

0:00

Guest.

0:02

He is the VP and GM of our database and business critical apps business.

0:08

I call him the wise one because he's always calm, never gets rattled, and he can talk databases nonstop.

0:15

He was at Oracle for over 14 years working in the kernel development and the server technologies team.

0:21

So please help me welcome Bala Kuchibhotla. Bala.

0:31

Hi, Bala, welcome.

0:32

Thank you.

0:33

How are you?

0:34

Well, Bala, we know that mission critical workloads are a very important workload that runs on Nutanix today.

0:40

So what's your perspective on the scope of this opportunity for our customers and for us?

0:44

Got it.

0:45

See, Monica, we represent business critical apps as well as databases and I manage it.

0:50

To me, business critical apps almost constitute 70 to 80% of the budget, right.

0:56

So it's funny to say that as a kid, my mom used to give me a candy to sort out BCA as ABCs, right?

1:04

But what I realized in enterprises is there are billions of dollars to sort out BCA problems. And we would like to make them is as easy as ABCs.

1:15

Right, now let's talk about data, right?

1:16

We both come from the database world.

1:18

We know that data is doubling every two years.

1:21

Everybody wants to mine insights from data, but I think there's some challenges for enterprises to actually get to the data in real time and be able to make smarter decisions.

1:29

Absolutely.

1:30

See, for business-critical apps, data is gold, right?

1:33

But to get to the process goal, we go through a lot of challenges.

1:37

Similarly, to get to the data, there are a lot of processes that actually we have to cross through the hardens to get to that kind of data, right?

1:46

The same thing that happens with it.

1:48

So what are those four pain points that we talk about?

1:50

Data to get to the process data, right?

1:52

So one is the provisioning.

1:54

I still see provisioning is one of the biggest pain point that it takes like three to five weeks to kind of even create a database for some of our big customers and it kind of spans across multiple teams.

2:06

The second problem that I see is the proliferation, right?

2:10

Proliferation is like creating copies of your databases.

2:14

This is like our copy paste problem.

2:15

Like I don't see anyone, I don't think anyone in this audience has never done a copy paste on their laptop, right?

2:21

Same thing happens for databases in your data center.

2:24

Only thing is this seemingly simple innocuous operation of our daily life becomes the most dreaded, complex, long running error prone operation.

2:33

I've seen customers spending like 3 days and having a dedicated job profile for cloning databases and then lineage is even more bigger pain for that.

2:41

The third Point is data production.

2:43

If you're talking about the business-critical apps and databases, mission critical databases, it's an absolute must that you must have your ability to go back to any point in time to protect your data.

2:55

And then it could be malicious data, someone, someone corrupted the data or it could be your own Butterfinger that drop the data kind of stuff.

3:02

Your ability to protect data is absolute must.

3:05

And the last one is the patching pain that we talk about it.

3:08

If you are managing databases like more than 100 databases kind of stuff, even even a few dozens, it's an absolute must that you have to have complaints and standardization to how the database are under control, kind of stuff like this is what we talk about.

3:22

I see.

3:22

So you mentioned the four PS, the challenges.

3:24

So what is Nutanix doing to help our customers get past these challenges?

3:28

So that is the journey that we started.

3:30

You know Nutanix is made its name and fame for simplifying the storage, simplifying the compute, simplifying the network and what we call as infrastructure as a service and we made a ton of name and fame and then we are today about 1.5 or \$1.6 billion bookings around that kind of stuff.

3:46

We want to extend the design and delight of for of this infrastructure that we have created for databases and that is what we did for with ERA.

3:55

We went, we went to GA last year at the .Next and EMEA conference.

4:01

And I'm actually very happy to kind of convey that we are seeing a huge fraction for ERA.

4:07

And we were able to cut down database provisioning times from days to minutes. And some of our customers able to cut down backup times from hours to seconds.

4:17

And literally they can create clones and refreshes them with a single click.

4:21

And these clones are 0 byte warhead clones kind of stuff.

4:24

And obviously we are following the big problem of Oracle patching and other for other databases.

