keep finding the opportunity inside each client's organization. Several of these clients are now embedding our GAIN platforms as core infrastructure in their own operations — not as a project tool, but as a sustained capability.

AI revenue reached 30.7% of the total company revenue in the second quarter — growing 54.6% year-over-year and crossing the 30% threshold for the first time. Two consecutive quarters of year-over-year growth above 50% tells us something important: this is not a spike, it is a sustained shift. The trajectory is clear, and we intend to build on it.

Driving this strong performance is a combination of multiple factors. Our GAIN platforms are winning wider enterprise adoption. Our clients continue to transition enterprise AI workloads from pilots to production. Our engineers are more deeply embedded inside client organizations. Bottom line, we are winning entirely new programs. That gives us confidence in the growth ahead.

AI-first Delivery is now the default, not the aspiration. Fixed-price is the preferred approach on new RFP responses. The productivity and margin gains are real. We are executing well and delivering projects successfully. Our focus on executing larger AI platforms is aligned with significant progress we are making in upskilling our engineering talent. By the end of October, we plan to have 90% of our engineers trained on AI-SDLC.

Our GAIN platforms have expanded LLM partnerships meaningfully this quarter. We are now working with several of the world's leading AI companies, including the top four frontier providers with whom we are under commercial agreements. This approach ensures our GAIN platforms stay aligned with the leading AI platforms, with broader reach across our enterprise client base. GAIN remains the backbone through which we bring AI capabilities to market. Its partner depth makes it stronger every quarter.

Our client relationships are evolving too. Clients who came to us for platform deployments now ask us to stay. They want us to be involved in advisory, execution, ongoing operations. This meaningful shift is opening a growth vector that didn't exist in our model two years ago.

On the partnership front, partner-influenced revenue reached 19.1% of total company revenue in the second quarter. This was driven primarily by our three core hyperscaler relationships with Google Cloud, AWS, and Microsoft Azure. A growing proportion of that revenue is coming from AI engagements.

We are running Agentic AI workshops across our Google Vertex AI Search customer base, converting search engagements into broader agentic commerce programs. We extended our Google partnership into Banking and Financial Services, closing our first joint win this quarter at a leading global bank. And we are deepening our AWS partnership around application modernization and Agentic AI in CPG, Manufacturing, and Financial Services.

Our NVIDIA partnership is gaining momentum across both Agentic AI and Physical AI. Our longer-term target remains 25 to 30 percent partner-influenced revenue, and we are confident of achieving this target.

Last quarter I introduced our Physical AI capabilities and our first commercial engagements in this space. Physical AI requires a deep understanding of multiple disciplines that include modeling

real-world robotic movements, digital-twins, verification in simulators, and integration with hardware systems.

Our active programs span humanoid robotics for pharmaceutical intralogistics, autonomous driving stacks for construction equipment, and policy control platforms for manufacturing clients. We signed a strategic partnership with Doosan, a leading robotics manufacturer this quarter, elevated our NVIDIA relationship, and opened an engineering office in Dresden Germany to support our European manufacturing clients.

Grid Dynamics enhanced its robotics offering by welcoming Ekumen, a leading robotics engineering team that joined us in May. Their expertise resides in the Robot Operating System, a foundational open-source standard that powers the vast majority of the world's industrial robots. Over the past decade, the company has built an enviable list of some of the world's most respected robotics companies.

Grid Dynamics brings advanced AI modeling, policy control, and enterprise-scale delivery capability. Ekumen brings deep knowledge of the foundational software layer that robot manufacturers depend on. Together, the combination is formidable — spanning the full stack from the foundational software layer through simulation, hardware integration, and enterprise-scale deployment. We believe no other services company in the market today matches this combined footprint and technical depth.

Now. let me pass on to Vasily Sizov, our Chief Revenue Officer, who will expand on key business aspects of Grid Dynamics client engagements.

###

Vasily Sizov, Chief Revenue Officer
Thank you, Leonard.

Let me begin with three demand trends we observed during the quarter. First, clients are prioritizing AI investments that deliver clear, measurable business outcomes. Second, as clients move from isolated use cases to enterprise-scale initiatives, they realize that the underlying technology layers must be modernized to support AI adoption. Third, clients increasingly recognize that successful AI transformation requires more than technology alone, driving interest in AI process consulting, performance benchmarking, and change management. These trends align closely with our strategy and the capabilities we are building. Let me discuss each of them in more detail.

First, the demand environment remains constructive, with clients directing AI investments toward practical applications with tangible business impact. We are seeing particular interest in AI-enabled automation that improves operating efficiency, scalability, and speed. Importantly, these investments are increasingly moving beyond experimentation, with clients deploying AI capabilities into production to automate complex manual processes, improve customer service, reduce operating costs, and create new sources of revenue.

Second, as clients move from isolated AI use cases toward enterprise-scale transformation, they are finding that their data, application, and core platforms must be modernized and made AI-ready. As a result, AI adoption is creating broader demand across the underlying technology landscape. This trend aligns closely with our core expertise in data engineering, application modernization, cloud and platform engineering, and reinforces the relevance of these capabilities in the AI era.

Third, we are seeing growing demand for AI process consulting, performance benchmarking, and change management as clients focus on converting AI investments into measurable business value. They need to identify the business processes where AI reengineering can create the greatest value, establish clear performance baselines, redesign those processes, build the technical enablers, and

drive enterprise-wide adoption. We have been deliberately strengthening these capabilities to help clients realize measurable value from AI across the enterprise.

These trends are reflected in our client work. Let me highlight a few engagements from the quarter that demonstrate how these capabilities are being applied in practice.

For a leading foodservice distribution company, we built and deployed an AI-powered product credit claims platform that automatically validates customer claims against photographic evidence. The platform cross-checks product, manufacturer-label, and shipping-label images against the claim's reason code in real time, replacing a fully manual, salesperson-mediated review process. In performance testing, the system processed approximately 400 claims, supported by 1,000 images, end-to-end in under 15 seconds per claim. The capability is now live in production, and the client has approved a long-term roadmap to further enhance the system and extend automated decision-making into more advanced credit-adjudication scenarios.

