

Arm technology development accelerated by SoftBank acquisition

Ian Thornton (Head of Investor Relations, Arm Limited)

December 5, 2018

Good morning everybody and thank you very much indeed for coming to this presentation today. As you know, Arm was acquired by SoftBank Group (“SBG”) in September 2016, just over two years ago. And since then Arm has been able to significantly increase its investments in R&D. This morning I wanted to just give you some insight into how we've changed our approach in investing in new technologies and to give you some visibility of some of the first products and technologies and services that are coming out based on those investments.

[#: indicates slide number]

#2

Before I get into the details of the new technologies, though, I wanted to give a brief update on to how Arm is doing financially and talk you through our revenues and our cost and profits. I also wanted just to touch on an update on the Arm's joint venture, which we set up in Q1 this year. And then, we'll start going into some of the new investments that we've been doing post-acquisition by SBG, particularly in some of the new technologies around laptops, servers, and also around new services business for the Internet of Things.

#3-4

First, starting with the finances. On the royalty side, as you can see from this chart, we've been growing at about a consistent 11% CAGR over the last few years. Sometimes a little higher, sometimes a little lower, but about 10%, 11%. In the period prior to 2015, from 2010 to 2015, our royalty revenues grew much faster. But that's because Arm has good exposure into the smartphone market, and for the years prior to 2015, obviously the smartphone market grew very strongly. But in recent years, it started to slow down a bit. The smartphone market makes up about 60% of our royalty revenues. For that part of our business to be growing more slowly, then it's obviously harder to sustain double-digit growth.

We are starting to see revenues from the Internet of Things and networking equipment start to come through. We are starting to see design activity and design wins around servers, which is very encouraging as well. And we are also signing contracts today for some of our new technologies that I'll be talking about in a moment. And many of these have much higher royalty rates associated with them, which, again, gives me encouragement that we'll see higher royalty revenue growth in the future. But the new royalties from the new technologies probably won't start to come through until 2020 or 2021 and won't start becoming particularly

meaningful for a few years after that.

#5

On the non-royalty side of the business, which includes technology licensing, and also our software and services business, that has been growing or has been around about 180 million dollars per quarter for the last few years. Both Q1 and Q2 2018, though, have both been below that rate. That's primarily because of some of the activity involved with the setting up of the Arm China JV, which I'll talk more about later. But, primarily, there has been a bottleneck in the licensing pipeline. As the Arm China JV novates or transfers contracts from Arm Limited over to the joint venture and that's just kept the sales teams and legal teams busy and so they've been unable to do so many commercial deals.

As I mentioned earlier, on the royalty side, we are now signing licenses for some of the new technologies that we've been developing over the last couple of years. Some of those are at higher values than previous processor technologies. But again, we expect to see that coming through next year. Although we are signing licenses today, because we can't deliver the technology until 2019, some of the revenue associated with that, those new technologies can't be recognized until we actually deliver the product.

#6

So, what does this mean for profitability? Going back prior to acquisition, we had about a 50% operating margin that we'd managed to gradually increase over many years. But with the increase in investments in R&D, therefore, our costs have gone up. And as our costs have gone up, our profitability has come down. This year, we expect to be around about 10% of operating margin which I think is what I told you it was going to be two years ago. So, it's good. But we expect that over the next few years, it's going to be around about 5%, plus or minus a bit. I'm describing this as a black zero as we are trying to reinvest everything that we make back into the business. But we don't want to make a loss. Therefore, we need to keep a bit of a buffer. Therefore, it's effectively, it's just to stay into the black, hence a black zero.

There's one thing just to point out here. As part of our employee remuneration, every employee gets a small amount of money each year to spend on personal development. You can go in a training course or learn a language, or I use mine for gym membership. But because you can pull down on that at any time, the entire years' worth of that remuneration all gets recognized in one quarter. So, it's a years' worth in one quarter. We can't spread it out through the year because of the accounting rules. Last year, 2017, is when we introduced it for the first time. And it was introduced in Q3, so it all happened in Q3. This year, though, it's been recognized in Q2 and so you can see there is a bit of a spike because of that, the full-year impact all happening in one quarter. So, it's likely that next quarter is going to be lower than this quarter and in fact it might even be lower than actually Q3 2017 costs that is. Last year, this was around about 33 million dollars. This year, it was about 40 million dollars' worth of extra one-off costs.

#7-8

A little bit on China. Probably worth just starting with a bit of background. Arm has been licensing our technology into China for about 17 years now. And we've been, we believe very successful. We have over 150 licensees. We estimate that about 95% of all complex chips designed by Chinese companies have an Arm processor in them. And if I add up all of the chips, Arm-based chips that have been shipped by Chinese companies, it's around about 10 billion over the last 17 years and growing by about a billion every year. Which, to a certain extent, begs the question: if Arm has been so successful, why change? Why introduce a JV at all? And the reason why we've introduced this joint venture at this time is because China is changing.

