



IBM Cloud Data Warehouse



Netezza on Cloud

Internal Enablement




Miran Badzak
Product Lead
Data & AI Cloud Databases

Netezza on Cloud

/ne-teez-a/

*a high performance, cloud data warehouse
running the next generation **Netezza database engine**
delivered through **Cloud Pak for Data***

IBM Netezza on Cloud




Easy to get started

Minimal touch to provision, scale, maintain



Runs Netezza

The obvious choice for Netezza upgrades



Scalable

Independently scale compute & storage



Highly available

Managed infrastructure with Red Hat OpenShift



Secure

Private deployment with encryption for data in motion/rest



Reliable

Backups, replicated across AZs

**Easy to get
started**
Minimal touch to
provision, scale, and
maintain

Certified cloud deployment

IBM takes the guesswork out of deploying Netezza to the public cloud. We've tested various IaaS configurations to determine the combination that gets the best price/performance.

Deploy to any public cloud data center on the planet

Choose the data center or availability zone closest to you. Set the number of compute units and amount of storage required to run. The service will take care of acquiring the hardware from the cloud vendor and laying down the software to run your database service.

Scaling

Increment or decrement the amount of storage or compute through a UI or REST API. The automation will acquire or release hardware as necessary. *Coming 2H 2020.*

Maintenance

Product updates pushed through our platform on a monthly cycle. Deploy to your environment when you are ready.

Runs Netezza

the obvious choice
for Netezza
upgrades to the
cloud

*Simple lift and shift to cloud
with
nzbackup/nzrestore*

100% compatible with existing TwinFin, Striper and Mako workloads

It's the same database engine you love, reconfigured and retooled for cloud.

Seamless lift and shift

It's not a migration — it's a lift and shift. `nzbackup` your appliance to IBM Cloud Object Storage or AWS S3 — `nzrestore` into your cloud data warehouse, or `nz_migrate`. Done.

Separation of storage & compute

Scale these two sets of resources independently — configure a compute-dense, storage-dense, or balanced system.

Use your nz-utilities in the cloud

All the same utilities that you've grown to love, now available as REST API calls in the cloud.

Access your cloud data lake with external tables

Reach into your IBM Cloud Object Storage or Amazon S3 buckets directly via Netezza external tables.

Scalable

with independent
scaling of storage
and compute

*Scale up your compute
during peak demand*

*Scale down your compute
when demand falls*


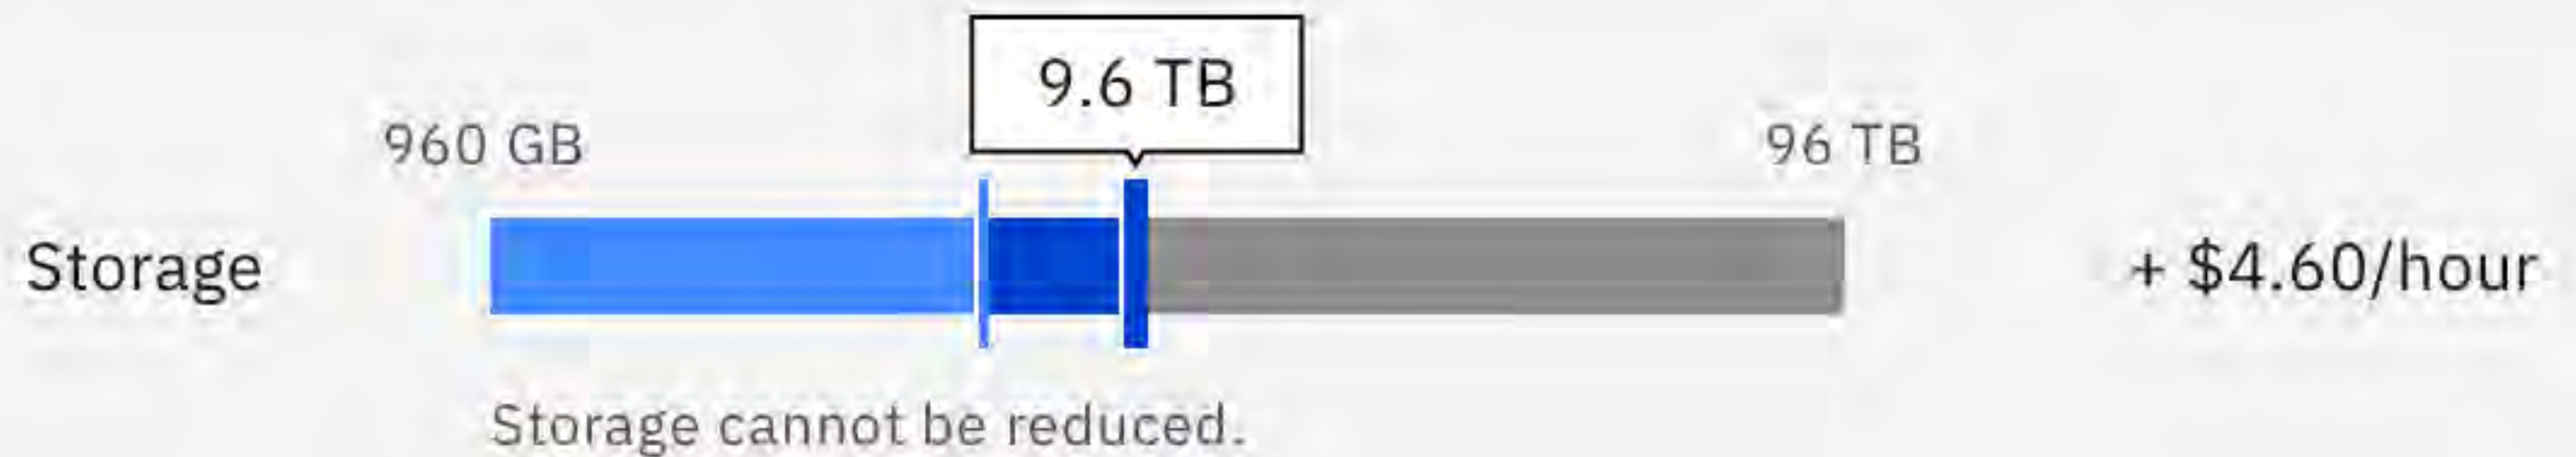
*Ramp up your data storage
as your data needs grow*

