

Nutanix Data Services for Kubernetes (NDK)

As the adoption of cloud-native applications continues to grow, there is a pressing need to address the challenge of providing application-aware data services for Kubernetes. Current tools for developers and administrators only address the container layer, necessitating multiple third-party tools to solve for the application/namespace layers.

Nutanix Data Services for Kubernetes (NDK) simplifies and unifies the process of provisioning and operating applications by extending enterprise data services to containerized applications. As a result, NDK delivers a faster time to value and reduces the operational risk for administrators, allowing a cloud operating model for developers while also increasing cost efficiency for business owners.

Faster Time to Value With Application-Level Data Services

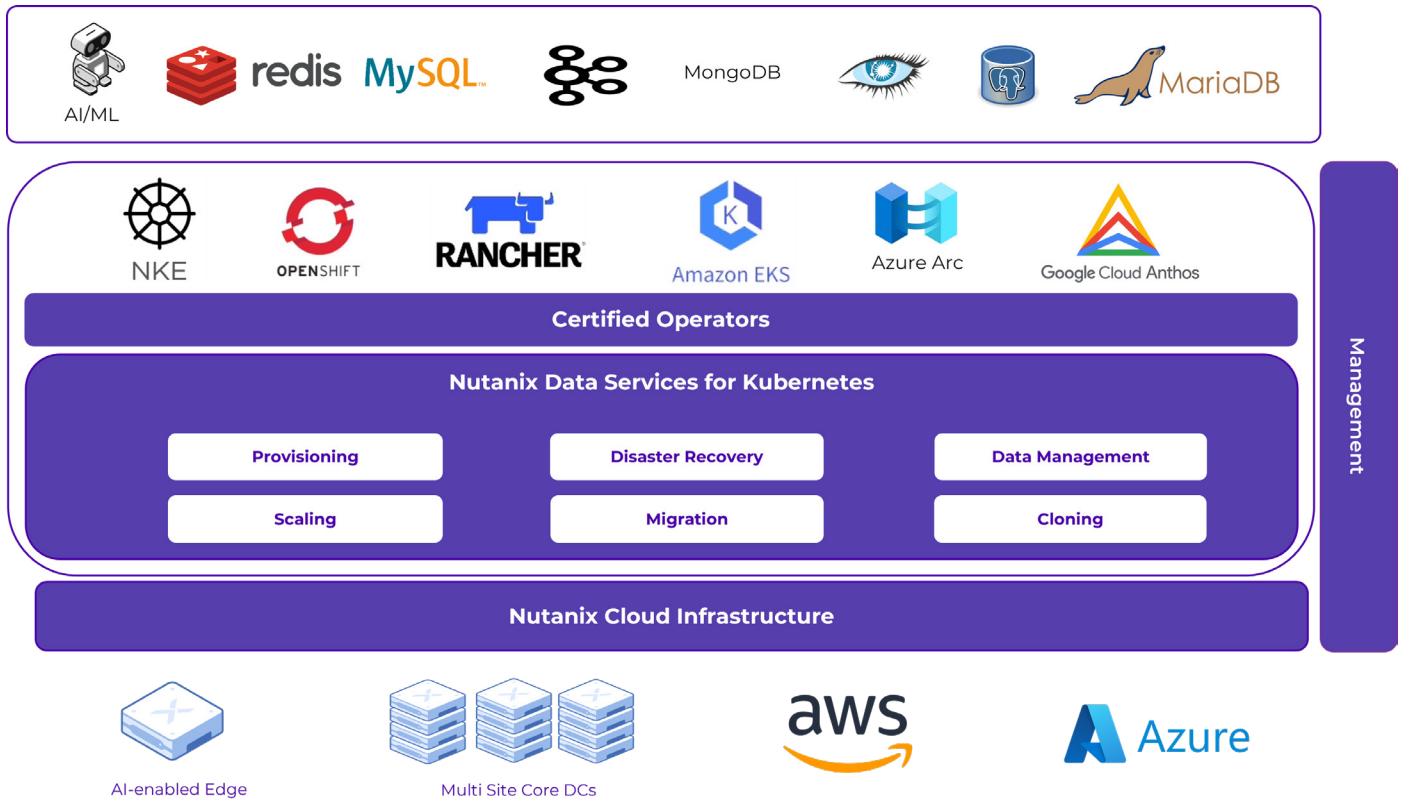
Kubernetes storage drivers today provision and manage at the container level while organizations struggle to provision and manage applications at the application level. This struggle is due to application/namespace level management requires third-party tools that increase time to manage, reduce efficiency, and cause migration, backup, and disaster recovery (DR) headaches.

