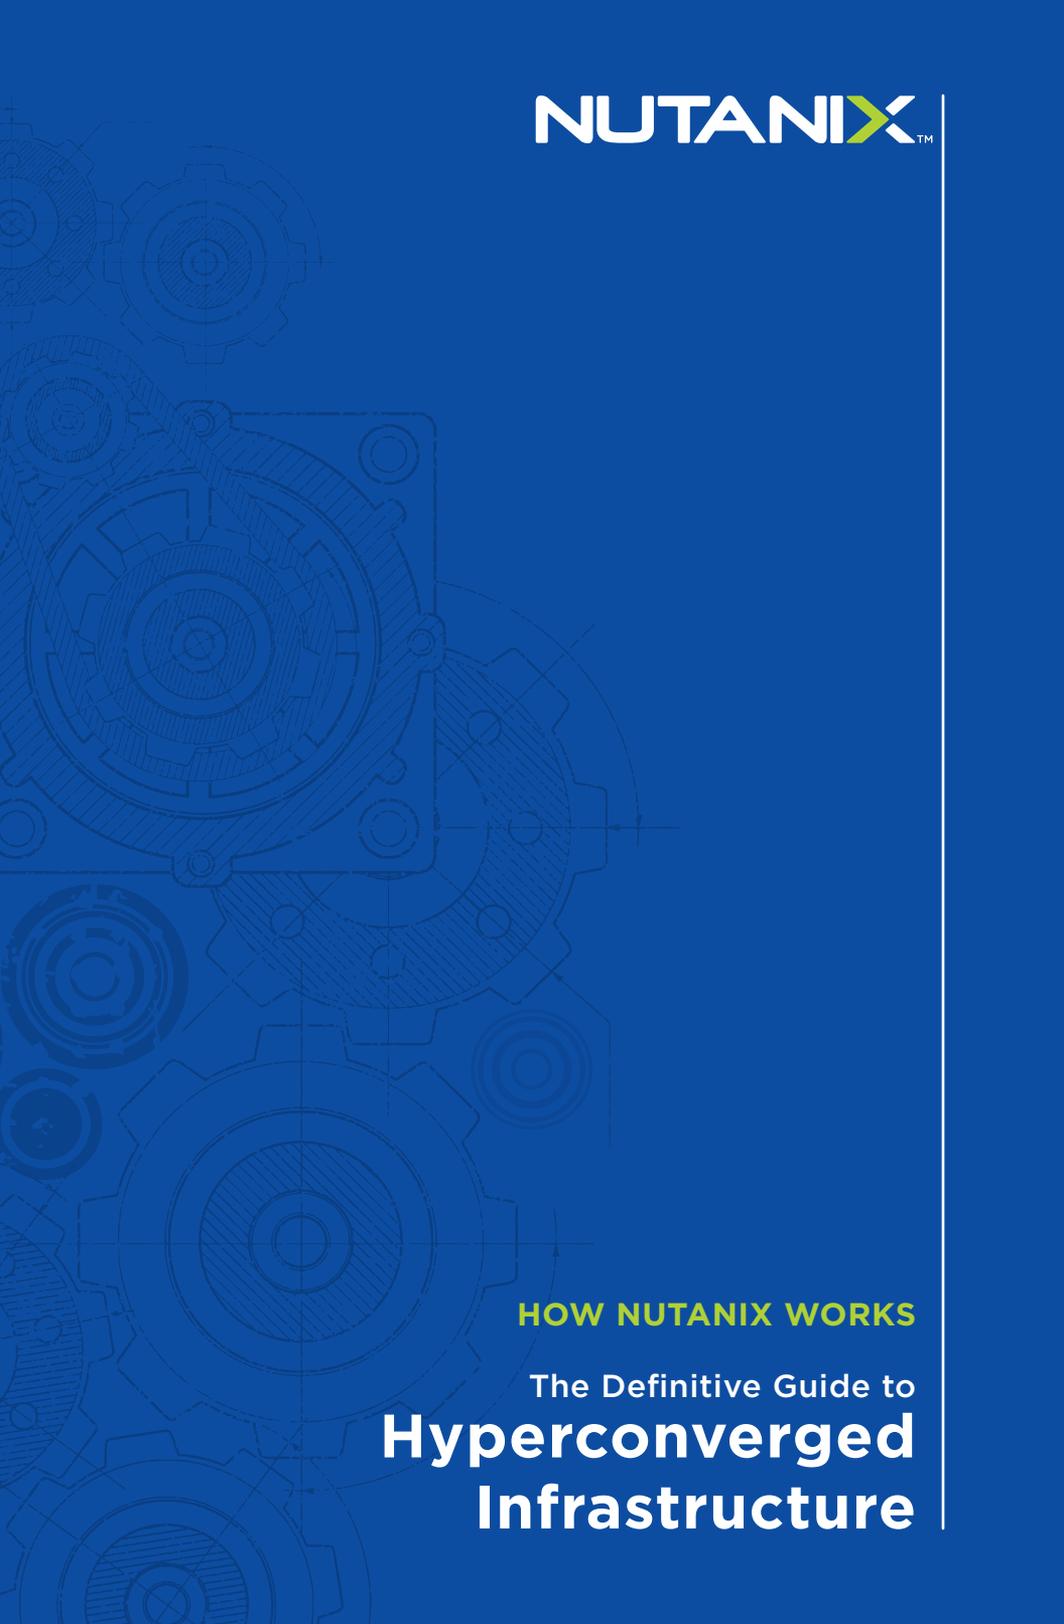


The Nutanix logo is positioned at the top right of the page. It features the word "NUTANIX" in a white, sans-serif font, with a green "X" that has a white diagonal line through it. A small "TM" trademark symbol is located to the right of the "X".

NUTANIX™

The background of the entire page is a deep blue color. Overlaid on this background are several white line-art drawings of interlocking gears of various sizes and orientations. The gears are rendered in a technical, blueprint-like style with fine lines and hatching to indicate depth and texture. They are scattered across the page, with some being larger and more prominent than others.