For a leading home improvement retailer, Grid Dynamics enabled next-day delivery by designing and deploying a high-load service that modernized the retailer's logistics operations. The solution includes an AI-powered routing capability that assigns fragile items to the appropriate vehicle types, eliminating hundreds of delivery errors each week. As a result, the solution cut average delivery time by more than half from 3.5 days and is expected to support up to half a billion dollars in incremental annual revenue for the client.

For a global technology company, we modernized large-scale data processing infrastructure, migrating more than 1,000 data pipelines to a serverless execution model. This reduced idle compute capacity, reduced infrastructure costs, and improved scalability. Our proprietary AI-powered automation accelerated the migration and established a reusable delivery approach that is now being applied across broader initiatives at this client.

Now let me turn the call to Yury Gryzlov, our Chief Operating Officer.

###

Yury Gryzlov, Chief Operating Officer
Thank you, Vasily

Let me build on the Physical AI and Robotics work Leonard introduced.

Physical AI needs a full technology stack, and we operate across everything between the robot and the enterprise — the devices themselves come from our hardware partners.

At the foundation is the robot operating system — ROS and ROS 2 — the open-source layer the majority of the world's modern robots are built on, connecting the hardware to everything above it. Through Ekumen, we're not just users of it; we're among its maintainers and a founding member of the alliance that governs it. On top of that sits the intelligence — the AI models that let a robot perceive its surroundings, generate its own motion, and handle real-world variability. We design and validate that in simulation, before it ever runs on a real robot. And our own platform, Incarno — our GAIN Platform for Physical AI — is where enterprises bring it all together: building manipulation and inspection workflows, deploying those models, and monitoring robotic lines with digital twins.
What unifies it is our focus on the enterprise — expanding this capability to the companies that have robots deployed at scale. Here are a few examples that illustrate our work across the stack.

For a leading manufacturer of construction and mining equipment, we're building a next-generation stack for autonomous driving, loading, and excavation. We're helping them design the platform, onboard the first use cases, and add capabilities like policy-based control. What began as our first commercial Physical AI engagement, is now a multi-year program across several regions.
With Ekumen, we've proven two-arm manipulation — grasping and assembly — trained entirely in simulation and then run reliably on the real robot. Bridging that gap, from simulation to the physical robot, is one of the hardest problems in the field.

Humanoids are the next step: a leading life-sciences company is piloting humanoid robots for intralogistics — moving and repacking containers of chemicals — work that was out of reach only a couple of years ago, and is now possible thanks to new AI models that generate motion. We provide the platform those robots run on, working with Wandelbots and on NVIDIA's stack. The customer calls it a lighthouse project for their industry, and it's the opening step in a much wider program.

We're also building the channels to scale. This quarter we announced a strategic partnership with Doosan Robotics, a global leader in collaborative robots deployed across 45 countries. It's a full-stack collaboration — our platform plus the foundational AI components, integration services, and engineering around it, paired with Doosan's cobots and our combined global reach. Together we can provide what traditional robotics software can't — dual-arm assembly, inspection of complex-geometry parts, and packing of deformable items. It sits alongside our elevated NVIDIA partnership, and we're in active talks with several more hardware and software vendors.

Considering the economics of software services in this space and our positioning, we are confident that we have a material market advantage. Reliable performance in the physical world takes engineers who understand simulation, control, and hardware variability — working through problems that have no templated solution and so can't be easily automated. This combination is hard to assemble: Ekumen's decade of foundational robotics depth, together with our strengths in AI modeling, simulation, and enterprise delivery. We don't believe another services company matches it today. Closing that gap isn't a matter of hiring a team; it's years of hard-won experience, which we're now putting to work for our customers.

In summary, robotics and physical AI is a growing market — measured in the trillions over the coming decade. Our expanded capability is helping us capitalise on the early traction we saw last year, reflected in a rapidly growing pipeline from both existing customers and new logos.
Another important part of my update is tied to our Capital Markets focus — where a similar pattern is playing out, in software rather than robots. As our banking clients push agentic AI deep into their engineering, the hard part is no longer producing code. It's doing it safely — with quality, security, and control they can provide to a regulator.

This quarter that showed up most sharply around security: banks want the speed of frontier models and AI-generated code without introducing new vulnerabilities. Our answer is spec-driven agentic engineering, led by Allium — part of our GAIN platform for AI SDLC — and it's exhibiting real traction across our banking clients.

The clearest example is at one of the world's largest banks, where Allium is being used to build new tools as part of a bank-wide initiative to modernise business operations. Working across London, New York, and India, we're bringing specification-driven development to both new and existing systems — starting with tools for AI-assisted productivity, and extending to agents that automate operational work.

Taken together — physical AI reaching the enterprise, and AI-native engineering scaling inside the world's largest banks — this is the frontier work that keeps Grid Dynamics differentiated.
Over to you, Eugene.

Eugene Steinberg, Chief Technology Officer
Thank you, Yury. Good afternoon.

Last year I described our AI strategy through three horizons. This quarter, I'll describe them by maturity: what has reached scale, and what is beginning to scale.

Horizon one, scaled: AI-first modernization and the agentic platform

Modernization remains the foundation of our business, but AI is changing how that work gets done. Agents can now accelerate work across most of the modernization lifecycle, particularly code generation and testing. The remaining work—business acceptance, production scaling, and complex coordination—still depends on human judgment and accountability. An agent can write the code; a person still makes the call and stands behind it.

We have invested in a set of GAIN tools that support this lifecycle. Rosetta governs how agents operate. Allium analyzes legacy systems to create reliable specifications for their replacements. And SpecFlow, our latest open-source contribution, uses those specifications to support autonomous feature implementation.

Rosetta has progressed from its first lighthouse clients to larger engagements across retail, financial services, and manufacturing. At a Fortune 30 U.S. home-improvement retailer, approximately 550 of the client's engineers are working with the platform. In one program, seven COBOL services were moved to a modern technology stack, with approximately 90% of the code generated by agents. All seven services entered production this quarter.

