Nova generacija IBM DWH rješenja

IBM Data & AI

Robert Božič
Data & AI Segment Leader, IBM CEE
robert.bozic@si.ibm.com

Simplicity in its Purity !

https://www.youtube.com/watch?v=FFy-HxcNUNI


arnold
bmc

Postoječi DWH appliance na BH Telecom-u
50% kapacitete 1/3 rack-a

	1/3 Rack	2/3 Rack	1 Rack	2 Rack	4 Rack
Servers	3	5	7	14	28
Cores	72	120	168	336	672
Memory	1.5 TB	2.5 TB	3.5 TB	7 TB	14 TB
User Capacity	64 TB	132 TB	196 TB	392	784
Optional Tiered Storage Flash + HDD : GA 1H 2018	TBD	TBD	TBD	TBD	TBD

Slika 56: IBM Integrated Analytics skalabilnost sistema


You’re a data
company now

You need a new architecture to be a data-driven organization: **a data fabric**

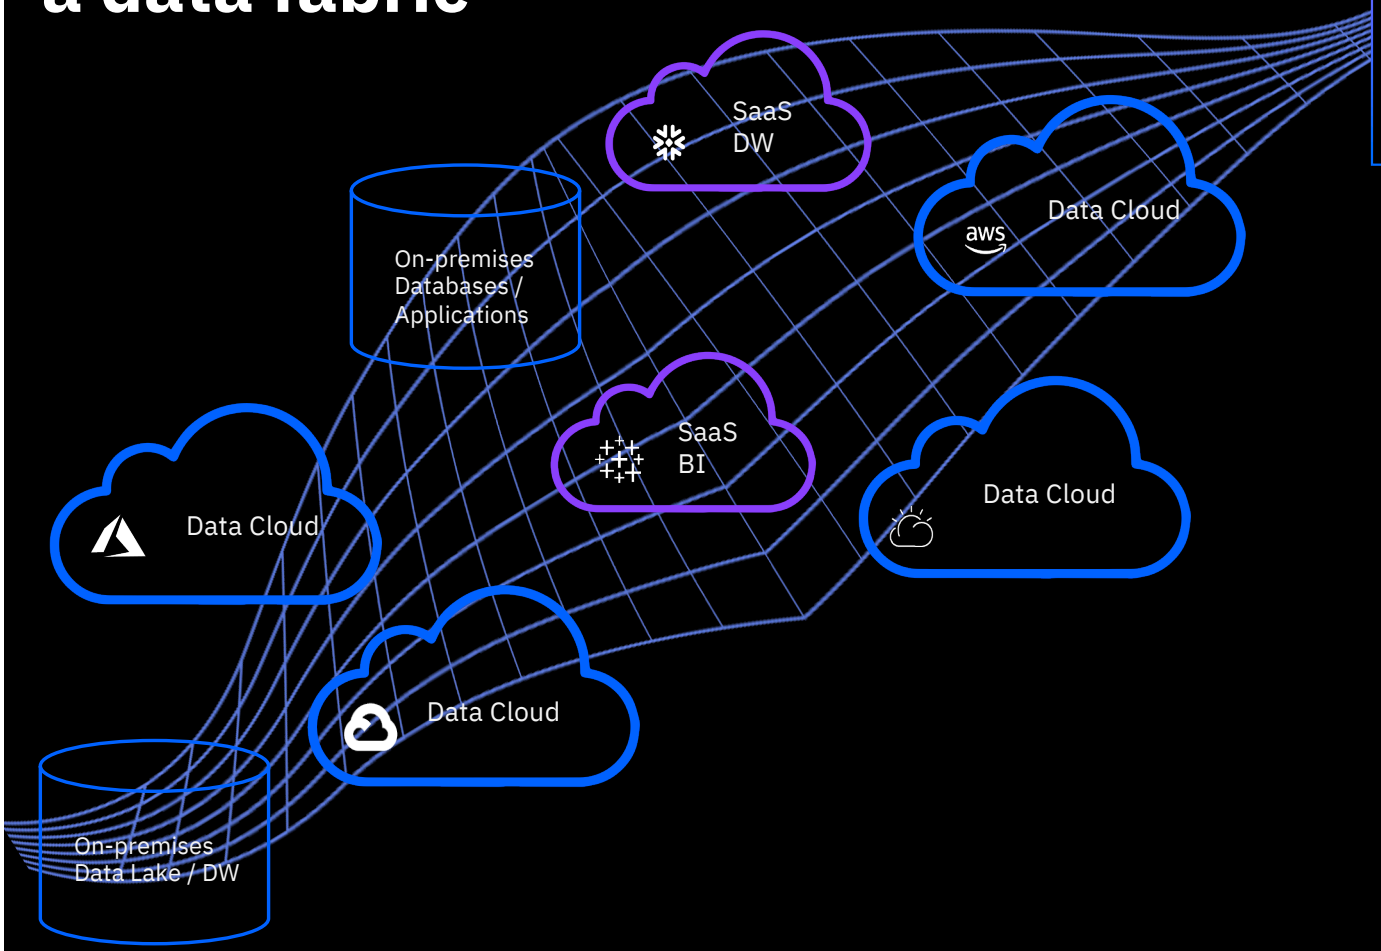

Data Fabric
“Design”

An abstraction layer to share and use data across a hybrid cloud landscape.