HOW NUTANIX WORKS

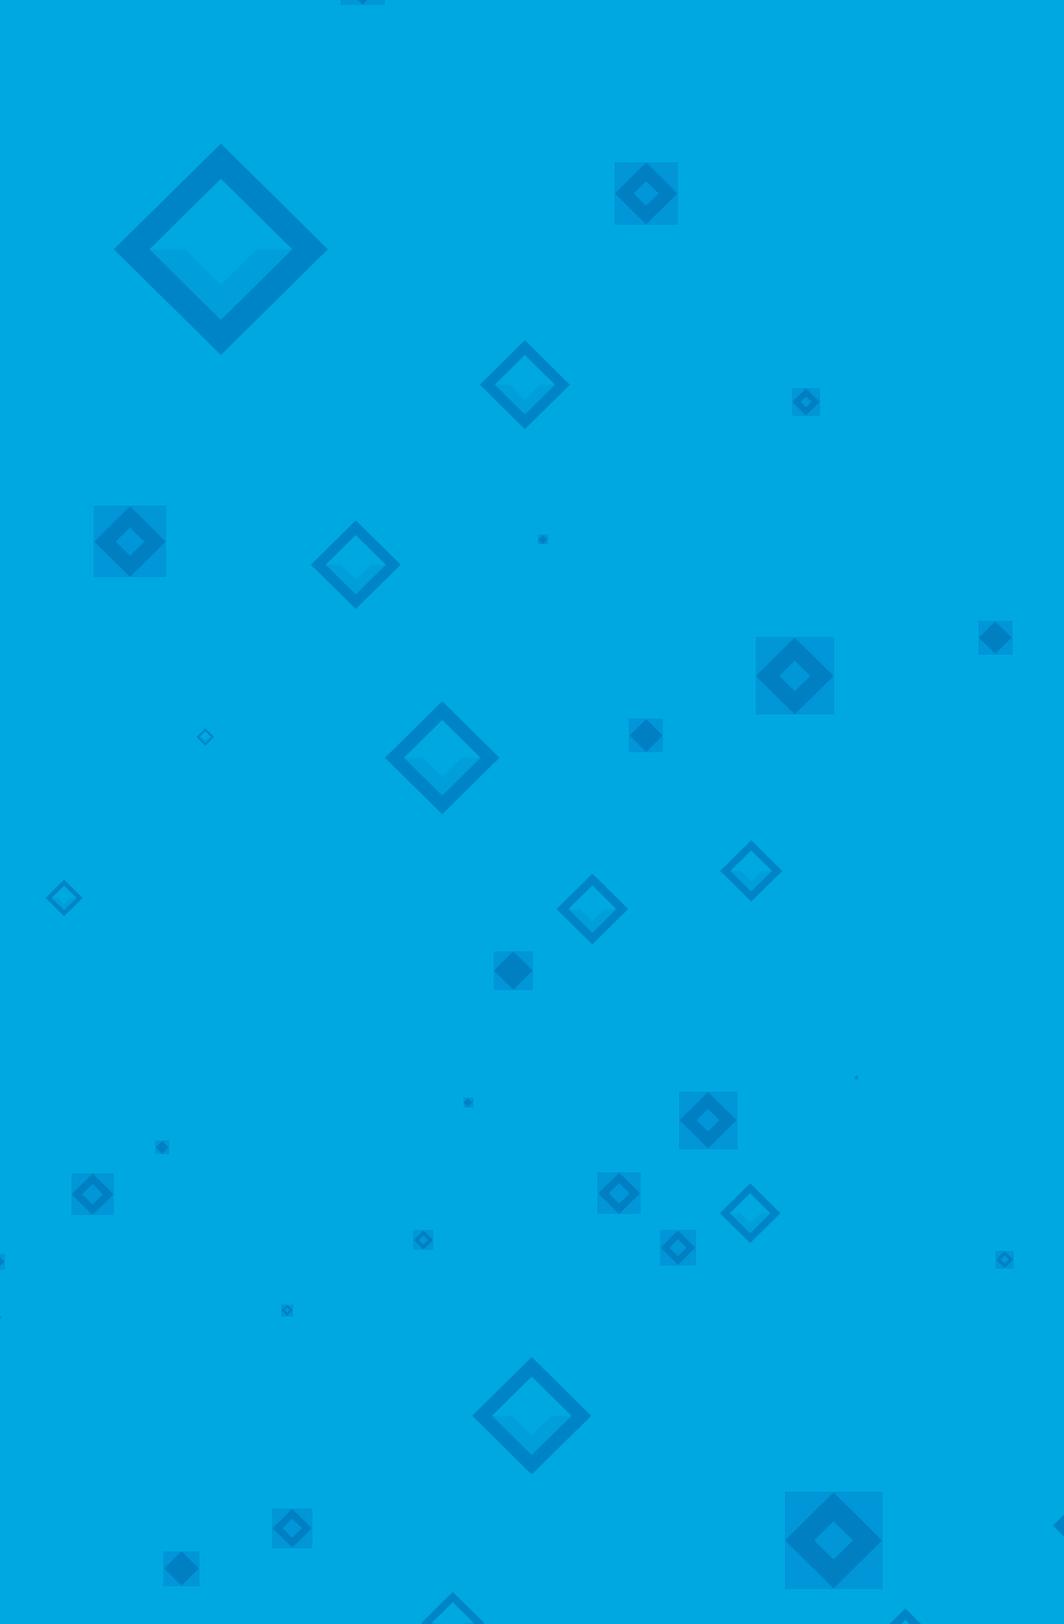
The Definitive Guide to
**Hyperconverged
Infrastructure**

Table of Contents

IT at a Crossroads	5
<i>Time for a Different Approach?</i>	6
<i>What is Hyperconverged Infrastructure?</i>	6
The Nutanix Solution	7
<i>Nutanix Community Edition and Community Edition On-Demand</i> ..	8
Prism and Acropolis	8
How Nutanix Software Is Deployed	9
<i>Nutanix Leads the Pack</i>	9
Acropolis	11
Distributed Storage Fabric (DSF)	11
Acropolis Hypervisor (AHV).....	11
App Mobility Fabric (AMF).....	11
<i>Distributed Storage Fabric (DSF)</i>	12
<i>Infrastructure Resilience</i>	13
Tunable Redundancy	13
<i>Replication Factor versus RAID</i>	13
Data Path Redundancy	13
<i>Nutanix Software Upgrades and Data Path Redundancy</i>	14
Integrity Checks	14
Availability Domains.....	14
Performance Acceleration	15
Intelligent Tiering	15
Data Locality.....	15
Automatic Disk Balancing	15
VM Flash Mode	16
Shadow Clones	17
Capacity Optimization	19
Deduplication	19
Compression	19
<i>Pro Tip: Compression</i>	20
EC-X.....	20

