

Meet the better, faster, stronger Netezza from IBM

Upgrade with a single
command and deploy on
the cloud or on premises
with 100% compatibility



Table of contents

3
Executive summary

4
Upgrade with a single command

5
Innovative and fast deployment with cloud-native containers on Red Hat OpenShift

6
Netezza with 5 key characteristics

7
Hyper-converged, cloud or both?

8
FAQs

9
Get started



Executive summary

Netezza® systems have always combined simplicity and performance. The next generation of Netezza warehouse—[Netezza Performance Server for IBM Cloud Pak® for Data](#)—continues to improve the performance of the Netezza engine while retaining 100% compatibility with IBM® PureData® System for Analytics and its models TwinFin, Striper and Mako.

The advanced technology in Netezza Performance Server fuses data warehousing and in-database analytics into a scalable, high-performance, massively parallel advanced analytic system that is designed to crunch through petabyte-scale data volumes with linear performance scalability. On the IBM Cloud Pak for Data platform, Netezza Performance Server benefits from a containerized “any cloud” architecture and full integration with other IBM solutions for business intelligence, DataOps, machine learning and AI.

As a modern, performant data warehouse, Netezza Performance Server is the best version of Netezza available today, as well as the best option to avoid any lapse in support for your existing Netezza workloads.



Upgrade with a single command

Upgrading to the newest version of Netezza is simple. Netezza Performance Server is cloud-native and 100% compatible with IBM® PureData® System for Analytics and its earlier models TwinFin, Striper and Mako. That means you can upgrade via a single command. Whether you choose to run on a [public cloud](#) or a [hyper-converged system](#) behind your own firewall, Netezza Performance Server offers the performance and simplicity you have always come to expect.

Netezza Performance Server offers:

- The ability to seamlessly upgrade with a single command: *nz_migrate*. Point your existing applications to the new Netezza server and they will work exactly as they did before.
- An optional two-step “lift and shift” option for moving from on-premises to the cloud: *nzbackup* your appliance to IBM Cloud® Object Storage, AWS S3 or Microsoft Azure, then *nzrestore* into your cloud data warehouse.
- Business continuity as older models go out of support. Striper ended support in June 2020 and Mako will follow in 2023.

./nz_migrate

Innovative and fast deployment with cloud-native containers on Red Hat OpenShift

With Netezza integrated into [IBM Cloud Pak for Data platform](#), running on Red Hat® OpenShift® Container Platform (see Figure 1), you gain better access to—and value from—all your data. For example, built-in data virtualization means you can treat many databases as one, querying any data source, anywhere. And powerful integrated capabilities such as IBM® DataStage® (for DataOps), IBM Cognos® Analytics (for business analytics) and IBM Watson® Knowledge Studio (for machine learning) help you infuse AI through your entire data analytics pipeline. The result is an intelligent, integrated view of all your data, no matter where it resides. You can more easily operationalize real-time analytic insight and make it available throughout your enterprise.

Advantages:

- **Microservices that run anywhere**
Microservices can run on the cloud or hardware of your choice, including clouds from other providers. You can also easily run a hybrid cloud environment if you need different infrastructures for different data.
- **A robust and flexible information architecture**
Data within IBM Cloud Pak for Data is pre-governed to help with data accuracy, organization, high availability and resiliency. Failure of a processing node or any element in the system causes no significant performance degradation. The integrated platform delivers not only a complete AI-ready architecture, but also flexibility to grow and scale as needed.
- **Unified data**
A unified platform makes all of your data available for analysis and AI, solving more business problems while saving costs that would have gone towards data migration. Furthermore, a unified interface makes it easier for employees to build skills and then retain the value of them for future AI projects even if they move to a different role.

