

## Summary Translation of Question & Answer Session at Briefing on AI-Driven Software Development for Media, Analysts and Investors

Date: February 17, 2026  
Location: In person and live-streamed from Fujitsu Marunouchi Office  
Presenters: Ted Okada, Head of AI Strategy & Business Development Unit, Fujitsu Limited  
Izuru Kokubu, Head of Measures for Specific Projects Unit, Fujitsu Japan Limited

### Questioner A

***Q1: I would like to ask about your AI-driven platform. At last month's financial results briefing, I think it was said that now around 60% of your projects in Japan are using AI to some degree, but could you explain the difference between that initiative and this new platform you have announced? In addition, at the time of the financial results briefing, it was said that the results from using AI are still minimal, so please tell us how you plan to expand those results.***

**A1 (Okada):** The AI-Driven Software Development Platform we are discussing today was designed for the AI to understand relatively older and complex large-scale systems, and the platform provides a way to automatically improve the system. With the use of AI as you pointed out, we expect the deployment of AI itself to result in higher productivity, but our approach with the platform that we announced today is slightly different. Rather than increasing productivity to parts of a project, we are fundamentally upgrading existing systems, so we think it is a much broader scope of the use of AI.

### Questioner B

***Q1: You said you are applying it to 67 software packages. Please tell us specifically how the updating process works. Do you apply it to each customer separately, including customization work? Or is it more like a general patch that can be applied to common parts of different systems?***

**A1 (Kokubu):** What we do first is, as you pointed out, apply the patch to the software package. Next, for each customer's system, we work on the patch issued on the software package side, preparing everything up to the adjustments required for each customer.

**(Okada):** Let me explain a bit further. There are two main reasons why we are applying it to the government administration and healthcare sectors now. One is that they are subject to relatively frequent legal or regulatory changes. The other is that Fujitsu has for itself the relevant systems. You also asked about customizing it for each customer and the systems environment each customer has, and of course we have considered this issue. To do that, however, we need authorization from the customer, so it will depend on the outcome of discussions with each customer.

***Q2: In terms of the benefits for users, is it mainly that they can get the patch earlier, for example, several months faster compared to the previous approach?***

**A2 (Kokubu):** That is exactly right.

### **Questioner C**

***Q1: My understanding is that, with the Takane-Driven Initiative, integrating tacit knowledge into a development tool is an extremely advantageous point. On the other hand, as with the example of the Anthropic Shock, in using AI directly in the customer's workflow, it may be broadly interpreted as able to do everything. I think that interpretation is incorrect, but I would like to hear your views. My understanding is that, because of the potential for hallucinations, generative AI is not suited for precise processes. Am I correct in thinking that, rather than using AI as-is, you should use AI to generate the program code itself? I would also like to hear your views on whether this thinking will persist into the future.***

**A1 (Okada):** First, we think Anthropic is a wonderful tool, and we are actively thinking about utilizing it. On the other hand, by supplementing that with data or the tacit knowledge within the data, as well as Fujitsu's longstanding insights into workflows and expertise in system development, we think we can generate synergistic effects. However wonderful Anthropic might be as a tool, we think there are places where it comes up a bit short, and that is where we can complement the process. Likewise, we think the implementation of AI can accelerate the process.

Another point is that it is debatable whether important decision-making and workflows should all happen automatically with no human intervention. Our explanation today covers up to the point where AI understands the legal regulations, defines the requirements, determines which changes need to be made to which system assets, implements the changes, and then performs integration testing. We are not saying that we will provide customers with the results created by the AI as-is. Our expert system engineers will evaluate whether the changes deemed correct by AI are indeed correct so that we can deliver a process that has been validated. In that sense, for the foreseeable future, it will still be important to have "humans in the loop" providing expert intervention.

***Q2: In other words, although some end-users have started to create OneConsole-like programs that can bring together and complete a variety of elements, projects that include complex areas cannot be handled using OneConsole programs because they require human intervention, including judgement decisions, and this situation is likely to persist for some time?***

**A2 (Okada):** Yes, that is right.

### **Questioner D**

***Q1: You explained that this is for existing systems that are upgraded to reflect regulatory changes, but would it also be possible to use this platform even for systems being developed from scratch? If so, please describe the process to us.***

**A1 (Okada):** The initiative described today is designed to fully update existing system assets. In other words, it includes Fujitsu's domain knowledge on healthcare and government administration as well as our knowledge in building systems. Based on this, our announcement is about how we are able to automate the updating process when there are any required changes.