4:29

This is what we did for Era and we got to do more in kind of stuff.

4:33

Yeah, I was going to ask you this seems like it's a fantastic product to simplify database operations.

4:37

So where are we headed next with this product?

4:39

Yep.

4:40

So I've been talking to a lot of enterprise customers in last one year and the journey that we started with ERA was ERA was running on a single cluster as of today.

4:49

But obviously, if you're talking with customers with hundreds of databases, they want ERA to be managing multiple clusters.

4:56

So this is our first journey towards the next big project for the ERA, which is we call ERA 2.0.

5:03

We are enabling like PC, like Prism Central, the way that manages multiple PES, Era will manage databases across clusters, ability to provision, ability to patch, ability to do that data management across the clusters.

5:15

This is the marking feature of era 2.0. But the second stuff that we are trying to do is obviously if you're managing data, if you're managing mission critical database systems, you got to have role-based access control, right?

5:29

And we're not typical role-based access control, but we are actually creating a unique feature here, something called data access management, like dam, right?

5:38

Similarly similar to how dams control the water flow in the river.

5:42

And then try to give the data to whom and who, who gets what kind of stuff we need to control the data.

5:47

We need to control the data from production to the QA staging, how they get it, where, where do they get it, all the complex kind of stuff.

5:53

That's what we're going to do in the data access management.

5:56

And lastly, we kind of promised more engines for error.

6:01

Obviously, there's a lot of demand and we're going to announce SAP HANA on ERA very soon and actually it will be the tech preview in the coming months.

6:08

And also we're going to extend ERA for no SQL space and we're going to start with MongoDB.

6:13

That sounds great.

6:14

Could you give us a sneak peek at some of this?

6:17

Happy to do that.

6:17

Thank you.

6:18

OK, let's do it.

6:21

All right, OK, let's see what we have.

6:24

I'm going to show three things which is rollback, roll, roll based access control, as well as the multi-cluster demo kind of stuff.

6:32

See we used to have error with only admin functionality.

6:35

Now we can do with users obviously.

6:38

And I'm logging in as Bala.

6:42

As you can see, this is our popular dashboard for Era.

6:47

And then let me show how what does it mean to take a multi-cluster demo for databases.

6:52

This demo you have seen it last year, but I'm going to show quickly and then launch on the other new features.

6:59

So let's pick on Oracle 2-node rack instance.

7:01

Like you want to create a 2-node rack instance.

7:05

If you see this is the point that we are talking about, you will be able to choose a cluster on which your rack databases you can create it.

7:13

In this demo I have 3 clusters.

7:15

1 is the Tomahawk prod and non-prod as well as Gemini QA and that's all it takes.

7:20

You will be able to choose the choice of your cluster and say create a database in that cluster.

7:24

If it happened to be AAG, you can even create the database across clusters too, not just rack database like that stuff in single cluster.

7:33

You will be able to spread the database itself across Nutanix clusters.

7:37

Let me quickly go through provisioning demo and then we'll talk about data access management, right?

7:44

So I will create a database name here, Sales database, let me create a system identifier, sales DB. Let me give a password.

7:57

And this is the database parameter profile.

7:59

I don't need to talk about it.

8:00

This is the size of the database.

8:02

We'll do all best practices that are needed for you with just one dimension, 1 parameter size of the database, right?

8:09

And finally, the popular time machine, you will be able to choose the different SLAs for your time machine.

8:15

Here I'm choosing the goal, which is 30 days of continuous recovery region and 30 days of daily backups and so on so forth.

8:20

One interesting thing is when we talk to a lot of customers, they won't be integrate with Service Now or any other choice of their automation tool.

8:28

So they want to kind of pick the APIs or the CLIs and then do what we need to do.

8:33

So this is what it is.

8:34

You don't need to use the UI as is.

8:36

You can just take the APIs, walk through the UIs and generate the APIs as well as CLIs and then just use them in your automation.

8:43

So let me go ahead and provision the database and sorry, I'm actually I need some water here.

8:50

I'm drying out.

8:51

Right there.

If you can help me.