The client already had capable engineers and access to many of the same AI tools we use. What it needed from us was domain knowledge, governance, and control—the capabilities that turn powerful agents into dependable enterprise systems.

This productivity is helping us expand client relationships. It is also creating opportunities to use more fixed-price and outcome-based commercial models when the scope and accountability are clearly defined.

Allium also reached an important milestone this quarter. It is being piloted across five major banks and has begun moving into its first commercial banking engagements. Allium analyzes legacy code to help establish reliable functional specifications for replacement systems. It also supports controlled migration and rollback, reducing the operational risk of moving critical applications onto modern platforms.

At one major North American bank, a one-hour GAIN demonstration in February led to a signed contract in April. The bank was managing 150 applications with limited test coverage and a growing security backlog. We translated identified issues into failing tests inside the bank's own tooling, allowing its engineers to independently reproduce and assess each finding.

At another Tier 1 bank, the same approach is supporting a security-modernization program spanning more than 100,000 systems, with part of the modernization work co-funded by the client's cloud provider.

Our differentiation is not limited to code generation. Our agents can also incorporate context such as security advisories, dependencies, and upstream changes. That broader context helps identify

problems that code-only tools can miss and provides the traceability and evidence regulated enterprises expect.

We deliberately make selected GAIN platforms open source. The immediate objective is adoption and technical credibility, not software-license revenue. Open code allows engineering leaders to evaluate our capabilities directly and strengthens our position when clients need help deploying those capabilities at enterprise scale.

The same pattern applies to data. Enterprise AI cannot deliver reliable results without accessible, well-governed data. That is increasing demand for data-platform modernization. Our new AI data-migration accelerator, released this quarter, is already being deployed in a data-lake modernization program for a global consumer-products manufacturer.

The second scaled component of Horizon One is GAIN Agentic Runtime. Enterprise agents need access to trusted data, evidence that their behavior is controlled, and governance over operating cost. For a global payments client, we brought these capabilities together as shared services, with the retrieval layer now supporting 25 enterprise consumers.

We also converted the client's disputes architecture—including fraud, chargebacks, and KYC—to configuration-driven workflows. A common foundation now supports four use cases. By automating much of that configuration, the program rebuilt a decade of business logic in six months and reduced integration and release-cycle times by 96%.

At our largest banking client, an internal platform built with our support now centralizes the registration, governance, and operation of AI agents across the organization. The client reports regular adoption by more than 80% of its employees across more than 80 markets. As adoption grows, we are also developing the operational tooling needed to govern and support the platform at that scale. Across these engagements, the pattern is consistent: AI accelerates production, but enterprise value comes from the domain knowledge, governance, and accountability required to put it into production responsibly.

Horizon two, scaling: harness engineering and physical AI

Horizon Two covers capabilities that are moving from research and internal validation toward repeatable client deployment.

The first is agentic harness engineering. Traditional agentic workflows are most effective when the task and sequence of steps are already known. Harnesses are designed for more dynamic work—situations in which an agent must select tools, adjust its approach, and respond to new information while remaining within defined controls.

The harness provides those controls. It records what the agent did, tests its output, manages exceptions, and introduces human review where accountability requires it. This allows enterprises to apply agents to more complex work without giving up oversight.

During the second quarter, our India engineering center developed nine harness-based solutions. Following our Client Zero approach, we are testing them first within our own operations. The objective is to establish evidence of reliability, define the necessary controls, and improve the solutions before introducing them into client environments.

The second area is Physical AI and robotics. We are investing here because the engineering challenge is fundamentally different from conventional software development.

A coding agent can generate software and test it in a digital environment. A physical system must also operate safely and reliably in the real world. It must account for geometry, motion, changing conditions, and the behavior of the physical environment. Validation therefore has to take place both in simulation and on hardware. A language model alone cannot close this loop.

Our research is focused on bringing physics, geometry, simulation, and continuous validation into the agent's operating environment. That is also the strategic rationale for the robotics engineering team we acquired in May. Members of this team have long contributed to core infrastructure in the robotics operating-system ecosystem, with particular expertise in simulation and validation.

Their capabilities are now contributing to GAIN for Physical AI, our platform built on Incarno. During the quarter, we released three new components: tools for composing robot policies, a continuous-improvement loop, and sandbox environments for controlled testing.

We are beginning to validate the platform through early client and partner deployments. A leading life science company is piloting humanoid robots in its warehouse operations using our platform. Separately, a robotics partner has incorporated the platform into its own offering, creating a distribution channel for our Physical AI technology.

Horizon Two is not yet at the same maturity as our modernization and agentic-platform business. Our focus now is to demonstrate repeatability, convert technical validation into production deployments, and establish scalable commercial models. The opportunity is to build differentiated intellectual property in areas where success requires not only generating software, but proving how that software behaves in the physical world.

Across both horizons, the pattern is clear: the cost of producing software is falling, while the value of governing it, validating it, and taking responsibility for it in production is increasing.

This quarter, more components of GAIN moved from tools and pilots into broader enterprise adoption. At the same time, our investments in agentic harnesses and Physical AI progressed from research toward controlled client deployment.

As these capabilities mature, they allow us to reuse more of our engineering, deploy solutions faster, and take greater responsibility for measurable outcomes. Our advantage is not simply that our agents can generate code. It is that we combine those agents with domain knowledge, operational controls, and the engineering discipline required to make them dependable at enterprise scale.

That is where we believe durable value will be created in the agentic era—and where Grid Dynamics is positioned to lead.

Anil, over to you.

###

Anil Doradla, Chief Financial Officer:
Thanks, Eugene. Good afternoon everyone.

Second quarter came in at $108.2 million, slightly above the higher end of our guidance range of $106.0 million to $108.0 million. That represents 7.0% year-over-year growth, including de-minimis contribution from Ekumen.

Grid Dynamics

Non-GAAP EBITDA was $14.7 million, or 13.6% of revenues and was closer to the higher end of our $14.0 million to $15.0 million guidance range.

