Event ID: 334285
Aetna's Commitment to IBM WebSphere DataPower Appliances Drives Faster SOA Adoption and Integration

Karen Bannan: Good morning, good afternoon or good evening, depending on where you are in the world. And welcome to today's webcast, Aetna's Commitment to IBM WebSphere DataPower Appliances Drives Faster SOA Adoption and Integration. This is the second part of a five-part series on Succeeding With Business Agility, Five Forward-Thinking Organizations, brought to you by InformationWeek, IBM, UBM TechWeb.

I'm Karen Bannan and I'll be your moderator today. To ensure that this is as interactive an event as possible I would like to make just a few announcements before we begin. At this time we recommend you disable your popup blockers. This webcast is designed to be interactive and makes use of popups or widgets to help you interact with the presenters. You can launch these widgets from the docking tray at the bottom of your console. Some will open in a separate browser window. You can resize the widgets by dragging the lower right corner, move them by dragging the top bar or minimize them by clicking on the blue button in the upper right hand corner.

During our webcast you can ask questions at any time, just type your question into the Ask a Question text area and then click the Submit button. The slides will advance automatically throughout our event. You can download a copy of our slides by clicking on the Download Slide icon which is located below the presentation window.

Towards the end of the program we will ask for your feedback in a separate popup form. If you experience any technical problems please click on our webcast help guide at the Help icon located in the far right of the docking tray at the bottom of the webcast console.

In addition, you can contact our technical support help line, the number for which is located in the webcast help guide as well.

And now, on to today's presentation, Aetna's Commitment to IBM WebSphere DataPower Appliances Drives Faster SOA Adoption and Integration. We have three distinguished speakers; Preksha Patel, Andy Grohman and Nimesh Shankaran.

Preksha Patel is a web engineer with Aetna Web Technology Engineering. Preksha started at Aetna in 2005 as an application development doing J2EE and IBM WebSphere DataPower Appliance development. In 2008 she moved to the Web Engineering team where she works with WebSphere DataPower Appliances from an infrastructure perspective.

Andy Grohman is a WebSphere DataPower Connectivity leader at IBM. As the WebSphere Connectivity Competency Leader for North America and a client technical professional for WebSphere DataPower appliances, Andy is based in Charlotte, North Carolina. Since 2007 he has worked closely with some of the largest DataPower customers in the world. Andy has 16 years of experience with IBM in a variety of development services and sales roles.

Nimesh Shankaran is a Senior Engineer at Aetna. Nimesh is a Senior Engineer at the Web Systems Engineering team at Aetna and has spent the past ten years working for the Architecture and Administration of the IBM WebSphere suite of products. He is currently Aetna's DataPower Technical Team lead, focusing on the architecture and implementation of DataPower technologies.

And with that, I'd like to hand the presentation over to Andrew so we can get started. Over to you Andrew.

Andy Grohman: Thanks very much Karen, hey everybody it looks like we've got a great turnout for today's presentation. So really thank everybody for taking some time out of their day to join this exciting conversation about IBM WebSphere DataPower Appliances, and more importantly the success that Aetna has had with DataPower over the years.

We really have a true historic relationship between IBM and Aetna over the years, with DataPower, and like I said, we're excited to be here today to talk about this a little bit.

So with that, let me introduce today's agenda, going to spend a few minutes just introducing the WebSphere DataPower appliances to make sure everybody is level set on what it is we're going to talk about, the technology involved here. I'll talk a little bit about the healthcare insurance industry challenges from IBM's perspective and then we'll turn it over to Nimesh and Preksha who will take us through Aetna's specific experiences with DataPower over the years. And specifically the specific challenges that they've faced, how DataPower has helped them overcome those challenges; the initial benefits they've seen; and most importantly the evolution that DataPower has experienced at Aetna.

We'll recap with some lessons learned and then we'll talk a little bit about the highly collaborative relationship IBM and Aetna has enjoyed. And then we'll open it up for Q&A.

So with that, what are the WebSphere DataPower Appliances? Simply put, it is a one-U or two-U rack mountable appliance that provides secure, connectivity and integration solutions for your infrastructure. We usually talk about DataPower in four important pillars. We talk about Security, Simplicity, Acceleration and Governance.

With respect to security it's important to understand that DataPower is truly built from the ground up to be a secure device. From its physical characteristics and tamperproof casing to all of its encrypted and signed firmware and protected images that we put on the appliance itself, it's really built to be very secure, built to run in your DMZ, built to run in internal and to lock everything down very securely.

We have lots of functionality within the appliances that help you integrate your service oriented architecture, provide Web 2.0 interfaces and restful interfaces, to emerging mobile applications. We do a lot of B2B focused integrations and security as well. And we can even provide some secure connectors to your cloud-based environments. So there's a lot of different ways that DataPower can help secure your infrastructure.

