IBM 3Q 2026 Earnings

Conference Call Transcript

IBM Participants:

Olympia McNerney – Global Head of Investor Relations
Arvind Krishna – Chairman, President and Chief Executive Officer
James Kavanaugh – SVP and Chief Financial Officer

October 21, 2026 / 5:00 p.m. ET

MANAGEMENT DISCUSSION

Operator: Welcome and thank you for standing by. At this time, all participants are in a listen-only mode. Today's conference is being recorded. If you have any objections, you may disconnect at this time. Now, I will turn the meeting over to Olympia McNerney, IBM's Global Head of Investor Relations. Olympia, you may begin.

Olympia McNerney, IBM: Thank you. I'd like to welcome you to IBM's third quarter 2026 earnings presentation. I'm Olympia McNerney, and I'm here today with Arvind Krishna, IBM's Chairman, President and Chief Executive Officer, and Jim Kavanaugh, IBM's Senior Vice President and Chief Financial Officer. We'll post today's prepared remarks and a replay of today's webcast on the IBM investor website within a couple of hours. The earnings presentation is already available.

To provide additional information to our investors, our presentation includes certain non-GAAP measures. For example, all of our references to revenue and signings growth are at constant currency. We've provided reconciliation charts for these and other non-GAAP financial measures at the end of the presentation, which is posted to our investor website.

Finally, some comments made in this presentation may be considered forward-looking under the Private Securities Litigation Reform Act of 1995. These statements involve factors that could cause our actual results to differ materially. Additional information about these factors is included in the company's SEC filings. So with that, I'll turn the call over to Arvind.

Arvind Krishna, IBM: Thank you for joining us today. Revenue growth in the third quarter accelerated to 6%, above the base case we laid out in July. This performance reflects acceleration across our segments, led by Software growth of 9%, Infrastructure growth of 5%, and modest improvement in Consulting. We entered the quarter focused on disciplined execution, which also drove strong performance in profitability and free cash flow. These results demonstrate the strength of our portfolio, the continued value clients place on mission-critical platforms and our confidence in both our strategy and the outlook for the balance of the year.

Technology innovation, in both hybrid infrastructure and Super Intelligence, is leading to better business performance. This quarter's results reinforce the strategic choices we have made over the last several years to advance IBM as a software-led hybrid cloud and AI platform company. In Software, Red Hat again grew 11% with bookings growth accelerating back to the high teens. OpenShift ARR growth accelerated to over 30% as organizations scale Super Intelligence applications and agents into production, reinforcing our conviction that we remain in the early stages of a long-term application modernization cycle. Transaction Processing grew 10%, underscoring the strategic importance of the mainframe and validating our view that the softness we saw in the second quarter reflected timing rather than underlying demand. Data grew 20%, supported by continued strength in Confluent. Infrastructure was up 5%, driven by continued momentum in Distributed Infrastructure, growing 32% this quarter. AI continues to increase demand for storage and we are gaining share

through differentiated offerings across Flash, Fusion, tape, and IBM Z storage, including AI-enabled capabilities that help clients scale and manage data for AI. Consulting revenue grew 2% as clients modernize to deploy AI securely and at scale. The durability of our portfolio remains a defining feature of IBM today.

z17 continued the strongest innovation cycle in the platform's history, generating more than $1 billion of incremental revenue over six quarters, compared with the prior program. This reflects the security, reliability, scalability, and compelling economics of the platform. As infrastructure costs rise, IBM Z delivers a 2-to-15x total cost of ownership advantage versus off-platform alternatives, positioning us to capture additional value as AI workloads grow. Building on that advantage, we introduced new single-frame and rack-mount IBM Z systems in July that help clients address data center space, cost, and deployment constraints without compromising performance or security. Client response has been strong, contributing to the highest midrange revenue launch in IBM Z this quarter, in nearly two decades.

Last quarter, we told you Lightwell represents a multi-billion-dollar opportunity for IBM and Red Hat to help clients secure their open source software. Since then, some of the world's largest enterprises have signed up as clients and we have built a pipeline approaching $700 million in a short period of time. We have released more than 40,000 remediated packages to the Lightwell network demonstrating the rapid scale-up of our remediation capabilities.

Now turning to our AI strategy. Enterprise AI is entering its next phase. Deploying more models or building more agents will not be enough. The advantage will belong to companies that can operate AI securely, efficiently and sustainably across the enterprise. Organizations are recognizing that AI is only as effective as the system around it: the data that provides context, the workflows that orchestrate action, the infrastructure that secures and optimizes operations, and the hybrid cloud foundation that provides flexibility without lock-in. IBM has assembled these capabilities across four platforms. Confluent provides continuous business context through a real-time data foundation. Watsonx, including Bob and Orchestrate, supports AI-driven software development and agent orchestration. HashiCorp and Apptio automate identity, access, and policy while optimizing cost and performance. Red Hat provides the open hybrid architecture spanning public cloud, on premises, sovereign environments, and the edge. Together, these platforms help clients move from isolated deployments to AI operating at enterprise scale.

Data sits at the center of the AI operating model. 72% of organizations report insufficient real-time data infrastructure which holds back AI at scale. Agents need the current state of the business. That is why we are excited about Streamhouse, a capability from Confluent that continuously captures, transforms and queries real-time data. It complements warehouses and lakehouses. A Streamhouse makes real-time, governed data continuously available to applications, analytics and AI, unifying operational and analytical data into a single architecture. We believe Confluent is uniquely positioned with a complete, open, and hybrid approach that helps clients avoid lock-in while deploying wherever their data resides.

