

White Paper

Nutanix Cloud Platform on Intel Technologies Provides Seamless Application Mobility for Hybrid Multicloud Environments

Sponsored by: Nutanix and Intel

Eric Burgener
May 2022

Natalya Yezhkova

IDC OPINION

There is no doubt that more and more enterprises are moving to a hybrid multicloud infrastructure model as they move through digital transformation. During this journey, most enterprises start with infrastructure modernization, and then move to increased automation and integration of public cloud technologies, and ultimately begin to look for a unified control plane that spans all of their on-premises and off-premises IT infrastructure. Hyperconverged infrastructure (HCI) is often deployed as part of infrastructure modernization, and the fact that its software-defined nature is very similar to that of the web-scale infrastructures deployed in the public cloud makes it easier to manage hybrid cloud environments. HCI can provide a seamless on-ramp to the public cloud, with enterprises moving applications into virtual machines (VMs) that run on premises and then ultimately moving some of them to the public cloud – provided there is a proven, automated way to migrate applications without code changes.

Nutanix, a leading vendor of hybrid multicloud solutions based on HCI architecture, addresses these needs with its Nutanix Cloud Platform (NCP). NCP is flexible; is simple to deploy, manage, and scale; and includes enterprise-class storage management capabilities that enable dense workload consolidation onto a software-defined architecture. Many of Nutanix's thousands of customers run mission-critical workloads on it as well. NCP also includes Nutanix Cloud Clusters (NC2), a solution that allows the full NCP stack, with its comprehensive set of capabilities, to be deployed in public cloud-based environments. NC2 provides the flexibility, simplicity, and cost efficiency needed to manage applications in private and multiple public clouds from a single pane of glass, all operated as a single cloud. It also provides the application mobility to move workloads between private and public clouds simply, reliably, and quickly – all without any application code changes – a fact that makes on-premises NCP an excellent on-ramp to the public cloud.

Software-defined infrastructure like NCP typically runs on Intel x86 servers, taking advantage of multiple generations of Intel Virtualization Technology to make migrating VMs between different generations of servers and different instances much easier. The broad usage of Intel technology in software-defined infrastructure offerings both on premises and in the public cloud provides the proven technology and consistent underlying foundation that infrastructure modernization requires.

SITUATION OVERVIEW

Most enterprises are undergoing digital transformation – the move to much more data-centric business operations – and IDC notes that enterprises with preexisting on-premises IT infrastructure tend to go through four developmental phases as part of this process.

- **Phase 1** is infrastructure modernization.
- **Phase 2** is a move to more automated IT operations.
- **Phase 3** is where organizations begin to leverage cloud technologies.
- **Phase 4** is where enterprises that now have a hybrid cloud infrastructure (with some on-premises and public cloud-based workloads) seek to unify those operations with a unified control plane that spans all deployment models and locations.

Improving infrastructure efficiency is an important consideration as enterprises move through each of these phases. Throughout the latter half of this process, most enterprises will have both on-premises and off-premises IT infrastructure that must be managed.

IDC survey data from 2021 indicates that almost 70% of digitally transforming enterprises will be refreshing their server, storage, and/or data protection infrastructure by 2023. In undertaking these projects, the typical enterprise has multiple goals, including increasing IT agility, positioning to be ready to take advantage of cloud technologies, increasing infrastructure efficiency, and providing support for new workloads that demand higher performance, availability, scalability, and security requirements from their IT infrastructures. At the same time, they also need to reevaluate legacy workloads to determine the best path forward for them as IT infrastructure evolves.

Almost 70% of digitally transforming enterprises will be refreshing their server, storage, and/or data protection infrastructure by 2023.