We talked about simplicity, being an appliance it dramatically reduces the total cost of ownership of traditional software-on-server solutions by reducing the amount of maintenance and overhead through our firmware driven approach. More importantly, our configuration driven approach for creating connectivity solutions makes it very easy to stand up DataPower and within minutes and hours you can implement some pretty advanced integration solutions. And to that end we focus on accelerating your time to value and we usually talk about acceleration from two aspects. Certainly everything that goes into the DataPower hardware is built to perform very, very quickly, we have optimized network interface cards, we have optimized XML processing chips within the appliance and all that combines together to give you a very, very high performance environment.

But we also talked about accelerating your time to value, and again, that goes back to our appliance drive approach and configuration driven approach. And hopefully through drag-and-drop interfaces you can very quickly stand up and like I said, complex integrations very quickly.

And then underscoring all of this stuff is governance, so DataPower provides an awful lot of capability for helping you manage and govern your IT infrastructure, whether it be integration with some of our service registering repository tools or other access policy manager tools to help you manage that security policy and access policy from other tools and use DataPower as the enforcement point to protect that traffic flowing through your network.

At the end of the day we see a DataPower appliance providing superior performance and hardened security, helping you improve your return on investment with a lower overall total cost of ownership. So specifically when we take DataPower and we start looking at the healthcare industry there are a number of challenges that Aetna, in particular, but the industry in general faces. We talk a lot about the need to secure the personal health information and personally identifiable information. So when you're dealing with this highly sensitive information it's important to have a very, very secure solution and DataPower provides that security in a variety of ways.

It's important to understand that a lot of this data lives in a variety of systems, very heterogeneous systems, some of them legacy, some of them brand new cutting edge stuff and so you'll need a solution that can speak that heterogeneous language and interoperate in that environment well.

Also important is health insurance is a very competitive industry and so you need products that provide market advantages, give your ability to respond quickly, to quickly changing factors and changing market place. And again, Aetna will take us through how DataPower helps them achieve those goals.

And like I mentioned, highly heterogeneous systems is the hallmark of the health insurance industry and I think that probably applies to most industries we see today. So with that, I will turn it over to Preksha and Nimesh who will take us through their specific experiences at Aetna. Preksha?

Preksha Patel: Thank you Andy and great job. Welcome everyone it's wonderful to see all these people here. And since we have folks from all over the place I wanted to do a quick overview of what Aetna is and who they are in the industry. Of course we're a healthcare company and we mainly deal with healthcare, dental, pharmacy, group life, disability and also employee benefits.

A quick overview of our membership, we are over 18 million medical members, over 13 million dental and over 8 million pharmacy. We offer, I guess, a plethora of different products to our members as well as our providers and we take pride in the products that we build and support here at Aetna. So just -- I'll keep the slide up for a quick second just because there's a lot of information. But that's Aetna in a nutshell.

So Andy covered a lot of the general healthcare industry challenges and a lot of the Aetna challenges flow along that line. They parallel exactly what Andy went over. Aetna is a very mainframe oriented, originally -- well all of our data lives on the mainframe. So we were looking for a way, when the web services exploded here, a way to easily share this information with web services, looking for data on the mainframe, how do we quickly transform these different various data formats as fast as we can and as easily as possible. So those were the things we were looking at. We had a different platforms here at Aetna, we have CICS, we have a web platform, .NET, everything, you name it we have it.

The other part that we had to deal with, with all this data going back and forth between one back end to another, we were concerned about security. We have a lot of important information about you, providers, and we do not want anyone getting a hold of that. Andy covered some of that but you know, we're looking to protect our PCI, our PHI, our PII data and that was important to do as far as exposing all these services we want to make sure we keep them secure.

What we did here at Aetna years ago when we were looking into solving this problem space was to do a technical evaluation and that's not different from any other company whatsoever. So we asked ourselves these questions, how can we quickly do a transformation, make it easier to do, mostly integrate this mainframe heavy environment and expose these web services in a way that was simple.

We wanted to do quick XML transformations and we also wanted to do apply security. Now, security in itself is a complicated space as well as integrating and building mass and doing them quickly. So all these questions came to the table as well as how can we standardize and build an environment that is simpler than what we have. Aetna before DataPower was building services using code and it was hard to say one area or domain of Aetna was trying to get at data on the mainframe and they built their own Java code and maps to do so.

Then another, say the member area, came and did the same thing. So you have five different domains doing five different implementations of Java code to get at some mainframe data and it's very difficult to troubleshoot, to help support that type of environment. And we were looking for a way to easily build a reusable environment and we looked at a couple of different options and DataPower was one of the ones that really stood out in a way that we really felt it was a good fit.

That being said, I'm going to hand these slides off to Nimesh and he's going to talk about DataPower at Aetna.

Nimesh Shankaran: Hello everybody and -- this is Nimesh, thank you all for joining this presentation and I welcome you all again. And thanks Preksha for going over all the challenges that Aetna have faced before we brought DataPower in. And to overcome these challenges is what we brought DataPower into Aetna for.

