The Nutanix Design Guide

FIRST EDITION

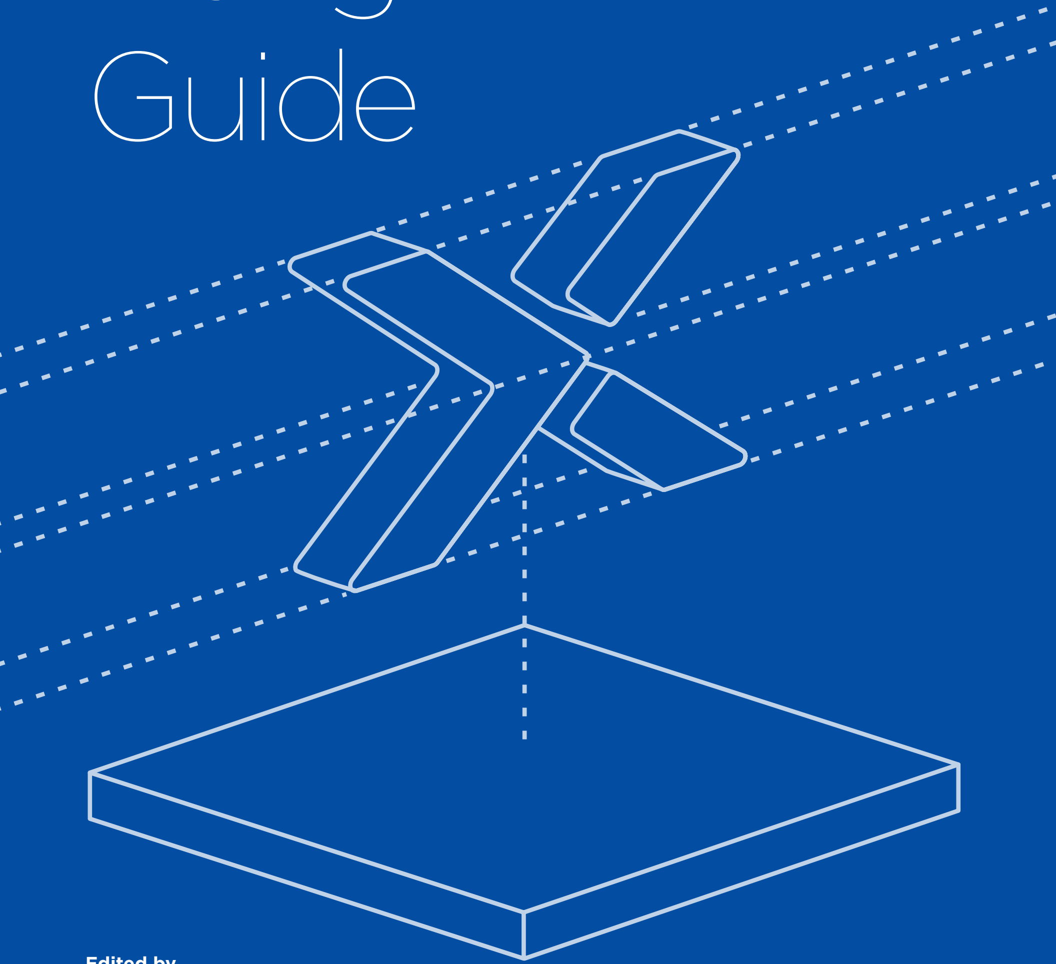

Edited by

Angelo Luciani, VCP

René van den Bedem,
NPX, VCDX4, DECM-EA

By Nutanix,
in collaboration with
RoundTower Technologies

Table of Contents

Table of Contents

1 Foreword

I am honored to write this foreword for ‘The Nutanix Design Guide’. We have always believed that cloud will be more than just a rented model for enterprises. Computing within the enterprise is nuanced, as it tries to balance the freedom and friction-free access of the public cloud with the security and control of the private cloud. The private cloud itself is spread between core data centers, remote and branch offices, and edge operations. The trifecta of the 3 laws – (a) Laws of the Land (data and application sovereignty), (b) Laws of Physics (data and machine gravity), and (c) Laws of Economics (owning vs. renting in long term) – is forcing the enterprise to be deliberate about its cloud journey. At Nutanix, we firmly believe that the private and the public cloud must mimic each other when it comes to ease-of-use and friction-free operations.

Crawl. Walk. Run. The Three Stages of the Enterprise Cloud Journey

Like most memorable and practical things in life, we break down the cloud journey for the enterprise into three achievable phases:

- **Modernize your Infrastructure**: with a hyperconverged architecture.
- **Build your Private Cloud:** with automation and an entirely software-defined infrastructure.
- **Simplify your Multi-Cloud:** with governance, application mobility, and location-agnostic services.

Modernize Your Infrastructure

The core foundation of any cloud is web-scale engineering and consumer-grade design. The enterprise needs to operate an infrastructure that allows for fractional consumption (“pay as you grow"), continuous consumption and innovation ("seamless upgrades"), and rapid time to market ("invisible operations"). We've made compute, storage, and virtualization invisible, and extremely bite-sized, allowing the enterprise to have a modern infrastructure to free them up to build a true private cloud.

Build Your Private Cloud

Once the enterprise has experienced hyperconvergence of compute, storage, and virtualization running on commodity servers, it is now ready to leverage deep automation and orchestration for a truly programmable infrastructure. With our app-centric approach to automation, and with our focus on software-defined security, networking, and file management, we enable the enterprise to own and operate an efficient and reliable private cloud.

Simplify Your Multi-cloud

The enterprise will surely have multiple clouds for multiple apps, just like it had multiple operating systems for multiple workloads in the past decades. Our aim here is to provide a set of cloud-agnostic SaaS and PaaS services that help our customers manage their multi-cloud operations with simple 1-click experiences for edge computing (IoT), cost governance and security, and management of desktops, databases, containers, and object storage.

I am glad that The Nutanix Design Guide is explaining the "Why" of Nutanix and providing a 40,000ft view of the ecosystem we have crafted. One of our core cultural principles is "Believe in Striving". We are a constantly learning, continuously improving, eternally evolving company with an immense respect for the law of small improvements.

Let us learn together.

Dheeraj Pandey, Co-Founder & CEO, Nutanix

2 Contributors

Angelo Luciani is the Nutanix Technology Champion Community Manager at Nutanix. He is VCP certified. He blogs at virtuwise.com (voted Top 50 vBlog 2018 at vsphere-land.com) and can be followed on Twitter at @AngeloLuciani.

René van den Bedem is the Master Architect and Strategist at **RoundTower Technologies**. He is also a Nutanix Platform Expert (NPX), a quadruple VMware Certified Design Expert (VCDX) and a Dell-EMC Certified Master Enterprise Architect (DECM-EA). René is a current Nutanix Technology Champion Elite (NTC Elite). In 2018, he was also presented the Nutanix NTC & Community award and the Nutanix Education & Certification award. He blogs at vcdx133.com (voted Top 10 vBlog 2018 at vsphere-land.com) and can be followed on Twitter @vcdx133.

RoundTower Technologies is a solutions provider that delivers innovative solutions and services in the areas of service management, data center infrastructure, hyperconverged platforms, cloud automation and orchestration, DevOps, data analytics, and cybersecurity. RoundTower (roundtower.com) is enabling its customers to drive positive business outcomes by becoming more agile, efficient, and secure using technology. RoundTower is a Nutanix Master Partner and the only Nutanix Partner in the Americas to have a Nutanix Platform Expert (NPX) on-staff. RoundTower can be followed on Twitter @roundtowertech.

The following people authored content for this book:

- **Magnus Andersson**, Senior Staff Solutions Architect, Nutanix. He is also a Nutanix Platform Expert (NPX) and a double VMware Certified Design Expert (VCDX).
- **Neil Ashworth**, Solutions Architect, Nutanix.
- **Kees Baggerman**, Technical Director, Nutanix.
- **Daemon Behr**, Solutions Architect, Scalar Decisions. He is also a Nutanix Technology Champion (NTC).
- **Chris Brown**, Technical Marketing Manager, Nutanix.
- **Mark Brunstad**, Director, Nutanix.
- **Wayne Conrad**, Consulting Architect, Nutanix. He is also a Nutanix Platform Expert (NPX).
- **Rohit Goyal**, Principal Product Marketing Manager, Nutanix.
- **Laura Jordana**, Technical Marketing Engineer, Nutanix.
- **Steve Kaplan**, Vice President, Customer Success Finance, Nutanix.
- **Gary Little**, Director, Technical Marketing Engineering - Core HCI & Performance, Nutanix.
- **Mark Nijmeijer**, Product Management Director, Nutanix.
- **Bas Raayman**, Staff Solutions Architect, Nutanix. He is also a Nutanix Platform Expert (NPX).
- **Michael Webster**, Technical Director, Nutanix. He is also a Nutanix Platform Expert (NPX) and a VMware Certified Design Expert (VCDX).
- **Greg White**, Solution Marketing Principal, Nutanix.

3 Acknowledgements

The following people reviewed and provided feedback for this book:

- **Kasim Hansia**, Staff Solutions Architect, Nutanix.
- **Michal Iluz**, Art Director, Nutanix.
- **Dwayne Lessner**, Principal Technical Marketing Engineer, Nutanix.
- **Jordan McMahon**, Senior Content Marketing Manager, Nutanix.
- **Alexander Thoma**, Senior Manager, Nutanix. He is also a VMware Certified Design Expert (VCDX).
- **Rosemarii van den Bedem**, Content Editor

4 Using This Book

This book has been written to provide a consolidated view of the Nutanix Enterprise Cloud eco-system. It primarily focuses on the "Why" of Nutanix. The references section found at the end of each chapter contains links to resources that explain the "What" and "How" of Nutanix.

Each chapter is designed to be read as a standalone artifact. And regardless of the reader's familiarity with Nutanix, there should be something for everyone.

Nutanix has the intent of renewing and updating this publication each year, to include the latest products and enhancements of the Nutanix Enterprise Cloud.

For more information on the supported third-party ecosystems, please refer to the appropriate vendor documentation.

Please note that some of the listed resources require a valid customer or partner login to my.nutanix.com.

5

Introduction to Nutanix

Author: Angelo Luciani

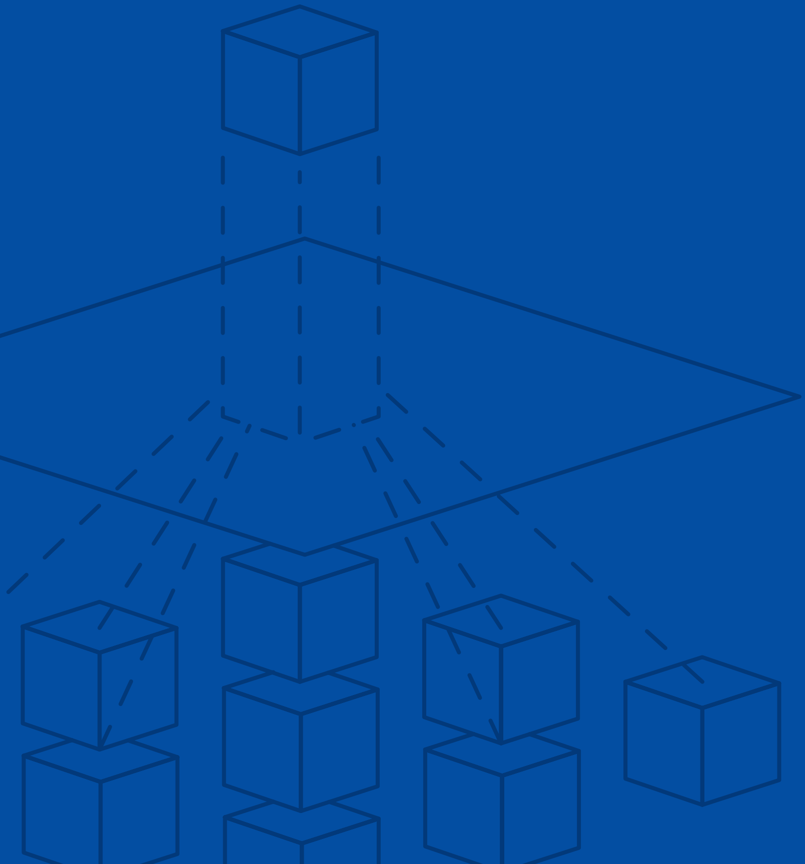

> "In twenty years' time, people will look back and shake their heads at the complexity of IT today. The future will be a utility model - connect to the cloud and consume a service. A bit like living in a city now and talking about running diesel generators and not using the power grid for your house."
>
> **– René van den Bedem, RoundTower Technologies**

The core of HCI is software-defined infrastructure, in that all data center devices need to move to pure software running on commodity x86 servers. Standardized hardware, a common operating system, consumer-grade design, and deep automation are the distinct virtues that make a true cloud (public or private). An open source movement around commoditizing private cloud – displacing large incumbents in virtualization (compute), storage, networking, and management – has largely fizzled out in the last decade, because it lacked integrity of execution.

Our company mission, since the beginning, has been to make infrastructure invisible. Naysayers scoffed at us when we were trying to make storage invisible. Most people believed that we had a niche market for the SMB, only to realize that the large enterprise had suddenly woken up to this simple yet powerful idea of software-defined infrastructure for almost everything. Converged infrastructure (CI) – a coalition solution of large compute-storage-networking incumbents, masqueraded as private cloud – is now considered a much smaller market than HCI. In fact, it would not be rhetoric to say that CI is dead as a segment.

In 2014, when we set out to build our own hypervisor (AHV), pundits gave us no chance in a saturated market of compute virtualization dominated by one or two large companies.

FIGURE 1

Nutanix Customer Journey

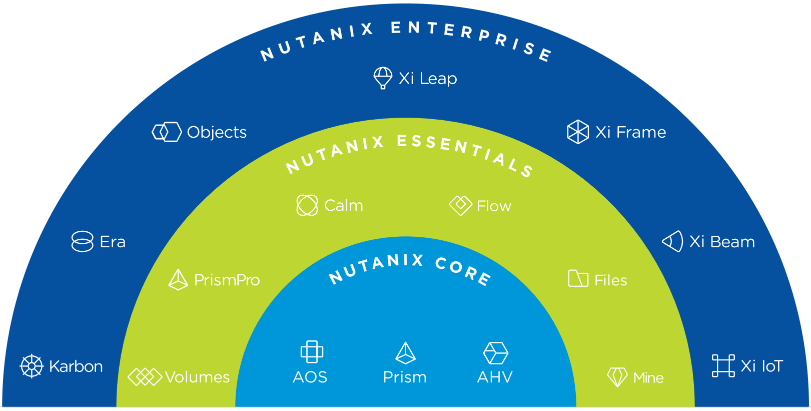

In the last 3 years, we have proved them wrong with the deep inroads that AHV has made with a large swathe of workloads in the enterprise. Industry watchers gave us no chance to shift from an appliance business model to a pure software business model as a public company.

Nutanix enables IT teams to build and operate powerful multi-cloud architectures. Our Enterprise Cloud OS software melds private, public and distributed cloud operating environments and provides a single point of control to manage IT infrastructure and applications at any scale.

Nutanix solutions are 100% software-based and are built on the industry's most popular hyperconverged infrastructure (HCI) technology, delivering a full infrastructure stack that integrates compute, virtualization, storage, networking and security to power any application, at any scale.

Nutanix software runs across different cloud environments to harmonize IT operations and bring frictionless mobility to all applications.

The Customer journey to the Nutanix Enterprise Cloud typically starts with Nutanix Core, before progressing to Nutanix Essentials and then Nutanix Enterprise. Nutanix is committed to making enterprise IT consumable as a utility service.

Software Options

5.1

Whether you choose Nutanix software as part of a turnkey appliance solution or to run on your own installed platform, the Acropolis and Prism editions provide a range of capabilities to match your needs.

The Acropolis Software Editions are:

- **Starter** – Core set of software functionality, ideal for small-scale deployments with a limited set of workloads.
- **Pro** – Rich data services, resilience and management features ideal for running multiple applications or large-scale single workload deployments.
- **Ultimate** – The full suite of Nutanix software capabilities to tackle complex infrastructure challenges ideal for multi-site deployments and advanced security requirements.

The Prism Software Editions are:

- **Starter** – A comprehensive systems management solution for the single and multi-site management of Nutanix clusters.
- **Pro** – VM operations & systems management with advanced machine intelligence, operations & automation capabilities.

Nutanix also supports capacity-based licensing to include cluster attributes such as the number of raw CPU cores and raw total of flash drive capacity.

Nutanix Calm is sold as an annual subscription licensed on a per virtual-machine (VM) basis. Calm licenses are required only for VMs managed by Calm, running in either the Nutanix Enterprise Cloud or public clouds. Nutanix Calm is sold in 25 VM subscription license packs.

Nutanix Era is a subscription term-based software license. This product is licensed based upon the concept of managed database server vCPUs. vCPU licensing is a consumption-based model that will allow customers to license just the database servers that will be managed by Nutanix Era. Licenses are sold in 1 to 5-year subscription terms.

Nutanix Flow is sold as an annual subscription licensed on a per node basis. Licenses are needed for all nodes in a cluster where micro-segmentation functionality will be used. This option requires a Nutanix cluster managed by Prism Central using the AHV virtualization solution. Licenses are sold in 1 to 5-year subscription terms. Prism Central with Starter license is needed to manage micro-segmentation policies.

References

5.2

What We Do:
https://www.nutanix.com/what-we-do/

What Is Hyperconverged Infrastructure?
https://www.nutanix.com/hyperconverged-infrastructure/

Hyperconverged Infrastructure: The Definitive Guide:
https://www.nutanix.com/go/what-is-nutanix-hyperconverged-infrastructure.html

Hardware Platforms:
https://www.nutanix.com/products/hardware-platforms/

Software Options:
https://www.nutanix.com/products/software-options/

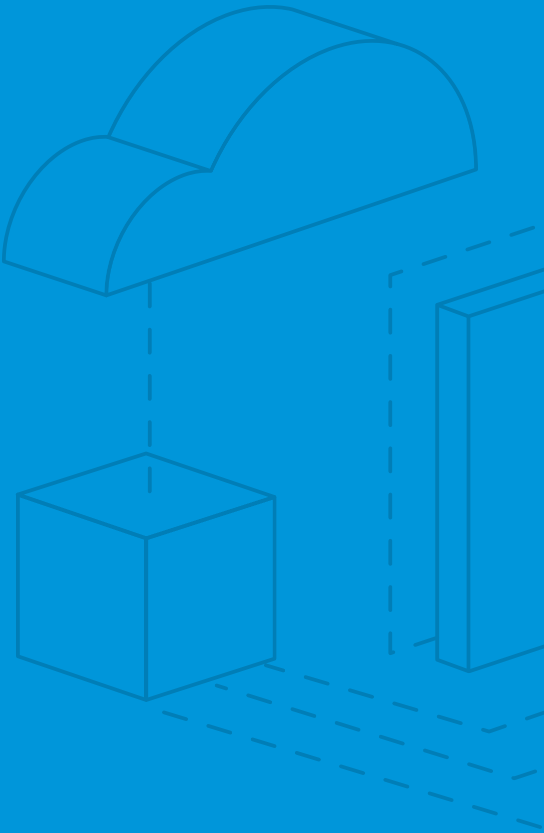

6

Why Nutanix?

Author: Steve Kaplan

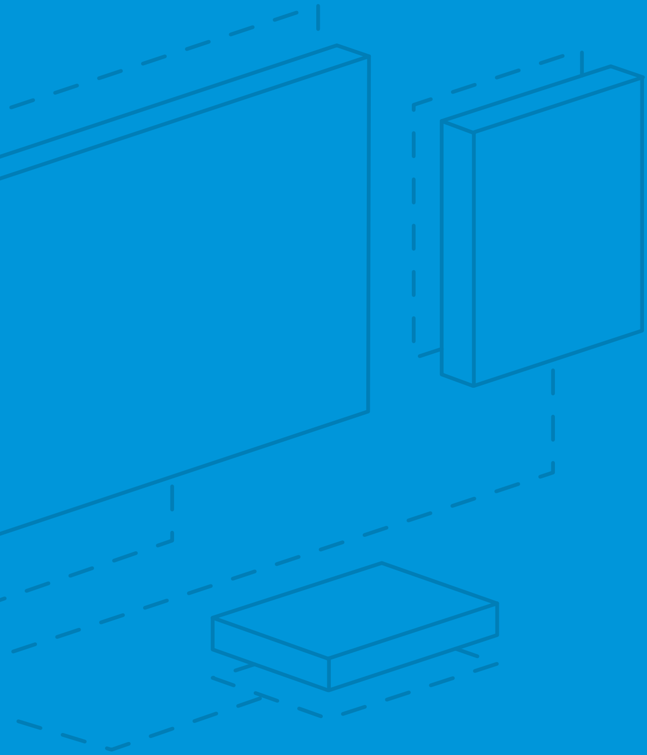

6.1 The Broken Legacy Data Center

When you think about it, "shadow IT" is a bizarre concept. You never hear, for example, about "shadow Human Resources" or "shadow Sales". Shadow accounting might occasionally be a thing, but then it is typically called "fraud". Yet "shadow IT" is nearly ubiquitous among larger organizations with legacy hardware-dependent IT infrastructures.

The preponderance of shadow IT testifies to the broken legacy data center model. The hardware defined nature of proprietary storage arrays limits most of the IT staff's time to infrastructure tasks and otherwise "keeping the lights on". This results in an inability for IT to rapidly deliver the innovative new services and offerings that the business demands for its customers, let alone leading the way to digital transformation. In frustration, the businesses take it upon themselves to fulfil a pressing need rather than wait for the "Department of Slow or No" to act.

Shadow IT is, of course, far from the only symptom of broken legacy IT.

The typical traditional data center is a hodgepodge of technology silos containing overlapping or redundant equipment that is both expensive and time-consuming to deploy and manage. In larger organizations, these silos are often accentuated with functional IT staff dedicated to domains such as servers, storage, network and virtualization. Specialists work independently and then collaborate out of necessity to cobble the individual results and initiatives together. In addition to processes simply taking more time, such an arrangement leaves the door open for human error, as things invariably get lost in translation.

This architecture of centralized shared storage, storage network and servers are known as “legacy 3-tier”. Not only is it complex, it does not scale well, is not natively resilient, and is expensive. Dedicated specialists configure LUNs, zone switches, manage RAID groups and rebalance hot spots; tasks that all disappear in the Nutanix software-defined architecture.

Legacy 3-tier infrastructure is expensive and complex in almost every way, making business agility very difficult to achieve. Some of the many drawbacks of legacy infrastructure include:

- Risk of overprovisioning because of large purchase increments.
- Multiple management systems and manual operations that impede flexibility and slow down deployments.
- Scaling limitations that result in outgrowing the solution too soon.
- Limited resiliency and other technical debt resulting from the lack of CapEx budget required to purchase multiple SANs.
- Multi-hop support and lack of end-to-end visibility that leads to operational firefighting.
- Complex, big data center footprint.
- Legacy IT organizations can take 40 days of approval processes to provision one VM.

Nutanix Software-Defined HCI Changes the Game

6.2

Andreesen Horowitz partner, Marc Andreesen, wrote a famous 2011 Wall Street Journal article titled, “Why Software is Eating the World”. Software is also eating the data center. The software-defined infrastructure terminology is not just marketing speak. As an analogy, think about what Apple did to phones, calculators, cameras, Rolodexes, the Sony Walkman and eReaders. The iPhone converged these individual technologies using a software-defined platform that changes the keyboard on the fly to match whatever functionality is accessed.

Nutanix built its Enterprise Cloud OS on top of software-defined Hyperconverged Infrastructure (HCI). Software-defined infrastructure, whether residing at an AWS, Azure or Google Cloud Platform data center, or on-premises in the form of an enterprise cloud, is necessary to provide "iPhone-like" consolidation benefits. And in the process, it reduces both cost and complexity. Instead of spending most of their time on infrastructure issues, IT staff can work more closely with the business. This allows them to leverage the software-defined infrastructure capabilities of speed and agility to achieve not just IT, but real business objectives.

Rather than utilizing proprietary storage arrays, the leading cloud providers instead provide storage as an application running upon millions of commodity servers. This is the same software-defined model employed by Nutanix. And like the public cloud, the Nutanix hyperconverged infrastructure approach slashes complexity and cost whilst dramatically enhancing agility and scalability.

6.3 The Financial Analysis Process

IT leaders across the globe struggle with making the optimal IT purchase decision, even when they intuitively know what it is. There may be a concern about lack of application vendor support, budget restrictions, security department resistance, internal politics, and so on. In present times, two primary pressures thwart rationalizing IT: A status quo bias and a public cloud bias.

How does a CIO rationalize, on one hand, a tendency to “hug” legacy infrastructure and on the other hand, pressure to “hug” the public cloud? A financial analysis, i.e. TCO or ROI depending upon the use-case, provides a framework for addressing both challenges. An analysis exposes the costs of all alternatives under consideration as well as help quantify business benefits that might otherwise fail to be considered. It helps the IT staff, as well as the other organizational stakeholders, make the optimal decision for the organization as well as securing the budget dollars required.

Whether TCO or ROI, it is essential to incorporate all variables that can affect the cost of any scenario being evaluated. This includes not just the up-front cost of hardware and software, but also the operational costs including rack space, power & cooling, administration, security, backup, disaster recovery, maintenance, support, networking, and so on.

One of the most important differentiating variables between legacy 3-tier infrastructure and HCI is growth. When SAN customers fill up an array or reach the limit of controller performance, they must upgrade to a larger model to facilitate additional expansion. Besides the cost of the new SAN, the upgrade itself is no easy feat. To try and avoid this expense and complexity, customers buy extra capacity and headroom up-front that may not be utilized for two to five years. This high initial investment cost hurts the project ROI. Moore’s Law then ensures the SAN technology becomes increasingly archaic (and therefore less cost effective) by the time it is utilized. And even in the best case, the upgrade cost is simply pushed off for 5 years.

Even buying lots of extra headroom up-front is no guarantee of avoiding a forklift upgrade. Faster growth than anticipated, new applications, new use-cases, acquisition of another company, etc. all can, and all too frequently do, lead to the under-purchasing of SAN capacity.

As shown in the figure below, the extra array capacity a SAN customer purchases up front starts depreciating on day one. By the time the capacity is fully utilized down the road, the customer has absorbed a lot of depreciation expense along with the extra rack space, power and cooling costs.

SANs lock customers into an old technology for several years. This has implications beyond just slower performance and less capabilities; it means on-going higher operating costs for rack space, power, cooling and administration.

Some of the newer arrays do an excellent job of simplifying administration, but even these arrays typically still require storage tasks related to LUNs, zoning, masking, Fiber Channel, multipathing, etc. And this does not include all the work administering and

FIGURE 2

Excess Capacity Depreciation Required for a SAN

Depreciation Cost per Consumed VM

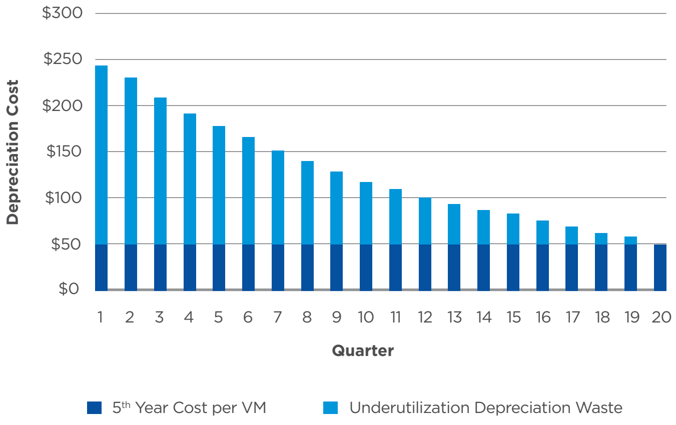

upgrading the server side, which can also be very significant. An August 2018 IDC study of eleven Nutanix customers who migrated from legacy infrastructure, reported a 61% decline in the cost to deploy, manage and support Nutanix HCI.

Business Outcomes 6.4

> "The purpose of IT is not to reduce the cost of IT"
>
> **–Steve Kaplan, Nutanix**

While a disruptive infrastructure solution such as Nutanix will reduce the status quo cost of IT, by far the more important outcome is typically going to be a change in the way the organization does business. Perhaps it will be able to utilize increased agility to boost sales, reduce customer turnover or get offerings to market more quickly. Indeed, it is the expectation of greater business agility that drives much of the public cloud adoption despite higher costs.

The analyst should clarify up-front whether business outcomes can be identified, quantified and considered as part of the analysis results. While the answer should be an unqualified, "yes", the reality is that many decision-makers are focused almost exclusively on reducing costs. While they might consider business outcomes in the case of a tie, in general they are far more focused on hard cost savings. In the following case study of a large healthcare organization, it was a desire to improve business outcomes that was the main driver for migration to the Nutanix Enterprise Cloud.

6.4.1 TCO Case Study of Nutanix vs. Legacy 3-Tier: Large Healthcare Institution

One of the largest healthcare institutions in the United States had over 900 Access Points but was behind on deploying new access points due to delivery and infrastructure complexity. The hospital needed better visibility into its VM environment. And in addition to the Nutanix technology solving the storage visibility issues, using Nutanix API automation and orchestration slashed provisioning times.

Nutanix analysts prepared a financial analysis for the organization utilizing numbers secured from the business for running an initial 1,400 VMs with anticipated yearly growth of 20%. The analysis reflected a capability of deploying new access points about 3 months faster, enabling doctors to see more patients. This boosted projected yearly revenues about $2.5M.

Infrastructure savings, as shown in the table below, were projected at an additional $10.8M over the 5-year analysis period. These results were very compelling for the hospital, which shortly thereafter became a Nutanix customer, and continues to expand out its Nutanix footprint.

TABLE 1

Projected 5-year TCO Savings for a Large Healthcare Organization

	3-Tier Legacy	Nutanix	Delta Legacy vs. Nutanix
Capital Expenses			
Compute Layer (Blades, Rackmount Servers) vs. Nutanix	$1,740,000	$6,240,000	-$4,500,000
Data Storage Services	$9,346,920	$0	$9,346,920
Storage Area Network Total Services	$343,392	$0	$343,392
SAN Ports & Cables	$40,768	$0	$40,768
Server Virtualization Software/Hypervision	$1,792,000	$1,064,000	$728,000
Capitalizaed Professional Services/ Installation	$676,516	$96,000	$580,516
Total Capital Expense	$13,939,595	$7,400,000	$6,539,595
Operating Expense			
Data Center Rack Space	$923,524	$86,857	$836,667
Power & Coding	$821,262	$173,665	$647,597
Post Warranty Support	$4,095,066	$6,249,579	-$2,154,513
Server Virtualization Software Support	$1,397,760	$618,240	$779,520
Administration FTE	$9,212,500	$5,005,000	$4,207,500
Operating Expense	$16,450,111	$12,133,341	$4,316,770
Total CapEx & OpEx	$30,389,707	$19,533,341	$10,856,366

Quantifying Virtualization Savings 6.5

Moore’s Law, which states that the number of transistors on a processor doubles every 18 months, has long powered the IT industry. Laptops, the Internet, virtualization, smart phones, cloud computing and hyperconverged infrastructure (HCI) are examples of technologies enabled by ever faster CPUs. There is no end in sight for the continued performance benefits of Moore’s Law, even though the ways in which that performance is achieved, such as using more cores, photonics and memristors, differs from the original precepts.

While VMware took advantage of increased CPU performance to launch ESX in 2001, the environment was of course much different than today. Network connectivity was at 100MB, Intel processors were running at 1.2Ghz – with only one core. And flash was not yet in use. As a result, VMware needed to add a lot of complexity to its virtualization environment, such as a separate vCenter Server management console in 2003.

Nutanix began selling its software 11 years later when virtualization was already the data center standard. Multi-cores were ubiquitous, connectivity was 10Gb Ethernet and flash was already becoming popular. As a result, Nutanix's system was natively clustered, and its software automates much of the complexity extant with legacy virtualization. Nutanix virtualization includes integrated management as part of every node that scales out with the environment, and which is also resilient using the same replication factor technology utilized by Nutanix for its operating system, as well as by the leading cloud providers.

Virtualization as a stand-alone product has had an incredible run, transforming data centers the world over into hosting environments for virtual machines. But virtualization, like deduplication and compression before it, has morphed from product to feature. Gartner has now even retired its Magic Quadrant for virtualization. And four of the leading public cloud providers, AWS, Google, Oracle and IBM, all use customized KVM variants for their hypervisors. This is not accidental. KVM has emerged as the optimal cloud hypervisor.