1. **Services ecosystem:**
Access and deploy an ecosystem of over 45 analytics services and templates from IBM and third parties.
2. **Data virtualization:**
Quickly query across multiple data sources without moving your data.
3. **Platform interface:**
Integrate data management, data governance and analysis for greater speed, efficiency and improved use of resources.
4. **Red Hat® OpenShift®:**
Build on the leading hybrid cloud, enterprise container platform.
5. **Any cloud:** Leverage a multicloud infrastructure.



High-performance Netezza with 5 key characteristics

Smarts

- In-database analytic tools and functions enable data mining, statistics, predictions, transformations, geospatial analytics and data preparation.
- Analytical infrastructure and extensive libraries of statistical and mathematical functions support a breadth of analytic tools and programming languages, including open-source R.
- Data science and AI tools enable data scientists to build, iterate and execute models from relevant enterprise data.

Scalability

- Hyper-converged option combines storage, compute, networking and software to reduce complexity and increase scalability in a [single system](#).
- Netezza Performance Server requires minimal human intervention to provision, scale and maintain.
- Independent scaling of storage and compute¹ through the REST API helps meet peak demand or save money when demand falls.

Simplicity

- Upgrade with one command: `nz_migrate`.
- Fully integrated system does not require indexes to be defined and optimized, or storage to be manually administered. Install it and be ready for data load in just hours with minimal ongoing maintenance.²
- One GUI access point helps you monitor system resources, administer database objects, configure workload management and view active sessions.

Speed

- Up to 3x SQL performance over previous generation systems.³
- The unique asymmetric massively parallel processing architecture (AMPP) and patented hybrid columnar acceleration assist crunches through massive amounts of data quickly.
- Leverage state-of-the-art hardware technologies such as faster cores and highly performant NVMe drives⁴ while minimizing data movement, enabling analytics workloads that support thousands of users.

Savings

- The total cost of ownership (TCO) for IBM Cloud Pak for Data System storage could be ¼ the TCO of public cloud storage.⁵
- Netezza Performance Server requires minimal ongoing maintenance, which in turn cuts down internal resources and implementation for a low TCO.



Hyper-converged, cloud or both?

You can deploy Netezza Performance Server on IBM Cloud Pak for Data (cloud-based), [IBM Cloud Pak for Data System](#) (hyper-converged), or both. The options are fully compatible with one another, and either can offer an economical and practical solution for your data warehousing and analytics requirements. There are many benefits that are common to both deployment models.

Benefits

- Because the containers are based on Kubernetes and Red Hat OpenShift, you can easily move them between private and public clouds as needed—including clouds from other providers such as AWS.
- Built-in containerized microservices that are independent of each other allow you to deploy multiple services—whether related or not—depending on your needs and resources.
- Compatibility and data virtualization allow you to place your data and run workloads where it is most advantageous. Furthermore, there are no ETL or migration costs when the data is needed elsewhere.
- Netezza simplicity and user-friendliness mean that a single user can undertake self-service data access, dynamic workload management and queries of virtually any archived data.

Benefits of using Netezza on a hyper-converged system

- Pre-configured appliances require less time and complexity with integrated hardware and software.
- Appliances are tuned and tested in advance. Instead of focusing on hardware, your team can focus on the system capabilities.
- Combine the flexibility of the cloud with the security protocols of your own data center.

Benefits of using Netezza on the cloud

- Scale compute and storage independently—and infinitely.⁶
- Pay only for what you use.
- Simple interface; scale by dragging a slider.⁷



Frequently asked questions

Is Netezza going away?

No. The Twinfin, Striper and Mako models will no longer be supported, but Netezza Performance Server carries forward the Netezza legacy with 100% compatibility and added benefits.

Does IBM help with migration and modernization of my data analytics platform?

Yes, IBM Expert Labs can help, but most clients find it easy to migrate themselves due to the 100% compatibility—all it takes is the `nz_migrate` command.

Will the data integration and business intelligence (BI) tools that I use today still work with Netezza Performance Server?