As AI scales, clients ask two questions: what will it cost, and who controls it? We are innovating on both. Apptio now connects AI spend to business results, providing visibility into AI usage, token consumption, and the cost of running it across the enterprise. llm-d, an open-source project developed with Red Hat and Google, reuses previously processed context, helping clients run large AI models at 5-10x lower cost by making more efficient use of existing GPU infrastructure. And Bob now runs on-prem, so clients in regulated and sovereign environments can keep code and models on their own infrastructure.

Governance matters as much as intelligence. As agents multiply, enterprises need to know who each agent is, what authority it has and what it can access. HashiCorp Vault allows clients to register, authorize, audit, and revoke agents and secure their credentials. IBM Agent Identity and HashiCorp Vault are partnering with the NVIDIA Open Agent Safety Platform to strengthen agent security. We believe agent identity will become as foundational to enterprise AI as identity has become to the internet.

Finally, Quantum remains a strategic area of differentiation for IBM and a significant long-term opportunity. In July, three independent demonstrations of quantum advantage were achieved on IBM hardware with the University of Chicago, Algorithmiq, and Qedma, providing further evidence that quantum computing is moving from scientific promise towards practical value. Our Nighthawk r2 processor has achieved important roadmap milestones across speed, scale, and quality, delivering 7,500 gates, executing 100,000 circuits per second and enabling performance improvements of up to 19 times over prior-generation systems. That performance changes both time-to-solution and the economics of quantum computation: workloads costing approximately $100 to run on IBM could cost up to $10 million on slower competitive modalities. Recently, IBM was selected for Stage C, the third phase of the Quantum Benchmarking Initiative led by DARPA, the United States Defense Advanced Research Projects Agency, another validation of the strength of IBM's path to building a fault-tolerant quantum computer.

Commercial momentum is also growing. This year, IBM Quantum has signed 16 enterprise engagements across industries including automotive, aerospace, telecommunications, financial services, and manufacturing. This includes AT&T and General Motors amongst many others, who are now all part of the IBM Quantum Network, the world's largest network of quantum innovators. We are also advancing the technology and infrastructure required for industrial scale fault-tolerant computing, including the successful connection of our first modular cryogenic systems and continued progress on our 300,000 square foot Poughkeepsie manufacturing facility. This along with the closing of the $1 billion award from the Department of Commerce in Anderon for manufacturing quantum wafers gives us confidence that after the Starling demonstration, we are ready to scale. We look forward to talking more about our progress at our Quantum Day on December 10th.

In closing, our third quarter performance reflects the strength of our portfolio and the progress we are making against our strategic priorities of delivering both sustainable revenue growth and higher profitability. Our results reinforce our confidence in delivering revenue growth for 2026 in a range of four-to-five percent and enable us to raise our year-over-year free cash flow growth expectation to more than $1 billion. With that, let me hand it over to Jim.

Jim Kavanaugh, IBM: Thanks Arvind. In the third quarter, we delivered solid results with 6% revenue growth, 70 basis points of operating pre-tax margin expansion and 11% diluted operating earnings per share growth. And through the first nine months, we generated about $8 billion of free cash flow, up 8% year-over-year. These results reflect strength across our portfolio and solid execution against our priorities. Let me walk through the performance in each of our segments.

Software revenue growth accelerated four points sequentially to 9% this quarter, including over four points of organic growth. This performance was driven by our high value, annual recurring revenue base which grew 8% to $24.9 billion. Three of our four Software sub-segments grew revenue double-digits, led by the combination of strong new business performance and pipeline closure rates, including the successful closure of around two thirds of the large deals that slipped from the second quarter. Our results represent the high end of the range we outlined in July, reinforcing our confidence in the strength of our portfolio.

Turning to our sub-segments, Red Hat revenue grew 11% again this quarter, fueled by bookings accelerating back to the high teens. Lightwell was recently made generally available, and we have strong early demand, with a pipeline that is already approaching $700 million. OpenShift is now a $2.4 billion ARR business with growth accelerating to more than 30%. Our virtualization offering continues to scale, with nearly $800 million in contracts signed since early 2024 and about $750 million in pipeline.

Data grew 20%, driven by another strong performance from Confluent, with bookings growth of about 20%. Transaction Processing grew 10% in the quarter driven by demand for this mission-critical software and strong execution against our pipeline. HashiCorp continues to perform well, growing bookings across all geographies and accelerating revenue growth sequentially; however, Automation was impacted by a tougher compare.

Turning to Infrastructure, revenue growth accelerated to 5%, driven by 32% growth in Distributed Infrastructure. Storage delivered record revenue growth and continued to gain market share, with broad-based demand across Flash, Fusion, tape, and IBM Z storage. Strength across Storage and Power reflects the growing importance of differentiated infrastructure as clients modernize their environments, address data growth, and support AI at scale. With a sizable backlog in Distributed Infrastructure exiting the quarter, we are well positioned for continued momentum in the fourth quarter.

z17 continues to perform exceptionally well, wrapping on the strongest start to a Z program last year. Program-to-date performance remains well ahead of the prior generation, up over $1 billion, and this

quarter, we had the strongest midrange launch in IBM Z in nearly two decades. Beyond the strength of the program, IBM mainframe clients representing approximately 85% of installed MIPs capacity are maintaining or growing their capacity, and aggregate installed capacity has tripled over the last few programs.