Looking at performance of our verticals:

TMT remained our largest vertical and accounted for 31.8% of total revenues for the quarter, with a growth of 11.7% sequentially and 36.4% on a year-over-year basis. The growth was primarily driven by our largest technology customers. We continue to benefit from vendor consolidation at these customers, which has driven increased wallet share across new and existing programs.

Retail contributed 26.5% of total revenues in the second quarter of 2026. The vertical was flat in absolute dollars on a year-over-year basis and grew 3.1% sequentially. The sequential growth was supported by demand from key accounts, including a major specialty retailer.

Our Finance vertical accounted for 22.9% of total revenues in the quarter and grew 1.2% on a sequential basis. Within this vertical, we witnessed solid demand from our fintech service engagements, including increased contributions from a major payments network, which helped offset the successful completion of engagements with insurance and data analytics and consumer credit reporting clients in North America. Looking ahead to the remainder of 2026, we remain bullish on our growth outlook within this vertical.

CPG & Manufacturing represented 10.9% of quarterly revenues and grew 2.1% on sequential basis and 4.2% on a year-over-year basis. Within this vertical, we are witnessing robust demand from a leading wholesale food distributor, along with growth from some of our manufacturing customers.

Turning to our remaining verticals:

Our Other vertical contributed 6.0% of second-quarter revenues, while Healthcare and Pharma accounted for 1.9% of our revenues for the quarter.

We ended the second quarter with a total headcount of 4,838 down from 4,964 employees in the first quarter of 2026 and from 5,013 in the second quarter of 2025. We continue to rationalize our overall headcount as we align our skill sets and geographic mix.

At the end of the second quarter of 2026, our total US headcount was 379, or 7.8% of the company's total headcount versus 7.2% in the year ago quarter. Our non-US headcount, located in Europe, Americas and India was 4,459, or 92.2%.
:
In the second quarter, Revenues from our top 5 and top 10 customers were 43.5% and 61.5%, respectively, versus 37.5% and 57.3% in the same period a year ago, respectively.

Moving to the income statement, our GAAP gross profit during the quarter was $39.6 million, or 36.6%, compared to $36.2 million or 34.8% in the first quarter of 2026 and $34.5 million, or 34.1% in the year ago quarter. On a non-GAAP basis, our gross profit was $40.0 million or 36.9% compared to $36.7 million or 35.3% in the first quarter of 2026 and $35.1 million or 34.7% in the year ago quarter. On a year-over-year basis, the increase in the gross margin percentage was primarily driven by revenue

growth outpacing delivery cost. On a sequential basis, the increase in gross margin percentage was due to a combination of working time and improved resource utilization.

Non-GAAP EBITDA during the second quarter that excluded interest income/expense, provision for income taxes, depreciation and amortization, stock-based compensation, restructuring, expenses related to the geographic reorganization, and transaction and other related costs was $14.7 million, or 13.6% of revenues versus $12.5 million or 12.0% of revenues in the first quarter of 2026, and was up from $12.7 million or 12.6%, in the year ago quarter. The sequential and year-over-year growth in EBITDA was largely due to a combination of higher revenues and strong operating leverage across our non-engineering overhead.

Our GAAP net income in the second quarter was $2.9 million or $0.03 per share, based on a diluted share count of 83.0 million shares, compared to the first quarter net loss of $1.5 million or a loss of $0.02 per share, based on a diluted share count of 84.7 million and net income of $5.3 million or $0.06 per share based on 86.4 million diluted shares in the year ago quarter.

On a non-GAAP basis, in the second quarter our non-GAAP net income was $9.0 million, or $0.11 per share based on 83.0 million diluted shares, compared to the first quarter non-GAAP net income of $7.5 million or $0.09 per share based on 85.9 million diluted shares and $8.3 million or $0.10 per share based on 86.4 million diluted shares in the year ago quarter.

On June 30, 2026, our Cash and cash equivalents totaled $298.4 million, down from $327.5 million on March 31, 2026.

Since our first quarter earnings call, we repurchased approximately 2.6 million shares for a total consideration of $17.3 million. Cumulatively, since our Board authorized the $50.0 million share repurchase program, we have repurchased approximately 4.4 million shares for a total of $30.8 million, reflecting our continued confidence in the long-term value of the business.

Coming to the third quarter guidance, we expect revenues to be in the range of $112.0 to $114.0 million. We expect our third quarter non-GAAP EBITDA to be in the range of $16.5 million to $17.5 million. For Q3 2026, we expect our basic share count to be in the range of 81.0 to 82.0 million and our diluted share count to be in the range of 83.0 to 84.0 million.

For 2026, we are maintaining our full year revenue outlook of $435.0 million to $465.0 million.
That concludes my prepared comments. We are ready to take questions. Cary ...

###

Questions and Answers

Cary Savas, Director, Branding & Communications:
Thank you, Anil. As we go into the Q&A session of this call, I will first announce your name. At that point, please unmute yourself and turn on your camera.

The first question today comes from Mayank Tandon of Needham. Go ahead, Mayank.

Mayank Tandon, Analyst, Needham & Company:
Great. Thank you. Congrats on the quarter, Leonard and Anil.

Anil Doradla, Chief Financial Officer:
Thank you, Mayank.

Mayank Tandon, Analyst, Needham & Company:
So there was a lot of detail around AI. So just to step back, Leonard, could you maybe talk about the AI efforts and the implications for both growth and profitability over the next, say, 12, 24 months? Maybe you can help reassure investors that AI will actually be a net positive for you because there is still a lot of skeptics out there that think it's going to be a net negative over time.

Leonard Livschitz, Chief Executive Officer:
Right! Thank you, Mayank. Well, it's a pretty comprehensive question, and if I answer all of the parts, there'll be probably nothing left for the other end of this. So I'll try to be concise in terms of the key elements, and then we can talk a little bit more in detail.