Nutanix also uses a customized KVM hypervisor, the Acropolis Hypervisor (AHV). Nutanix built AHV from the ground up to leverage the software intelligence of the hyperconverged architecture. AHV changes the core building block of the virtualized data center from hypervisor to application and liberates virtualization from the domain of specialists; making it simple and easily manageable by anyone from IT generalists to DevOps teams and DBAs.

License savings is only one of the drivers of customers moving to AHV. It is the integration of virtualization into the software-defined infrastructure and the resulting simplicity enabled that is truly compelling. The Nutanix Management platform, Prism, provides a single pane of glass for managing the entire infrastructure stack whether in a single data center or spread throughout data centers and offices globally. AHV deploys, clones and protects VMs holistically as part of the software-defined hyperconverged architecture, rather than utilizing disparate products and policies.

TCO Case Study: Nutanix AHV vs. VMware vSphere 6.5.1

In 2017, the U.S. government, Nutanix's largest worldwide customer, ran AHV on 74% of the Nutanix nodes it purchased. One agency still running vSphere on Nutanix, asked Nutanix analysts to provide a TCO comparison versus running AHV. This case study example is based upon that analysis.

The environment consists of 5,000 VMs running on 200 Nutanix nodes spread across multiple geographies and incorporate use-cases such as production, test, development, VDI, and so on. VMware vCenter Server is redundant using the virtual appliance, meaning that no copies of SQL Server or Oracle are required to enable the redundancy. Using Nutanix Move, it requires an average of 24 minutes to migrate each VM from vSphere to AHV at a fully burdened hourly rate of $50. Each 2-CPU version of vSphere costs $7,000, while each of ten vCenter Servers costs $6,000 per 2 CPUs. SnS (software and support) averages 20% per year.

Deployment and planning requires 8 hours of planning plus 8 hours per vCenter Server primary instance at a frequency of 1.5 times per year. We do not calculate the normally considerable vSphere upgrade time since Nutanix simplifies vSphere upgrades through the One-Click functionality.

The table below shows the projected 5-year TCO savings of $6,838,000 from switching from vSphere to AHV including the estimated $100,000 for migration to the Nutanix hypervisor.

TABLE 2

Sample vSphere Savings from Migrating to AHV

	VMware vSphere	Nutanix Acopolis (AVH)
Capital Expense Equation		
Virtualization software license costs	$1,400,000	$0
+ Virtualization management software license costs	$60,000	$0
= Total CapEx Costs	$1,460,000	$0
Operating Expenses Equation		
+ Virtualization software support costs	$1,460,000	$0
+ VM migration costs (if applicable) - Using Xtract	$0	$100,000
+ vSphere Upgrades ave 1 time per yr - Not included	$0	$0
Deployment Planning & Installation vCenter (ave 1.5 times/yr)	$18,000	$0
+ Security hardening ($4K per year per virtualization host)	$4,000,000	$0
= Total OpEx costs	$5,478,000	$100,000
CapEx & OpEx		
= Total CapEx & OpEx costs	$6,938,000	$100,000

6.5.1.1 Security

Each VMware product, and each version of said product, requires a separate hardening guide. The vSphere 6.7 Update 1 hardening guide alone includes 50 tasks, and these are not trivial tasks. The hardening guides additionally do not operate in isolation. Changes in hardening one product line can adversely affect another.

After administrators go through the weeks or months of applying hardening policies, they often need to be validated by an

Information Assurance (IA) team that then engages in an iterative process with the administrators. Once this process is completed and the entire system tests out, it needs to be documented in terms of the issues that came up, the resolution, mitigation, etc. All this work can equate to a great deal of time. It then must be closely monitored and manually mediated when "drift" occurs from upgrades in any of the individual products, or from administrator changes to any of the hardened configurations.

Nutanix AHV is hardened, tested via both Retina and Nessus vulnerability scanners, and validated out of the box. It eliminates all the VMware-required manual methodologies of going through each setting one by one. AHV eradicates all the hours spent applying VMware controls by hand and testing to see if it breaks. Administrators do not have to write documentation for IA since the automated STIG (Security Technology Implementation Guide) takes care of the documentation report. Administrators or IA can log in and run a STIG report or run Security Configuration Management and Automation (SCMA). The automated process runs on the cluster and self-heals the security baseline, eliminating the problem of drift.

Quantifying virtualization hardening savings with AHV varies greatly depending upon the organization and its security policies, and is reputably millions of dollars per year per VMware instance. A military branch has reduced man-hour costs by about $150K per year in managing the STIG for 36 Nutanix nodes running AHV. The $150,000 the military organization saves per 36 Nutanix nodes equates to a little over $4,000 per vSphere host (node) per year. Using this figure saves the organization in the figure above $4M over the 5-year analysis period.

Micro-segmentation 6.5.1.2

As Nutanix increasingly evolved its HCI technology to a comprehensive enterprise cloud platform, its engineers knew that VM-based security would be a requirement. For most use-cases, building an overlay network was not an efficient way to solve the problem. Nutanix Flow eliminates the need for overlays by implementing a distributed firewall built into the AHV kernel. Flow enables security, automation and network visualization without the massive complexity of building and managing a virtual network. Existing or upcoming API-integration with programmable switches including Arista, Mellanox, BigSwitch, Juniper and Cisco allows network automation in response to what the application needs.

Rather than managing a separate virtualized network with its own control and data plane, administrators can simply focus on the applications that business units want. No one cares what network hardware is utilized or what firewall is deployed, they just want application uptime, security and automation. Nutanix Flow removes networking knowledge as a requirement to getting things done. Nutanix Flow is enabled via a one-click workflow. Unlike VMware NSX Data Center, there is no upgrading of the environment to make it "Flow-ready." On day one, administrators can begin configuring policies.

6.6 The Public Cloud Alternative

In the staid world of IT, public cloud has gained momentum incredibly quickly as organizations have both the motive in the form of digital transformation, and now the means to escape the lack of agility and inefficiency of legacy data centers. But organizations often march to public cloud without fully understanding the financial implications.

An IDC study: Private vs. Public Cloud, for example, says that predictable workloads (which typically account for many applications) on average result in costs more than twice those when running on-premises with Nutanix HCI. A July 2018 IDC survey of 400 organizations, Cloud Repatriation Accelerates in a Multi-Cloud World, found that 80% of organizations in the study had repatriated at least some applications out of public cloud back on-premises, and that 50% of all public cloud applications installed today will move back on-premises, over the next two years.

If you are going to use a car a few weeks out of the year, it would be silly to purchase a vehicle as it would be far more expensive. If, however, you are going to use the car most of the time, it is far less expensive to purchase it rather than rent it year-round. The same type of logic applies to a public cloud. Elastic, burstable workloads make all kinds of economic sense to run in the public cloud. But customers can typically run predictable and persistent workloads at a much lower cost on-premises with Nutanix.

Migrating to public cloud requires expertise for security, redundancy, backup, specific tool sets and so on. The variable nature of public cloud charges means a lack of cost certainty and a risk of overspending. A May 2018 ZDNet article headlined, Cloud Computing Sticker Shock is Now a Monthly Occurrence; public company shareholders relying on CFO quarterly or annual reporting tend to dislike these potentially large unexpected costs.

Case Study Nutanix vs. Public Cloud: International Real Estate Company 6.6.1

The Chief Cloud Officer, Simon, of a large international real estate company is responsible for cloud services. Simon, after being exposed to Nutanix said to Tim McCallum, a Nutanix Business Value Analyst, *"Tim, Nutanix is interesting, and I get the whole thing about bringing cloud agility and simplicity on-premises, but there is not* *really any way that you can be less expensive than AWS. AWS is really cheap."*

Tim responded, *"Well, we do a lot of financial analyses and we find that for predictable workloads, AWS is generally about two to three times the cost of Nutanix Enterprise Cloud. I can help show you this using your own data if you are up for it."* Simon replied, *"Tim, I tell you what, I will send you an RVTools output that AWS used to size our next workload environment. You can price it as well, but do not get your hopes up."*

Tim responded later in the week with a TCO analysis. The summary slide is shown in the table below. *"Simon"*, he said, *"We have mapped all the AWS costs and, in fact, about 58 of them were*

TABLE 3

TCO Summary of Nutanix versus AWS

5 Year TCO Financial Summary

	Option 1: Amazon Web Services	**Option 2: Nutanix**
Capital Expenses		
Compute Layer	$0	$271,173
SAN Ports & Cables	$0	$480
Capitalized Professional Services/Installation	$0	$4,800
Sub-Total Capital Expense	**$0**	**$276,453**
Operating Expense		
Cloud Instances	$1,074,885	$0
EBS Storage	$163,665	$0
AWS Storage	$105,534	$0
Data Center Rack Space	$0	$6,857
Power & Cooling	$0	$8,191
Post Warranty Support	$0	$39,016
Administration LOE	$13,359	$31,250
Sub-Total Operating Costs	**$1,357,543**	**$85,314**
Total CapEx & OpEx	**$1,357,543**	**$361,767**

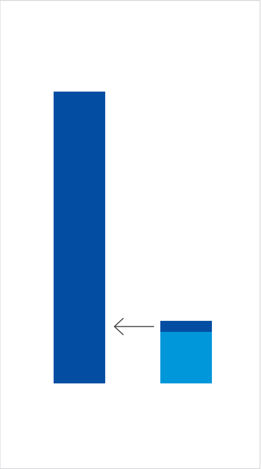

marked as not requiring more than 36% monthly activity. The results are a 72% reduction with Nutanix, less than one-third the cost of AWS." Six weeks later the real estate company was a Nutanix customer.

The real estate case study example is hardly uncommon, especially when it comes to evaluating public cloud. IT leaders, even those who have formal finance backgrounds, often get caught up in the excitement of cloud or in the comfort of legacy infrastructure, and make assumptions about costs without going through a financial rigor. This can easily lead to decisions that are less than optimal, particularly when dealing with disruptive infrastructure solutions such as HCI and cloud, whose cost models are dramatically different from traditional solutions.

Multi-Cloud First Strategy 6.7

While both public and Nutanix enterprise clouds provide the agility necessary to achieve digital transformation, the cost and complexity of the public cloud can make "cloud first" a very expensive strategy. In addition to the rental cost of putting workloads in the public cloud, the time to make the transition can take years. Meanwhile, the organization must still pay for its on-premises infrastructure and probably for most of its on-site IT staff. It must also hire new staff, or contract with consultants who have the expertise to implement the specialty backup, redundancy, and security required for the public cloud.

Depending upon application mix, it typically makes sense to have a multi-cloud first strategy rather than simply cloud first. An enterprise cloud embraces both the private cloud for the control and customization customers need, and the public cloud for cloud-native and elastic workloads. This hybrid approach provides the best of both worlds. IDC's Private vs. Public Cloud study referenced earlier in this chapter concludes that a multi-cloud environment is now, "...the norm for enterprise organizations."

This conclusion is supported by a mid-2018 survey conducted of 350 IT decision-makers by the Enterprise Strategy Group (ESG), Tipping Point: Striking the Hybrid Cloud Balance. The study showed that around half of respondents (49%) plan to run most of their applications/workloads in their own data centers, while another 43% plan to evenly split applications between their own data centers and public cloud.

Nutanix began publishing the Enterprise Cloud Index in late 2018 which consists of VansonBourne conducted research of 2,300 global IT decision-makers. Organizations spend 26% of their annual IT budget on public cloud, according to the survey results, with this percentage set to increase to 35% in two years' time. Only 6% of organizations that used public cloud services said they stayed under budget, while 35% overspent. The study showed that whereas today 36% of enterprise workloads are running in both private and public clouds, the number is expected to jump to 56% in 24 months. Most respondents (91%) identified hybrid cloud as the ideal IT model.

Any IT infrastructure decision, whether legacy, public cloud, or enterprise cloud, should be carefully evaluated within the context of the organization's long-term business objectives and application mix, and all the relevant variables should be quantified and compared. This is the best way to ensure an organization selects the optimal architecture for enabling success.

For further reading, see Steve Kaplan's book – The ROI Story: A Guide for IT leaders, available now from Amazon.

References 6.8

Wall Street Journal, Why Software is Eating the World:
https://www.wsj.com/articles/SB10001424053111903480904576512250915629460

IDC, Nutanix Pricing versus Traditional Infrastructure TCO ROI Report:
https://www.nutanix.com/go/nutanix-pricing-vs-traditional-infrastructure-tco-roi-report.html

IDC, Private versus Public Cloud:
https://www.nutanix.com/go/multicloud-architectures-empower-agile-business-strategies.html

IDC, Cloud Repatriation Accelerates in a Multi-Cloud World:
https://www.idc.com/getdoc.jsp?containerId=US44185818

ZDNet, Cloud Computing Sticker Shock is Now a Monthly Occurrence:
https://www.zdnet.com/article/cloud-computing-sticker-shock-is-now-a-monthly-occurrence-for-many-companies/

Enterprise Strategy Group (ESG), Tipping Point: Striking the Hybrid Cloud Balance:
https://www.esg-global.com/research/esg-master-survey-results-tipping-point-striking-the-hybrid-cloud-balance

Nutanix & VansonBourne, Enterprise Cloud Index:
https://www.nutanix.com/enterprise-cloud-index

7

The Nutanix Eco-System

Author: René van den Bedem

The figure below provides a 40,000-foot view of the Nutanix Enterprise Cloud eco-system.

FIGURE 3

Nutanix Enterprise Cloud Eco-System

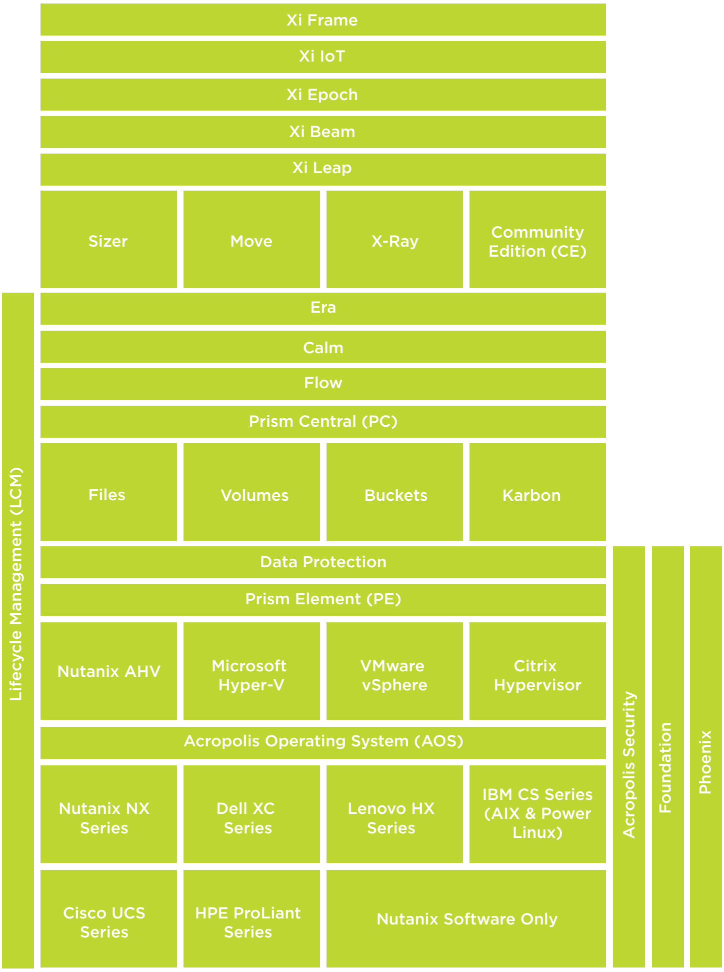

Since Nutanix exited stealth start-up mode in 2011, the Nutanix offering has come a long way. Nutanix has transitioned from the Virtual Computing Platform (VCP), to the Extreme Computing Platform (XCP) and now to the current Enterprise Cloud offering in 2019.

Each component has the following function of the Nutanix Enterprise Cloud OS:

- **Xi Frame**: Run full Desktops and Applications in your browser.
- **Xi IoT**: Edge Computing for Internet connected sensors.
- **Xi Epoch**: Observability and monitoring for Multi-Cloud Applications.
- **Xi Beam**: Multi-Cloud Optimization to reduce cost & enhance Cloud Security.
- **Xi Leap**: Disaster Recovery Service to protect applications running on Nutanix.
- **Sizer**: Tool to create design scenarios, size workloads and download the proposal & Bill of Materials.
- **Move**: VM migration tool.
- **X-Ray**: Automated test tool used for proof of concepts and benchmarking other technologies against Nutanix.
- **Community Edition**: Free version of Nutanix software for the community.
- **Era**: Database Lifecycle Management PaaS.
- **Calm**: Application Lifecycle Management and Cloud Orchestration.
- **Flow**: Advanced Networking and Application Centric Network Security.
- **Prism Central**: Manager of clusters, advanced management and planning, also hosts Calm and Flow.

- **Files**: Filer services.
- **Volumes**: iSCSI Block storage services.
- **Objects**: Object storage services.
- **Karbon**: Kubernetes container storage services.
- **Data Protection**: Backup, Recovery and DR Orchestration.
- **Lifecycle Management (LCM)**: Software and firmware lifecycle management.
- **Prism Element**: One-click infrastructure operations.
- **AHV**: Cloud optimized hypervisor based upon KVM.
- **Acropolis Security**: SCMA and STIG based security lifecycle management (SecDL).
- **Acropolis Operating System**: Software-defined storage layer.
- **Nutanix NX Appliance**: Hardware appliance built on Supermicro.
- **Nutanix Software Only**: Option for using commodity hardware from the Nutanix HCL.
- **Foundation**: Nutanix software imaging service used by partners and Nutanix to bring Nutanix nodes into service.
- **Phoenix**: Bare-metal recovery and imaging service.

Please refer to each chapter for a more detailed breakdown on the customer use-cases for each technology stack, including the advantages and disadvantages of those features.

References

7.1

Nutanix Product Overview:
https://www.nutanix.com/products/

Nutanix Core:
https://www.nutanix.com/products/core/

Nutanix Essentials:
https://www.nutanix.com/products/essentials/

Nutanix Enterprise:
https://www.nutanix.com/products/enterprise/

Nutanix Community Edition:
https://www.nutanix.com/products/community-edition/

Nutanix Test Drive in the Cloud:
https://www.nutanix.com/test-drive-hyperconverged-infrastructure/

8

Certification & Training

Author: René van den Bedem

Nutanix has four certification and learning tracks, with the Nutanix Platform Expert being the premier level of certification (refer to following chapter for additional detail). The tracks are Sales, Systems Engineer, Services and Technical.

FIGURE 4

Nutanix Certification by Role

Sales Role	Systems Engineer Role	Services Role	Technical Role
NCSR Level 1	NCSE Level 1	CCIC	NCP
NCSR Level 2	NCSE Level 2	NCPI	NCAP
NCSR Level 3	NPX	NCS	NPX
NCSX			

Most of the learning content is available online via the Nutanix NuSchool platform. Nutanix also has a series of classes and bootcamps delivered by Nutanix and learning partners, delivered on-site.

Most of these tracks are only available to Nutanix partners, however the Technical track (NCP, NCAP, NPX) is available to customers also.

The NPX and NCSX have an in-person panel defense component that must be met before the certification can be awarded.

The NCP, NCAP, NCPI, NCS and NCSE exams are proctored certifications.

The Nutanix Certification acronyms are:

- **NCSR** – Nutanix Certified Sales Representative
- **NCSX** – Nutanix Certified Sales Expert
- **NCSE** – Nutanix Certified Systems Engineer
- **CCIC** – Nutanix Core Competency – Install and Configure
- **NCPI** – Nutanix Consulting Partner Installation
- **NCS** – Nutanix Consulting Specialist
- **NCP** – Nutanix Certified Professional
- **NCAP** – Nutanix Certified Advanced Professional
- **NPX** – Nutanix Platform Expert

The Nutanix Partner Program has specific requirements for Nutanix Certification when achieving the Pioneer, Scaler and Master Partner levels. Refer to the Nutanix Channel Charter chapter for additional information.

8.1 References

Nutanix Virtual Technology Bootcamp:
https://www.nutanix.com/bootcamp/virtual/

Nutanix NuSchool:
https://nuschool.nutanix.com

Nutanix Partner Network:
https://www.nutanix.com/partners/

Nutanix Partner Network – Learn by Role:
https://nutanix.portal.relayware.com/?eid=2155

Nutanix Partner Network – Search for a Class:
https://nutanix.portal.relayware.com/?eid=2175

The Nutanix Bible:
https://nutanixbible.com

9

Design Methodology & The NPX Program

Authors: René van den Bedem & Mark Brunstad

"Complex is competent.
Simple is Genius."

– Binny Gill, Nutanix

The Nutanix Design Methodology is about simplicity. Nutanix products are not complicated and never will be. However, Nutanix solutions do need to be integrated into the enterprise data center, which is usually intricate and convoluted.

For Nutanix solutions to be successful, meeting the business requirements of the customer with the minimum of risk, the Nutanix Platform Expert (NPX) program was developed.

The NPX program is a peer-vetted, hypervisor agnostic certification designed for veteran Solution Engineers, Consultants, and Architects. In accordance with the program goals, every NPX will be a superb technologist, a visionary evangelist for Web-scale and a true Enterprise Architect, capable of designing and delivering a wide range of cutting-edge solutions; all custom built to support the business goals of the Global 2000 and government agencies in every region of the world.

As a customer, you should be working with Nutanix Master partners that have NPX certified individuals on-staff to oversee your Nutanix solution projects, ensuring business success with minimal risk in a timely fashion.

References

9.1

Nutanix Platform Expert (NPX): Why We Built It and Why It Matters:
https://www.nutanix.com/2015/03/20/nutanix-platform-expert-npx-why-we-built-it-and-why-it-matters/

Nutanix Platform Expert (NPX) Certification and Directory:
https://www.nutanix.com/npx-certification/

NPX Link-O-Rama:
https://vcdx133.com/2015/03/06/nutanix-platform-link-o-rama/

Nutanix NPX Community Forum:
https://next.nutanix.com/nutanix-platform-expert-npx-37

The ROI of Nutanix Platform Expert (NPX) certification:
http://bythebell.com/2016/04/the-roi-of-nutanix-platform-expert-npx-certification.html

NPX Design Review Preparation Guide:
https://www.nutanix.com/go/nutanix-platform-expert-npx.php

Silicon Angle Interview .NEXT 2016:
https://siliconangle.com/2016/06/21/top-1-of-it-architects-wanted-top-tier-certification-with-npx-nextconf/

10

Channel Charter

Author: René van den Bedem

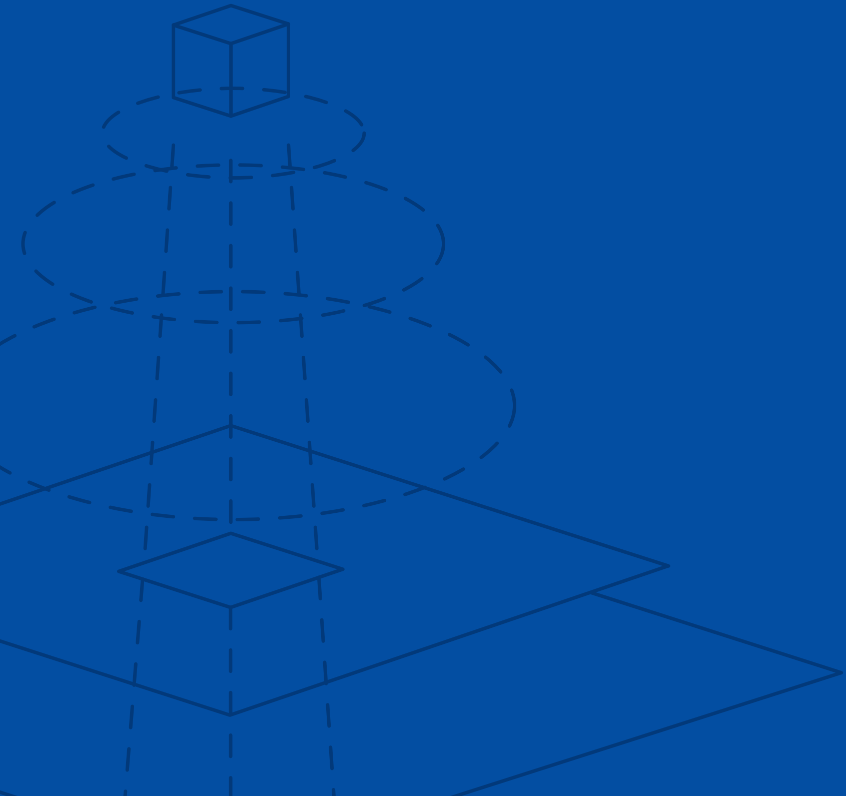

There are levels to this game, and if you have a strategic project that you want done right, it makes sense to align yourself with the correct Nutanix partner.

The Nutanix Channel Charter has three Partner levels:

- **Pioneer** – Fundamental sales and technical proficiencies in Nutanix core products.
- **Scaler** – Develop integrated solutions around the Nutanix Enterprise Cloud OS ecosystem.
- **Master** – Sell the full Nutanix portfolio consistently with an established service practice that has advanced sales and technical staff that are highly qualified.

The table below lists the criteria for each Partner Level in developed Zone 1 countries.

TABLE 4

Nutanix Channel Charter Requirements (Zone 1)

Requirement	Pioneer	Scaler	Master
Closed Deals	2	9	30
Transformational Deals	0	1	6
NCSR L1-L3	2	4	5
NCSX	0	1	2
NCP	1	2	4
NCSE L1-L2	1	2	4
NPX	0	0	Optional
NCPI or NCS	0	1	2

Refer to the Certification & Training chapter to understand the certification abbreviations.

Some of these listed requirements are not currently being enforced but will be in the future. Check the Nutanix Partner Portal for the latest requirements matrix.

Deals that include Objects, Calm, Era, Files, Flow, Xi Beam, Xi Epoch, Xi Leap and Xi Frame are counted as Transformational deals.

The table above lists the requirements for developed Zone 1 countries. Countries in Zones 2 and 3 have a lower number of requirements to be met.

References 10.1

Nutanix Partner Program:
https://www.nutanix.com/partners/

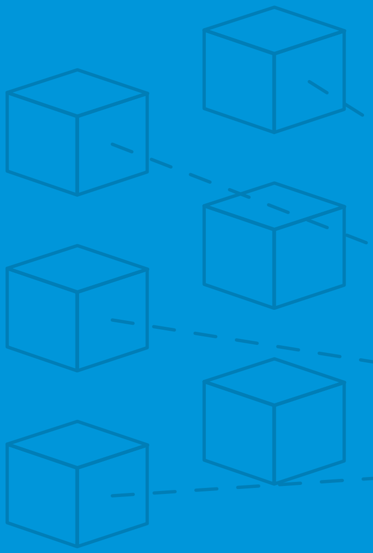

11

Mission-Critical Applications

Author: Michael Webster

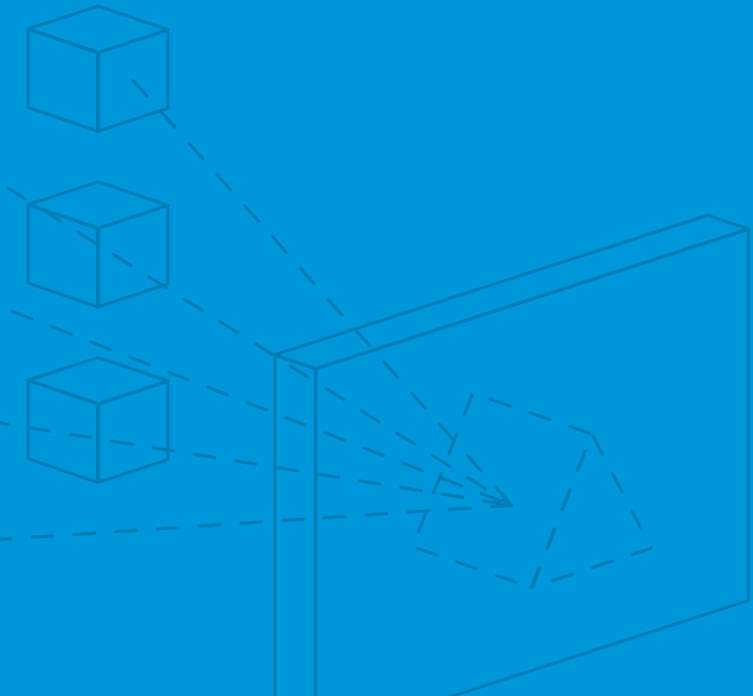

Mission-Critical or Business-Critical Applications are any applications that could have a material impact on the reputation, productivity or financial viability of an organization, if it were to become unavailable or experience severe performance degradation for an extended period. Performance degradation could be as severe as a complete outage and can cause the need to activate business continuity or disaster recovery plans.

The cost of unavailability can often be measured in millions of dollars per hour, and therefore reducing risk and risk management are of critical importance. Careful planning, processes, testing, monitoring, operations, and design are therefore required to ensure the required experience on any platform, especially when migrating to a cloud-like platform such as Nutanix.

Since going public (Nasdaq: NTNX) in 2016, Nutanix has always included the use-case distribution information in the infographic available with every earnings release. The proportion of Mission-Critical and Business-Critical apps on Nutanix is approximately 50% and has been consistent since the reporting began.

The main driving factor for this is how the Nutanix architecture reduces risk, improves predictability and performance consistency from day 1, during growth, and when disasters strike. The figure below displays the workload use-case proportion from Q1 FY2019, which was included in the earnings infographic (see references for full infographic).

FIGURE 5

Nutanix Enterprise Cloud Use-Case Distribution Q1 FY2019

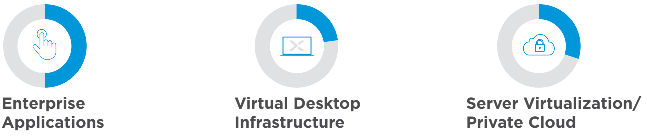

The Mission-Critical and Business-Critical Apps are categorized into the following types:

ERP systems, such as SAP, and supporting databases and middleware – the beating heart of most large organizations, are usually interconnected to every other system within the enterprise

- Pros: Reduced downtime risk and no single point of failure, reduced complexity due to less components, predictable performance and scalability, more accurate non-prod environments lead to lower risk of defects being found in production.
- Cons: Some adjustments to OS and App configurations may be required to get the best possible performance.

Process Automation and Control Systems – SCADA, across many industry sectors including Utilities, Oil and Gas, Manufacturing

- Pros: Small initial deployment size requirements mean small isolated environments are easy to deploy and manage and can increase availability while reducing overall risk, start small and grow, and easily support multiple fully isolated systems.
- Cons: Dark sites and completely isolated networks are harder to support and update as upgrade bundles need to be transported manually after being validated for authenticity, proactive support systems require Internet access.

Financial systems, payment processing, online banking

- Pros: Low latency and high throughput with high availability, flexible availability domains allow for large scale and increased failure tolerance as the environment grows, very low overheads compared to bare metal.
- Cons: Lack of pass-through network device support on some hypervisors to guest VM’s.

Middleware, ESB, Messaging systems, the translation and communication channels between different applications and organizations

- Pros: Low network latency, easy to scale app instances and infrastructure, easy configuration of network micro segmentation for improved security, data locality for persistent low latency message storage.
- Cons: Additional configuration required when using external physical load balancers.

Billing systems, there is no cashflow without billing and invoicing

- Pros: Easy to scale up on demand for cyclical or periodic application peaks and scale down again afterwards, making efficient use of the infrastructure deployed, data locality provides low latency and high throughput storage to reduce billing cycle times.
- Cons: Cyclical peaks must be included in capacity for infrastructure from the beginning and monitored as the environment grows.

Customer-facing online systems

- Pros: Start small with predictable scale as growth requires, extremely agile cloud like infrastructure that allows both traditional applications and cloud native applications to coexist side by side and be deployed on demand, small initial deployments make it cost effective to have completely isolated storage and networking for DMZ and other secure zones, making compliance with standards such as PCI DSS easier.
- Cons: Hosting multiple applications with different compliance requirements on the same consolidated infrastructure may mean additional audit logging is required to prove isolation and compliance with the policies.

Virtual Desktop Environment – If it supports all users

- Pros: Predictable linear scale out performance, High user density per node, Predictable failure characteristics, Flexible deployment options including cloud and on site, Reduced deployment time.
- Cons: Not all hypervisors or brokers support the same features or end user devices, which requires careful design and sizing considerations and selection of components.

