



Dell XC Series Appliances – Reference Architecture for Deploying VMware vRealize Automation

Dell Storage Product Group
August 2015

Revisions

Date	Description
August 2015	Initial release

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Copyright © 2015 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.



Table of contents

Revisions.....	2
Executive summary	5
1 Introduction.....	6
Audience.....	6
Purpose.....	6
2 Overview.....	7
Solution Architecture	7
1.1 Application Overview	10
1.2 What is vRealize Automation?	10
2.1 vRealize Automation Capabilities.....	10
3.1 vRealize Automation Center the Nutanix Way	11
3.2 Data Tiering and Management	12
3.3 XC Appliances and vRealize Automation.....	14
3.4	
5 Architecting vRealize Automation on XC Series appliances	16
Architectural Elements.....	16
5.1 VMware Architectural Elements.....	17
5.2	
5.2.1 vCenter Server	17
5.2.2 vRealize Application Services.....	18
5.2.3 vRealize Business	18
5.2.4 vRealize Orchestrator.....	18
5.3	
5.2.5 vSphere Cluster	18
6.1 Infrastructure	18
6 vRealize Automation Solution Architecture.....	20
Deployment components	20
6.1.1 Dell XC System	20
6.1.2 XC Storage Pool	20
6.2	
6.3 6.1.3 Storage containers.....	21
6.1.4 vSphere Cluster Including XC appliances.....	22
6.1.5 Dell XC Platform and vSphere Cluster Relationship	22
Deployment Logical Component Communication Structure.....	23
vRealize Automation Center Component Placement.....	23



7	Deployment profiles.....	25
	Deployment Profile: Proof of Concept or Test Environment.....	25
	Deployment Profile: Small Production	27
	Deployment Profile: Medium Production	28
	Deployment Profile: Large Production	32
8	Scalability.....	35
7.1	Scalability Example	35
7.2	Scalability Input Parameters.....	35
7.3	8.2.1 Software-Based	35
7.4	8.2.2 Corporate Data center	36
8.1	8.2.2 Corporate Data center	36
8.2	9 vRealize Automation Deployment Tests	37
	General Test Setup.....	37
	9.1.1 Hardware	37
9.1	9.1.2 Software.....	37
	9.1.3 Virtual machine.....	37
	Clone Test	38
9.2	9.2.1 Results	38
9.3	Linked Clones	38
	9.3.1 Linked Clone Test	38
9.4	9.3.2 Linked Clone Test Results	38
	Test Results Summary	39
10.1	9.4.1 Provisioning Time	39
10.2	Deployment Profile Hardware Requirements	40
	vRA Components	40
	Additional Components.....	41
11	Conclusion	42
A	Appendix.....	43
	A.1 References.....	43
	A.2 List of figures.....	43
	A.3 List of tables	44
	A.4 Further Research	45



Executive summary

Dell XC Series powered by Nutanix hyper-converged infrastructure with web-scale technologies is a scalable virtualization platform for cloud, desktop, business-critical applications and big data deployments. A Dell XC Series solution contains a minimum of three nodes or physical servers running a hypervisor that includes both compute and storage resources. The Nutanix web-scale software architecture greatly reduces costs while increasing performance and scalability, empowering you to deliver a reliable and consistent user-experience.

Utilizing the Dell XC Series platform alleviates many of the complications of traditional storage. All storage management is virtual machine-centric. There is one shared pool of storage that includes flash-based SSDs for high performance and low latency, and high-capacity HDDs for affordable capacity. Data is automatically tiered across the different types of storage devices in the storage pool using intelligent placement algorithms. This ensures the most frequently used data is available in memory or in flash for the fastest possible performance.

This document offers recommendations for implementing VMware vRealize Automation* (vRA) deployments on the Dell XC Series of web-scale converged infrastructure. It demonstrates architectures that take advantage of the combined values of Nutanix scaling, performance, and simplicity, while utilizing the capabilities of vRA when providing datacenter infrastructure-as-a-service. To this end, we present deployment scenarios at test, small, medium, and large scales, indicating for each the type and amount of hardware, software, and virtual machines needed.

Note: This reference architecture draws on VMware's vRealize Automation reference architecture throughout, including, for example, such parameters as deployment types and the required number of vRealize Automation management components.

