IBM Digital Video Services

Cloud_on_System_z_audio

Mike: Hi. I'm Mike Baskey. I'm the Distinguished Engineer responsible for our System z management strategy and architecture. Today, I'm going to be describing the Cloud on System z Whiteboard.

This whiteboard is intended to help our sellers better understand and discuss the importance of having the right systems infrastructure to take advantage of cloud computing. The intended audience for this is C level executives, line of business execs, high level IT managers, and our IBMers that are particularly on the client teams. This whiteboard is not meant to be a presentation but rather to set up an environment where you can have a meaningful discussion with your client that not only helps them understand the key challenges and issues with respect to the cloud but also enables you as a seller to describe at a high level why and how System z is the best platform for them to be able to build a cloud infrastructure. And that's for specific workloads and why System z is the key element in their cloud initiatives.

This whiteboard can take you through most of the issues which might come up, but it also can be very effective if you choose to tailor it to meet your own client's unique needs. So, after you've learned and practiced the whiteboard, feel free to modify it to fit your customer's specific situation.

While this whiteboard was developed with the sales rep in mind to do the presentation, depending upon the customer it might be useful to have an architect in the audience in case the discussion turns more technical. Finally, this whiteboard is meant to be used early in the decision process. This will allow you to have a better chance of influencing the decision to use System z as part of your customer's cloud implementation.

Pay attention to how I talk through the whiteboard and then start by either downloading the coaching guide or the sales kit from the Whiteboard Sellers Workplace and begin presenting this whiteboard to your customer. Good luck and good selling.

The journey to cloud computing starts with virtualization. As a base platform that many clients have already started with virtualization on their x86 platforms, for example on Linux, it gives you a base to be able to abstract your existing infrastructure and be able to start your journey towards the cloud. Most cloud initiatives are focused around the applications or the workloads. So, what we typically use is a workload-centric and what I'll talk about in the future is a fit for purpose algorithm to help clients understand the benefits that cloud computing can offer to a specific workload and why platform selection is an important criteria. So the first question that comes up when people who have started with x86, especially VMware virtualization and beginning the journey toward cloud computing is how do I get there. Shortly after they believe that they can achieve benefits from using cloud computing, the question about service requirements comes up. And by service requirements, I mean things like service level

agreements, or qualities of service. Many applications that start on a fairly simplistic platform with very simple requirements can grow quickly into an enterprise-wide mission-critical application that accesses sensitive data and so qualities of service becomes important.

Another consideration is scale. As the number of users for this cloud application grows, how many virtual machines are required to meet the capacity requirements and the throughput that this application requires? So if you look at the applications, their initial use might be in an application development environment. But then what might happen is questions about availability. Is this a 24-by-7 application? Also performance: What kind of response time is going to be expected for these users? In an internal private cloud that's only used internally for your developers, performance may not be a key issue. But if you look at the enterprise-wide requirements once you deploy this in production, all of a sudden performance and service level agreements are critical.

What about security? Many enterprises have shown the risk of having their data exposed publicly on the cloud and if it's compromised, those intrusions become front page news.

What about efficiency? What are the cost parameters in terms of being able to scale up this environment to be able to meet the kind of throughput requirements that your customers are expecting?

Openness and flexibility: Many clients are tempted by the promise of public cloud computing but also those tend to become lock-ins for a specific vendor. So while we might think that today's cloud provider is going to be there into the future, what happens if a new technology comes available? You're locked in and not able to exploit the benefits of the new providers or be able to compose multiple providers to provide the best solution. So what about the ability to compose multiple vendors in an open and standard-based manner?

And then finally, cost: Who has the most cost effective solution, both for an entry cost and as you grow through the evolution of your cloud workloads?

So let's look at the System z and how it addresses these service requirements and these requirements for scalability. The System z is a very resilient platform and I think as many of you understand it has efficiencies in how it delivers that resilience, up to three times more efficient in providing the infrastructure and the ability to maintain availability levels over the course of years.