And in particular the two most important challenges, which is implementing WS security for web services being the -- one of the challenges. And the second challenge as Preksha said, was to transform data from XML to CopyBook and CopyBook back to XML. XML being the format that Java applications prefer and CopyBook is the -- CICS Cobol applications would understand the CopyBook format and to do this transformation.

So I've been at Web -- my team, Web engineering we are the ones who roll out all the infrastructure when any web related product comes in and in the last ten years' experience I've seen that the biggest challenge when we bring a critical product in is adoption. I mean nobody wants to be the first ones in the line to adopt a brand new technology and they want it to be more baked and they don't want to take the risk of adopting something brand new and they're willing to spend a little more money or resources at something that's baked in. And we assumed that that's exactly what we will have with DataPower as well.

And we were surprised because you know, we did have adoption challenge but of a different kind. It seemed like everybody wanted to jump into this. Every single product that came along the way had something to do with DataPower or another. And I'll give you -- build this in terms of numbers so this would make more sense. So we -- the first services we -- or real services that we put into DataPower was like probably three years back. So, in the last three years at -- where almost nothing was flowing through our systems to today where we have more than 8 million requests flowing through our systems every day.

And we started it off with like four boxes, we have 16 today in production. And with just a few handful of services that we piloted, today we have more than 200 and to do a rough math, we have one major release every quarter, so that's four releases a year and 200 services in three years, that's about 17, 18 services every release.

And to imagine putting 17, or 18 services every release, now we are not counting the core fixes here, just brand new services, it's virtually impossible with the APLC processes and everything that we here at Aetna. It's almost impossible to have those numbers in place without DataPower or some other such tool that we have. And what -- and there are many reasons why we have such an exponential growth of web services once we had DataPower and just to zone in one the one important reason, it's because of the fact that we can deploy web services at a much quicker rate with DataPower than without it.

And there are many more reasons and hopefully at the end of this presentation those will be evident as well.

So we started off with just the XI50 boxes to do the transformations and security and we soon realized that that would not cut it for us so we have now two completely different tiers. One tier that does security, and those are the XS40 boxes, and the other tier that does transformation which is the XI50 boxes. And you know, on the security devices that's where we put in our WS proxies to enforce WS security, the chosen mechanism that we have at Aetna is to use the X509 certificates to do that. And we also do XML validations, in force SLMs security and error validations, logging and those types of stuff.

And I have another couple of charts down the line that goes into these details. At Aetna we have a shared environment and what that means is that across all the different business domains we have the same infrastructure for DataPower that caters to all services cross-domain.

And this has really helped us a lot because we don't have to reproduce the different environments that we have across each and every one of these domains. Like we have all the way from lab to production like most companies would. We have a very robust lab and our lab has got at least one service that pertains to all of the different patterns of things that we do in production. Our lab has got everything from WebSphere, the different flavors of that, .Net all the way through to mainframe, CICS and MQ. So whenever a new pattern evolves we make sure that we test this out in the lab and we also test out all the services when we have a formal upgrade and such. So that when something hits development it's well baked in. All the code and instrumentation and all the surrounding scripting and everything is in place before something moves into development.

I'm rushing through, sorry, apologize about that. I have just a lot of material to cover so that's why I'm speaking really fast.

And we also have a very good performance lab. Performance lab was a little bit difficult in the beginning because we had to don a different hat. We are very used to performance testing off web applications and we had to switch gears when we are performance testing a web service. So we had to, for instance figure out how to do tests of multiple services at the same time using SOAP UI and things like that. So we got through that learning curve and so now by the time a service gets filtered from lab all the way through production our production is pretty solid.

Now, what I'll do is I'll go through some of the request flow patterns that we commonly have for the most important use cases here at Aetna. What you are seeing on the screen right now is a high-level topology diagram. And the vertical lines that you see is the firewall and everything on the left of the firewall is the external network and everything on the right of the firewall you're seeing, is our internal network and we have the DMZ in the middle. And the yellow boxes over there, that's the ones that we mentioned is doing the security and everything on the right side of those yellow boxes, the XS40s, that's where the web servicers reside.

So this is the request flow for getting to a web service that's on the back end. And the backend typically is a WebSphere application server on Linux or mainframe, you've got IAAS, you've got .Net applications, or you've got CICS, as Preksha said, that's where we have our core business running and we have -- that's the most important use case for us where we want to expose out these business services to the Java front end applications.

And you see that you have internal clients and external clients and these internal clients are again, IIS or .Net, WebSphere, or CICS. And you would see that everything is flowing through these boxes, these yellow boxes. The reason for that is that these boxes are not just for security but they also double up to do what we internally at Aetna call the routing devices. Because the reason why we call it the routing devices is because everything flows through these boxes and it's these boxes that look at a service and say, hey, where is the back end, let me send it to the correct back end. Okay, this has to go to WebSphere so let me send it there. Or oh, look at this, this is going to go to CICS, but how do I send it to CICS? I'm sure it's going to CIIS it would need some form of transformation because what is coming in to me is SOAP or HTTP and CIIS does not understand that so let me send it to the transformation devices, that's another set of devices that will take this information and transform it into the CopyBook format and send it to NQ which is picked up by CIIS on the back end. And I've got a little bit of a zoned in view of this in my next slide.