Data Protection	21
<i>What are RTO and RPO?</i>	21
Converged Local Backups With Snapshots And Time Stream	22
Integrated Remote Backup and DR Using Async Replication	22
<i>Self-Service File Restore</i>	22
Cloud Connect	23
Metro Availability and Sync Replication	23
Security	25
Data-at-Rest Encryption	25
Two-Factor Authentication & Cluster Lockdown	26
Security Automation with SaltStack	26
Hypervisors & Application Mobility	27
<i>It's a Multi-Hypervisor World</i>	27
Application Mobility Fabric	27
<i>Hypervisor scalability limits</i>	29
AHV Live Migration	30
AHV Data Protection	30
VM High Availability (VM-HA)	31
AHV Networking	31
<i>High Availability Out of the Box</i>	31
Datacenter Management with Nutanix Prism	33
<i>Prism is Highly Available by Design</i>	33
The Prism Approach	33
<i>Prism Keyboard Shortcuts</i>	36
<i>Performing Software Upgrades</i>	36
<i>Pro Tip: Prism Central</i>	37
<i>Who Would You Rather Call?</i>	38
About Nutanix	39





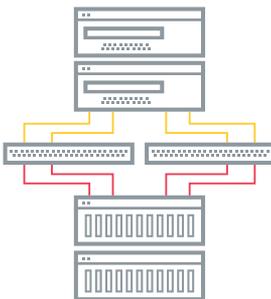


1.

IT at a Crossroads

IT is increasingly being asked to spend less time on infrastructure and allocate more time and budget to application services that add business value. Despite a continuous stream of IT hardware and software enhancements, the infrastructure challenges faced by IT teams continue to rise. The IT infrastructure and virtualization software required to meet the needs of business is complex and expensive, and datacenter management has become painful. Far too much time and effort are focused on just keeping the lights on.

Legacy infrastructure—with separate storage, storage networks, and servers—is not well suited to meet the growing demands of enterprise applications or the fast pace of modern business. The silos created by traditional infrastructure have become a barrier to change and progress, adding complexity to every step from ordering to deployment to management. New business initiatives require buy-in from multiple teams, and IT needs have to be predicted 3-to-5 years in advance. As most IT teams know, this is almost impossible to get right. In addition, vendor lock-in and increasing licensing costs are stretching budgets to the breaking point.



1. INHERENT COMPLEXITY
2. INEFFICIENT SILOS
3. FORKLIFT SCALING
4. PAINFUL MANAGEMENT

FIGURE 1: Challenges of legacy three-tier infrastructure

Time for a Different Approach?

Enterprise IT teams today are looking for ways to deliver on-premises IT services with the speed and operational efficiency of public cloud services such as Amazon Web Services (AWS), Microsoft Azure and Google Compute Engine.

Taking cues from web giants, hyperconverged infrastructure combines x86-based compute and storage resources with intelligent software to eliminate common pain points associated with legacy infrastructure.

Nutanix delivers a comprehensive enterprise cloud platform that bridges the wide gap that exists between traditional infrastructure and public cloud services. The solution delivers turnkey infrastructure that integrates servers, storage and virtualization along with end-to-end systems management and operations management capabilities. This allows enterprises to deploy infrastructure in minutes and shift the focus to applications that power the business.

WHAT IS HYPERCONVERGED INFRASTRUCTURE?

Hyperconverged infrastructure combines x86-based compute and storage resources with intelligent software to create flexible building blocks that replace legacy infrastructure consisting of separate servers, storage networks, and storage arrays.

While hyperconvergence is not an end point in itself, it is the fundamental building block for enterprise cloud. This book gives an overview of the Nutanix hyperconverged solution and walks through how different features and functionality work to provide a fast, highly scalable and efficient datacenter solution for enterprises of all sizes.



2.

The Nutanix Solution

Nutanix converges the entire datacenter stack including compute, storage, storage networking, and virtualization. Complex and expensive legacy infrastructure is replaced by simple 2U appliances that enable enterprises to start small and scale one node at a time. Each server, also known as a node, includes Intel-powered x86 hardware with flash SSDs and HDDs. Nutanix software running on each node distributes all operating functions across the cluster for superior performance and resilience.

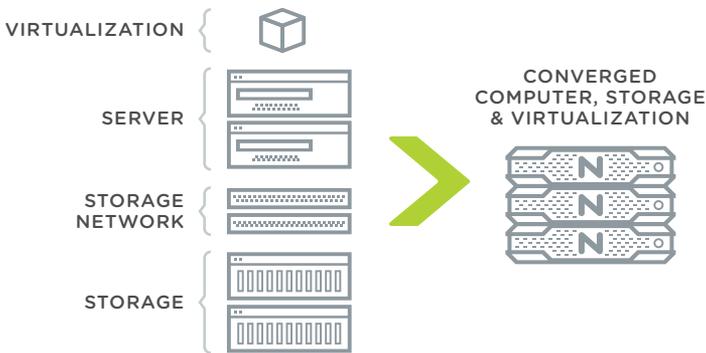


FIGURE 2: Nutanix converges compute, storage, and virtualization in simple, scalable building blocks

A single Nutanix cluster can have an unlimited number of nodes. Different hardware platforms are available to address varying workload needs for compute and storage. Nutanix software is hardware agnostic and runs on Nutanix hardware, and also hardware from OEMs such as Dell and Lenovo.

***NUTANIX COMMUNITY EDITION AND
COMMUNITY EDITION ON-DEMAND***

[Community Edition](#) is a free, 100% software solution that lets enterprises easily evaluate the latest Nutanix technology at zero cost on existing hardware.

Hardware platforms include “compute-heavy,” “storage-heavy,” and “storage-only” options. All nodes include flash to optimize storage performance, and all-flash nodes are available to deliver maximum I/O throughput with minimum latency for all enterprise applications.

Prism and Acropolis

Nutanix software has two components: Acropolis and Prism. Acropolis is a distributed data plane with enterprise storage and virtualization services, and the ability for applications to move seamlessly across hypervisors and in the long run cloud providers. Prism is a distributed management plane that uses advanced data analytics and heuristics to simplify and streamline common workflows, eliminating the need for separate management solutions for servers, storage networks, storage and virtualization.

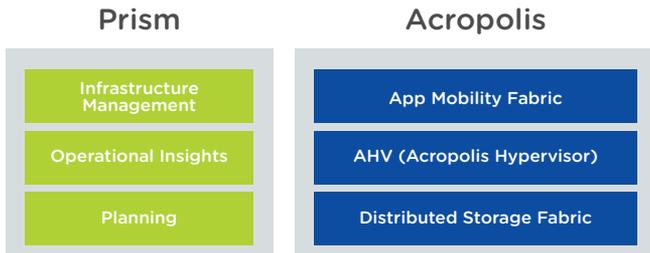


FIGURE 3: Key functions of Acropolis and Prism

How Nutanix Software Is Deployed

A Nutanix cluster is 100% software defined. Each node in a cluster runs a hypervisor (VMware ESXi, Microsoft Hyper-V, or Nutanix Acropolis Hypervisor), and the Nutanix software runs as a virtual machine called the Controller VM (CVM) that runs on every node in the cluster. This allows Nutanix systems to be hypervisor agnostic. The CVM includes Prism management functions and Acropolis data plane functions.

