



Introducing

Cloud Pak for Data System

IBM Performance Server^N

Hemant R. Suri

Program Director, Offering Management

Data & AI, IBM Analytics

USA



Welcome to Virtual Enzee

Today's Topic:

Introducing Cloud Pak for Data System with IBM Netezza Performance Server

Virtual Enzee Schedule and
Webcast Replay links are available on:



www.enzeecommunity.com



Call Logistics



Audio is broadcast via your computer speakers (no dial-in)

Ask questions via the Q&A button

Q&A

Experiencing technical difficulties? Let us know via the Q&A button

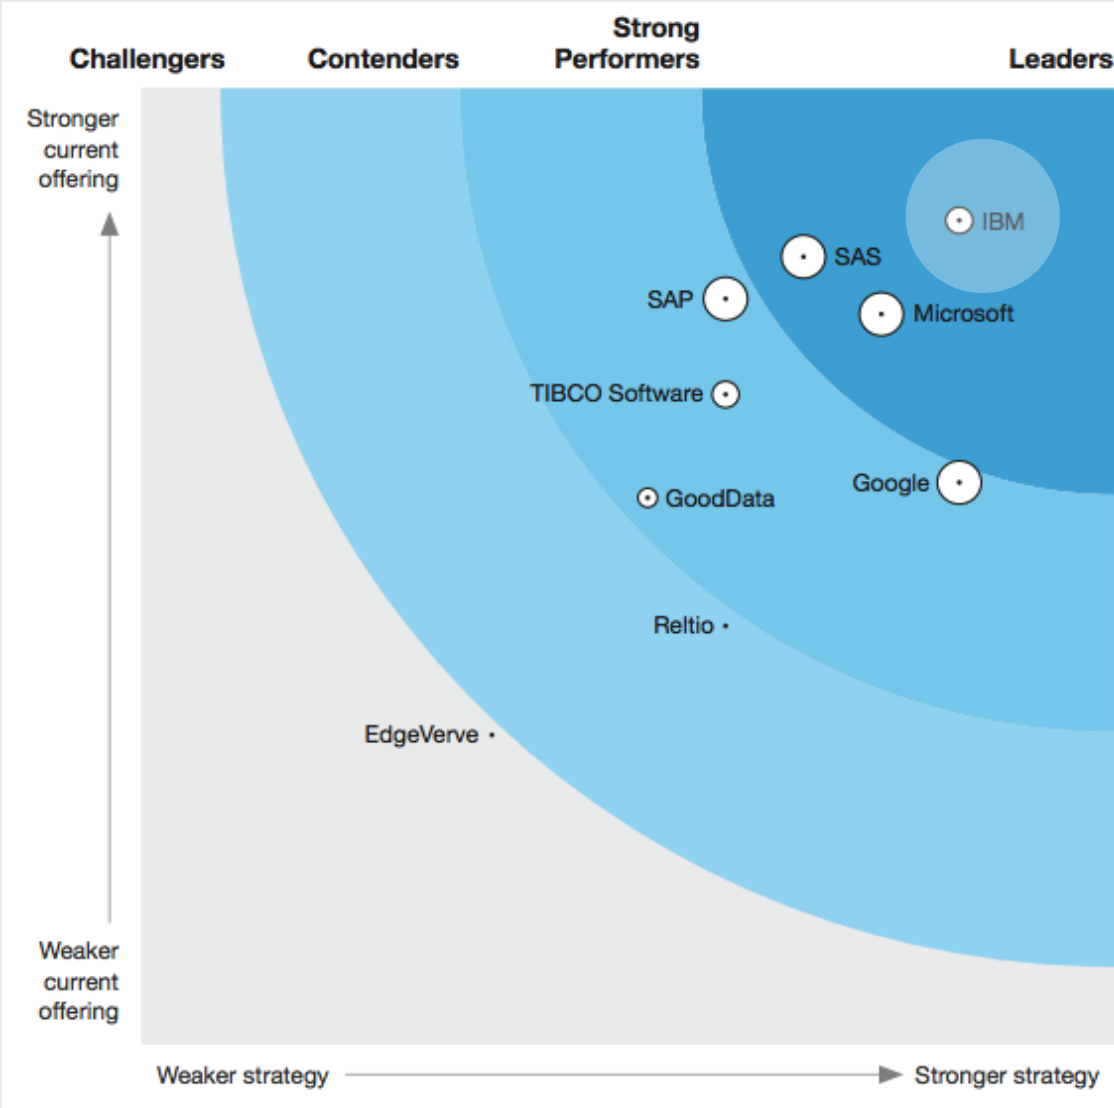
Virtual Enzee: webcast series sharing information to maximize your success