This picture is an evolution of some of the standards that we put in place and one of the important standards that we have is that every -- is that we have a single point for entry for any web service. And that's an outcome of another standard that enterprise architecture had, which basically states that we -- I missed my train of thought, but, yes, this is the single point of entry called web services and -- I forgot.

But anyways I'll skip over that, now this standard that we have is that we don't have any distinction between an external request and an internal request. So we consider all requests that come in for a web service with the same security level and implement the same security standards regardless of whether it's an external service or it's an internal service.

And in this diagram what you would see is that you would see that there is no DataPower boxes in our DMZ and we initially started off with having DataPower boxes in the DMZ and we pulled it out because the bunches that we got out of putting a box on the DMZ was not enough to work on some of the challenges that we internally faced. I'm thinking that could potentially be particular for Aetna, we have our backend services being consumed both by internal and external services. So we had to duplicate out our security in both places, in the DMZ as well as in the internal network. And we replaced out these boxes with the existing web service that we already had. So we created a few more instances of the web server, IHS instances and basically proxied the request into the internal network.

So what you're seeing is that all the requests that are coming in, on this chart at least, these are all in SOAP or HTTP and I'll move over to the next chart that I said, where I'm zooming into the next view and here what you're seeing is that you have these transformation boxes and the other mechanism of -- or the other protocol that this borders, MQ, and that's when, you know, a CICS service for instance wants to call, sorry, a CICS consumer wants to call a service that's living out on WebSphere for instance. So it puts something on MQ which DataPower picks it up and passes it on to WebSphere after doing protocol transformation and data transformation. So that's the second request flow pattern that we have. And the important thing that you see here is that the backend is completely abstracted so that's one huge plus with DataPower where the consumer is not aware of where the web service is served out from. It could be from .Net or it

could be from CICS and the consumer does not need to make any code changes or code anything different depending on the back end.

I think I've covered all the most important standards and patterns that we have except the one thing that I lost my train of thought on. So, with that I'll hand it over to you Preksha for going over some of the benefits that we have derived out of using DataPower.

Preksha Patel: Hi everyone. So initially we brought DataPower in as, we've been talking about, an integration. There was a lot of these transformations which we've talked about, a lot, already. And it was a -- the fact that DataPower could allow us to quickly integrate and get these transformations done. And originally I said we were doing code-based transformations, so here at Aetna and I don't know about other companies that when you make a change to an application EAR you have to repackage it and deploy it out through your environment and retask and I mean there's still some of those items of migrating out your changes in testing, but with DataPower and the way we've implemented the map transformations it's much simpler.

We push out a map, DataPower picks it up and that is done through a simpler process outside of a huge code lifecycle process. This is a piece of code, but the way we handle it is much more simpler. We do integrate it into our source repository but the way we're -- you know, we're not having to rebuild and repackage and go through the overhead of an application deploy which can have its own issues in itself. The EAR could have other code in it, application related code, so this simplifies it and separates out that map transformation away from application code.

It's fairly easy to build a transformation map. There -- IBM offers another great piece of product that is the design studio which you can use to build these maps. And along with the integration piece of DataPower there is also, I'm sure you guys have heard, enterprise service bus and when you hear them out solo, you hear about the enterprise service bus.

With Aetna we had our own enterprise service bus that were implemented here previously to DataPower arriving and what was interesting there is that the fact that DataPower fit very nicely into that world. We were to take those standards that we started with the service bus and getting traffic and data moved around to DataPower. And those standards have since been moved over so the DataPower piece has been a nice clip into our Aetna internal service bus because we take those standards that we build and say, oh, well if you want to get to something that is reachable through the service spots DataPower can do this and it's simple.

So, some of that really helps. It helped control all of the different exposure to your services, the providers say they're mainframe in this case, we have a lot more of mainframe providers, is that you can take -- you can easily say, hey, here's how you call a service, you don't -- like Nimesh said, you don't have to worry about who you're talking to. And also the way we brought DataPower in and broke it up within Aetna as a department. We have a one-stop area of expertise that deals with the transformation maps and how they expose these mainframe services.

So it made things a lot easier, and easier to go to one place. So the agility and flexibility is there with DataPower. A lot with the fact that you can code your own excess LTs and build your own -- use the DataPower extensions to build what you need and make it as reusable as possible.

As far as like a return on investment benefit, DataPower was initially brought in to solve a problem and with all problems, people don't generally just jump into a solution unless it does bring a return on investment. I don't have exact numbers but I can say that with the fact we see the return on investment with every project that comes through.