This process will continue to evolve. In the future, we think we can input our accumulated domain knowledge and Fujitsu's expertise in developing systems to also build new systems with new architecture. We are also pursuing the R&D for this at the same time. Accordingly, because we have the domain knowledge, we think it would be possible to build new systems from scratch.

***Q2: Are we correct to assume that the current scheme cannot be applied as-is to new systems, but this is the first step, and that, in step two and beyond, you may create a way to apply to the development of new systems?***

**A2 (Okada):** Yes, that is correct.

***Q3: Other companies that develop systems are also working to apply AI to automate certain phases of the development process. Some are calling for an "AI native" approach, where they reset all of the existing processes, and redesign all processes from the start to make it easier to apply AI. By contrast, I think Fujitsu is rather unique in automating while keeping everything in place, from the existing requirements and definitions to the design. Please tell us about the merits of your approach and what results you expect to achieve compared to other companies.***

**A3 (Okada):** At present, we think there are basically three types to AI-driven system development. There is vibe coding, in which natural language can be used to create systems, agentic development, and specification-driven development. If you want to build something quickly, vibe coding is suited. Specification-driven development is for complex system development projects involving multiple people.

The approach we took this time was first to have the domain of existing systems fully prepared, and then use that as a foundation for automation. By doing that first, we can also advance toward the next phase, which is specification-driven development. Fujitsu has accumulated operations knowledge that it has refined since the 1950s and 1960s, and we think there is value in starting from an approach in having AI be able to understand that knowledge. Rather than from a tools-based approach, we seek to achieve breakthroughs from operations knowledge.

**(Kokubu):** I would like to add to that from a business perspective. Right now, Fujitsu has 4,000 to 5,000 system engineers who, in such areas as government administration, healthcare, and education, spend each year chasing the inevitable regulatory changes and revise the software packages to reflect those changes. We think this new initiative will free up those system engineers from maintenance operations.

In relation to the many societal issues we still face, those freed-up system engineers will be able to shift their focus to human interaction with customers to help them use AI to address those societal issues. That will enable us to expand our business to resolve the issues customers are facing. That is what we would like to do.

**Questioner E**

***Q1: If you achieve what you have described today, I think it will change the way you have approached projects up to now, and I think it will also change your contractual relationship with customers. In particular, I think that, for big corporate projects, the relationship with partnering companies, such as second-tier subcontractors and third-tier subcontractors, will also change, and I would like to hear your thoughts about how you think things will change. In addition, amid these changes, there is a possibility that the role of your system engineers and project managers may change, as well as the number of people required. How are you training them and preparing them for the role of AI-Ready Engineering?***

**A1 (Okada):** I cannot discuss details because we are in the midst of preparing ourselves for April. As you point out, however, we think business conducted and invoiced in terms of man-months will go away. Work like writing simple source code will be converted to tools like the ones we discussed today. When that happens, we expect that billing will shift to a model in which customers are charged for the value they have perceived. For example, if there are several releases per year, those releases themselves are the value we will be providing to customers. In terms of calculating that value, we think that billing will be in the form of a usage fee for the technology used each time. For example, the basic charge may be like a public utility charge, which covers updates several times a year, and anything beyond that would be charged based on the volume of additional services needed. I think that would be a reasonable approach. In other words, we think a reasonable approach would be, the more modifications made, the more value we provide, and we charge based on their usage fee.

Regarding the shift in resources, we currently face an urgent shortage of human resources. We think there are tremendous needs to determine how to deploy agentic AI. We think we are still in the very early stages of this development, but I think there are many discussions on how AI can raise productivity, depending on the partial workflows to which AI can be used. We think what is most needed now is an understanding of how workflows can be transformed by being AI-driven. Instead of development work being contracted by customers, what is needed by the market is how to implement the idea of how AI-driven innovation can address the issues customers are facing. I briefly talked about the key word Forward Deployed Engineer, with the acronym FDE. In other words, we want to have a deep understanding with the customers of the problems they face and how we can quickly resolve those problems using technology. In the past, this was what system engineers did, but now we think these should be addressed by specialized engineering talent. I think this represents a swing of the pendulum. When we continue our FDEs type of business, we will understand the commonalities, leading to the development of packages in the fields of government administration and healthcare. Then there will be a time when AI will automate the analysis of packages from 30 years ago with domain knowledge enabling the experts to have free time and they become FDEs again. I feel that we are again at a point in time when this pendulum is swinging back.

