

WHAT DOES IBM'S NEW TAILORED FIT PRICING FOR MAINFRAME MEAN FOR CUSTOMERS?



In this Run and Reinvent podcast I chat with Steve Solomon, product manager for mainframe cost optimization here at BMC, about IBM's new pricing structure for its mainframe offerings. Below is a condensed transcript of our conversation.

Stephanie Grubbs: Talk to us a little bit about the value you've seen the mainframe provide.



Steve Solomon: The mainframe is one of the most trusted platforms out there, and it's really critical and an integrated part of the digital economy. And folks don't really realize it when they're using a mainframe. So, I gathered a couple of statistics for today's podcast. Did you know that more than 30 billion – that's with a B – transactions are run on the mainframe every day? That's 10 times more than Google searches. So, it's pretty amazing. 92 of the top 100 banks in the world rely on the mainframe. 80 million ATM transactions are processed each day, worth nearly 5 billion dollars a day. If you're moving money, you want to have a mainframe behind you. 87 percent of all credit card transactions – that's nearly 8 trillion dollars in payments a year – are all managed by mainframes.

The mainframe's here to stay. We don't hear about it a lot with some of the newer platforms, but we continue to invest in it. Companies like BMC and IBM – we invest millions of dollars a year in new R&D research to keep it as a core computing platform and keep modernizing it.

Stephanie: You mentioned innovation on the mainframe. Can you tell me a little bit more about

what BMC is doing in the innovation space?

Steve: BMC continues to invest in R&D and launching and creating new products, not just maintaining the ones that we've had for many years. In 2018, we launched a new product line called AMI, or automated mainframe intelligence. You may have heard of some of the catchy terms in the industry today, like AIOps, Artificial Intelligence Ops. In fact, that's what we're investing in for the mainframe, and we have all the latest trends. We have AIOps, [DevOps](#) for development and operations, and even agile development on the mainframe.

And the mainframe is just like a big server these days in your environment; it has a lot of the same functionalities. Speaking to BMC AMI, or automated mainframe intelligence, we use automation and machine learning and predictive analysis to deliver on the promise of a self-managing mainframe. Now, this is really critical because we want to take all those alerts and do all the automation and reduce the cost that it takes people to run the mainframe even more. So, it provides the ability to manage, optimize, and heal the mainframe and IT systems, and oftentimes, give us advance warnings with the mainframe and all of the things that are connected to the mainframe in your environment.

And this is especially valuable to our customers because many of them are hiring millennials or even Gen-Z programmers and technicians to run their mainframes. I'm in the class of the 30-plus year experience, and we can see retirement coming. So, we're more invested in trying to build the tools to help the next generation of mainframers be productive. And building that automation, we're having some great successes. And our goal at BMC is to have a mainframe that can predict issues, heal problems, and in a secure and cost-efficient way.

Stephanie: I heard recently that IBM is messaging the Z platform as part of this new hybrid cloud strategy. What are your thoughts on this?

Steve: Our customers are focused on [digital transformation](#). And we at BMC call it "run and reinvent." We're running our business and we're reinventing it as we go, and as a matter of fact, we continue to invest and reinvent with our customers to help them drive them forward. So, part of this digital transformation – there's mainframes, there's distributed, there's cloud, and there's hybrid clouds. And depending on what solution is right for our customers, we go support that and have the mainframe be an integral part of that digital transformation strategy. So, in order for us to help companies move faster, be more agile, be more dynamic, drive their business forward, we want to make sure that that mainframe is a core part of those hybrid multi-cloud strategies. So, it's very critical that workloads that span clouds, hybrid clouds, on-prem, distributed, all work together in a safe, secure, and cost-effective way.

As a matter of fact, when we think of the mainframe, we often hear rumors that the mainframe is dead. That is absolutely not true. As a matter of fact, BMC – we do a survey of our customer base every year, and we go through and we find all of the mainframe users out there, and we ask them questions about what are they doing, what are the latest trends, who's investing in the platform.