API and UI driven

Available in 2H 2020

Compute & storage

Show settings ×



Total cost change: + \$13.48/hour or \$9,840.4/month

Notification

Notify me when scaling is done:

Email Slack

Show API ▾

Highly Available

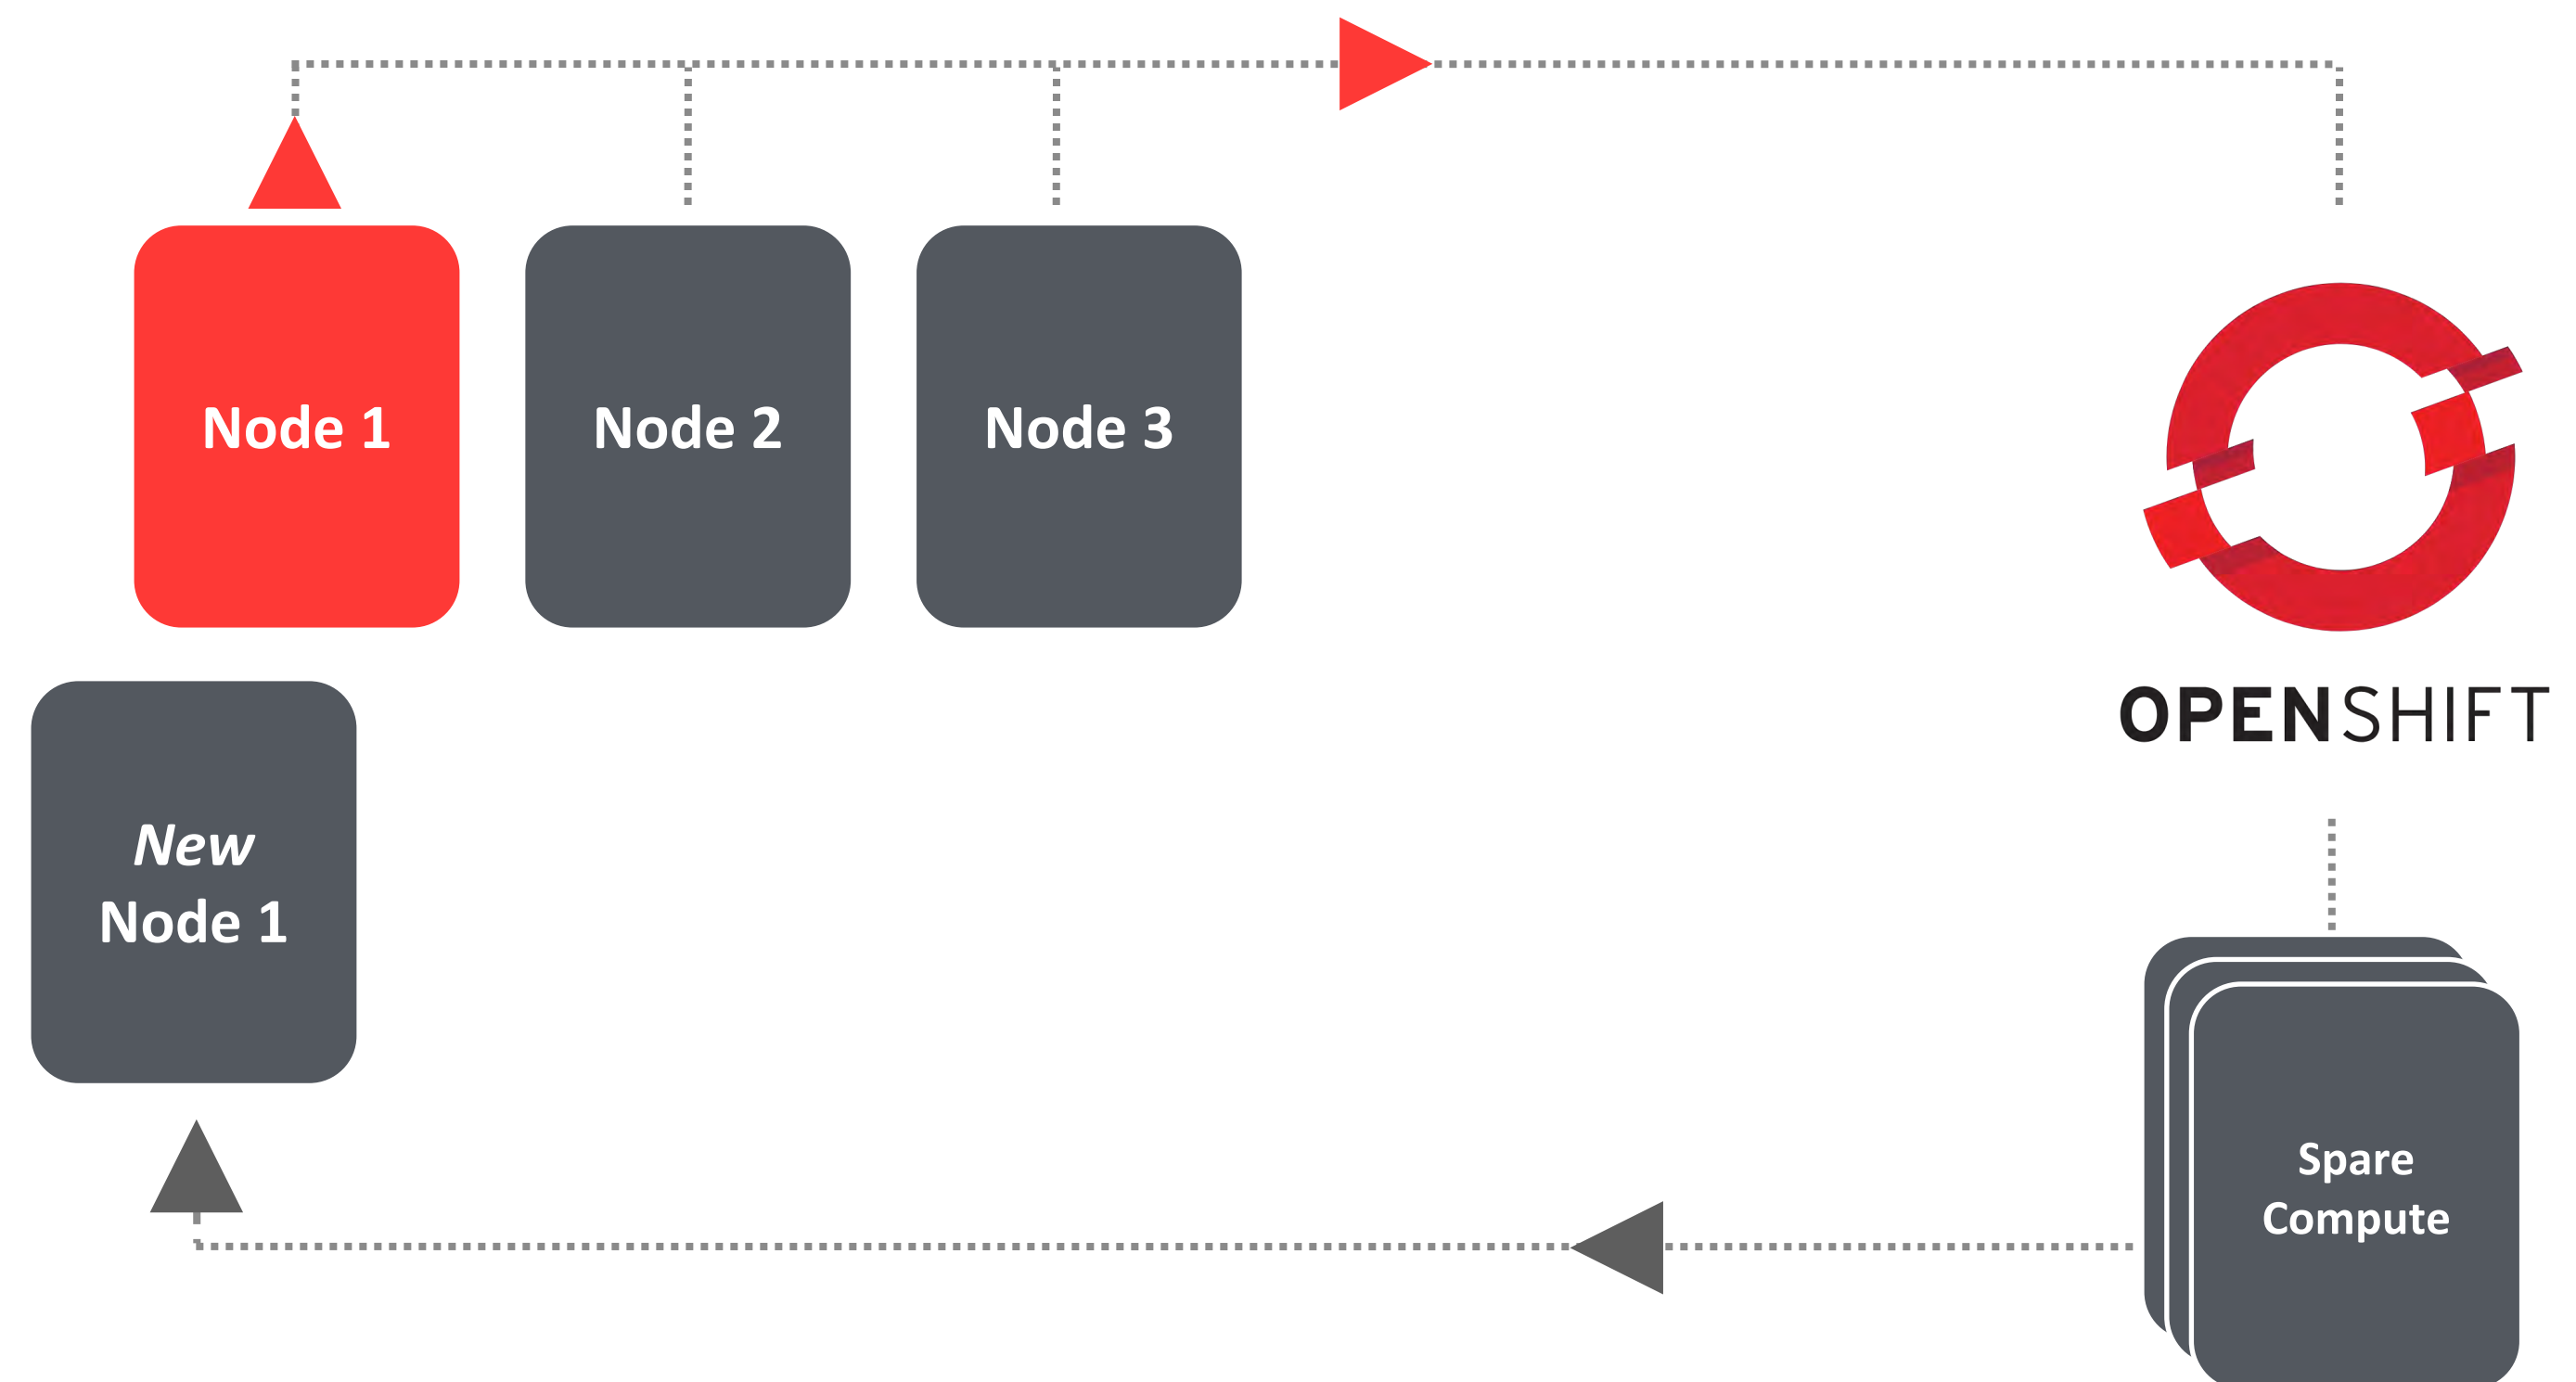
with Red Hat OpenShift

Compute

Containerized database engine deployed to compute nodes, **orchestrated** by Red Hat OpenShift. **Failure detection and recovery** provided by OpenShift – once an unhealthy node is detected, it's swapped in with a new healthy node from a standby pool.

Storage

Cloud-managed SSD storage in a redundant, performant, high available configuration.