Use-Cases

11.1

The following use-cases drive design for Mission-Critical and Business-Critical Applications:

- Reducing risk of downtime, performance degradation.
- Improve predictability of performance and scalability.
- Ensure consistency, both during operations and during failure and maintenance events.
- Increased business agility for traditional apps and faster time to market without having to redevelop everything for a cloud native environment.
- Disaster recovery, to be provided by the infrastructure, or the application layer, or a combination. When designing a metro cluster environment, you still need to have backup and DR, since metro is disaster avoidance.
- Removing limitations of non-production environments, such as dev & test, allowing for exact copy of production quickly, without traditional limits, mean more valid testing and lower rate of defects found in production.
- Improved default security with automated configuration drift management for the infrastructure, secured by default, compliant by default, continuously monitored and remediated.

- Significantly lower total cost of ownership compared to traditional infrastructure solutions with more predictable and smaller growth increments as environments scale.

11.2 Design Considerations

These are some of the design considerations for Mission-Critical and Business-Critical Applications with Nutanix solutions:

- Migrating from traditional Mainframe and Unix systems to x86, converting from big endian to little endian and reducing migration downtime required, while allowing roll back where possible.
- Nutanix, Application Vendor, and Hypervisor guidelines and recommendations when virtualizing mission critical and business critical systems. The guidelines are based on significant testing, validation and real-world environment experience and are the baseline that all systems should meet to reduce risk and provide the best possible performance.
- Single threaded performance or single compute unit (SCU) performance. Some applications benefit greatly by having higher clock speed and therefore higher single threaded execution performance, whereas others are better with many threads, even if they are lower speed per thread. Care should be taken to provide the right clock speed, especially for older and single threaded applications and processes.

- Application and Database licensing. Application and database licensing have an extremely high impact on infrastructure design, failure domains, recovery and availability design. Any application or database that is licensed per processor should have an infrastructure optimized for high clock speed and low core count, to keep the per processor licenses as low as possible.
- Scale up versus scale out. Does the application only scale up, or can you scale it out and add multiple components to balance the load across a data center or multiple data centers? In a virtualized environment having more smaller VMs can achieve higher performance and better load distribution than fewer larger VMs. In many cases better than bare metal performance can be achieved with an optimized VM design due to more efficient processor scheduling.
- Availability and Disaster Recovery at the infrastructure as well as at the database & application level. For extreme high availability requirements additional vendor components may be required, especially if non-disruptive, cross data center automated failover is required. The more automation that is required, the more solution testing and training that will be required for operations staff.
- Operating system limitations do not change when you virtualize an application. Queue depth and OS schedulers remain limiting factors. Applications do not know they are virtualized. With the right design you can make the best of the infrastructure and the application within known limits.

FIGURE 6

App Architecture for Mission-Critical System with Automated Failover

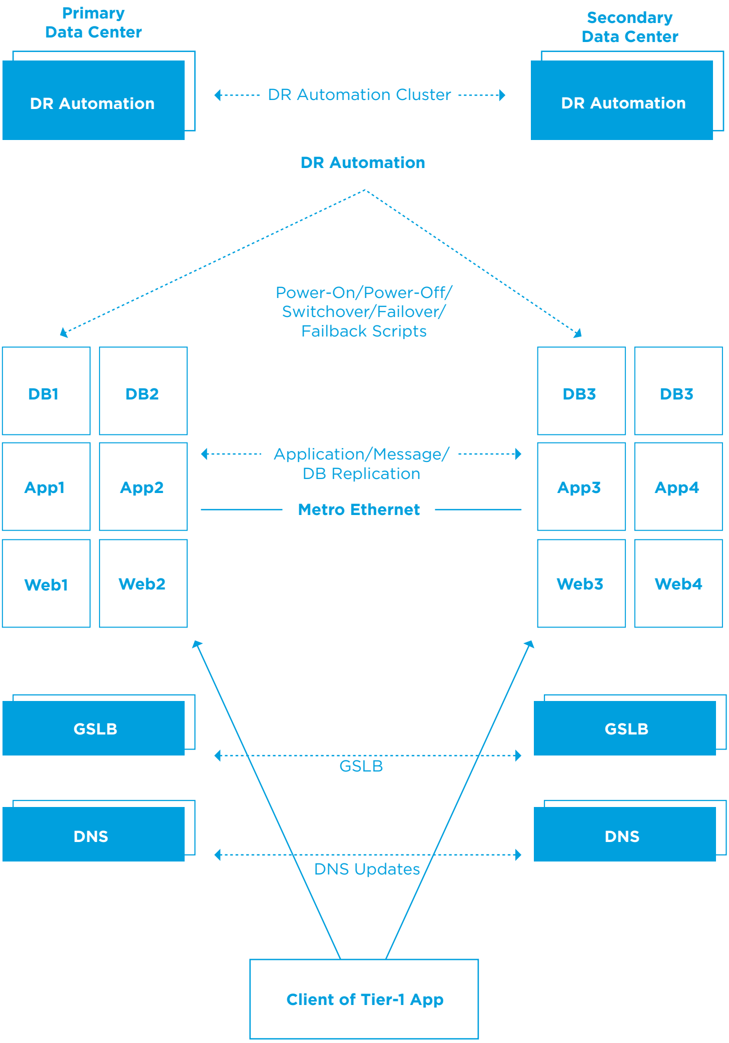

Risks

11.3

These are some of the risks associated with Mission-Critical and Business-Critical Applications:

- Lack of planning and validation leading to business requirements not being met. If you fail to plan, you plan to fail. This happens far more often than it should for critical apps. Diligence, care and attention to detail in validating all business requirements are essential to a high-quality project delivering the desired business outcomes.
- Virtualizing Mission-Critical and Business-Critical apps for production is not like virtualizing dev and test, or less critical apps. Aggressive over commitment of resources is likely to lead to project failure. Resources should be guaranteed for critical applications, so you know for sure the business requirements will be met.
- Non-production environments (development/test) for mission critical applications can be almost as critical as production and require careful consideration and planning, especially when there is a high cost to productivity loss, or when development and testing is a major part of your business, especially when outside business partners also integrate with these systems and you have contracts around the SLAs.
- If you do not have objective validated baseline performance information and business metrics before a migration project, you will not be able to determine if it meets your performance requirements, and it will make troubleshooting afterwards incredibly difficult.

11.4 References

SAP NetWeaver Certified:
http://scn.sap.com/docs/DOC-8760

SAP HANA Certified:
https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/hci.html

Microsoft SVVP Certified:
https://www.windowsservercatalog.com/svvp.aspx

Exchange ESRP:
https://technet.microsoft.com/en-us/office/dn756396.aspx

Enterprise Applications with Nutanix:
http://www.nutanix.com/solutions/enterprise-applications/

1 Million IOPS in 1 VM on Nutanix:
http://longwhiteclouds.com/2017/11/14/1-million-iops-in-1-vm-world-first-for-hci-with-nutanix

Nutanix – Oracle Platinum Partner and Exastack Ready Solution:
https://solutions.oracle.com/scwar/scr/Solution/SCSP-MJCIZLCF.html

Best Practices Guide: Oracle on AHV:
https://www.nutanix.com/go/optimizing-oracle-on-ahv.html

12

SAP on Nutanix

Author: Bas Raayman

As explained in the previous chapter, a mission-critical or business-critical application is any application that could have a material impact on an organization's reputation, productivity or financial viability in case of unavailability, or severe performance degradation, for an extended period.

A prime example of a software vendor whose software is generally considered to be among the most critical workloads inside of a company is SAP. Founded in 1972 by five former IBM employees, the various solutions offered by SAP touched over 77% of the world's transaction revenues.

SAP offers a variety of products used by over 425,000 customers worldwide, including some well-known ones such as:

- **Enterprise Resource Planning** - ERP/ECC
- **Business Warehouse** - BI/BW
- **Customer Relationship Management** - CRM
- **Supply Chain Management** - SCM/APO
- Various databases such as Oracle, Microsoft SQL Server, Sybase ASE, and notably the in-memory SAP HANA database

A typical SAP setup consists of several distinct layers. There is a hardware layer which is designed to be highly available, typically sized for peak business workloads such as year-end closing, the number of users working on a system simultaneously, and the agreed upon service levels. Keep in mind that the existing hardware layer is not necessarily running on an Intel x86 system, this could very well be something like an IBM Power or IBM Z system, HPE Superdome or Oracle SPARC systems. On top of the hardware, we tend to either see a bare-metal operating system installation or a virtualization layer with a guest operating system. Running inside the operating system is usually one of three different layers:

- The database layer, which is responsible for the reading and writing of data.
- The application layer, which processes data using application logic.
- The presentation layer, which presents the processed data to the user.

Since the availability of these SAP systems and their entire landscape is vital, and we want to avoid any changes to the underlying business logic to impact a running system, we frequently see a generic SDLC model followed, a so-called "three-system landscape" in SAP terminology, in which multiple clients are set up:

- **DEV** - A development system
- **QAS/TEST** - A quality assurance system
- **PROD** - A production system

Other systems such as testing, training, prototyping, and more, depending upon the testing methodologies employed, can be added as desired. Any changes or customizations to objects are released in a change request and then transported to the next client, for example from the DEV to the QAS client to then have the changes tested by some key users. Only once transportation of the changes and customizations is completed are the modifications visible in the target system. Terminologies in SAP such as system, instance, and client can sometimes be interchangeable colloquially and must be set in the right context with each customer to understand the impact of sizing and overall solution.

SAP was first certified to run on Nutanix in 2015. In 2016, Nutanix announced their certification of the Nutanix AHV hypervisor for SAP Business Suite powered by SAP NetWeaver, and in 2018, was the first hypervisor certified for production SAP HANA on Hyperconverged Infrastructure (HCI).

Implementing SAP on Nutanix results in the following benefits:

- Lowering TCO for SAP deployments.
- Reduction of complexity for HA and DR.
- Efficient and fast clones.
- Improved day-2 operations.
- Reducing guesswork and administrative overhead.
- Predictable growth & performance.
- Dramatic reduction in rack space.
- Focused root cause analysis.
- Quicker time to value.

12.1 Use-Cases

The following use-cases drive SAP design:

- Reducing the risk of downtime and performance degradation.
- Improve the predictability of performance and scalability.
- Ensure consistency, both during operations and during failure and maintenance events.
- Increased business agility and faster time to market.
- Disaster recovery provided by multiple layers such as the infrastructure, and the application layer. When designing a clustered environment, do not just protect against physical errors and outages, also protect against logical errors.
- Removing limitations of non-production environments, such as dev/test, allowing for exact copies of production environments quickly and without traditional limits. Also, at points in time that are considered compliance relevant such as year-end closing.
- Improved default security posture with automated configuration

drift management for the infrastructure, secured by default, continuously monitored and remediated.

- The significant lower total cost of ownership compared to traditional infrastructure solutions with more predictable and smaller growth increments as environments scale.

Design Considerations 12.2

These are some of the design considerations for SAP on the Nutanix platform:

- Your system could be migrating from traditional mainframe or Unix systems to x86, converting from Big Endian to Little Endian, non-Unicode to Unicode, from bare-metal to virtualized or any of the above combinations. Aim to reduce migration downtime required, while allowing rollback where possible. Also, ensure a proper test plan which includes performance and regression testing.
- If you are performing a heterogeneous system migration or copy, involve your SAP partner early since these migrations tend to be very complex and lengthy. Analyze the business impact of downtimes and phase freezes.
- Try to understand the best practices of the Nutanix Enterprise Cloud platform, Hypervisor, SAP application, and database guidelines and recommendations when virtualizing the SAP system. The guidelines are based upon significant testing, validation, and real-world environment experience and are the configuration baseline for all systems to reduce risk and provide the best possible performance. Each customer environment and requirements are different; hence it is always a good idea to create a best practice matrix to understand what is applicable and its correlation.
- Refer to relevant SAP Notes while designing, planning and before a scheduled activity. SAP Notes give you instructions

on how to remove known errors from SAP systems and may include workarounds, correction instructions, and links to support packages that solve problems. Ensure that the SAP Support Portal login information is available and has the necessary authorizations to plan and execute the various stages of the project.

- Consider your old workload. Transactional systems rely on low response times, whereas analytical systems emphasize throughput of the IO-subsystem. When designing around this, do not ignore network design architecture.
- Take CPU generations into account. Older x86 CPUs did not have many cores to accommodate hyperthreading but ran at very high clock speeds. Modern CPUs come with a large number of cores but have the tradeoff that as more cores are available on a CPU, the clock speed is lowered. Some workloads benefit significantly by having higher clock speed and therefore higher single-threaded execution performance, whereas others are better with many threads, even if they are lower speed per thread. This metric is called SCU (Single Compute Unit) in SAP. In general, higher frequency per core is favorable for application processes in SAP.
- Availability and disaster recovery at the infrastructure as well as at the application level. For extreme high availability requirements, additional vendor components could be required, especially if non-disruptive, cross data center automated failover is required. The more automation that is required, the more solution testing and training required for operations staff. Also, take into consideration at what level you want to achieve resiliency. Resiliency can be achieved at the infrastructure level but doing this at the application level can potentially make a rollback easier during implementation scenarios and offer more flexibility and a reduction in components, which in turn reduces the number sources for potential issues.

- Prepare for your new environment. Settings and designs that worked on the old environment might change in your new environment. Ensure you understand how the platform behaves to maximize your gains. To leverage the CPU example from before; just because a CPU is newer, it does not mean it is automatically faster. Your traditional storage system might have been leveraging an entire array of hard drives to push performance to a single virtual disk, whereas the Nutanix architecture scales by utilizing multiple virtual disks.
- Size for peak workloads without running into limits. Peak workloads tend to happen infrequently, but their duration can be over a longer time span. When this situation occurs, you do not want to have the system running at 99% utilization for hours or days.
- Account for infrastructure outages. Ensure resiliency and redundancy is in place, and size for N+X where X can accommodate your most significant workload.
- Adhere to NUMA boundaries. Non-uniform memory access or NUMA is a method of configuring a cluster of microprocessors in a multiprocessing system so that they can share memory locally, improving performance and the ability of the system to be expanded. Not adhering to NUMA boundaries can have a severe impact on performance, and in some cases is not even allowed.
- Do not ignore performance requirements for ancillary systems such as SAP Solution Manager, Web dispatcher or Content server among many others. While they may not be resource intensive generally, they can quickly become a bottleneck if not sized correctly or appropriately designed. Any system that is part of a critical business transaction or process needs to be treated as such.
- Build a clear integration architecture. Business critical systems, such as SAP, do not exist on their own in any organization.

Dozens, if not hundreds of systems, are interfacing with SAP either via a middleware layer or directly through various system calls. It is easy to overlook many things in this complex multi-vendor product matrix. These systems need to be accounted for in the design, and security and printing processes become extremely critical in this area, so ensure that the application security aligns with the overall security architecture of the setup, including that of your cloud platform, such as Nutanix.

12.3 Risks

These are some of the risks associated with running SAP:

- Failing to meet the business requirements through lack of planning and validation. If you fail to plan, you plan to fail. Not meeting business requirements is far too common for business-critical applications. Diligence, care and attention to detail while validating all business requirements are essential to a high-quality project delivering the desired business outcomes. Properly documenting design decisions, acceptance criteria and responsibilities are vital. Consider using a RACI matrix to describe participation by various roles.
- Avoid contention at all costs for all production workloads or landscapes. Often, with general workload virtualization, overcommitment of resources is acceptable. This is something to avoid with productive SAP instances, and for some solutions is not allowed at all. Often the same approach is used for the quality assurance environment, but development environments allow for certain amounts of overcommitment.

- A non-production environment (dev/test) can be almost as critical as production and require careful consideration and planning; especially when there is a high cost to productivity loss, or when development and testing is a significant part of your business, particularly when outside business partners also integrate with these systems and you have contracts around the SLAs.
- If you do not have objective and validated baseline performance information and business metrics before a migration project, you are not able to determine if it meets your performance requirements and it makes troubleshooting afterward incredibly difficult. Establish that there is a performance test phase in every project plan before a go-live of any project or roll-out, to ensure that business are not impacted during production operation. As an end customer, explore if you can automate these tests.
- Be aware of tradeoffs and constraints when designing the solution. For example, when integrating into an existing Nutanix cluster setup with an SAP HANA database system design, you would need to ensure that the existing cluster consists of Skylake hosts only. Because Nutanix AHV automatically reduces the amount of available CPU instruction sets for a VM in a cluster with mixed CPU generations, we would not be able to install the SAP HANA database on a VM, which requires an Intel Skylake CPU.
- Do not design for cost first. Obviously, nobody has an unlimited budget. So, while the budget is a valid constraint, always start with the (business) requirements, then if required, optimize for the budget. Reversing this order is very common, and more frequently than not, fails to meet business requirements.

12.4 References

Best Practices: SAP on Nutanix:
https://www.nutanix.com/go/virtualizing-sap-on-nutanix-best-practices.php

77% of the world's transaction revenue touches an SAP system:
https://www.sap.com/documents/2017/04/4666ecdd-b67c-0010-82c7-eda71af511fa.html

Nutanix Announces AHV Certification for SAP® Business Suite Powered by SAP NetWeaver®:
https://www.nutanix.com/press-releases/2016/11/10/nutanix-announces-ahv-certification-sap-business-suite-powered-sap-netweaver/

The Only Hypervisor Certified for Production SAP HANA on Hyperconverged Infrastructure (HCI): Another First for Nutanix AHV:
https://www.nutanix.com/2018/08/28/hypervisor-certified-production-sap-hana-hyperconverged-infrastructure-hci-another-first-nutanix-ahv/

13

Hardware Platforms

Author: Wayne Conrad

Nutanix Acropolis Operating System (AOS) supports a wide variety of vendors and hardware form factors, especially with the move towards software only, that vary greatly in size, performance, cost and every other variable. Here are the factors you should consider when picking a platform for a Nutanix cluster. The table below summarizes the supported vendor hardware.

TABLE 5

Hardware Platforms supported by Nutanix

Vendor	Model	Use-Cases
Nutanix	NX	Business-Critical Apps, VDI, Compute intensive, Data intensive, ROBO, SMB
Dell Technologies	XC, PowerEdge	
Lenovo	HX	
Cisco	UCS	
HPE	ProLiant, Apollo	
Intel	SU2600	
Inspur	NF5280M5	
Hitachi	HA8000V	
Huawei	FusionServer 2288H V5	
IBM	CS	AIX, PowerLinux
Klas	Voyager 2	Rugged, MIL-spec
Crystal	RS2616PS18	

Note that the vendor support agreements for these platforms are different. Dell Technologies (XC), Lenovo and IBM are OEM agreements, where the customer contacts the vendor for support (not Nutanix). All other platforms are third-party platforms consuming the Nutanix Software-Only model, where Nutanix provides direct support for the software only. Nutanix maintains a Hardware Compatibility List for the supported vendor hardware.

Only the Nutanix NX platform provides complete support through Nutanix directly. This is an important consideration for customers who want to leverage the value of the Nutanix NPS score of 90+ year after year.

Hardware Performance and Capacity

13.1

Anyone familiar with virtualization should hopefully be familiar with CPU and memory sizing at this point, but we will briefly consider the traditional virtualization sizing considerations. Nutanix does add a wrinkle most of us have not considered in years - local storage.

CVM Overhead

13.1.1

The CVM will typically use 32GB of RAM and between four and eight cores of CPU. The CPU usage of the CVM is primarily driven by random IO and storage features like compression or software encryption use. The memory of the CVM may need to be increased with very large active set sizes or heavy utilization of deduplication.

Remember that the CVM is pinned to the first CPU, and the CVM vCPU size may cause co-stop with other large CPU count VMs as they may not be able to run side by side with the CVM.

CPU Considerations

13.1.2

The industry has standardized on dual socket for almost every use-case, with a few exceptions. Cost sensitivity in ROBO might mean single socket, and the needs of some high-end business critical applications like SAP HANA, Oracle RAC, Epic Hyperspace and Intersystems Cache might need four or more sockets. Nutanix supports four-socket nodes, other HCI vendors only support a maximum of 2-socket nodes.

Generally, when running virtualized workloads or most applications, more cores is better. However, there is still a surprising amount of applications bound by single thread performance, especially legacy applications or in the VDI space. There are also applications out there that now charge per core instead of the traditional per socket model.In both of those cases, CPUs with less cores but higher GHz is a much better idea.

Modern CPUs generally increase performance per thread, which should be taken into account in sizing, and don't forget that Meltdown, Spectre and other CPU security vulnerabilities had a much worse performance impact on older platforms.

13.1.3 RAM Considerations

Generally, RAM is pretty easy to size based upon capacity, not performance, but there are a few edge cases to be aware of. The Intel Skylake platform has an unusual number of memory channels and DIMMs, so 768GB is the new 512GB style sweet spot for memory.

Secondly, you do not want to have a VM spanning across two sockets or using memory from more than one socket if you can avoid it. In modern systems, each memory bank belongs to one processor, and accessing memory across the other processor has a significant latency and thus performance penalty. This is called non-uniform memory access, or NUMA, and applications need to be NUMA aware to perform well in these scenarios. While many common off the shelf business critical applications are now NUMA aware, there are always exceptions like Microsoft Exchange.

In general, Nutanix NX nodes do not recommend or support mixing of 1) Memory Speed (e.g. 2400 and 2666), 2) Memory Vendor (e.g. Samsung and Hynix), and 3) Memory Capacity (e.g. 16GB and 32GB) in the SAME memory channel.

13.1.4 Storage Considerations

Storage performance on a server is something that most of us have not thought about in ten years since the switch from local RAID and direct attached storage to centralized SANs. Basically, your storage performance and capacity will both be driven by the number of hard drive slots on your server. Additionally, in hybrid storage, you have got to worry about the amount of SSD space compared to your traditional HDD.

The corollary to “You always run out of RAM first so you should buy as much as you can afford” in virtualization, is your SSD space in hybrid storage. The working set size will almost certainly increase over time as software continues to increase in size and their patches get larger and larger. Consider the growth in size of your server gold images over the last ten years - did you start with 20GB Windows 2003 images that are now 50-60GB on Windows Server 2016? One of the easiest ways to size your hot tier is by looking at the change rates on daily backups.

All SSDs are 2.5-inch, 3.5-inch really does not bring any benefits, but 3.5-inch traditional hard drives can provide a lot more space in hybrid nodes at a much lower price.

Deep storage nodes provide a lot of slots for either 2.5-inch or 3.5-inch hard drives. Deep storage nodes are generally used for Files and Objects use-cases. Some deep storage nodes are more performance oriented for business-critical applications, like large databases.

Nodes with a single SSD may be more cost-effective but have more risk of performance issues with SSD failure. Exercise caution with the intended use-cases when purchasing single SSD nodes.

Self-encrypting drives (SEDs) are more expensive than the normal model of the same capacity SSD or HDDs, but offer better storage performance than software encryption with reduced CPU overhead.

Network Considerations 13.1.5

The clear majority of Nutanix nodes ship with dual 10GbE NICs. 1GbE is usually seen in remote office locations, the extra pain and expense of upgrading is not deemed necessary yet. The declining cost of 10GbE switches suggests that 1GbE NICs are an endangered species, and will not be seen for many more years.

IPMI ports, also known as iDRAC, iLO, or generically lights out management ports are almost always 1GbE and cabled to a separate switch, surrounded by firewalls. Nutanix NX nodes, and other manufacturers, support failing over from the dedicated 1GbE management port to a production NIC if necessary, but this is rarely seen in the wild as IPMI is only used to troubleshoot failures. It is extremely unlikely that a 1GbE switch failure would happen at the same time as a host hardware failure, or that the dedicated 1GbE NIC would fail yet the host remain usable and need to be remotely accessed via the failover port.

Extra NICs, either PCIe added or on-motherboard ports, are primarily used for physical network segmentation. This is usually seen in DMZ or air gapped shop floor and SCADA style industrial control applications where traffic needs to be physically separated for regulatory or security comfort reasons.

For most workloads, 10GbE NICs are sufficient thanks to data locality, but on faster all flash platforms with a lot of disk or NVMe devices, extra NICs (25GbE or 40GbE) can help drive everything at full speed. Remember that without LACP, each VM NIC will only use the maximum bandwidth of one physical NIC, so extra NICs tend not to provide too much benefit. With LACP, each VM can use multiple physical NICs, but each IP flow can only use one NIC maximum bandwidth.

13.1.6 GPU Considerations

GPUs for VDI can significantly decrease user density, but are the only way to make some applications work, or support users with high resolution multi-monitor requirements. Consult the NVIDIA documentation for the current GPU profiles. GPUs are bulky and put out a lot of heat in a chassis, so larger 2U or 4U chassis hold more GPUs. Note that NVIDIA GPUs will not work for VDI without a separate license and license server VM running.

Logistics, Deployment & Support 13.2

OEMs have varying global distribution networks for parts or installation. An OEM with truly global reach may be needed to support remote sites in smaller countries, for instance. Some OEMs or their resellers can perform integration work, such as loading custom images onto servers or racking all hardware, cabling up switches, and configuring everything for drop in installation into your data center.

Also consider the level of vendor support required for Day-2 operations. Check the NPS scores for OEM vendors. If this is a critical requirement, then Nutanix Support is premier.

13.3 Ruggedization

Nutanix has several partners who specialize in ruggedized servers that can be used in extreme thermal, humidity and vibration environments. This is seen primarily in Oil & Gas, manufacturing, and defense applications (MIL-spec).

13.4 Compliance

Not every manufacturer holds every certification for sales to government entities or may not hold the correct regulatory certifications for some workloads.

13.5 AHV Compute Only Nodes

New to AHV is the support for compute only nodes, which do not run a CVM or have local storage. The obvious downside of this new option is the lack of data locality, so what use-cases make sense for compute only nodes? Compute only nodes allow us to use every CPU clock cycle and all the memory without worrying about CVM overhead.

Good use-cases for compute only nodes are monster CPU and memory business critical application VMs, especially those with license per server or per code. If your applications are not CPU or memory bound, you may be better off with partially populated all flash or traditional configurations that have data locality.

Poor use-cases for compute only nodes include cluster expansions where additional storage is judged unnecessary, VDI, stateless cloud apps and containers, general virtualization or just about anything else. Compute only nodes are really targeted only at the very specific monster VM use-case above.

Compute only nodes may be utilized with storage only nodes to provide the best possible protection against aggressive vendor licensing practices. With storage-only nodes and compute-only nodes, it is easy to prove which hosts can use the licensed software. This works especially well with using low core count, high performance per thread CPUs for compute, to get the most value possible if software is licensed per core versus per thread.

13.5.1 Compute-Only Node Requirements

- Four HCI or Storage only nodes
- Two or more HCI or Storage Only nodes per Compute Only node.

- Minimum of Two Compute Only nodes
- Two or more 10GbE interfaces for Compute Only nodes
- Two or more 10GbE interfaces for HCI or Storage Only nodes
- vCPUs assigned to CVMs on HCI nodes must be greater or equal to the total available cores on all Compute Only nodes
- HCI / Storage Only networking bandwidth greater or equal to double the total available bandwidth of CO nodes

Since compute only nodes are for monster VMs, please consider 25GbE or better networking, and/or LACP unless you know the storage and network throughput expected.

Other suggestions include maximizing the CVMs on the HCI or storage only node to use all the processor cores and all the RAM on the processor.

13.6 References

Nutanix Hardware Platform:
https://www.nutanix.com/products/hardware-platforms/

Dell EMC XC Series Hyper-Converged Appliances:
https://www.dellemc.com/en-us/converged-infrastructure/xcseries/index.htm

Lenovo ThinkAgile HX Series:
https://www.lenovo.com/us/en/data-center/software-defined-infrastructure/ThinkAgile-HX-Series/p/WMD00000326

Best Practices Guide: Nutanix on Cisco UCS® C-Series:
https://www.nutanix.com/go/nutanix-hyperconverged-ucs-c-series-best-practice.html

Best Practices Guide: Nutanix on Cisco UCS® B-Series:
https://www.nutanix.com/go/nutanix-ucs-b-series-best-practice.html

Best Practices Guide: Nutanix on HPE® ProLiant®:
https://www.nutanix.com/go/hpe-proliant-best-practices.php

Compatibility matrix:
https://portal.nutanix.com/#/page/compatibilitymatrix

Nutanix Compute Only Environment Minimum requirements:
http://www.joshodgers.com/2019/02/20/nutanix-compute-only-environment-minimum-requirements/

Sizing a Nutanix Cluster:
https://vcdx133.com/2015/12/15/npx-sizing-a-nutanix-cluster/

14

Sizer & Collector

Author: René van den Bedem

Nutanix Employees and Nutanix Partners use the Nutanix Sizer to quickly scope and design Nutanix solutions.

The Sizer allows the creation of Workloads, which are then translated into a hardware Bill of Materials, where the vendor hardware model can be selected. Sizer also generates the Rack Layout which includes rack units, typical power, typical thermal and weight data.

Note that Nutanix Sizer does not currently support non-x86 IBM CS hardware.

The Nutanix Collector is a data collection agent that provides a file that can be imported into Nutanix Sizer to define the Workload profiles. Nutanix Sizer also supports the import of RVTools and AWR output files for defining workload profiles.

14.1 Design Considerations with Sizer

When the workloads are defined, ensure that the correct Resiliency Factor (RF) is selected, Compression is disabled, Deduplication is at 0% savings and Erasure Coding is disabled.

If you have hard data on the expected data reduction ratios, use them here, otherwise they are assumptions, which introduces risk and should be avoided.

Architects who cannot access Sizer, can use the Nutanix Storage Capacity Calculator in conjunction with the hardware Data Sheets to size a solution manually. The data reduction recommendations mentioned previously apply to the Nutanix Storage Capacity Calculator also.

References

14.2

Size and Design Your Web-Scale Data Center with Nutanix Sizer:
https://www.nutanix.com/2014/11/10/size-and-design-your-web-scale-datacenter-with-nutanix-sizer/

Make the Move to Hyperconverged Infrastructure:
https://www.nutanix.com/go/size-your-data-center.php

Nutanix Sizer:
https://sizer.nutanix.com

Nutanix Storage Capacity Calculator:
https://services.nutanix.com/#/storage-capacity-calculator

Nutanix Collector:
http://download.nutanix.com/documentation/Documents_ANY_Version/Nutanix-Collector-User-Guide.pdf

15

IBM Power Systems

Author: René van den Bedem

The Nutanix Enterprise Cloud platform supports running IBM AIX and PowerLinux on the IBM POWER8 hardware with a specially compiled version of Nutanix software.

For any company that is wrestling with the transformation of non-x86 IBM Power Systems to x86 platforms, this option can simplify the consolidation process. This removes the need for Application refactoring, Database and Guest OS migrations and becomes a Hypervisor/Hardware migration exercise, which reduces risk, complexity, timeline and cost. In addition, the Nutanix Enterprise Cloud platform eco-system tooling would be used for those IBM Power Systems running AHV.

15.1 Use-Cases

The following use-cases drive AIX and PowerLinux with IBM Power Systems on Nutanix:

- **Cost Reduction** – Reduce licensing and support costs for IBM and Oracle software.
- **Management simplicity** – Replace IBM SystemDirector and HMC with a management eco-system that aligns with an existing Nutanix investment.

15.2 Design Considerations

These are some of the design considerations for IBM Power Systems on Nutanix:

- **Management & Control Plane** – SystemDirector, LPARs, VIO Server and HMC are no longer required and are replaced by AHV, Prism Element, Prism Central and IPMI as the management, control and virtualization plane.

- **IBM POWER8 hardware** – CS821 or CS822 node types.
- **Hypervisor scheduling** – AHV schedules threads for x86 and cores for POWER8.
- **Threads per core** – Intel CPUs have 2 threads per core, IBM POWER8 CPUs have 8 threads per core.
- **Workload types** – POWER8 processors excel at computationally intensive workloads, in particular high throughput databases and cognitive/AI workloads.
- **AHV version** – Run Acropolis Operating System (AOS) 5.2.1.1 and Acropolis Hypervisor (AHV) 20170331.78 or later. AHV and Acropolis (CVM) have been compiled for the POWER processor.
- **AIX version** – Run AIX 7.2 with the 7200-02 Technology Level with Service Pack 7200-02-02-1810 and APAR IJ05283 or later.
- **PowerLinux OS and version** – SLES 11 or 12, Ubuntu 16.04 or 17.04, CentOS 7 and RHEL 7 are supported.
- **App/DB binary porting** – Ensure app/DB binaries are ported to ppc64le (PowerLinux) and ppc64be (AIX) to run on the IBM CS platform with AHV.
- **Image Services** – NIM/mksysb are supported.
- **Testing** – Nutanix X-Ray 2.3 or later supports the IBM CS platform as test targets for scenario-centric performance testing.
- **Workload Migrations** – Big endian/little endian is not applicable here, since it is a Power System to Power System migration. Data can be migrated using similar techniques available to the Nutanix x86 platform, there are no custom migration tools available from Nutanix.
- **Support process** – Customer contacts IBM for support.
- **Implementation process** – Implementation services via IBM Lab Services.
- **Nutanix Licensing** – Current cluster licensing is valid for non-x86 and x86 clusters.