And on top of availability in many environments an increasing requirement is for disaster recovery. In many zEnterprise environments disaster recovery is provided for 100 percent of the infrastructure resources with examples like of our GDPS solution. In many x86 or distributed environments, that could be as little as 20 or 30 percent. And in the case of z, the clients exercise each year that they're able to

not only replicate but also be able to restore the environment back to its original state with recovery time objectives that are measured in terms of hours, zero downtime. Not only is the processor and infrastructure availability measured in years, but also the ability to dynamically relocate a running virtual server with things like Livecast Relocation is another example of how System z is designed architecturally to avoid any outage.

So as an example of how a client is achieving this type of availability, let's look at ZIVIT in Germany. They have more than 400 _____ systems today on z/VM and they're able to leverage the availability and reliability of the platform to achieve recovery of any failure through co-automation in less than an hour with zero data loss.

Let's talk about performance. Not only is the ability to meet stringent SLAs in terms of response time critical but also to be able to maintain those response times as the system and the workload grows. System z typically runs above 80 percent in terms of utilization, where for the distributed systems those are typically at least two to three times less than that.

One example of a client achieving significant volumes is Every(?) in the Nordics where they're running over 100 million transactions a day but they push that during the peak holiday season to over 128 million a day. Their quote was they handle this peak in transaction volumes without breaking a sweat. So, our clients are able to leverage the platform for dynamic growth in their workload volumes and continue to manage their client expectations for a responsiveness and scalability without a hiccup.

Predictability: Your clients expect the same experience, no matter what time of day or day of the week or what month of the year they access your system, especially around the holiday season. Now, you don't want to have your enterprise being overloaded and your users becoming dissatisfied. So, although the cloud offers very attractive pricing models, can those cloud computing cloud providers meet the kind of response time that your business requires, especially in these high volume times? System z, not only are the microprocessor speeds are unparalleled in the industry but also our I/O subsystem has bandwidth that is measured up to five times what's available on the distributed systems. So, for data intensive workloads that I/O subsystem is nice to be able to provide that level of predictable throughput that, again, your business requires.

What about security? One of the primary considerations that many enterprises address when they look at cloud and concern themselves with not only is the security built in but it's pervasive. Security is designed into the architecture of the z ecosystem so things like hardware encryption is an assumption. It's a given. The ability to protect your data through hardware and trusted key(?) use(?) both within the system itself or on the storage and be able to manage those keys effectively and securely, again the system is designed 100 percent versus the kind of best

efforts that typically accompany many distributed systems. As well as performance: Because we designed the encryption hardware into the processor and into the I/O subsystem we're able to maintain high levels of performance by using hardware not only for security but for scalability, where in many cases performance of a function like SSL significantly impacts the CPU from a distributed system that's trying to provide the level of performance that you need.

As an example of how some of our -- one of our clients is achieving the benefit of this hardware encryption solution, if you look at Vantive, a US company, they run 19 billion credit card transactions a year and with the hardware crypto they're able to do over 100 million cryptographic operations on a daily basis and they see no limit to the scalability of that solution. So, for them, for their purposes this allows them to seamlessly leverage hardware encryption to keep up with the demands of their business and improve the reliability of their offerings to their clients.

Efficiency: Think about the time it takes to provision new capacity. In the System z we have a feature like capacity upgrade on demand that can be turned on when you need it and is operational within minutes. Compare that to the time it takes to provision new servers and procure the software and deploy them in a private cloud with distributed systems.

An example of a client that's been able to realize the benefits of the cloud technology on zLinux is Nationwide. By using pattern technology, they're able to simplify all of the application development platforms that they support and be able to have the application developers create job applications on the xLinux platform and for the zLinux platform and use a fit-for-purpose model to deploy them where they make sense based on a workload profile. And as an example, they were able to deploy 15 new applications into their cloud in one weekend. Again, dramatic savings in the efficiency of that application development process and the speed of deployment.