Secure

Private deployment
with encryption for
data in motion/rest

Private network deployment

Cloud Pak for Data and your Netezza data warehouse are both deployed inside the customer's own cloud account, giving them full control over the underlying infrastructure, resources, networking and security

Encryption for data in motion

Secure, SSL-enabled connectivity for data moving over the wire

Encryption for data at rest

Encrypted storage volumes (by cloud provider) ensure your data is protected at all times

Reliable

with nz-based backups fed to object storage

Schedule your backups to run when it's most convenient for your business

Partially or fully restore your database from backup

Scheduled or ad-hoc backups

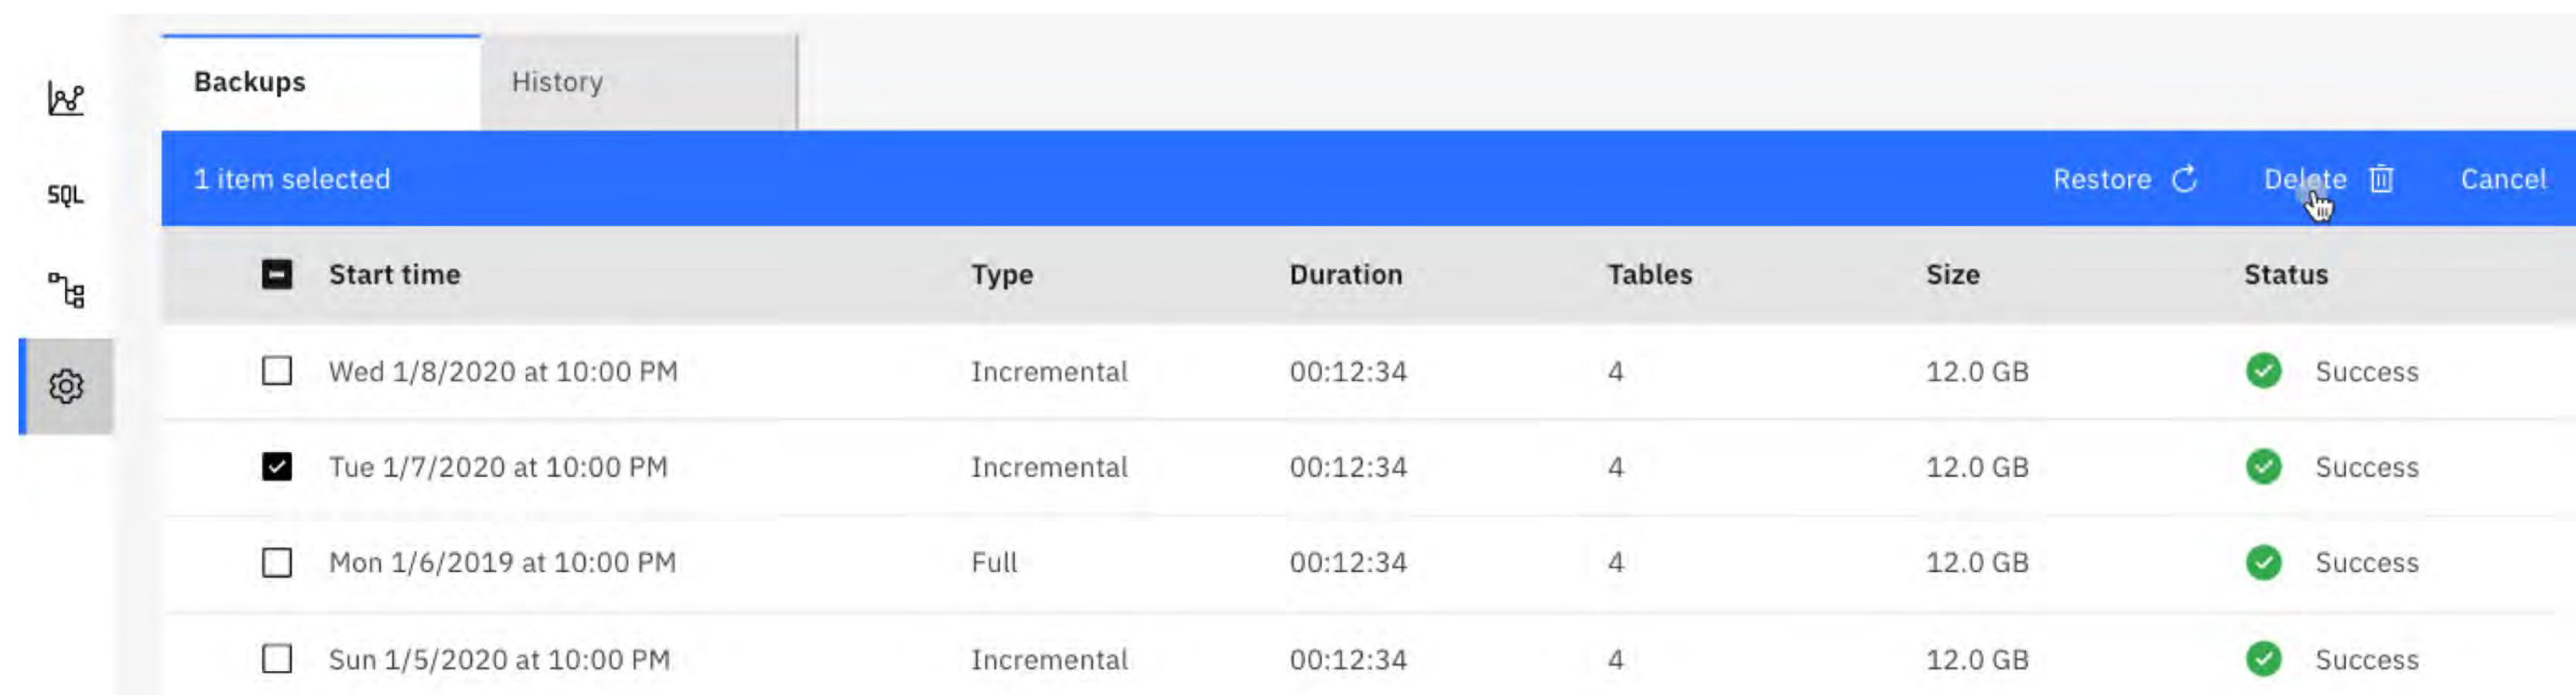
Choose the time of day for your backup to occur or run a backup on-demand via the console or REST API. For scheduled backups, we take 1 full and 6 incremental backups every week.