- **Prism Central** – Separate non-x86 and x86 clusters supported within the same Prism Pro instance.
- **Nutanix feature support** – non-x86 and x86 AHV/CVM code follows the same feature roadmap. All of the features in AHV/CVM are available to the IBM Power Systems cluster.

15.3 Risks

- By avoiding application refactoring, there is still a need to maintain operations and administration staff that can manage AIX and PowerLinux. However, the management and monitoring complexity will be reduced by leveraging the Nutanix Enterprise Cloud Platform eco-system.
- Workloads that require the resources of big-iron Power Systems will not be able to run on the supported IBM CS821 and CS822 hardware. Those workloads would continue to run on Frame-based hardware.

References

15.4

IBM® Simplifies Data Centers with AIX and Linux on Nutanix:
https://nutanix.com/ibm

IBM Hyperconverged Systems product page:
https://www.ibm.com/it-infrastructure/power/hyperconverged

Oracle on IBM Hyperconverged Systems with AIX Best Practices Guide:
http://download.nutanix.com/solutionsDocs/BP-2104-Oracle-IBM-Hyperconverged-Systems-AIX.pdf

IBM Hyperconverged Systems powered by Nutanix:
https://www.ibm.com/blogs/systems/ibm-hyperconverged-systems-powered-by-nutanix/

IBM AIX on IBM Hyperconverged Systems powered by Nutanix:
https://www.nutanix.com/documents/solution-briefs/ibm-aix-on-ibm-hyperconverged-systems.pdf

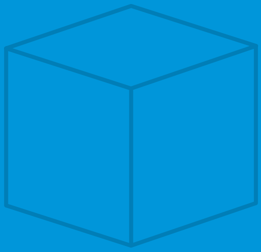

16

Remote Office, Branch Office

Author: Greg White

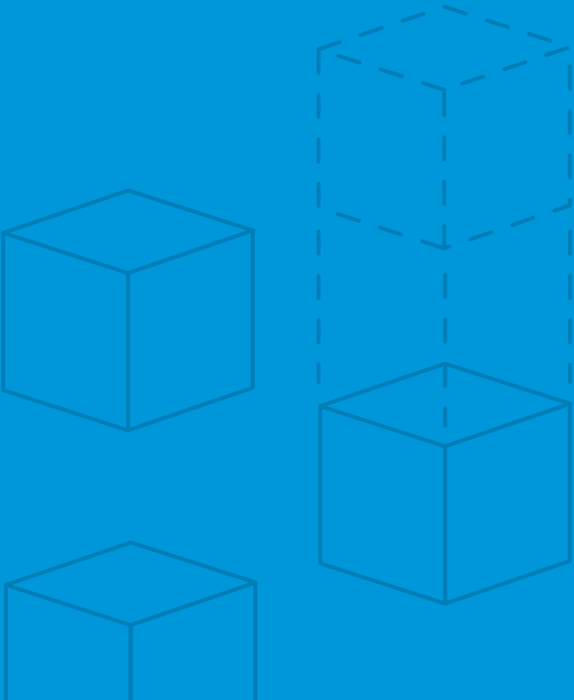

Remote and branch office (ROBO) and edge locations have historically been challenged by traditional infrastructure offerings. The lack of dedicated local staff; cost, power and space restrictions; connectivity, sizing and scaling challenges; and data protection, DR and security limitations have caused data center IT teams major inefficiencies and countless headaches. HCI found an early foothold in these environments due to immediate benefits around:

- Management,
- Sizing and scaling,
- Data protection, and
- Providing a single platform for multiple, different workloads.

Nutanix continues to build on initial successes in ROBO and edge by adding capabilities to address the challenges and required efficiencies these environments face. We will discuss these innovations, including 1 and 2 node clusters, scheduled upgrades, file services, cluster tagging, backup target and more in this chapter.

16.1 Use-Cases

Due to the ability to architect solutions using 1 or more nodes, Nutanix is able address a wide spectrum of use-cases for ROBO and edge sites in all key verticals. Whether it is a retail store, restaurant, manufacturing site, bank branch, drill rig, ship, clinic or other location where latency, connectivity, data locality or business reasons dictate a need for local compute and storage resources, the flexibility of the HCI-based Nutanix Enterprise Cloud software, and variety of hardware platforms and configurations, ensure that unique needs can be met.

Three-Node Clusters 16.2

Three-node (and more) clusters provide the broadest set of capabilities and resources. Able to typically handle more than 15 VMs, RPOs in minutes and data rebuild times, that occur without user intervention in as low as 60 seconds, they typically support larger sites and where more applications and important operations are kept local.

A self-healing Nutanix three-node cluster also obviates needless trips to remote sites. It is recommended that these solutions be architected with enough capacity to handle an entire node going down, which would allow the loss of multiple hard drives, one at a time. Because there is no reliance on RAID, drives can be lost and healed, one after the other, until available space runs out. For sites with high availability requirements, or that are difficult to visit, it is recommended to add additional capacity above the N +1 node count. Also, ensure that there is 5% capacity reserved for cluster processes in the event of a failure to ensure that a full cluster cannot perform necessary rebuild operations. Three-node clusters can scale up to eight nodes with 1GbE networking, and up to any scale when using 10GbE and higher networking. They also support a wider variety of hypervisors by including Nutanix AHV, VMware ESXi, Microsoft Hyper-V and Citrix Hypervisor for VDI. Lastly, three-node and above configurations are well-suited for adding local file data using Nutanix Files.

Two-Node Clusters 16.3

Two-node clusters offer reliability for smaller sites that must be cost effective and run with tight margins. These clusters use a witness only in failure scenarios to coordinate rebuilding data and automatic upgrades. You can deploy the witness offsite up to 200ms away and multiple clusters can use the same witness for two-node and metro clusters. Metadata is maintained in RF4 with 2 copies on each node and data is RF2 with a copy on the other node. In a failure scenario, the operational node will create a 2nd (RF2) copy of lost node's data. During this rebuild, additional writes are held.

Once the other node is online again, the metadata is restored to the other node back into RF4 across both nodes. Data is rebuilt to RF2 across the nodes as it is accessed for the restored node. The witness VM requires, at a minimum, 2 vCPUs, 6 GB memory, 25 GB storage and must reside in a separate failure domain, which means that there must be independent power and network connections from each of the two-node clusters. Nutanix recommends locating the witness VM in a third physical site with dedicated network connections to avoid a single point of failure. Two-node clusters can only run ESXi or AHV.

16.4 One-Node Clusters

One-node clusters are a perfect fit if you have low availability requirements and need strong overall management for multiple sites. One-node clusters provide resiliency against the loss of a hard drive while still offering great remote management. Additional considerations for these deployments include understanding local needs and impacts from upgrades and drive failures.

Upgrades of one-node clusters are disruptive, requiring planning for downtime, and in the event of a drive failure, the node goes into read-only mode until the drive is replaced and data can be rebuilt from the RF2 copy on the node. The read-only mode can be manually over-ridden. Nutanix supports one-node with ESXi and AHV only and recommends reserving 55 percent of useable space to recover from the loss of a disk.

Backup and Disaster Recovery

16.5

Remote sites are great candidates for storing another copy of native Nutanix snapshots for recovery and DR purposes. Configuring backup on Nutanix lets an organization use its remote site as a replication target to retrieve snapshots from it to restore locally, but failover protection (that is, running failover VMs directly from the remote site) is not enabled. Backup also supports using multiple hypervisors, for example ESXi at the main site and AHV at the remote site. Configuring the disaster recovery option allows using the remote site both as a backup target and as a source for dynamic recovery so that failover VMs can run directly from the remote site. Nutanix provides cross-hypervisor disaster recovery between ESXi and AHV clusters. Hyper-V clusters can only provide disaster recovery to other Hyper-V-based clusters.

For sites where there is a large volume of data, higher latency and/or a demand for a faster RTO, the one-node backup target is available to keep a protected copy of data locally off the active cluster. This target, using native Nutanix snapshots, can also be used to store a replicated, off-site copy of data from other sites or the main data center. There are limited compute resources available on the one-node target for supporting additional VMs beyond the backup and replication processes.

This can be reserved to run a key VM or two in the event of a failure on the cluster or can be used for lightweight VMs that are not critical, like a print server, until the resources are needed for a recovery. Best practices for one-node backup targets for three-node and larger clusters (consult the best practices guide for specifications for one and two-node clusters):

- All protection domains combined should be under 30 VMs.
- To speed up restores, limit the number of VMs in each PD.
- Limit backup retention to a three-month policy.
- Recommend seven dailies, one monthly, and one quarterly backups.
- Map a one-node target node to only one physical cluster.
- Set the snapshot schedule to six hours or more.
- Turn off deduplication.

16.6 Prism Central Management

The additional capabilities available from Prism Central Pro for managing ROBO sites include:

- Customizable dashboards,
- Capacity runway to safeguard against exhausting resources,
- Capacity planning to safely reclaim resources from old projects and just-in-time forecasting for new projections,
- Advanced search to streamline access to features with minimal training, and
- Simple multiple-cluster upgrades that can be scheduled. There is also the capability to schedule upgrades to run in certain windows, as well as exclude times when they should run. For example, you could designate upgrades not run after 8 AM when employees arrive for the work day. In addition, upgrades can be run all at once (simultaneous) or one-at-a-time (staggered.) Consider these best practices for upgrade planning: If WAN links are congested, pre-download your upgrade packages near the end of your business day and perform pre-upgrade checks before attempting the full upgrade.

Prism Central also provides labeling and tagging for VMs and clusters. Tag VMs with labels to easily sort and find the ones associated with a single application, site, business owner, or customer. Tag clusters for similar needs in order to quickly identify clusters by size, geography or other characteristics you specify. You can perform operations or actions, like an upgrade, on multiple entities at the same time.

Built-in Hypervisor, File Services, and Micro-segmentation 16.7

AHV is a good fit and proven hypervisor option for ROBO and edge sites due to the enhanced efficiency and reduced complexity that can be realized with built-in services. AHV as an enterprise-grade hypervisor included with Nutanix AOS provides the foundation for remote HCI deployments. It is required to take advantage of Nutanix Files and Flow.

Often remote and branch sites need local file server capabilities. Nutanix Files provides the ability to have local SMB and NFS/Linux file data. As little as one node can be enabled for file data, and it can then be expanded either scale-up or scale-out. Consult the Files chapter for more information on design considerations and limitations.

Security at remote and branch sites and at the edge can also be a challenge due to the lack of attention and resources available locally. Nutanix added Flow to provide built-in micro-segmentation capabilities. This can be enabled at remote sites to protect east-west traffic in the event of a breach. Consult the Flow chapter for more information on design considerations and limitations.

16.8 References

ROBO Solution page:
https://www.nutanix.com/solutions/remote-and-branch-office/

ROBO Deployment and Operations Best Practices Guide:
https://www.nutanix.com/go/deploy-and-operate-robo-on-nutanix.php

ESG Solution Showcase: Nutanix Simple Turnkey ROBO IT:
https://www.nutanix.com/go/simple-turnkey-it-infrastructure-for-remote-and-branch-offices.html

17

Xi Frame & EUC

Author: Kees Baggerman

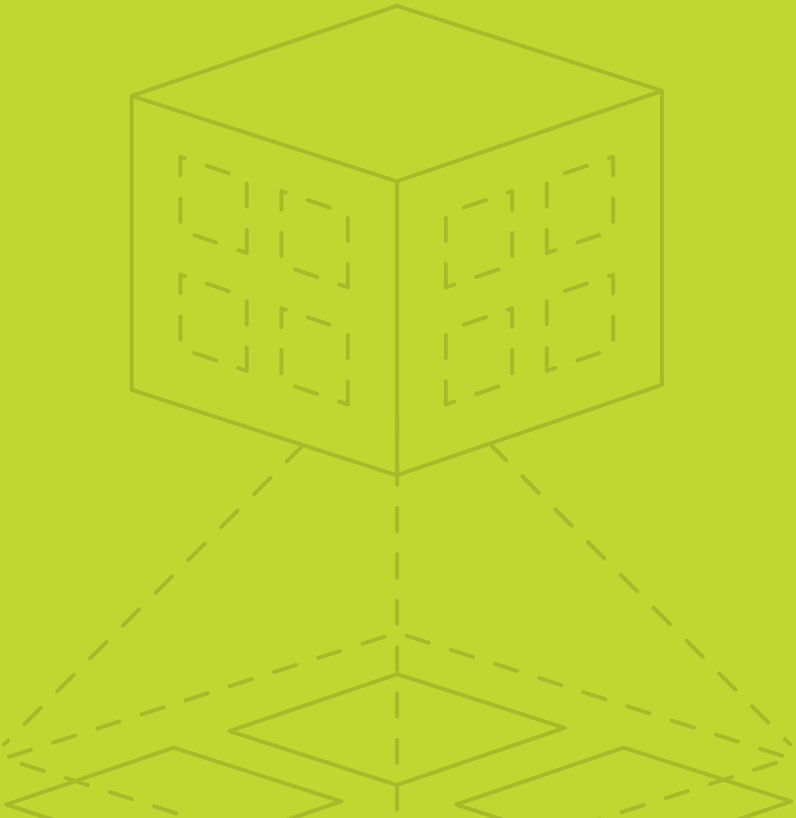

Nutanix empowers IT with a software-defined platform ideal for organizations seeking efficient and cost-effective alternatives building the next generation data center. The Nutanix Enterprise Cloud Platform liberates end-user computing projects from expensive and hard to manage traditional server and compute infrastructure.

The Nutanix solution was designed for simplicity, enabling administrators to deploy View, or XenDesktop less than one hour after racking, significantly increasing time to value and speed to market.

The Nutanix Enterprise Cloud Platform can dynamically grow or shrink with a single click, without any downtime to end users or application availability. Unlike with traditional infrastructure, linear compute and random I/O scalability from Nutanix eliminates the dependencies on extensive capacity planning and forecasting. Simply add more nodes when additional storage or compute resources are needed.

End User Computing (EUC) comes in two major form factors; Virtual Desktop Infrastructures (VDI) and Server Based Computing (SBC). VDI typically has a 1:1 ratio of VM to Users where SBC has a multiuser to VM ratio.

A VDI/SBC solution is a desktop virtualization solution that transforms desktops and applications into a secure, on-demand service available to any user, anywhere, on any device. With VDI/SBC, you can deliver individual Windows, web, and SaaS applications, or full virtual desktops, to PCs, Macs, tablets, smartphones, laptops, and thin clients with a high-definition user experience. VDI/SBC provides a complete virtual desktop delivery system by integrating several distributed components with advanced configuration tools that simplify the creation and real-time management of the virtual desktop infrastructure.

Considering that the EUC platform often is the first and most visual point of access for end users to get to backend services provided by IT, it is considered a mission critical application set. When the EUC environment becomes unavailable due to unplanned downtime or misaligned technology choices on business requirements, this immediately takes away the primary point of access for access to backend services, resulting into direct impact for end users. Even performance degradation will be instantly visible to your end users as the EUC environment is their access platform resulting in additional service calls and escalations.

The cost of unavailability or performance degradation can be considerable, meaning risk reduction and management are crucial. Setting up the EUC platform is only a small percentage of the implementation, the main points of consideration, and thus most time consuming, are:

- Availability/DR
- Apps/Data/End User Personas
- Predictable performance

The risk in all of this can be reduced by carefully designing, planning, testing and building out the platform.

Since going public (Nasdaq: NTNX) in 2016, Nutanix has always included the use-case distribution information in the infographic available with every earnings release. The proportion of EUC on Nutanix is approximately 26% and has been consistent since the reporting began. The main driving factor for this is how the Nutanix architecture reduces risk, improves predictability and performance consistency from day 1, during growth, and when disasters strike.

As outlined in the Mission-Critical Applications chapter, End User Computing is included in this category:

Virtual Desktop Environment – If it supports all users

- Pros: Predictable linear scale out performance, High user density per node, Predictable failure characteristics, Flexible deployment options including cloud and on site, Reduced deployment time.
- Cons: Not all hypervisors or brokers support all the same features or end user devices, which requires careful design and sizing considerations and selection of components.

Implementing EUC on Nutanix will result in the following benefits:

- Lowest TCO for persistent and non-persistent desktops.
- Amazing user experience – low boot time and application response.
- Elastic deduplication engine to boost performance for persistent virtual desktops.
- Efficient clones:
  - VAAI/VCAI (vSphere)
  - OCX (Hyper-V)
  - Storage native technologies (AHV and Citrix Hypervisor)
- Predictable growth & forecasting.
- Up to 10x reduction in rack space.
- 5x Reduction in project completion times.
- DR built-in: enables VM-level disaster recovery for virtual desktops.

17.1 Use-Cases

The following use-cases drive design for End User Computing environments:

- Reducing risk of downtime, performance degradation.
- Improve predictability of performance and scalability.
- Ensure consistency, both during operations and during failure and maintenance events.

- Organizational User Personas.
- The move towards digital transformation, enabling end users to work the way that fits their needs best vs IT prescribed methods that are outdated by the time they get implemented.
- Significantly lower total cost of ownership compared to traditional infrastructure solutions, with more predictable and smaller growth increments as environments scale.

Xi Frame 17.2

Xi Frame is a secure, high-performing platform that allows delivery of Windows apps to users on all connected devices. Customers select a leading cloud infrastructure to run their workloads, such as Amazon Web Services or Azure. This cloud-based platform, which is called the Hosted Service, integrates with external services for authentication and storage. For maximum flexibility and customization, Frame provides an API and documentation to developers.

FIGURE 7

Xi Frame Components

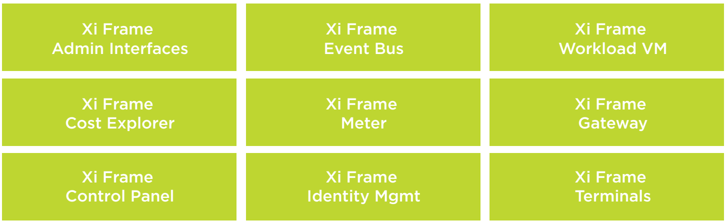

Xi Frame Admin Interfaces

Xi Frame Event Bus

Xi Frame Workload VM

Xi Frame Cost Explorer

Xi Frame Meter

Xi Frame Gateway

Xi Frame Control Panel

Xi Frame Identity Mgmt

Xi Frame Terminals

17.3 Supported VDI brokers

Nutanix supports Citrix XenDesktop and VMware Horizon View.

17.4 References

Xi Frame Product Page:
https://www.nutanix.com/products/frame/

Frame Test Drive:
https://fra.me/test-drive

Best Practices Guide: Citrix XenApp and XenDesktop on Nutanix:
https://www.nutanix.com/go/citrix-xenapp-xendesktop-best-practice.html

Citrix Runs on Nutanix:
https://www.nutanix.com/solutions/vdi/citrix/

Virtual Desktop Infrastructure (VDI) Solutions:
https://www.nutanix.com/solutions/vdi/

18

Xi IoT

Author: Rohit Goyal

In 2017, 3 billion industrial edge devices generated 256 zettabytes of data. That is over 30 times more data than what was stored across cloud and private data centers. As the number sensors and devices increase, the amount of data produced will continue to grow at a staggering rate. According to Gartner analysts, more than 50 percent of IoT projects will use edge devices for analytics by 2022.

Most organizations deal with these oceans of data by processing it all in the cloud, an approach that causes significant IT and business challenges, such as bandwidth congestion, lack of scalability, processing delays, and compliance and privacy issues.

Traditional architectures were not built to accommodate edge workloads, and efforts to employ them in this new context result in poor performance, disabling complexity, and untold lost opportunities, afforded by real-time intelligence at the edge.

FIGURE 8

Classic IoT Model

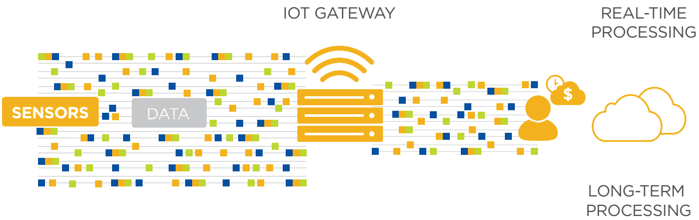

While IoT devices have been around for years, making sense of the data generated from these devices has not been a top priority for many organizations, largely due to complexity and cost. With the right edge computing and IoT platform, however, deploying planet-scale edge intelligence can be straightforward, cost-effective, and a path to unprecedented innovation within the enterprise.

The Nutanix Xi IoT platform is a 100% software-defined solution that delivers local computing, machine learning and intelligence for your IoT edge devices, converging the edge and your choice of cloud into one seamless, delightful application development platform. IoT eliminates complexity, accelerates deployment and elevates developers to focus on the business logic powering IoT applications and services.

FIGURE 9

Nutanix Xi IoT - Move Real-time Processing to Edge and Gain Faster Insights

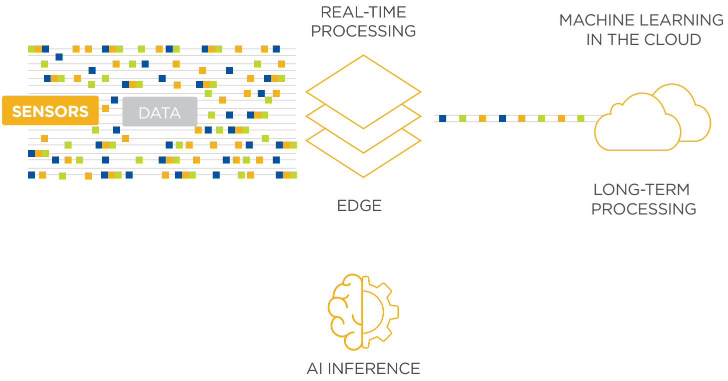

Nutanix Xi IoT is comprised of a SaaS infrastructure and application lifecycle management plane and Xi Edge running on an edge device. The SaaS management plane provides an end-to-end platform that is centrally managed from the cloud, through a user-friendly interface for application development and operations, to easily deploy thousands of edge locations.

Xi IoT has the following benefits:

- **Freedom of Choice:** The IoT platform can be delivered as a virtual machine, on standard Nutanix HCL hardware, or specialized edge hardware and seamlessly connect to any cloud.
- **Infrastructure & Application Lifecycle Management:** The end-to-end platform is centrally managed from the cloud and provides a user-friendly interface and SaaS based control plane for application development and operations.
- **Deploy Complex Applications at Planet-Scale:** The edge PaaS supports easy-to-use developer APIs, reusable data pipelines, and pluggable machine learning architecture to enable rapid development and global deployment of modern IoT apps.

The Xi Edge platform leverages Kubernetes, which allows you to consolidate traditional IoT applications as well as enable new-generation, data science-based applications in containers with the following benefits:

- Edge computing stack for real-time processing
- Centralized planet-scale ops and app management
- Data pipeline to converge edge and cloud

The Xi Edge platform provides secure access to IoT data sources with data pipelines all the way from the edge to the cloud, including AWS, Azure, GCP, and managed/on-prem private clouds. It also provides seamless data mobility between edge and cloud, which lets users send metadata and build ML models in the cloud.

Use-Cases 18.1

Manufacturing 18.1.1

Increase efficiency and maximize productivity by using edge intelligence to predict equipment failure, detect process anomalies, improve quality control, and manage energy consumption. Real-time analysis reduces decision latency and minimizes costly production delays

FIGURE 10

Xi IoT Manufacturing Use-Case

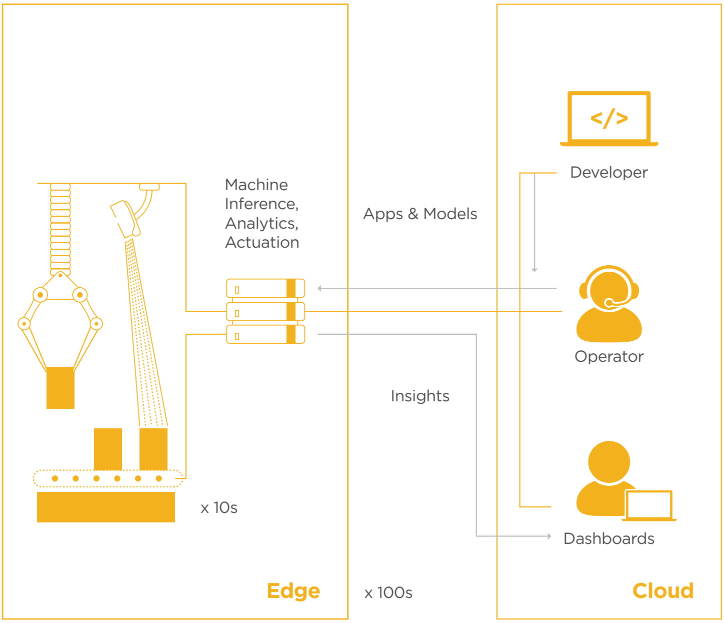

18.1.2 Retail

Deliver unique customer experiences by leveraging data at the edge to personalize offers, build an omnichannel customer relationship, and streamline the purchase process. Edge data can also improve inventory management, ensuring product availability and easing supply chain strains.

FIGURE 11

Xi IoT Retail Use-Case

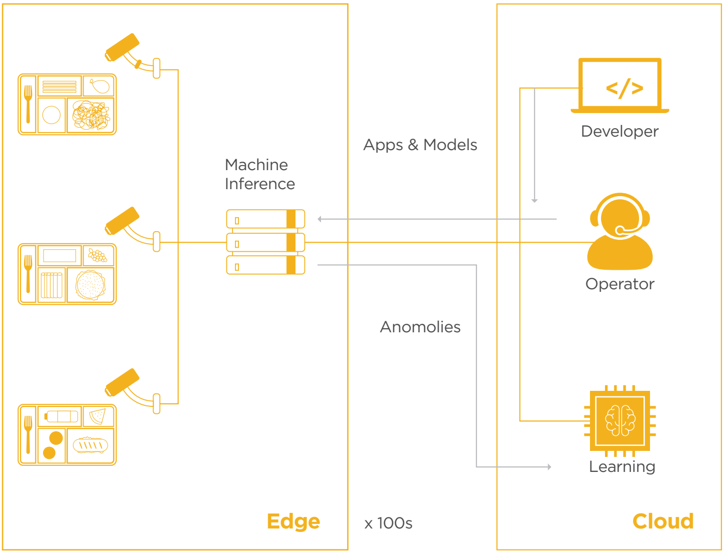

18.1.3 Oil and Gas

Transform upstream and downstream operations with edge intelligence. Real-time analysis of well sites can optimize extraction processes, and analysis at retail locations can identify trends to maximize revenue.

Healthcare

18.1.4

Edge-based diagnostic equipment and monitoring tools bring processing and analysis closer to the patient, improving care and services without compromising patient privacy. Realtime detection and diagnosis can make a significant impact on patient outcomes.

Smart Cities

18.1.5

Connected city services can dynamically improve traffic flow when trouble spots appear, dispatch emergency personnel quickly, and detect issues with utilities before they become problems. With the amount of data involved from all devices and sensors across the city, computing at the edge is the only viable approach.

Design Considerations

18.2

Technology Challenges with Edge Computing:

- **Point solutions versus Platform Approach:** Solving a single problem versus utilizing a platform to solve many problems.
  - **Consideration:** It is important to think about an architecture that enables multiple IoT applications that can be deployed at the edge versus solving a single pain point. Additionally, a new platform is required to ingest data from devices (e.g. imaging) or sensors in real-time. It is not easy to predict future business challenges today. However, it is important to choose an architecture that is flexible enough to handle those challenges as they arise.
- **Bandwidth Congestion:** As more and more internet connected devices arrive on edge networks, sending all this data to the cloud creates bandwidth congestion and increases costs.
  - **Consideration:** It is important to make sense of the data where it is generated but is not important to send all this data

from the edge to the cloud. Intelligently analyzing the data at the edge and only sending relevant data to the cloud for long term processing will not only save on bandwidth costs, but reduce application contention.

- **Lack of Scalability:** The scale of deployment, frequently involving hundreds to thousands of locations makes it even more challenging.
    - **Consideration:** How does the solution scale from a single site to managing thousands of locations with simplicity and ease? The setup and maintenance of the edge platform must be centralized and simple enough to start with a single site but flexible enough to scale to thousands of edge locations with a few clicks. It should not require heavy lifting from remote staff.
- **Processing Delays:** Imagine a scenario where it takes hundreds of milliseconds for data to travel from the edge to cloud and back, only to find out the problem at the edge location has now taken down a production facility. The ideal situation is to know about a problem as soon as it occurs.
    - **Consideration:** Enterprises should easily be able to shift processing from the cloud to the edge to reduce latency between data generated and problem alerted, to a few milliseconds. Processing data in the cloud is not “real-time” and should be avoided for fast results. However, processing data in the cloud is great for long term deep learning.
- **Compliance and Privacy Issues:** Sending data outside the enterprise and/or the country is not always permitted to meet compliance or privacy requirements.
    - **Consideration:** The edge platform should provide a simple way for developers to connect the edge to multiple public or private cloud options with a matter of a few clicks.

- **Protocol Diversity:** Previously, an edge cloud (with local appliances connected to sensors) was very difficult to operationalize due to the diversity of sensors, which communicate via protocols like Modbus, CAN bus, PROFINET, and MQTT, and require different physical interfaces.
  - **Considerations:** To work with many existing environments, it is essential to have an edge platform that enables secure and easy connectivity between IoT sensors or devices. The platform should easily be able to ingest data from multiple sources and run machine inference based on requirements set by the enterprise.
- **Ability to Support Cloud Native Applications:** Next-generation cloud native applications require new constructs and AI (Artificial Intelligence) frameworks. Applications need to run on a range of devices with different types of CPU, as well different types of GPU, ASICs, FPGAs, and add-on cards from various vendors.
  - **Considerations:** When designing applications at the edge, it is vital to leverage hardware components to increase processing capabilities. However, it is not always easy to figure out how to leverage it. The edge platform should make it easy to automatically leverage the hardware based on the application requirements.
- **Adapting to New Staff Requirements:** The human-element of IT: operational technologist, developers, and data scientists, all need to come together to operate IoT applications.
  - **Considerations:** Developers and data scientists should be able to bring their own cloud and machine learning models from any domain and access rich data and runtime services to execute AI at the edge. Developers should also be able to leverage APIs and integrate into existing CI/CD workflows for easy debugging.

18.3 Risks

These are some of the risks associated with edge computing solutions:

- Ensure that the business creates metrics to focus on when designing edge computing. If metrics are not clearly defined, the scope of the project can drastically change and reduce the chance of success.
- Many enterprises do not realize the benefits of analyzing data close to where it is being generated and often do not leverage the data.
- Properly designing out an edge solution strategy is important and working with system integrators will reduce the overall effort required from the enterprise. It will also increase the chance of success. It is also important to understand all the business requirements and ensure they map to the provided solution.
- If the enterprise does not have data scientists on staff or expertise in image analytics, it is always good to bring in a partner that can solve those problems, otherwise it could take a longer time to solve edge use-cases.

References

18.4

Xi IoT Edge Product Page:
http://www.nutanix.com/iot

Xi IoT Solution Brief:
https://www.nutanix.com/documents/solution-briefs/xi-iot-sb.pdf

Xi IoT Retail Solution Brief:
https://www.nutanix.com/documents/solution-briefs/nutanix-iot-retail-brief.pdf

Xi IoT Oil and Gas Solution Brief:
https://www.nutanix.com/documents/solution-briefs/sb-oil-gas-iot.pdf

19

Xi Leap, Data Protection, DR & Metro Availability

Author: Mark Nijmeijer

The Nutanix architecture around the Distributed Storage Fabric is designed with data protection in mind. It provides high resiliency with full protection against any kind of component failure: SSD, HDD, node, block, cluster, and rack. This chapter will describe the various capabilities within AOS that help you protect all your applications against any kind of outage. It will also provide guidance on how to choose the best options for your organization's applications.