Simplicity: One of the key benefits of the zLinux environment is that we can grow dramatically on a single footprint. Many of our clients tell us that they have constraints within the data center, limits on energy supplies, as well as the personnel that's available to manage the infrastructure. With zLinux on z/VM we're able to scale to hundreds or thousands of virtual servers on a single footprint, whereas in distributed systems you typically tend to need to provision new server footprints, which change your energy consumption, change the number of system admins you require in order to be able to maintain the levels of support that your users expect.

So for a couple of examples in terms of the data center simplicity and the savings, let's look at the Japanese supermarket Asumia(?) where they were able to consolidate 115 distributed servers onto a single server footprint, which reduced the amount of square foot in their data center from 1,800 square feet down to 180 square feet, a 10X savings in floor space. And then Nationwide, by going to this

cloud solution, the overall _____ and floor space requirements were reduced by 80 percent and saved them an estimated $46 million over the previous distributed environment, pretty dramatic savings.

Administration and management costs, total cost of ownership: We have studies that show that running hundreds of thousands of virtual servers with Linux under System z is dramatically more efficient and your people are more productive than on distributed systems, partially due to the fact that we've got reduced numbers of images and simplified software license management.

Finally on OPEN and flexibility: Linux. Linux is a great equalizer across multiple architectures. As you develop Linux applications, they may start on x86 or even on Power. That Linux application can run on any platform. And so, by having Linux on System z you're able to get the best of both worlds. You get that OPEN operating system that is supported by many Linux distributors like Red Hat and SuSE available on the System z with the appropriate device drivers and kernel mods to achieve the value propositions of availability, performance, scalability on the z architecture.

There are also many other examples on Linux in terms of programming languages, scripting languages, tools that are available on the Linux for System z as well as thousands of ISV applications like SAP and Oracle which are typically found in many large enterprises. All of those are fully supported on the System z Linux system. Again, the primary benefit here is that you can deploy on one single System z versus many distributed systems for the same level of throughput and the same level of availability of the applications. So your applications run on the platform which best fits the workload profile and you can then take advantage of the System z qualities of service.

So I think if you look at the overall picture you'll be able to convince your client that System z is truly the most cost effective platform for deploying your private cloud or becoming a member of your hybrid or heterogeneous cloud infrastructure that you can take advantage of those qualities of service for those workloads that require and demand them.

There are many options to get started with our business class machines with Linux on System z that are competitive. And as you grow, you improve your overall total cost of ownership on the management side as well as on software licenses.

So the net result is as you grow, once you get beyond a few hundred virtual servers, the cost advantages of a public cloud, which is basically flat as you grow the number of VMs compared to you'll see with x86, which starts out at a higher cost and then becomes more competitive growing over time to the System z, which eventually gets to a point where it is the lowest total cost of ownership for

any option, public cloud or distributed private cloud, that you can find in the industry.

So, as the System z is brought into the cloud, you're able to achieve the kind of volumes of service that your enterprise requires, at the same time being able to offer the flexibility across all the dimensions that many of our current private clouds and what we're calling next gen cloud version 2 are demanding in terms of the requirements. And there's a study from Forrester that articulates many of the qualities of service that these next gen clouds are starting to realize are critical for running enterprise-wide.

So, in summary, System z addresses all of these service requirements as well as the scalability in terms of the most cost effective solution for private or hybrid clouds and we believe that the case is pretty compelling. And what we recommend is that you work with your clients to schedule a Fit for Purpose Workshop for those workloads that are matching these types of requirements. So as you identify the need for availability, performance, scale, efficiencies, etcetera, the Fit for Purpose Workshop will help you understand the tradeoffs between what's available on the distributed systems and what's available on System z.

With that, I'd like to close the discussion and wish you all well. And I'm always available for any questions that may come up or go through the website, the Cloud System z Whiteboard, to provide the answers you're looking for. Thank you.

END