Moving Toward Optimal Workload Placement

As they craft the right infrastructure for the digitally transformed future, CIOs will consider the optimal deployment models for various workloads. Most new applications will be architected in a cloud-native manner, which supports application mobility, and many of them will be developed in the public cloud, regardless of where they are ultimately deployed. DevOps is often implemented as enterprises move to more agile development models as well – a capability that improves business agility. In the long run, organizations generally will choose from among three deployment models – traditional three tier IT infrastructure, on-premises private cloud, and public cloud infrastructure.

Traditional three-tier IT infrastructure is defined as separate server, storage, and networking components that are purchased separately, integrated by the customer on premises, and dedicated (not multitenant infrastructure like much of the public cloud is). It also tends to be of the more hardware-defined variety and hence is generally less flexible, less scalable, and more costly than software defined. On-premises private cloud refers to more software-defined architectures (hyperconverged infrastructure is the fastest-growing segment of this market) that are much more flexible, scalable, and easier to deploy and manage by IT generalists. Public cloud is generally software defined.

While most new workloads will be developed as cloud-native applications, strategies around legacy workloads typically focus on one of six possible dispositions:

- **Retain** them in existing infrastructure as is
- **Rehost** them onto more software-defined (and, in many cases, cloud-based) infrastructure

- **Refactor** them for better operation in the cloud
- **Rearchitect** them for optimal operation in the cloud
- **Replace** them with software-as-a-service (SaaS)-based versions (e.g., a move from on-premises Microsoft Exchange servers to Office 365)
- **Retire** them

For legacy workloads, there are time to deployment and cost implications with each of these decisions. Retention, replacement, or retirement can have lower risks, while refactoring and/or rearchitecting them can introduce long lead times and development costs along with higher risk. In fact, the risk introduced by refactoring and/or rearchitecting increases the chance the cloud projects can fail. Rehosting by encapsulating a workload into a VM is a very common approach that is generally low risk yet still provides an application mobility that did not exist before. That mobility is important as it gives IT managers a deployment agility, in both on-premises and off-premises locations, that can be critical in optimizing workload placement over an evolving application life cycle.

With every enterprise modernizing its infrastructure, IT managers are faced with developing a coexistence strategy for legacy and new workloads. Implementing private cloud infrastructure on premises through a software-defined approach like HCI can provide a platform that makes it very easy to support data sharing and process interactions between legacy and new workloads. This is where the flexibility of HCI to support a variety of different deployment options – VMs, containers, bare metal, different hypervisors, and on premises or off premises, all of which are managed from a single control plane, can become very important.

Hybrid Multicloud: The Way IT Infrastructure Will Be Built Going Forward

While enterprises typically begin their foray into public cloud with a single provider, it has become clear over the past several years that most enterprises actually use more than one public cloud provider. In fact, in IDC's *IaaSView Survey* conducted in December 2021, 64% of enterprises confirmed that they use more than one public cloud provider. This not only creates a sense of competition and avoids comprehensive cloud provider lock-in but gives enterprises access to different services that may not be available from just a single cloud provider. Each deployment model – traditional IT, on-premises private cloud, and off-premises public cloud from potentially multiple providers – generally requires its own management interface, increasing complexity. It is at this stage that enterprises seek to unify hybrid, multicloud operations with a unified control plane that brings all deployment models and locations under a single pane of glass.

