

IBM Podcast

[MUSIC]

MATHENY: Welcome to this IBM podcast, What's New in IBM Rational Cloud Computing Solutions. I'm Angelique Matheny with IBM. Achieving effective software delivery while keeping risk and cost in check is a difficult balancing act.

With Cloud computing, you can reduce the high deployment cost of delivering software while maintaining quality, improve control and governance while speeding time to market and enhance the flexibility of software development and delivery processes.

And in today's podcast, we'll discuss new capabilities from IBM Rational around cloud computing, including Rational Software Delivery Services for Cloud Computing, Virtual Load Testing and Modeling and IBM Cloud and Deployment Planning and Automation.

Joining me for today's discussion is Steve Weaver, IBM Rational Marketing Manager, Cloud Computing Solutions; and, Jamie Knowles, Product Manager, Rational Cloud Solutions. Hi Jamie and Steve. Welcome to the podcast, thanks for joining us today. And Steve, I'll turn it over to you.

WEAVER: Okay, thanks, Angelique. Thanks everybody, for listening. So as you may be aware, over the past year, Rational has begun to roll out a number of new capabilities focused on leveraging the cloud to transform software delivery.

As Angelique mentioned, in March, we announced Rational Software Delivery Services for Cloud Computing, which gives organizations the ability to run several Rational products including Team Concert, Quality Manager, Asset Manager, Build Forge and Requirements Composer on the cloud for the first time.

With support for the IBM business and development test cloud and the IBM CloudBurst appliance for private cloud, and for IBM Smart Business Development and Test on the IBM cloud, Rational Software Delivery Services reduces the infrastructure, maintenance and configuration costs associated with application lifecycle management.

And we recently announced that we're expanding our cloud portfolio to include two new offerings: load testing on the IBM cloud and deployment planning and automation, both of which help to flush out our strategy around helping our customers to plan, build and deliver cloud services. So Jamie, let's talk a little bit first about load testing. What are some of the challenges associated with running

large-scale performance tests?

KNOWLES: Well, the first issue is in order to apply a load to an application, you need raw processing power to simulate the actions of a user. So, to simulate a small number of users, let's say one or two hundred, we can get away with a fairly low-cost machine to generate that load.

But if you want to generate larger loads, then you need more muscle, so you need to buy more machines. You then need to set those machines up with the required software, and then get them connected up to the test software, the central console or workbench application.

So for a very large load, let's say you want to run, you want to test your tax return system and find out what happens the day before everybody's tax returns have got to be in, you might want to simulate 10 or maybe 100,000 users hitting the application at the same time.

So for this, you might need many more machines -- you might need 10 or 100 machines just sitting there generating that load. So that represents in itself quite an acquisition cost, and then cost in time to set and configure those machines, connect them up. And then, ongoing costs to maintain those machines, which you might only use once in a while when you're generating these large tests.

So the first issue is around acquiring that raw processing power to simulate the actions of the user; and the second issue is around the licensing. So most load testing software will charge a licensing fee based on the number of users that you want to simulate.

So if you're generating very large loads, the costs can get very high. And so for those large loads that we just mentioned, it might come to some tens or hundreds of thousands of dollars for software licenses that again, that you're going to use only once in a while. So, generating large loads is expensive business.

WEAVER: Yes. It certainly is. And I think I've seen some figures that show that your typical load testing infrastructure is only used about 10 or 15 percent of the time in any given software development lifecycle.

So, and you have a ton of infrastructure investment that sits out there not being used, and you're really not getting good return on investment for that infrastructure. Turning to our specific load testing on the IBM cloud solution, can you talk a little bit about what it is and how it works?

KNOWLES: Yes. So again, there's two parts to it. It's set up the workbench software, so that within the

application itself, you can define what sort of load do you want to generate. And you might have some of the load generated from traditional physical hardware. You might then want to say, okay, we want some large loads, I don't have the infrastructure for this. So I want to provision that muscle on the IBM cloud.

So we have an image on the IBM cloud that can generate that load, so we can provision one or more machines on the IBM cloud. And against each of these virtual machines, we have a pay as you go licensing model for it.

So within the workbench software, I can say, yes, I'd like to simulate 100,000 users. The workbench software will then say, okay, so for that we're going to need 100 or 1,000 machines depending on the type of test I want to apply.

Now, at the point that you're going to start the test, the software will then automatically provision those machines on the IBM cloud, connects them up to the workbench software, allocate the tests to them, the tests then run. At the end of the tests, the workbench software then de-provisions those machines, and the test is complete. And you only pay licensing costs and the infrastructure costs as you go, so it's a pay as you go model. So it can be a lot cheaper.