* VMware vRealize Automation was formerly marketed and sold as vCloud Automation Center (vCAC).



1 Introduction

Audience

This reference architecture document is intended for use by individuals responsible for architecting, designing, managing, and/or supporting VMware vRealize Automation (vRA) on Dell XC Series appliances. Users of this document should be familiar with concepts pertaining to VMware vSphere, VMware vRealize Automation, and Nutanix software.

1.1

The document is presented in sections to highlight deployment scenarios and addresses key items for a successful design and implementation of vRA on Dell XC Series appliances.

Purpose

1.2

The focus of this reference architecture is the development and deployment of vRA management components on the Dell XC Series powered by Nutanix platform. For easier reference and comparison, it utilizes the same deployment profiles as the VMware vRA Reference Architecture [1]

This document will cover the following subject areas:

- Dell – "Nutanix" solution overview
- vRealize "Automation" overview
- Architecting a vRealize "Automation" solution on "Dell" "XC" "Series" appliances

In addition, vRealize Application Services (vRAS) (formerly vCloud Application Director) will be included from a sizing perspective, but not discussed in detail regarding functionality.



2 Overview

Solution Architecture

The Dell – Nutanix Virtual Computing Platform is a scale-out hyper-converged system of high-performance nodes, or servers, each running a standard hypervisor and containing processors, memory, and local storage (consisting of SSD Flash and high capacity SATA disk drives). Each node runs virtual machines just like a standard virtual machine host.

2.1




Figure 1 Node Architecture

In addition, local storage from all nodes is virtualized into a unified pool by the Nutanix Distributed File System (NDFS). In effect, NDFS acts like an advanced networked or shared storage system that uses local SSDs and disks from all nodes to store virtual machine data. Virtual machines running on the cluster write data to NDFS as if they were writing to shared storage.

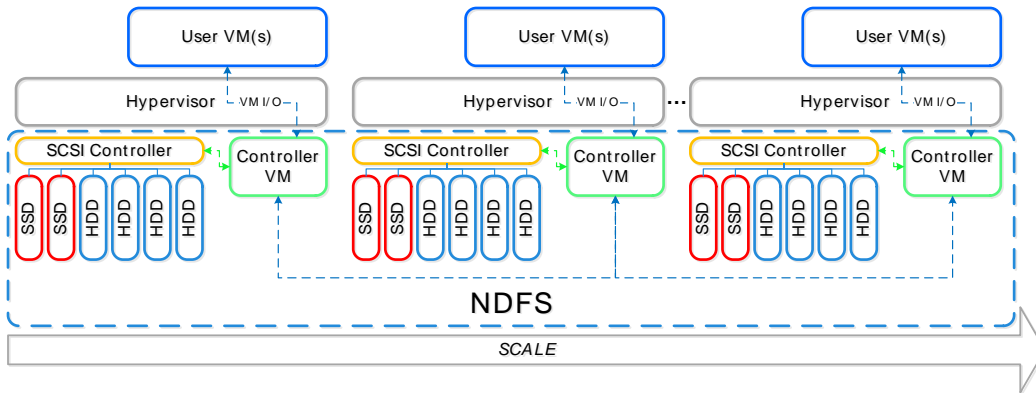


Figure 2 Dell – Nutanix Solution Architecture

NDFS is VM-centric and provides advanced data management features. It brings data closer to virtual machines by storing the data locally on the system, resulting in higher performance at a lower cost. Dell XC Series infrastructure can horizontally scale from as few as three nodes to a large number of nodes, enabling organizations to scale their infrastructure as their needs grow.



The Controller VM and NDFS deliver a unified pool of storage from all nodes across the cluster, leveraging techniques including striping, replication, auto-tiering, error detection, failover, and automatic recovery. This pool is then presented as shared storage resources to XC Series nodes or ESXi hosts for seamless support of features, including vMotion, HA, and DRS, along with industry-leading data management features. Additional nodes can be added in a plug-and-play manner in this high-performance scale-out architecture to build a cluster that scales to meet the needs of the business.