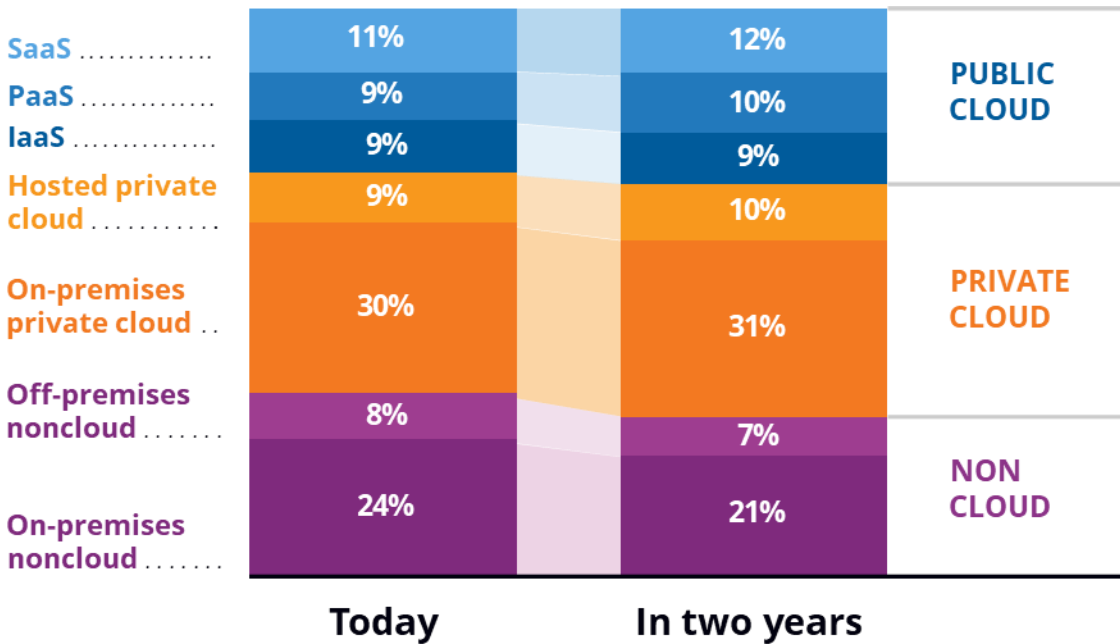
In the new hybrid multicloud world, application mobility is a critical capability in optimizing workload placement decisions. Rehosting legacy workloads and placing them on modernized, more software-defined on-premises infrastructures make it much easier to ultimately leverage public cloud-based technologies. In addition to providing new geographical locations from where it can make sense to run workloads, cloud services can dynamically provide additional resources to scale an application to better meet seasonal demands (i.e., cloud bursting), can provide a disaster recovery option that can help get operations back up and running quickly in the event of an outage, and may be the best way to run applications in various stages of their life cycle (e.g., development and testing of applications in the public cloud with production deployment on premises and initial deployment in the cloud with a repatriation to on-premises infrastructure due to evolving access patterns).

Recent IDC research shows that enterprises are clearly shifting more and more workloads into cloud-based environments (see Figure 1).

FIGURE 1

Enterprise Workloads Increasingly Shifting to Cloud-Based Environments

Q. Thinking of all your organization's applications, what percentage are currently deployed in the following venues today? (For 1Q21 sample: What about in two years?)



Note: For more details, see *Surviving and Thriving in a Multicloud World* (IDC #US48845822, February 2022), sponsored by Nutanix.

Source: IDC, 2022

A unified management interface that spans on-premises and off-premises software-defined infrastructure is a key enabler of hybrid multicloud environments. As enterprises modernize their infrastructure, they will need to support bare metal, VM, and container-based deployment models as well as different hypervisors, and the flexibility to support this variety should be a key consideration when crafting an optimized hybrid multicloud strategy. Other features that can be attractive here also include the ability to choose from a range of capex and opex consumption models for a particular platform and transferrable product licensing (between use of a software product on premises and in the public cloud) so that moving a workload does not require relicensing.

HCI: Leading the Evolution to Software-Defined Infrastructure


Traditionally, IT has been a very siloed environment with different storage infrastructure required to support different workload types and requirements. One of the key reasons software-defined infrastructures are growing so rapidly is because they provide configuration flexibility that allows multiple workloads with very different I/O profiles to be run on the same shared infrastructure while still

meeting performance, availability, and cost objectives. Denser storage workload consolidation is a goal of many infrastructure modernization projects and a factor in driving software-defined storage (SDS) deployments. Within the SDS market, HCI is the fastest-growing segment, growing at an 11.8% compound annual growth rate (CAGR) through 2025 when it will hit \$15.6 billion in revenue.