Full or partial restore

When needed, you have the option to restore the full database or select specific schemas/tables to restore.

Backups stored in object storage

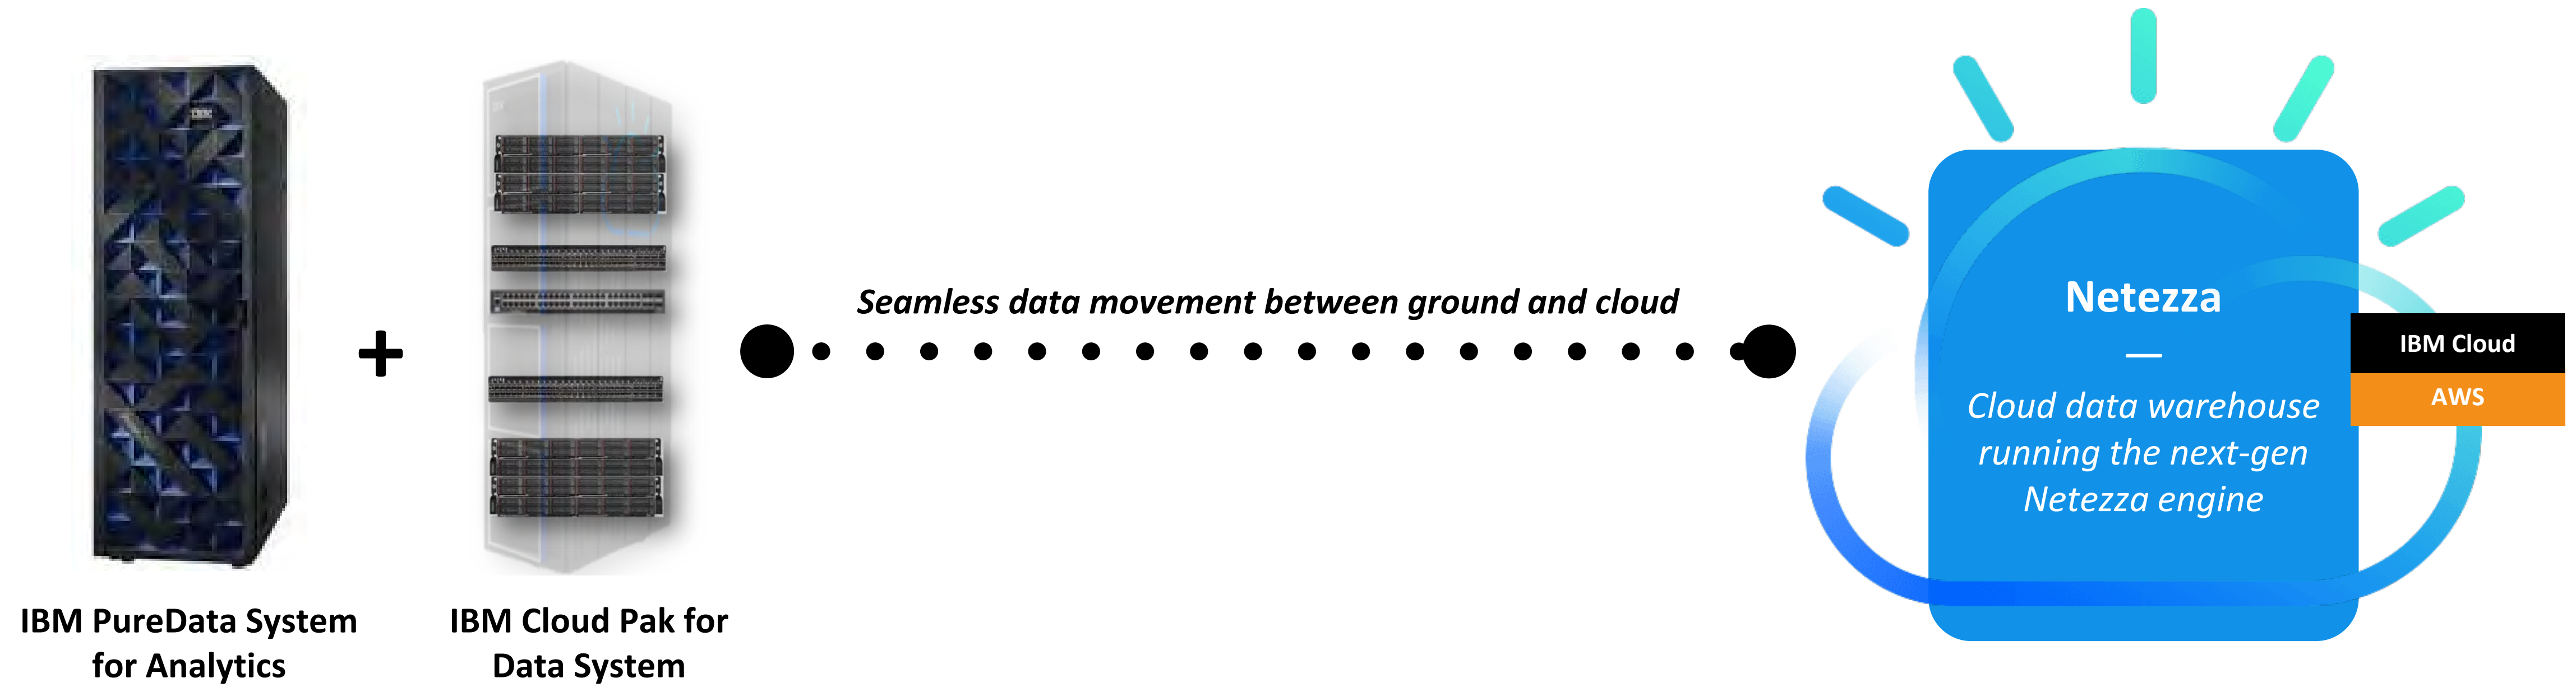
Backups are stored in either IBM Cloud Object Storage or AWS S3, giving you an inexpensive data store with infinite scalability and cross-AZ availability



The screenshot shows a web interface for managing database backups. It features a 'Backups' tab and a 'History' tab. A blue header bar indicates '1 item selected' and contains 'Restore', 'Delete', and 'Cancel' buttons. Below is a table with columns for Start time, Type, Duration, Tables, Size, and Status. The table lists four backup records, all with a 'Success' status.