Introducing Cloud Pak for Data



Forrester Wave for “Enterprise Insight Platforms”

IBM Cloud Pak for Data is #1 (Top Ranked Leader)



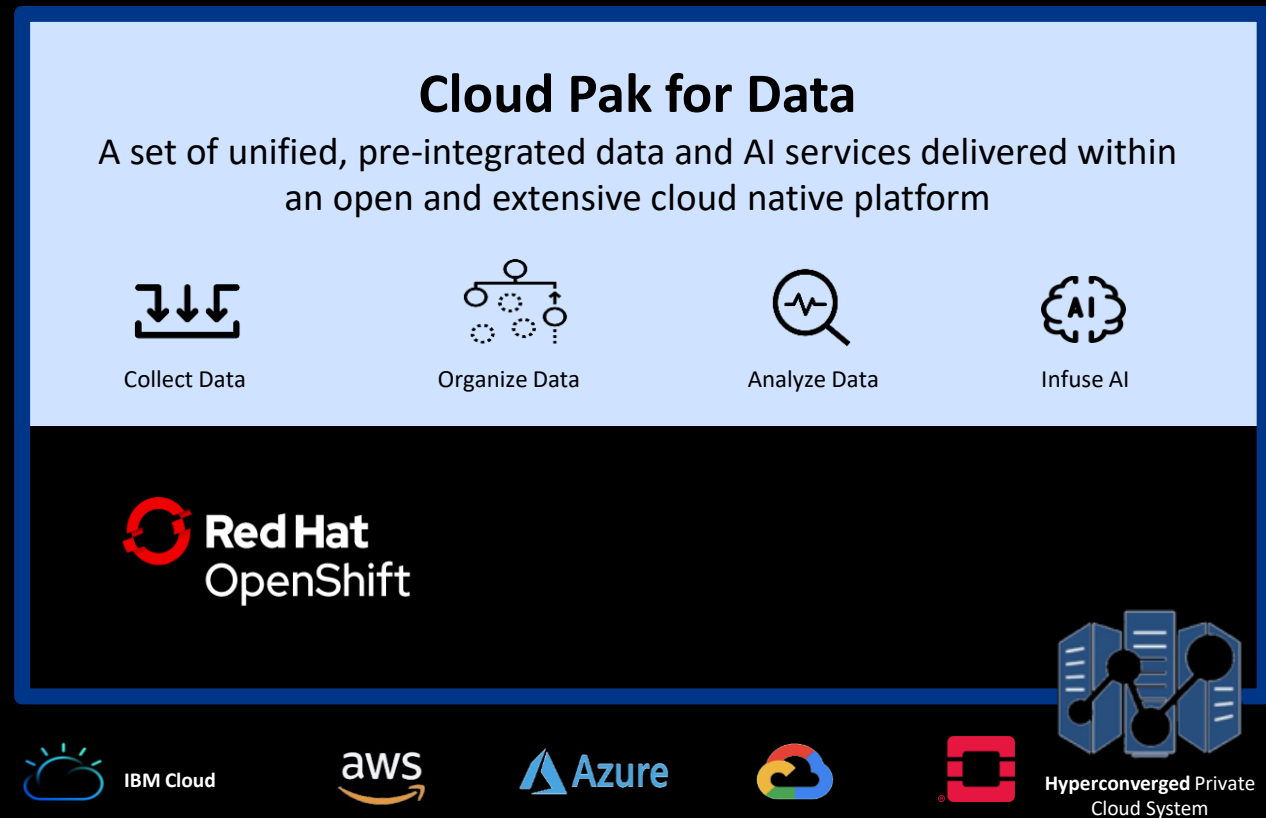
	Forrester's weighting	EdgeVerve	GoodData	Google	IBM	Microsoft	Reltio	SAP	SAS	TIBCO Software
Current offering	50%	1.62	2.86	2.94	4.31	3.82	2.19	3.90	4.12	3.40
Data management	25%	0.75	1.80	2.00	5.00	3.50	3.10	4.30	4.30	3.80
Analytics tools	35%	1.80	2.80	2.40	4.40	4.20	0.70	4.80	5.00	4.20
Insight app development tools	15%	2.60	3.40	3.40	4.00	3.70	1.70	2.60	3.60	2.60
Platform	25%	1.66	3.66	4.34	3.66	3.66	3.66	3.00	3.00	2.34
Strategy	50%	1.80	2.60	4.20	4.20	3.80	3.00	3.00	3.40	3.00
Product vision	20%	3.00	5.00	3.00	5.00	5.00	3.00	1.00	3.00	5.00
Past performance	20%	0.00	3.00	5.00	1.00	5.00	3.00	1.00	5.00	3.00
Innovation road map	20%	3.00	1.00	5.00	5.00	1.00	3.00	3.00	3.00	3.00
Delivery model	20%	3.00	3.00	3.00	5.00	3.00	5.00	5.00	3.00	1.00
Partner ecosystem	20%	0.00	1.00	5.00	5.00	5.00	1.00	5.00	3.00	3.00
Market presence	0%	1.00	2.00	5.00	3.00	5.00	1.00	5.00	5.00	3.00
Revenue	50%	1.00	1.00	5.00	5.00	5.00	1.00	5.00	5.00	3.00
Number of customers	50%	1.00	3.00	5.00	1.00	5.00	1.00	5.00	5.00	3.00