Recent survey data from IDC indicates widespread usage of HCI for mission-critical workloads: 74% of enterprises have replaced at least one SAN/NAS array with HCI on technology refresh and over 56% of IT organizations are running mission-critical workloads on HCI (see Figure 2). Several characteristics of HCI are driving this rapid growth. Its software-defined nature gives it great agility and makes it look very similar to public cloud-based infrastructure, a factor that makes it easier to migrate workloads between the two deployment models. HCI is easy to deploy, easy to manage, and easy to scale. And it offers better economics than more hardware-defined infrastructure that lacks flexible deployment options.

FIGURE 2

HCI Deployment Patterns in Enterprises



Source: IDC's *HCI Survey*, January 2022

HCI has a significant impact on application life-cycle management:

- **Day 0 operations**, which is the "design" stage, are simpler because they are typically done by carving resources out from an existing HCI cluster and because it is easy to adjust resource allocations virtually if they need to be modified. Most enterprises run multiple workloads on their HCI infrastructure and are able to logically apportion resources as needed.
- **Day 1 operations**, which encompass deployment, are simple because the required resources can easily be assembled into the components needed (compute, storage, and networking) from logical pools, and expanding HCI clusters to accommodate additional workloads is as easy as adding one or more nodes.
 - Application mobility is a consideration that can come into play during deployment, making it easy to move a workload around within a cluster for workload balancing or to move it to a different deployment location (either to or from the public cloud).
- **Day 2 operations**, which are the maintenance stage, are also simple because of HCI's software-defined orientation and modern design. HCI's single management interface allows individual workloads to be monitored, managed, diagnosed, and upgraded – rapidly and nondisruptively – based on a global view of all workloads.

Cloud Migration Use Cases and Business Justifications

Public cloud makes IT infrastructure resources available on demand for a number of use cases (see Figure 3). Disaster recovery was one of the original use cases for public cloud. It gave enterprises an offsite location to host backup data and also provided a remote site from which to recover operations in the event of a disaster that did not require that IT organizations own and maintain their own secondary datacenter. This represents immediate and very large cost savings for enterprises that are looking to consolidate operations and close down datacenters or for those that would otherwise be faced with opening up a new datacenter.

The public cloud also provides access to pay-as-you-go infrastructure on a short-term basis. Prior to the availability of the public cloud, enterprises had to size their IT infrastructures to meet the highest annual demands, an expensive approach that often left resources underutilized most of the time. Businesses that faced high seasonal demands, such as retail during the holiday season, were particularly affected by this issue. But with on-demand access to cloud-based resources, enterprises could size their on-premises infrastructure to accommodate lower demand, "bursting" to the cloud to meet short-term demands as necessary.

Enterprises only pay for these added IT resources when they are being used, an approach that saves them considerable cost in on-premises infrastructure spend. This IT elasticity also allows enterprises to rapidly provision infrastructure to support short-term and/or new or longer-term projects, paying for that infrastructure on an opex rather than a capex basis. For instance, hosting virtual desktop infrastructure (VDI) in the public cloud to scale desktops as new employees are added across a variety of different locations or leveraging the public cloud for rapid infrastructure buildout during supply chain shortages that limit the scaling of on-premises infrastructure are both recent examples of drivers of expansion in the public cloud.

This subscription-based pricing flexibility provides new financial options in paying for various project costs, regardless of whether those are for short-term projects or for what an enterprise wishes to permanently move to a more opex basis. It also gives enterprises the option to offload the costs of on-premises infrastructure management on an application-by-application basis. This might allow, for example, an enterprise to decide to keep more strategic applications on premises while moving others that don't require such close management to the public cloud.