Start time	Type	Duration	Tables	Size	Status
<input type="checkbox"/> Wed 1/8/2020 at 10:00 PM	Incremental	00:12:34	4	12.0 GB	Success
<input checked="" type="checkbox"/> Tue 1/7/2020 at 10:00 PM	Incremental	00:12:34	4	12.0 GB	Success
<input type="checkbox"/> Mon 1/6/2019 at 10:00 PM	Full	00:12:34	4	12.0 GB	Success
<input type="checkbox"/> Sun 1/5/2020 at 10:00 PM	Incremental	00:12:34	4	12.0 GB	Success

All Cloud or Hybrid



All in on cloud

Drama-free lift and shift with either `nz_migrate` or `nzbackup/restore`

Dev/test environment

As easy as `nzbackup` to cloud object storage and restore into the cloud data warehouse

Cloud data mart


Localize/isolate data sets for specific lines of business

Disaster recovery for private cloud deployment


Upload backups to cloud object storage service, restore to your cloud data warehouse, and switch over if disaster strikes

Deploying the first cloud data warehouse

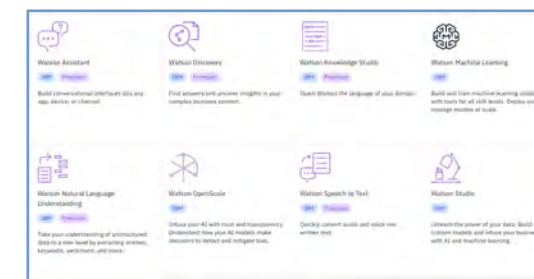
1 Deploy Cloud Pak for Data platform



Start the **Cloud Pak for Data installer** from laptop or desktop

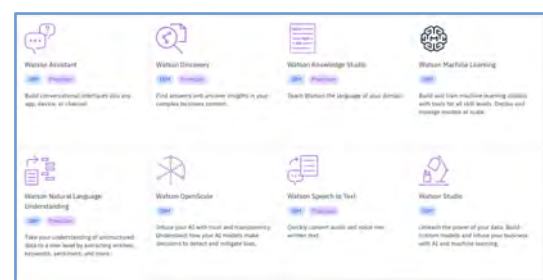


Installer procures the **Cloud infrastructure** and installs the Cloud Pak for Data Software



Cloud Pak for Data is deployed and ready to create a new Netezza instance


2 Deploy Netezza



From Cloud Pak for Data catalog, **choose Netezza** and click Provision



Choose the compute & storage allocation, or start with a template based on existing Netezza footprint (i.e. quarter rack, half rack,...)



Using the customer's cloud account, **provisioner will procure the necessary infrastructure** needed to support the chosen configuration



Netezza is up and running, accessible via console or external tooling

Netezza on Cloud Lift & Shift

Migration CLI

The same tooling you've always used to move between Netezza appliances works with your cloud data warehouse – `nz_migrate`.

```
nz_migrate  
-shost <src_host> -thost <tgt_host>  
  -sdb <src_db> -tdb <tgt_db>  
  -t <table1, table2, ...>  
-suser <src_user> -spassword <src_password>  
-tuser <tgt_user> -tpassword <tgt_password>  
  -cksum fast -genStats Full  
  -TruncateTargetTable YES
```

Backup & Restore

A simple 2-step process to move to cloud. Take a backup of your Netezza appliance, streaming the backup files to your IBM Cloud Object Storage or Amazon S3 bucket. Then, restore to your cloud data warehouse.

```
nzbackup  
-db <db> -u <user>  
  -connector s3  
-connectorArgs ACCESS_KEY_ID=<access_key>:  
  BUCKET_URL=<bucket_url>:  
  SECRET_ACCESS_KEY=<access_key>:  
  DEFAULT_REGION=<s3_region>:  
  UNIQUE_ID=<unique_bucket_id>
```

Then, jump into the **console**, find the backup, and hit **Restore**.

Mass data migration service

IBM Cloud offers the Mass Data Migration storage appliance; Amazon offers Snowball. **Copy** your `nz` backup to the appliance, then **ship back** to your cloud vendor. Backup files will be loaded into object storage (IBM Cloud Object Storage or Amazon S3). Then, **restore via the built-in console**.

Happy beta customer

“I’ve done more POCs than I can count, but moving Netezza on-prem to Netezza on Cloud was the most seamless, free of hiccups, and in exceeding expectations, than any POC I’ve ever done, or even heard of. Add to this IBM’s aggressive, proactive support, and the overall experience was outstanding.”

–David Birmingham
Principal Solutions Architect for Sirius Big Data and Analytics