We found that from a perspective of trying to get a map rolled out, a change done, like I said before, it's -- we can quickly get a change to a map out faster than a deploy of an EAR which may have to go through a regression test and things like that. I mean we obviously do a set of regression tests but in this case we're dealing with a regression test for that specific map change outside of the entire EAR.

So in that case the return on investment is the speed of delivery that we can get certain solutions out and expose these web services. So, a lot of that comes from individual projects that make up a big return of investment and with -- from a security perspective we see a lot of that. So as DataPower evolves at Aetna, and Nimesh said, we started with a couple of XI50 devices and then we brought in some XS40s, like we built up to this whole tiered environment and how we handle our traffic.

With the evolution of which hardware we brought in, we also had evolution of how DataPower was used. We talked a lot about security previously, as well. We have found that DataPower has brought in quite a bit of help and ease to our security. When we have these projects that come out, we today use a lot of -- we started with simple things like X509 for verifying signing and DataPower does that very easily, it breaks down these components of a big picture for you to understand. Here's what makes up a verify, here's what makes up a sign, here's what makes up SSL, and it simplifies these difficult concepts of security. So we -- DataPower definitely helps in that space. A lot of it has to do with the GUI DataPower has for configuring security, it makes those things very easy.

So because that's the case a lot of the security usage really blew up here. And we've had -- gone from simple verify and sign to starting to expose better -- securing our SSO solutions. DataPower is just not a -- for data, but for security, it's not just an integration point for that. So it's a lot there.

One of the big projects that really DataPower was such a simple thing that added such a big benefit was we had an outside customer who was starting to send from their SAML assertion with a lot more private data. And we had to figure out a quick way to be able to encrypt that data, get it through Aetna and get the customer what they wanted.

And this solution we implemented into DataPower, we put DataPower in front of this -- our SSO application in the backend and we were able to implement this out in literally no time. In this case, the developer and folks that were working on this project didn't know enough about security, they also didn't have enough time to learn it and implement it through a code implementation. So it was a really big benefit to us, for this, we actually gained customers and members from being able to get this solution out as fast. It's actually become a great strategic solution as we have gotten a lot more customers involved in doing this. So it was great.

Other side ones are Kerberos, single sign-ons, we already talked about so -- and before I also said that the XSLTs and extension functions that DataPower has, it makes the usage much easier. So I guess I can't say enough for that. But as you saw in Nimesh's slides too, we've evolved and things are changing and they're still changing. Mobile is coming down the road. Aetna is really hands on with those things and we just really feel that DataPower seems to be flexible enough to get into those spaces and it's very beneficial.

So I would -- I did want to move on to Nimesh, he's going to talk about some of the lessons learned related to DataPower. Nothing is smooth, so we want to share this information with you.

Nimesh Shankaran: Thanks Preksha, and hello again. So yes, this slide is about some of the lessons that we learned at Aetna in the four years that we've had experience with DataPower. Now, this would probably resonate well with people who are in the nascent stages of implementing DataPower. What did we do right, what did we do wrong and what can we learn out of it? And one thing that immediately -- that we saw is that your use cases will grow, because typically when we bring a product in it's for a particular business unit who wants to use it in a particular fashion and we bring a product in to -- Like for instance, we want to expose business service out to the Java world, right, so to do transformations is the use case, but the product itself is having so many capabilities and we would soon find that those capabilities are done better on DataPower than on products that are already available in-house, that does the same thing, that more and more people want to use DataPower. And we cannot say no because it's the business that funds the project ultimately.

So we can't say no because we are not prepared for a particular use case, so we will have to think about this ahead of time and see how this device is going to evolve two years down the line and plan today for that. Like I said, you feed the box data of type a, you get data of type b out of it and we think that the box is an island but actually it's the opposite of it. The box is so well connected that we realize that the success of DataPower lies in the support that we get from all the different teams that are supporting this.

Like for instance, here at Aetna, web engineering owns DataPower and we are a logical fit because we already own all the other web products and we are already interfacing with a multitude of teams. On the one side we are interfacing with the development community on a day-to-day basis and then we turn around and interface with all the infrastructure teams like performance engineering, like we are really tight with enterprise architecture in terms of whenever a new pattern evolves, whenever a new standard needs to be established, we are tight with the enterprise architecture team to do all of that. With security engineering and networking and our team dons the hat of so many teams all at the same time, because DataPower is classified -- IBM classifies it as software but in all reality it's a hardware device and you need to have a little bit of a networking background, a little bit of storage management background. The team wouldn't have to do scripting on SOMA, Java, Pearl and the nature of such. So placement of the device in the correct organization is I'd say quite critical in the success of DataPower.