Cloud Pak for Data

Simplifies, unifies, and automates your journey to AI

Extended Services
A growing ecosystem of Open Source, Partner, and IBM services

Base Data Services
Integrated capabilities to climb the AI ladder

Cloud Pak Control Plane

Red Hat OpenShift
A streamlined hybrid-cloud

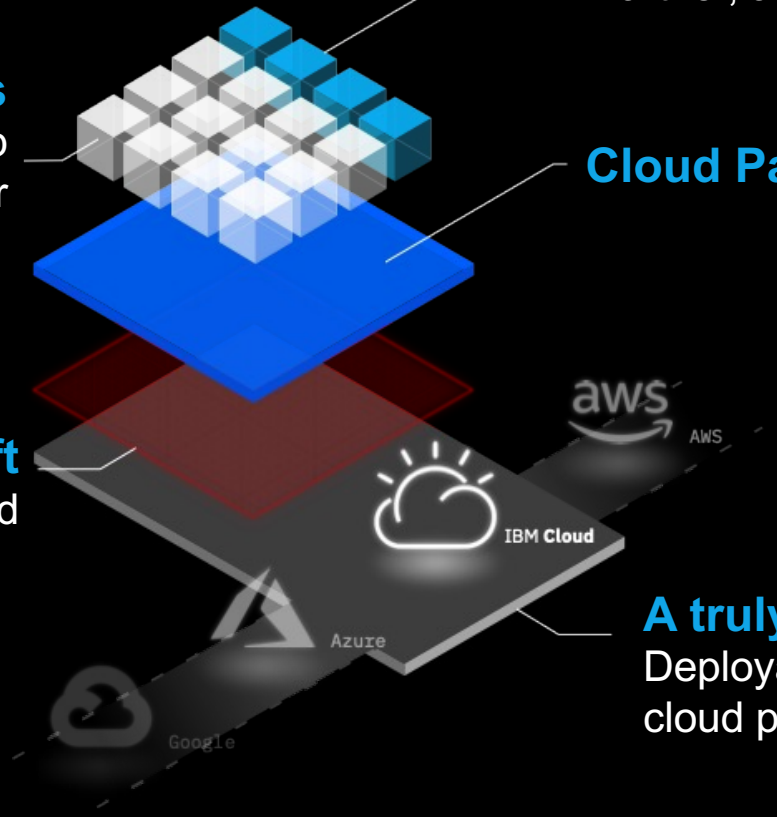

A truly Hybrid-Cloud world
Deployable on-premises or on any vendor's cloud platform; also available as-a-Service

Cloud Pak for Data
Unified, modular, deployable anywhere

App Developers | Business Analysts | Data Engineers | Data Stewards | Data Scientists | Business Users

Integrated user experience

Extensible: APIs, partner ecosystem, accelerators, and solutions

Data management

- Watson Query (data virtualization)
- Db2 and Db2 warehouse
- Netezza warehouse
- OEM databases (Cloudera,…)

Data fabric

- Watson Knowledge Catalog (governance)
- Watson Studio (data science platform)
- Match 360 (self-service data matching)
- DataStage (data transformation)
- Data Replication (data movement)

Business analytics

- Cognos Analytics
- Planning Analytics
- Cognos Controller

Customer Care

- Watson Assistant
- Watson Discovery

Cloud Pak core services
Security, Administration, Operations

Red Hat OpenShift

Netezza is a key performant data source to any Data Fabric Architecture

Data Sources, Types and Domains

Any Data, Any Cloud, Anywhere

Data and AI Outcomes

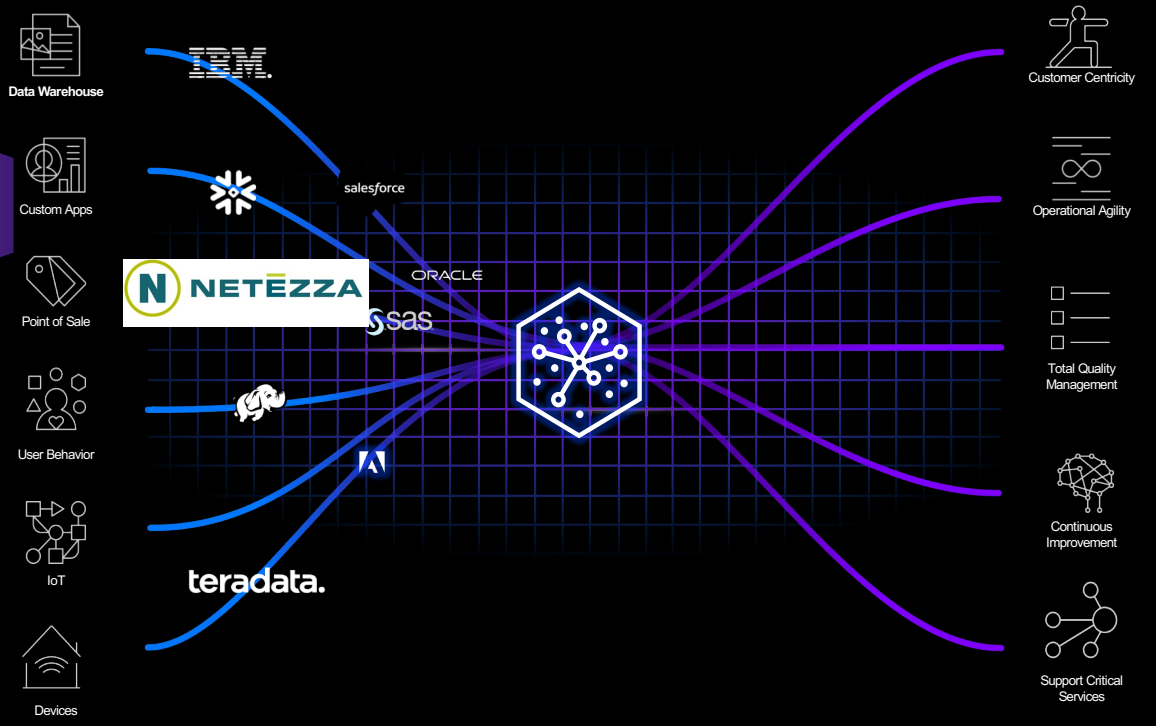

IBM Netezza's revolutionary approach

The Appliance for Analytics

Simpler, Faster, more accessible advanced analytics

"This is what Netezza has done in the data warehousing market: It has **totally changed the way we think** about data warehousing."

- Philip Howard, Bloor Research

No **RDBMS** inside !!
No Tuning required !

Hardware Acceleration
Simplicity

Anno 2016 – What about the Next Generation Analytic Appliances ????