WEAVER: Okay, so that's actually a unique new service I

think from IBM Rational in that what we're doing is providing this ability to only use what you need.

And in conjunction with the GTS -- Global Technology Services -- we provide the IBM cloud capability through the IBM smart business development and testing on the IBM cloud, we've given the opportunity for our customers really to take advantage of this unique model and save a tremendous amount of money doing the performance testing piece.

So as Jamie said, IBM's Rational Performance Tester is really the workbench piece, and then the image on the IBM cloud is the second piece. So Jamie, I understand this is available now or is going to be available later this year as a technical preview. Is that correct?

KNOWLES: Absolutely. So this stuff is brand new. So we're going to run a technical preview which will be free of charge in terms of licensing costs. We're hoping to run this for three months the start of next year. So again, the unique proposition is the automation capabilities and the workbench tool to automatically provision the agents and then the pay as you go model. As I said initially will be free of charge.

WEAVER: Okay, terrific. So let's move on now to the second capability that we've announced recently around our

cloud solution, and that's really around deployment, planning and automation.

So can we talk just a little bit about some of the issues that...or, the challenges that organizations are facing today around deploying and testing really complicated applications, and really talk about how we can leverage the cloud to improve the ability of organizations to deploy complex apps.

KNOWLES: Absolutely. So as you say, deploying large, complicated applications, there's a lot to do, there's a lot of steps, activities that need to be worked through to deploy that application. And generally, there's a lot of different groups that might get involved in this process.

So, where companies are developing apps and testing, testing is really big area for this. So in every release cycle, the application has to be tested. So we need to deploy that application again and again and again. And during the test phase, hardware is used, but outside of a test phase, that hardware will be sitting there dormant.

Then the point of deployment in a production environment, then we have to go through that process once again. And quite often, the production environment may be different to the test environment. So we're finding a lot of customers

are telling us that they're finding new issues, new user observed issues as a result of the deployment problem so that the production environment might be different for that test environment.

So we've got a lot of complicated steps that may be manual and time consuming, maybe error prone. We've got a lot of different groups that are getting involved in this. Communicating those activities is different. And then, we have the hardware required to provision that application, to deploy that application, may change across different environments we want to provision to. So those are some of the issues that our customers are facing with application deployment.

WEAVER: Okay, thanks. So in a lot of cases, we actually have existing technology that we've been using for a while that helps with some of these processes. And what we're doing in the case of our deployment and automation solution is really tying all these pieces together, providing some new integration to really help automate a lot of that flow.

So Jamie, I think, we can start with the idea of discovering your application's topology, understanding the environment you already have with some of the Tivoli tools like TADDM, right?

KNOWLES: Absolutely, yes. So I think that there's two things that can really help with application deployment. The first is a really good way of doing this. So as you say, we've got a series of tools and a series of steps that will take you from discovering an existing environment that we need to provision to.

So, let's say that we have a production environment, a physical infrastructure. You might want to discover what does that environment look like, so we have the Tivoli discovery tools to do that.

We then want to start to design an environment that perhaps you want to use virtualization. So to save a lot of those costs that we're talking about, that hardware sitting there dormant, we want to be able to serve up environments on a self service basis.

Then we might want to use virtualization technology, you might want to use a cloud environment to run our application to. So we want to model what does that cloud environment look like and make it as similar as possible to our production environment. So first step is about that discovery process.

The second step is then the planning. So we want to be able

to model out what is that desired topology of hardware look like, and against those topologies we might want to start to specify automation instructions.

So, there's been a trend in IT deployment -- so in the olden days, all these activities would be manual. Somebody would be sitting there with perhaps instructions written out in handwritten documents, documents in all sorts of shapes and forms. So it was a manual operation.

And then came along automation tools. So Rational has the Build Forge and RAFW, the Rational Automation Framework WebSphere toolset, okay. So a fairly technical guy can start to write out a series of activities to form an automation workflow for a particular load.

So what we can do in this solution is to make it much easier to build out that workflow. So within a graphical environment, we can draw pictures of what this environment is going to look like. We can start to select from a series of templates or reusable patterns that we've done before, and start to select those reusable patterns and apply them to our notes and our new topology.

So we've tried to bring into this the notion of architecture. So you might have tools like Rational Asset Manager, which is going to hold a library of predefined

patterns. So we might download these topologies before, so we want to reuse those successful patterns.

We might have a corporate architectural unit that might define, this is what a standard topology for this type of application might look like. And these are the sorts of workloads that you want to carry out to define applications to it. And so we've got this library of reusable patterns.