8:52

While it is happening.

8:53

Sure.

8:54

The cluster database is getting created.

8:56

We are creating virtual machines, we are creating storage, we are attaching the storage, we are creating the rack cluster database, so and so forth.

9:02

And we are doing the database right.

9:03

So let me go to the time machine.

9:10

So this is our time machine.

9:13

Most of the people know here and it is, it has been in this particular cluster.

9:17

I've been operating this time machine for more than a year now.

9:21

It has an SLA of gold, which is 30 days of continuous recovery region, 30 days of daily backups and so forth.

9:27

Again, you see this cluster concept is coming into the time machine too.

9:31

You'll be able to see the time machine capability on this particular cluster Tomahawk prod.

9:36

But what it means for this data to be available on a Gemini QA cluster even though I'm the super administrator?

9:43

So if it is a QA cluster to adhere to the Sarbanes Oxley or any kind of compliance requirements, I cannot create the production data on this particular cluster.

9:53

So I can only see the sanitized snapshots.

9:55

So my ability, even though I'm the super administrator, I can only create the data.

9:58

Now how do we do it?

10:00

What are the policies to manage it?

10:01

Let's just go through the administration here.

10:04

You can create masking policies and then see whenever I'm going to create a sanitized snapshot, run this masking script and then make sure that I have 4 copies of that and I can tell where the data needs to be available.

10:16

I can tell that on Gemini QA, I just want sanitized snapshots.

10:20

I don't want any production data to be available.

10:23

And then let's see, I also can pick and choose who can get what.

10:26

In this case, Mr. David can get sanitized copies of this time machine on Gemini QA cluster.

10:33

So let me just log out and log in as Mr. David and see what he can do, right?

10:40

Look, I told you he can talk for hours about databases.

10:43

That's what I said. Bala can do that.

10:45

Oh, here is.

10:46

Here is a view of the time machine for Mr. David, right. The rule, rule based access control.

10:53

Obviously he doesn't have any production databases.

10:55

He does have access to the time machine.

10:57

Let's look at the time machine.

10:58

What he's trying to do it.

11:00

Obviously he doesn't have any access to the production data.

11:03

He only has access to the sanitized snapshots.

11:05

That's your rule based access.

11:06

Access.

11:07

Yeah.

11:07

And he can select any snapshot and then basically say that I want to create a clone of it.

11:12

Let's just create. Even his privileges of what he can do is also restricted.

11:16

He cannot do backup management or anything of that sort.

11:19

He can just do create a clone and then let's try to create a clone and rewind the demo and standard input parameters.

11:28

We have seen it last time last year.

11:30

And then let me give you the password again.

11:33

You can do the API equivalent and then clone.

11:36

This is pretty much it.

11:37

What it takes for a QA administrator to consume the production data in a sanitized shape, and they can refresh the data at will too.

11:44

Once they have copies, they can always do it.

11:47

And I will sum it up with one click demo truly for patching and then we'll take it.

11:52

So let's go to databases, go to the clones, and as a QA administrator, I have been running some databases and let's see, my manager asked me to do a update management.

12:03

I want to do the change management testing.

12:05

All I need to see is there is an update available for my server.

12:09

I can click on it and Era can give you the latest profile that you can apply to your databases.

12:14

And here I'm on the July PSU with one offs and I can go to the October PSU and then I can just say patch now or patch later and submit it.

12:24

Thank you, Bala.

12:25

I think this is a great example.
This is what it takes.

12:26

Yeah, this is a great example of us going beyond a platform, which is the resilient platform for data services, but providing the value for database administrators to simplify the operation operations of databases.

12:39

But why don't we walk back and sum it up?

12:41

So is there anything else you'd like our audience to remember before?

12:46

Basically, I, I, I think that we are the true honest database virtualization, and we are clubbing with cloud, cloud, cloud experience for the databases.

12:54

I can feel confident that we can do it at half price and provide 2X value for our customers.

12:59

Sounds great.

13:00

Thank you so much, Bala.

13:01

I really appreciate it.

13:02

All right, all right.