19.1 Data Protection Options

Nutanix AOS provides multiple data protection options that each provide different characteristics in terms of RPO, RTO and retention:

- Asynchronous replication and Disaster Recovery
- Synchronous Replication and Metro Availability

To choose the correct technology, you need to identify your business needs for protection by talking to the business owners for each application. You should determine the following factors:

- **RPO** – Recovery Point Objective. Typically expressed in units of time. The RPO denotes the amount of data the business can afford to lose. This translates in how often the system should snapshot (and replicate) the data. For instance, if you have an RPO of 15 minutes, it means that the system should take a snapshot at least every 15 minutes.
- **RTO** – Recovery Time Objective. Also expressed in units of time. The RTO denotes the amount of time the business can afford to be down, or data to be unavailable. This typically means how fast the business application needs to be restarted in a different data center. This includes the time that the business takes to do a failover (WRT). For simplicity, we are

including Work Recovery Time (WRT) and Maximum Tolerable Downtime (MTD) in the RTO calculation. Normally, MTD = RTO + WRT

- If you have an RTO of 2 hours for a particular application, it means that that application needs to be fully accessible within 2 hours.
- **Retention** - Denotes how far back the system should be able to restore to. If you have a retention goal of 3 months, you should be able to restore your application in the state it was three months ago.

The Nutanix data protection family have the following options, listed in the table below.

TABLE 6

Data Protection Options Nutanix

Method	RPO	RTO	Retention	Typical Use-Cases
Async & DR	1min –12mths	Minutes	No limits	Any app
Sync & Metro	Zero	Minutes (automated with witness)	No limits	Critical apps (< 5ms RTT)

Nutanix offers a choice in how you want to protect your applications and data:

- **On-premises** – you can use Nutanix to Nutanix replication and disaster avoidance functionality to protect your applications. Disaster Recovery functionality is licensed via your AOS licensing.
- **Xi Leap** – a Nutanix DR-as-a-Service offering to protect your applications using the Nutanix Xi Cloud. This has the benefit of not having to own or rent a secondary data center to install and run another Nutanix cluster. This option is licensed as a per-protected-VM per month model.

19.2 Nutanix Distributed Storage Fabric Snapshots

All Nutanix Data Protection Capabilities are based on the efficient VM-Centric snapshots that the Nutanix Distributed Storage Fabric provides. These snapshots are amongst the most efficient snapshots in the industry. As Nutanix manages the entire virtualization stack from workload down to the storage, it is entirely aware of what storage is in use by each virtual machine and uses this information to create snapshots at the virtual machine (or really, at the virtual machine disk) level. These snapshots are entirely meta-data based and provide byte-level incremental storage allocation for changed blocks after the snapshot has been taken.

Once a snapshot has been taken, it can be transferred to another Nutanix Cluster or to the Nutanix Xi Cloud. These transfers leverage the byte-level incremental nature of the snapshots, so only changed data will be transferred over the wire.

The schedule by which these snapshots are taken and replicated is defined by Protection Policies. These policies are defined by the customer and contain information about the business SLAs for your applications: RPO and retention goals. The admin can then create Protection Rules that automatically will apply the Protection Policy to workloads that conform to that rule. This allows the system to provide automatic protection and replication of new applications and virtual machines.

For instance, you can have a ‘Gold’ Protection Policy that states a 15-minute RPO and a 1-week retention goal. There is a rule defined that ties the Gold Protection Policy that ties it to any VM that has been configured with a Nutanix Category “protection-level equals mission-critical”.

The Nutanix DSF provides a wide range of supported RPOs. The minimum RPO as of the AOS 5.10 release is 1 minute, and the RPO can be configured as high as one year.

See The Nutanix Bible for more detailed information on the DSF snapshot technology.

Async Replication and Disaster Recovery

19.3

On top of the efficient Nutanix snapshotting and replication capabilities, Nutanix provides workflows that allow the admin to configure the system to provide disaster avoidance and recovery capabilities. The vision is to shift all of the hard work to a period of time that is typically not high-stress, and provide capabilities to help you ensure you can quickly and successfully recover from any kind of outage, with intuitive workflows and the right level of information to keep your organization abreast of progress towards and ETA of a full recovery.

The admin can define Recovery Plans for each of his applications. This Recovery Plan contains all information that is necessary to migrate that application to another data center, or to provide a failover after an outage occurs. In particular, a Recovery Plan contains the following information:

- **Boot ordering** – the admin can create groups of virtual machines that should start together. Dependencies between groups can be expressed through the definition of delays between both phases. This allows virtual machines in a particular boot phase to start up and be entirely functional before the next phase boots.
- You can use Nutanix Categories to add dynamic groups of virtual machines to a particular boot phase.

- **IP address management** – when failing over into a different data center, you typically have to re-IP your virtual machines to ensure they can communicate with the networks in the failover data center. In the Recovery Plan, you can create mappings between the vSwitches as well as define mappings to specify which IP address ranges should be used for re-IPing after a failover. You can also indicate which virtual networks or VLANs to use for the Test networks (see below).
- **Script execution** – allows the admin to specify a script that should be run as part of a failover operation. This can be used to re-configure a setting that is managed outside of Nutanix, such as a desktop broker or a global load balancer.

The recommendation is to create one Recovery Plan per application. This allows you to manage protection, replication and failover decisions at a very granular level.

To help ensure that a failover of a particular application will be successful, Nutanix provides two ways to help with this:

1. **Validate Recovery Plan** - This is a fast and efficient operation that checks whether all required resources are available and accessible for a successful failover. Examples of these resources are:
    a. **Licensing** – are all involved clusters licensed at the appropriate level?

    b. **Snapshots** – are there snapshots available for all virtual machines in the Recovery Plan?

    c. **Compute resources** – is enough CPU and memory available on the failover cluster to run the application covered by the Recovery Plan?

    d. **Networks** – are all networks defined in the Recovery Plan available and accessible?
2. **Test Recovery Plan** - This operation will instruct the failover

location to start a clone of the application covered by the Recovery Plan in its own networking bubble. It does this by cloning the applicable virtual machines from the most recent snapshots, and configure the virtual machines to be connected to the Test networks as defined in the Recovery Plan. Once all virtual machines have been registered, they will be started according to the specified boot order.

It is important to note that this operation does not interrupt any ongoing replication, and because the application is connected to the Test networks, production networking traffic is not impacted at all.

Clean-up of this test deployment is a one-click operation to ensure no unnecessary resources are being consumed.

Metro Availability and Synchronous Replication 19.4

Nutanix Synchronous Replication provide a 0 RPO data protection solution for those applications that require the highest levels of data protection. Any application write that the system processes will be acknowledged by at least 2 nodes in the local cluster (assuming RF2) and at least 2 nodes in the remote cluster before that write gets acknowledged back to the application's VM.

To avoid excessive IO overheads, Nutanix requires the latency between the 2 clusters to be below 5ms RTT (Round Trip Time). This latency is enforced when the replication is started, but the system will not interrupt replication if the latency spikes above 5ms.

Nutanix Metro Availability leverages these synchronous replication capabilities and integrate with the hypervisor's stretched cluster support to provide additional capabilities:

- **Cross-cluster live migration** - because the hypervisor cluster is stretched across the two physical Nutanix clusters, VMs can migrate transparently between the two clusters.
- **High Availability** – Applications can be restarted automatically if the clusters where the application is currently hosted goes down.

Synchronous protection is done at the container level, so any VM disk that is placed in the protected container is automatically replicated to the paired cluster.

The Metro Availability failover behavior can be configured in one of three modes:

- **Manual** – any cluster failure (primary or failover) or cluster communication failure results in the cluster pausing all IOs until the admin takes action. This mode is only recommended when the requirement for having two full datasets is more important that the requirement for availability of the protected application.
- **Automatic** – the system will automatically resume the application on the primary cluster in case there is a secondary cluster failure or network communication failure between the two data centers. The system will pause all IO for 20 seconds to wait for communication to be restored to account for temporarily blips, and once the 20 seconds has passed the system will resume the applications on the primary cluster without replication to the secondary.
- **Witness** – in the witness mode, the system will automatically handle any system outage and determine what the best location is to either keep the applications running or to start a failover. In case of a primary cluster outage, it will automatically start a failover to the secondary cluster.

The witness itself is a small VM that can be run on a vSphere or AHV node (in case of vSphere, it is supported to run on a non-Nutanix server). The witness is a passive entity in the system, meaning the witness can go down temporarily without impact on the running systems, assuming there is no primary, secondary or networking failure while the witness is unavailable.

Any required data resyncs when re-protecting the applications after an outage will be based on the most recent automatic snapshot. These snapshots get taken automatically every four hours, so the maximum amount of data the system needs to replicate is 4 hours' worth of data.

Protection Domains versus Xi Leap 19.5

Nutanix provides two ways of managing the Disaster Recovery configuration. There is the newly-released orchestration that is used for Xi Leap and on-site Leap part of AOS 5.10 and later, and there is the legacy method using Protection Domains.

Refer to the table below for a comparison between the two methods.

Licensing for the Nutanix data protection functionality can be consumed in two ways:

On-Prem Disaster Recovery, this functionality is licensed via the three AOS license levels:

- **Starter** – gives access to basic DR functionality. You can protect your applications with a minimum RPO of one hour. With On-prem Leap, you can create Recovery Plans with one boot stage and no IP management.
- **Pro** – gives access to Self-Service Restore capabilities
- **Ultimate** – Gives access to advanced Leap functionality:
    - **Multiple boot phases** – if you have dependencies between

TABLE 7

Protection Domains versus Xi Leap

Functionality	Xi Leap	Protection Domains
Management	Managed at the data center level, via Prism Central	Managed at the cluster level via Prism Element
Dynamic/Static	Very dynamic management. Protection Policies can automatically be applied to new VMs. Recovery Plan can automatically include new VMs	Being part of a Protection Domain is static. Admin must manually manage membership of VMs in a Protection Domain
Reusability	Protection Policies can be re-used to protect applications that need the same protection specs (RPO, RTO, retention	Schedules must be defined on each Protection Domain.
Scope	Managed at the data center level via Prism Central	Managed at the Nutanix cluster level through Prism Element
Licensing	On-prem Leap Basic functionality is included in AOS Starter and Pro Edition. Advanced functionality (multiple boot phases, IP address management, testing) is available in AOS Ultimate licenses. Xi Leap Xi Leap is licensed per-VM per month. AOS Starter license for on-prem clusters is sufficient	Async DR is in AOS Starter and Pro Ultimate provides Metro, NearSync (RPO < 1hr) and multi-site replication support

Virtual Machines that together form an application or service (e.g. a three-tier application with database, app and web-tiers), you can use boot phases to ensure that certain VMs are booted and available before other parts of the application boot up.

- **IP address management** – used to indicate whether different IP addresses should be used for a Test of Failover. Typically, different data centers use different IP address ranges, meaning the VMs need to use different IP addresses when they boot for a Test or Failover operation.
- **Script execution** – this executes an admin-defined script as part of the Test or Failover sequence. This script can be used, for example, to change configuration files, or re-program external entities like a global load-balancer.

- **Testing** – Testing of recovery plans to a Sandbox environment.

Xi Leap - this functionality is licensed at a per-protected VM per month basis. This consumption model is independent of the on-site AOS license used, meaning you get full access to all Disaster Recovery functionality even with AOS starter licenses.

References 19.6

Xi Leap Product Page:
https://www.nutanix.com/products/leap/

Xi Leap – The First No-Install DR:
https://www.nutanix.com/2018/11/28/xi-leap-first-no-install-dr/

The Xi Leap Service:
https://www.nutanix.com/documents/solution-briefs/leap.pdf

Backup and Recovery:
https://www.nutanix.com/products/acropolis/backup-and-recovery/

DR Orchestration:
https://www.nutanix.com/products/acropolis/dr-orchestration/

The Nutanix Bible:
https://nutanixbible.com

20

Cloud Management & Automation: Calm, Xi Beam & Xi Epoch

Author: Chris Brown

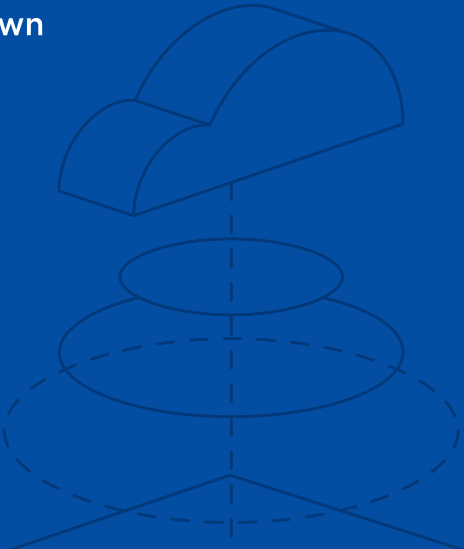

Nothing is ever 100% in IT. Even the most ardent Dell-EMC fan has at least some Netapp running in their environment just in case something goes wrong with a Dell-EMC patch. Clouds are the same. Using a single cloud introduces a new single point of failure, but the friction of maintaining policy, governance, and control across clouds makes it difficult to use more than a single public cloud at a time.

FIGURE 12

Multi-Cloud Operations

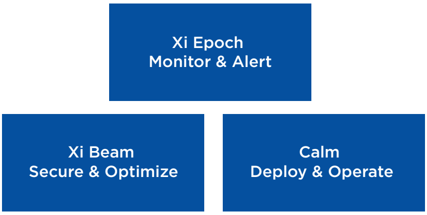

As we at Nutanix have shown, differentiation in the market today is provided through software not hardware. An enterprise's ability to compete in their market - no matter the industry - is tied inextricably to their IT team's ability to deliver applications faster than ever before. Even business problems that, on their surface, do not look like IT problems are driven by IT's speed. IT used to be seen as just a cost center - a price required to do business. In the Internet age this has been flipped on its head. IT is now a driver of value and a key component of just about every business problem experienced. Automation is the key to deliver the speed required to compete in the Internet era.

IT in the Always-On World 20.1

With the advent of online services - such as Facebook, Netflix, and YouTube - customer expectations have changed. Any time Facebook or Twitter go down - products we do not even pay for - people lose their mind. If someone wants a new game, they can open Steam and download just about anything. 15 years ago, if someone wanted to know how far the Earth was from the Sun they had to find an encyclopedia; today they can just ask Alexa. There is no more waiting 30 minutes for a taxi to show up – Uber arrives in 5 minutes.

This is what we all expect from our consumer services. Why should anyone wait a week for IT to provision a VM? Why does it take longer to get a simple app deployed then it takes Amazon to deliver a life-sized big-foot statue? Why are free services more reliable than infrastructure that costs more than a house?

IT, always a slow and cautious bunch, need to adapt to this new world. Automation is the key to meeting these demands. Automation never types a command wrong. Automation does not sit in a queue or backlog waiting for someone with the right expertise to be available. Automation does not take vacation. Automation does not forget a step. Automation is always on.

20.1.1 From Monolith to Scale-Out

Just like Nutanix evolved massive storage arrays into small, modular blocks that you can add as needed, applications have gone through their own shift. In the past, an application might entirely run on a single box or VM (or perhaps 2 for HA). As they evolved, we broke out components. For a simple example of this, look at databases. They do not run on the box that needs the database; they run in their own cluster, own machines, under their own management.

Today these applications are even more distributed, and microservices take this even further. By taking each part of an application and breaking it out into their own machine, IT can gain incredible flexibility when it comes to patching, scaling, and growing the specific part of an application at the cost of supreme complexity as the number of machines under management grow.

Breaking an app out into 10 different components give you incredibly granular control over the application, at the cost of 100x the complexity (what version of code is running? What machines do this depend on? How can I update this application? How can I be sure I completely deleted an application?). It is the difference between updating an app on your phone and updating the phone OS.

Automation closes this gap by tracking all of this for you. No matter how many components in an application, an automated upgrade never grows in execution complexity. Automation remembers where everything is, what it depends on, and what needs to be done. Automation never forgets.

20.1.2 Allure of the Cloud for Users

We often talk about cloud adoption in terms of OPEX/CAPEX and saving by closing data centers, but why do end users create their own cloud account? In AWS, spinning up a new VM is only a few clicks (and credit card swipe) away. The AWS marketplace has over 5,000 different applications, ready to launch when you are. It does not require you to fill out a standardized form that does not really meet what you need, does not require waiting, does not require interacting with another person. One-click and you are off. How can IT compete with that? In the past, their response has been to ignore the cloud, and their users responded by adopting the cloud on their own (Shadow IT).

Automation closes this gap. Automation allows IT to provide this exact same experience to their users while at the same ensuring that corporate policy is properly applied. That security rules are followed. That IT is still in control. Automation is the key to a true cloud-like experience on-prem.

Use-Cases 20.2

The following use-cases drive design for automation and multi-cloud management tools:

- Application Marketplace or Catalog for Self-Service with centralized control across public and private clouds.
- Reducing risk of downtime or mistakes with rigorously tested automation.
- Streamline daily operations and eliminate time wasted on repetitive tasks.
- Cloud Bursting to handle unexpected surges in demand.
- Cloud Cost Visibility and Optimization.
- Centralized Financial Governance across clouds.
- Unified Cloud consumption planning to identify the most cost-effective resource to use.
- Unified Cloud Security and Compliance.
- Cloud Optionality.
- Multi-Cloud Governance.

Calm 20.3

Calm is a multi-cloud application management framework delivered by Nutanix. Calm provides application automation and lifecycle management natively integrated into the Nutanix Platform. With Calm, applications are defined via simple blueprints that can be easily created using industry standard skills and control all aspects of the application lifecycle, such as provisioning, scaling, and cleanup. Once created, a blueprint can be easily published to send users through the Nutanix Marketplace, instantly transforming a complex provisioning ticket into a simple one-click request.

Calm uses application blueprints to model details of an entire application running on the cloud. Reading that definition answers what the overarching goal of Calm is, but does not get into the deeper question - how does Calm model applications? Now we are going to dive into exactly what a blueprint is, and how Calm models application with blueprints. First, we need a new, more tangible definition of a blueprint, one that explains what it truly is.

FIGURE 13

Calm Blueprint Components

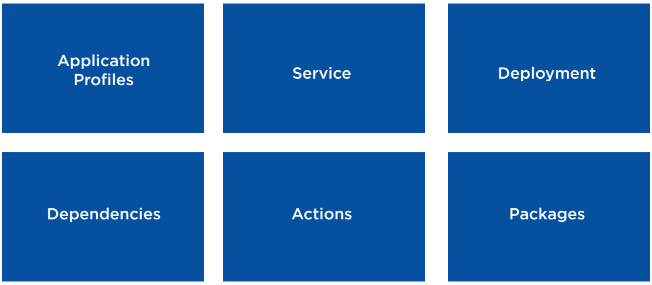

Application Profiles

Service

Deployment

Dependencies

Actions

Packages

Blueprints are Application Recipes. These recipes encompass Application Architecture and Infrastructure choices, Provisioning & Deployment steps, Application Binaries, Command steps, Monitoring endpoints, Remediation steps, Licensing & Monetization, and Policies. Every time a Blueprint is executed it gives rise to an Application.

Calm uses Services, Packages, Substrates, Deployments and Application Profiles as building blocks for a blueprint. Together they fully define applications. By encoding these into a blueprint, Calm can understand the application and properly automate the life cycle.

One-Click Application Provisioning 20.3.1

Fully automate the way you provision and scale both traditional multitiered applications and modern distributed services, using pre-integrated blueprints that make managing applications in private and public clouds extremely simple.

For example, IT managers can access the Nutanix Marketplace, choose a pre-integrated application blueprint, and, in a single click, deploy the application. Organizations can choose from a growing number of application blueprints, including Active Directory, Citrix XenDesktop, Microsoft SQL Server, and MySQL, among many others. With Calm, infrastructure teams can dramatically reduce the time it takes to provision applications, allowing them to invest more resources in driving high-value activities.

Automated Self-Service and Governance 20.3.2

Empower different groups in the organization to provision and manage their own applications, giving application owners and developers an attractive alternative to public cloud services, while elevating the role of infrastructure manager to that of a cloud operator. Nutanix Calm provides powerful, application-centric self-service capabilities, while providing role-based governance to maintain control. Administrators can limit user operations based on user role, such as IT operator, developer, or managers for approval. Additionally, Calm logs all critical activities and changes for end-to-end traceability, aiding security teams with key compliance initiatives.

For example, Calm lets you enable employees on the development team to create, scale, and destroy test and development environments without the need to file IT ticket requests. Development teams benefit from rapid provisioning times, while IT maintains control, traceability of user operations, and visibility into resource consumption.

20.3.3 Microsoft SQL Deployment Use-Case Example

Microsoft SQL Server is a common application in traditional IT organizations. This use-case example looks at the overall experience and the behind-the-scenes activity for a one-click mirrored SQL deployment.

The Calm user clicks on a blueprint in the Marketplace. Marketplace prompts the user to fill in a few fields with runtime variables to be used as part of the deployment process. For this example, Marketplace asks the requestor to choose the destination for their deployment; this destination can be either an on-prem cluster or a public cloud instance. Marketplace may ask the requestor to provide IP addresses or instance names if the preexisting blueprint does not contain automation processes for these configuration points.

After the requestor provides the required input, the blueprint begins the automated provisioning process. This example is for a mirrored SQL install, which requests and provisions a pair of VMs. These two VMs are instantiated based upon the template image approved by the creator of the blueprint. The blueprint names and assigns an IP address for the VMs either based on requestor input or by utilizing automated methods with callouts that are part of the blueprint design.

Once the VMs are prepared, the blueprint installs Microsoft SQL Server into each VM by accessing install media from a shared repository and following the configuration specifications contained within the blueprint. This process includes mirroring the SQL instances, applying best practices for SQL deployments, and assigning administrator rights for default groups and the requestor.

This one-click deployment results in a Microsoft SQL Server installation that the requestor can consume without any delay or additional effort from external teams.

Xi Beam 20.4

Many application and technology budget owners are surprised by the unexpectedly high costs of their cloud services. To prevent uncontrolled cloud spend and enable more accurate resource planning, cloud teams need better visibility of actual service consumption across all cloud environments.

Nutanix Beam is a multi-cloud cost optimization service delivered as part of the Nutanix Enterprise Cloud OS. Beam provides deep visibility into consumption patterns in a multi-cloud environment, helps with intelligent purchasing decisions and enhances security compliance of cloud resources.

Unlike other cloud expense management solutions, Beam provides a single pane of glass to optimize your cloud spend and monitor security compliance checks, along with one-click remediation. Cloud operators are empowered with intelligent planning capabilities across multiple clouds to streamline purchasing decisions based on business needs.

Cost Visibility and Optimization 20.4.1

Beam tracks cost consumption across all cloud resources at both aggregate and granular levels - per application workload, team and business unit. Beam identifies underutilized and unused cloud services and provides one-click remediation, empowering cloud operators to realize cost savings immediately and set policies to continuously maintain high levels of cloud efficiency.

20.4.2 Centralized Financial Governance

As cloud environments grow, the need to centralize control across multiple teams becomes critical. Cloud operators and business owners need a systematic way and appropriate tools to track all cloud spend and map consumption to business units. Beam visualizes resources by groups and departments, empowering cloud operators to manage their usage. Beam provides policy-based reporting and chargeback, so that teams can ensure consumption is within budget and aligns with business objectives.

20.4.3 Intelligent Consumption Planning

Cloud providers offer multiple purchasing options that can yield significant savings when utilized effectively. However, navigating the complexity of multiple options across a number of cloud accounts using variety of services can be challenging. Beam makes this planning process easy using machine intelligence and recommendation algorithms that analyze workload patterns and continuously suggest optimal purchasing decisions.

20.4.4 Cloud Security and Compliance

Beam automates cloud health checks so that you can easily monitor and ensure security compliance. You can gain insights into your multi-cloud environment based on over 250 health checks and security best practices. Beam enables continuous security management using built-in templates that certify and maintain industry standards such as PCI-DSS, HIPAA, CIS, SOC-2, NiST and ISO.

20.5 Xi Epoch

Businesses are increasingly adopting distributed application architectures with multi-cloud flexibility to foster rapid innovation. The shift from monolithic to distributed architectures has resulted in an explosion in the number of service dependencies and application health metrics from short-lived instances. Operations teams that need to ensure application uptime are also using a wide variety of languages and frameworks making it difficult to standardize on one monitoring tool. This makes it challenging to troubleshoot quickly, leading to prolonged outages.

There is an urgent need for an application monitoring service that provides visibility into the health metrics without relying on code instrumentation. Real-time, auto-discovered service dependency maps and golden signals of application health (such as latency, throughput, error rates, etc) can help to greatly reduce the average time-to-resolution by providing much needed visibility and key health metrics.

Nutanix Epoch is the observability and monitoring service for distributed applications and multi-cloud architectures. Epoch simplifies application observability with auto-generated maps that eliminate traditional requirements for code instrumentation. Operations teams are empowered with instant visibility into service interactions and continuous monitoring of service level objectives (SLOs) that truly impact end-user experience.

Epoch delivers a robust analytics engine that renders millions of data points in real-time to accelerate outage investigation. As a result, teams can quickly test failure hypotheses using sub-second queries and application drill-downs, leading to dramatically lower mean-time-to-resolution (MTTR) and increased application uptime.

20.5.1 Instant Observability

- Auto-generated application maps provide instant visibility into application health.
- Quickly generate a complex service dependency maps without code instrumentation.
- Monitor traffic flows and service interactions, not just basic metrics of individual components.
- Complete application monitoring - APIs, DNS, Databases, VMs, Containers, etc. as well as HTTPS traffic.

Nutanix Epoch leverages network as the vantage point to deliver a low-friction, framework agnostic observability and monitoring service. Live application maps generated by Epoch help to quickly figure out what part of the application is being affected. Epoch also gathers metrics for the golden signals of application health and integrates with several common protocols, such as REST, HTTP/S as well as specific ones such as DNS, MySQL, Thrift, EC2 etc., to provide complete application monitoring. With Epoch, you get instantaneous visibility into your application health without any code instrumentation. This helps to reduce the average time-to-resolution and improves application uptime.

20.5.2 Smart KPIs

Key benefits:

- Out-of-box alerting for “golden signals” such as latency, error rates, and throughput.
- Continuous monitoring of key service level objectives (SLOs) rather than individual instances.
- Reduction in “alert fatigue” with aggregated KPIs, rather than thousands of low-level notifications.

Most multi-cloud applications today are built using hundreds or more services and run on infrastructure that has short lifetime.

Legacy monitoring systems that are built to alert on low-level infrastructure issues result in increased alert noise by sending alerts on issues that do not affect end customer. With Epoch's query-centric interface, you can create custom metrics that help to set up alerts based on service level objectives (SLOs) that impact business value. You can use PagerDuty, email or webhooks integrations to send out alert notifications.

Rapid Outage Response 20.5.3

- Quickly test failure hypotheses using sub-second queries and multidimensional application drill-downs.
- Lower mean-time-to-resolution (MTTR) with real-time analytics engine that accelerates outage investigations.
- Tailored dashboards and alerts for thousands of application metrics and KPIs.
- Utilize time-travel feature to replay application performance indicators and topology changes.

Epoch comes with a powerful analytics environment that allows you to query multi-dimensional data in real-time and create custom metrics that fit your business needs. You can derive valuable insights using aggregation, transformation functions, mathematical expressions and queries on time-series data. The analytics engine in Epoch allow you to analyze metrics from past or present, to understand application changes and failure progression.

The ability to run sub-second queries and dashboards that can render millions of data points in real-time, to help you quickly get the answers you need.

20.6 References

Nutanix Calm Product Page:
https://www.nutanix.com/products/calm/

Nutanix Calm: Application Centric Automation:
https://www.nutanix.com/documents/datasheets/calm.pdf

Nutanix Beam:
https://www.nutanix.com/products/beam/

Nutanix Beam: Multi-Cloud Management and Optimization:
https://www.nutanix.com/documents/solution-briefs/sb-beam.pdf

Xi Epoch:
https://www.nutanix.com/products/epoch/

4 Golden Signals of Application Health & Performance:
https://www.nutanix.com/go/golden-signals-of-application-health.php

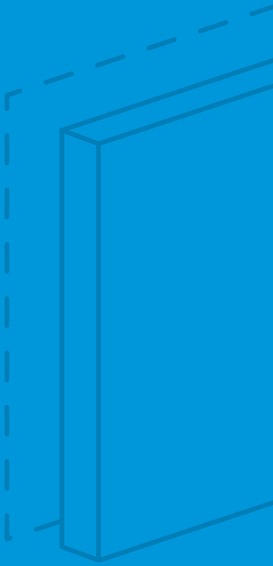

21

Era

Author: René van den Bedem

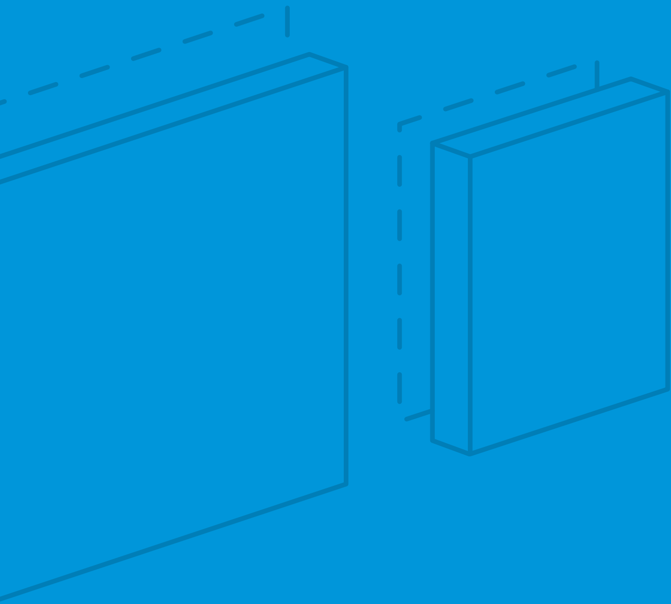

Nutanix Era is a DBaaS software suite that automates and simplifies database administration, enabling DBAs to provision, clone, refresh, and backup their databases to any point in time.

21.1 Design Considerations

- Nutanix Era supports:
  - Oracle 11.2.0.4, 12.1.0.2, 12.2.0.1 and RHEL 6.9,
  - PostgreSQL 9.x and 10.x,
  - Microsoft SQL Server 2008 R2, 2012, 2014, 2016 and 2017
  - MariaDB
- Tech Preview features: Support for Single Instance Provisioning of SQL Server, Support for Provisioning and Copy Data Management for MariaDB Database and Support for SQL Server Authentication
- Nutanix Era does not support provisioning and cloning of databases across multiple clusters. Install all the components (source database and database server, target database server, and Nutanix Era) on the same Nutanix cluster.
- Nutanix Era does not support databases running on platforms other than Nutanix. To clone databases that are running on other platforms, you must first replicate the source database to a database VM by using a tool such as Oracle Data Guard. This database VM must be running on a Nutanix platform.
- Nutanix Era supports multiple databases only on a single Microsoft SQL Server instance. Only a single database must be running on a database server. That is, you cannot create or clone multiple databases on a single database server VM in this release. This limitation applies to Oracle, PostgreSQL, and MariaDB database server VMs.

- Nutanix Era does not support database servers protected by NearSync and Metro Availability.
- Nutanix Era does not support time machine for Oracle 12c Container Databases (CDB) and Pluggable Databases (PDB).
- The Nutanix Era software is available only in English.
- The Nutanix Era software supports the source databases only in en_US.

References 21.2

Nutanix Era Product Page:
https://www.nutanix.com/products/era/

Nutanix Era Solution Brief:
https://www.nutanix.com/documents/solution-briefs/nutanix-era.pdf

Nutanix Era Version 1.0.1 Release Notes:
https://portal.nutanix.com/#/page/docs/details?targetId=Release-Notes-Nutanix-Era-v101:Release-Notes-Nutanix-Era-v101

Nutanix Era Version 1.0 User Guide:
https://portal.nutanix.com/#/page/docs/details?targetId=Nutanix-Era-User-Guide-v10:Nutanix-Era-User-Guide-v10

22

Karbon

Author: René van den Bedem

Nutanix Karbon, formerly known as Acropolis Container Services or ACS, is a curated turnkey offering that provides simplified provisioning and operations of Kubernetes clusters. Kubernetes is an open-source container orchestration system for deploying and managing container-based applications.

Karbon leverages the CentOS and Ubuntu Linux-based operating systems for Karbon-enabled Kubernetes cluster node creation. Linux containers provide the flexibility to deploy applications in different environments with consistent results.