The public cloud also provides options to continue to grow IT infrastructure for enterprises that are experiencing physical resource constraints in their datacenters around floorspace, cooling capacity, or energy consumption limits. IDC's 2021 *Datacenter Operational Survey* of United States-based organizations showed that one in four report delays in IT deployments due to power and/or space constraints. IDC is seeing this issue pop up more often for enterprises whose existing datacenters have reached (or are about to reach) capacity limitations in those areas.


Workload rehosting is another common use case. Many of the newer applications are being developed using public cloud infrastructure but are then moved to on-premises infrastructure for production deployments. New projects that haven't yet been funded as a long-term commitment may also start in the public cloud and then move back on premises as they move out of the pilot stage. Similarly, other workloads may start small in on-premises infrastructure but, as they grow, move to the public cloud to take advantage of higher capacities for more data-intensive applications. Workloads often evolve over time,

One in four United States-based organizations report delays in IT deployments due to power and/or space constraints.

and the optimal deployment location may change, driving a need to move a workload to the public cloud or back on premises and driving a need for bidirectional mobility.

FIGURE 3

Cloud Migration Use Cases



Source: IDC, 2022

Application blueprinting can help make the migration of applications much easier. An application blueprint maps application requirements to a resource template, including factors such as dependencies, connectivity, performance, and security requirements, effectively providing a model that makes it easier to move a workload to another deployment model like the public cloud. Legacy workloads can be blueprinted to make it easier to move them to modernized infrastructure, and new applications will often be developed using an application blueprint so that they are portable right from the start.

The capabilities of infrastructure on premises and in the public cloud differ, offering different generations of hardware with differing underlying capabilities in terms of performance, security, and cost. It's important to examine and evaluate those capabilities to ensure the infrastructure can support the application blueprint, wherever that workload is placed or moved over its lifetime.

Nutanix: The Foundation for Hybrid Multicloud Environments

Nutanix is a \$1.5 billion enterprise cloud infrastructure provider, headquartered in San Jose, California. Founded in 2009, the public company provides enterprise-class, software-defined infrastructure that spans both on-premises and off-premises deployment options in hybrid multicloud environments. Its core platform is based around HCI, but the vendor differentiates itself from the competition through its wide-ranging deployment flexibility, simplicity of operation and management, and outstanding customer experience (CX). The Nutanix Cloud Platform is a comprehensive solution that spans private and public cloud; supports block, file, and object storage options; accommodates bare metal, VM, container, and cloud-based application deployment as well as different hypervisors; and delivers the configuration flexibility to host literally any application – all managed under a unified control plane called Prism. The vendor's mission is to make clouds (both public and private) invisible, freeing IT organizations to focus on business outcomes.

Nutanix has thousands of enterprise customers that have consolidated datacenter, DevOps, and desktop services onto NCP. Supported workloads include:

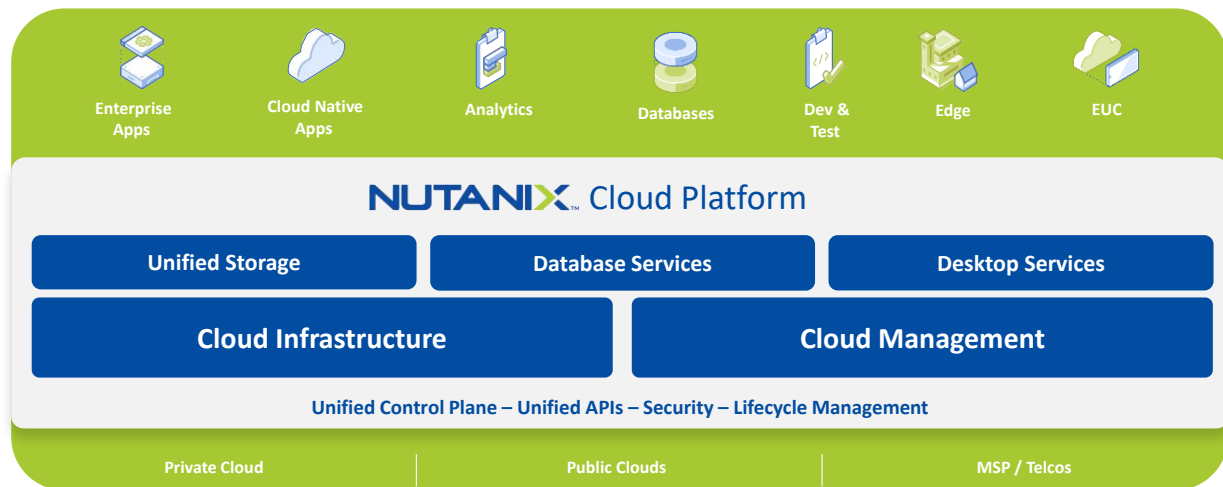
- Enterprise applications
- Cloud-native apps
- Big data and analytics
- Artificial intelligence and machine learning
- Databases
- Test and development
- IoT edge
- Remote office back office (ROBO)
- End-user computing (VDI)