NUTANIX LEADS THE PACK

Nutanix is now the acknowledged leader in hyperconverged infrastructure and Gartner positions Nutanix as a leader in its 2015 Magic Quadrant for Integrated Systems



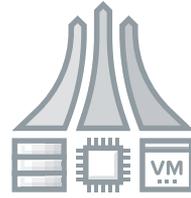
“

“My key requirements were to have something that was simple, easy to manage, and ideally a single pane of glass. I wanted a solution that was very powerful and also very versatile. For me, Nutanix ticked all of those boxes.”

PURDIP BAHRA

IT Manager, Joseph Chamberlain College

3.



Acropolis

Nutanix Acropolis has three major components:

DISTRIBUTED STORAGE FABRIC (DSF)

- Enterprise storage services for applications, eliminating the need for separate solutions from vendors such as NetApp, EMC, and HP
- Includes a comprehensive set of capabilities for performance acceleration, data reduction, data protection, and much more
- Full support for VMware® vSphere, Microsoft® Hyper-V and Nutanix Acropolis Hypervisor

ACROPOLIS HYPERVISOR (AHV)

- Natively built virtualization solution
- Based on the proven Linux KVM hypervisor, AHV is hardened to meet the most stringent enterprise security requirements
- Integrated management through Prism

APP MOBILITY FABRIC (AMF)

- Intelligent virtual machine (VM) placement, migration, hypervisor conversion, and cross-hypervisor high availability for maximum flexibility.

Distributed Storage Fabric (DSF)

The Acropolis Distributed Storage Fabric is designed to simplify storage and data management for virtual environments. By pooling flash and hard disk drive storage across a Nutanix cluster and exporting it out to the virtualization layer as iSCSI, NFS and SMB shares, DSF eliminates the need for SAN and NAS solutions.

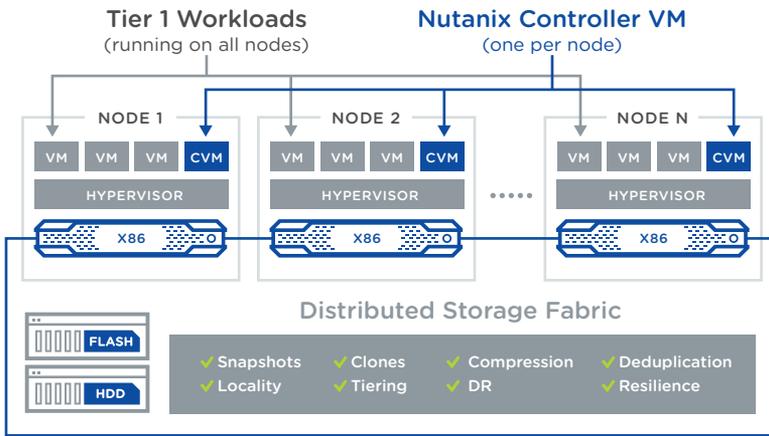


FIGURE 4: Acropolis Distributed Storage Fabric joins HDD and SSD resources from across a cluster into a storage pool.

How DSF Organizes Data

There are a few key concepts that are important with regard to how DSF organizes data:

Storage Pool. A group of physical storage devices including SSD and HDD devices across the entire cluster. The storage pool spans multiple Nutanix nodes and is expanded as the cluster scales.

Storage Container. A logical segment of a Storage Pool. Containers typically have a 1-to-1 mapping with a VM datastore.

vDisk. A vDisk is any file over 512KB on DSF including .vmdk files and VM hard disks. vDisks are composed of extents which are grouped and stored on disk as an extent group.

Infrastructure Resilience

The Nutanix platform is fault resistant, with no single points of failure and no bottlenecks.

TUNABLE REDUNDANCY

With Tunable Redundancy, each Nutanix container is configured with a replication factor (RF) of two or three. RF=2 ensures that two copies of data are maintained at all times, allowing the cluster to survive the failure of a single node or drive. When RF is set to 3 (RF=3), three copies of the data are maintained in a cluster, providing resilience against two simultaneous failures.

REPLICATION FACTOR VERSUS RAID

RAID has been a popular way of protecting against drive failures while limiting the extra storage capacity required. Rebuilding a multi-TB drive can take days to complete, creating a risk of data loss should further failures occur. RAID has gone from single to double and even triple-parity to try to reduce this risk.

Nutanix Replication Factor (RF) eliminates reliance on RAID, the need for expensive spare drives that sit idle, and the performance penalty that comes with multiple parity calculations.

DATA PATH REDUNDANCY

Data Path Redundancy ensures high availability in the event a Nutanix Controller VM (CVM) becomes unavailable or needs to be brought down for upgrade. If a CVM becomes unavailable for any reason, Nutanix CVM autopathing automatically re-routes requests to a “healthy” CVM on another node. This failover is fully transparent to the hypervisor and applications.

Data Path Redundancy is possible because every node in a cluster has access to all copies of data—I/O requests can be serviced immediately by any node in the system.

NUTANIX SOFTWARE UPGRADES AND DATA PATH REDUNDANCY

Nutanix software upgrades take advantage of reliable data path redundancy. While the local CVM is unavailable because of software upgrade or a failure, VMs running on the node use data path redundancy to satisfy I/O through a CVM on another node—transparent to users and applications.

INTEGRITY CHECKS

Acropolis has a variety of features to proactively identify and fix issues related to data consistency and integrity, bit rot failures, and hard disk corruption.

- Detection of silent data corruption and repair of data consistency errors
- Automatic data integrity checks during every read
- Automatic isolation and recovery during drive failures

AVAILABILITY DOMAINS

Availability Domains offer greater protection from hardware failures by allowing Nutanix clusters to survive the failure of a node or block (multi-node enclosure). Availability domains are created based on the granularity at which failures are likely to occur.

With DSF, data replicas will be written to other blocks in the cluster to ensure that in the case of a block failure or planned downtime, the data remains available. This is true for both RF2 and RF3 scenarios, as well as in the case of a block failure. An easy comparison would be “node awareness”, where a replica would need to be replicated to another node which will provide protection in the case of a node failure. Block awareness further enhances this by providing data availability assurances in the case of block outages.

4.