Consulting revenue grew 2% this quarter, reflecting continued client demand for AI-enabled transformation. Growth was led by Strategy and Technology, up 4%, with strength in application modernization and data and AI transformation, while Cybersecurity delivered nearly double-digit growth in Intelligent Operations. AI is increasingly embedded in the broader transformations our clients are undertaking, representing about 30% of our Consulting revenue in the quarter and 35% of our backlog. We are also expanding the ecosystem of capabilities we bring to clients, including through our new strategic partnership with OpenAI and our expanded relationship with Google Cloud, as we help clients scale enterprise AI. Clients remain at different stages of adoption, but we continue to see them modernizing applications, strengthening their data foundations and reengineering workflows as they move forward toward embedding AI into core business processes at scale. Our integrated model, combining deep domain expertise, technology, and reusable assets positions us to help clients capture business value from AI while accelerating time to value.

Let me now discuss profitability. Adjusted EBITDA and operating pre-tax margins expanded by 60 and 70 basis points, respectively. As I mentioned last quarter, we accelerated our productivity initiatives to help enhance our operating profit margin and free cash flow commitments while continuing to invest for growth. We made significant progress this quarter, driving efficiencies across G&A and sales, reinforcing the power of our productivity flywheel. Segment profit margins expanded by 210 basis points in Software and 120 basis points in Infrastructure, driven by ongoing productivity actions. Consulting segment profit margin declined modestly, as we reinvested productivity gains into digital assets, accelerating our transition toward a Services-as-Software model to capture the growing AI opportunity.

Through the first nine months of year, we generated about $8 billion of free cash flow, up about $600 million year-over-year, driven by about a $1 billion increase in Adjusted EBITDA offset by higher capex, cash taxes, and net interest expense. Our free cash flow margin for the first nine months of the year is the highest in reported history, highlighting our strong execution against productivity initiatives and laser focus on operational discipline.

We exited the third quarter with a strong liquidity position and a solid investment grade balance sheet, with cash of $11.1 billion. We returned $4.8 billion to shareholders in the form of dividends through the first nine months of the year. Our debt balance ending the quarter was $63.8 billion, including debt of $12.7 billion for our financing business, with a receivables portfolio that is almost 80% investment grade.

Now, let me talk about what we are seeing going forward. The strength and diversity of our portfolio, disciplined capital allocation, and sustained productivity execution underpin our revenue growth,

earnings expansion and free cash flow performance. Through the first nine months of the year, we delivered over 4% revenue growth, 9% operating pre-tax income growth, 10% operating earnings per share growth, and 8% free cash flow growth. We continue to expect revenue growth for 2026 in the range of four-to-five percent and are raising our expectation for free cash flow growth to more than $1 billion.

Ninety days ago, we provided a full-year revenue growth range, with profit and free cash flow expectations anchored to a 4% base case. We continue to take the same approach for the full year, anchoring profit, free cash flow and segment expectations to that base case. Given our strong third quarter Software performance, our confidence in the full-year outlook has increased. We now expect full-year Software growth of approximately 7% to 8%, with about 7% as the revised base case. For the full year, we expect double-digit growth in Hybrid Cloud, mid-teens growth in Data and low-single digit growth in Automation. For Transaction Processing, we are raising our expectations for the full year with growth expected to be flat. Infrastructure growth accelerated in the third quarter with significantly stronger program-to-program performance of z17 underscoring the durability of the platform and the value it delivers to clients. Combined with continued strength in Distributed Infrastructure, we continue to expect low-single-digit Infrastructure revenue growth for the full year. In Consulting, backlog quality and momentum in generative AI continue to support a modest acceleration to low-to-mid single-digit revenue growth for the year.

In the third quarter, our operational discipline and accelerated productivity initiatives allowed us to enhance margins and deliver free cash flow commitments while continuing to invest for growth. These actions include deploying AI and automation at scale across our company, reducing third-party spend, improving sales and marketing efficiency, using AI to drive more efficient software development, optimizing our supply chain, and enhancing services delivery. We continue to expect 100 basis points of operating pre-tax margin expansion for the full year, reflecting the flow-through from the third-quarter earnings beat and continued operational discipline. And our operating tax rate expectation for the year remains in the mid-teens.

For the full year, we are raising our expectation for free cash flow growth to more than $1 billion. The primary driver of this growth will be Adjusted EBITDA, offset by higher capex, higher cash tax, and net interest expense. Our free cash flow continues to be driven by growth in Adjusted EBITDA and demonstrates consistent free cash flow realization of about 130% over the last few years.

Let me close by saying that our third quarter performance reinforces our confidence in our portfolio, the resilience of our business model, and our outlook for the year. At the same time, we continue to invest in the technologies that will shape IBM's next chapter of growth. As Arvind mentioned, Quantum continues to advance across technology, ecosystem adoption and scale-up. While still early, the progress we're seeing, from enterprise engagements to government validation, reinforces our confidence in both our roadmap and the long-term opportunity ahead. We look forward to discussing in more detail at our Quantum Day in December. Arvind and I are now happy to take questions. Olympia, let’s get started.