***Q2: You used the term FDEs. Is my understanding correct that these are an evolution of the pre-sales personnel?***

**A2 (Okada):** I think they are closer to what we used to consider system engineers to be.

**Questioner F**

***Q1: Regarding your open platform, is my understanding correct that you will make the foundation of your platform available for customers to use through external sales and joint usage?***

**A1 (Okada):** We used the word “insourcing,” but ultimately, we believe that there will be cases in which Fujitsu will supply this platform to customers and create a world that the customer is able to operate themselves. In addition, we also believe that there are many of our business partners that are developing software packages businesses similar to ours. I believe that we will prioritize these areas. We also have no intention of dismissing ideas that are brought to us when we hold events.

What we wish to do is not simply the optimization of systems. Rather, we would like to become the standard for new system development, in order to change the way of working for system engineers who have struggled with application maintenance.

***Q2: Is this platform something that you envision that Fujitsu will disseminate within a community, rather than use to proactively generate business?***

**A2 (Okada):** Of course we will use this to generate business. We will use this platform to take Fujitsu’s business in a new direction. Rather than only Fujitsu using the platform to increase productivity, we are aiming for business in which everyone will be able to use it.

***Q3: I anticipate that customers will ask you to lower the cost of your platform, as it no longer requires as many man-hours per month. I believe that similar discussions arose during the implementation of tools up until now, but because human involvement will be even further reduced with this platform, such demands could grow even stronger. If this were to happen, how would you communicate to customers to pitch the value of this initiative?***

**A3 (Kokubu):** In the areas of local government municipalities and healthcare, we recognize that SaaS business has already started to penetrate and the business model for services is already somewhat established. There may also be an aspect in which it would become even easier for us to win business by lowering the cost of our services.

But, what we would most like to convey is that this dramatic increase in productivity allows for continuous system modifications. The laws and societal issues will continue to change. We want to increase the speed at which systems are able to catch up with these changes. This means that instead of a one-off update once or twice a year, systems will be updated on closer to a monthly basis, and will become more convenient. Once this occurs, the conventional fixed-price business style will no longer be viable, and we believe that the large-scale legacy systems that we handle will also be transformed by service-style and subscription-based business.

***Q4: Is my understanding correct that, even if your top line decreases with each project, your operating profit margin will increase, and that, through repeating these types of projects, your sales revenue will exceed its prior levels?***

**A4 (Okada):** Allow me to clarify the meaning of the previous answer. We believe that our top line will by no means decrease. This is because, up until now, the annual contract value has been fixed. Due to this, there were many cases in which customers wanted to improve resident services, but have not been able to. If the market accepts this new business model, then customers will request more modification projects from us, which will result in our top line increasing.

***Q5: I understand that is your perspective when looking at the overall picture. But, if you were to view things in units, such as a single project, I believe that your top line would decrease. What are your thoughts on this?***

**A5 (Kokubu):** As was mentioned in the prior endorsement comments, among our customers, there are many who require speed and accuracy and that is what they expect from Fujitsu. Out of the customers that we spoke with for this initiative, there was not a single company that only mentioned cost. If we are able to provide customers with an even greater number of services, then we believe that it is entirely possible to charge a higher price for them. We believe that how we utilize our productivity will be a key factor.

#### **Questioner G**

***Q1. It was mentioned that, with the method you are using, you will automate the process from the requirements and definition stage. But I felt that it is only possible for you to do so because the design specifications already exist. I believe that, among existing legacy systems, there are many cases in which there are only program codes, and the design specifications are so old that they do not exist. Would this method not be unsuitable for carrying out the migration of such legacy systems?***

**A1 (Okada):** Yes, that is correct. The method that we use assumes that the systems will have the documents necessary for the software packages business. Despite this, there were many systems that did not have these documents.

In the areas of modernization and migration that you mentioned in your question, it is often the case in which, as you said, the running environment does not match the documentation, or that such documentation does not exist in the first place. Although it is not in the scope of today's presentation, there already exists a "system specifications reversal" sort of function in the technology at Fujitsu Research. We believe that, it will be necessary to add that function to the platform. This is because, if we did not add it, then we would be unable to help with systems that have become so old that no one knows the information for them anymore.

***Q2: Is my understanding correct that your AI-Driven Software Development Platform essentially assumes that there will be design specifications for the software packages?***

**A2 (Okada):** Rather than design specifications, it assumes that, in a broad sense, the domain knowledge of the system will be properly compiled.