Mission-critical applications are widely deployed on NCP as the solution supports the enterprise-class features (erasure coding [EC], snapshots, encryption, quality of service, replication, etc.) those workloads require.

As modernized infrastructure, NCP is often deployed to replace more hardware-defined platforms on technology refresh. One-click operations, an automation capability that continues to differentiate Nutanix, make administration fast, simple, and reliable. Once customers have NCP deployed, they can use Nutanix Self-Service to automate the provisioning and orchestration of day-to-day operations as well as new application deployments. They can consolidate multiple workloads with a variety of different deployment requirements (bare metal, VMware ESXi, Hyper-V, KVM, containers, different access methods, etc.) onto their NCP cluster, defining storage pools to meet individual application requirements. In 2015, Nutanix introduced its own hypervisor (AHV) as an additional no-charge virtualization deployment option, lowering costs in particular relative to other hypervisor-based vendor offerings (see Figure 4).

FIGURE 4

Nutanix Cloud Platform



Source: Nutanix, 2022

IDC has noted that it is easier to manage hybrid cloud environments when modernized, software-defined infrastructure exists in both on-premises and off-premises locations. The vendor's software-defined NCP has much in common with the web-scale infrastructure many cloud providers use – it uses the same hardware architecture including x86 processors and high-performance storage device types such as NVMe – a fact that makes it easier to build hybrid multicloud infrastructure. Hardware component-level commonality is also important when considering application and operating system portability, such as having common Intel Xeon Scalable processors between on-premises and public cloud infrastructures that eliminate the need to refactor. Prism, the vendor's cloud-based unified management console, provides visibility into all deployment locations, centralizing management onto a single pane of glass that provides a common dashboard from which to monitor and manage all IT infrastructure.

As a software-defined platform, NCP can be deployed on a variety of different on-premises and off-premises infrastructure. With the NCP software stack running in on-premises HCI, AWS EC2 bare metal instances, and Microsoft Azure dedicated hosts (slated for calendar 2H22) and/or service providers, customers have more flexible, easier-to-manage, and lower-cost on-premises infrastructure (relative to SAN and NAS arrays) for their enterprise workloads with seamless application migration to the public cloud. When coupled with a consistent underlying hardware infrastructure based on Intel architecture, this enables existing applications to migrate to the cloud without any code changes and enables the next generation of application development to be built alongside existing applications with tight integration into cloud-native services.

Nutanix has consistently produced one of the highest Net Promoter Scores (NPS) of any enterprise cloud software provider ever. NPS is an independent, de facto standard for measuring the quality of the customer experience across more than 220 different industries, based on a simple ongoing sampling of a vendor's current customers. Nutanix has produced an unheard-of score of 90+ (out of a possible 100) for the past seven years running, indicating an extremely happy customer base willing to

enthusiastically recommend Nutanix to their peers. For more information on NPS, see Figure 5, or *Net Promoter Score Becoming an Important Metric for Enterprise Storage Managers to Understand* (IDC #US43896818, June 2018).