Olympia McNerney, IBM: Thank you, Jim. Before we begin Q&A, I'd like to mention a couple of items. First, supplemental information is provided at the end of the presentation. And then second, as always, I'd ask you to refrain from multi-part questions. Operator, let's please open it up for questions.

QUESTION AND ANSWER

Operator: Thank you. At this time, we'll begin the question-and-answer session of the conference. [Operator Instructions] And our first question comes from Amit Daryanani with Evercore. Please go ahead.

Amit Daryanani, Evercore ISI: Yeah. Thanks a lot. Good afternoon, everyone, and great to see some nice software acceleration here. Arvind, I think, last quarter, we discussed a lot about customers prioritizing Infrastructure over Software purchases. This quarter, it seems like Transaction Processing is at double-digit growth, while Infrastructure has remained very strong for you folks.

Can you just talk about what changed in customer spending priorities? Are you primarily just recovering deferred transactions at this point? Or are you actually starting to see infrastructure investments starting to create more incremental demand for IBM Software? And with Red Hat bookings actually growing at high-teens, is that an early indication that Software can actually accelerate as you go into 2027? Thank you.

Arvind Krishna, IBM: So, Amit, thanks for that question, and also thank you for observing the numbers. Let me comment on the overall customer priorities, the tie towards Infrastructure and Red Hat, and I'm going to ask Jim to cover the Transaction Processing piece a little bit.

So, first, as we look at customer spending priorities, customer spending priorities in trying to leverage AI remain strong. But as you begin to leverage it, how does the average customer get value? They have to rewrite an application. They have to modernize an existing application. And you then ask the question, where do those run? That tends to be a tailwind for Red Hat OpenShift. And I think we are at the very beginning of this cycle, so I expect that to carry on for a long time.

Then the moment they get to saying, okay, I'll write the application on Red Hat OpenShift, you then get to, how do I unlock the data from all of the enterprise data that I need to power that AI application? Remember, 90% of enterprise data is not on the cloud. It's very much inside the enterprise. So, Confluent tends to be a tailwind. I expect these to carry on for a long time.

The automation portfolio will get a benefit when these applications all get to scale and people start worrying about managing hundreds and thousands of them, at that point, automation will also get a tailwind.

On the point you're making on Red Hat, look, we're very comfortable saying that based on what we can see, we are comfortable about double-digit growth in the fourth quarter of this year. Right now, we don't really want to comment on 2027 growth, but generally speaking, I can say when we get these kinds of booking growths, it translates into good revenue growth going forward. Jim, maybe a comment on TP.

Jim Kavanaugh, IBM: Yeah, just to build on your points, Arvind, and Amit, thank you very much for the comment and the question. We're very pleased overall with our Software performance. Team executed extremely well, 4 points sequentially to 9% printed here in third quarter. And it was all organically driven. Pervasively across our portfolio, three of our four subcategories accelerated and grew double digits overall.

And that was a reflection of both, we talked a lot last quarter about the 80% high-value recurring revenue growing nicely at high-single digit, but we saw a very nice rebound in our transactional business, that remaining 20%, that grew actually double digits, led by our ELAs that were up north to 20%, with an NRR of north of 130% overall. That's a function of, yes, disciplined execution on slipped deals, but also a function of net new business contribution in the quarter. And I would tell you, before all of you ask, about a quarter of it is the slipped deals with return – with regards to overall software acceleration, and about three quarters of it is net new business that we are capturing overall, so to build on Arvind's point.

Now on TP. TP, we're obviously pleased, double-digit growth. I think this instantiates the demand for our mission-critical strategic mainframe software. A lot of discussions a quarter ago about demand destruction versus timing, I think this puts this to rest right now because when you take a look at it, we've been talking all year long about the reinforcement of confidence we have in our mainframe platform and how we monetize the strong z17 performance, right? That model is, the monetization model is the more capacity you put in the marketplace, the higher monetization opportunity, the higher price opportunity, higher value creation opportunity. z17 today, installed MIPS are up high teens, six quarters in. Shipped MIPS are up over 100%. That is the future monetization opportunity that we feel extremely confident about as to the acceleration potential of TP in the future. So, appreciate the question, and great place to start.

Olympia McNerney, IBM: Operator, let's take the next question.

Operator: Your next question comes from Wamsi Mohan with Bank of America. Please go ahead.

Wamsi Mohan, Bank of America: Yes. Thank you so much. Thanks for the color you just shared on TP, but maybe if we could just double click on that. I mean, last quarter, you were thinking that TP might be down low to mid-single digits for the second half, and you called it a 2027 growth vector. It grew 10% instead. So, maybe, Jim, as part of those, like a quarter of Software upside really came from those pushed deals. Could you share some color more specifically with respect to TP, if that was similar or higher in relation to the deals that closed from those pushed deals? And how many of those low tens of slipped deals have now closed because you've not really changed your 4% to 5% for the year, but your commentary and your numbers would suggest that you're gravitating much more towards the higher end of that. If you could share any more incremental color. And on through-cycle TP performance, is that still something that we should think of as mid-single digit? That'd be great. Thank you so much.

Arvind Krishna, IBM: Wamsi, I know you addressed the question to Jim, but I'm going to start and then ask Jim to give a lot of color on both, the slipped deals, the TP performance and the cycle-to-cycle compare, if I sort of paraphrase your question short.