The Chinese government now wants the semiconductor industry to be of national importance. They're investing 160 billion dollars in developing a domestic semiconductor industry. And within that, the Chinese government have identified certain market segments where all of the technology associated with that segment has to be developed in China. And so, these are particularly things around things like government servers, telecommunications networks, smart grid, and smart meters. Not consumer electronics but electronics required to run the basic infrastructure of China as a country. As a consequence, processors developed by Arm in United States or in Cambridge are not Chinese. So, they couldn't be used on these projects. If we had done nothing, then those projects are still going to go ahead but they'd go ahead without our technology, with somebody else's technology and maybe we'd be forcing Chinese companies to look for alternative processor designs. We did not wish to do that. A while ago, we set up our own engineering team in China. It's about 100 to 150 engineers. They are developing Arm processors to be used in these projects and programs. So, we have Chinese nationals, working in China, developing technology for Chinese projects. Unfortunately, that doesn't satisfy the requirements because it has to be as part of a Chinese-owned and controlled company.

In April, we set up Arm China, as a company. And in June, we sold 51% of Arm China to a consortium or syndicate of local private equity investors. And so, that, then makes this venture "Chinese enough" to sell its technology into China.

#9

In the first half of this year, we've been very busy, trying to novate those 150 customers who today, or in April, all had contracts with Arm Limited. Those contracts had to be novated across to Arm China. This is not difficult, this is not controversial. It's just very, very time-consuming. Basically, it meant that for the first couple of quarters, Arm's sales team and legal team, and our customer's legal teams were busy with contract novation, rather than signing commercial deals. At the start of Q1, in April, all 150 Chinese customers had contracts with Arm. By the end of Q1, 30% of those had been novated across to Arm China. By the end of Q2, 90% of contracts had been novated across to Arm China. The vast majority of them are

now completed. Of the remaining 10%, some Arm Limited is going to keep hold of because they have some bad debt associated with them and we don't want to transfer the bad debt to Arm China. Once that bad debt has been resolved, then we will novate those contracts. Hopefully this gives us, gives Arm China a good platform to now develop and build their business. However, Arm China still has other operational issues, just to sort out. Although we are hopeful that in Q3 their license revenues will be nearly back to normal, if not completely back to normal, but certainly, by the end of Q4, we would expect that Arm China will be fully operational and business as normal.

#10-12

Since being acquired by SBG, as I mentioned earlier, Arm has significantly increased the amount of investment in R&D. Where previously we had a 50% operating margin, we are now pretty much investing everything back into the business. This has enabled us to take a new approach to investing in processes and investing in technology. We are hopeful that this new approach should enable us to charge a higher license fee and also to achieve a higher royalty.

But first, I want to explain how we used to do things, and then I'll explain how we're going to do things in the future. These are our customers. We have about 550 of them. And from the outside, our customers may look all very similar. They all make chips. All chips look pretty much the same, little black boxes with legs and pins. But, actually, many of the companies within the semiconductor industry make chips for just one or two end markets. You have companies like Qualcomm that predominantly make chips for mobile; companies like Broadcom, that predominantly make chips for networking equipment; there are some companies, like Renesas, that make chips for multiple markets. Renesas is both embedded and on automotive. But by and large, one or two markets.

And when developing a processor, Arm talks to the thought leaders within a particular market segment. We try to understand how they think the future of technology is heading and what types of technologies they are going to want in the years to come. You have to remember, it takes Arm two to three years to build a processor, then takes our customers two to three years to build a chip. So, we have to stare a long way into the future. And so, therefore, we need to talk to the thought leaders, the companies who will, sort of be setting the pace for new technology. Following the conversations, we then come up with a specification for a new processor design and what should it do? How should it look? How should it operate? And then, we build a processor for a mobile phone. And this is our sort of mobile-first strategy. And this is because Arm has a 100% market share of the smartphone market and those royalties from that market generates 60% of our royalty revenues. So, we absolutely have to make sure that we secure future smartphone design wins, to secure Arm's future business. Then, we would take the processor that we had designed for a smartphone and then license it to everybody else as well. We would help our customers who are designing chips for other end markets on how to take a mobile phone processor and deploy it into a television or into

a server or into network equipment. But, basically, it's a smartphone processor going into a non-smartphone product.

#13-16

Now, you may think that, why would our customers be happy with this? Why would they not use something else or develop their own processor? Primarily, it's because these are very advanced processors. They are very expensive to go build. They're very, very difficult to build. And okay, it's a smartphone processor. But actually, it's good enough for most of these other markets. So, we've ended up with lots of design wins from mobile phone technology in non-mobile markets.