FIGURE 5

The Net Promoter Score

Based on a customer's willingness to recommend a business to a friend that correlates with financial performance



$$\text{NPS} = \% \text{ PROMOTERS} - \% \text{ DETRACTORS}$$

Note: Scores span a scale from -100 to +100.

Source: IDC and www.netpromoter.com/know/, 2022

The Benefits of Intel Architecture in NCP Environments

Intel's server-based storage nodes were the first platforms on which Nutanix software-defined infrastructure ran, and the two companies have had a long-standing partnership. Intel's Xeon Scalable processor forms the heart of these hardware platforms for NCP today and includes technologies like Intel Virtualization Technology (Intel VT) that are specifically designed to improve the portability of VMs between clouds and across hardware generations to maximize resource utilization in hyperconverged environments. Intel Xeon Scalable processors also are the mainstay of public cloud-based services, with more than 50,000 different cloud instance (i.e., service) types available globally across the leading cloud service providers, offering a more flexible migration path through a broad range of performance options, integrated AI technology for deep learning, and security technologies integrated down to the silicon layer to meet a variety of different business requirements.

Intel VT is a portfolio of features, now in its sixth generation of production availability, that provide hardware assist to the virtualization software, driving down cost and complexity and improving infrastructure efficiency. The current portfolio includes CPU, memory, I/O, graphics, security, and network virtualization enhancements, and while it offers a very broad value proposition that impacts many areas, it is important to note in this discussion around application mobility that it enables both easier migration between deployment models (e.g., on-premises HCI to public cloud) and live migration of workloads from one Intel CPU generation to another (which makes technology refresh and other upgrades easier and less risky).

Intel also offers high-performance, tiered memory and storage devices that are fully supported by Intel Xeon Scalable processor-based servers. Branded as Intel Optane Technology, these devices include both persistent memory (Intel Optane Persistent Memory) and storage-class memory (Intel Optane Data Center SSDs) devices, offering performance and cost that lie between expensive, low-capacity but very low-latency DRAM and less expensive, higher-capacity but higher-latency NAND flash-based SSDs. These devices interface with standard Intel Xeon Scalable processor-based servers through either DIMM or PCIe/NVMe interfaces and can add significantly higher memory capacity at a lower cost or high-performance and low-latency storage for those workloads that may benefit.

Nutanix Cloud Clusters

The Nutanix Cloud Clusters solution allows the entire NCP software stack, with its full complement of enterprise-class capabilities, to be deployed in public cloud-based environments. The Nutanix offering includes other features of interest such as single console management using Prism, transferrable licensing applicable to both on-premises and off-premises deployment models, and a one-click "hibernate and resume" capability in NC2 on AWS that allows administrators to quickly and easily start or stop cloud instances. On the public cloud, hibernating instances lowers costs while still providing the ability to rapidly spin resources up on demand if and when they are needed.

Simply put, NC2 dramatically reduces the operational complexity of migrating, extending, and/or bursting workloads (including both applications and data) between on-premises and off-premises deployment locations. With NC2 running both on premises and in the public cloud, applications can be quickly and easily migrated between locations without any retooling or code changes, making it easy to burst additional capacity to meet seasonal needs, spin up a disaster recovery in the wake of a site failure, or just generally move workloads around for rebalancing or rehosting. NC2 features built-in integration with public cloud networking, promoting operational simplicity when moving workloads in addition to providing frictionless access to native cloud services for application modernization.

Nutanix customers use native migration tooling to rehost tens of thousands of workloads to new locations each year, moving from traditional IT infrastructure or cloud-native VMs into NCP, both on premises or on NC2 in the public cloud. For planned migrations, downtime can be limited to literally just a few minutes. Preparatory steps include data seeding where data is copied in the background to a new location and mapping target to source networks. Once that is complete, the workload can be migrated using a one-click migration cutover. Applications are not available during the migration, but downtime is minimal and Nutanix lets administrators pre-validate migrations through testing on isolated networks to ensure the fastest cutover.