Yes. The drivers used in the previous generation of Netezza products will be compatible with Netezza Performance Server. Any best-of-breed BI or analytic tool with standard interfaces can connect to and run against Netezza data.

How does Netezza Performance Server compare to competitors?

A recent study from Cabot Partners analysts showed that Netezza Performance Server outpaced competitors in key areas such as modernization, infrastructure, enterprise readiness, data management and analytics. This led Cabot to conclude that “NPS which includes IBM Cloud Pak for Data System is outperforming competitors and fast-tracking the delivery of client value in their advanced analytics and AI journey.” The full study is available [here](#).

How can Netezza benefit data scientists?

When you buy Netezza you get Cloud Pak for Data VPCs that you can use towards Data Science and Watson Studio, enabling your data scientists to quickly create, train and deploy in-database machine learning models. Using cognitive machine learning, Netezza provides your data scientists with in-database analytics, enabling them to collaborate inside one unified platform. Embedded Spark and support for Python and R enable rapid and sophisticated adoption of data science and machine learning at scale.

Is Netezza secure and resilient?

Yes. Built-in automation and auto-recovery (available on both IBM Cloud and Amazon Web Services) provides a highly available and fault-tolerant Netezza deployment with minimal human intervention, ensuring your data warehouse can run 24x7x365. Backups can also be scheduled at your convenience and kept in object stores as well as replicated to multiple availability zones. In addition, the NVMe drives in IBM Cloud Pak for Data System are self-encrypting drives (SEDs) that can be turned on or off based upon user requirements.⁸

How do I migrate my current Netezza system to the public cloud?

To migrate data from IBM PureData[®] System for Analytics to Netezza on Cloud, use the backup and restore method, or the `nz_migrate` command.

Visit the [IBM Knowledge Center](#) for more information on both of these methods.

Note: To migrate or connect to Netezza on Cloud from Mako, make sure you are on version 7.2.1.9 P1 or lower.

Get started

Whether you're using Netezza already or looking to add Netezza to your architecture for the first time, you owe it to yourself and your users to be on the best version of Netezza available: Netezza Performance Server. Fortunately, upgrading to Netezza Performance Server is easy and can actually save you money. Contact your IBM representative to get started or learn more at the links below.

See our demos

to learn how you can easily migrate to Netezza Performance Server.