And in the success of DataPower lies the success of SOA at Aetna for us. So that's probably the first couple of things that we learned. And the other thing, like Preksha mentioned was the

whole thought about centralization. So like other organizations we also have these business units, like plan, member, claim, provider and we are not just centralized in terms of an infrastructure for DataPower but we are also centralized in terms of the development of code on DataPower and the management of DataPower itself. So like this is contrasting to what we currently do with other technologies. Like for a WebSphere and we have developers developing Java across all of the different organizations and across different geographic locations. But then, DataPower when it came in it was a brand new technology and there was definitely a learning curve. So we imposed this learning curve on one centralized team and we gave the benefits of DataPower to all the architects of -- who are across domain. So the demand for having a service developed came from the architecture teams and it's the development teams that -- it was a single development team that already had the expertise to develop the services. So that was a big reason why people were not hesitant to move to the new technology because hey, they don't have to develop it.

So that is one reason why things really grew at an exponential pace. So yes, business is always good so growing and having more and more services, enterprise business services, as we call it, right. Because these are across domain, across all the platforms and there is a lot ofcross-domain calls, a service calling another service, an application and domain one calling a service in domain two. So that's the essence of the whole SOA, shared, reuse and stuff like that. So centralization, that's another core takeaway for -- if we were to -- and lucky enough we started off in the right path so we reaped the benefits of this as well.

The next lesson learned I would say is the establishing of standards and patterns and best practices ahead of time, especially if you are going to be using this as a security device. Double your security, there is so many questions around this, you need to answer these questions ahead of time. How are we going to do the double your security, are we going to use -- what are the different patterns, are we going to use X Power 9] are you going to use user name token profile? Okay, so we are using X Power 9 how do we exactly code this, when we want to verify, when we want to sign and WebSphere and IAS and so what do we do? And it is always easier to prototype all your patterns ahead of time and publish them out somewhere that the development community can use.

And -- because these questions are going to be asked over and over again, especially security related ones. How do we do security? In terms of, okay, are we having mutual authentication, do we support Kerberos, and if all this is published out, it's easier to point to a website and say, hey, that's how we want to do it rather than reinventing the wheel each time.

And ahead of time is the essence of this whole thing because we had an issue where when we brought the first service in we had not really baked the pattern and we decided to sign on DataPower, I don't want to go into the nitty gritty details, but basically we realized after the fact that okay, if we put the servers in place and there are five consumers now and we have security hold. And to address that this is a huge project because we have to change the code in so many different consumers and also the provider and the DataPower in the middle. And all this has to be happening synchronously.

And it was quite difficult to retrofit and anything to do with security, retrofitting it, is going to be really difficult after the fact. So that's another takeaway.

And the -- another -- and of course, you know, you have to take it for granted that the rate of adoption is going to be phenomenal. And when you have this a rate of adoption, when you go into the details, an explosion of the number of web services, depending on your usage, what does that mean to you? Does that mean that you're going to have two years down the line, are you going to have 200 domains, that's not going to be a supportable model.

Are you going to have 200 web servers in the domain, again, that's not going to be supportable. Or, if you're going to do transformations, are you going to have 200 multiprotocol gateways? And you know, again, that's not going to be supportable. So you will have to see the end point and like for instance, you may have to come up with common framework services that is -- and these things are more geared towards your necessities. Like for example, we wanted to collapse a bunch of multiprotocol gateways into one and we have this common framework that looks at an incoming request and says, hey, okay, so you are belonging to this a pattern so I'll send you here and the framework looks at it and says, hey, okay this is the transformation map that I'm going to apply for you and that's the MQ that I want to send the requests to and this is the MQ that I'll listen for the response.

And those type of -- so what we essentially have done with the frameworks is collapse like 50 different services that belong to a particular pattern into one multiprotocol gateway, otherwise it's going to be -- we cannot sustain the growth rate of web services.

And these are important standard decisions that we had to make and we had to work very closely with IBM on all of these, bouncing ideas back and forth, getting feedback. And to go over the relationship that we had with IBM through the whole process I'll hand it back to you Preksha, thank you.

Preksha Patel: Okay, so this is really quick but, we are -- we work with IBM a lot. We have quite a bit of collaboration with them starting from product support, of course, to purchase, but also lab advocacy programs where we work with a lab advocate that is intimate with DataPower at Aetna, the implementations and the use.

So that is very beneficial and we meet on a -- you know, it all depends on what's going on at that time of year, how busy we are. But what is important at the time and what we're trying to learn, so we use that a lot.

Aetna provides a lot of information to IBM as a whole, the DataPower team, about our implementations, not just the lab advocate, so we can help them build a robust testing environment for DataPower appliances. As you -- I look -- some of you may know, DataPower can do a lot of things and can be used for many, many different things. I mean, across 100 companies I'm sure there's so many different uses.

So it's very important to share those uses with IBM so they can test firmware better and deliver a better product to everybody. And that's what we try to help do. The most beneficial that we've

seen with the lab advocate program and also with -- across our product discussions to help integration between all the IBM products it's very beneficial to have those types of meetings. We have group touch points with Tivoli, West Durham Q and Application Server because DataPower is an integration product, it needs all these things.

So what I'm going to do is pass this on really quick to Nimesh and he'll go over quickly what some of that cross-product integration has provided.