The Nutanix Elastic Deduplication Engine is a software-driven, massively scalable and intelligent data reduction technology. It increases the effective capacity in the disk tier, as well as the RAM and flash cache tiers of the system, by eliminating duplicate data. This substantially increases storage efficiency, while also improving performance due to larger effective cache capacity in RAM and flash. Deduplication is performed by each node individually in the cluster, allowing for efficient and uniform deduplication at scale. This technology is increasingly effective with full/persistent clones or P2V migrations.

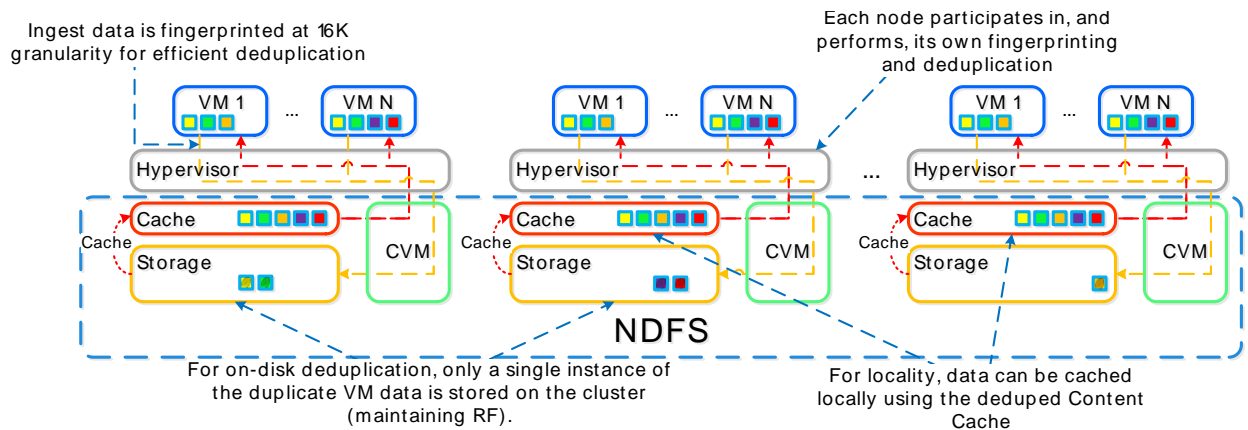


Figure 3 Elastic Deduplication Engine

The NDFS Shadow Clone feature allows for distributed caching of vDisks or VM data that is in a 'multi-reader' scenario including deploying private cloud environments. This will allow VMs on each node to read the Base VM's vDisk locally instead of forwarding read requests to a master 'Base VM'. In the case of VDI, this means the base disk can be cached by each node and all read requests for the base will be served locally. In the case where the Base VM is modified, the Shadow Clones will be dropped and the process will start over.



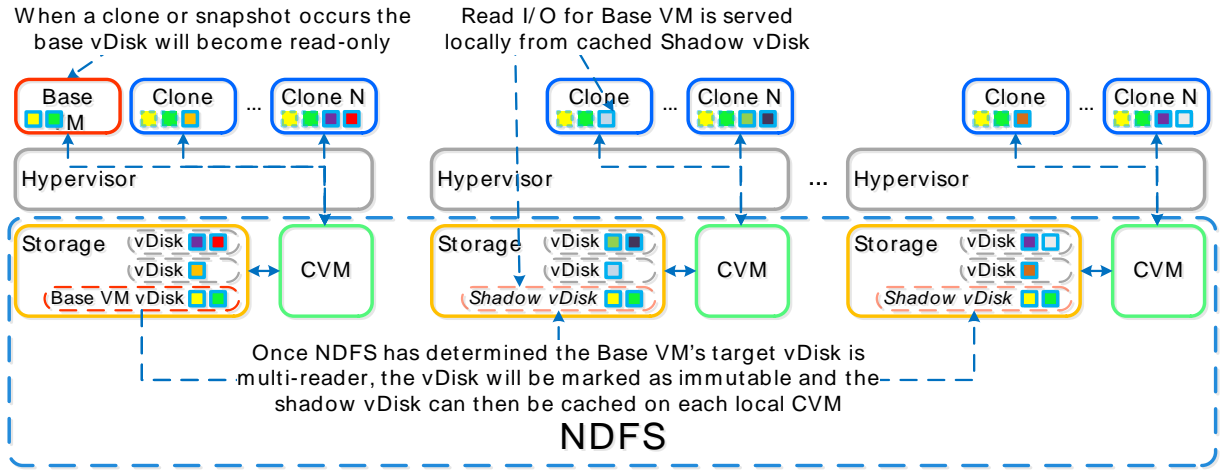


Figure 4 NDFS Shadow Clones



3 Application Overview

What is vRealize Automation?