So first of all, we are reaching many aspects of AI implementations. We talked about it in the past. We're adding those features now. We're talking about... directly or indirectly about forward-deployed engineers. We make announcements, we train a substantial number of the people in the workforce, and these people are basically driving a new way of implementing our solutions because as we tend to get more focused on fixed bid and fixed budget projects, it helps us to identify not only the execution of the various modernization projects, but also create a technology consultative. So that's with respect of the people and why it's accretive to us.

When it comes to Agentic AI and a part of the... again, implementation of the suite of our solutions, we are driving our customers to adapt our GAIN platform model. All the elements of the model are driven by internal testing and developments but also tailored to our customers' needs. So they will need to adapt the solution where they see the most fee for themselves. But also, we guide them through the process to create the best ROI for that. So that's the second part. And before I talk about the physical AI, I want to bring the interest to your point in terms of net positive versus net negative.

If you look at the increased growth in just these two areas, that substantially exceeds the sum of the aged businesses, which would eventually drop out because the gloom and doom from many facets were about that engineering and consultancy is less relevant. Moreover, people would say it's easier to train FTEs. We embrace FTEs, we embrace our clients. At the same time, as many of the leaders in industry are saying, we can do more work. We can do more engagements, which we prove with all the listed examples. And I'm not going to go through all of that because we have a lot of people who can give you more details. So, as a consolidated effort, as we go today through further discussions, we demonstrate on specific examples where this accretiveness works. But I want to emphasize forward-deployed engineers and Agentic AI. And the third part, which is also super critical for us, which actually drives the adoption and partnership enhancement of our relationship to the next level is actually our preparation for our physical AI work.

We've not just made a small acquisition. We've not just made announcements about opening additional robotics labs. We're working... we've been working with our clients for a long enough time to understand what it means for their own platform. What it means for their application and solutions for various other verticals. From industrial, from modern machinery to logistics companies, to even work in the industrialization of various new solutions. So, the material side, the remuneration for physical AI is still to come. But now we have evidence of substantially more players looking at Grid Dynamics again in a leadership role by expanding our capabilities to the practical world of their usage. So this is pretty much a summary, and of course, we'll go in more detail.

Mayank Tandon, Analyst, Needham & Company:
That's very helpful. I'm sorry if you can't see me. I'm having an issue with my video. Maybe you can help with that. So I'll probably get that fixed eventually. And then just a very quick follow-up, Anil, for you, in terms of the guide, just want to get a sense of the visibility that you have today versus last quarter. And what I mean by that is, is the pipeline now converting faster? Have you seen evidence of that? Does that maybe give you more confidence in the sustainability of growth acceleration once we get beyond fiscal 26 into fiscal 27?

Anil Doradla, Chief Financial Officer:
Oh, you're talking about next year. Well, let's talk about this year, and then we'll get to next year. When you go into the second half, Mayank, you see, if you look at our visibility and our second half, there are a couple of factors. Number one: Remember the 85 10 5? Most of our revenue comes from customers who've been with us for two years and beyond. That formula more or less stays well intact. And that you're seeing in the top 5, top 10 customers, right? Because most of the absolute dollar and year-over-year growth is coming there so that stays intact. As you go into the second half, there are three layers as you go. First is the working time, the second half is higher than the first half.

The second thing is the billable headcount. So we're seeing new programs kicking in. So, maybe without addressing your pipeline question directly, indirectly is that yes, we're seeing an increased billable headcount as we go into the second half. And the third thing is that, you know, we are planning some acquisitions. So all these three add up to layers.

Now, when you look into 2027, I think I'll let the business guys chime in here, but from my point of view, I see two things that are very interesting. Number one, the relationships that we're having with our technology customers, our financial customers, our top 10 and 20 customers, is going deeper and deeper. Things that we've not done, we're doing. Application modernization programs which we've not done, we're addressing. The addressable market that we're going after is larger. And I overhear these conversations week after week, which leads me to believe, as you go into 2027, if we continue winning at the rate that we're winning, it should play out incrementally positive. But I don't know, Vasily or Yury, whether you want to add anything to that.

Vasily Sizov, Chief Revenue Officer:
Yeah, let me chime in. So yes, I would say that our position with most of our biggest clients is...has been strengthening over the last few years by... through vendor consolidation. And what we see is that we should benefit in the coming years from this consolidation, which means bigger programs would come our way by customers cutting loose their long tail of vendors, which are no longer relevant. And given our strong technology positioning in Agentic AI, which is a very, very hot topic for most of our customers, we are really well positioned to benefit from that.

Mayank Tandon, Analyst, Needham & Company:
Terrific. Thank you so much. Congrats.

Cary Savas, Director, Branding & Communications:
Thank you, Mayank. The next questions come from Bryan Bergin, TD Securities. Go ahead, Bryan.

Bryan Bergin, Analyst, TD Cowen:
Hey y'all, good afternoon. Thanks for taking the question here. So, maybe just to start a follow-up on that last question as it relates to kind of that second half, more of a near-term question, just as it relates to... you give us a 3Q guide, implied 4Q is still a decent ramp. Are you seeing a broadening of momentum in other sectors? Obviously, you're doing quite well in technology. Are you seeing a broadening of momentum elsewhere that gives you that confidence? And then, as it relates to kind of, you know, potentially some M&A requirements, any way you can share with us how you're thinking about maybe the organic contribution remaining versus any needed M&A that you have to go get?

Anil Doradla, Chief Financial Officer:
Right. So Bryan, let me point out that, you know, as you know, there's a certain seasonality in our business, right? As we go into Q3 Q4, that's well established. As I said, there are three levels at which we're operating. Number one is just the working times of the second half of the year, and you guys know it's better. The second thing is that the billable headcount and the trends are positive, and all our prepared commentary should lead you to include that. And the third thing is that there is a certain amount of acquisition, and we do have a pipeline. It varies, you know... I always joke, right? An acquisition is not done till the money is transferred to their bank, right? We've seen acquisitions that we thought are not going to happen, they happened. We've seen acquisitions that were locked and loaded, and we just are not able to.