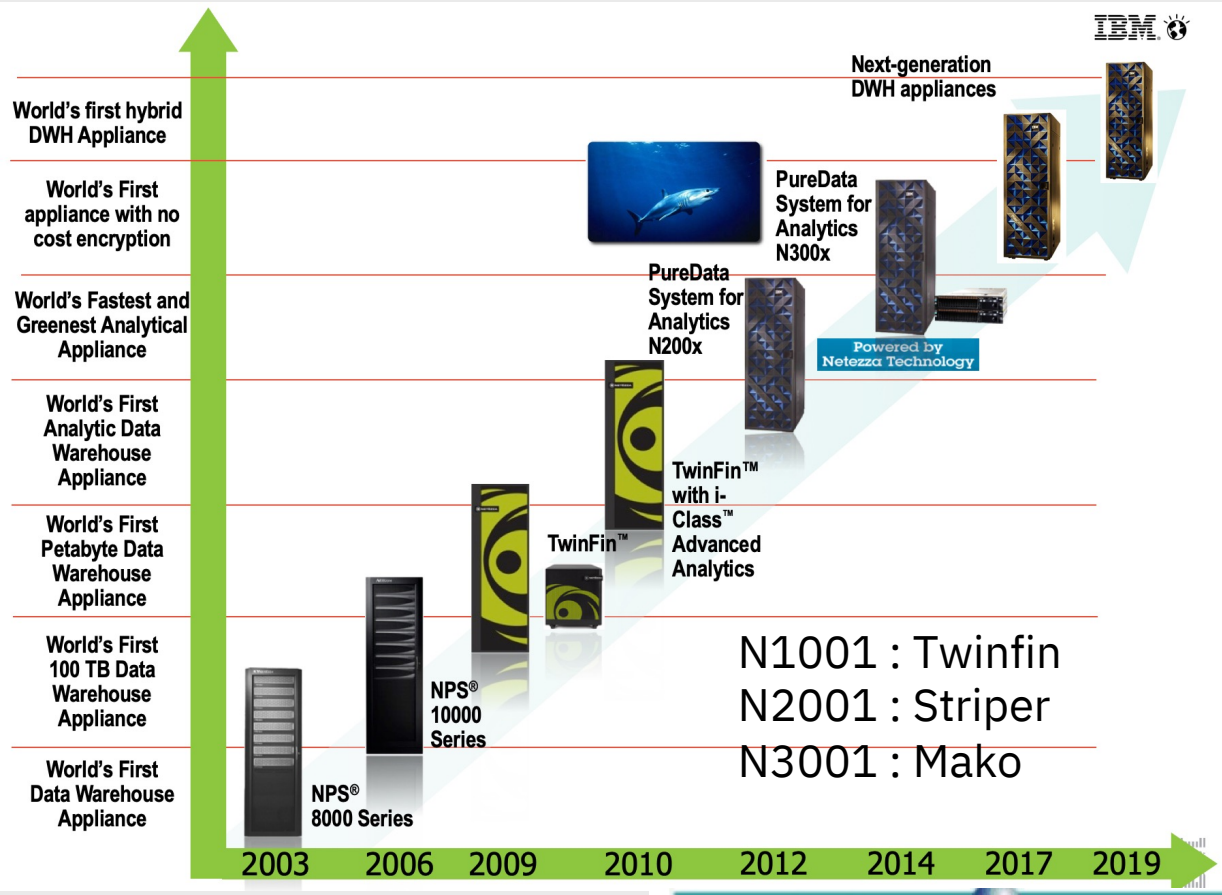

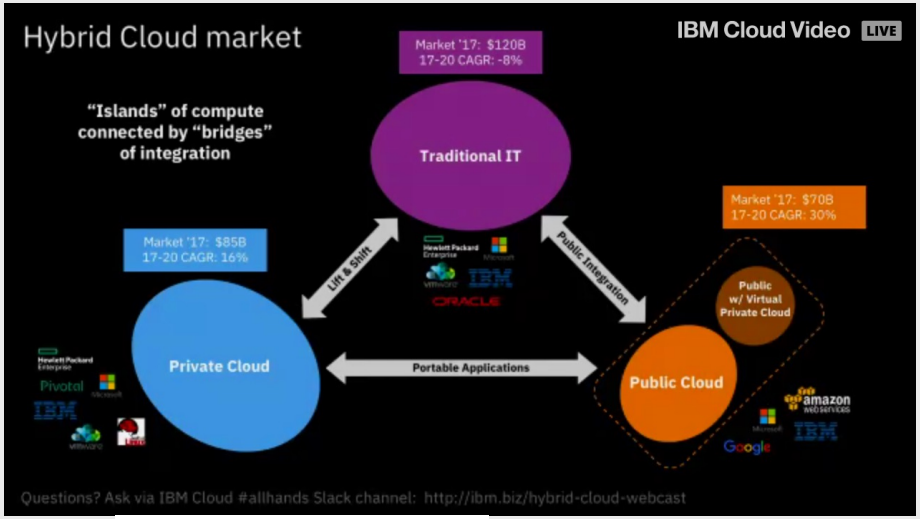

lenovo

Mako → Swordfish → BLU Marlin R1 → BLU Marlin R2

IBM Netezza Performance Server (NPS)

The foundational component of analytics, a data warehouse pulls together data from many different sources into a for-purpose data repository that enables sophisticated analytics, performant, and at scale.