And the yet to – unreleased survey for 2019, 62 percent of our respondents came back and said they were continuing to invest in the mainframe, and they're planning on expanding their MIPS or MSU consumption. So, MIPS are millions of instructions per second, and MSUs are millions of service units. They're both ways that we measure capacity. So, basically, our customers are telling us they're continuing the mainframe, and they're growing their mainframe capacity. And modernization is important, so as we transform, they can be successful.

To help customers do this, what we needed to do was align pricing with IBM mainframe software with some of the hybrid clouds. And in order to do that, IBM, just recently this month, has given us two new software pricing models to help customers build software models that match this hybrid cloud. And so, it's a strategy that BMC is working on also, and we're also working on our BMC AMI efforts to tailor to that. So, I think it's good to talk about it as part of a hybrid cloud strategy because they don't have to pick between cloud, premise, mainframe. Everything should work together to make sure that it's enabling the business to go faster.

Stephanie: Explain a little more about what these new pricing models are that IBM announced?

Steve: A lot of it has to do with software. And software pricing – however, you can combine hardware and software in an AKA cloud-like model if you so choose. So, IBM has revamped their mainframe software pricing models under a nomenclature called [Tailor-Fit Pricing for IBM Z](#). IBM customizes what software you have from IBM and is deployed on your mainframes. And many independent software vendors, or ISVs, like BMC, then work with customers to help build and support those same models.

As a matter of fact, John McKinney, our VP of strategy for BMC, was at a recent Gartner show this month with IBM and several of our partner ISVs announcing this new change to tailor-fit pricing for IBM Z. There are two new pricing options that are delivered as part of this. One is called enterprise consumption solution, and the second is the enterprise capacity solution. Those two are added to two previously added solutions known as that development and test solution and the new application solution. So, customers have a choice, and there are different pricing options and different cost saving options, depending on what models they go to.

Stephanie: What is BMC's view on the new released of Tailored-Fit Pricing?

Steve: We have to go through these cost optimization strategies to continue to make the mainframe viable, cost-effective, and successful as part of the hybrid cloud. So, needless to say, we're glad our customers have more options, and we're looking forward to working with them to ensure our products and software continue to meet their needs.

Stephanie: Can you walk us through each of the new solutions that came out under Tailored-Fit Pricing?

Steve: It's a bit complex, and I apologize if I confuse folks, but I'll try to make it simple. So, enterprise consumption allows you to pay for the millions service units. Million service units is a measure of capacity consumed in an hour, and you pay for that and you contract for how many service units you wanna buy, and you pay for that as you consume them. It's a new model.

Previous, in the first 20 years or so of the mainframe, we used to license everything on capacity models, and all of software was based on how big your machine is, even if you didn't use it. Which was always kind of a complex model, because what would happen is, as a mainframe developer, I would be working along, and we'd get a new mainframe, we would have lots of power, lots of extra capacity, everything would be zooming along. And mainframes are big, powerful boxes. But as the years went by, we'd at new work, we'd release new things, and we would run into that high capacity. So, if their capacity would start hitting above 90 percent, 100 percent – you can run a mainframe at 100 percent for months. It runs just fine. But then we would have to wait for the next upgrade, because we'd get slower.

So, IBM introduced something called sub-capacity pricing, where you could license less than the full capacity of your machine so you could have that growth and it wouldn't penalize your folks and slow

your system down. Well, that four-hour rolling average, is what it's known as, for sub-capacity, was a bit difficult to manage, because you had to look at what your current consumption was and they smoothed it by using an average of four hours. So, there was a bit of complexity in how things actually got built.

So, IBM came out and said, okay, we're going to eliminate that complexity, and we're going to tell you what we're going to do – is we will just charge you by when the MSUs are consumed. And that takes the peak, or the four-hour rolling peak, out of the billing situation. So, it reduces the complexity. It also changes some of the models around, so if you guys are listening to that and you're considering your IBM software, please be careful and look at all the different options and evaluate them, because the [MLC model](#), as we call it, monthly license charge model, is still available to you and it still may be right for you.