Going forwards, we are still talking to the thought leaders of the semiconductor industry. But we're taking their input and building some base technology. A set of functions and capabilities, which we then can optimize for individual end markets. So, we are now developing processes that are optimized for mobile but also now optimized for servers, optimized for automotive and safety-critical applications as well. Now, when we are licensing these technologies, we are licensing something that is more appropriate for the end market. Therefore, we are very hopeful that we can charge a higher fee for these because you are now licensing something that is a design that's specifically designed for the products that you want to go and build. Our goal here is to charge about an extra 30% higher fee for that.

So, all of our customers should now be happy. Well, not quite everybody. We also sell licenses into, technology into markets that are very small, very fragmented. It's too difficult to build processors for every single end market. If you're building a chip to go into a digital camera, we are not designing a processor specifically for you. But at least you now have more choice, in terms of the types of processors that you might choose to want to use. So, we now have multiple roadmaps of processors for different end markets. But if you look in the detail of these roadmaps, you'll find there's a lot of similarity, a lot of common features because they are all based on the same underpinning technology. In our roadmap for smartphones and laptops, you'll find a processor for 2019 that has a lot of similar features to the, so, I said it's Deimos in 2019 for mobile, areas for networking and servers. And then, in 2021, we've got some technology coming along for networking and there'll be a similar spec'd processor for mobile. But they are optimized for the particular end markets. They just share a lot of base technology in common. So, this makes it still very efficient to develop this. It probably costs us about an extra 20%, we estimate, to build processors optimized for multiple end markets rather than building just a mobile-first processor and then selling that processor to everybody. We still expect to be able to cover the cost of developing the technology in about the first 10 licenses. We will be charging higher license fees because of this. But what we are really trying to do here is to increase the royalty rates going forward. If we can increase the royalties by 30%, or 40%, or 50%, because of the recurring nature of royalties, that will help drive our profitability in the future and for many years to come.

#17-8

And now to dig into the detail a little bit for some of these end markets. This is Arm's new roadmap for laptops and computing. Arm processors have always been known for being very low power. Because our processors can support many cores, we've also been very good at multicore computing. If you want more performance, you can always add more cores. And that's helped us do very well. But we've always been less good or less good than Intel's PC chips at what's called single-threaded performance which is that burst of speed you need when opening up a new application or opening up a large PowerPoint file or large Excel spreadsheet, that's when you need very, very high single-threaded performance and we've always lagged the Intel roadmap here. Over the last few years, though, we've been catching up and I'm pretty sure that the processor in your smartphone today probably has the similar levels of performance as what's in a low-end laptop. In fact, if you have a new Apple iPhone, as the benchmarking tests have demonstrated that that has higher performance than a MacBook Air. So, we are already beginning to see how Arm processors have been catching up. But the latest generation of processors that we've been developing have significantly accelerated the performance and the capability of Arm in computing. This presentation, today, is being run off an Arm-based laptop, running Windows 10, running PowerPoint. It works just as well as my old laptop. In fact, I've had it now for a week and I can't find any difference between my old laptop and my new one, with one exception. This one has a two-day battery life. My old one had about a three- or four-hour battery life. In fact, I don't even bother turning this one off. I just put it in my bag and walk away, and it still has a two-day battery life. It's very surprising.

Going forwards, we will have even greater compute capability. The Cortex-A76 that we have developed this year and will be in chips next year has a greater single-threaded performance than a Core i5 from Intel. That's the chip that goes into about 80% of the world's laptops. So, that's now bringing Arm right up to the same level of performance than the majority of PC chips. With Microsoft having ported Windows 10 and all the office applications and making all of the apps on the Microsoft store all available on Arm, then we are hopeful that we'll start to see some significant design wins in PCs as well, going forwards. Just to compare, for the benefit of this roadmap, what's in your phone today with what we are developing right now and will be in computers in 2021, we are looking at around about 2.5 times increase in performance, which will be quite significant.

#19-20

Servers. Arm-based servers have been very controversial for many years. But we are now developing processors specifically for the server markets although these also can be used in networking, as well. And again, we are focusing very much on increasing the performance of the processor, so that we can better compete with Intel's server chips. You may have seen that there was an announcement by Amazon recently. This is Peter DeSantis, Vice President of Infrastructure at Amazon AWS. Last week, he stood up and announced a

new chip from Amazon. You may remember, I've been here now a few years, giving this presentation and over the last few years, I've talked about the progress that Amazon is making in developing their own chips for AWS. In 2013, they bought a semiconductor company, and in 2014 and 2016 announced new chips that Amazon had developed. The chip in 2016 was a chip that offloaded tasks from the main Xeon processor, and therefore made it more efficient. The chip they announced last week actually replaces the Xeon processor entirely. They also announced a business model around it. These are available today. They announced the pricing and the Arm-based AWS platform is 45% lower cost than the equivalent Intel chip, for you as a user to use. They're able to charge lower costs for a couple of reasons. Firstly, because they're building this themselves, they don't have to pay Intel's high prices. The Intel chips were about 1,000 dollars and they have about a 60% gross profit in there as well. So, not having to pay somebody else's profits immediately reduces the cost. And also, Amazon had built this chip specifically for the task they are trying to solve. Intel builds general-purpose chips that they sell to everybody. Amazon have built a chip that's specific for the AWS and is optimized for that purpose. We don't have the details of exactly how much Amazon are buying their chips for from TSMC or exactly what the power consumption is. But, if they are able to offer AWS services at 45% lower cost than they were previously, it must be much, much cheaper.