Performance Acceleration

DSF includes a number of capabilities that enhance performance:

INTELLIGENT TIERING

DSF continually monitors data access patterns and optimizes data placement on either the SSD or HDD tier, achieving the best performance without administrator intervention.

The SSD tier provides maximum performance for hot data and random I/O, while the HDD tier provides maximum capacity and economy for cold data and sequential I/O.

DATA LOCALITY

DSF ensures that as much of a VM's data as possible is stored on the node where the VM is running. This negates the need for read I/O to go through the network. Keeping data local optimizes performance and minimizes network congestion.

Every VM's data is served locally from the CVM and stored preferentially on local storage. When a VM is moved from one node to another using vMotion or Live Migration (or during an HA event), the migrated VM's data automatically follows the VM in the background based on read patterns.

AUTOMATIC DISK BALANCING

Automatic disk balancing ensures that data is distributed uniformly across the entire cluster. Any node in a Nutanix cluster can utilize storage resources across the cluster, without requiring time-consuming and error-prone manual rebalancing.

Automatic Disk Balancing reacts to changing workloads and allows heterogeneous nodes to be mixed in a single cluster. Once utilization reaches a set threshold, disk balancing keeps it uniform among nodes.

VM FLASH MODE

VM Flash Mode pins specific VMs or vDisks to the cluster-wide SSD tier, so that IOPS and latency-sensitive workloads can be mixed with other workloads in a single cluster, without compromising on performance. VM Flash Mode gives fine-grained control over I/O performance. For instance, all database transaction logs can be pinned in flash. Or all financial data can be pinned in flash during quarter-end reporting.

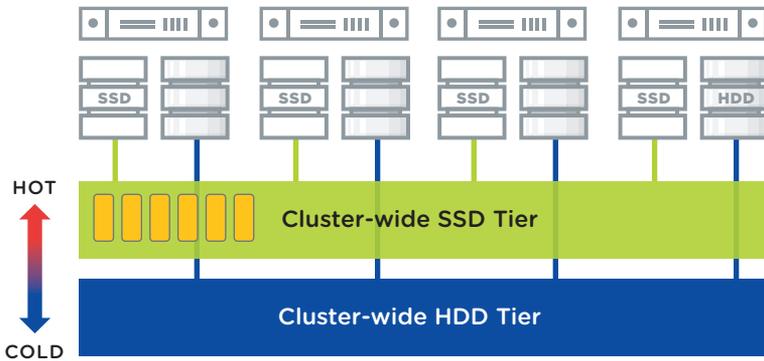


FIGURE 5: VM Flash Mode allows individual vDisks to be “pinned” in the cluster-wide SSD tier for maximum IOPS and low latency.

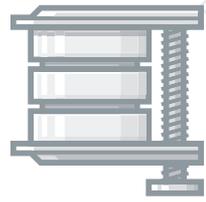
SHADOW CLONES

Shadow Clones significantly improve performance by caching virtual machine data across a Nutanix cluster. Unique to Nutanix, Shadow Clones benefit scenarios where there are multiple VMs reading a single source of data, such as deployment servers and repositories. VDI deployments, where many linked clones forward read requests to a central master (e.g., Citrix MCS Master VM or VMware View replica disks), are an ideal example.

With Shadow Clones, Nutanix actively monitors vDisk access trends. If there are requests originating from more than two remote Controller VMs (CVMs), as well as the local CVM, and all of the requests are read I/O, the vDisk will be marked as immutable. Once the disk has been marked immutable, the vDisk is then cached locally by each CVM so read operations are now satisfied locally by direct-attached storage resources.



5.



Capacity Optimization

DSF incorporates a wide range of storage optimization technologies that work together to make efficient use of the available capacity in a cluster.

DEDUPLICATION

Nutanix delivers two types of data deduplication to accelerate application performance and optimize storage capacity. Performance-tier deduplication removes duplicate data in the content cache (SSD and memory) to reduce the footprint of an application's working set. In addition, global post-process MapReduce deduplication reduces repetitive data in the capacity tier to increase the effective storage capacity of a cluster. Both forms of deduplication can be easily configured and managed at vDisk granularity.

When deduplication is enabled, data is fingerprinted on ingest using a SHA-1 hash. Deduplication operations are software-driven and leverage the hardware-assist capabilities of the Intel chipset for the SHA-1 fingerprint generation. Because SHA-1 is a strong hash, deduplication is performed based on a fingerprint match.

COMPRESSION

Data can be compressed inline as it is written to the system, or post process after the data has been written. Inline and Post process compression is intelligently determined based on sequential or random access patterns to enable optimal performance. Post-process compression is executed as a series of MapReduce jobs.

DSF uses the Google Snappy compression algorithm, providing good compression ratios with minimal computational overhead and extremely fast compression and decompression rates.

PRO TIP: COMPRESSION

Use inline compression most of the time; it will not impact random write performance. Inline compression pairs perfectly with erasure coding.

EC-X

Nutanix systems include an innovative erasure coding technology—Nutanix EC-X—that provides resilience and can increase usable capacity by up to 75%. EC-X reduces the capacity cost of replication factor (RF) without taking away any of the resilience benefits and with no impact on write performance.

EC-X encodes a strip of data blocks on different nodes and calculates parity. In the event of a disk or node failure, parity is used to calculate any missing data blocks. DSF uses an extent group as the data block, and each data block in a strip must be on a different node and belong to a different vDisk. The number of data and parity blocks in a strip is configured based on the desired number of failures to withstand.



6.

Data Protection

Nutanix offers natively integrated data protection and continuous availability at the VM level. A range of options is available to meet the recovery point objective (RPO) and recovery time objective (RTO) of different applications.

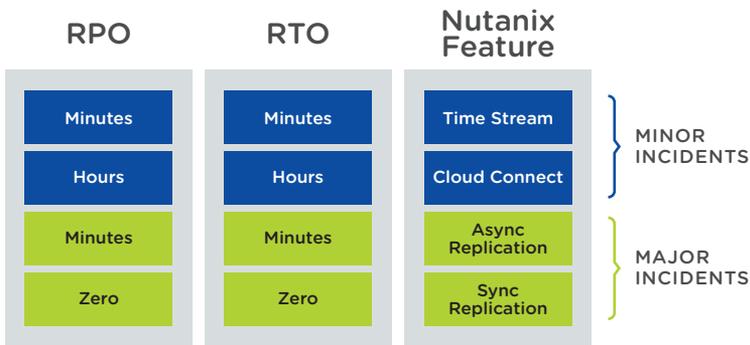


FIGURE 6: Nutanix data protection options

WHAT ARE RTO AND RPO?

Recovery Time Objective (RTO) defines how much time you have to recover if an IT failure occurs.