All scores are based on a scale of 0 (weak) to 5 (strong).

Cloud Pak for Data

Delivers the foundational platform for deploying an information architecture for AI, on any cloud

- Eliminate data silos, connect all data
- Automate and govern the data & AI lifecycle
- Operationalize AI with trust & transparency
- Avoid lock-in, run anywhere with agility




Cloud Pak for Data

Cloud-Native by Design

Agility ● Efficiency ● Cost Savings


Cloud Pak for Data

Microservices




An architecture of loosely coupled data services, easily refactored to create containerized workloads

Containerized Workloads



Stand-alone workloads composed of microservices & data that are flexibly deployed, orchestrated and managed

Multicloud Provisioning



Agile provisioning of containerized workloads in multicloud environments and consumption of cloud services



Built for Multi-cloud

Avoid vendor lock-in and get started on your cloud journey today

Pick Your add-on

Containerized Services

Customize & Extend with add-on microservices

Db2 DW

Netezza

Mongo

Cognos

IBM Watson

ISV

OS

Custom

Data Platform

#1 Ranked by Forrester

IBM Cloud Pak for Data

Unified platform of foundational Data & AI
cloud services

Collect

Organize

Analyze

Infuse

Kubernetes Based

Containerized, easy to manage



Pick Your Deployment

Private or Public



On-Premise



Cloud Pak
For Data
System

Expanding 3rd Party Eco-system

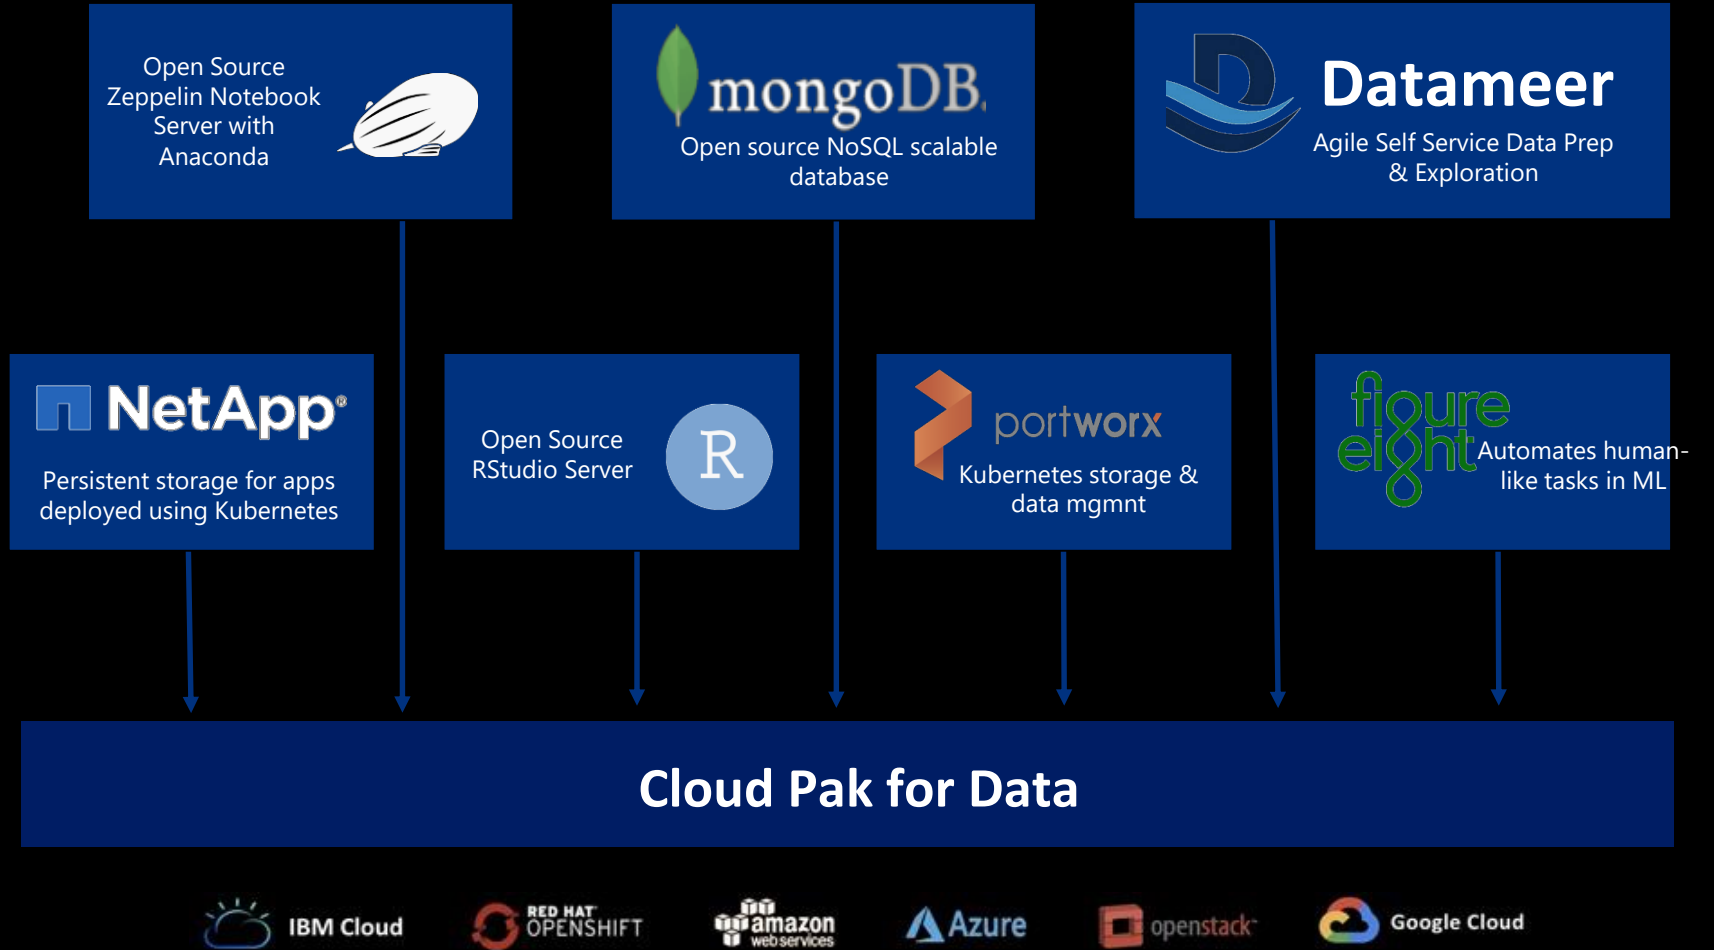


An expanding ecosystem of
Partner add-ons

Analogy



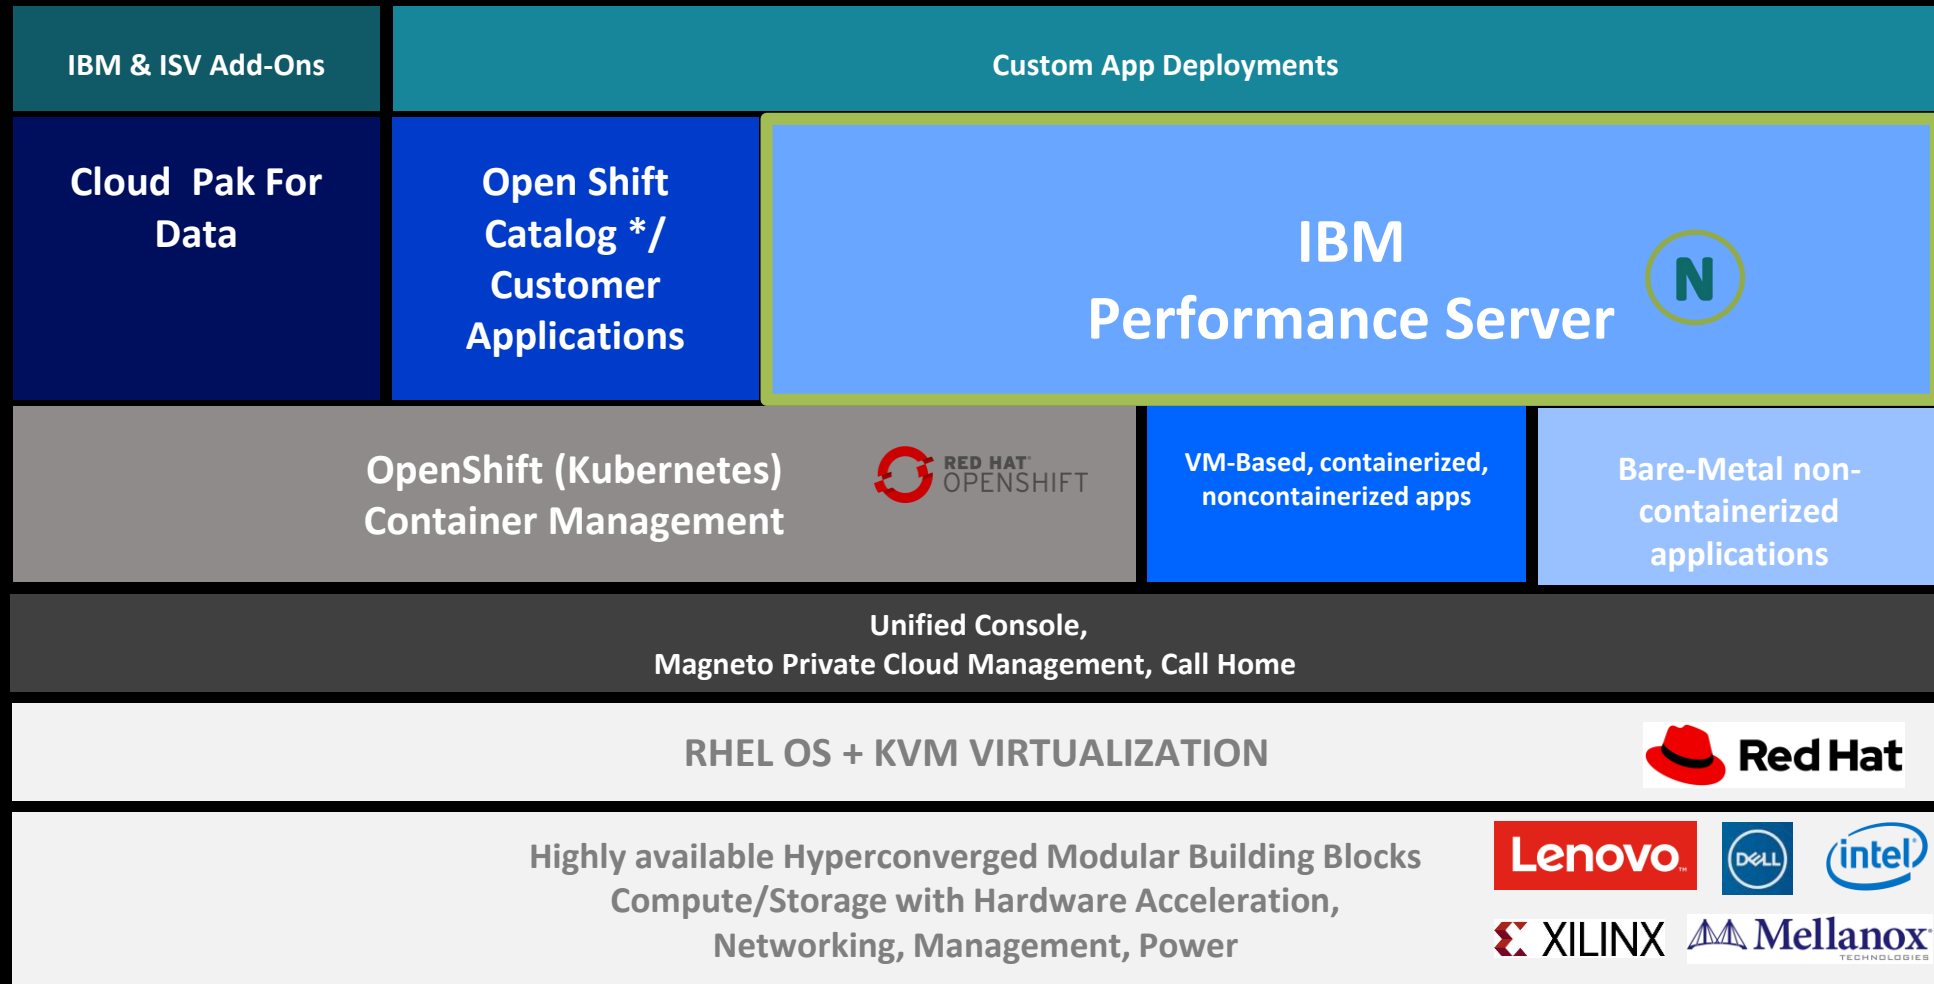
IBM Cloud / © 2018 IBM Corporation



Cloud Pak for Data System IBM Performance Server^N



Cloud Pak for Data System: A Hyper Converged System Architecture





IBM

Cloud Pak for Data System

IBM Performance Server

*fully (100%) compatible with Netezza/Pure Data System for Analytics line of appliance offerings,
all nz* utilities compatible, Synergy with data integration & BI tools, “fork-lift” existing Netezza systems, ‘nz_migrate’ and go!*

Cloud in-a-box

available in Cloud Pak for Data System and 100% compatible with Netezza with deep 3rd party tool integration

Blazing-fast

built over open source and optimized for High Performance Analytics with built-in hardware acceleration and NVMe drives

Simple

load and go, minimal administration and tuning

Scalable & elastic

multi-tenant, start small, grow as you need on storage and compute