The cloud compute market is very competitive. Microsoft, and Google, and Amazon have been fighting this battle for many years. They all buy the same chips from Intel. They all build their own servers, but they all have them built by Foxconn. So, effectively, they all have the same cost base and as a consequence they all have pretty much the same pricing, except Amazon's now just lowered their prices by 45%. It's going to be very interesting to see how Google and Microsoft respond. Amazon's been working on this for five years. Does this mean it's going to take five years for Microsoft and Google to respond? I hope not. Qualcomm and Cavium are offering similar chips as merchant silicon. So, Microsoft can go and buy those now. In fact, I think when I was here last year, I told you how Microsoft was trialing the chips from Qualcomm and from Cavium and how they'd said that around about 50% of workloads in Azure can be, could be run using an Arm-based chip. Since then, they've done more trials, but they haven't yet rolled out a service yet. Hopefully, this might encourage them to start rolling out the service now. So, this is really big news for Arm. Plus, we are now providing to Amazon, chips specifically designed for servers. So, even more efficient ones. The processor in this new chip is a Cortex-A72. It's the same technology that's in many of your phones. As I mentioned earlier, in our mobile-first strategy, it was designed for a mobile phone. It's a 16-core processor, running at about 2.3 gigahertz. Most of your phones probably have an eight-core processor running at about 2.3 gigahertz. So, it's very similar technology to what's in a phone today. When we provide a processor optimized for services, we would expect it to be even more capable. Therefore, Amazon to be able to deploy this even more widely. Today, they've just started deploying it into the East Coast United States and into Europe, but we expect this technology to be deployed across AWS globally.

#21-22

I'd now like to talk a little bit about the new business that we've been investing in for, it's an Internet of Things software as a service business. I've been talking about this for about a year or so now, I think. And so, you may be familiar already with Pelion. It's a cloud-based, secure IoT platform and basically it's formed of five parts. The first part is, the first job we need to do is, we need to make sure that every single IoT chip on the planet has an Arm processor in it. That's our first goal. We are doing pretty well so far. We estimate that about 90% of all IoT devices chips so far have been Arm-based. So that's helpful. We need to make sure that all of those IoT devices are running our MBED operating system. That's a secure operating system and communications stack that enables an IoT device to connect to a cloud service. We need MBED OS in there, for Pelion to work.

The first thing is to make sure that all IoT devices are Arm-based. The second thing is connectivity management services. In Q1, this year, Arm acquired Stream Technologies and Stream is a connectivity management aggregator. They have agreements in place with around 600 operators be they cellular or satellite, and there's IoT-specific providers in there as well for technologies such as LoRa. Stream enables IoT devices to roam. So, it basically means that an IoT device can be installed anywhere in the world, and as long as it can see a cellular or satellite connection, then by Stream it can then connect to that network.

The third part is our Arm MBED cloud. That's a technology that I know I've spoken about before. That manages a device through its lifetime. It helps a new device be provisioned on to a network for the first time. It enables a device to be configured. If there is a need for a software update, if your connected streetlight has a security flaw and it needs a software patch, then that's the sort of thing that the device management software will do. Also, at the end of life, when your sensor or connected device needs to be taken down and thrown away, you don't want somebody to be able to take that out of the bin and to be able to reconnect it to your network. So, at the end of life of an IoT device, you need to be able to remove it from your network and remove its permissions to be able to reconnect to your network.

The fourth part is data management. In Q2, Arm acquired a company called Treasure Data, and Treasure takes unstructured data from many, many different sources and they sort it, structure it, store it, allow it to be combined with other data and then pass it on to the owner of that data for data analytics to be run. IoT technology is particularly unstructured. If you imagine a factory with 100 robots, all working autonomously, each robot might have multiple cameras, multiple sensors, multiple motors, actuators. All of the data from those devices will come back in and needs to be associated with the right robot. So, all the information for the first robot needs to be associated together. All the data for the second one needs to be associated together. You don't want to muddle up your data. This structuring of data is the primary thing that Treasure Data does.

The last part is, we created some interfaces for other companies to base their products and services onto Pelion. Pelion is very suitable for things like asset tracking, fleet

management, these sorts of IoT dependent technologies. We are providing an interface to enable third party companies to build their products based on Pelion.