So, if you look at that second half, I don't want to comment too much about Q4, other than saying that, look, we have a seasonal pattern for the year, but as we go from the low end of our full year guide to the high end of the guide, the first component of working time stays intact. The second component of billable headcount, we have variable calculations, and the third component perhaps, picks up a little bit more, is the acquisitions.

Bryan Bergin, Analyst, TD Cowen:
Okay, okay, understood. My follow-up is kind of a margin and a tie-in with the delivery model question. So you reiterated the confidence in the 300 basis point expansion. So that's good to hear. I'm just curious how much of this margin improvement is coming from structural changes, you know, automation and efficiencies in the delivery versus traditional kind of cost control/cutting measures. And I think it's notable yet, like, 7% revenue growth while headcount was down 3. Is there a lasting change? I know you're saying you're going to add billable headcount, but is there a lasting change in this delivery model? So just maybe talk about that AI-driven efficiency and delivery that you're seeing.

Anil Doradla, Chief Financial Officer:
Right. So there are 3-4 parts of the question. Let me take the first part. And then when it comes to some of the AI trends, I'll pass it on. But, when you look at what we set out to do, we said that on a year-over-year we're going to deliver 300 bps margins on a Q4 by Q4 on a year-over-year basis.

Part of that effort is efficiency—it's just the way we're organized. As you know, we've ramped from a handful of countries to 19 countries. We've got many incorporated entities. So there's a little bit of a efficiency that we brought in, and some of those are one-time. But we operate at a certain level, right? The second part that we are seeing here is we're embracing a little bit more change in the way we're doing business, whether it's AI, whether it's fixed price, whether it's a greater...embracing more tools. And that is creating a certain level of, I would say, it's not so visible now, but over time you'll see a nonlinearity, perhaps that you'll see.

The movement that you've seen on the headcount right now was largely driven by efficiency improvements on non-engineering headcounts. So people should not worry. It's not that we let go of billable headcount. No, it's just non-engineering, non-billable headcount. We cleaned it up. But from this point onwards, beyond the 300 basis points that you'll have from Q4 to Q4, as you go into 2027, there is a plan for us to leverage more of these tools. There is a plan of bringing a certain level of nonlinearity. We have the plans. The clients have to accept it, and we have to proceed with that.

Leonard Livschitz, Chief Executive Officer:
Yeah, let me do a couple of things. First of all, just a complete answer on the very first question of yours about diversification of the platforms. I think it's very critical to understand that this is not overnight that we started diversifying the verticals. First and foremost, we've been in the payment system, financial services, industrial modernization. The second of all is, what you can actually see from the previous comments about us, what Vasily said, replacing some incumbent vendors is because we're playing in a big boys league, in a higher level. I mean, in the past, look there were always couple top guys and a couple mid level vendors. Now we only compete with the top guys. And the reason being is, I think AI adoption and technology implementation equalized the field a bit.

So we've always been prepared for the big tasks and a big program, big transformational solutions. But we've also gained a reputation of this consultancy part. So as we get more admittance to the bigger projects, inevitably what happens with it is better visibility, better projection, better position. So we're saving with the tools, we're adding more capabilities. And we're looking back. Then we say, what are these internal systems which we had for a long time, are less efficient.

We accepted to lead in a world of uncertainty. That's very important. We don't see the world changing so dramatically. We'll go back immediately to the luxury of being very consolidated in very few locations. We're adding not only India, and we're adding investment into India and the technology capability, but also LATAM. And we... as we do more, we create a global platform internally to optimize this efficiency. So it's a cost structure, it's performance-based, it's tooling, it's removing redundancies from the past. I hope Bryan, I covered a lot.

Cary Savas, Director, Branding & Communications:
Thank you, Bryan.
The next question comes from Puneet Jain of JP Morgan.

Puneet Jain, Analyst, JP Morgan:
Hey, thanks for taking my question. So, how are your AI and robotics partnerships different from your traditional hyperscaler relationships, like with Google, AWS, Microsoft Azure that generate much of your 19% of partnership revenue? So the partnerships you got with NVIDIA, model companies, do they differ,

or do they offer, like a different revenue trajectory potential or client ownership structure than your other partnerships?

Vasily Sizov, Chief Revenue Officer:
All right. Thank you so much for the question, Puneet. Let me address this question.
So we definitely value our relationships with NVIDIA and believe that's a great partnership to build a pipeline of future opportunities on. As you understand right now, the industry, the manufacturing is going through a massive transformation, and new tools like Agentic AR, physical AI definitely brings new technology to more kind of traditional manufacturing.

And we see this as a great opportunity to build a new pipeline of opportunities and new type of engagements which would help us to transform those manufacturers on a bigger scale. So just an example. So, for example, right now we have an active engagement with one of the world's largest industrial equipment manufacturer on building an Agentic AI platform which allows to manage a fleet of autonomous vehicles and deploy physical AI capabilities on the edge devices. And we see more and more interest to such opportunities. So it's definitely one of the top priorities for us to grow.

Puneet Jain, Analyst, JP Morgan:
Got it. And I'd like to follow up on the prior question, specifically around headcount. I noticed, like, your non-U.S. headcount was down despite, like, Ekumen, which probably contributed employees in Argentina. And the US headcount by comparison was up on a sequential basis. So should we expect this remix to continue, like, as you do more AI-based services? Will that require more on-site headcount or U.S. headcount compared to in the past? And if that's true, what does that mean for margin and change management within your employee base?

Leonard Livschitz, Chief Executive Officer:
Very good. Well, Mayank, Puneet, what you said, it's music to Eugene's ear because he'd been the one who is architecting the acceleration of some of the US-based presence, both from the technology office perspective, but also from the consultancy, technology consultancy with the clients. I'm not saying there's more shift toward onshoring as a trend. I think if you look back pre-COVID days, our onshore presence between onshore technology people, as well as some of the offshoring engineers who would come on the long-term projects reached almost close to 20%. It's never been so low. And then, when the onshoring presence pulled back due to the inability to work directly with the clients, a lot of work has been going on offshore.