#### **Questioner H**

***Q1. Recently, in AI coding, work to verify programs that were created through the use of vibe coding has enormously increased. This has led to the issue of people no longer being able to do the work that they originally wanted to do. How are you working to address this problem with the workload of manually checking coding on your platform?***

**AI (Okada):** We have two approaches for this. The first is the multi-layer quality control multi-agent function that we mentioned during the presentation. Currently, in the physical world, outside of the work conventionally performed by system engineers, there are quality assurance departments and quality assurance professionals. We used what these people do in their work as reference for this function. For example, AI will first conduct an audit to see if the output generated by the requirements definition agent is correct while referring to such intricate rules as what Fujitsu's system development guidelines are and what should normally be done in the requirements definition, or what sort of information must be written in them. The fundamental mechanism is that, if an output passes the audit after completing multiple loops, it is then passed on to the next agent. After this process has been implemented, the output will finally be sent to humans who will serve as experts. These experts will have also read over 700 pages worth of information on requirements and rules and regulations, and have answers for what should be done. The answers from the experts and the results generated by AI will then be compared, and a final check will be performed to see if the results are of a high enough quality that they can be provided to customers.

We believe that it is only through having both a mechanism with which AI can audit AI and a human in the loop element, in which experts check information in critical stages, that we will be able to guarantee quality.

In addition, there is also our second approach. As for what customers have done up until now, for example, the revision of the medical payment system in Japan generally takes place around January. Because this is merely an overview, the particulars will be filled in through repeated questions and answers. After the details have been somewhat clarified around February or March, some reworking of the revision may occur. When this happens, one must go back to revise the design specifications and requirement definitions to what they were before. We believe this will be the critical component to reduce the time to market. The operations of our customers will also change, and I believe that they have had the experience of a system arriving with different specifications than they were originally told while they are in the middle of needing to make various preparations, which resulted in them becoming even more swamped with work. There were also cases where customers checked to see if there were any issues with the requirement definition and design documents and then sign off on them. With this platform, we anticipate there will also be a change in the way of looking at this process, as the documentation containing the changes that have been made will show what has changed in the completed system and what sort of test to run to repair what asset, which will allow customers to view everything at once.

### **Questioner I**

***Q1. You mentioned that the platform will achieve a 100-fold increase in productivity (what used to take three man-months will take four hours). Of this, how much time will humans spend being involved with the process? Also, you said that you will open up this platform as an***

*ecosystem, including to your partners. I think that it would be incredibly difficult to bring an existing system onto this platform. I believe that you are already exploring partnerships, but how much time do you believe it will take to do so?*

**AI (Kokubu):** Regarding your first question, humans are not involved in the scope of work that was shown in the demonstration. In order to create this situation, however, humans will need to spend a considerable amount of time cramming on AI-Ready Engineering. In addition, there will also be the time spent by humans auditing the logs and design document details after they have been generated by AI. Currently, we are continuing to develop this technology through trial and error, so it is difficult to provide you with an answer of the precise amount of time this takes. We are using this mechanism to handle the revision of the medical payment system in Japan for fiscal 2026, and are at the phase of testing how much time it will actually take, so I believe we will be able to share the information with you at another time.

**(Okada):** Regarding your second question about the preparation period, it took one year to achieve bringing an existing system onto the platform. This is due to repeating the trial-and-error process a multitude of times. As we explained, we did repeated fitting work to detect tacit knowledge. When we had AI to work on change requirements for hundreds of items from fiscal 2024, it gave wrong results. There are many different types; one example would be experts being unable to find what systems data set needs to be fixed, even after searching for it, and then, when looking for the cause of this, found that there was not a document to show the dependency of a certain file, and needed to create one. Because the experts needed to find the issue and perform the task manually, it took an incredibly large amount of time. We believe that, going forward, in order to disseminate our AI-Ready Engineering, we will need to use AI to optimize this preparation process as well. We will apply technology such as the multi-AI control that we mentioned. We would generate ideas about where information should be stored, so the human experts can then check like to make it so that AI will be able to explain what is missing when an issue is discovered and the AI's output and enter the information. In terms of our goal, it is our perception that if this preparation cannot be completed in approximately three months, it would have too great of a cost for business use. We believe our initial target will be customers who are able to determine how many years it will take to receive a return on their investment. We think that, for areas that struggle to handle frequent amendments in the law, such as government administration and healthcare, this mechanism should be implemented after it accumulates domain knowledge.