#23

The key question, then, is all sounds great, but how do you make money on this? Here is an Internet of Things device, a connected washing machine. Maybe it's from Samsung, or Hotpoint, or someone. They make washing machines that need to be shipped all over the world. They need their washing machines to be able to roam. Inside this washing machine, you're going to have some technology that's going to enable it to connect to a network. So, there'll be a smart chip with some wireless technology in there. Hopefully, that smart chip will have some Arm processors in it, to both control the washing machine and also control the communications channel as well. Arm has developed a technology, an embedded SIM card technology called iSIM. It works a bit like a SIM card or if you are familiar with eSIM, it's a type of eSIM. But this can get integrated into the microcontroller itself. It means that the SIM card can go into the same chip as the main controller. Hopefully, this will be running our MBED operating system as well. So, how do we make money from this? Well, Arm obviously gets a royalty per chip that contains our technology. For something like this, we would typically get a 1% or 2% royalty. These are very cheap chips. So maybe they're only going to be two dollars or three dollars each. Maybe there's two chips, say, in a connected washing machine. So, maybe we are getting 5 to 10 cents per washing machine, for the semiconductor IP.

#24

That's our traditional business model. It doesn't sound very much but, if there are millions of washing machines that are connected, eventually we expect there to be billions of IoT devices. So, all those 10 cents should add up to be a very nice revenue stream over time. But this is one washing machine. I'm sure that Samsung want to sell a lot more than one. Eventually, you end up with lots of washing machines, all connected to a network. And when the consumer takes a washing machine home, plugs it in for the first time and switches it on, our view is that you could have that connect to your Wi-Fi. But the problem with a washing machine is, it doesn't have a keyboard, it doesn't necessarily have a screen. So, entering your Wi-Fi password is always challenging for something like a washing machine. It's much better for it to be able to do that for itself. With a washing machine with a cellular connection, and a SIM card, that, as long as it can see a 5G or LTE network, then you can connect. So, when this washing machine gets turned on, if it can see a cellular connection, it will connect to the operator, the operator will come through to connect to Arm and basically say, does this have permission to come onto the network. We'll say yes and then that will then, with the embedded SIM technology, the operator can then configure the SIM card for that network. Embedded within the price of the washing machine was a fee, and then that fee gets passed from the OEM, so in this case, from Samsung via Arm, to the operator to pay for the lifetime

data costs of the washing machine. Washing machines don't generate that much data. Maybe, let's call it 10 dollars and maybe Arm takes a small slice of that, of that fee as it gets passed through to the operator. You can see this today in some products already. The Amazon Kindle, when you buy one of those, has about a 10-dollar fee that connect, that pays for the lifetime data cost of the Kindle to enable you to download books over a cellular network. And we are extending that concept for all forms of IoT devices.

So, that's your setup. You now have a connected IoT device. And then, what Arm will also do is provide the software updates, configuration updates, changes, and also extract the data from all of these washing machines to pass back to Samsung. In return for that, we charge a monthly or annual fee per device. And the magnitude of that fee will very much change by volume. But we are anticipating that for companies with a small number of connected devices, maybe we get something like 10 cents per device per year. For somebody with large numbers of devices, maybe it's one or two cents per year. But it's a recurring fee that gets paid every single year for every device in that network that we manage.

#25

Because we acquired a couple of companies recently, we already have a business with revenues as well. And so today we are managing around about 30 petabytes' worth of customer data as part of Treasure Data. And they're adding to that about 2 million bytes of data every single second. So, it's a rapidly growing amount of data. We have 800 customers so far across our businesses. About 140 companies developing technologies to fit onto the interface I was mentioning earlier. And between them, they have somewhere in the region of 350,000 software engineers working on that. We recently had a deal put in place with KEPCO. KEPCO provides nearly all of South Korea's electricity. Apparently, South Korea has amongst the cheapest electricity in the world, not because of South Korea has large amounts of cheap electricity to produce but because it's heavily, heavily subsidized. And, as a consequence of that, lots of Koreans leave the lights on at night. As those subsidies are going to be removed, so, the South Korean government wants consumers to become more efficient. And so, KEPCO is rolling out smart meters and to every home and business in South Korea. And we are involved in the first part of that trial, or the first part of that rollout. And hopefully, if this is successful, we'll be involved with the subsequent rollout as well.

#26

Just to finish off, then, we can hand over to your questions, Arm has been investing, increasing our investments over the last few years resulting in our 50% operating margin going to very close to zero. Over the next few years, we expect that margin to remain in the sort of 0 to 10% range before we start to see new revenues from the new technologies that we've been developing start to contribute meaningfully. Then we should expect to see profitability returning in the sort of 2021, 2022 timeframe. This extra investment in R&D is enabling us to build a broader range of products than we've been able to do before. Because

they're more optimized for end markets, it's enabling us to charge more for these products, as well. It's also enabling us to develop new technology to go into brand-new markets that we've not been able to address before. So, multiple benefits from that. The combination of all these new technologies, we believe, are going to be driving our future revenue and profit growth in the years to come.