Recovery Point Objective (RPO) defines the maximum amount of data you are willing to lose.

CONVERGED LOCAL BACKUPS WITH SNAPSHOTS AND TIME STREAM

Nutanix Time Stream can create unlimited local snapshots—with VM and application-level consistency—and recover data instantly to meet a wide range of backup and data protection requirements.

Time Stream uses VM-centric snapshots to provide production-level data protection without sacrificing performance. Nutanix utilizes a redirect-on-write algorithm that dramatically improves system efficiency for snapshots.

Commvault IntelliSnap integration combines Commvault backup capabilities with enterprise storage features from Nutanix.

INTEGRATED REMOTE BACKUP AND DR USING ASYNC REPLICATION

Nutanix DR and replication capabilities are built on snapshot technology. VM snapshots can be asynchronously replicated or backed up to another datacenter based on a user-defined schedule. Replication topologies are flexible and bi-directional, enabling one-to-one, one-to-many, and many-to-many deployments. During replication, data is compressed and replicated at the sub-block level for maximum efficiency and lower WAN bandwidth consumption.

Nutanix Prism interface offers a simplified view of all local and remote snapshots, allowing administrators to restore a VM from a snapshot with a single click. In case of disaster, can failover to the secondary datacenter can be done with a single click.

SELF-SERVICE FILE RESTORE

Acropolis data protection includes self-service file restore, which allows users to recover individual files from VM snapshots without getting an administrator involved.

CLOUD CONNECT

Nutanix Cloud Connect lets enterprises use public cloud services, such as Amazon Web Services (AWS) and Microsoft Azure, as a long-term backup destination for all types of workloads, making them a logical extension of their own datacenters.

Cloud Connect for AWS provides a live Nutanix cluster in the cloud running on EC2 instances and using the Elastic Block Store for metadata and S3 storage for backups. For Microsoft Azure, the Nutanix software runs on Azure Compute and storage is from Azure Page Blob. Data transfer is WAN optimized, reducing the storage footprint and networking bandwidth by over 75%. Support for Amazon Virtual Private Cloud (VPC) and Azure Virtual Network (VNET) allows secure data transfer over an IP connection.

METRO AVAILABILITY AND SYNC REPLICATION

For critical workloads requiring zero RPO, and near-zero RTO, Nutanix offers Metro Availability ensuring continuous data availability across separate sites within a metro. Metro Availability is simple to set up and manage using Prism.

Metro Availability can be set up bi-directionally between two sites connected over a metro area network. The only network requirement is a round-trip latency of less than five milliseconds. Data is written synchronously to both sites, so it is always available to applications in the event a site fails or needs to undergo maintenance. Virtual machines can be non-disruptively migrated between sites for planned maintenance events or other needs.



“

“We were focused on flexibility and innovation. We were looking for a partner who would be able to understand our business needs. With Nutanix, there was a willingness to listen and propose an innovative solution.”

LAURENT PERRIAULT

Director of Operations, Claranet

7.



Security

Nutanix Acropolis is hardened by default. It utilizes the principle of least privilege, and delivers a true defense-in-depth model. Its custom security baseline exceeds the requirements of the U.S. Department of Defense.

Nutanix combines features such as two-factor authentication and data-at-rest encryption with a security development lifecycle. This is integrated into product development to help meet the most stringent security requirements. Nutanix systems are certified across a broad set of evaluation programs to ensure compliance with the strictest standards.



DATA-AT-REST ENCRYPTION

Data-at-rest encryption is delivered through self-encrypting drives (SED) that are factory-installed in Nutanix hardware. This provides strong data protection by encrypting user and application data for FIPS 140-2 Level 2 compliance. Acropolis interfaces with third-party key management servers using the industry-standard Key Management Interface Protocol (KMIP) instead of storing the keys in the cluster.

TWO-FACTOR AUTHENTICATION & CLUSTER LOCKDOWN

Nutanix solutions enforce two-factor authentication for system administrators in environments requiring additional layers of security. When implemented, administrator logins require a combination of a client certificate and username/password.

Nutanix also offers Cluster Shield, which restricts access to a Nutanix cluster in security-conscious environments such as government and healthcare datacenters. Cluster Shield not only disables interactive shell logins automatically but can also enable more restrictive access based on those keys.

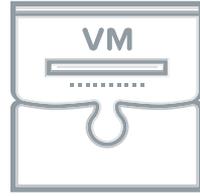
Nutanix uses a unique, well-defined Security Development Lifecycle (SecDL) to incorporate security into every step of the software development process, from design and development to testing and hardening. Threat modeling is used to assess and mitigate customer risk from code changes. SecDL testing is fully automated during development, and all security-related code modifications are timed during minor releases to minimize risk.

The Nutanix Security Technical Implementation Guide (STIG) is written in the eXtensible Configuration Checklist Description Format (XCCDF), allowing it to be read by various automated assessment tools, such as Host Based Security System (HBSS). This provides detailed information on how to assess a Nutanix system to determine compliance with the STIG requirement, cutting down the accreditation time from 9-12 months to a matter of minutes.

SECURITY AUTOMATION WITH SALTSTACK

SaltStack is a robust, open-source automation and management framework that provides a simple way to check and fix a system baseline. Acropolis uses SaltStack to self-heal any deviation from the security baseline configuration of the operating system.

8.



Hypervisors & Application Mobility

Acropolis provides an open platform for virtualization and application mobility by taking advantage of the same underlying web-scale architecture.

IT'S A MULTI-HYPERVISOR WORLD

According to IDC, more than 72% of enterprises in 2015 are using more than one hypervisor, up from 59% in 2014.

APPLICATION MOBILITY FABRIC

The Acropolis App Mobility Fabric (AMF) is a collection of powerful technologies built into the Nutanix solution that allows applications and data to move freely between runtime environments. AMF includes a broad range of capabilities for migrating between different environments, including:

- Non-Nutanix infrastructure to Nutanix systems
- Between Nutanix systems supporting different hypervisor environments
- Nutanix to a public cloud infrastructure

Acropolis App Mobility Fabric Features:

Nutanix Sizer: Select the right Nutanix system & deployment configuration to meet the needs of each workload

Foundation: Automatically install the hypervisor of your choice on a Nutanix cluster

High Availability: Automatically restart virtual machines on healthy nodes in case of a node failure in a Nutanix cluster

Resource Scheduling: Place virtual machines intelligently on nodes in a cluster based on rich analytics to balance compute and storage utilization

Cloud Connect: Use built-in hybrid cloud technology for seamless data backup and disaster recovery to public cloud services

Cross-hypervisor Backup: Backup application data on remote clusters running different hypervisors and recover quickly with only a single click.