Now, we're not saying that work is no longer relevant. There is more and more demand to be present on-premises with the clients to work together on these complex cases, because the rapid change of transformation sometimes catches the clients a little bit through uncertainty, right? We're talking about two basic approaches to their mental and budgetary resolution of the projects. One of them is more like a status quo. Let's see and tell what's going to happen. Obviously, they don't need as much of onshoring presence, and some of them demand very rapid acceleration. But they're concerned with some of the spendings. As you know, around tokens and other things, which definitely create the pressure. So that's where our headcount onshoring technology-wise is coming.

As I mentioned to Bryan, some of the reduction of offshoring headcounts comes from non engineering and non, I would say forward looking specialties. So there is a difference between the headcount and the contribution of this headcount. So, from the budget perspective, it's a little bit less clear that these

people were extremely expensive, but just the infrastructure of these people would no longer be needed for us to serve the markets better.

So to answer your question, we do see some additional growth on onshoring. The ability of us to prove that our margin expansion will continue to grow is vastly driven by how much of the fixed bid fixed budget projects we can adopt, how much of our internally developed tools are accepted by the clients, how much of the nonlinear value we're bringing to the party. And I think we're quite growing with those elements. So just to conclude on that from my side, and if people want to add. I think you picked the right trend. I don't think the legacy... some of the people are a sign for concern because, you know, the majority of them come from Central/Eastern Europe. And I think this is all by the book. So we are really moving forward with a clear plan on continuing to have margin improvement.

Puneet Jain, Analyst, JP Morgan:
Got it.
Thank you.

Leonard Livschitz, Chief Executive Officer:
Thank you, Puneet.

Cary Savas, Director, Branding & Communications:
Thanks, Puneet. The next question comes from Matt Dezort of William Blair. Go ahead, Matt.

Matt Dezort, Analyst, William Blair:
All right! Thanks, guys. Congrats on the results. I wanted to see if you could double-click on this new consultancy practice that you're talking about. Can you discuss more of how you see this business developing? I know you talked about activities like change management, but what sort of opportunities are you seeing in the pipeline build there? Who are you going up against in these bake-offs, and how is the competitive environment different from your traditional work, maybe?

Leonard Livschitz, Chief Executive Officer:
Yeah, so I will start very briefly, and then with Vasily will actually expand on it. So there are two parts of it. The first part is, there's no change of our purpose. Consultancy has always been a part of our DNA. Nothing is like, you know, earth-shattering because our clients consider us to be technology consultants. And that's why we're able to compete against the big firms. What has changed is the distribution of that kind of offering. And just the previous question with Puneet was about, you know, onshoring presence, right? And the people who we hire—they're extremely technical, but they're also customer-oriented, and that kind of work is very important because we are expanding the purpose of consultancy from pure technology consultancies, and now adding AI infrastructure consultancy, hardness selection consultancy, tool selection consultancy. And to some extent, getting more into the sacred world of business consultanties. So listen.

Vasily Sizov, Chief Revenue Officer:
Yes, so think about business consultancy as a natural extension of our technology enabler build-out capabilities. Essentially, the focus of the customers is shifting from just the creation of a system, but for creation of systems to change the business processes they have. Therefore, they would like to analyze first which business processes are the best candidates to improve which value is hidden there,

then to build a technical enabler to reveal this value, and then adopt that technical enabler on enterprise-wide scale. And that's exactly where the focus of our consultancy is. Not only to create a technical enabler, but also to help get all the value on the enterprise scale from this... from this change. So that's the essence. And right now, we have several active engagements on that. Specifically, on the front of consulting change management, which goes along with technical enablers. And we see this opportunity having... for a great growth in the future.

Matt Dezort, Analyst, William Blair:
That makes sense.

Eugene Steinberg, Chief Technology Officer:
I can... I can... I can add to that that many of our customers observe the performance and productivity of our delivery teams using our GAIN platforms. They become interested, and they want those platforms and those methodologies inside their own software factory, and we are helping them to establish the tools, methodology, and change management, which is required to gain a similar productivity in their broader organization.

Matt Dezort, Analyst, William Blair:
That's a good segway, Eugene, for my follow-up on GAIN adoption and just the S-curve that implies. As you accelerate GAIN rollout, how should we think about you know, that adoption curve and AI... pure AI revenue? Is it likely to scale linearly? Or you're talking about wallet share gains from AI. Is there a way we could see some exponential growth, and how could you drive more sharper inflection in that AI penetration with GAIN?

Eugene Steinberg, Chief Technology Officer:
So what is interesting about our GAIN strategy is that we are consolidating all our IP from multiple accounts, from multiple practices under the same umbrella. And AI helps us to do that very, very rapidly and quickly. And our embedded engineers, forward-deployed engineers, are all tasked to bring back the learnings, the ideas, what works and what not works back to the GAIN platform. And part of the GAIN platform is also open source that helps to drive the insights from the broader community and put the GAIN platforms in front of many leaders. At this point in time, we observe a growth of the revenue direct-revenue from GAIN platforms, but much more important, we observe the growth of the overall connection and expansion of our relationships inside our accounts, and in new accounts, which are driven by these platforms. So we see many of the inbound interests and conversations, which result in new leads, new opportunities, and new converted business from GAIN platform. This is what is happening right now.

Cary Savas, Director, Branding & Communications:
Thank you, Matt.
Our next question comes from Surinder Thind of Jefferies.

Surinder Thind, Analyst, Jefferies:
Thank you, guys.

I'd like to start with a question just around this idea of there's a bit more excitement around moving from proof of concept to maybe the actual implementation projects. And that commentary seems to be

a bit more universal. From your perspective, can you maybe talk about what the revenue journey for that looks like, meaning, how big a proof of concept project would be if, you know, if it's a few hundred thousand dollars, does that become a $2 million project, or what's kind of the range of outcomes that we can expect here? As we think about more of those proof of concepts coming and how that would impact the growth rate?