Nimesh Shankaran: Thanks Preksha. And this is a narration, an interesting narration of the experience that we have, or had, at Aetna. And in the beginning when we had, as I mentioned, we were surprised with this growth of services and we were not fully prepared for it. And our on call pager was quite scary. You know, we jump onto the call and we are asked basic questions that we are expected to answer. Hey, are web services slow in the backend or somewhere, can you tell us how much time it took on DataPower for that web service to perform? Or, can you tell if all the boxes are slow or just that one DataPower box is slow? Or things like that, which we were expected to you know, provide immediate answers.

We did not have the answers for that. And what -- some of the things that really helped us was -- because DataPower throws a lot of introspection. We came to know that, you know, they have something called the latency logs that is basically a logging that can be turned on, on DataPower, which throws a ton of information in space dealing with the numbers. And it's quick to glean that information and -- but the problem is that it is not easily human readable so we will have to take that out, parse it out through a parser, and then we'll know exactly what happened within DataPower. And we thought that that an information we should be able to -- we already have a working Tivoli infrastructure, here in-house, we have TEP, Tivoli Enterprise Portal that can parse these information for MQ folks, for WebSphere, so why not for DataPower.

So we got in touch with our onsite deployment consultant at IBM who set up a call with DataPower product development and Tivoli product development and we provided this feedback. And I'm sure other people have provided it too because now we have this product that it's a part of Tivoli, it's the ITCAM for DataPower is what it's called. And it's capable of sucking in the latency information and provide wonderful graphs and charts that we can quickly use and jump into a production call. In a couple of clicks we know exactly what's happening. We can get a top-ten vendor list, for instance. So which are the services that are running the slowest, top-ten in production, in a click of a couple of buttons.

And then we can narrow down -- yeah, this chart you are seeing now, for us, in production, this is probably from production or QM, I'm not sure, but this is what you would see in a push of a couple of buttons. And then we can drill down into the service that we are more interested in and we get a chart like this and what this chart really tells you is that the time it has taken for this particular service averaged out, and the last one hour is like a millisecond, within DataPower, so this transformation of data took like a millisecond and then DataPower is waiting on the backend for like 14 seconds for a response.

So we go on to this immediate response team, IRT call that we have, troubleshooting this, and I send this out to people and that sends -- pushes the problem away from DataPower at least. So the on call itself is starting to look less scary right now.

So that's -- this -- all this is the outcome of some of the relationships that we've had with DataPower. It's mutually fruitful and in return we get products like these.

So next to talk about what's coming down the pipeline for us, I'll hand it back to you Preksha.

Preksha Patel: Okay, so quickly, because we're running short on time here, we're recently working on hardware migrations. This is hardware and it does need to be updated and we are moving from our current XI50 hardware to a new hardware and working on migration plans and things like that. As part of that we're trying to fit in some new lessons -- from our lessons learned, how we can better implement DataPower through our environment. So that's some of the stuff we're working on. Testing all the different hardware types and comparing performance data and it's really cool to see the benefits of what they've changed.

So as far as integration, we're looking at integrating with WSRR and we're working with a couple of teams here at Aetna to do that, as well as IBM. And we're also, there's always new security patterns coming down the pipeline. You never know what crazy map transformations may come down the pipeline too, so we're always working on those things and so really quickly I'll hand it off to Andy to close us out.

Andy Grohman: Great, thank you very much. That's just fascinating stuff, Nimesh, Preksha, I can't thank the two of you enough to take your time and present this really fascinating integration and evolution you've had with DataPower at Aetna. Hopefully this helps open the eyes of some of the audience today on the benefits and possibilities we have here with DataPower.

I think at this point I'm going to open it up for Q&A, is that right Karen?

Karen Bannan: It is, but actually before we get to the Q&A we would just like to ask everyone to please fill out a feedback form that's opening on their computer.

To complete the feedback form please press the Submit Answer button at the bottom of the page. Thanks so much in advance, for filling out our feedback form, participation in this survey allows us to better serve you.

And with that, on to our question-and-answer portion of our event. As a reminder, to participate, just type your question into the question box and click the Submit Question button.

And let's get right to it. We have a question here, this question is from Christopher, he would like to know, so, for the Aetna folks, where did you see the majority of your cost reductions come from?

Preksha Patel: On -- I think we talked about that, a couple of slides back, but a lot of it was after DataPower came here and we saw how quickly we were able to get services rolled out for a

project. So originally because it was code based it took a lot longer and there was no real standard people could follow. So a lot of that was seen in the whole -- just the projects that came through with different customers, internal projects and that's where we saw a lot of it. I don't know if Nimesh has any --

Nimesh Shankaran: Oh, you nailed it. Yeah that's exactly what it was, you know, the quick response time for customers and the quick speed at which we were able to create the web services which we could not have done without DataPower, with just WebSphere -- a Java on WebSphere or CICS, it would not have been possible.