Acropolis AMF relies on the underlying Acropolis Distributed Storage Fabric in (DSF) to provide data services such as VM-centric provisioning, snapshots, clones, data protection, resilience and availability for all applications. Acropolis Hypervisor

Acropolis Hypervisor (AHV) was built from the ground up to provide a much simpler and more scalable hypervisor and associated management platform by leveraging the software intelligence of the hyperconverged architecture. AHV changes the core building block of the virtualized datacenter from hypervisor to application and liberates virtualization from the domain of specialists – making it simple and easily manageable by anyone from DevOps teams to DBAs.

Hypervisor scalability limits

	VMWARE ESXI*	MICROSOFT HYPER-V**	AHV
Max Cluster Size	64	64	Unlimited
Max vCPUs per VM	128	64	160
Max Memory per VM	4TB	1TB	2TB
Max VMs per host	1,024	1,024	Only limited by RAM
Max VMs per Cluster	8,000 (2,048 per datastore w/ HA)	8,000	Only limited by RAM
Offloads	VAAI support	ODX support	Full integration
Monitoring and Mgmt	Monitoring	Monitoring	Entire lifecycle

*VMware vSphere 6.0 **Hyper-V 2012 R2

AHV is based on the proven Linux KVM hypervisor to ensure support for all popular workloads, and is hardened to meet the most stringent enterprise security requirements. (See Section 7 for more on security.) The converged storage and virtualization stack dramatically simplifies virtualization.

It is fully supported by Nutanix, which means enterprises get full infrastructure and virtualization support from a single vendor. AHV is included with Nutanix purchases with no additional licensing costs.



Simplicity

- VM-centric policies
- Single pane of glass
- Single vendor support
 - High automation



Scalability

- VM management scales in step with the cluster
- Complements pay-as-you-grow economics



Security

- Security built into SDLC
 - Custom STIG
- Self-healing with SaltStack



- First "truly distributed" virtualization management
 - Built with web-scale technologies
- Purpose-built for hyperconverged architecture
 - Included at no additional cost

FIGURE 7. Advantages of the Acropolis Hypervisor

AHV is designed to take advantage of the intelligent storage services provided by Acropolis DSF. Because DSF is optimized for use with server virtualization, it provides full data resiliency, and data services such as snapshots, clones, provisioning operations at VM granularity. As a result, AHV is leaner and focused on delivering secure virtual compute services and high availability.

AHV LIVE MIGRATION

The ability to move running VMs between hosts has become a must have feature for all virtual environments. AHV Live Migration is analogous to Hyper-V live migration or VMware vMotion. AHV Live Migration enables Admins to move a user VM from one host to another while the VM is powered on. It follows similar resource rules to those used by VM-HA to determine if migration is possible.

AHV DATA PROTECTION

Ensuring data protection for VMs running on AHV is a simple process. Each VM is automatically protected according to a designated schedule that can include local snapshots as well as replication to a remote site. AHV has full access to all the data protection capabilities of Acropolis as described earlier.

VM HIGH AVAILABILITY (VM-HA)

With AHV, Nutanix takes high availability for virtual machines to a new level. VM High Availability (VM-HA) works in conjunction with the resiliency built into the Distributed Storage Fabric, Prism, and other Nutanix services. Out of the box, a Nutanix cluster running AHV is pre-configured to provide “Best Effort High Availability.” It automatically responds to node failures by restarting VMs on healthy nodes as long as the cluster has available capacity.

AHV NETWORKING

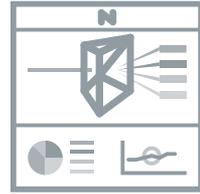
Acropolis Hypervisor uses Open vSwitch (OVS) to connect the Controller VM, hypervisor, and guest VMs to each other and to physical networks. OVS runs on each AHV node and starts automatically, greatly simplifying networking within a Nutanix cluster.

HIGH AVAILABILITY OUT OF THE BOX

AHV delivers high availability with minimal configuration and no additional software purchase. When a node fails, VMs are automatically restarted on other nodes in the cluster.



9.



Datacenter Management with Nutanix Prism

Nutanix Prism provides an easy way to manage Nutanix environments end to end. Prism combines multiple aspects of datacenter management into a single consumer-grade product that lets IT admins manage infrastructure and virtualization, gain access to operational insights, and fix problems all with a few clicks.

Just as Acropolis creates a data plane that spans the entire cluster, Prism creates a cluster-wide management plane. High Availability is built right into the product removing the need to have to manage the management solution.

PRISM IS HIGHLY AVAILABLE BY DESIGN

Because Prism runs on every node in a cluster and there are no external servers or databases to configure.

In addition to the Prism interface, all management capabilities are exposed through REST APIs, PowerShell, and through the command line interface (CLI) to facilitate integration and automation.

THE PRISM APPROACH

Prism is designed for an uncluttered experience with an intuitive user interface that simplifies and streamlines common datacenter workflows, eliminating the need to have different management tools

for different tasks. Prism enhances productivity through features such as:

- Instant Search: Integrated search to query and perform actions quickly.
- Capacity Planning: Prism’s predictive analysis engine forecasts the capacity needs of applications running on a Nutanix cluster, giving the IT team the ability to proactively understand and to plan for infrastructure needs.
- Customizable Operations Dashboard: Visual dashboard gives an at-a-glance summary of application and infrastructure state
- Integrated management with One-Click Simplicity: Infrastructure management, operational insights, and problem remediation can be accomplished easily

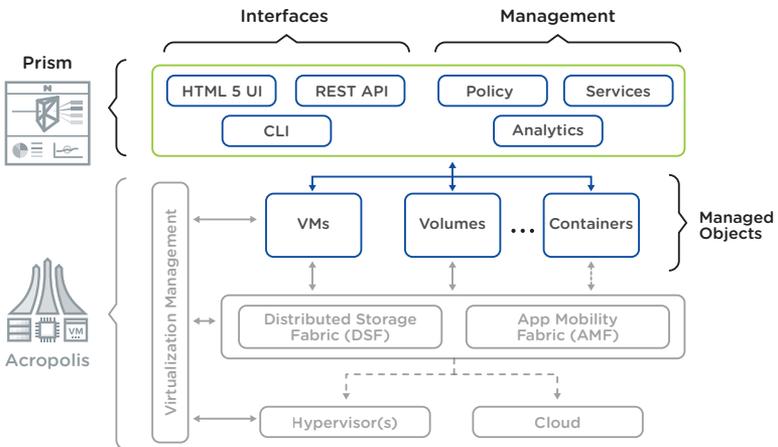


FIGURE 8: High-Level Prism Architecture

The Prism UI contains the following main pages that provide monitoring and control over all the capabilities described in this book.