VMware vRealize Automation (vRA) helps information technology (IT) departments with delivery and ongoing management of infrastructure, applications, and custom services, which can lead to lower operating expense (OpEx). vRA can be used to build both enterprise and service provider cloud solutions.

3.1

Supporting implementations for private clouds, public clouds, and hybrid clouds, vRA can be used in conjunction with existing or new business processes and tools to leverage a very flexible, highly automated and intuitive self-service experience for the end user.

This makes it easy to maintain control over provisioned physical and virtual workloads via its centralized provisioning structure, governance, and infrastructure management capabilities. For example, lifecycle management of virtual machine workloads is controlled via vRA through built-in processes for:

- Request
- Approval
- Provisioning
- Management
- Reclamation
- Decommission

This document will focus on the infrastructure-as-a-service portion of vRA by covering the components listed above. vRA can also be used to deliver both infrastructure-as-a-service (IaaS) and platform-as-a-service (PaaS), when used in conjunction with vRealize Application Services (vRAS). Additionally, the vRA portal offers you the ability to connect to other systems and to create everything-as-a-service (XaaS). This document will solely focus on the IaaS part of vRA but include XaaS in figures to illustrate vRA's full capability.

3.2

vRealize Automation Capabilities

As a converged infrastructure solution, the Dell – Nutanix solution simplifies the datacenter and eliminates the complexity of legacy storage architectures.

vRA enables IT teams to lower OpEx and gives end users the ability to rapidly provision virtual machines in a secure, standardized manner.

The following VMware vSphere capabilities are supported by Nutanix:

- Support for key VMware vSphere capabilities include VMware Distributed Resource Scheduler (DRS), VMware High Availability (HA), VMware vMotion, Performance and Resource Optimization, VMware vSphere Storage API for Array Integration (VAAI).
- NFS protocol support.



- Converged infrastructure eliminates the cost and complexity of a separate SAN or NAS-based storage.
- Provides enterprise-class storage features, including thin provisioning, snapshots, fast clones, compression, and de-duplication.

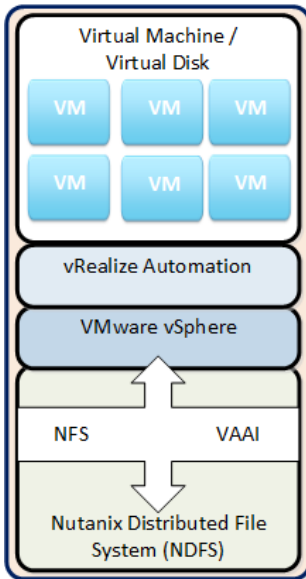


Figure 5 Logical overview of Dell XC Series

3.3

vRealize Automation Center the Nutanix Way

A 2U XC appliance contains one high-performance server node with local storage to run and store virtual machines. SSD Flash combined with intelligent tiering of data deliver the high performance of SSDs at the cost of hard drives to tackle vRA performance requirements during cloud service peak hours. This design greatly reduces overall cost and complexity while increasing performance and scalability in a fully integrated multi-workload solution.

A linear scale-out architecture enables customers to select any initial deployment size and grow in more granular data and compute increments. This removes the hurdle of large up-front infrastructure purchases that customers will need many months or years to grow into, ensuring a faster time-to-value for the implementation.