Watch now →

Read this analyst report

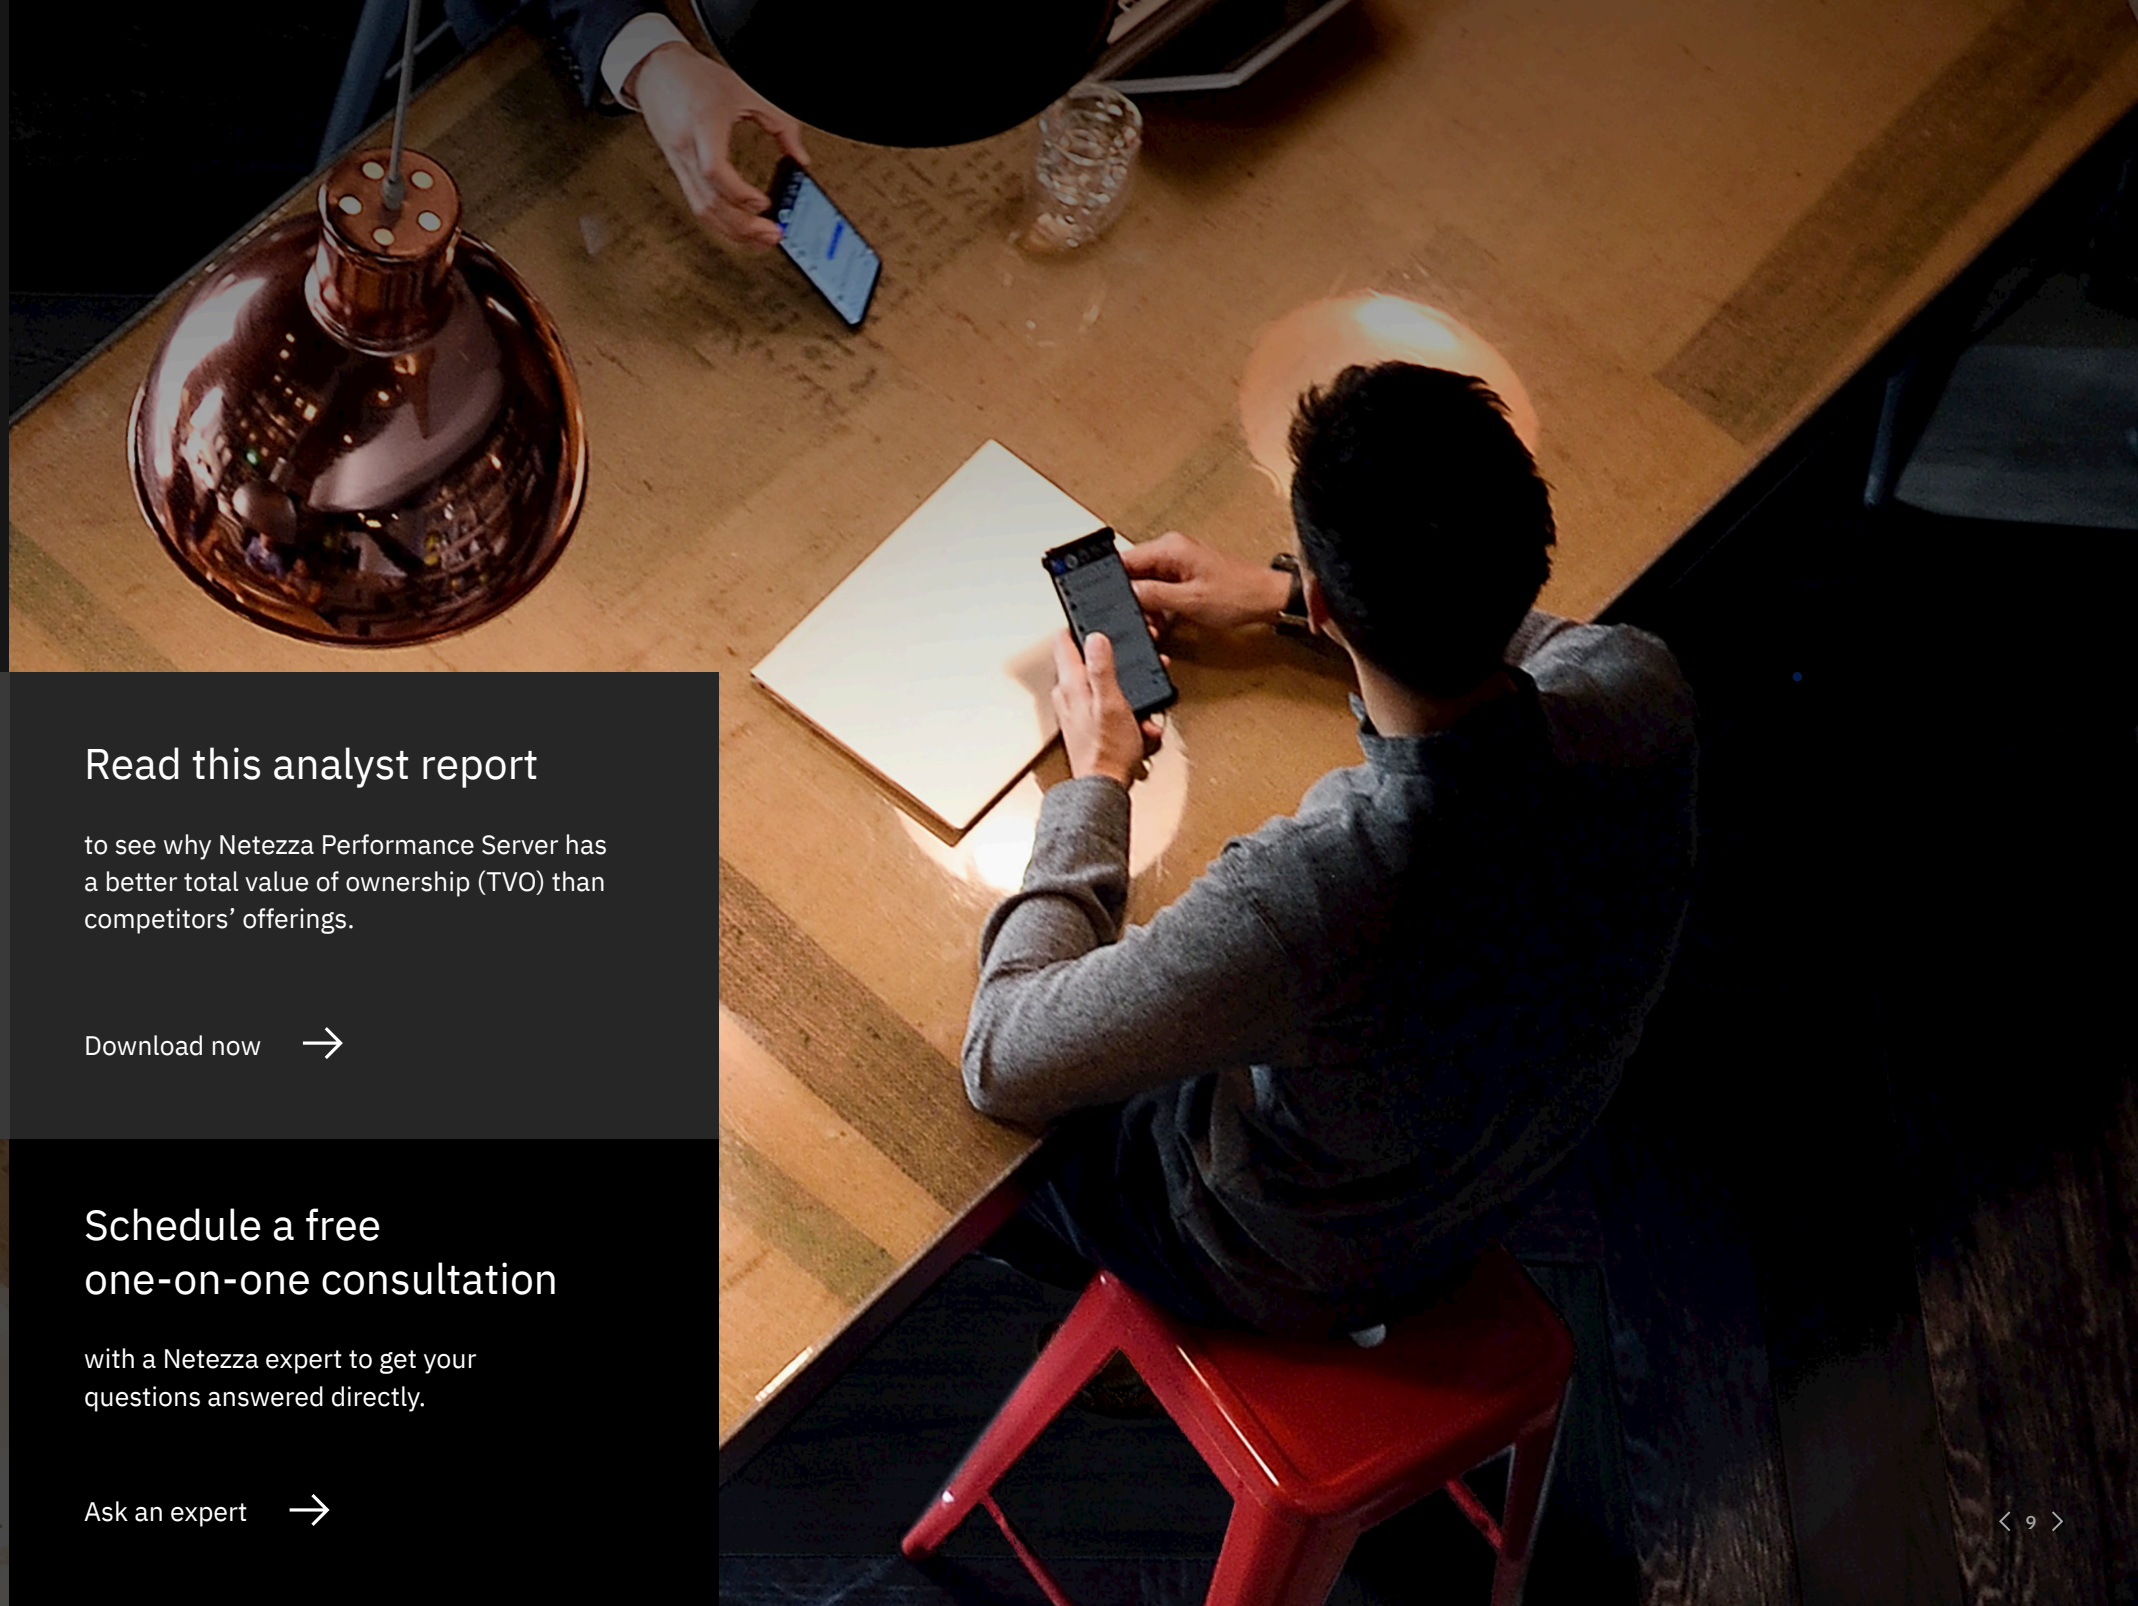
to see why Netezza Performance Server has a better total value of ownership (TVO) than competitors' offerings.