Home — Local cluster monitoring dashboard includes detailed information on alerts, capacity, performance, health, tasks, etc.

Health — Provides environment, hardware, and managed object health and state information. Also includes NCC health check status.

VM — Full VM management, monitoring, and create, retrieve, update, and delete (CRUD) for Acropolis; VM monitoring (non-Acropolis)

Storage — Container monitoring and management including all create, retrieve, update, and delete actions

Hardware — Server, disk, and network management, monitoring, and health. Includes cluster expansion as well as node and disk removal.

Data Protection — Disaster Recovery, Cloud Connect, and Metro Availability configuration. Management of protection domain objects, snapshots, replication, and restore.

Analysis — Detailed performance analysis for cluster and managed objects with event correlation

Alerts — Local cluster and environment alerts

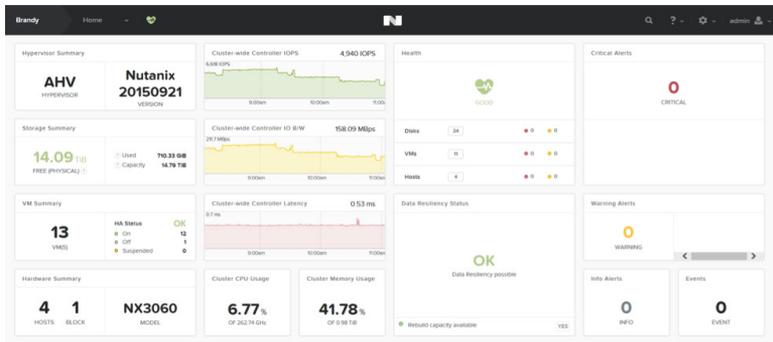


FIGURE 9. Prism Element — Dashboard

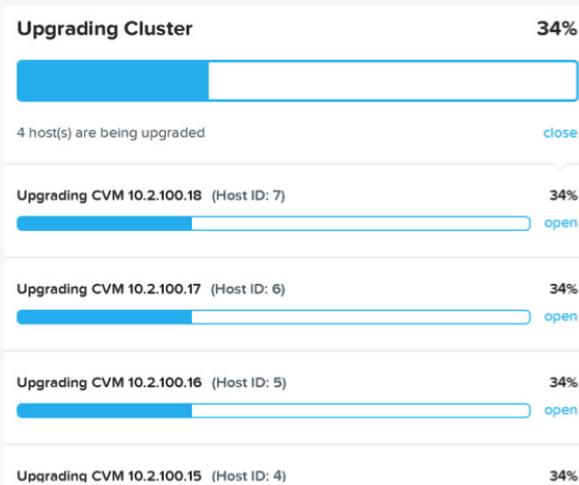
PRISM KEYBOARD SHORTCUTS

Accessibility and ease of use are of critical importance. Prism includes a set of shortcuts that allows users to do everything from the keyboard, simplifying and accelerating operations. You can use shortcuts to quickly change the view, see activities and events, and access menus.

PERFORMING SOFTWARE UPGRADES

Performing software upgrades is a simple and non-disruptive process and can be performed with a single click. This includes upgrading Nutanix software, Hypervisor and system firmware.

To accomplish an upgrade, all that is required is to select “Upgrade Software” from the Prism dashboard, and download the desired software version from the cloud. The new software installation is automatically orchestrated across all nodes.



That's it. Three steps, regardless of cluster size.

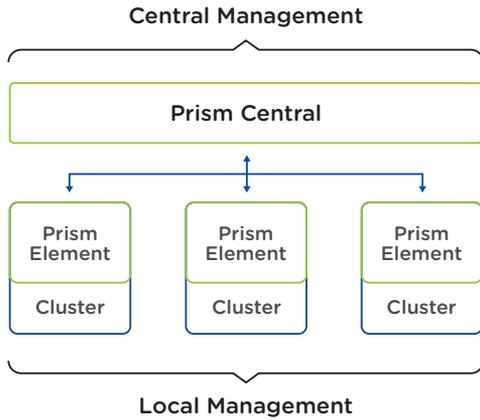


FIGURE 10. Prism Central

PRO TIP: PRISM CENTRAL

Prism Central is recommended for larger or distributed deployments (more than one Nutanix cluster or multiple sites) to simplify operations and provide a single management UI across all clusters and sites.

With the Prism Central dashboard, multiple clusters can be monitored and managed including consolidated alerts, available storage, performance in terms of both bandwidth and IOPS, and more.

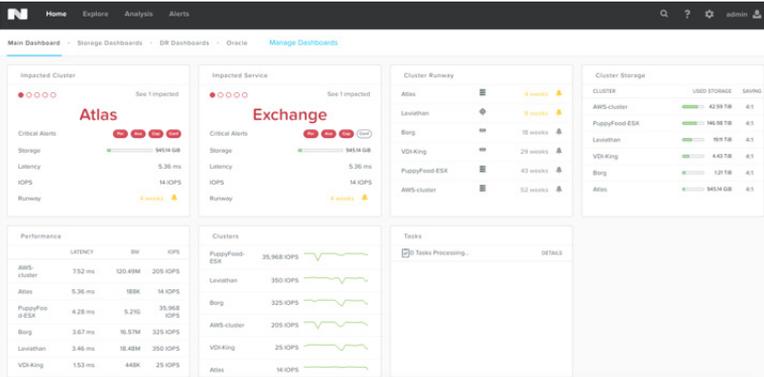


FIGURE 11. Prism Central – Dashboard

WHO WOULD YOU RATHER CALL?

Nutanix wields [support as a competitive advantage](#) with an industry-leading 90+ Net Promoter Score. Nutanix support covers the entire infrastructure stack—compute, storage, and virtualization.

Ready to learn more about hyperconverged infrastructure? Contact us at info@nutanix.com, follow us on Twitter @nutanix, or send us a request at www.nutanix.com/demo to set up your own customized briefing and demonstration to see how validated and certified solutions from Nutanix can help your organization make the most of its enterprise applications.

Stay engaged with Nutanix experts and customers on the Nutanix Next online community (next.nutanix.com).

ABOUT NUTANIX

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power their business. The Nutanix enterprise cloud platform leverages web-scale engineering and consumer-grade design to natively converge compute, virtualization and storage into a resilient, software-defined solution with rich machine intelligence. The result is predictable performance, cloud-like infrastructure consumption, robust security, and seamless application mobility for a broad range of enterprise applications.