Karen Bannan: Okay, we have another question here. And this question is from Joel, and Joel would like to know, he says, please tell me more about the IBM Lab Advocacy Programs, what has Aetna's participation with the program helped it to achieve?

Preksha Patel: I think that the program is great because we get to be in touch with someone at Aetna who has feelers out -- I mean, sorry, someone at IBM, who has feelers out within IBM to pull out resources for all the different products. We use -- we have a lot of different IBM products here at Aetna and we integrate with a lot of them. So when we have one person we can go to to pull those subject matter experts in it's been really helpful.

The other thing is, any time we're starting to move down a new path, firmware upgrades, new functionality and it's not always a problem per se, but we're trying to figure out how to do something in a certain way you don't necessarily want to open a PMR ticket. It's -- sometimes you just need someone to tell you, here's the best way to do it and here's where, and that's where I think we've seen the most benefit. And we have a good relationship with ours.

Andy Grohman: Yeah, and this is Andy so I could speak to that shortly as well. I would just say that it's mutually beneficial. It's great for IBM to get some of our key development leaders out of the lab and into customer-facing situations where they can really understand firsthand the challenges that our clients are using as they implement our software and what the problems -- the problems that they're facing, the problems they're trying to solve. So it's really usually beneficial to both parties, both on the client side and on the IBM side. And I encourage everyone on the call to speak with their WebSphere brand leaders if they have interest in the lab advocacy program.

Karen Bannan: Okay, we have another question here. This is from Joe, Joe would like to say what is the new hardware? So I guess he's asking what the new hardware was that was installed?

Preksha Patel: Oh, we haven't done it yet, and I'm not sure if we can share some of that internal information. But I will tell you that we looked at the ones that are out there right now and tried to figure out what the best fit from a replacement perspective was. So maybe Andy you could talk about some of this actually. We did work with him on this, there's a what -- the blades the V and the XI52, right?

Andy Grohman: That's right, there's been some interesting evolution of the hardware over the past 12 to 18 months, I'd say. First, we released a blade-based version of our XI50, what we

called the XI50B that is the appliance form factor in a blade form factor that you can insert into a blade center. Provides higher networking connectivity, 10-gig access, things like that. Easier manageability, able to do dark datacenter management, things like that.

That led to our introducing the XI50z which is the XI50 meant to live in a zBX container, which is the blade center that can attach to a mainframe z196 and so now you're harnessing the integration connectivity power of DataPower, but sitting right alongside the XI50Z and managed -- or, I'm sorry, the System z environment and managed within all the confines of z with its higher levels of availability and things of that nature.

And then most recently we released the XI52 which, if you noticed the screenshots, that were scattered throughout the presentation, at the beginning, at the end, those are the black appliances. That's a new, 2U version of the XI50 that we released in June. It offers, again, much more manageability, much more serviceability. We now have eight 1-gig ports instead of just four, we have two 10-gig ports on it for the first time. Much more -- a higher performance, everything is just bigger, faster, more, on it. But what's nice about all of these different appliance models that we provide is that they all run the same firmware. So if you have configurations today on a given appliance model it's really just an export/import. There's no refactoring of any of your configuration required across these different hardware models.

Karen Bannan: Okay, we're actually coming to the top of our hour so we have time for just two more questions. This question actually comes from Randy, Randy would like to know, do you have any plans to use the X150z on the System z, zBX frame?

Preksha Patel: Well, we've actually, we've looked at it. As far as plans, I --

Nimesh Shankaran: A bit neutral on that right now, we are still in the process of evaluating various hardware so we don't have an answer, concrete answer on that right now.

Preksha Patel: But you never know.

Karen Bannan: Okay, so one more question here. We have a question, this is from Uriah, and they would like to know do you have a load balancer in front of your architecture or an SSL load balancer?

Nimesh Shankaran: Well, we do have a -- it depends on where the request is coming from. We do have F5 to send the request out to our DataPower boxes, but once we have it at DataPower we use the DataPower's built in load balancers to go to other back ends from there. So we do have a mixture of F5 and DataPower's load balancing capabilities.

Karen Bannan: Okay, well with that we've come to the end of our Q&A session. There were plenty of questions that we just didn't get to. But the wonderful thing is the folks at IBM and Aetna will be able to follow up and there are actually are email addresses right there on that form, so you can actually reach out and email any of the folks who were speaking today and thanks to all three of them.

And thanks to everyone out there, for attending today's webcast, Aetna's Commitment to IBM WebSphere DataPower Appliances Drives Faster SOA Adoption and Integration which has been brought to you by InformationWeek, IBM an UBM TechWeb.

For more information about today's webcast, please visit any of the resource links that are open before you. Within the next 24-hours you will receive a personalized follow-up email with details and a link to today's presentation.

Additionally, you can view today's events on demand by visiting www.netseminar.com.

On behalf of our guests, I'm Karen Bannan thanks for your time and have a great day.