With NDK, organizations can extend enterprise data services to modern applications by bringing application-aware data services to Kubernetes, allowing administrators to:

- move operations up to the application level
- accelerate application provisioning and streamline operations of data services
- define, manage, protect, and migrate applications and namespaces
- unify application management at scale and define application-level attributes to persistent storage features
- gain resiliency of performance, availability, encryption, and compression
- utilize 3-step deployment including defining the app/namespace, the governance policy, and a protection plan

Key Benefits

- Stop managing containers: Define, manage, protect, and migrate applications and namespaces.
- 3-Step deployment: Provide application-level data services in minutes
- Delegate infrastructure responsibility: With policy-based guardrails, enable Kubernetes administrators to automate workflows
- Developer self-service: Create application policies to drive complex self-service operations
- No new developer tools: NDK extends existing Kubernetes CLI (kubectl) and APIs, removing the need for additional training and tools.



Reduce Risk Through Simplification and Unification

Operational risk increases with complexity and Kubernetes is a dynamic, distributed platform that becomes difficult to manage quickly. In addition, management of Day 2 operations and monitoring is a constant challenge. Lastly, one-to-one mappings of storage to a container with storage drivers create an incomplete solution for application-level protection.

To solve all these problems, NDK extends Nutanix enterprise data services to any Kubernetes environment.

- Reducing operational risk at scale with NDK's policy-driven management for stateful applications. Easily ensure governance and standardization across all applications.
- Simplification through Prism Central, one unified dashboard that includes application/namespace level observability and capacity planning, Abstract multiple Kubernetes clusters and Nutanix storage services into a single management plane.
- With no new developer tools, NDK extends existing Kubernetes CLI (kubectl) and APIs, removing the need for additional training and tools.

Enable Cloud Operating Model

Application and operations teams increasingly require on-premises infrastructure to operate like the public cloud. Traditional infrastructure often doesn't allow for role-based delegation and developer self-service of modern applications. Organizations require the ability to automate operations and location transparency, regardless of whether data is on-premises or in the cloud.

For both Kubernetes administrators and developers, NDK enables a cloud operating model.

- Increase application development velocity so IT administrators can delegate infrastructure responsibility to Kubernetes administrators and developers with policy-based guardrails.
- Enable automation workflows that reduce the dependency on IT administrators.
- Enable self-service and create application policies to drive complex self-service operations for backup and DR.
- Application/namespace mobility from on-premises to the cloud without the complexity of multiple tools and synchronization, and manual restoration of metadata.

Increase Cost Efficiency

Business owners commonly struggle to properly implement modern applications across the multicloud. This inability to properly implement applications coupled with a lack of visibility into current and future resource consumption leads to wasteful expenditures and mismanaged infrastructure budgets.

Nutanix Data Services for Kubernetes streamlines operations, reduces friction, and optimizes cost efficiency. With NDK, businesses can:

- abstract away infrastructure dependence, on-premises and in the cloud
- reduce risk, optimize resource allocations, and save money with streamlined operations
- lower management costs by simplifying the operations of cloud-native applications at the container, namespace, and application level
- instill a cloud-like user experience freeing up customer budget and time for other strategic products and projects (Nutanix provides infrastructure that is scalable, resilient and operates in a cloud-like model)

Click here to learn how Nutanix can help you accelerate your cloud native journey <https://www.nutanix.com/solutions/cloud-native>



NDK extends enterprise data services to modern applications by bringing application-aware data services to Kubernetes.

NUTANIX

info@nutanix.com | www.nutanix.com | [@nutanix](https://twitter.com/nutanix)

©2023 Nutanix, Inc. All rights reserved. Nutanix, the Nutanix logo and all product and service names mentioned herein are registered trademarks or trademarks of Nutanix, Inc. in the United States and other countries. All other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s). filename here including version followed by date (datasheet-template MMDDYYYY)