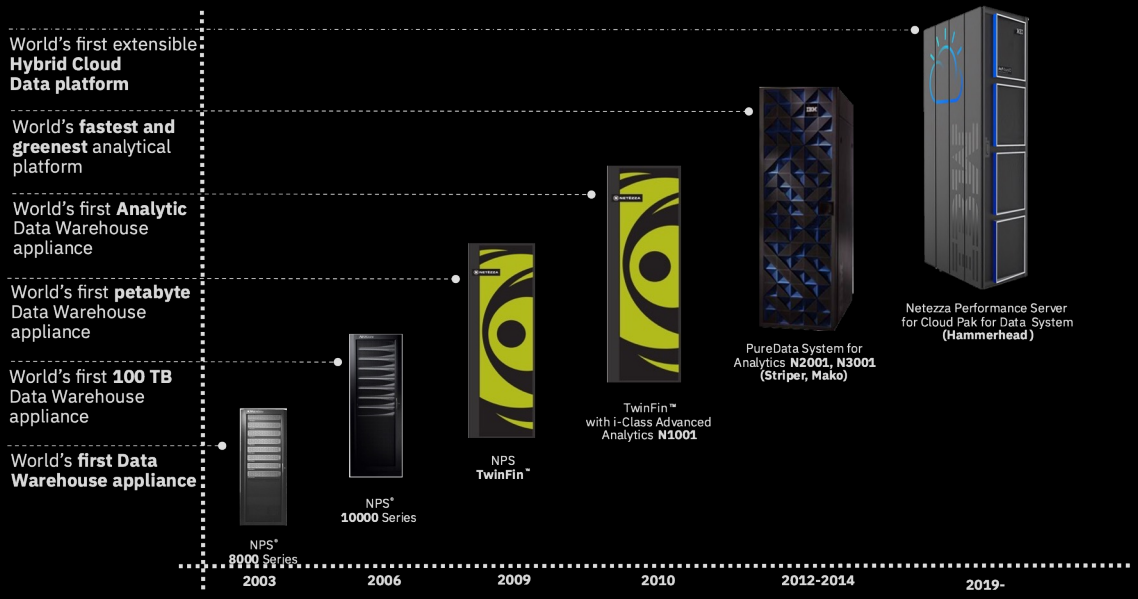

On Premises ⟵ Hybrid ⟶ Public Clouds

A **common database engine** deployed in any **public, multi-cloud**, **on-premise** (Netezza for Cloud Pak for Data System) or a **hybrid cloud** environment

Netezza Performance Server

A hybrid, multi cloud, data & analytics platform for actionable insights

Simplicity Drives **N**etezza

Accept no substitutes ! Same engine, same experience but modernized and cloud ready

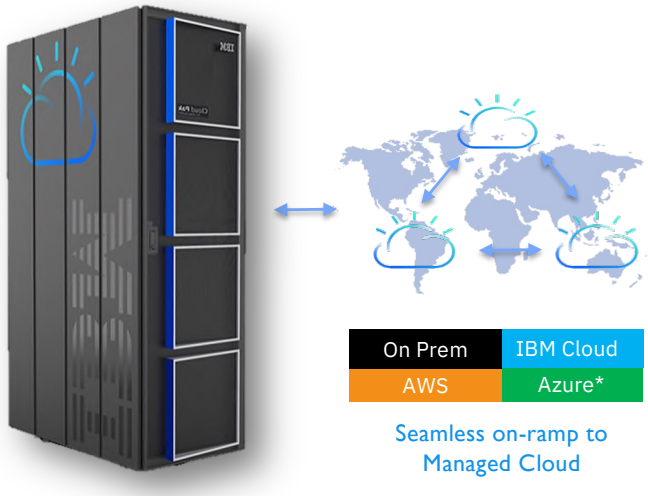

On Prem	IBM Cloud
AWS	Azure*

Seamless on-ramp to Managed Cloud

Simplicity
Minimal *administration & tuning*

Speed
Blazing *Performance*[1]

Organize Analyze Collect

Extensible
Hybrid Data Platform

Portable
frictionless ***lift-and-shift***
Risk free *upgrades*

Scalable Hybrid Analytics
Multi petabyte with ***elastic scaling***

Sophistication
In-database *advanced Analytics with built-in* ***DS/ML***

[1] 10x-400x faster than traditional custom systems based on IBM customers' reported results. "Traditional custom systems" refers to systems that are not professionally pre-built, pre-tested and optimized. Individual results may vary.

Please Allow Me To Introduce Myself : **Netezza V^{Next}**

Trade name: **Netezza Performance Server**

Outside View

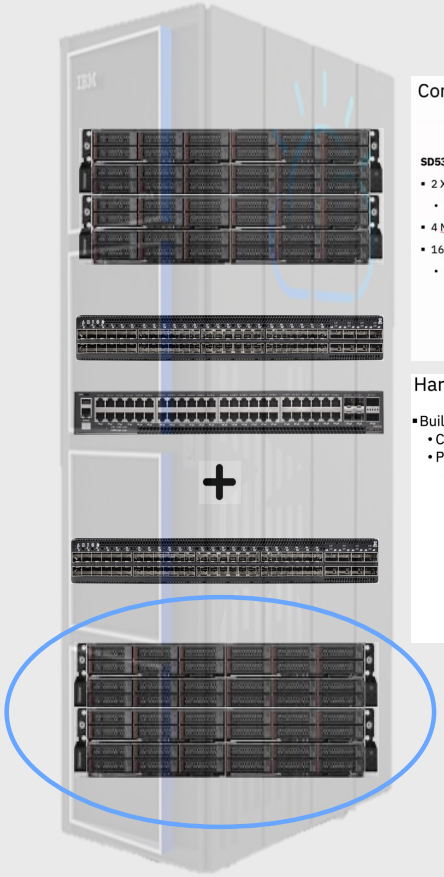

Inside View

Compute System Architecture

SD530 Server

- 2 Xeon Sockets
  - 16 Cores
- 4 NVMe Drives
- 16 DIMM slots
  - 12 memory channels

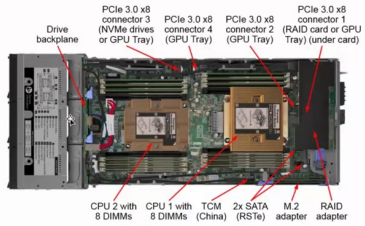

Hardware Acceleration

- Build on Mellanox Innova-2 NIC w/ Xilinx FPGA
  - ConnectX-5 ASIC for networking
  - Performance Server Query Acceleration
    - 8 PAX4 Engines
      - Decompression
      - Restriction
      - Row-Projection