With that, I'll hand over to your questions.

Q&A

Q1: About IoT business. I think in the past, you were saying that the penetration or shares of the IoT was 75 or potentially a little bit lower. Now, it's 90%. Could you talk about how you reached the number?

A1: I think the 75% was for, sort of more general embedded processes. The 90% is for IoT which could be managed, have a managed service associated with it and therefore is suitable for the Pelion platform. So, the 90% and 75% are slightly different numbers, between slightly different markets. The 90% is for IoT that will have a managed service associated with it.

Q2: About the business model of IoT. My understanding is that there is not much of the license revenue for the IoT businesses. If that's correct, it's a more leveraged model compared to traditional model that you show the chart. Now, what would you expect in terms of profitability in the long run for this IoT business? Will that be better or worse than the current business model? Say, such as for the smartphone?

A2: If I look at the two areas of revenue, revenues generated from the sale of chips and related technology, the license fee is very low for that. But we tend to sell lots of licenses. Because our costs today are all developing future technology and our license revenue is all for technologies that we've developed in the past, effectively, every license is a 100% margin. Once you've paid for the cost to develop the processor, then every new dollar is a new dollar of pure profit. Because we sign lots of licenses for our technologies that go into IoT devices, even though the individual license fee is quite low, it still more than pays for the cost to develop the processor. In fact, we've announced a new licensing model called DesignStart Pro, where we've actually removed the upfront fee completely for our two most popular processors and we've been able to do that because of those two processors that have now signed 200 licenses. We probably covered the cost to build in the first 10. It's now fabulously profitable, to the point where we don't need to charge another license fee. And what we've seen is, by allowing the technology to be accessed for free, we've seen a huge explosion in the number of companies developing IoT products based on that process because it now has no upfront fee. There is no barrier to me accessing it. You all could go home tonight and

download it from our website. There's no fee. There is, however, a royalty, even if there is no upfront fee. Our hope is that that will help drive new design, more design wins and therefore more royalties in the future.

The IoT market is probably going to be smaller than the smartphone market, in terms of dollar values. Therefore, if your question is not necessarily on profitability, but on revenues and contribution, it'll probably be smaller. There aren't many markets as big as the mobile phone market, or the smartphone market. So, they'll all be smaller than that.

On the second part of the business, around the Pelion services business, our expectation is that over time this could grow to be the same size as our business today, and our business today is 2 billion dollars. We can get there by, not by making any bold assumptions about the market size or our market share. If IoT happened, we think it will be billions upon billions of devices. If there is a management fee of one or two cents per device, we don't need a big share to get to 2 billion dollars' worth of revenue. The math is very compelling. We do expect there to be more competition. We do expect this to be more like a typical software as a service business. If I look at Salesforce and Norton and other companies that develop a software as a service business, most of them have a 20% to 30% operating margin. So, my best guess is, at some point, a 2 billion dollar business with a 20% to 30% operating margin. So, 500 million dollars' worth of profit per year. That seems like a pretty good starting point. What I can't tell you is when that's going to happen. One of the interesting things we found around the Internet of Things market is it's more like a things market than an Internet market. Internet services generally, if you have a new idea, you can have a go and develop it. You can trial it with 10,000 users and then you can roll it out globally in a few months. The Internet market is really, really fast. The Internet of Things market doesn't move that quickly because the things have to be manufactured, the things have to be distributed, the things have to be installed and if you are installing connected streetlamps, then that means somebody has to get on a ladder and climb up to the top of the lamp post and screw it together. It's a very slow-moving market, we've found. Even though everyone wants it to move fast, because it involves physical product being designed and shipped and manufactured and whatnot, it just takes a long time to distribute. Therefore, it's a slower market than other Software as a Service businesses.

Q3: I just have one follow-up question. How large is the addressable market, in terms of percentage of this Pelion platform business model out of the total IoT market?

A3: I think that's going to change over time. I mean, most IoT devices today aren't managed. I mean maybe, your Google Nest is, your Amazon Alexa is, your car might be able to connect to the Internet. But it probably isn't a managed device. Over time, maybe it's 100%. Right now, maybe it's only 50%.

Q4: For the sake of specification and customization for the processor for each end-user, you

suggested 10% to 20% increase of the cost for longer-term. And your assumption is that cost will be well covered by processor license fee increase for higher rate than maybe, more than 10%, 20%. And how, actually would it be defined; I want to know the calculation formula or driver for that take rate.