Availability

Red Hat Open Shift (Kubernetes)-container managed, highly available

Support & usage

production ready and fully supported by IBM, unlimited use as with Db2 Warehouse

Native Netezza Compatibility

❑ Simple

- Load and Go, Minimal administration and Tuning


❑ Compatible

- 100% Netezza (Mustang/Skimmer/TwinFin/Striper/Mako) compatibility
- "Fork-lift" existing Netezza/PDA systems
- Simply use familiar 'nz_migrate' and go!
- Existing UDFs / UDXs work (*requires re-compile*)
- Bi-directional D/R Replication, BnR with older Netezza/PDA

❑ Smart with built-in FPGA Performance Accelerator

- Data Stream processing
- Machine Learning
- In-place Advanced Analytics & Geo Spatial

❑ Familiar interface and tools



Netezza Database Engine Enhancements

IBM Performance Server

☐ Modernization

- Containerized for microservice architecture and deployment in kubernetes
- UI upgraded to REACT based UI with same look and feel as Cloud Pak for Data
- Native Go (Lang) Driver implemented for modern day app connectivity
- PostgreSQL compatible

☐ Serviceability

- Maintenance prediction engine

☐ Hardware

- Upgraded to 64-bit engine to take full advantage of memory
- Upgraded to full NVMe support with NVMeoF implementation for fast performance
- FPGA ready

☐ Differentiation

- In-database Analytics
- Multiple databases with Cross-db writes
- Out-of-box Security
- Near linear performance in scalability
- Ability to take advantage of NVMe Drives
- Hardware acceleration through FPGA

Built-in Data Science and Machine Learning

Powered by state of the art Data Science Hardware Performance Accelerator

In place Machine Learning
+
High Speed Data Accelerator (FPGA)



Extensible
Add-on
AI Services

Foundational
Data & AI
Services

*Included with
Performance Server for
PostgreSQL

Watson Studio & Machine Learning

Data Science Premium
Tools to design, build
and deploy AI models



Watson OpenScale

▪ Operational optimization services

Watson Assistant

▪ Conversational AI services

Watson APIs

▪ Interactive AI services (e.g., speech)

*Cloud Pak for Data

Open Source frameworks to build, train and
deploy AI models powered by **Watson
Studio and Machine Learning**

Jupyter
Apache Spark
ML
Caffee
Apache
Zeppelin

TensorFlow
Rstudio
Python
Keras
Torch
XGBoost
Scala

- Data visualization
- Machine learning learning
- Model build & deploy
- Model management
- Dashboards & reporting