Customers can migrate single VMs or group VMs and migrate them together. VM migrations to NC2 are based around a Nutanix feature (Nutanix Move) that has been in use for over four years and has proven its ability to reduce risk during application migrations. Some Nutanix customers use NC2 as an on-ramp to cloud-native computing. They modernize their infrastructure by moving to HCI on premises, rehost applications into VMs and/or containers, and then ultimately move some of those workloads to the public cloud, gaining public cloud benefits immediately without requiring any application code changes. In this way, enterprises do not have to refactor and/or rearchitect applications to gain the benefits of public cloud deployment.

Choice of Consumption Models

With NCP deployed across both on-premises and off-premises environments, it is important to Nutanix customers to be able to choose from several different licensing models. NCP purchase options include outright purchase of right-to-use licenses, on-demand consumption using a pay-as-you-go model, and a cloud commit approach that requires a term-based minimum annual commitment. Licenses are transferable between on-premises and off-premises deployment models so customers can migrate a Nutanix workload from a private cloud to the public cloud without having to relicense Nutanix software. Since Nutanix licenses by physical CPU cores, if customers increase capacity after a move, they would need to adjust their licensing.

Nutanix products are available from several of the major storage OEMs that have their own subscription-based licensing models. Under programs like HPE GreenLake, and Lenovo TruScale, enterprises can also provision NCP resources using operational rather than capital expenditure models.

CHALLENGES/OPPORTUNITIES

Nutanix's portfolio of products and services provides a smooth transition to a hybrid multicloud strategy, but there are still some workloads that enterprises prefer to retain on disaggregated storage infrastructure. As newer, faster compute and storage technologies get integrated into HCI offerings, there are fewer and fewer workloads where this needs to be true, and the flexibility, ease of use, and compelling economics of HCI make it a very attractive option. Intel delivers a regular cadence of technology advancements, and HCI solutions like Nutanix running on standard servers are able to rapidly take advantage of the new capabilities. Nutanix customers already run many mission-critical workloads on NCP, and NC2 lets customers deploy and/or move those workloads to the public cloud without having to give up the enterprise-class capabilities they demand. Nutanix needs to ensure that lack of awareness is not a reason for enterprises not to consider a move to software-defined modernized infrastructure offerings based around NCP.

CONCLUSION

As enterprises undergo digital transformation, they are modernizing their IT infrastructure, using more automation, leveraging public cloud, and looking for unified management planes that will give them a single pane of glass from which to manage their hybrid multicloud infrastructure. Enterprises are also thinking more about the optimal deployment model for their workloads, and since this may change over the life of an application, they are taking advantage of the software-defined flexibility of modernized infrastructure platforms like NCP to simplify application mobility. NC2, a capability within NCP, meets this requirement with a solution that delivers operational simplicity, seamless application mobility, and cost efficiency. Key use cases for NC2 and application mobility in general include cloud bursting, disaster recovery, and workload rehosting, and it is important that NC2 enable all of this without requiring any application code changes. NC2 ensures that the same enterprise-class capabilities provided by NCP are available to workloads regardless of deployment location, providing not only the features needed to densely consolidate workloads of all types but also a centralized dashboard with Prism for unified management.

Enterprises using HCI platforms like NCP are clearly consolidating workloads and care about the infrastructure efficiency and security improvements Intel brings to the table with the latest generation of Intel Xeon Scalable processors, including Intel VT, as well as Intel Optane memory and storage products. Together, Nutanix and Intel technologies provide highly efficient modernized infrastructure for today's on-premises and multicloud-based workloads cooperating together in hybrid cloud environments.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2022 IDC. Reproduction without written permission is completely forbidden.