A4: I mean our working principle has always been we want to cover the costs to build something in about the first 10 licenses. And effectively, the offer to our customers is, we can offer it to you at a fraction of the cost that it would cost for you to build it yourself. Now our customers don't exactly know how much it costs us to go build something. But what they do know is that now they're getting something that is far more optimized than anything they've had before. So, they know they are going to be paying more. And so, our goal is to charge them 30% more, even if it's only cost us 10% to 20% more. We have started signing licenses for the first generation of new processors in this area. The signs are very encouraging that we are going to achieve that. We may, in some cases, achieve more and, certainly, the focus is trying to achieve a higher royalty, if we can. Even if the license fee ends up being the same, if I charge a 50% increase in royalty, I'll be very happy with that. I think it's too early to say that everyone's going to end up paying more, though. We haven't signed enough contracts to be able to declare it as a trend. But the first contracts with the first leap partners are very much in line with a, on average, a 30% increase in pricing.

Q5: On page nine of your presentation, regarding Arm China. The contract transfer has been completed 90% and the remaining 10% is still there because of bad debts and that's why it's not transferred to Arm China. Does that mean that these customers are not paying? And if that's the case, you could also say that 100% of the transfer has been completed for those that can be transferred. And another question is some operational challenge. By the end of fourth quarter, you expect the operation to normalize. Can you please elaborate on these operational challenges? Also, is there any risk that it will not be addressed by the end of the fourth quarter?

A5: Of the remaining 10%, so, of the 150, around about 12, I think, are not transferred. Some have bad debt. Some are just awkward companies who are being stubborn and not wanting to sign their novation. But I think there's maybe one or two companies that are still in process. And so, hopefully, should be done soon, very soon. The vast majority of the ones that we care about have now been novated, have been transferred. On the other operational issues, it's things like ... Arm China needs to hire a CFO. We need to set this up as a business. And there are some things that were missing when we set it up so, they need to hire a CFO, they need to hire a new head of sales. It's just a few bits and pieces that they need to put in place which we didn't quite manage to get done by April. Some of their IT systems are still dependent upon Cambridge because we weren't able to transfer everything over by April. And we are most of the way there, doing it. Is there a risk? There is a risk that some of these

details aren't done by the end of the year. I would be really disappointed, though if, in Q1, I was still saying that the reason why we've missed the license revenues is because of Arm China. That would be very disappointing. I'm hoping that from Q3, not only are they back to normal operation, or from a financial point of view, but we start to see some of the licenses that weren't signed in H1 signed in H2. In some of those, the contracts that we expected to sign in H1 probably have gone away. You know, the customer has changed their plan, they've used a competitor's technology, they've gone bust, some of them. But, some of those companies still want to license Arm technology from H1. And so, that should end up going into the second half of 2018. Not only do I want to see Arm China hit their budget, I want them to exceed their budget because they missed Q1 and Q2. I feel owed.

Q6: On page 25. I have a question about these 30 petabytes and these are really big numbers. Compared to the numbers in the past, how much growth have you achieved? Also, as a business model, what does your topline depend on? Is it number of customers or is it dependent on the size of the data that you manage?

A6: Good question. We acquired Treasure Data and Stream Technologies in the first half of this year. Both of those were operational businesses. Treasure Data has been operational for nearly 10 years. Stream has been going for 20 years. So, they have built up real businesses and some of this data is based on businesses that are many, many years old. The 30 petabytes of customer data is from Treasure Data. They not only manage IoT data, but they also manage consumer data. So, companies like Muji and Subaru are their customers. Fujitsu as well. Treasure Data is a half-Japanese, half-Silicon Valley company. So, they have a lot of Japanese customers. What they do is, they take data from websites, you know, the companies' websites, they take data from in-store transactions and they combine that data together to associate it with an individual consumer, a consumer journey so, they can tell when someone's gone into the website and looked at a product. Maybe they've gone into store to check the price, and then maybe they bought it somewhere else. They try and associate all that data together, so that Muji can track the consumer as they make their purchases over weeks, months, years. And so much of that 30 petabytes of data is traditional web-generated, consumer-generated data. But, if I carry on with the Muji example, they are now also installing beacons in their shops, so as you walk in with a smartphone in your pocket, if you got their app running as well, they'll send you sales information and vouchers. But, at the same time, they'll track you through their stores across Japan, or even globally. They really can track their customer's journey.

And then, that can get combined with other data as well, that they have, in order to really get a lot of insight into their consumers. Actually, the Treasure Data system can also interface into other databases. So, Salesforce, and SAP and other CRM systems. You can combine multiple data inputs. That's all associated with particular consumers or stores or whatever. That gets, then structured and then passed back to the OEM, or the service provider, or Muji,

or Subaru for them to go and do their own data analytics. Just to be clear, we are not doing the data analysis part. This is structuring of data. We are not here to compete with IBM Watson, or Microsoft Azure. This is very much focusing on the bits of the data, rather than the analysis of the data. And, in terms of the, sort of the scale of the business, for the quarter just reported, it was about 12 million dollars' worth of revenue from ISG. About 10 of that was from Treasure, 1 from Stream, and then 1 from MBED Cloud. And so, on an annualized basis, we are probably looking at about 50 million dollars per year. Next year, we expect that to grow. But I'm not really here to tell you how much that 50 million is going to grow by. It's too small of a number at the moment to be interesting to our overall business. My intention is that when ISG, when Pelion becomes about 10% of revenues, that's when I'll look about breaking that out as a separate business. Probably not in 2019. But, hopefully, by the time we get to 2020, ISG should be big enough for us to actually start to report separately.