Happy beta customer

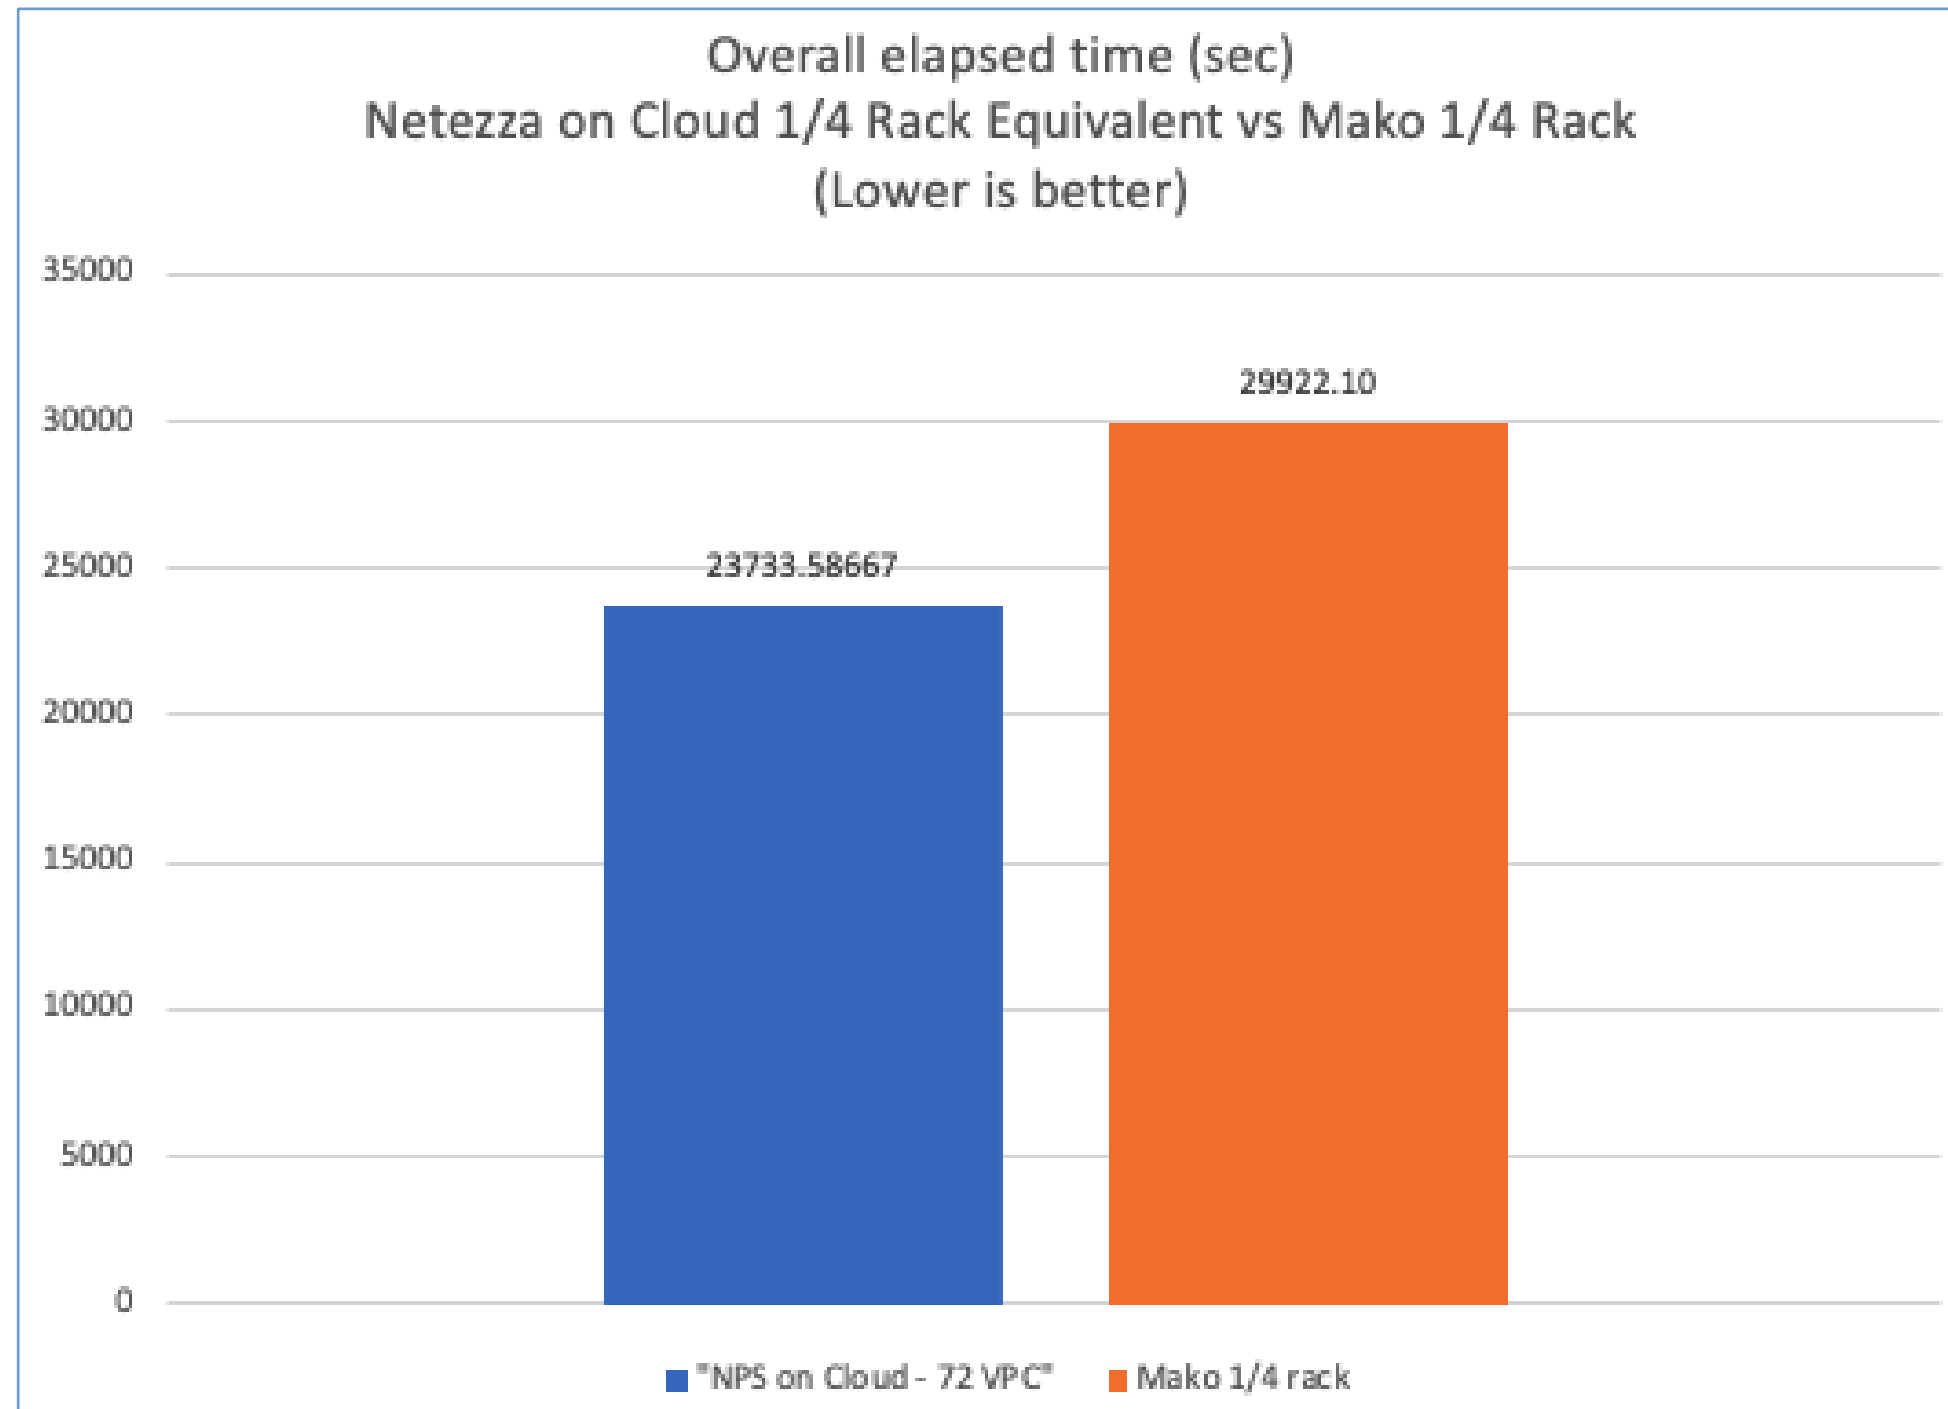
*“Every nz_migrate test I’ve done has worked without issue.
300GB, 1.5TB, 3TB and 6TB”*