The Karbon web console simplifies the deployment and management of Kubernetes clusters with a simple GUI and built-in event monitoring tools. Kibana, the built-in add-on, lets you filter and parse logs for systems, pods, and VMs. Karbon also leverages Pulse, Prism's health-monitoring system, which interacts with Nutanix Support to expedite cluster issue resolutions.

22.1 Design Considerations

- Nutanix Karbon version 0.8 is currently in Technical Preview and should not be used for production systems.
- Must use Prism Central 5.9 or later. Multi-node Prism Central is not currently supported.
- Must use Prism Element 5.6.2.x, 5.8.2.x, 5.9.x or later.
- Must use the Nutanix AHV hypervisor.
- Cluster Virtual IP and iSCSI Data Services IP addresses must be configured.
- Cluster must be registered with Prism Central.
- Cluster and Prism Central time zones must be synchronized.
- NTP and DNS must be configured.

- IPAM or DHCP enabled network with Internet access required.
- If a Web Proxy is used, the following domains must be whitelisted: hub.docker.com, gcr.io, k8s.gcr.io, quay.io & docker.elastic.io.

References 22.2

Karbon Product Page:
https://www.nutanix.com/products/karbon/

Nutanix Karbon: Enterprise-grade Kubernetes Solution:
https://www.nutanix.com/2018/11/27/nutanix-karbon-enterprise-grade-kubernetes-solution/

Karbon Community Forum:
https://next.nutanix.com/kubernetes-containers-30

23

Acropolis Security & Flow

Author: Neil Ashworth

Thinking back, over the history of how vendors address security concerns, there has been a change; we have evolved from a preventative, isolationist strategy, transitioning into prevention by means of Detection and Response. Detection and response, solely as a means to mitigate the exploitation of potential vulnerabilities, is now proving ineffective with, as of late, the shifting types of vulnerabilities we've seen emerging. From a few years ago where we were predominantly seeing vulnerabilities in applications such as Apache or Java, we are now seeing more sophisticated exploits coming out, affecting more difficult areas of the data center to solution and impacting those areas of the data center that we've trusted for years.

These new areas of exploitation are some of the most difficult to mitigate, due, in part, to their lack of transparency, reliance by vendors on microcode fixes which can, when implemented, heavily impact performance and the dark nature of processor technology which is facilitated by a general lack of visibility in the industry. It is not like open source where you have an ecosystem of support looking at the product, looking at code; processors are closed and subsequently, much more difficult to identify these types of vulnerabilities, mitigate them and strategize around safely reducing their threat to the environment.

All of these Side-channel exploitations, (Spectre, Meltdown, L1TF etc.) and BMC exploitations that we saw in 2018, are a few examples of the new avenues of attack that can be leveraged by an attacker, which in most cases can negate the traditional framework for prevention purely by means of detection and response.

Attackers leverage weaknesses in an organization to gain access, and those weaknesses are sometimes found in failures to properly secure or harden the IT fabric, the network, the endpoints, servers and cloud. This process of system and security hardening organizational IT, when conducted manually and in isolation for each of the various fabrics, often belies the goal of achieving adequate Operational Security (OpSec). Achieving a robust security posture in heterogeneous environments, each with their own operating systems, kernels, firmware and management control is immensely complex - complexity breeds inefficiencies, and inconsistencies and those lead to vulnerabilities.

The goal then, should be to first assist Organizations in their effort to achieve good OpSec should be to first, attempt to remove some of that complexity, and second, where possible, adopt automation in order to reduce the capacity for human errors in configuration that undermine the efforts of an otherwise well-rounded security posture.

Security Development Lifecycle (SecDL) 23.1

"An ounce of prevention is worth a pound of cure."

- Benjamin Franklin

Rarely do Organizations ask of their vendors the processes by which they create their innovations, or integrate acquired technologies into the existing framework of the platform or product they hope to peddle. It is important to understand the ethos vendors have surrounding code development, as that can also be a potential attack vector. Do the development lifecycles of the vendor meet a repeatable and predictable security baseline in the product?

Is the vendor building a product that has a specific set of criteria or controls and requirements which are delivered in a predictable manner? Are they building a product with a framework wrapped around it that speaks well to your compliance needs, that speaks well to your cybersecurity posture? Whatever that may be.

When you look at the last few years and how cyber has become much more prevalent, in not just intelligence and federal communities, but public and private sector communities; Home Depot, Sony, Target, Experian, the DNC, with each breech seemingly more damaging than the last, providing endless media junkets and sound bites. The impact is not only measured in dollars and revenue lost, but in reputation and public perception.

It is with this mindset then that Nutanix believes vendors, such as ourselves, should be following responsible production guidelines. Hardening the code for a start, is something we can no longer afford to be complacent about. The practice at Nutanix is to wrap all this together, driven by our full stack Security Development Lifecycle (SecDL). Not to be confused by similarly named, and very well known SDLC process for systems and software engineering. The Nutanix SecDL process begins by forcing developers, QA and Test to all operate in the same locked down, hardened environment that have our full Information Assurance (IA) posture placed upon them.

This ensures the creation of known good within our stringent security control set. It continues with security best practices, keeping security as principal to product development, such as; the regular use of code and system vulnerability scanners, removal of superseded or superfluous code, thorough testing for interoperability, automated testing, and using an Agile delivery method which allows us to reduce the capacity for Zero-day threats to be exploited within our environment.

FIGURE 14

Security Development Lifecycle

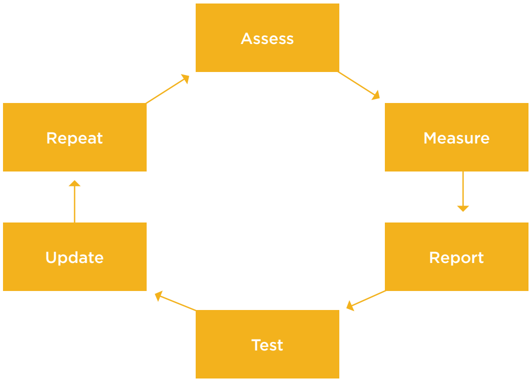

Another consideration when assessing vendor solutions is, are they thinking Cloud centric or Silo centric? This world of Public, Private, Edge cloud and Hybrid cloud models requires that we look at things from an application perspective, because at the end of the day when you look at what your end users are using, they are not using hypervisors, or storage protocols, or VLAN segments, they are using Apps, so why then are we spending so much time with network administrators and system administrators in order to deploy these applications? Why are the Apps not configuring the Firewall appropriately? Why are Apps not establishing if it is a public or private network that is needed to attach this application to? Why is the application not driving the ownership and permissions models? These are all questions vendors should be asking to help drive a different sort of philosophy, one that Nutanix is actively driving towards.

This philosophy is especially prevalent when you speak with application developers - the challenge is changing that mentality that the infrastructure, from an invisible perspective, has to be driven by the apps and not limited by the individual components.

A final thought on the principles of software development that Nutanix holds true, and it is to expect that external dependencies will be compromised. There are many technologies that make up the modern data center, components that make up your management tiers, and network tiers, and it is essential that we do our part as responsible vendors to ensure if there is an attack or compromise to a particular piece of technology that we are making sure that, this does not impact and effect the rest of the infrastructure.

23.2 Security Technical Implementation Guide (STIG)

> “What goes into it, effects what comes out of it.”
>
> **– Unknown**

When Organizations attempt to solve a problem they might be facing, they might invite various vendors to pitch their solutions which more or less meet some, or perhaps all, of the requirements they’ve set. After purchasing the solution that Organization could spend weeks, or even months, in Security Operations (SecOps) getting that product rolled out across their environment and meeting the appropriate IA posture necessary for the environment. This process can be expensive, time consuming, and as we've discussed, prone to potential human configuration errors.

Ultimately, what this boils down to is this; does the software vendor that is writing this product that you are using to solve a problem, understand how you plan to use it? And this is from a security perspective. Does the vendor you are speaking to actually understand the vertical that you are in? The compliance requirements that you have and the controls that you may have to adhere to?

Another way a vendor can make your life easier in this regard is best practices and standards baked into the product. Nutanix calls it the intrinsic method. Understanding what it is you need to do to the product and then baking it into the product. Instead of writing a product generally for the masses and then creating documentation and procedures for you to use on your own later, with your own resources and dollars, why not just understand those requirements and develop a product that has all that good to go? Baked it into the system and ship it to the customer in a hardened configuration state.

The way we deliver this intrinsic method of a secure configuration control set on Nutanix platform, is by way of Security Technical Implementation Guides (STIGs). Nutanix STIGs are based on common National Institute of Standards and Technology (NIST) standards that can be applied to multiple baseline requirements for the DoD and are equally prevalent for frameworks such as PCI-DSS, HIPAA, CIS etc.

The comprehensive STIGs are written in eXtensible Configuration Checklist Description Format (XCCDF) in support of the Security Content Automation Protocol (SCAP) standard. This machine-readable STIG format automates assessment tools and eliminates time-consuming testing. Because the STIGs are machine-readable, they are ideal candidates for third-party apps that probe for deficiencies in a system configuration.

Note: The XCCDF XML format is highly efficient for conversion from a manual process to machine automation. Designed specifically to meet the SCAP standard, the XML format is future-proof, in that it supports the transition to DoD DIARMF (Risk Management Framework) for continuous monitoring. Any third-party system that understands XCCDF XML style formatting can consume the STIGs.

One of biggest benefits from using machine-readable STIGs to perform system and security hardening, is time to accreditation. Previously, it could take countless hours to manually check files or find obscure settings. Even worse, administrators had to track any aspects that could not be automated in static spreadsheets. As a result of automating these testing tasks, the accreditation process time for the DoD Information Assurance Certification and Accreditation Process (DIACAP) has been shortened from as long as a year to less than half an hour. This speed allows you to dynamically check an ever-changing baseline.

23.3 Security Configuration Management Automation (SCMA)

> “ Automation is good, so long as you know exactly where to put the machine.”
>
> **– Eliyahu Goldratt**

Securing the code, then establishing a secure configuration framework, is a sound security strategy. Making sure that the code, at every layer, has hardening applied to each component and every layer that we are providing within the data center is vigorously secure and configured correctly is, to us, a glass half full scenario. This is because all of that good stuff we have completed so far is a point in time evolution. It is an individual snapshot from a given point in time of when you (the customer) have conducted a particular audit or produced a compliance baseline. Organizations put these systems into production for between 3-5-year life cycles, in most cases, so how can Nutanix help you keep and preserve that security and that IA posture, over that lifecycle?

Acropolis framework has what we call a Self-healing, or Self-remediating capability. By leveraging the power of configuration management automation, we give customers the opportunity to run on an hourly, daily, weekly, or monthly basis. It is a backend configuration management automation framework that checks to make sure that all of the IA components that we've embedded in our code remain compliant.

SCMA is a Saltstack configuration daemon that runs periodically to address what we call drift. Drift can happen throughout the environment for any number of reasons, admins adjusting settings temporarily and forgetting to revert them, a software patch, or perhaps even a bad actor lessening controls. These deviations are identified by SCMA logged and reverted to the secure configuration state that we support and provide with the Nutanix platform.

With Nutanix SCMA, organizations can alleviates that point in time compliance story and turn it into a continuous monitoring discussion, where Nutanix is continuously monitoring all those controls from Day 1 through Day 365.

Note: By default, SCMA runs daily, for organizations that are willing to accept the performance impact, you can change the SCMA setting to hourly in the aCLI. Guidelines are available on the Nutanix portal.

23.3.1 Authorization and Authentication

Controlling who can do what, who can see what, is a cornerstone of security on any platform. Prism supports three distinct authentication methods:

- Local user authentication
- Using a Directory Service for authentication, such as Active Directory or OpenLDAP.
- And Security Assertion Markup Language (SAML) authentication. Users can authenticate through a qualified Identify Provider (IdP) such as okta or ADFS, when SAML support is enabled for Prism Central.

Local user authentication, although an option on Nutanix platform, is not recommend for wider scale use. It is deemed best to restrict this type of authentication for limited cases, such as initial configuration, and account recovery.

The most accepted method of authentication on Nutanix platform, is using a directory service such as Active Directory or OpenLDAP. The benefits of this are obvious - a secure, centralized repository of user credentials, one place to authenticate against and one place to manage password policy, down to the management interface.

SAML is an open standard for exchanging authentication and authorization data between two parties, a Service Provider (SP), in this case being Prism Central (PC), and an Identity Provider (IdP) which creates, maintains and manages identity information. SAML can also enable enhanced security functions like Multi-Factor Authentication.

Within any Enterprise cloud platform the necessity for Role Based Access Control (RBAC) is apparent. Separation of duties is the concept of having more than one person required to complete a task. In business, the separation by sharing of more than one individual in one single task, is an internal control intended to prevent fraud and error. It is also often a security compliance requirement. Nutanix enables a form of RBAC via Self-Service Portal (SSP), SSP gives customers the capacity to build attribute centric controls to users or groups. The administrator sets up a “Project” with all the resources that may be needed to run and manage the Project VMs, including networks for your VMs, images with which to create VMs, and user permissions. The administrator can then invite users to the self-service portal, whereupon they can log on to use and manage the assigned project VMs and allocated resources.

This method of access control is more fine-grained, which allows for more input variables into an access control decision. Attribute Based Access Control (ABAC), seen in SSP, can be used by an administrator to set available attributes by themselves, or in combination to define the right filter for controlling resource access. ABAC is both more flexible and more secure than RBAC, and can control access based on differing attribute types, such as: Subject attributes, System attributes or Environmental attributes.

Encryption 23.3.2

Making information indecipherable as a means to protect it from falling into the wrong hands is not anything new. As far back as 600 BCE the ancient Spartans use a device called a scytale to send secret messages during battle. Modern cryptography uses an algorithm, a mathematical cipher to encrypt or decrypt data, turning plaintext into ciphertext.

Today Nutanix offers three methods of encrypting your data-at-rest, which helps us compete in security sensitive markets, such as US Federal, Healthcare and Financial:

- The first method is via Hardware with the use of Self Encrypted Drives (SEDs). Key Encryption Key management is done via an External Key Manager (EKM) sometimes referred to as a Key Management Server (KMS). Our system treats data in an encrypted system much the same way as it treats data in a non-encrypted system. Encryption happens when data lands on SEDs. When a client reads data from SEDs, non-encrypted data is returned.
- The second method is Software driven, and this happens natively within our AOS software stack. Key management is handled entirely by AOS as a Local Key Manager (LKM). In a Nutanix cluster, regardless of the number of nodes, each node runs a standby instance of every service necessary for a cluster to operate. This ensures you have a highly resilient and available service to your end user. The LKM is structured in the same manner, where each node can function as the LKM for the entire cluster. Each node functioning as an LKM for the entire cluster gives Nutanix the ability to ensure your data remains available. The advantages are obvious:

    a. No premium for SEDs required.

    b. No delayed delivery due to often higher lead times of SEDs.

    c. More media choices - SEDs are often available on a select set of drives only. Doing it in SW ensures that even customers needing DAR Encryption have all the available media offerings from Nutanix.

    d. Time to Market - Customers can take advantage of the latest HDDs/SSDs available from Nutanix, without waiting for the SED equivalent SKU to be available.

    e. No premium of External Key Manager

- The third method is an amalgam of the previously mentioned two. Dual encryption using both SED's and SW Data-At-Rest Encryption (DARE). This method requires the use of an EKM for Key management.

Since the second method allows for encryption without yet another silo to manage, customers looking to simplify their infrastructure operations can now have one-click infrastructure for their key manager as well. Key management and properly storing secret keying material is the centerpiece of the Nutanix design. In the LKM we use a mathematical method referred to as Shamir Key Splitting. This allows us to securely store only portions of each private key per node, requiring a quorum of nodes to be present in order to reassemble the key and decrypt the data. This ensures that drive theft, and node theft, are covered use-cases.

Data is encrypted using a data encryption key (DEK). The native LKM service uses the FIPS 140 Crypto module to keep all the DEKs safe. No separate VMs are needed to support the native LKM. Every storage container has its own DEK, which is typically then encrypted by a key encryption key (KEK) that is sent to an EKM. Now that Nutanix supports its own native LKM, Nutanix also takes the KEK and wraps it with a 256-bit encryption key called the machine encryption key (MEK). The MEK is distributed among the CVMs in the cluster using the Shamir splitting algorithm.

Since the MEK is shared, each node can read what other nodes have written. In order to reconstruct the keys, a majority of the nodes need to be present. We use the equation K = Ceiling (N / 2) to determine how many nodes are required for the majority. For example, in an 11-node cluster (N = 11), we would need 6 nodes online to decrypt the data.

FIGURE 15

EKM & LKM Workflows

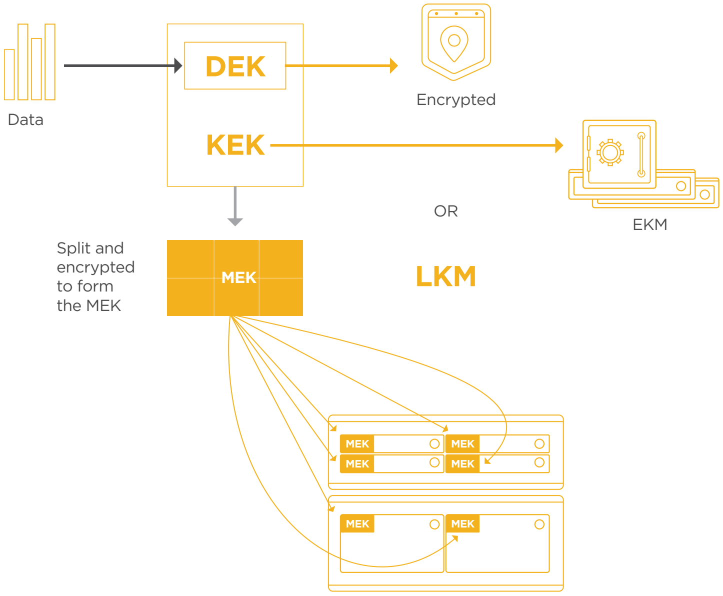

Backing up keys with Nutanix and Prism is also seamless. Each storage container has a DEK, so when a new storage container is created, an alert is generated encouraging administrators to make a backup. The backup is password protected and should be securely stored. With the backup in hand, if a catastrophic event happens in your data center, you can replicate the data and re import the backup keys to get your environment up and running.

Software DAR Encryption (SWDARE) uses the Intel AES New Instructions (NI) encryption instruction set, improving upon the AES algorithm and accelerates data encryption. Supporting AES

NI in software gives customers flexibility across hardware models, while reducing CPU overhead. The default encryption setting is AES-256.

For customers potentially concerned at this point that SWDARE is an AHV only product, you can rest assured. SWDARE is available across Hyper-V, ESXi, and AHV for x86 platforms. For ESXi and Hyper-V, software DARE operates at the storage container level, and you can move data from unencrypted to encrypted containers. Container-level encryption must be turned on when the container is created. With ESXi, Hyper-V, and AHV, you can also decide to encrypt the entire cluster.

Encryption is often a compliance requirement for organizations handling sensitive data. This requirement can often be a burden to implement. Nutanix has taken away the complexity surrounding this often daunting process, replacing it with a joyful experience that meets or exceeds all the requirements you may have to meet in those security compliance frameworks, and giving piece of mind that data-at-rest is properly and effectively secured.

23.4 Micro-Segmentation with Flow

Network Security is big, complicated, requires specialist training, and can often necessitate years of experience to deploy in enterprise organization. On top of that, in order to architect appropriate solutions for isolating environments, it takes careful planning, resources and time. Re-architecting an environment can be even more problematic if not precisely carried out. Many network engineers might jokingly state “it is always the network,” yet they know the consequences of failure being; potentially breaking existing applications and functionality, or worse, exposing network vulnerabilities.

Networking is made even more complicated in modern data centers using virtualized environments and building applications for both on-prem and cloud deployments.

In the legacy data center, external traffic bound for the database, as indicated by the green line in Fig 3, is only filtered via a perimeter firewall - this is considered North - South traffic. In the virtual data center, the attempt of East - West traffic of the two VM's passing data, as indicated by the blue line in the figure below, to be properly

FIGURE 16

Legacy Traffic Flow

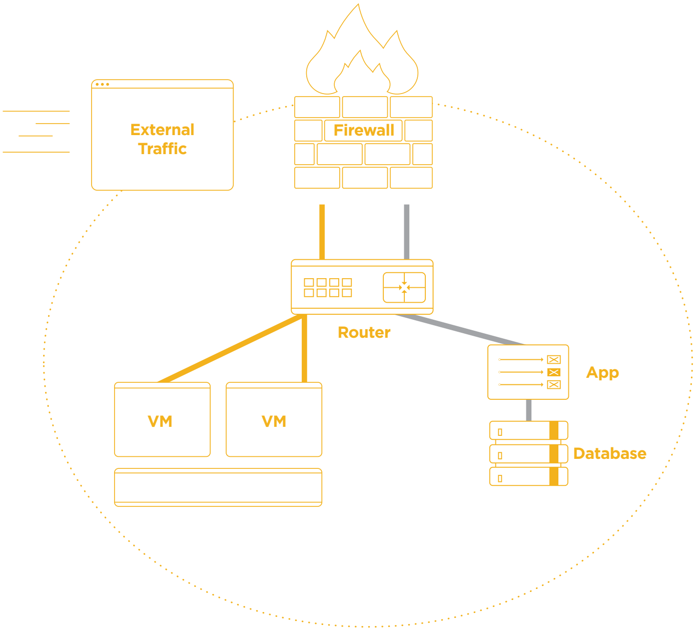

inspected by the perimeter firewall, creates a hair-pinning effect.

The issues faced with this networking approach is that once an intruder has breached the perimeter firewall they effectively have free reign to move laterally throughout the environment. An admin could introduce additional firewalls within the data center to inspect East - West traffic. However, given the dynamic nature of virtualization, probable subnet/vlan migrations, and with the dawn of the cloud the potential to migrate platforms, this approach would be very costly, add latency to the data flow, and be extremely difficult to manage numerous 5-tuple firewall rulesets individually.

The next evolution in enterprise IT was Network Virtualization and API programmable switches. This allowed for physical hardware to become more responsive to application needs dynamically. Whereas previously it could take days, or sometimes weeks, to provision the necessary network framework for your new application, these new, Software Defined Networks could be provisioned in minutes. Physical networking, just like physical servers, was an IT bottleneck and virtualization, again, was the resolution.

Nutanix was a pioneer in the space of hyper-converged infrastructure (HCI). Customers loved how we captured web-scale principals to build a platform free from the legacy thinking, complicated SAN arrays, NAS, HBAs, storage network switches, siloed infrastructure, all replaced with commodity servers with directly attached storage running our intuitive, intelligent and delightful software. Those same customers asked many times if we could direct our attention to the emerging capabilities brought about through Network Virtualization, and we did.

When some companies “innovate” in an attempt to take advantage of new capabilities, they are sometimes guilty of simply replacing old concepts with modern interpretations, like a modern cover of a classic song. They are not new or fresh ideas, they are usually decent at achieving a nostalgic semblance of the old method, but not exactly forward thinking. When Nutanix approached the subject of integrating some network capabilities within the platform, we first had to understand a few things: What do customers want from Software Defined Networking? And how can this support the way customers will want to build applications in the future?

Addressing the latter question first: Consider for a moment the rapid adoption of cloud over the past decade. Organizations are rallying to cloud platforms in droves, AWS is alone responsible for more than 30% of the market, riding a remarkable $10 billion dollar run rate. The reasons for clouds mass adoption is speed, simplicity, cost and reaping the benefits of continued innovation. Simply put, customers can quickly realize the benefits of their applications without the legacy burdens of old, such as archaic cumbersome processes, and engaging with multiple stakeholders, (i.e. the Network guy, the Storage guy, the Database guy, the Server guy, the virtualization guy etc.).

Rather, they simply swipe a credit card, build their application and meet the business need. They do not care about the underlying infrastructure because the application is what the end users are touching, the application is what generates revenue. Given this information then, in the cloud centric world, why are we not allowing the application itself to manage some of its own infrastructure and security needs? The key is to innovate solutions around cloud-centric application needs, rather than outdated silo-centric infrastructure capabilities.

To address the former question, “What do customers want from Software Defined Networking?” Simple query and analysis is used to determine the prevailing use-case, it being micro-segmentation. Also, given what we learned earlier, in that the focus is now on the application, it should be no surprise that a network overlay as an abstract representation of the entire physical network into a virtual layer, is not necessary. Nutanix focussed on delivering the networking attributes that matter: Application uptime, security, visibility and automation. A network overlay is not an efficient means of providing these capabilities.

Nutanix Flow is a dynamic, policy driven, intelligent way to create secure zones for Virtual Machines and in the future, Containers running within Nutanix Enterprise Cloud. Our realization of micro-segmentation in Flow, provides granular control and governance of all traffic into and out of a VM or groups of VMs. It ensures that only permitted traffic, between application tiers or other logical boundaries, is allowed and protects against advanced threats propagating within the virtual environment.

But more than that, the experience of setting up and operating Flow is, as you may have come to expect with Nutanix, delightful, intuitive and simple. In developing security policies for Flow, the virtual machine or the container is the first-class citizen. Less concern is placed on a subnet, VLAN, or even a specific IP address.

All of this is native to our AHV virtual networking and is based upon Open Virtual Switch (OvS). There is no additional software or controller to install to leverage the functions of Flow. To achieve

FIGURE 17

Flow OvS Bridge Chain (Micro-Segmentation)

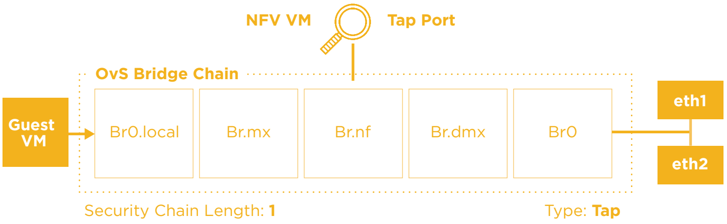

micro-segmentation, we leverage the distributed firewall built into the AHV Open vSwitch as seen in the figure below.

For environments requiring additional functionality such as Application based Firewalls, network threat detection (ie IPS/IDS), or general application network diagnostics, Nutanix utilizes service chaining. These services are inserted in-line, and can be easily enabled for all traffic, or deployed only for specific network traffic. With the ability to redirect only VM traffic on certain ports, Flow can also reserve the resources of more expensive virtual appliances.

To help you, the reader, connect to the capabilities of Flow and Micro-Segmentation, let us go through the process of isolating an application in Development (Dev) environment from the same application in a Production (Prod) environment. This is all achieved in Prism Central - no “installation” is required. You only need license the feature and enable via the dashboard. There is no upgrading of the environment to make it “Flow-ready.” On day zero, administrators can begin configuring policies.

All Flow security policies are constructed utilizing categories. Categories are a text-based method of organizing VMs into groups as they relate to function, location, environment, etc. A number of predefined categories and category types exist, and administrators can create their own categories with just a few clicks. Of the predefined categories types, the Environment, AppTier, and AppType are the most prevalent when implementing micro-segmentation with Flow.

Once categories are created, administrators then apply them to Virtual Machines. For example, in Figure 5, the “ex-Mbox” virtual machine exists in the Production Environment

(Environment:Production), is a part of the Application Exchange (AppType:Exchange), and within the Exchange Application a part of the Application Tier Exchange_Mbox (AppTier:Exchange_Mbox).

FIGURE 18

Categories Applied to a Virtual Machine

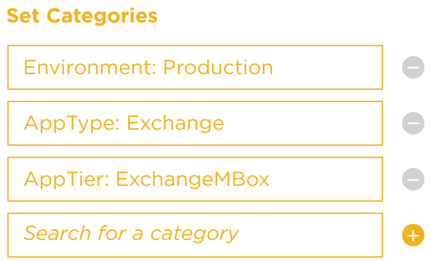

Set Categories

Environment: Production ⊖

AppType: Exchange ⊖

AppTier: ExchangeMBox ⊖

Search for a category ⊕

Extrapolating this out for your environment, you can quickly build layered categories all your applications across environments.

Once categories are set and applied to all the relevant VMs, an administrator can begin constructing policies. Flow maintains three different policy types: Quarantine, Isolation, and Application.

In a scenario where admins are required to completely segment traffic flow between Production and Development environments, an Isolation proves very beneficial. In order to create an Isolation Policy, an administrator needs only select the two Categories he/she wishes to completely segment (see Fig 6). In the case of the

FIGURE 19

Flow Isolation Policy

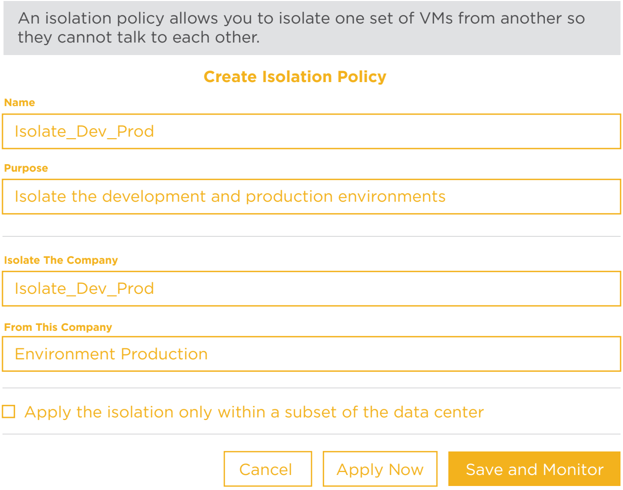

An isolation policy allows you to isolate one set of VMs from another so they cannot talk to each other.

Create Isolation Policy

Name: Isolate_Dev_Prod

Purpose: Isolate the development and production environments

Isolate The Company: Isolate_Dev_Prod

From This Company: Environment Production

□ Apply the isolation only within a subset of the data center

policy defined in Figure 6, all traffic between virtual machines with the category “Environment:Dev” applied, and all virtual machines with the category “Environment:Prod” applied, will be denied.

Another key benefit in Flow is the ability to monitor a policy prior to applying it. Before completely applying a new policy, and accidently denying critical traffic, an administrator place the policy in “Save and Monitor” to review the traffic flow. Once a new policy is saved,

FIGURE 20

Flow Policy Visual

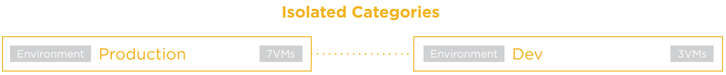

an admin can select the policy again to view it, and by hovering a mouse over the traffic flow visual, identify the quantity of violations to this policy (measured in bytes & "flows") over the past 24hrs. Applying the policy will prevent this traffic flow from occurring.

Hopefully this brief example has impressed upon you the power of simplicity embedded within this design. Note how we did not have to carefully plan IP addressing, or VLAN allocation, we did not have to think about VLAN IDs, ALCs or subnets. Quite instinctively and naturally we identified the two environments we wanted to isolate, assigned those environments categories and effected segmentation of those environments with a simple policy, which will now maintain those dynamic environments automatically as VMs are added and deleted.

23.5

References

Flow Product Page:
https://www.nutanix.com/products/flow/

Datasheet: Nutanix Flow:
https://www.nutanix.com/documents/datasheets/nutanix-flow.pdf

Application Centric Security with Nutanix Flow:
https://www.nutanix.com/go/application-centric-security-with-nutanix-flow.php

Acropolis Security:
https://www.nutanix.com/products/acropolis/security/

Tech Note Information Security with Nutanix:
https://www.nutanix.com/go/information-security-with-nutanix.php

24

Files

Author: Wayne Conrad

Nutanix Files, formerly known as Acropolis File Services, or AFS, is the Nutanix way to create standard SMB and NFS file shares, to replace traditional file servers or NAS filers.

The traditional use-case for Files was VDI, but now with NFS support and various improvements, Files is now ready to tackle many of your traditional use-cases.

24.1 Use-Cases

Nutanix Files provides file shares in the two most common network file access protocols, Microsoft Windows SMB and Linux/Unix NFS. Nutanix Files evolved to provide for VDI profile data for what was our most common use-case, VDI, without having to create file shares or buy NAS filers from 3rd parties.

Nutanix Files has evolved rapidly since its launch two years ago, scaling in performance, features and share size so that it can now take on general purpose file shares, or some application level transactional storage for workloads, like containers.