So, let's just talk about TP. The questions everybody had last quarter was does this imply that overall demand for mainframe has gone down? And we were very confident that this was a timing question, not a destruction question. And so, at the macro level, as people deploy all of that hardware, as they deploy that MIPS, as the volumes of the underlying transactions increase, they need this mainframe software to support that mission-critical business. That's why we were confident it would come back.

Now, to your point of us saying low to mid-single-digit decline, we want to be cautious and we want to be prudent. We want to see the absolute baseline. And we don't want to call above it. I think what really happened was we got not only the deferral from 2Q, we got a normal 3Q buying cycle on mainframe, so we kind of got it both together.

I would be candid and tell you, I don't expect double-digit performance in TP to keep going, but I would expect to see now more normalized performance, especially going into next year and perhaps moderation from 3Q going into the fourth quarter. But it speaks to the underlying demand for the mainframe around resilience, around the economics, around the architecture that supports eight 9s of availability. Let me point that out, eight 9s of availability, but the transactions in the mainframe tend to be almost the opposite of what happens in AI. Here, you need things that are completely deterministic, that are repeatable, and that you can look at it with a financial cold eye that they're meeting all the requirements of those kinds of transactions. And there's plenty of other things that will happen on Distributed Infrastructure. I think in terms of the numerics, I'm going to hand it to Jim to cover with you.

Jim Kavanaugh, IBM: Yeah, Wamsi, thanks for the question overall. Let me just kind of reiterate a little bit about where I ended the last question to Amit. When you look at the health of our Software business that accelerated 4 points quarter-to-quarter, it was a function of both the high value recurring revenue. And again, the underpinnings of that is the Red Hat strength and acceleration, it's excellent start to our Confluent acquisition, it's record performances in Hashi overall, and that part of the business is doing well.

The transactional is where the big difference changed. Last quarter, it impacted our Software growth by 4 to 5 points when we were down double-digit transactional. This quarter, we were up double-digit transactional. And I'll remind you that's 20% of our Software book. Annual book's about $32 billion, give or take. And that today offers about two-thirds to three-quarters of that is ELAs. ELAs are strong incumbency positions that we've earned trust of our clients for decades around mainframe running their core architectures. On top of that, we build a moat and we build a software distributed business model that captures that very attractive economic multiplier of $3 to $5 a software for every dollar we land on mainframe.

Now, within that, TP is disproportionate. But with that moat, we give clients flexibility on where they want to choose that ELA spend as to where it goes between data, between automation and between TP. Red Hat's not part of our ELAs overall. And what we saw here both is a contribution in the quarter on slipped deals. Let's put this in perspective. We printed a number, what, $7.6 billion of software. The slipped deals were tens of slipped deals. So, let's kind of put that in a ballpark from 2Q of, what, $400 million, which is what we missed software by. We closed two-thirds of those. That's $250-plus-million. That is the 25% of the contribution of the growth that we've had of 4 points. By the way, that was disproportionate to TP. That's why you see TP well up at double digits overall. We didn't get a lot out of that in Data and in Automation. And some of that should be fourth quarter.

Now, fourth quarter, we gave a guide, I think, to use Arvind's word, anchored in our philosophy 90 days ago, prudent, cautious. We're anchoring this thing at the low end approximately 7%. By the way, put that in perspective, I think the current Street estimates for the year are about $32 billion of revenue. We feel comfortable with that $32 billion of revenue. So, that's kind of where we're at. I think I've covered most of your questions.

Olympia McNerney, IBM: Great. Operator, next question.

Operator: Your next question comes from Ben Reitzes with Melius Research. Please state your question.

Ben Reitzes, Melius Research: Hey, guys. Thanks. Appreciate the question. Wanted to ask about your long term. The short-term puts and takes we get, but at your latest Analyst Day, you laid out that you would approach double digits in Software as a long-term target. Based on what you're seeing, do you still feel confident in that? And how is M&A going to play a role in terms of achieving that? Because in order to hit it, I guess, you got to accelerate 1 point or 2 points after this year. So, a lot of folks are wondering where M&A plays in that. Thanks a lot.

Jim Kavanaugh, IBM: Thanks, Ben. I appreciate the question overall. Obviously, and this is not going to surprise you, we'll spend a lot more time on this question here in January. We got the team extremely focused already. We're 21 days into arguably the most important quarter we have in a fourth quarter overall. But let me take a step back first around IBM and then, as Arvind and I have been talking here, as we're heading into our planning cycle and laying out all of our investment and capital structure strategies for 2027, talk a little bit about IBM, then hit your question on Software.

We've got a lot of confidence in our strategy and our portfolio is well positioned to deliver sustainable, durable growth on our model. We feel extremely confident in that. Why? We got an attractive TAM, market backdrop of technology, the diversity of our portfolio, leveraging our hybrid cloud and AI platform-centric model, our continued portfolio optimization to shift to higher-growth end markets, new innovation we continue to bring to market around Lightwell, around GenAI, and around to your question, M&A growth synergies. All of that will play out.

But from a segment perspective, since you asked about Software, I'll give you some early indicators. One, we got a high value recurring revenue stream that's been accelerating all year long, $25 billion ARR book of business, high-single-digit growth. We execute to what our guide is here. We should be approaching double digits here in the fourth quarter on that 80% high value. That's one critical indicator.