But customers that are growing, that 62 percent of our customer base that said they're committing to growth and expanding their MSU consumption, these are great programs for them. Now, remember, a million service units is just how we measure capacity, and you determine with IBM how many you want to contract for, and you'll contract for those and you can bill them at any time. Companies that have spikes in workload — so, let's say that the last two days of the month when you're running big billing cycles, you use twice the capacity. It used to cost them more under the four-hour rolling peak because it would have this sudden peak for two days, and then they'd have this excess capacity the rest of the time. Customers like that are perfect for this MSU consumption because it doesn't matter when in the month that you use your MSUs.

It is a little bit different in that you contract for a number of MSUs to consume during the year and you pay for one-twelfth of those contracted MSUs monthly. So, it's also very easy to convert from your monthly budgeted process. If you've set up a budget and you've done a three-year contract with IBM, you have a good idea of what your current charges are under the MLC model. When you flip to enterprise consumption, you want to continue to match those budgets that you've had planned out, and that's a good way to move that across.

All right. It's important to note that this is a usage-based pricing model. Therefore, you're paying for what you consume. It's a big change in the mainframe, because now every hour of the month matters because everything you consume. It used to be just in the peak, you only focus on those four hours in the peak and try to reduce workloads then to save money. Now, as we move, we have to be aware of everything that runs on the mainframe all of the time. With the model, there are some great benefits; there are other things we have to focus on and take care of for operational excellence. On the flip side of the coin of the changes, it means controlling consumption is really, really critical for operational excellence, not just at the peak.

Stephanie: BMC has a workload capping solution that's worked really well for our customers. Do you want to dive into that a little bit, this messaging from IBM about no longer needing capping?

Steve: This one is going to be difficult even for long time mainframers, unless you're in the operations space. So, how capping works on a mainframe is, if you're running in a sub-capacity model, which means you're running at less than full capacity, you can set different caps. And remember, under the MLC model, you paid for your peak usage. So, what we would do with a capping utility — we would set what the customer wants for their peak and when it hit the peak, we would put in a cap, and what the mainframe does — it's very unique — it will slow down your non-critical workloads.

So, with a capping utility, you can say this workload is important and we don't want to slow it down

at all. These other workloads are less critical, and when the mainframe dispatches cycles, it will slow them down. It doesn't stop them, it just doesn't dispatch as many cycles to them, and it smooths out the peaks so your low-priority workload runs longer, but it uses the same amount of MSUs. However, since we cut the peak, customers can save hundreds of thousands or millions of dollars a year by smoothing their peaks. So, it's kind of interesting on the capping.

IBM has come out and said, "Okay, with enterprise consumption, it doesn't matter when you use those MSUs. Or with the capacity, it doesn't matter."

Well, there are other things that capping does, and it depends on the capping solution that you have. You may have one from us, you may have one from some of the other ISVs that are out there. But before you come out and say, "Oh, we don't need capping anymore," please take a look at your environment and understand the other benefits your tool is doing. The BMC product looks at a lot of priorities and can move consumption around, depending on what the priorities are.

Secondly, there are a lot of ISV contracts — that's the independent software vendors like BMC — that have based their software contracts on the four-hour rolling advertising window to match your IBM contract. So, if you get a message from your procurement people and it says, "Oh, we've got this new model, it doesn't matter," don't just go turn it off. Please make sure you do an investigation. Talk to me, talk to your BMC account team, and let us make sure it's not going to have an adverse impact on you, either your ISV agreements or your system or the other performance items that are impacted by these capping utilities.

One of the reasons IBM came out with these new models is we believe that the independent vendors have been very, very successful at helping customers save money and now IBM has changed and said, "Hey, we'll support that. Let's do it so that we don't have to cap things and let things run faster at full potential, and not when they smooth the MSUs." So, it's a great way to do it. But we'd like to say that you need to be somewhat conservative with it. Just don't think I'm going to pull the plug on the tools. There are other things besides just smoothing the peak that are important in your environment.