New Features in Files:

- SMB v3 and NFS v3 support (Note, check documentation for specific features of the protocols)
- SMB / NFS hybrid shares with access via both protocols, allowing your *nix and Windows Servers to easily share files
- Massive performance improvements, up to 20K user home directories tested

Design Considerations

24.2

- Files has two share types, General or Home Directory. Home Directory shares have all subdirectories automatically sharded on to different Files VMs. General shares do not.
- 2,000 connection per File Server VM at the maximum size means you should carefully plan the use of General vs Home Shares to keep the maximum connections under the maximum supported size.
- Poorly built or buggy applications may keep multiple connections open per user, swamping the Files VMs. Always check the number of connections on existing file servers or NAS filers during the design phase.
- Migration of file shares is significantly simpler with the use of Windows DFS-Namespace (DFS-N) to abstract away the underlying shares from the links, mappings and file share names that end users remember. If you are not using a DFS-N in front of your file shares today, that should be your first planned exercise during the migration effort.
- Fully qualified domain names instead of hostnames should be used for all shares in all cases of mapped drives, links, etc. Windows and some applications may use NTLM style authentication on SMB with hostnames, significantly slowing performance and increasing chatter. NTLM style authentication vs Kerberos style authentication involves significantly more chatter over the network and is less secure. One of the first steps of performance troubleshooting is to ensure you are not using NTLM anywhere.

- Files uses Volumes, and thus has a 60-minute minimum RPO and does not support synchronous replication.
- If you need synchronous active / active Files across sites, consider the use of Peerlink by our partner Peer Software.

Files can be scaled out to more VMs, or scaled up to bigger VMs, or a combination of both with the following trade-offs:

- Scale up increases are 100% non-disruptive, as hot CPU and memory are added to VMs. Scaling out adds additional VMs, and shares are moved. This should not disrupt access to any files except files that are continually accessed via an open handle.
- Scaled out VMs have more overhead, as each VM has an OS and various other processes that take CPU and memory overhead.
- Scaled out VMs have more exposure to the risk of hardware failure. Scaled up VMs have a larger failure domain that would affect more users if there was a failure.

References

24.3

Nutanix Files Product Page:

www.nutanix.com/products/files/

Nutanix Files Datasheet:

https://www.nutanix.com/documents/datasheets/nutanix-files-ds.pdf

White Paper: Reimagine File Services With Nutanix Files:

https://www.nutanix.com/go/reimagine-file-services-with-nutanix-files.html

Transform File Storage with Nutanix Files:

https://www.nutanix.com/go/transform-file-storage-with-nutanix-files.php

25

Volumes

Author: Wayne Conrad

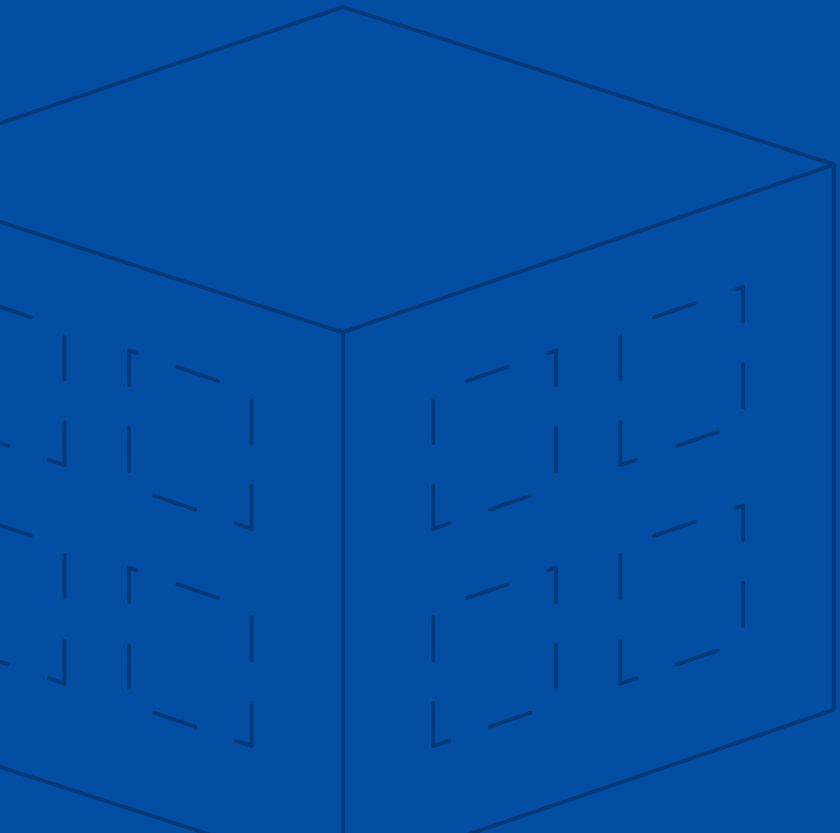

Volumes, formerly known as Acropolis Block Services, or ABS, also known as Volume Groups, allows presenting Nutanix storage as iSCSI disks to VMs inside the guest OS or to physical servers, rather than traditionally presenting them as disks at the hypervisor layer.

Why would you want to do this? The answer is simple. Shared disks at the hypervisor layer for clustered workloads like databases have always been painful and complex to setup.

Volumes is designed to make shared storage for Microsoft Clustering, Oracle RAC, and other shared disk solutions much simpler. It also supports using your Nutanix storage to run your physical database servers, or big iron Unix or mainframe boxes.

25.1 Use-Cases

Nutanix Volumes are iSCSI disks attached inside the OS versus at the hypervisor layer. This avoids the worst problems associated with VMware raw device mappings and allows shared disks without pain and suffering. Nutanix volumes also supports physical servers, with a large range of supported OSes. One use-case that is NOT supported however is attaching Nutanix volumes as VM storage to non-Nutanix hypervisor hosts. Nutanix has not built the vSphere, Hyper-V or other hypervisor level plugins required to make Volumes work for hosting VMs.

Design Considerations 25.2

- iSCSI generally does best over layer 2, not layer 3, so placing the VMs on the same VLAN as the CVMs is recommended.
- iSCSI seriously benefits from jumbo frames, and is the one plausible use-case for jumbo frames on the Nutanix platform
- The volume group load balancer is a double-edged sword for VMs running on Nutanix. While it massively increases throughput by allowing every CVM to equally participate, by design it also eliminates data locality increasing read latencies.
- Volume groups are currently only capable of 60-minute RPO replication or synchronous replication.
- Many or most Linux guests do not enable SCSI unmap commands by default, which can inflate the size of disks over time. Ensure that all attached OSes support trim and have it enabled for all Volumes.

References 25.3

Nutanix Volumes Product Page:
https://www.nutanix.com/products/acropolis/block-services/

Nutanix Volumes Datasheet:
https://www.nutanix.com/documents/datasheets/DS_ABS_web.pdf

Best Practices Guide: Nutanix Volumes:
https://www.nutanix.com/go/nutanix-volumes-scale-out-storage.php

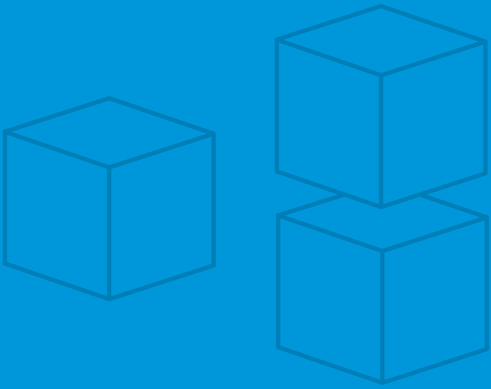

26

Objects

Author: Laura Jordana

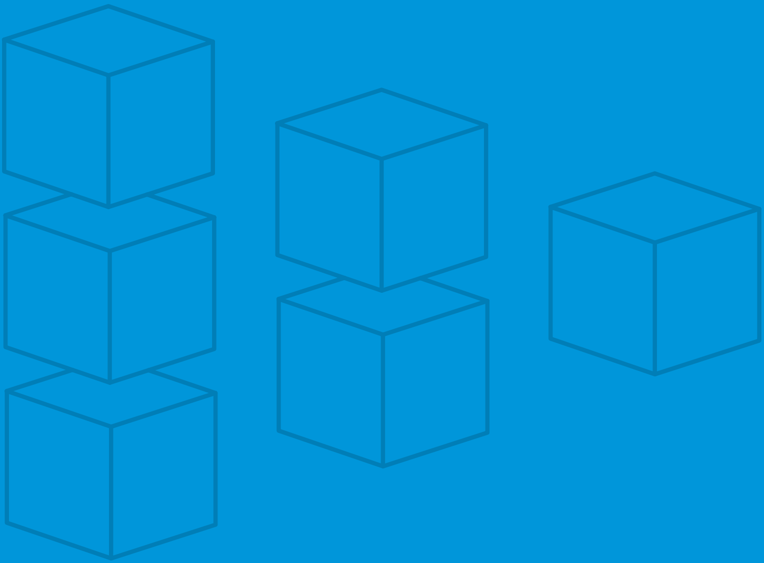

Nutanix Objects allows the creation of S3 API object storage. Over the last 10 years, the Amazon S3 service has driven the popularity of API accessible file storage and become the de facto standard API for object storage. Nutanix Objects allows for the creation of multi-petabyte scale S3 storage to run applications that need the S3 API on in-house and more excitingly, hybrid cloud storage where some storage may live on prem and some live in the cloud.

The amount of data being created is growing at an exponential rate. It has been said that the majority of data that has ever been created was created in the past few years. With the rise of IoT, sensors, and other machine-generated data, this number will only continue to increase. According to the IDC, by 2020 the total data volume will be more than 40 zettabytes, with almost 63% of that data being unstructured.

Nutanix Objects is an S3-compatible scale-out object storage solution for storing this unstructured data. Being built on the Nutanix Distributed Storage Fabric allows Objects to take advantage of existing storage features within the Nutanix software, such as encryption, compression, and erasure coding, as well as the built-in resiliency and scalability that is required of any cloud platform.

26.1 Use-Cases

26.1.1 DevOps

The way IT departments deploy applications is rapidly changing. With the emergence of containers and other cloud native technologies, users need a scalable and resilient storage stack which is optimized for the world of cloud computing. Nutanix Objects was built for cloud native environments and is the optimal solution for next generation applications that could be running anywhere, whether on-prem or in the cloud.

DevOps is the integration between Development and Operations to unify software development and software operations. For DevOps engineers, automation is key when working with object storage. Nutanix Objects provides important features required by DevOps including:

- **Single global namespace** - for collaboration across regions for engineering teams spread around the world.
- **S3 support** - S3 is an industry standard and widely used API with a very well documented interface. Many DevOps engineers are already using S3 in their scripts, so much of their existing code could be reused.
- **Performance** - time-to-first-byte of 10ms or less.

Long Term Data Retention 26.1.2

Depending on the industry, users may have to comply with state and federal laws and follow regulations that dictate how long they must keep data and what type of storage the data has to reside on. Some of the features that Nutanix Objects provides to help users meet regulatory compliance:

- WORM (write once, read many) compliance demands that data cannot be changed or altered. When policies are applied to entities such as buckets, objects or tags, this prevents data from being changed, tampered with or deleted. Nutanix Objects features S3 support for WORM on a bucket.
- Object versioning allows the upload of new versions of the same object for required changes, without losing the original data.
- Lifecycle policies dictate when old objects and versions should be deleted (WORM policy will take precedence if enabled).

26.1.3 Backup Target

Nutanix Objects will support 3rd party backup solutions, providing:

- Consolidation of all backup environments to Nutanix Objects.
- Standardized protocol making your backup data cloud-ready.
- Scale - Ability to support multiple backup clients simultaneously.
- Ability to handle small and really large backup files simultaneously with a key-value store-based metadata structure and multi-part upload capabilities.

26.2 Design Considerations

Object storage uses a flat hierarchy and is designed around large-scale data sets and is not suitable for low-latency workloads.

Objects are considered immutable and hence cannot be updated partially. This makes object storage suitable for data that needs to be retained for a long period of time.

References

26.3

Nutanix Objects Product Page:
https://www.nutanix.com/products/acropolis/object-storage-service/

Nutanix Objects Datasheet:
https://www.nutanix.com/documents/datasheets/buckets.pdf

Reimagine Object-based Storage in a Multi-cloud Era:
https://www.nutanix.com/2017/11/08/reimagine-object-based-storage-multi-cloud-era/

Object-based Storage Defined: Why and When You Need It:
https://next.nutanix.com/blog-40/object-based-storage-defined-why-and-when-you-need-it-28178

Nutanix Object Storage Service for the Enterprise Cloud:
https://www.youtube.com/watch?v=4TM7KWRN6FU

27

Prism

Author: Wayne Conrad

The Nutanix Prism UI lives at two layers, Prism Central, which manages multiple clusters, and Prism Element, running on each cluster.

27.1 Prism Element versus Prism Central

Nutanix's Prism Element UI and APIs are natively built into AOS at the cluster level and are suitable for day-to-day operations of tasks at the hardware and VM level.

Prism Central is a separate appliance or scale-out set of appliances. Prism Central has role-based access control for VMs, the V3 restful APIs, capacity planning, VM sizing analysis, and the UI for many Nutanix products like Calm, Karbon, and Flow that are intended to be used with multiple clusters.

Capabilities of Prism Central:

- SAML broker authentication
- Calm, Karbon, Flow
- Capacity Planning
- VM right sizing
- Reporting, Dashboards, Scheduled reporting
- V3 APIs

27.2 Prism Central Design Considerations

- Scale Out Prism Central requires running Prism Central on a Nutanix cluster. Single Prism Central servers may be placed on any other virtualization platform.
- Several Nutanix Prism Central services like Calm also require running Prism Central on a Nutanix cluster.

- The Nutanix cluster dependency is due to the use of Volumes for data storage in Prism Central, so make sure the Nutanix cluster has a data services IP.
- Scale Out Prism Central is three VMs and can tolerate one failure. Scale Out Prism Central only resides on a single Nutanix cluster and cannot span clusters or data centers.

Prism Element Design Considerations 27.3

- Some UI features are only on some hypervisors, such as network visualization, which is only on AHV.
- When using ESXi, vCenter registration allows the control of VMs on the Nutanix platform via the Prism UI equivalent to AHV.
- Prism Element is accessible via the IP of any CVM and the shared cluster IP.
- Microsoft IE11 and Edge have issues uploading large files to the Prism interface. We recommend the use of a modern browser such as Chrome or Firefox.

References 27.4

Prism Product Page:
https://www.nutanix.com/products/prism/

Nutanix Prism Datasheet:
http://go.nutanix.com/rs/nutanix/images/Prism-Data-Sheet.pdf

Tech Note: Prism Element:
https://www.nutanix.com/go/infrastructure-management-operational-insights-with-prism.php

The Nutanix Bible:
https://nutanixbible.com

28

Lifecycle Management

Author: Cameron Stockwell

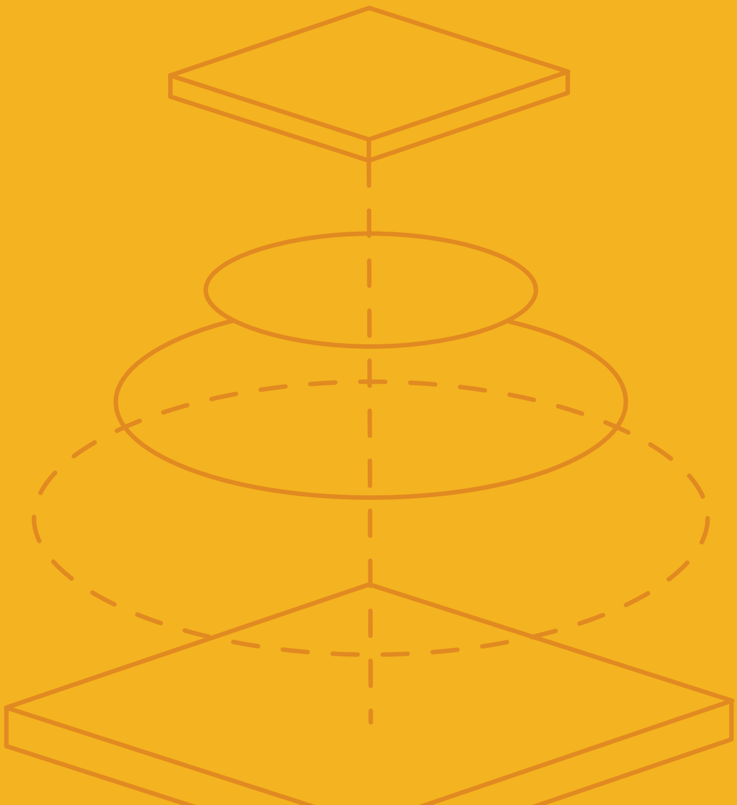

28.1 “One Click Upgrades” and Lifecycle Management (LCM)

‘One-Click upgrades’ are one of the most popular features of the Nutanix platform. From inception, our mission has been to deliver non-disruptive, yet simple upgrades at the click of a button. The premise concept being that administrators should not have to worry about keeping clusters updated with the latest patches whilst having to juggle maintenance windows or expensive labor costs.

This has a significant benefit for not only the administrator, but for the business by keeping clusters up to date for both software and firmware.

In the early days at Nutanix, upgrades were limited to AOS and the hypervisor, but over time the platform has grown to support more Nutanix software products and the ever-expanding list of supported hardware from different manufacturers. This evolution of the one-click upgrade story is continuing, with the recent introduction of the Lifecycle Management (LCM) framework to cater for this growth in the platform portfolio.

28.1.1 Software Upgrades

Software upgrades are a common operational process that should be simple and reliable in order to eliminate bugs and security flaws, and to deliver new features quickly as they become available to increase the business value from the solution investment.

Traditional software upgrades still come coupled with complex operational procedures or dependencies. This usually leads to expensive consulting costs in gaining the correct advice and project delivery to get software ‘up to date’ and to reduce any perceived operational risk.

‘One-Click Upgrades’ offer a simplified process for various software components on the Nutanix platform to address the shortcomings seen in traditional infrastructure software.

The list of software supported in the Nutanix ‘One-Click’ upgrade framework includes:

- AOS
- Hypervisor (AHV, ESXi, HyperV)
- Foundation
- NCC
- Files
- Calm

The list of Nutanix software that will be supported by LCM will continue to grow with each new release.

Nutanix software can be upgraded with no disruption to the operational status of the cluster. Even if a particular software component requires the reboot of a node. For example, upgrading the hypervisor, where all the associated maintenance mode or VM migrations are handled automatically by the ‘One-Click’ process.

It should be noted that AOS, Foundation, NCC, Files and other Nutanix software do not require any node reboots.

28.1.2 Firmware Upgrades

In traditional environments, firmware upgrades were often ignored after initial deployment until a critical problem arose which needed immediate addressing. Often, the same firmware would remain over the lifecycle of the equipment. The pain of such firmware upgrades for servers, SAN fabrics, network fabrics and others was simply not worth the risk for the perceived lack of business value.

This is no longer acceptable in the modern datacenter, given the increase in security concerns and such industry-wide reactions to the Spectre/Meltdown Intel microcode incidents.

On the Nutanix platform, many different hardware vendors are now supported under the same non-disruptive 'One-Click' functionality for firmware such as BIOS, HBAs, Disks, NICs and many others.

Currently, Nutanix firmware upgrades via LCM is supported on:

- Nutanix NX appliances
- Dell XC / XC Core appliances
- Lenovo HX / HX Ready appliancess

Not only are the supported hardware vendors increasing, also the matrix of components and hypervisors as well to support the vendor firmware. Ultimately, the goal is to allow your business to not be 'locked in' to any one hardware manufacturer as new purchasing cycles change, and still provide the operational simplicity for lifecycle management of firmware across those different manufacturers.

28.1.3 Requirements / Design Considerations

Like all good Nutanix designs, one of the core aspects to achieve non-disruptive One-Click upgrades is an "N+1" philosophy within the cluster. Design the cluster to be able to handle at least a single largest (in terms of CPU/Memory and Storage Capacity) node being temporarily offline while certain upgrades are undertaken, ensuring that remaining nodes can absorb the temporary increase in load due to the absent node.

The Upgrade Process - AOS Example 28.1.4

The component or entity being upgraded will determine the exact process. AOS upgrades are the most popular, which we will focus on here.

Once the upgrade binary has been downloaded, some pre-checks are done. These cover items such as version compatibility, cluster status, free space checks on the CVM (the AOS control plane runs in this virtual appliance), and internal service health checks.

The new AOS binary is then uncompressed on all nodes in parallel, however each node will not restart CVM services until the 'upgrade token' is received by each. Essentially, only one node is therefore upgraded at any single point in time (which is why the "N+1" design rule is important).

When the CVM restarts its services and associated post operational checks are complete, only then is the token released onto the next node in the cluster to be upgraded; and the process repeats throughout the entire cluster to completion.

During the CVM service restart process, storage services as seen by the local hypervisor on that same node are temporarily redirected to other 'healthy' nodes in the cluster, and the hypervisor on the node being upgraded can therefore still function as normal without any VM migrations required (storage remains accessible during the CVM temporary restart via replica copies on other nodes).

28.1.5 Summary

Nutanix has designed upgrades to be simple, reliable and scalable so the performance of workloads on the clusters are not impacted. Having Nutanix LCM managing multiple different hardware manufacturer's firmware via the same operational interface is industry-unique. For more information on Upgrades and a deep dive, refer to the Tech Note on Upgrades and LCM.

28.2 References

AOS 5.0 New Feature: Life Cycle Manager: https://next.nutanix.com/blog-40/aos-5-0-new-feature-life-cycle-management-17322

Tech TopX: Life Cycle Manager: https://www.youtube.com/watch?v=CftB7LhStnQ

29

AHV

Authors: Wayne Conrad & Magnus Andersson

Nutanix AHV embraces simplicity and turns on many features by default, vastly reducing the configuration complexity for the end user. We'll discuss the configurable options and a few requirements below.

29.1 Design Considerations

- AHV cannot overcommit memory. All memory is reserved.
- CPU masking between CPU generations for live migration is turned on by default. If older CPU nodes are added to a cluster, VMs will not be able to live migrate to them until after a VM restart. However, nodes with newer CPUs are automatically available.
- AHV cannot live migrate VMs between storage containers, or live migrate VMs between clusters.
- If possible, make sure your Nutanix clusters have access to the Internet sites listed in the AHV firewall and proxy KBs and documentation. This eases support by providing detailed statistics sent over Pulse "call home" and makes it much easier and faster to download new software versions.
- Network NIC load balancing on AHV must be configured by command line. Out of the box, Active / Passive is used.
- Nutanix AHV uses Open vSwitch (OVS) to provide network functionality. OVS natively supports Balance-TCP with LACP which is the recommended load-balancing option when you need more bandwidth than provided via the default active/backup configuration. Previous limitations around maintenance activities with LACP have been solved.
- VM high availability is turned on by default and comes in two modes listed below.
- VMs will restart from a failed AHV host as long as any AHV host has enough available resources to satisfy the memory requirement.

- A single click in Prism under “high availability” reserves enough resources to guarantee that all VMs from a failed AHV host will be restarted on other AHV host. By default, memory equivalent to one host is reserved across an RF2 cluster, and two hosts worth are reserved in an RF3 cluster. If needed, the VM High Availability memory reservation capacity can be changed via cli.
- Acropolis Dynamic Scheduler (ADS) - VM placement and load balancing.
- Leverage VM – VM Anti Affinity rules to make sure VMs does not run on same AHV host. These “should” rules can be violated for a short period of time during AHV host failure.
- Leverage VM – Host Affinity when one or more VMs needs to run on a set of AHV hosts for e.g. application licensing purposes This “must” rule is strictly respected by the ADS but can be overridden by AHV cluster administrators if needed.
- Never Schedulable node - An AHV host can be added as a never schedulable node meaning no VMs will ever run on this host. This can be used if ADS does not provide enough VM to AHV host placement guarantee to satisfy licensing requirements.
- Cluster lockdown mode secures access via SSH only to key authentication instead of passwords. The risk in a disaster is that SSH keys may be lost and admins cannot get in, but some high compliance and high security environments prefer securely stored keys to passwords.
- The Nutanix security operations guide contains a few additional hardening settings required by Federal, defense and other extremely high security environments. They are not considered generally necessary for typical compliance such as PCI, SOX, HIPPA, GDPR. These settings may have performance or other tradeoffs and should be carefully considered before implementation.

29.2 Pros and Cons of Nutanix AHV

Nutanix AHV Strengths:

- Natively clustered at the management level. Individual hosts cannot be misconfigured.
- Extremely simple to operate and quick to setup. AHV clusters can be built and running workloads in less than an hour.
- HTML5 interface, no plugins or Flash required.
- Can provide the complete physical CPU topology, including hyper-threads, to a VM.
- One-click simple upgrades.
- One-click micro-segmentation via Flow (Additional license required).
- Native open stack drivers.
- One-click Kubernetes via Karbon.
- Supports Citrix XenApp and XenDesktop.
- Excellent network, storage, and live migration performance.
- Supports higher density VDI than other hypervisors.
- Clean REST APIs.
- Included in Nutanix AOS. “One throat to choke” support from the NX node hardware to storage to hypervisor.

Nutanix AHV Weaknesses:

- Limited support for virtual appliances or certified applications.
- Much smaller ecosystem of 3rd party integrated products than vSphere.
- Limited support for role-based access control and multi-tenancy at the cluster level.
- Limited GUI support for unusual network topologies like NICs attached to different switches for different use-cases.
- No cross-cluster live migration support.
- No metro-clustering support.
- No memory overcommitment.
- No QOS for CPU performance.

Other Supported Hypervisors 29.3

In addition to AHV, Nutanix supports three additional hypervisors: VMware vSphere ESXi, Microsoft Hyper-V, and the Citrix Hypervisor.

Citrix Hypervisor is only supported on Nutanix for Citrix XenApp and XenDesktop use-cases.

29.4 References

AHV Virtualization Product Page:
https://www.nutanix.com/products/acropolis/virtualization/

Best Practices Guide: AHV:
https://www.nutanix.com/go/ahv-best-practices-guide.php

AHV: Virtualization Solution for the Enterprise Cloud
https://www.nutanix.com/go/ahv-a-virtualization-solution-for-enterprise-cloud.php

Nutanix Test Drive:
https://www.nutanix.com/test-drive-hyperconverged-infrastructure/

Best Practices Guide: Docker Containers on AHV:
https://www.nutanix.com/go/docker-container-best-practices-guide-with-AHV.html

30

Move

Author: René van den Bedem

Nutanix Move, formerly known as Nutanix Xtract, is a cross-hypervisor migration solution to migrate VMs with minimal downtime. The downtime is incurred during the cutover from a VMware ESXi or Amazon Web Services (AWS) source to the AHV target.

30.1 Design Considerations

- Nutanix Move migrations from VMware ESXi to AHV support: AOS 5.0.x-5.10.x, ESXi 5.5-6.7 and vCenter Server 5.5-6.7.
- For ESXi 5.1, use Move version 2.0.2 instead.
- Nutanix Move migrations from VMware ESXi to AHV Guest OS support: Windows 7/8/8.1/10, Windows Server 2008 R2/2012/2012 R2/2016, CentOS 6.3-6.9/7.0-7.4, RHEL 6.3-6.9/7.0-7.5, Ubuntu 12.04.5/14.04/16.04/16.10, FreeBSD 9.3/11.0, SUSE LES 11/11 SP1-4, 12/12 SP1-3, Oracle Linux 6.x/7.x, Debian 9.4
- Nutanix Move migrations from AWS to AHV Guest OS support: Windows Server 2012 R2/2016, CentOS 6.8/6.9/7.3-7.5, RHEL 6.8-6.10/7.3-7.5, Ubuntu 14.04/16.04/18.04
- VM preparation for ESXi Guest VMs can be completed automatically or manually.
- Replication data (target write) is automatically balanced across all nodes of a multi-node AHV cluster. This provides efficient CVM stress management and better write performance. This also removes the VM soft-limit recommended in previous versions.
- Dynamic tuning of the Move configuration, including compression, is used to improve the migration speed of VMs between ESXi and AHV.

- If migrating 32-bit Windows VMs, install the Nutanix VirtIO drivers first and then set the SAN policy.
- Two migration methods are supported: Full Migration and Data-Only. Data-Only is used when the Guest OS requirements are not met.
- Full migration support for Windows Guest OS requires UAC to be disabled.

References 30.2

Nutanix Move Product Page:
https://www.nutanix.com/products/move/

Nutanix Introduces Application Mobility from Public to Private Clouds:
https://www.nutanix.com/2018/05/09/nutanix-introduces-application-mobility-from-public-to-private-clouds/

Nutanix Move Overview in 90 Seconds:
https://www.youtube.com/watch?v=IN4koYLD_cI

31

X-Ray

Author: Gary Little

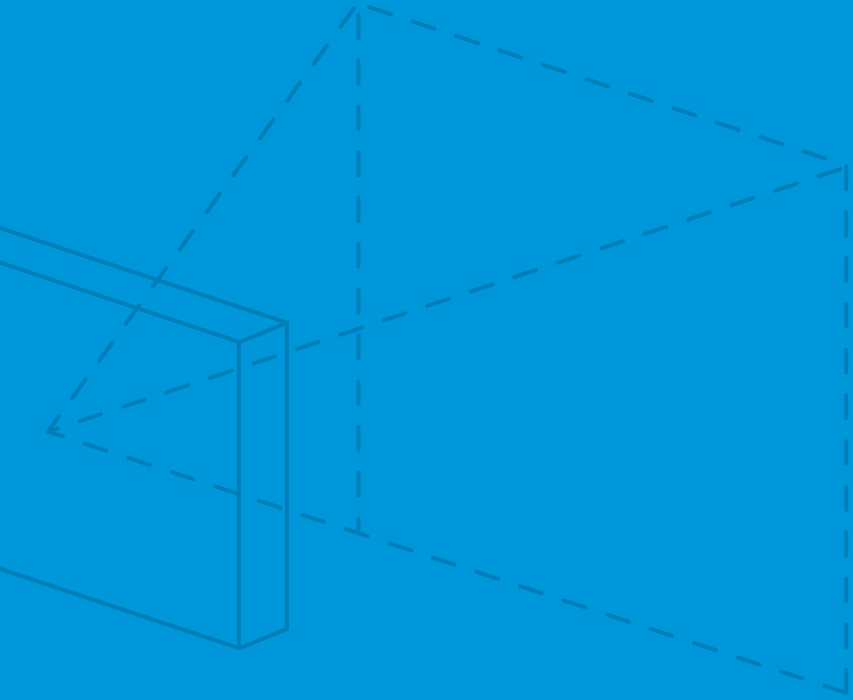

Benchmarking can be an important process when considering an enterprise cloud, based on Nutanix HCI. Adopting a new architecture typically involves proving that the performance is at least as good as the old infrastructure. It is hard to imagine how a simple, flat, cloud style architecture can compete with custom architected environment, with multiple switches, disk shelves and storage controller heads. A well-executed benchmark exercise can be a great proof point for the applicability of HCI.

FIGURE 21
Amdahl's Law

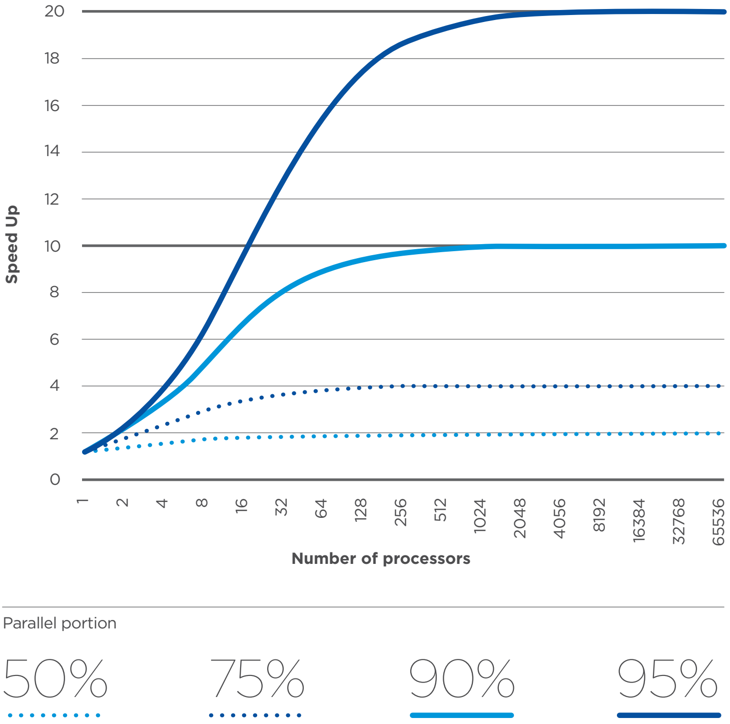

The desire to benchmark and execute Proof of Concept activities is as old as computing itself, Amdahl’s Law was theorized in 1967, and Little’s law, a key simplification of queueing theory dates to the mid-1950s. Amdahl’s law states that the improvement (speedup) of adding additional processors is inversely proportional to the sequential fraction).