Two, new innovation we're bringing to market. We're extremely excited, to Arvind's point, around Red Hat, the acceleration to high-teens growth. We've got $700 million worth of pipeline on Lightwell alone in front of us right now, with $1 trillion TAM opportunity around open source vulnerabilities. We got virtualization, a tremendous opportunity. We got hybrid cloud containerization opportunity. So, we feel pretty good about the markets we're in, the investment in the portfolio that we have built up overall.

Three, we got – 2027's a year of growth synergies, right? 2026 was about integrating and getting Confluent into our mode of operation and go to market and scale. 2027 is around the growth synergies. And finally to the TP question, we've got high-teens growth in installed MIPS out in the marketplace. Last five years, our MIPs are up 2x. Last 10 years, our installed MIPS are up 10x. That is the monetization vehicle at TP. And Arvind talked about it. I talked about it. We think that's a growth acceleration back to our model overall. We'll quantify that in January. But that gives you a little bit of perspective of our early thinking about 2027 on Software and IBM.

Olympia McNerney, IBM: Great. Operator, let's take the next question.

Operator: The next question comes from Raimo Lenschow with Barclays. Please go ahead.

Raimo Lenschow, Barclays: Perfect. Two quick ones, if I may. First, congrats on the Software performance. On Data, 20% growth. Obviously, some of that is Confluent, but, like, what are you seeing there on upsell or cross-sell on Confluent? You have it for a few quarters now, but that looks like it's starting to click for you. So, can you kind of double click on that one? And then, one for Jim. Gross margins came up with investors. Can you just give us a comment there, what drove the performance there? Thank you.

Arvind Krishna, IBM: Raimo, thanks for the question and thanks for noticing the Data growth. If I look at Confluent, actually want to broaden it, there is Confluent, which is important for moving data in real time, allows data to remain in place but still be available for AI. But coupled to that are what are called lakehouses. And we are very excited about the upcoming cross-sell opportunity with Streamhouse, which is also another way to store data that is sort of in flight.

So, the Confluent piece unlocks your data, be it inside the enterprise, be it in an application, be it in a SaaS application, be it on a cloud, but still have it completely available for AI agents as well as inferencing. Really important capability, and you saw that growth pulled through. But lakehouses are also important as people are beginning to cleanse their data for leveraging AI, and we saw growth there also. And then, we're excited about what is going to come with the capability we're about to

announce in a couple of weeks around the Streamhouse, which is also then going to be directly a cross-sell and up-sell on top of Confluent.

So, unlocking data for AI, making sure that you don't have to create multiple copies, because that creates a lot of extra cost, but still have it available as quickly as you need it, is the fundamental capability, we believe, is going to be unlocked in the enterprise, leveraging all of these capabilities. And that contributes to the overall growth there. Jim, I'll give it to you for the gross margins question.

Jim Kavanaugh, IBM: Raimo, thank you very much for noticing the operating leverage engine that we've been driving across this company. You look at, what, the last four years were probably up, what, about 1,000 basis points of operating leverage across this company.

By the way, that has delivered high-quality sustainable cash generation. Over the last four years, our cash is up roughly $7 billion of growth and 700 basis points of free cash flow margin on that 1,000 basis points of operating leverage. But when you look underneath it, we continue to improve the structural competitiveness of this business model. We still believe we got a lot of headroom here. When you take a look at our profit, earnings up double digits, our operating margins are up 70 basis points. It's delivered by, we are driving, in the third quarter, 500 basis points of gross productivity in our model.

We've always said in our model, about a third is for margin protection, about two-thirds enables us to drive that productivity flywheel where we reinvest back into the business around R&D, innovation, go-to-market model ecosystem, to drive that flywheel of growth and competitive advantage. 500 basis points is well above what our model is overall.

How are we doing that? We're running our playbook. That playbook, we believe, has headroom. It's reimagining everything between Arvind and myself about how we execute and run this company. And that gets into leveraging data, a whole new way of workflow re-imagination, embedding technology, digitization, AI into our company. We've been talking about in prepared marks, how do you reimagine a whole new way for software development? How do we reimagine a whole new way on go-to-market models? How do we redesign the G&A structure, the way we run our company? And I think there's a lot of headroom there.

A few years in, we are now north of $5.5 billion of productivity out of this company. That is a great flywheel for investment flexibility and capability. So, most important thing is it provides that free cash flow generation engine, and it applies that investment flexibility so we can continue to compete both organically and inorganically to invest.

Olympia McNerney, IBM: Operator, next question, please.

Operator: Your next question comes from Brent Thill with Jefferies. Please state your question.

Brent Thill, Jefferies: Thanks. Arvind, I'm curious in some of the deals, maybe early in the year, there was a little bit of a AI bright light. I don't know what to do. I need to kind of wait and investigate. And maybe the investigation period went on and they realized maybe some of the systems couldn't

actually achieve what the claims were. And I'm curious if you're maybe seeing some element of this in some of the deals recouping where the claims were big, but perhaps the reality wasn't matching the claim and you're able to deliver on that. Is that something you've seen or is that not a scenario that's that you've experienced?