So when we're going through the planning phase, we can reuse a lot of this information. So once we've planned what our desired topology is going to look like, we've planned out which automation workflows and activities we want to apply to it, then we can publish those models to automation infrastructure. So we have the RAFW server sitting there, we can publish that model to that server.

So we also want to send those instructions to our cloud platform. So the points where we provision a set of virtual machines, so, we want to be able to activate those automation instructions. So we want to keep a base set of images that is as small and simple and reusable as possible.

And the idea is, we can publish this deployment topology to our cloud platform, and then when a user or maybe our automated testing toolsets provisions the machines in that topology, so let's say we've got banking application version

1.0.1, then the cloud platform itself will then provision the machines, see that there are some automation activities listed in the service definition and then make a call back to our automation toolset to then carry out those automation activities.

So we go through a process of discovering where we want to go, what we want to produce. We can plan out against the virtual environment what are the activities that we want to perform to provision and deploy our application. We can give them that process by reusing existing fragments of model and deployment workflow. And then, we can automate that process as much as possible. So as soon as we provision that in the cloud, that automation workload is carried out.

WEAVER: Yes, that's terrific. So what we're really doing is reducing the risk in this whole process and making it much more repeatable when it has been in the past. So we're saving a lot of time in terms of trying to configure these deployed systems whether they're test systems or production systems.

We're saving a lot of time in terms of automating this process. It was very manual intensive before and we're reducing the number of errors, lots of errors because it's a complex process and lots of errors are typically going to

creep in just because it's such a complex process.

We already have tools across this whole deployment, planning and automation. And for discovery we talked about the Tivoli Application Dependency Discovery Manager or TADDM -- and you've got to love IBM acronyms.

For the planning piece, we have Rational Software Architect with the extension for deployment planning that really helps you to architect your desired deployment policy, pulling in templates from Rational Asset Manager.

And from the automation perspective, we have the Tivoli Service Automation Manager including things like Tivoli Provisioning Manager that can automate the process of actually creating these virtual machines out on our cloud with links back to Rational Automation Framework for WebSphere to actually deploy the application.

So we have a lot of these capabilities today, and we're actually building out more capabilities as we move forward.

So Jamie, where can a customer to get a hold of some of these inspirations that we've talked about, specifically between the Tivoli and Rational group.

KNOWLES: Yes, so we're introducing capabilities into the individual tools. On the Tivoli integrated service

management library -- it used to be called OPAL -- we're including some entries in there that list out the bill of materials, products that you need to implement this solution, along with some accelerator packs which introduce these capabilities.

So in Q4 this year, we have our Rational Software Architect Version 8.0.1 which has an integration with RAFW. Again, this quarter, we're releasing an integration between Tivoli provisioning manager, part of the TSAM set, which allows that automatic call back thing when you provision virtual machine in the cloud to Rational Automation Framework for WebSphere.

So, have a look at the Tivoli integrated service management library, and you'll see an entry in there which tells you exactly what you need to solution.

WEAVER: All right, that's perfect, thanks Jamie. So to wrap up, what we're really beginning to see here is the fruits of Rational strategy to help organizations plan, build and deliver cloud services.

And our solutions now range all the way from enterprise architecture management to plan the transition to cloud for an entire IT organization. The deployment management, planning and automation, which we just talked about, through

to ensuring the security and compliance of cloud applications with tools like Rational AppScan and Rational Quality Testing.

Application lifecycle management with software delivery services for cloud computing that we launched back in March; and finally, the large-scale performance testing with a new load testing on the IBM cloud.

And the key here is that we're really leveraging the cloud to transform software delivery and really transform what we called at [Innovate], the software common metrics, right, the ability to allow organizations to innovate and not be held down by a lot of the issues that we've seen in the past around fixed costs.

So we're trying to make more variable costs, lower costs, reduce risk and increase productivity for development and testing. So with that, we'll wrap up, and Angelique, I'll turn it back to you.

MATHENY: Steve and Jamie, that was really a great discussion. I enjoyed listening to that, so thank you so much for sharing your time today, we really appreciate it. That was Rational's Steve Weaver and Jamie Knowles, discussing what's new in IBM Rational Cloud Computing Solutions.

To share this podcast with your colleagues, or if you're interested in more podcasts like this one, check out the Rational Talks To You Podcast Page at www.ibm.com/rational/podcasts. We'll post a link to IBM Rational Solutions for Cloud Computing to help you get started.

This has been an IBM podcast. I'm Angelique Metheny. Thanks for listening. Keep tuning in as Rational Talks to You.

IBM Podcast

[MUSIC] [END OF SEGMENT]