Data Collection Services

Data Organization Services

Data
Preparation




Predictive
And
Prescriptive
Analytics



Geospatial
Analytics



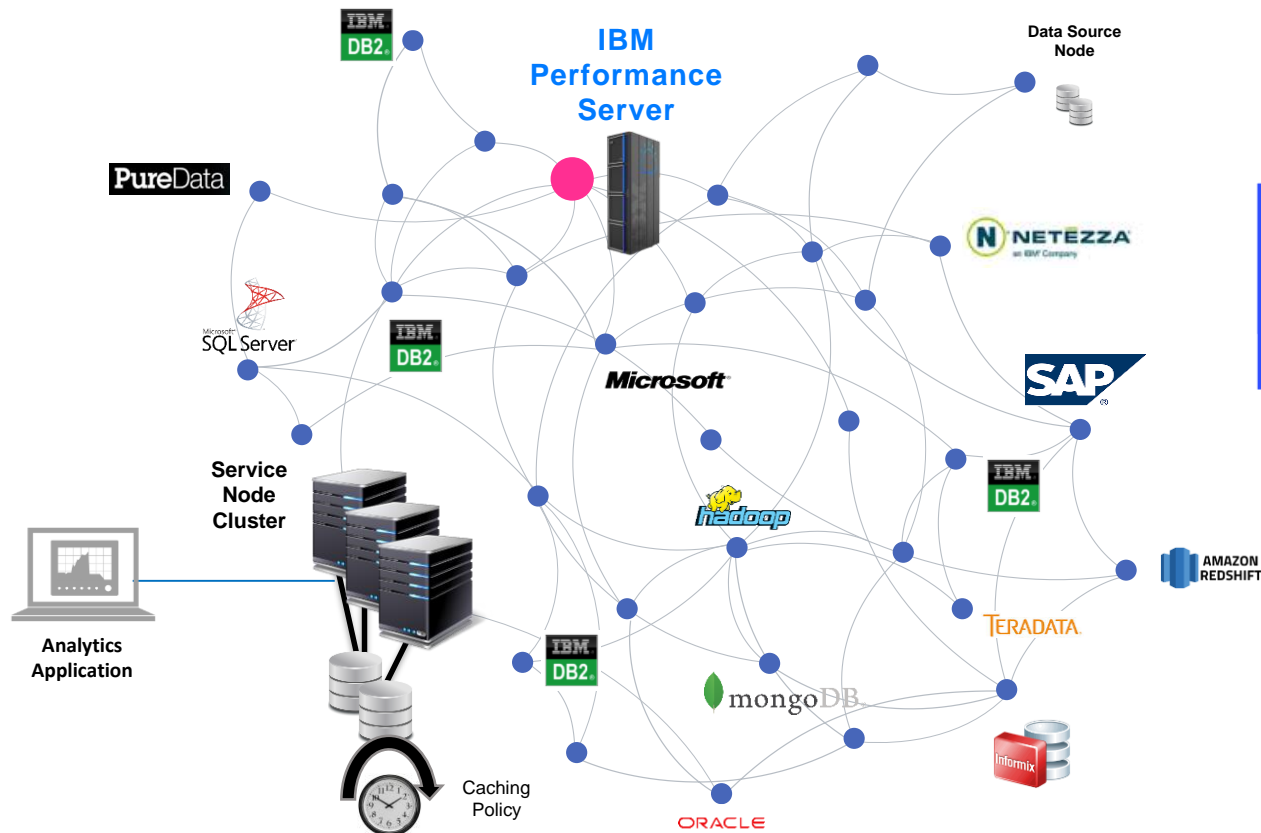
Advanced
Statistics



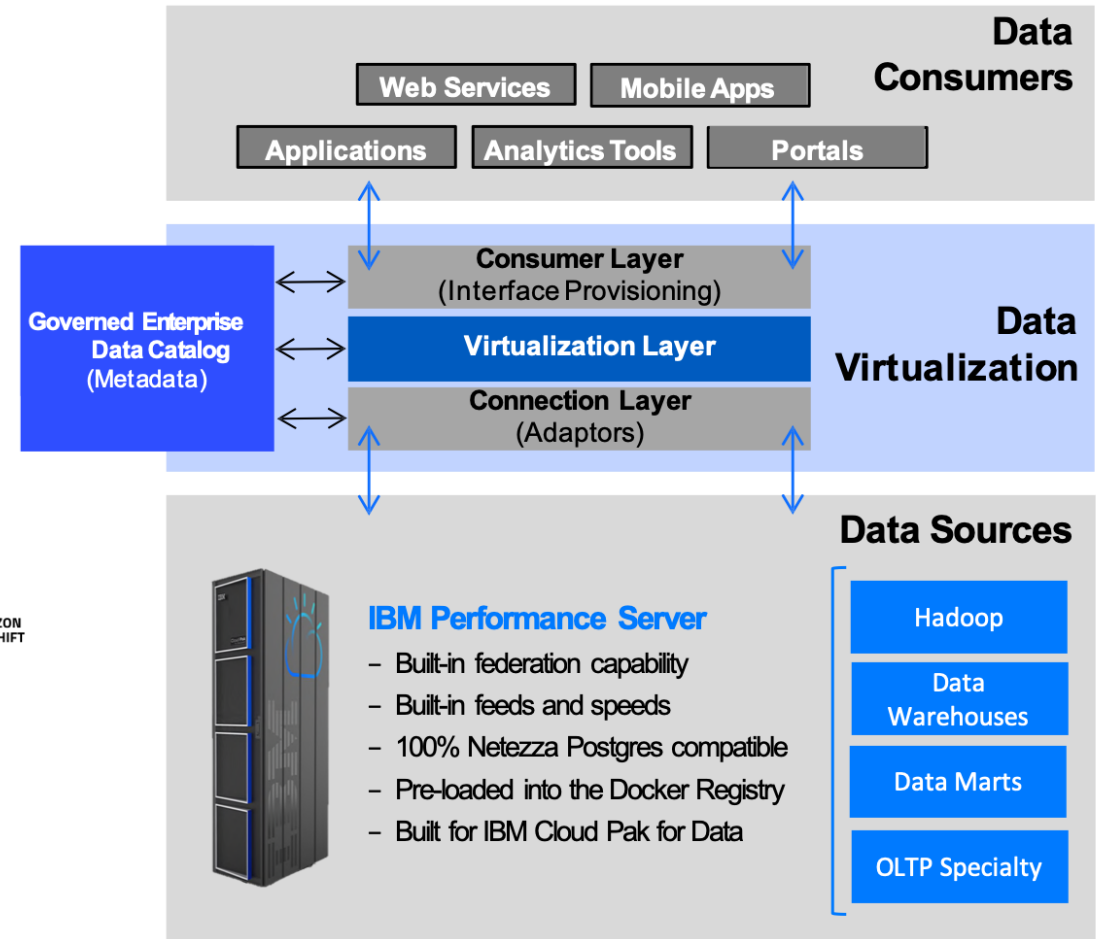
Data Virtualization

Key enabler for Hybrid Data Management

- High performance / Enterprise Ready
- Hybrid data store with extreme functionality




- Single data fabric for self-service
- DV gateway through Fluid Query



Synergy with Data Integration and Reporting & Analysis Tools

Data Integration¹

- IBM
 - Information Server
 - InfoSphere Streams
- Ab Initio
- Hadoop
- Informatica
- Microsoft
- Oracle
- SAP
- SAS
- Others using standard ODBC/JDBC/OLE-DB/SQL



Reporting & Analysis¹

- IBM
 - Cognos
 - SPSS
 - Campaign
- Hadoop
- Information Builders
- Microsoft
- MicroStrategy
- Oracle
- SAP
- SAS
- Tableau
- Others using standard ODBC/JDBC/OLE-DB/SQL

¹100% native compatibility with Netezza/PDA

Back up & Restore

100% Compatible

*With existing Netezza/Pure Data
System for Analytics*

***Lightning fast restore from
backup if you should need it***

- Build on External Table framework
- Backup and restore user data using [nzbackup](#) and [nzrestore](#)
- Supported backup types
 - Full
 - Differential
 - Cumulative
- Connectors Supported
 - File System Connector
 - Veritas NetBackup Connector
 - IBM Tivoli Storage Manager Connector
 - EMC NetWorker

Back up & Restore


Backup Method contd.,

User data backup method

Option	Description	Command
Full	Full database backup	default
Schema Only	Only saves database objects (tables, views, stored procedures, privileges ...) but no user data.	-schema-only
Differential	Saves only the data, that has changes since the last backup (full, differential or cumulative)	-differential
Cumulative	Back-ups everything that has changed since the last full backup	-cumulative

database backup with table-level restore possibility

Less Admin, More Analytics



NO requirements for armies of administrators

Load data, **start querying within hours**, not days

NO “trapping” your data on the platform

Built-in data migration tools

NO proprietary-only integration

Single management console

NO storage administration

System is **performance optimized** for analytics

NO physical data modeling needed

Hyper converged scale out system

NO indexes and tuning necessary


Single support number

Cloud Pak for Data System - Performance Server *vs everyone else*

Benefits of upgrading

100%

Netezza
cross-generational
compatibility



IA for AI
(Information Architecture for AI)




Migration effort
(frictionless upgrade)



Blazing fast
(up to 3X faster)



In-place expansion and
hardware based multi-
tenancy



Seamless on-ramp to
Managed Cloud
(1H 2020)

0%

0% risk



Built-in Security

Questions?

Type your question in the
Q&A panel on your screen.