Leonard Livschitz, Chief Executive Officer:
Surinder, so again, I will give you a little bit of a high-level. And I think because it's all revenue touched, Vasily will give you a little bit more color. So, there are different proof of concepts. The definition of proof of concept could be quite stretched. Both from intent and then dollars associated with that and follow-ups. When we looked at proof of concepts as a result of our partnerships, for example. That resulted in some of the very meaningful programs where the customer embraced not only our partner solution, but our offering, which was in conjunction with this partnership.

When we look today, and specifically at the suite of GAIN productivity, it's actually very interesting. Eugene mentioned about the inbound interest. As you know, for us, for Grid Dynamics, and our size and capabilities, visibility is very critical. So, the customers would reach to us with something we still call proof of concept, but those are substantial projects because the measurement of proof of concept sometimes is driven today not by the amount of dollars—could be quite more substantial in many cases—but the time to implement. The whole short-term engagement definition, which used to be followed or preceded by the proof of concept, becomes the proof of concept itself, and then it's a major rollout. So this is conceptually changing the definition of proof of concept and revenue associated with that. But I'm sure that Vasily will give some more details.

Vasily Sizov, Chief Revenue Officer:
Yes, so many customers start, definitely, with implementations of some smaller pieces of business cases, which have tangible business results in order to demonstrate it for their boards, for their management. And then using that as an example, essentially request more investment into that, which eventually gets converted into platform build out, to build AI harness, and etc. And this trend definitely persists. That's what we see with our customers. And I can state that for platforms work, this work essentially is much more sticky and longer-term in nature than the PoCs. And having built the platform, of course, there is a growing appetite to build more and more business cases on top of this platform. So it grows like a snowball, and that actually is what's reflected in our pipeline

Yury Gryzlov, Chief Operating Officer:
And I think I just wanted to add that it also depends on the industry, right? We mentioned today about physical AI and robotics. Definitely, there is a lot of proof of concepts in those areas, but at the same time, it also depends on how deep you are in your relationship with the customer and other programs around/outside even of those areas, right? And that's where those proof of concepts could be actually quite significant. Sometimes it could be just maybe a few weeks of small engagement, sometimes it could be six months plus. And going back to the GAIN model and our platforms in the GAIN, I think this is where also we try to condense the snow in a way to speed up this implementation as much as possible for a customer. And that also contributes to the ratio of the proof of concept revenue versus the longer engagement implementation revenue.

Leonard Livschitz, Chief Executive Officer:

So just to summarize for Surinder, the PoCs associated with FDE consultancy, GAIN model modernization subjects, and Agentic AI overall are very substantial from get go. The physical AI part is what, traditionally, we would call proof of concept, because it's such an innovative way to modernize modern productivity and interface between humans and robotics. So these type of PoCs are more traditional way, and their revenue will follow with a scale which you typically expect from PoCs.

Surinder Thind, Analyst, Jefferies:
Cool. And then maybe thinking about the data and AI practice and the really high growth rate that we're seeing there, the 30% of revenues, can you help me understand what's going on in the other 70% when I do the math? I get to roughly about a 10% decline in that other 70% of revenues. Is that how much of that is just cannibalization by the data and AI component? Because I assume every new piece of work probably falls into that bucket. And then is there a component that are maybe in that legacy, I'll call it legacy bucket for lack of a better word, other elements to that, such as pricing compression, or just other factors to think about?

Anil Doradla, Chief Financial Officer:
Very good question, Surinder. I'm actually going to make it much simpler. It's not even that complex. In our case, when you look at our business trends, from time to time, we might see a significant customer. I mean, loosely using that word, maybe a top 30 customer, top 40 customer, right? have some changes. Maybe there's a change in strategy, or a big project is completed, or something changes, and we can have some volatility there. If you go back over the past couple of quarters and look back at Leonard's commentary, he talked about the sensitivity of brick-and-mortar retail, for example.

And you see, you create some volatility there. So very simply put, if I were to extract some of these volatilities there, we would see a better growth pattern. The second point is, your question is whether there's cannibalization. What we see is, and Eugene and Vasily can back me up on this, when we get into our clients, especially in our top 20 clients, we're going deeper and deeper. So, all the work is incremental AI work that we typically see.

And the third thing that we see here is on the pricing. We are not seeing pricing pressures. As a matter of fact, when you go into the AI world, obviously, there's a premium, but when you look at what we are doing over the past couple of years, and I look at a certain grade in a certain country, and whether there's pricing pressures, the answer is absolutely no. We're not seeing that. Now, we can argue whether there's a pricing increase. That's a different story, but there is no pricing declines.

Vasily, I don't know whether you want to...

Vasily Sizov, Chief Revenue Officer:
Yes, so I think it's a question of semantics, right? The cost per project, or cost per functionality, or piece of scope is definitely getting reduced because of the higher productivity and the shortening of the timelines for the delivery. That's kind of what's happening. But that leads to more work and more projects, rather than reductions. So I think that's a very important color.

Leonard Livschitz, Chief Executive Officer:
And then finally, you can look at the revenue per person. Again, we need to have some history to prove that the trend will continue to grow. So we're not trying to defend legacy. I think Anil was quite clear that some businesses fall off, right? And you're absolutely right. The AI content brings more new

business. But there is one element, which is not there: is us trying to retain some legacy business and compressing our margin. That's just not part of it.

Surinder Thind, Analyst, Jefferies:
I mean, it does sound like there's a definitional component here, right, of how you, to your earlier point of how you divide up the buckets and trying to look at it collectively, plus all of the kind of the noise of project starts and stops, and things like that. So I appreciate that. Thank you.

Leonard Livschitz, Chief Executive Officer:
Thank you, Surinder.

Cary Savas, Director, Branding & Communications:
Ladies and gentlemen, this concludes our Q&A session for today. I will now pass it over to Leonard for closing comments.

###

Closing Comments

Leonard Livschitz, Chief Executive Officer:
Our top accounts are expanding. Our AI programs are moving consistently from pilot to enterprise-scale deployment. And our platform portfolio is deepening — both organically and through the capabilities we have added in robotics and Physical AI.

I am confident in the second half of 2026. The strategy is working, the momentum is building, and the team is executing.

###