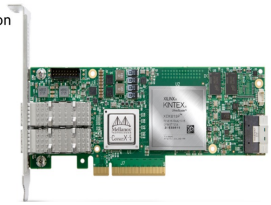

Introducing : Next Generation Netezza On Steroids

Trade name: **Netezza Performance Server**

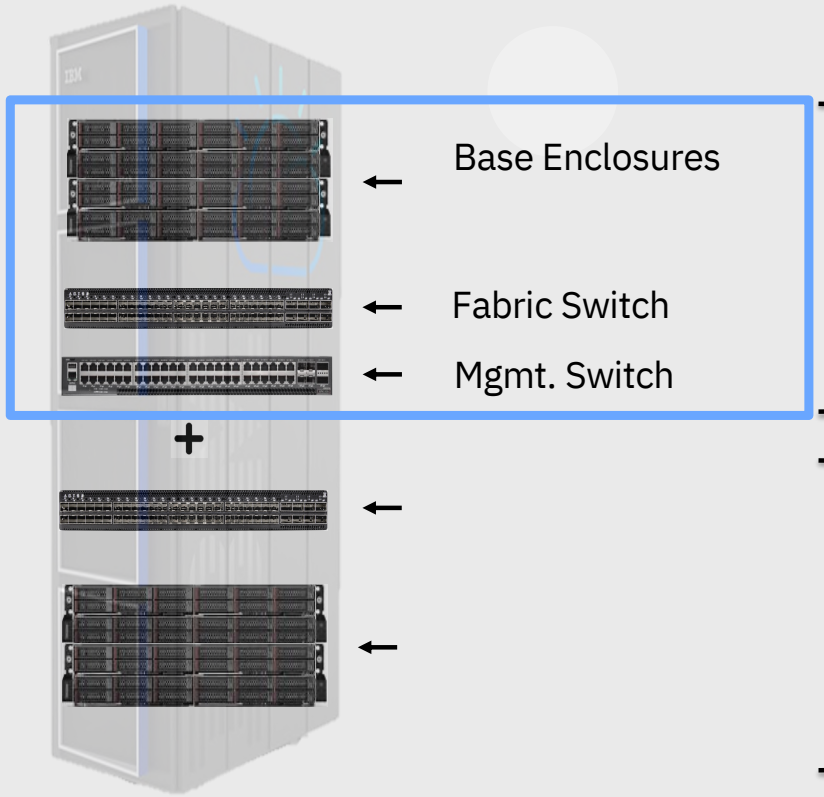

Cloud Pak for Data System
(CP4D) Starter Bundle
Base

IBM Performance Server
Addon Expansion Units
+2

Hyper Converged System
(In place, Scale out expansion architecture)

Netezza Performance Server - Configurations

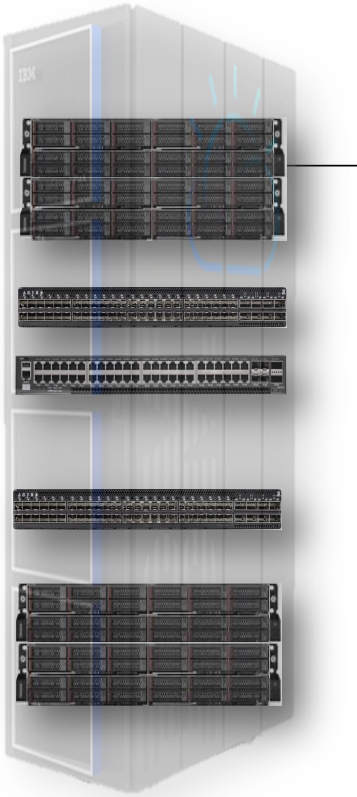

Cloud Pak for Data System Base (128 cores) Base 2x 2U enclosure + Performance Server Nodes

Hyper Converged System, rack mountable, 4TB NVMe self encrypting drives

25Gbps NIC Mellanox RDMA capable dual port card with FPGA Accelerator

	Base + 1	Base + 2	Base + 4	Base + 6	Base + 8
Performance Server Blades	4	8	16	24	32
Cores	64	128	256	384	512
Memory	0.76 TB	1.5 TB	3 TB	4.5 TB	6 TB
Available User Space[1,2,3]	68TB (17TB)	136TB (34TB)	272TB (68 TB)	408TB (102 TB)	544TB (136 TB)

[1]Assume up to 4x average compression to calculate user data (pre-load uncompressed user data).
Example a Base+2 (Mako full rack equivalent) user data capacity = 4 x 34TB = 136 TB

Netezza Performance Server – Value proposition

Deployment Options to allow customers to continue to run NPS workloads with the predictable performance they expect as they evaluate their journey to cloud