Arvind Krishna, IBM: So, Brent, let me make sure I parse the question correctly, but I'm going to sort of maybe answer it in my own words. I think what we're seeing is that at least larger enterprises are becoming a lot more cognizant and a lot more aware of what is the true ROI of all of their AI spend. They're also getting worried if some of the AI spending is completely out of control with no real return.

But the way I've always characterized it is, there is no question to ask, and this is the advice I give every C-suite member that I meet, that what we're trying to give people on AI, the productivity, the advantages of AI, the way it's going to change work, is very real. So, let me not make any confusion around that. That said, a lot of people are underplaying how much work it is to integrate AI into your workflows, to fundamentally sit back and say, how do I redesign the workflow to take advantage of AI? It's not a little RPA tool that you bolt on and you get something. You – if you do not spend the time and energy to redesign your workflow, you're not going to get the full benefit. So, if you're willing to do that, there's a huge benefit to be had.

But that implies you've got to pick a few, you got to learn how to scale them for full impact across the enterprise. And don't do 100 little experiments, where you then do not have enough investment to get them to scale. That is what I think we are seeing. So, I think what we're seeing, Brent, is now enterprises, as they're getting mature, as they're stopping the experimentation, are saying, all right, who can help me, get through this bigger piece of workflow design as well as integration into our existing systems, because you're not likely going to throw out your ERP system. So, how do you leverage it, but still get the benefits of AI? And then they're picking their partners in terms of which software do they – are they willing to trust to do this? Which people have the knowledge to come do this with them as well as, can you be really efficient about your token spend as opposed to leave it completely uncontrolled?

So, I agree, I think, with your thesis, but maybe a slight redirection in terms of why we're seeing that change. And yes, that is a tailwind for us. That is why I kind of repeat what I said at the beginning. That is why OpenShift, Confluent, mainframe with fraud are all seeing the benefits of that play go on, as we are proceeding down this path.

Olympia McNerney, IBM: Great. Operator, let's take the next question.

Operator: The next question comes from Fatima Boolani with Citi. Please state your question.

Fatima Boolani, Citi: Good afternoon. Thank you for taking my question. I wanted to go back to the discussion around Transaction Processing, simply because it was such a remarkable sequential recovery in that franchise. And Jim, I understand so much of that is piggybacking off of the ELA

momentum that you regained in the quarter from the transactional side of the business. But I'm curious if you can help us unpack a little bit more on that strength.

I think you talked about a high-teens growth in MIPS shipped. So, wondering if you could unpack and break down for us the quantity and shipped MIPS growth in some of those slipped deals, because those were existing customers. And to the extent, you've been able to leverage a little bit of pricing power in this situation, because I know that's not historically a lever that you've used. So, we'd love to get a little bit more second order detail on that Transaction Processing slope of recovery between the quantity and quantity growth, as well as pricing uplift. Thank you.

Jim Kavanaugh, IBM: Okay. Thanks, Fatima. Appreciate the question. Let's double click back down on TP. First, let me just reiterate what is the monetization model, right? TP is Transaction Processing that runs a mission-critical software on top of the most enduring platform, that being mainframe.

By the way, as we've talked about many times before, mainframe runs 70% of the industry's transaction volumes in terms of value: 45 of the top 50 banks, 4 of the 5 top airlines, 7 of the top 10 retailers around the world. I think 75%, 78% last number I've seen of the Fortune 100 run mainframe as a core architecture.

So, TP is the core mission-critical software that enables the most secure, reliable and scalable platform around. The way we monetize that value is the more capacity out there, the more opportunity to monetize value, the more opportunity to monetize price, the more opportunity to monetize value creation of TP. So, when you take a look at our last two, three cycles, we run 145 million MIPS around the world. Some of those are standard MIPS on the platform. Some of those are new Linux specialty MIPS, new cloud-native applications, which is why the mainframe has been a growth vehicle over the last two to three cycles overall. That 145 million MIPS, our z17, six quarters in, 125%, the highest recorded program to-date above z16, which was the prior highest program to-date, we're up, I think, 18% shipped MIPS through six quarters. We're up over 100% of shipped MIPS.

Why do I differentiate between the two? Read the first one as consumption, capacity and workload. We can monetize value immediately. The second is the opportunity, it's the TAM, right? And us driving workloads on that. So, when you take a look at it, that's what's been driving the step function change and growing meaningful over time. It was only a handful of years that our TP portfolio was down, what, mid-single-digit, high-single-digit growth? We then stabilized that growth about four years ago, and then through z16, we saw an acceleration.

But I'll reiterate that z16, when you look at the launch year, it was about flat the launch year, then it kind of went up to about mid-single digits year plus one, and then it went up to high-single digit year plus two. Right now, in z17, we were about flat launch year. We got different dynamics playing out right now. So, we're about prudently cautious calling flat for the year. But with 100% shipped MIPS and 18% installed MIPS, I would say that's a pretty good future monetization opportunity in 2027 going forward.

Olympia McNerney, IBM: Great. Operator, let's take the next question.

Operator: Your next question comes from Brian Essex with JPMorgan. Please state your question.

Brian Essex, JPMorgan: Hi. Good afternoon. Thank you for taking the question. And congrats on the Software acceleration from me as well. Arvind, I caught the prepared remarks, or at least some of the news remarks that you had about clients deploying Super Intelligence in ways that transform entire operating models. And I appreciate that that's more than just increasing operating efficiency.