–Daniel Hancock
IBM Technical Sales Specialist

How much faster is **Netezza on Cloud** than **Mako**?

Wrong question.

Netezza on Cloud vs Netezza Mako



Netezza on Cloud (\$474,132 3-Year TCO)

- 72 VPCs, 7.2TB storage
- Clients can reduce costs by scaling down when capacity is not needed
- Assuming 3-year Reserve Instance pricing with 20% platform discount (typical) + 86% CTL discount (roughly 50/50 split between infrastructure and license cost).

Netezza Mako Quarter Rack (\$443,345 3-Year TCO)

- Using average discount of 65% and including data center costs + DBA personnel
- Mako priced at \$217,490, including S&S payments

RESULT

Netezza on Cloud is 1.3x faster than a Netezza Mako
using TPC-DS benchmark

How to buy

1

Software Licenses (VPCs)

Purchase a set of Cloud Pak for Data Virtual Processor Cores (VPCs) to run both the **base** Cloud Pak for Data (6 VPCs) and however many VPCs will be need to run **Netezza** on the Cloud (N VPCs)

Software Cost = # of VPCs needed = $6 + N$
(N = VPCs required to run Netezza)

2

Cloud Infrastructure (IaaS)

IBM has certified specific hardware on each supported cloud provider. Infrastructure will be procured **using the customer's cloud account.**

Hardware Cost = (6 VMs for CP4D) + (N compute nodes) + (M storage volumes)

Total cost = Software Licenses + Cloud Infrastructure

Let's configure your Netezza cloud data warehouse

Your Netezza on Cloud is composed of two parts: the Cloud Pak for Data Platform and the Netezza database engine. We'll size both.

New workload or Netezza upgrade | **VPCs** | **Storage** | **CPD license type** | **License discount**
 New or Mako Half Rack - N3001-C | 144 VPCs | 24 TB | CTL Perpetual | 85 %
Cloud provider | **Region** | **Infrastructure payment term** | **Platform discount**
 IBM Cloud | AWS Amazon Web Services | US East (N. Virginia) | 3-Year Reserve Instance (RI) | 20 %

Summary Details

1 Base Cloud Pak for Data

Infrastructure and license necessary to deploy the Cloud Pak for Data and Red Hat OpenShift control plane.

3-node HA configuration

	Qt	Price
IBM Cloud Pak for Data License		
License Part # D264MLL	6	\$6,048
Amazon Web Services Infrastructure		
Cloud infrastructure estimate		\$6,867
Total per year		\$12,915

Summary Details

2 Netezza Performance Server

License and infrastructure needed to deploy your Netezza database engine.

144 VPCs 24TB storage

	Qt	Price
IBM Cloud Pak for Data License		
License Part # D264MLL	144	\$145,152
Amazon Web Services Infrastructure		
Cloud infrastructure estimate		\$155,866
Total per year		\$301,018

Summary Quote

3 Totals

Price estimate for Netezza cloud deployment is an approximation and does not include taxes or additional platform fees.

AWS US East (N. Virginia) 3-Year Reserve Instance

Summary	
IBM Cloud Pak for Data License	
License discount	85%
Total VPCs	150
Price	\$151,200
Amazon Web Services Infrastructure	
Platform discount	20%
Price	\$162,733
Grand Total per year	\$313,933

Size & Price

www.ibm.biz/clouddb
(Netezza > Size & Price)

Start with a brand new configuration or choose an existing Netezza model as a template. When choosing a Netezza appliance as template, sales rep can increase/decrease compute & storage as needed.

Choose between CTL and Perpetual license for Cloud Pak for Data (CDP) license. And, rep can set a discount percentage for that CDP license.

Next, choose between IBM Cloud and Amazon Web Services (AWS). Depending on choice, list of regions will be populated. Different regions carry different infrastructure rates.

For compute infrastructure, user can choose between On-Demand, 1-Year or 3-Year Reserve Instance (RIs). The longer the term, the better the discount on compute.

Data presented as summary view or detailed, full bill-of-materials view. When rep is ready to quote, hit the Quote button for an exact part number, quantity and discount that needs to be quoted.

Sample configurations

DEMO

Netezza on Cloud

	At GA	Post GA
Product	Up to 1-rack Mako equivalent 14-204 VPCs compute cores, 2.4-48TB of storage	Mako 2-rack 4-rack, 8-rack equivalents and beyond (petabyte scale) – will be rolled out starting in July
Scaling	Not available	Independent scaling of storage and compute with automatic data redistribution when number of partitions changes (2H 2020)
Migration	Over the wire: Lift and shift via existing migration CLI tooling or backup/restore (backup appliance, restore to cloud)	Storage appliance transport: leverage services like IBM Cloud Mass Data Migration or Amazon Snowball to lift and shift massive quantities of data
Cloud Support	IBM Cloud and Amazon Web Services	Azure (late 2020, early 2021)

Slack: #winroom-netezza-cloud

IBM