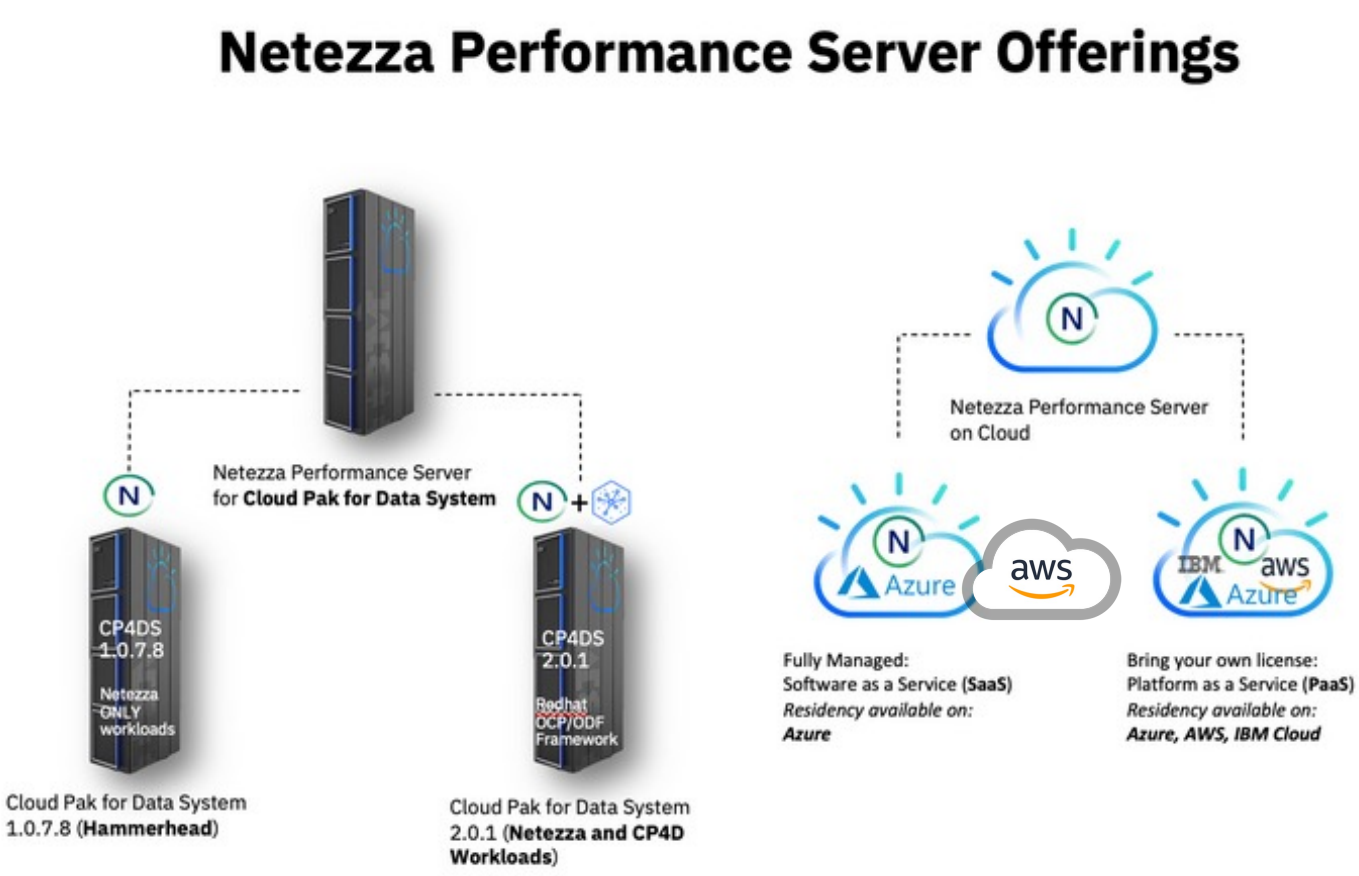

IBM Cloud Pak for Data

A truly hybrid-cloud world.

Building on the streamlined hybrid-cloud foundation of Red Hat® OpenShift®, IBM Cloud Pak for Data takes advantage of the underlying resource and infrastructure optimization and management. The solution fully supports multicloud environments such as Amazon Web Services (AWS), Azure, Google Cloud, IBM Cloud™ and private cloud deployments.