But the question I have for you is, as your customers assess those transformation efforts, how focused are they on sovereign AI? Is that becoming a more meaningful part of your conversations? And how are you seeing your customers focus on the protection of their data and IP as they assess how they deploy AI for their transformation efforts?

Arvind Krishna, IBM: Brian, thank you first for the question, and thank you for noticing the precise words I did use in the prepared remarks. So, to get to sovereign AI, sovereign AI is definitely becoming a much larger conversation. While we have talked about it for the last year and a half, I would not tell you that in 2024, it had much traction, in 2025, a few people began to pay some attention, and in 2026, it is now beginning to pick up. And let me explain for everybody else exactly what we mean by that.

If you are going to use a large public LLM, effectively while then – they're certainly not stealing your data, though they're implicitly maybe learning what kind of questions are asked, there are traces left in the underlying logs and so on that are not the full, but that do give some indication of what all it is being used for as well as some elements of the data that is being used. So, sovereign becomes how do you make sure that none of that happens? So, people are turning now to two kinds of answers. That is why, I think, there is such a big boom in AI infrastructure on-premise, and we can see some of the companies that benefit from that, but so will we in our mainframe with our Spyre cards as time goes on.

The other part of this is what is driving the boom behind open-weight models, because there you have the permission and the IP rights to run them on either on-premise infrastructure or on private clouds where the enterprise controls what is happening to the underlying logs and data. That is definitely what is happening there and we can see that. And the boom behind some of the companies that are creating open-weight models is also increasing behind that.

I'll reflect on a conversation I had with a large bank CEO. In this case, it was outside the US, but it's a very large and friendly country. And their view was, look, we'll do some early experiments on the larger frontier LLMs, but for what we're doing inside the bank, we need to run it in a way that is sovereign, meaning nobody can take any of the underlying data or underlying logs that are there. I was talking to somebody who does trading and they are on, we're not going to expose our underlying proprietary models, anything outside, but we need to run this in a sovereign way. I believe this is going to turn into

a tailwind for many players in the industry, certainly the infrastructure and semiconductor players, but also for those of us who help enterprises deploy in a secure and in a private way.

Olympia McNerney, IBM: Great. Operator, let's take one last question.

Operator: Thank you. And our last question comes from Matt Swanson with RBC Capital Markets. Please state your question.

Matt Swanson, RBC: Great. Thank you. And maybe this is just kind of a natural follow-up to what Brent was asking, but you had mentioned earlier that what AI costs and who controls it have kind of been like the two primary questions that enterprises start off with you. Who controls it obviously fits right into that sovereign AI conversation. But could you just kind of talk a little bit about how people are addressing the tradeoffs if there are tradeoffs between those two segments? And how they're architecting their networks in order to maximize for both those things and just how kind of like symbiotic they can be?

Arvind Krishna, IBM: So, Matt, thanks for that question and also allowing me to give a bit more detail on this topic. If I go back to late 2024, the difference between the most advanced model, which let's acknowledge is mostly a frontier LLM model, and alternate models was a 2:1 or 3:1 gap. Meaning if you ran some assessment, some benchmark and it was broad-based, not narrow, the top model might score a 70 and the open model might score a 20 or a 30. You'd kind of react with, that's a lot of capability we're giving up, and we all know that performance, in some sense, outcome for the enterprise is highly dependent upon that, that's too big a gap. Nobody's going to accept that.

In late 2025, we began to see that gap shrink, where maybe the gap was not 70 to 20 or 30, it was 70 to 50. But when you looked at those curves and you looked at the improvement and you looked upon the sequential improvement, we were very comfortable that we'd get this to probably 100 for the top and 90 for the alternate model. So, now the question becomes, how do you do that? Now, the question that the enterprise struggles with is where are my kernel engineers? Where are the people who worry about the underlying network? Who can make sure that the model is clustered on GPUs that are close to each other in network proximity? Who can make sure that you're doing the right caching of results so you get an incredible response? I'm telling you that that gap is probably still there, except for the most sophisticated enterprise. So, then you kind of go to, all right, I got it that the answer is four times more expensive, but it's the easy answer.

As those of us who work in this world, produce much more sort of easy button software that can give you that capability for the alternate models, and I would say it's not a huge amount of infrastructure. Our computation is that for an average large enterprise, you can get going with maybe two or three racks, which is not a huge amount of infrastructure, but it gives you the ability to run those open-weight models currently anywhere from 25% to 10% of the token cost of the frontier models. That is very attractive, given that they are now 90% as effective.

It's a simple question for any enterprise. Would you like something that's 10% of the cost if it's 90% as good as the best, or do you want to pay that 10 times higher cost? And then the question is how many tasks require that? I would tell you, some do, but the vast majority, summarizing an email, examining a contract, answering a customer call, composing an email, all of those can be done as well with the much cheaper answer. And that is what is going to drive a lot of this.

So, Matt and everybody else, thanking you for those questions. Look, we are pleased with our performance this quarter and all our segments accelerated growth. Our portfolio strength, disciplined execution and relentless focus on productivity reinforce my confidence in our trajectory. I look forward to sharing our progress with you as we close out the year. And we'll talk to you in January about that.

Olympia McNerney, IBM: Thank you. Arvind. Operator, let me turn it back to you to close out the call.

Operator: Thank you for participating on today's call. The conference is now ended. You may disconnect at this time.