Download now →

Schedule a free one-on-one consultation

with a Netezza expert to get your questions answered directly.

Ask an expert →



IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
September 2020

IBM, the IBM logo, ibm.com, Netezza, IBM Cloud Pak, IBM Cloud, DataStage, Cognos, IBM Watson, and PureData are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Red Hat® and OpenShift® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions. It is the user’s responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs. THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

The client is responsible for ensuring compliance with laws and regulations applicable to it. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the client is in compliance with any law or regulation.

Statement of Good Security Practices: IT system security involves protecting systems and information through prevention, detection and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed, misappropriated or misused or can result in damage to or misuse of your systems, including for use in attacks on others. No IT system or product should be considered completely secure and no single product, service or security measure can be completely effective in preventing improper use or access. IBM systems, products and services are designed to be part of a lawful, comprehensive security approach, which will necessarily involve additional operational procedures, and may require other systems, products or services to be most effective. IBM DOES NOT WARRANT THAT ANY SYSTEMS, PRODUCTS OR SERVICES ARE IMMUNE FROM, OR WILL MAKE YOUR ENTERPRISE IMMUNE FROM, THE MALICIOUS OR ILLEGAL CONDUCT OF ANY PARTY.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

- 1 Expected availability in 2H 2020
- 2 Based on IBM experience with Netezza migrations
- 3 Based on internal and customer-reported tests
- 4 IBM Cloud Pak for Data System: The Leading Hyper-Converged Platform for Data and AI, by Ravi Shankar and Srini Chari, Cabot Partners, April 2020.
- 5 SwiftStack, five-year TCO
- 6 Expected availability in 2H 2020
- 7 Expected availability in 2H 2020
- 8 IBM Cloud Pak for Data System: The Leading Hyper-Converged Platform for Data and AI, by Ravi Shankar and Srini Chari, Cabot Partners, April 2020.

WEK04KWM