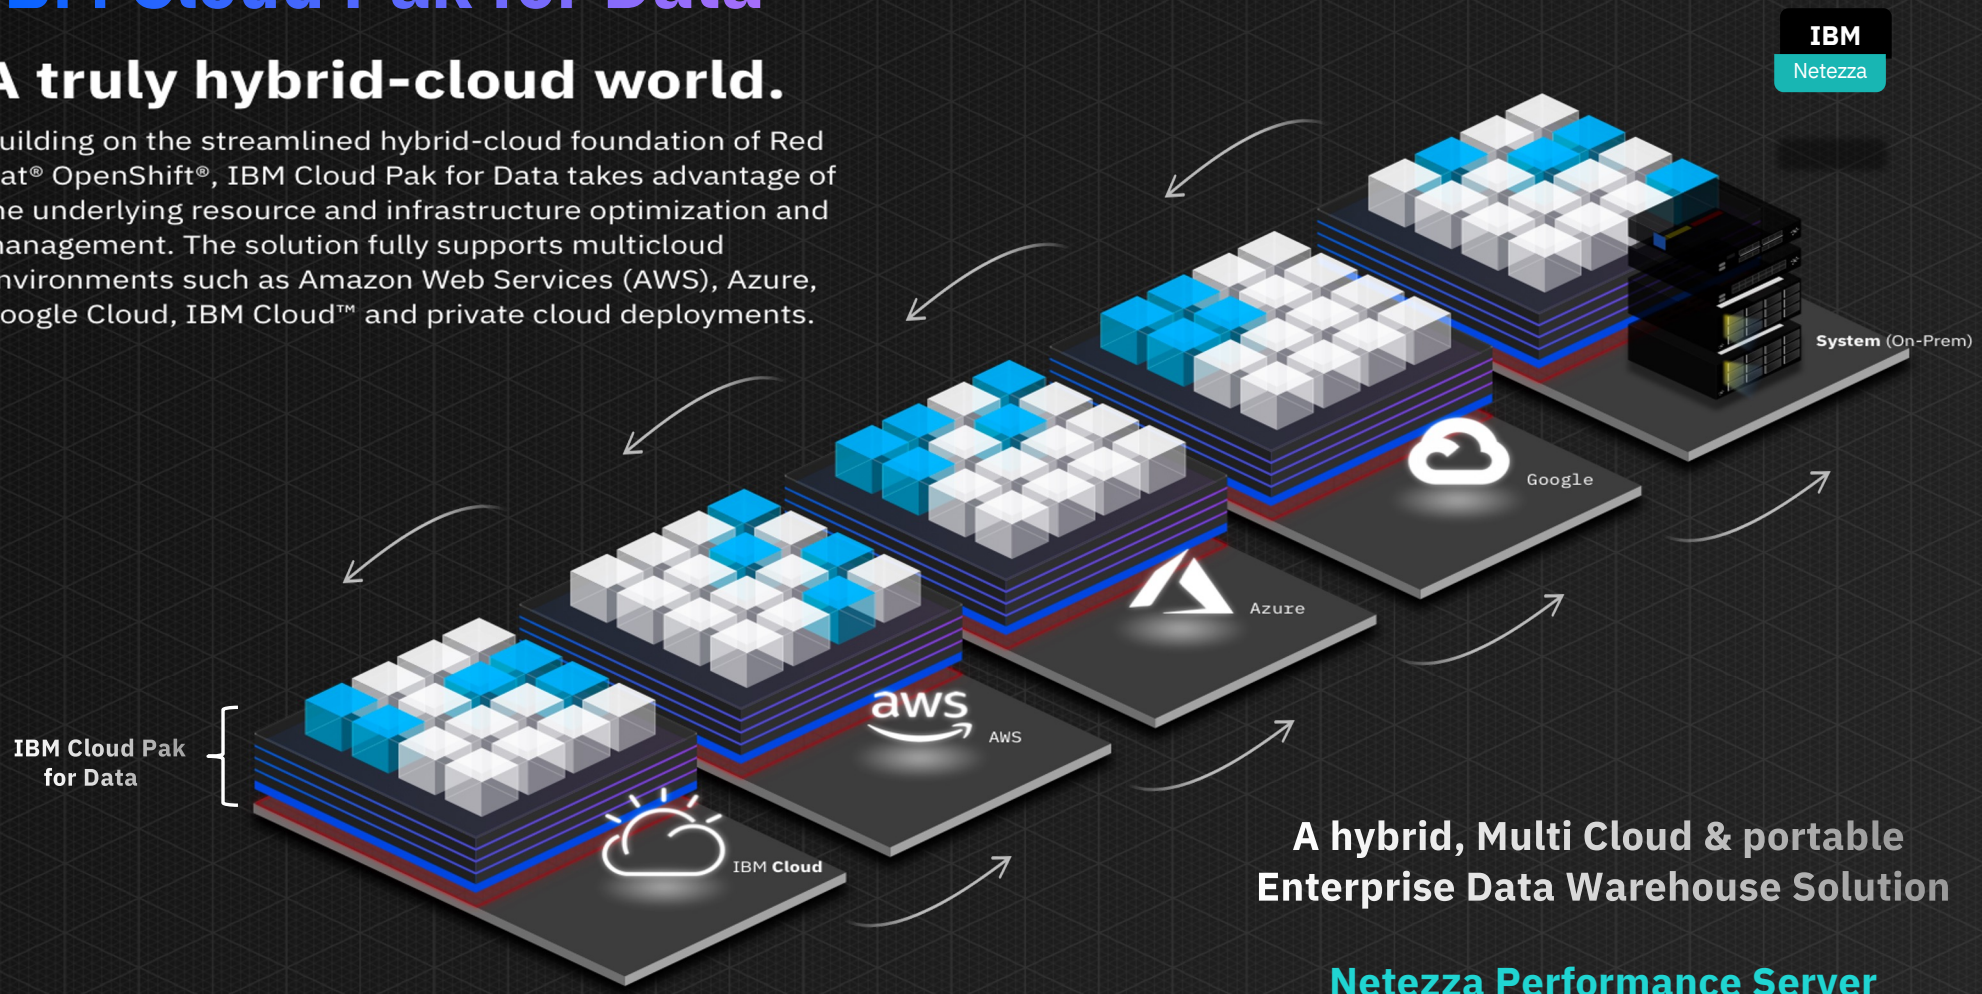

A hybrid, Multi Cloud & portable Enterprise Data Warehouse Solution

Netezza Performance Server

Netezza Invigorated for the Cloud

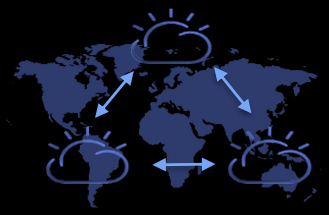

Fully leverage Netezza DNA infused with cloud native capabilities.
Performance, Simplicity, and Sophistication for actionable Analytics.

Flexible deployments
A **common database engine** deployed in any **public, multi-cloud**, **on-premise** (Netezza for Cloud Pak for Data System) or a **hybrid** environment

Same Netezza
Petabyte scale Netezza Technology for your deep analytics and AI workloads.

Risk free frictionless upgrades

Granular Elastic Scaling
Independently scale and manage compute and storage when you need & pay for only what you use, **pause/resume** as needed

Highly available and Reliable
Cloud-managed compute and **highly available** cloud storage, managed by a hybrid multi-cloud foundation of **Red Hat OpenShift**

Engine Sophistication
Mature MPP data warehouse with In-database analytical capabilities that work at scale along with seamless integration with IBM's analytical ecosystem

*A high-performance fully managed cloud data warehouse service running the next generation **Netezza database engine***

Fully Managed, Fully Netezza

Fully leverage Netezza DNA infused with cloud native capabilities.
Performance, Simplicity, Value, and Sophistication for actionable Analytics.

Scale

Granular elastic scaling with cost predictability. You pay for only what you need when you need it

Simplicity

Load & Go, minimum administration, no tuning, no indexing, risk free frictionless upgrades