Looking back in time, whenever we see a major shift in architecture we often see an accompanying benchmark standard.

- 1980s - 1990s first wave of mini-computers and early RISC architectures
  - TPC benchmarks for database workloads
- 1990s - 2000’s rise of shared external storage,
  - SFS for file services and later SPC-1 for database workloads
- Mid 2000’s initial phases of virtualization.
  - SPECVirt and VMmark gained limited traction demonstrating VM density
- Mid 2010’s early big-data and noSQL
  - YCSB

Moving to a hyper-converged environment presents a challenge for architects and teams who need to prove that HCI systems can run enterprise workloads. Most commonly the area of concern is around the capabilities of the HCI storage stack. This makes sense, because most enterprises are already familiar with running virtual workloads on stand-alone hosts.

Often the concerns are variations of the following themes:

- Does HCI architecture have the raw performance I need for my most demanding applications?
- Can HCI architecture give me the same level of resilience that I am used to with HW based failure protection?
- Will HCI retain consistent performance in the face of demanding multi-tenant applications?

By expressing concerns directly and clearly, we can devise a benchmark test-plan to address the them. A specific plan, with success criteria is usually more successful than creating a large matrix of test with rows for multiple IO sizes, read/write mix, randomness and queue-depth - then attempting that matrix to reverse-engineer a success criterion.

The simplest approach is to use Nutanix X-Ray to evaluate against these, or similar criteria. Nutanix engineering uses X-ray to validate their code against similar criteria prior to release.

31.1 Criteria 1 – Raw IO Performance

Does HCI have the raw IO performance I need for my most demanding applications?

The answer is almost certainly, “yes”. SSD performance is orders of magnitude faster than HDD. For instance, a single SSD can provide the same IOP performance as 500 HDDs. Furthermore, in an HCI environment the CVM is providing only the performance demand of guest VM’s running on the same host.

How do I know how much IO performance is needed?

For active database workloads - the IO demand per-node tend to be in the region of 20,000 IOPS per host before CPU cycles on the host. When we run the TPCx-HCI benchmark with 8 Postgres database VMs per host - the IO workload generated is 20,000 - 21,000 IOPS per host and CPU per host is almost saturated.

We saw similar patterns from Microsoft SQL Server driven by the HammerDB database benchmark. The HammerDB workload driver, attempts as man transactions as possible - hence more vCPU generate more transactions, which in turn creates more IO demand.

TABLE 8

Database Transactional Workload I/O Requirements per Host

Transactional Workload	I/O Requirement Per Host
4 vCPU SQL Server	5K – 10K IOPS
6 vCPU SQL Server	10K – 15K IOPS
8 vCPU SQL Server	15K – 20K IOPS
12 vCPU SQL Server	20K – 30K IOPS

For VDI workloads some rules of thumb are:

TABLE 9

VDI Workload I/O Requirements per VM

User/Worker Type	Applications Open	VM Config	IOPS
Task-based (Light)	Limited (1-5 open)	1 vCPU, 1GB RAM	3-7
Knowledge (Medium)	Standard office (5+ open)	1 vCPU, 1GB RAM	8-16
Power User (Heavy)	Compute intensive (5+ open)	1 vCPU, 2GB RAM	17-25
Power User (Heavy)	Compute intensive (5+ open)	2 vCPU, 2+GB RAM	26+

Using the guidelines above it is possible to measure (a) whether the HCI infrastructure can provide the necessary performance. Moreover, by creating an IO model (using fixed rate workloads based on the above) we can assess the likely impact during failure. If we accept that each host should supply 20,000 IOPS even under failure - we can devise tests that prove - or refute that.

Simply measuring how fast our storage will run under ideal conditions tells us nothing about how applications will be impacted in failure or multi-tenant environments. In almost all cases the raw performance exceeds typical demands.

31.1.1 Some benchmark pitfalls

One case where we frequently see confusion is when a Nutanix HCI cluster is compared to an existing SAN by running a single workload on a single VM within the cluster. I call this the “you are here” problem. Like any performance test, a single VM, single disk workload running in a cluster will tell you something about the performance of the cluster - but perhaps not what you think.

In this example, the single VM, single disk test yields 2,500 IOPS - but the total capacity of the cluster is around 600,000 IOPS. The reason for the discrepancy is that a single VM on a single node cannot drive all the performance from the entire cluster. A Nutanix cluster is designed to provide consistent performance to multiple VMs running on multiple hosts. In the example chart below - we see that the amount of IOPS that the cluster delivers increases as we add load, until it reaches saturation at around 600,000 IOPS on this particular cluster.

If very high performance is needed from a single VM - consider Nutanix volumes (ABS). However, with many VMs running on many hosts (as is the normal case) the cluster can deliver many hundreds of thousands of IOPS in standard HCI configuration as seen below.

FIGURE 22

Cluster-wide Performance

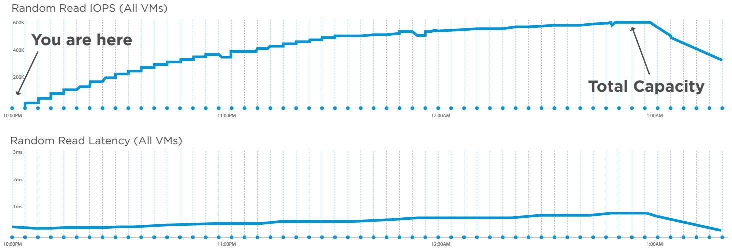

Total IOPS capacity easy-button

31.1.2

To avoid the single VM bottleneck, ensure that all nodes of the cluster are generating work, otherwise the capacity of the cluster will be severely under-reported. One simple way is to use the X-Ray 4-Corners test. X-Ray automatically provisions a VM to each host in the cluster, then uses the Linux “fio” generator to issue IO to the cluster.

FIGURE 23

Four Corners X-Ray Test

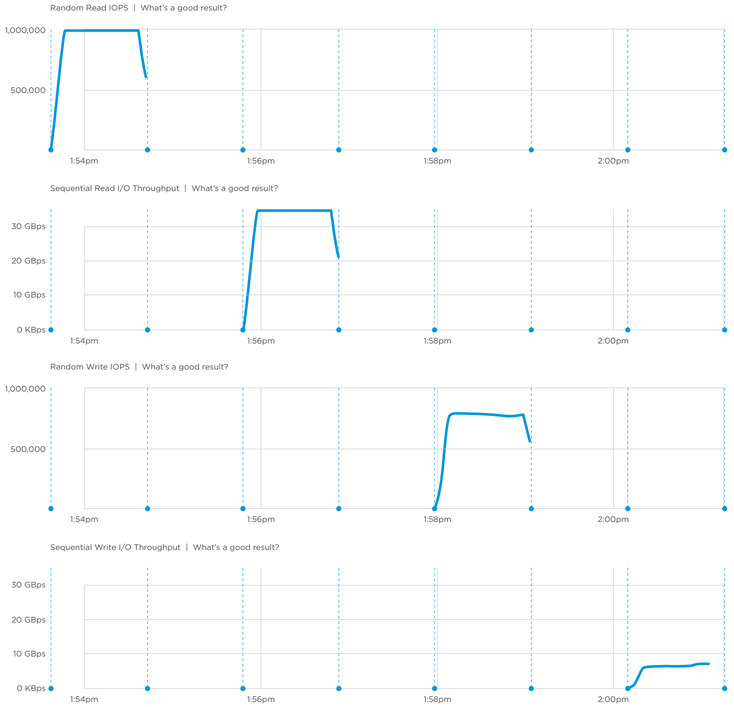

It is possible to use other workload orchestrators such as "HCIbench" - or even to use "fio" or "vdbench" generators directly - X-ray simply provides a nice wrapper around "fio" to deploy and manage the workloads. As well as a nice UI and reporting.

For full-disclosure, we use an 8KB IO size for the random workloads, and 1MB for the sequential workloads. The goal of the 4 Corners test is to optimize for the IOPS and throughput. In each case we use a high degree of concurrency, thus response time is expected to be high. The four corners test can be inspected and modified within X-Ray, should you wish to change IO sizes, concurrency values etc.

31.2 Criteria 2 – Resilience

Can HCI architecture give me the same level of resilience that I am used to with Hardware based redundancy?

As with raw performance, it is not obvious how an HCI cluster is able to achieve the same degree of resiliency as a traditionally architected deployment which uses multiple redundant hardware connections etc.

In many cases testers fall back on the "Disk Pull" test during a workload. While this experiment will reveal what happens when a disk is pulled from a running system, it does not accurately simulate a disk failure. The disk enclosure firmware will treat a disk-pull differently to a disk failure - which tends to degrade over time anyway.

For small clusters that typically are used for POC - the largest failure domain that can be sustained is an entire node. We can use the X-Ray extended node failure test to show what happens to the remaining nodes on the cluster - and the time/impact to rebuild the data.

In this test, X-Ray connects directly to the IPMI port on the cluster hardware and issues a power-off command (not shutdown) without giving the cluster software any warnings. The reason we choose to fail a node is a) larger domain and b) more like a real failure. c) Using IPMI we can fail any sort of cluster node that supports IPMI because IPMI is a public interface. It is possible therefore to compare the failure handling between nodes running Nutanix and other vendors HCI implementations.

In this test, power is removed from Node0 at around the 30-minute mark - and is not re-applied until after the test. We would expect to see an initial drop in performance as data is re-replicated, after which performance continues as normal.

Criteria 3 – Consistent Performance

31.3

Will HCI retain consistent performance in the face of demanding multi-tenant applications?

This test is all about multi-tenancy. HCI vendors make a variety of design decisions on multi-tenancy. One choice is to rely entirely on some sort of QoS, as is common practice in the traditional storage architecture. In the Nutanix world, workloads that are running on separate nodes should naturally have very little cross interference.

When architecting a real-world solution, it is very unlikely to co-locate both reporting and OLTP workloads on the same HOST because we want to avoid clashing CPU resources. With Nutanix, by virtue of data-locality there is a natural separation of IO, even though the storage is shared - as it must be to support VM migrations etc.

In the below X-Ray experiment, an OLTP workload is started in isolation on Node-A and no other workloads run on the cluster. After 30 minutes we start two additional reporting workloads on separate hosts. The reporting workloads are read-intensive and sequential. They will drive a lot of work to the storage. Without either QoS or data-locality it is likely that the OLTP workload will be negatively impacted by the addition of the reporting workloads. Since X-Ray can be run on multiple hypervisors, it can be interesting to compare the data-locality/QoS between different HCI imple mentations.

FIGURE 24

X-Ray Consistent Performance Graph

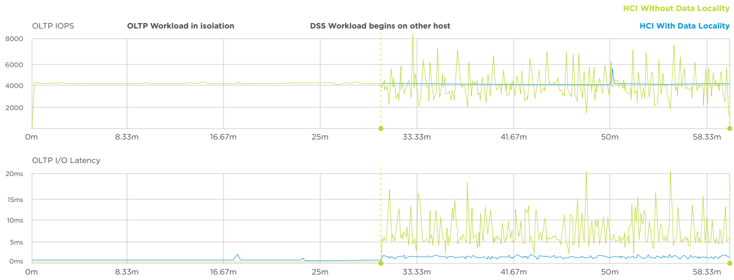

31.4 Other Criteria

- Ability to recover from total power loss
- Ability to ingest large amounts of data
- Predictable scaling

References

31.5

X-Ray Product Page:
https://www.nutanix.com/products/tools-and-technologies/x-ray/

X-Ray Datasheet:
https://www.nutanix.com/documents/datasheets/nutanix-x-ray-datasheet.pdf

Numbers that matter - performance requirements of databases on HCI:
https://next.nutanix.com/blog-40/assessing-hyperconverged-performance-the-numbers-that-matter-part-1-14347

HCI Performance Testing made easy Part 1:
https://www.n0derunner.com/2018/09/hci-performance-testing-made-easy-part-1/

HCI Performance Testing made easy Part 2:
https://www.n0derunner.com/2018/09/hci-performance-testing-made-easy-part-2/

HCI Performance Testing made easy Part 3:
https://www.n0derunner.com/2018/09/hci-performance-testing-made-easy-part-3/

HCI Performance Testing made easy Part 4:
https://www.n0derunner.com/2018/09/hci-performance-testing-made-easy-part-4/

X-Ray Community Forum:
https://next.nutanix.com/nutanix-x-ray-18

32

Foundation

Author: Wayne Conrad

Nutanix Foundation is one of the most powerful features of Nutanix and allows the setup of large clusters in approximately 2 hours or less.

Nutanix Foundation is hardware and hypervisor agnostic, allowing the setup of Nutanix clusters on all supported hardware platforms and hypervisors.

Nutanix Foundation comes in the following flavors:

- A Java applet supporting Nutanix NX and OEM hardware that already have a Nutanix CVM running. All the Java applet does is discover the IPv6 address of a CVM and forward traffic to the CVM to use the baked in version of Foundation in the CVMs. If you manually IP a CVM, you'll get the same experience.
- A Foundation VM supporting bare metal installs to hardware that is not preloaded
- A standalone Foundation application for Windows and MacOS, currently in tech preview.

Nutanix Foundation requires being on the same broadcast domain as the appliances CVM you are installing, so most people use a flat switch attached to a laptop to simplify the deployment. Bare metal Foundation attaches an ISO file to the out of band management of server hardware, so the out of band network needs to be plugged in and either IP addressed, or be accessible via IPv6 with the MAC address.

Important considerations for install:

- The Nutanix Portal has the Foundation Preconfiguration tool, which allows a configuration file to be generated with the cluster configuration. This is then uploaded into Foundation during install.

- Some new hardware does not support 1GbE adapters in 10GbE NICs. I suspect you do not have a spare 10GbE switch, so you'll need to Foundation using the top of rack switches you'll use for production.
- Set aside IP addresses for cluster growth later.
- Confirm you have got the correct hypervisor ISO for your chosen hardware if you are not using AHV. Nutanix published a whitelist of approved ISO files. Check the MD5 sum of your ISO file matches as vendors have been known to update their ISO files silently without changing a build number.
- Nutanix publishes wiring diagrams for NX nodes showing which NIC port is the IPMI failover so you'll only need to wire a single cable.
- Native VLAN tagging generally makes installations easier. If you are going to install on a flat switch, then connect to a switch with no native VLAN tag, as you will want to stop the cluster and tag all traffic.

References 32.1

Nutanix Foundation Demo Video – From Bare-metal to Production in Minutes:
https://www.nutanix.com/2014/07/15/nutanix-foundation-demo-video-from-bare-metal-to-production-in-minutes/

The Nutanix Bible:
https://nutanixbible.com

33

Data Center Facilities

Author: René van den Bedem

The Data Center Facility is very often overlooked in infrastructure design. This is because it is assumed that sufficient space, power, cooling, cabling and perimeter security will be available. The biggest risk with Nutanix solutions and Data Center design are the increased power and cooling requirements for enterprise grade solutions. A 42 Rack Unit Server Cabinet with forty NX-8035-G6 nodes will demand 26 kW of power (maximum) and 17 kW (average). A legacy Data Center will typically have a limit of 5 kW to 8 kW per rack.

The Data Center Facility is categorized into the following types:

"Bricks-and-Mortar" – You construct a concrete building that will operate as a Data Center.

- Pros: You own it, Reduced OPEX.
- Cons: Increased CAPEX, Increased time to design, order and implement, Space restrictions/wasted space.

Co-location – You rent rack-space from a Data Center Facility service provider that is responsible for the facility; you just have to provide the active equipment to be installed in the racks provided and be responsible for operating it.

- Pros: Reduced CAPEX, Reduced time to implement, Facility is not your responsibility.
- Cons: You do not own it, Increased OPEX.

"Pre-Fabricated" within an existing building – You have a building where a pre-fabricated Data Center is constructed.

- Pros: You own it, Reduced OPEX, Scalable/Modular, Reduced time to implement.
- Cons: Increased CAPEX, Substantial lead-time to order and deliver.

"Performance Optimized Data Center" (POD) – You have a plot of land with perimeter security and a vendor delivers shipping container-size (20-foot, 30-foot or 40-foot) modular Data Centers. You stack them across the site, just like Lego blocks when you were a child.

- Pros: You own it, Scalable/Modular/Mobile, Reduced OPEX, Reduced time to implement.
- Cons: Increased CAPEX. Substantial lead-time to order and deliver.

Use-Cases 33.1

The following use-cases drive Data Center Facility design:

- **Governance and Compliance** – Do you have any regulations regarding the locality of your Customer data that you must adhere to? How will you prove compliance?
- **Cost to implement (CAPEX) and operate (OPEX)** – Full investment upfront or PAYG? Do you have the budget?
- **Time to implement** – "Ready to go" or years of construction and project management? Do you have the leisure to wait? Or is your business requirement urgent?
- **Delivery method** – "Turnkey solution" or "Piecemeal"? Piecemeal projects with contractors and sub-contractors are a risk to the project schedule.
- **Green Energy** – Do you have a PUE compliance requirement? If yes, build your Data Center in a cold location.

Design Considerations 33.2

These are some of the design considerations for Nutanix solutions and Data Center Facilities:

- **Uptime Institute Tier Rating** – Tier-1 to Tier-4? Approximately US$100K to certify your design.
- **Power** – UPS 1+1, UPS N+1 or Radial (Generator and UPS are combined into one unit), Auto Transfer Switch (ATS), External Generators, Fuel Cells, Distance from Building Transformers? Supply voltage?

- **Flooring** - Raised or not? With "In Rack" cooling and overhead cabling and overhead pipes, a raised floor is not necessary.
- **Load rating** – What is the load rating of the data center floor, ramps and elevator?
- **Cooling Method** – Compressor or Chiller? Compressor is used for small to medium size Data Centers; Chiller is more expensive, but scalable for large Enterprises.
- **Cooling** – "Traditional" Raised Floor cooling (CRAC pushes cold air under floor to create a plenum, cold air is forced through tiles in the "Cold Aisle", active equipment sucks cold air in and pushes hot air out into the "Hot Aisle", hot air rises and returns to CRAC intake for cooling), Hot Aisle Containment ("HAC" – rigid enclosure containing hot air exhaust of two rows of equipment), Cold Aisle Containment ("CAC" – rigid enclosure containing cold air intakes of two rows of equipment). HAC and CAC use "In Rack" cooling solutions to increase the cooling mass per rack.
- **Power and Data Cabling** – Underfloor or Ceiling Suspended trays? Distance to Top of Rack, End of Row or Central Rack Access/Leaf Switches?

Monitoring – IP-CCTV, PUE monitoring, Temperature, Humidity, Water Detection, Smoke Detection (part of Fire Suppression)

- **Fire Suppression** – Legacy FM200 or Novec 1230? Protect every room in the Data Center?

Physical Layout – Single white-space (with functional rows) or separate functional rooms (UPS, Generator, Server, Recovery, Archive, Security, Network, ISP/DSP, etc.). Reserved areas for future requirements (10-year, 20-year plans)?

- **Physical Security** – Locking system: doors and racks? Data Center – above ground, underground, heavy equipment access, operator access, etc.? Data center entry requirements?
- **Location** – Perimeter of a city (cheap land) with dual electrical

sub-stations, multiple POPs from ISPs/DSPs, away from major thoroughfares and flight paths, no history of natural/ man-made disasters.

Risks 33.3

These are some of the risks associated with Nutanix solutions and Data Center Facilities:

- Non-technical people have no concept of how complicated a Data Center is. They envisage picking up a laptop and moving it from one desk to another. Set the right expectation; make sure you fully explain and communicate the risks, budget and project timelines involved from the start.
- There is no “right solution”, there is only the solution that fits your business requirements, budget and timeline. Talk to the experts that specialize in this field, get quotations and advice and then select the best strategy for your company.
- Legacy Data Centers typically have a limitation of 5-8kW of cooling per rack. Nutanix solutions require 20+kW peak per rack. Avoiding Legacy Data Centers is recommended for enterprise solutions, consider a modern Co-Location service instead.

References 33.4

Data Center Knowledge:
https://www.datacenterknowledge.com/manage

Schneider-Electric:
https://www.schneider-electric.com/en/work/solutions/for-business/data-centers-and-networks/

Google Container Data Centers:
https://www.youtube.com/watch?v=zRwPSFpLX8I

34

People & Process

Author: René van den Bedem

Successful HCI projects are bulit upon a very close collaboration between the server virtualization, network and storage teams. In fact, it makes more sense to merge these three teams into one “Cloud Infrastructure” team. It is also very important to cross-skill these team members and let them evolve into “Cloud Architects”, “Cloud Administrators” and “Cloud Operators”. However, make sure you keep your Backup/Recovery/Archive responsibilities separate.

When considering failure scenarios for Business Continuity and Disaster Recovery, the biggest risk is not a natural disaster, but the disgruntled rogue administrator or the incompetent administrator, who has the keys to the kingdom, taking out every system. With HCI and the “Cloud Administrator”, this risk is compounded. It is very important to separate the administration and operations responsibilities for operational data and backup, recovery and archive. This way if either one is wiped out across all data centers, you still have the other to recover from. Apply this concept to physical data center security as well.

The storage processor of legacy storage arrays has now become a virtual appliance running (Nutanix CVM) on the host itself. Make sure administration, operations staff and monitoring systems understand the importance and give it the respect it deserves.

Moving from legacy, 3-tier infrastructure to HCI is a big change, so do not underestimate or ignore the imperative to update all relevant processes and procedures. HCI will simplify and improve the infrastructure stack, consequently simplifying the standard operational procedures, but change will be required with respect to people, process and technology.

References

34.1

What is a cloud architect? A vital role for success in the cloud: https://www.cio.com/article/3282794/cloud-computing/what-is-a-cloud-architect-a-vital-role-for-success-in-the-cloud.html

Analyzing the Role and Skills of the Cloud Architect: https://www.gartner.com/binaries/content/assets/events/keywords/catalyst/catus8/analyzing_the_role_and_skills_of_cloud_architect.pdf

35

Risk Management

Author: Daemon Behr

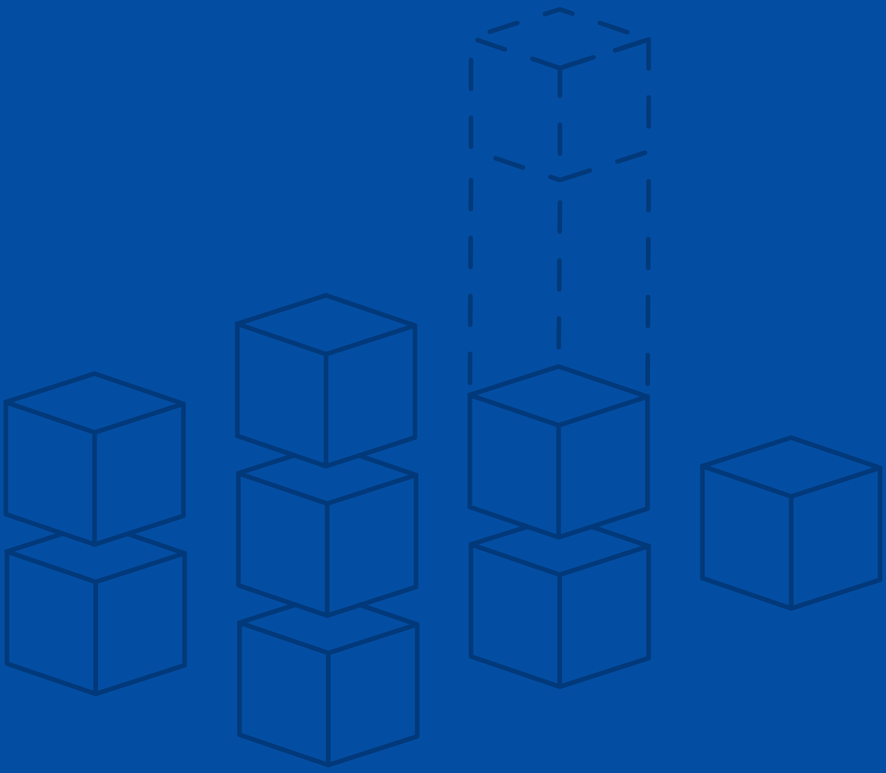

Risks are inherent in every aspect of infrastructure design and operations. They are the states that exist whether they are acknowledged or not. They are dynamic and change as both time and other aspects of the environment also change. Risks can be ignored, but their effects will often cause undesirable effects that are more difficult and costlier to repair. Knowledge is the power to affect outcomes and control the future. Understanding and managing risk is the easiest way to do this.

If you were commanding a military force and about to go into battle, you would want to know:

a. The actors, both friendly and unfriendly.

b. The active conflicts occurring.

c. The movement of troops, vehicles, etc.

d. Perceived and verified strategies of the above items.

e. Your immediate and long terms objectives.

f. A strategy for completing your objectives.

g. Antagonists of your objectives.

h. Possible negative outcomes from actions / events of each antagonist.

i. Timelines for each negative outcome to manifest if nothing is done.

j. Actions that can be taken to proactively or retroactively protect against antagonists.

k. The change in timelines for each negative outcome to manifest after actions have been taken.

l. The likelihood of each antagonist action / event occurring.

m. The impact of each antagonist action / event on infrastructure, personnel and operations.

n. The reliability of the information you have to make decisions.

How do you get this information and how do you organize it into and actionable strategy? In this chapter, we will explore some methods to obtain the operational intelligence required for making design decisions based on identified risks.

Risk can be defined as something negative that may happen and will have ramifications. Or said another way; probability and impact. If either part increases, then the overall risk severity will increase. Risk is dynamic and is monitored by Key Risk Indicators (KRIs). A KRI helps track and identify risks. A database of all identified risks is called a risk register. How much risk someone is willing to take is called risk appetite. An action that is to be taken for a risk, is called the risk treatment. The remaining risk after a treatment is performed is called the residual risk.

FIGURE 25

Risk Management

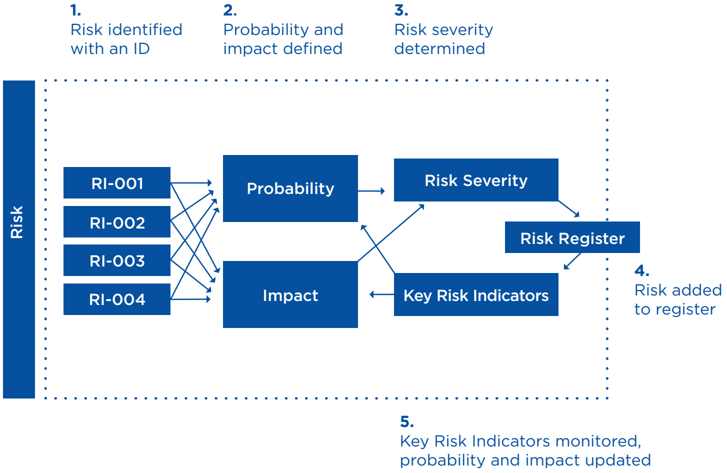

> "There are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns—the ones we don't know we don't know."
>
> **- US Secretary of Defense Donald Rumsfeld**

The terms, 'known knowns, known unknowns, and unknown unknowns' (KK-KU-UU), were popularized by Donald Rumsfeld in the early 2000s, but are also commonly used in project management, and analytical sciences. It is a simplified explanation of the framing of information into categories.

This can be used as a means of sorting information based on the completeness of its understood risk exposure.

35.1 Risk Categories

If you simply add risks into the risk register ad-hoc, then you will gain some benefit, but not its full potential. This is because the risks identified should at least cover the following areas:

a. **People** – All stakeholders, relevant business units, clients, operational staff, vendors and 3rd party professional services.

b. **Process** – Including operations, architecture design, 3rd party engagements, incident response, and disaster recovery

c. **Technology** – All relevant technologies in scope and their design qualities; (AMPRS) availability, manageability, performance, recoverability, and security. Technology can also have sub-categories such as hardware, software, configuration.

There are many more types of risks that can be considered, such as budget, competition, compliance, force majeure, integration, procurement, resource, strategic, etc. This chapter is focused on infrastructure risk and the technology category. Below is an example of how to determine the technology risks, based on its design qualities.

Availability 35.1.1

This pertains not only to operational states, but also transition states. This includes periods during upgrades, migrations and DR. Availability should be considered on an application or workloads basis. The availability requirements of each workload or application need to be determined in co-operation with the responsible business units.

A good method for determining availability requirements is determining where it fits in the business workflows. If, for instance, applications are client facing and need to be available to take orders, then it would have a higher availability requirement than a back-end system that only gets used a few times a month.

35.1.2 Manageability

This includes aspects such as who, how, when and from where will someone perform management operations on a technology.

35.1.3 Performance

This includes understanding the performance requirements and SLAs of all the workloads in the environment and what the ramifications are for not meeting them.

35.1.4 Recoverability

This includes understanding the various states that the infrastructure can be in when a failure occurs.

35.1.5 Security

This includes knowing the attack surfaces in your environment, the vulnerabilities, and the proactive and reactive actions for incidents.

35.2 Identifying gaps

A gap is the specific part of a risk that can be reduced by a treatment. It is a representation of the current state in comparison to the desired state. An example of a gap would be:

The password policy is set to a maximum age of 90 days, but most users have not changed their password in years and they are manually set to never expire.

35.3 Recommendations

Recommendations are the suggested risk treatments that consider the organizational risk appetite, the severity, and the operational capabilities to initiate a treatment. A recommendation based on the above password example would be:

Remove all manual password age exemptions and force a change of all passwords. This should be in concert with established organizational security policies.

Severity index 35.4

The severity index is a number that is obtained by multiplying the risk probability by the impact. For example, if both have a rating of 1-10, then the severity index would be a multiple. An example would be that a risk has a probability of 8 and an impact of 7. The assigned severity would then be 56.

Prioritization 35.5

The prioritization is based on the risk severity, business objectives and treatment. It can be listed in a spreadsheet in the order of priority, with risk ID, risk treatment recommendation, and severity. The risk ID can hyperlink to the full description in the risk register. This document is essentially a distilled action item list for risk treatments with justifications. See example below:

TABLE 10

Risk Prioritization

Priority	Risk ID	Recommendation	Severity
1	RI-020	Close all internet facing services that are known to be insecure.	90
2	R1-011	Configure ACLs between VLANs.	80
3	RI-013	Implement security strategy to assess, secure, monitor environment for indicators of compromise.	72

35.6 Risk Register

Below are some examples of fields that can be used in a risk register. This combines all of the areas that were covered in this chapter.

TABLE 11

Risk Register

Category	Risk ID	Gap	Recommendation
Network-config	RI-001	Using this network architecture for HCI will limit the scalability, increase complexity and potentially cause bottlenecks.	Implement a new spine-leaf network topology.
Network-config	RI-002	This will inhibit deployment and add additional costs to remedy.	Verify available ports on existing switching. Add additional switches if necessary.
Storage	RI-003	If the current methodology is to use native storage array replication, then this will need to be redesigned to accomodate the new infrastructure.	Verify current replication mechanism and design a replication strategy supported by the new environment.

Risk Description	Probability	Impact	Severity (probability x impact)
3-tier network architectures use inter-switch links to provide network connectivity across access layers segments. Link oversubscription will arise when the spanning tree blocks redundant links to prevent network loops on the L2 segments.	7	7	49
If the existing network infrastructure is used, there may not be a sufficient number of ports available for the proposed.	5	7	35
If a new infrastructure is being built in parallel to supersede the existing environment, then the role of replication target for South America needs to be considered and created.	8	8	64

To recap, there were two types of documents presented in this chapter:

a. The risk register. This can be a spreadsheet or a database, depending on the size and how it is shared and accessed. This will be continually updated as the environment changes or it will be reviewed on a schedule. Design decisions and risk treatment prioritizations will be outputs from reviewing the risk register.

b. Risk prioritization document. This is a simplified action item list that outlines risk treatments to be performed in order of priority, with links to a detailed overview in the risk register.

35.7 References

Designing Risk in IT Infrastructure, by Daemon Behr:
http://designingrisk.com/buy/

Insights:
https://portal.nutanix.com/#/page/insights

Field advisories:
https://portal.nutanix.com/#/page/static/fieldAdvisories

Security advisories:
https://portal.nutanix.com/#/page/static/securityAdvisories

End of Life Bulletin:
https://portal.nutanix.com/#/page/static/endOfLife

Failure analysis technical guide:
https://go.nutanix.com/failure-analysis.html

NUTANIX™

RoundTower®