Q7: In that sense, in the end, size of the data that you have collected, size of the data that you manage is the most important factor. Is that correct?

A7: It's more customers, more connected devices, more data. Right now, because we have just acquired two businesses, we have three business models. It's not yet consistent. But we expect, in the future, we are going to get two revenue streams from Pelion. One is for every new device connected to the network. So, every time that there's a new washing machine switched on, that lifetime connection fee, we would like to get a slice of that. Maybe, say, if it's 10 dollars, I'd like a dollar, please. And then, if we can also get a management fee, a monthly or annual management fee which would be related to the amount of data and related to the number of software updates being done, the number of configuration changes. So, how much of a service are we providing. But clearly, one of the big services we are providing is the collection and structuring of the data. So, yes. More devices, more data, please.

Q8: About the timeline of the forecast of the margin. For the time being, 5% to 10%, this margin, for short-term. But when is it going to hit the bottom, and when is it going to become 10%? And when are you going to enjoy that full effect the new approach? At that time, is the margin rate going to be exceeding 50%, which was the conventional margin rate?

A8: For the next few years, we would expect the operating margin to be around about 5% plus or minus. I'm saying plus or minus because as you saw from our license revenues, those are very volatile. And so, if we have a quarter when the license revenue is high, maybe it's a 10% margin. The next quarter, it'll probably be lower and, therefore, it might be a 0% margin. And that's why we are aiming for what I described as being a black zero, so that we've got a little bit of a buffer, so that if we are in the low quarter for licensing, we don't go into red. Our CEO is very keen not to have to explain to Masa why we are making a loss. And so, therefore, we're going to keep a little bit of a buffer. The goal is to keep it at that level for the next few

years. And it's only when the new revenues from these new technologies start to come through in a meaningful way that I would expect us to start to see any profits coming through to the bottom line. My best guess at the moment is, that's going to be around about 2021 to 2022. And I think it's interesting to note that Masa has been talking about a potential IPO for Arm in five years' time, which is 2023. I know it's very fashionable these days to IPO a lossmaking business. But I think it would be easier to IPO a company that is making some kind of profit. And so, therefore, that's probably why he's picking 2023 as a potential date for the IPO.

Q9: About Pelion and data management. OEM companies would like to do data management directly, as well, I think. Is that a risk that OEM makers directly does data management? And what's the strength of Arm, or advantage of Arm against the OEM companies doing data management?

A9: I guess we need to distinguish between data management and data analysis, and when does the collection of the data become the analysis of the data. So, we're not proposing that we'll be doing the analytics, that we'll have the AI or machine learning algorithms that are going to give the OEMs the insight into all the data that they're collecting about their products, and their consumers, and their robots, and their streetlights. We're not going to be that far up the value chain. We need to be careful about competition with other cloud companies. We know that Amazon, and Microsoft, and Google will be offering IoT services, and that they will be building those IoT services on their expertise around data analysis. Google collects all the data, all the clicks, and all the searches that we make, and they put that into their AI algorithms. We are not expecting to compete with that.

But what we've found from the cloud companies is that they get very confused when they start to get close to the devices. The vast majority of Internet data today comes from consumers typing into keyboards, clicking on links, it comes from PCs, and servers, and mobile phones, things which look very similar from the cloud. As you get closer to IoT devices, the diversity, the number of different types of devices goes up significantly. We have had conversations with all of the cloud companies where they say to us, can't you make them look the same? We want them to all have the same interface, to work in the same way. And the answer is, no, you can't. You can't make a streetlight look the same as a washing machine, look the same as a sensor in a field, look the same as a robot, look the same as an autonomous car. They will look different. They will look differently. They will generate data in different ways. So, what we want to solve is that connection to the thing because we understand the things better than anybody else does, in our view. And we don't want to go compete with the big cloud companies, because that's what they know well. But, if we can find a way of combining our ability to connect and collect the data from the things and then pass that on to the cloud companies to do the analysis, then I think we might have a winner. But, it's very early days. We don't know, and we don't know how big this market will be. We

don't know where the greatest revenues will be. We don't know where the greatest profits will be. And actually, we don't want to be in the area where the greatest profits and greatest revenues are. We've been successful by being able to take a piece of technology and commoditize it, to make it expensive to build, but very cheap to buy. If everybody makes lots of profit on this, then everyone will want to differentiate and invest in this area. We wanted to remove the desire to invest in this area and have them invest in areas where their differentiation can generate profits for them. We're trying to aim for a particular part of the market and hope to be successful there.