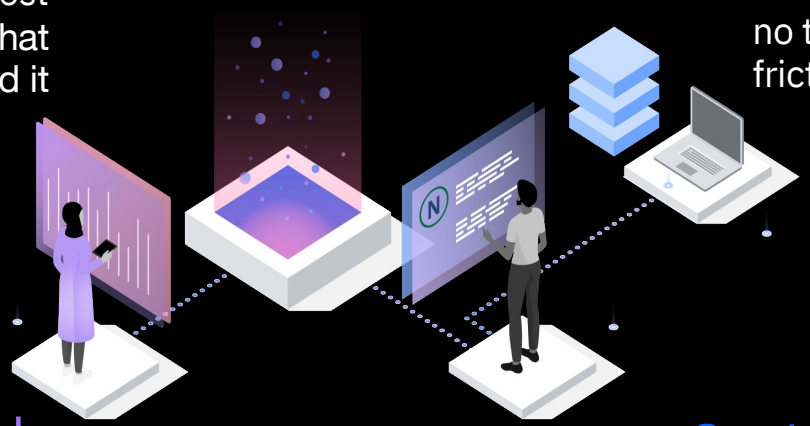

Speed

Faster Insight with vastly superior Price-Performance

Sophistication

Data Science with In-database Analytics at Scale

Scale

Granular elastic Scaling

independent scaling of performance and storage needs

Scale up during peak demand

Scale down when demand falls

Pause/Resume

Ramp up your data storage as your data needs grow

API and UI driven

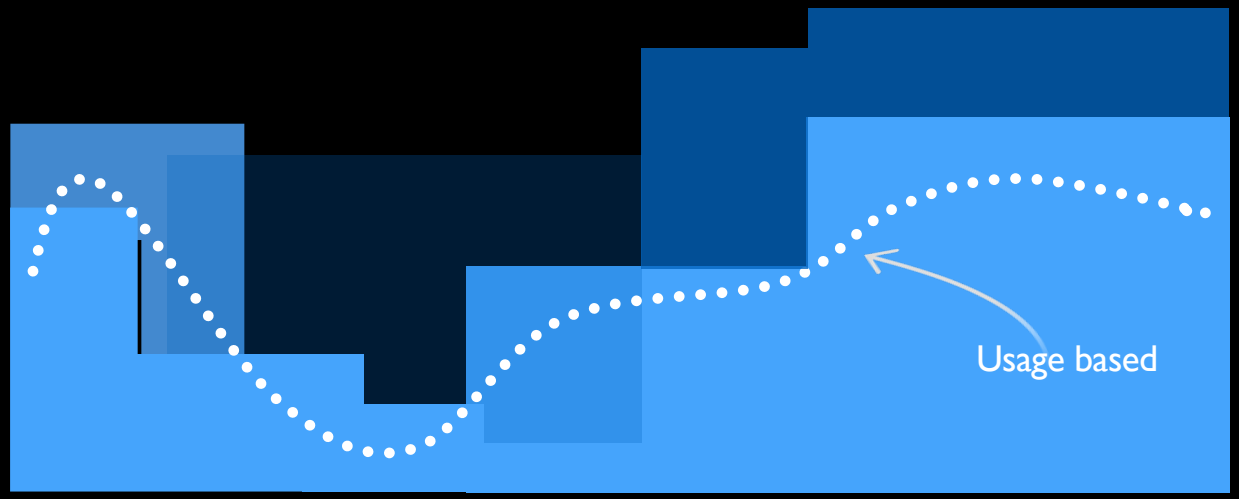

Pay for what you need when you need it!

Granular scaling : Adjust system power incrementally

Performance Profiles optimized for costs and performance

Granular Scaling

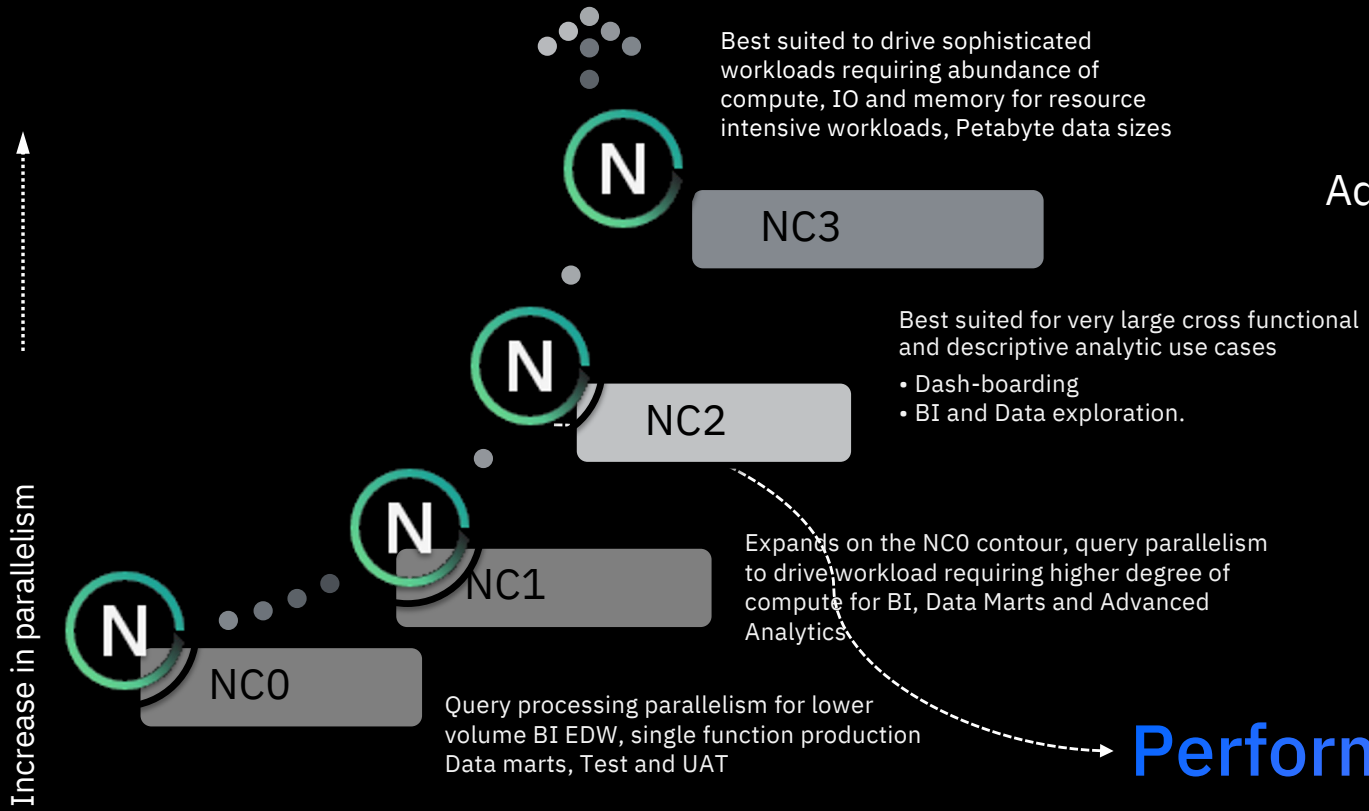

Adjust the power for within any contour

Tune power up and down

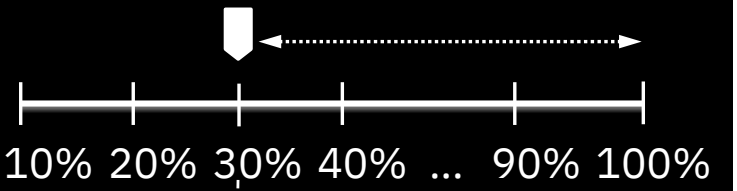

Performance profile : NC2-30


Whatever you need for Data & AI or you just want to have a chat
about innovating your business ☺?
+386 40 456 745 (call, text, whatsapp, viber, slack)
robert.bozic@si.ibm.com