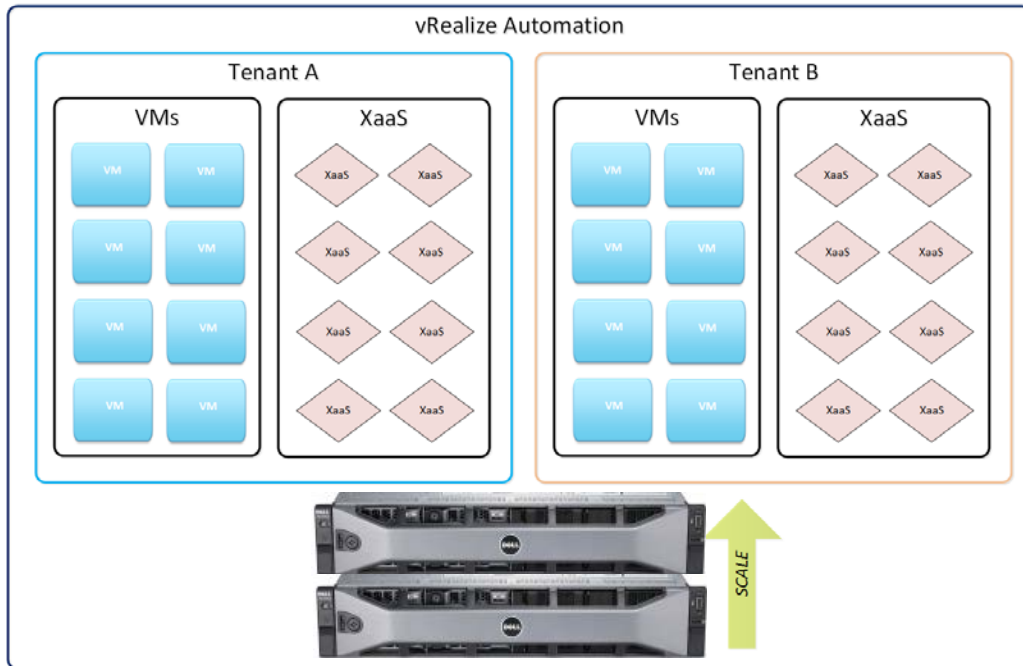


Figure 6 vRA Conceptual Architecture

The Dell – Nutanix solution is fully integrated and supported with the VMware APIs for Array Integration (VAI) to enable you to provide the best possible experience to the end user with the flexibility of a single modular platform.

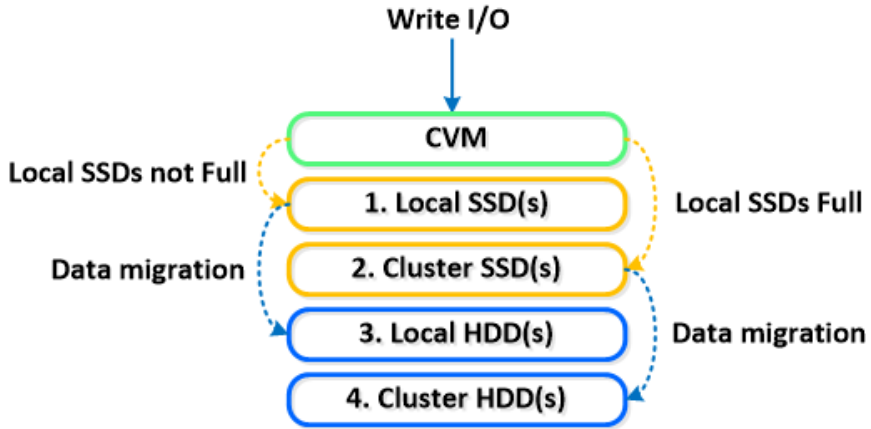
3.4

Data Tiering and Management

The Nutanix Distributed File system (NDFS) has a built-in intelligent tiering process that automatically handles data placement. All hot data is automatically written to the high performance SSD tier. Both hot and warm data sit in this tier to be readily accessed and provide the highest performance. An in-memory read cache is leveraged to cache frequently accessed data from all tiers.

As virtual machine data becomes less frequently accessed, Nutanix data tiering will automatically detect the data cooling down and migrate the data from SSD to the higher-capacity HDD tier in order to free SSD space for new hot data. This data can then be automatically compressed after a specified period of time to allow for increased storage capacity. In the event this data becomes accessed again, it will automatically become un-compressed and placed on the appropriate tier. This keeps the most highly accessed data in the cache and/or highest performance tier.

A local ESXi host's SSD tier is always the highest priority tier for all I/O generated by virtual machines running on that ESXi host. However the entire cluster's SSD resources are made available to all ESXi hosts within the cluster. The storage tier prioritization can be visualized at a high-level by the following diagram:



*NOTE: Sequential IO can be configured to bypass SSD and be directly written to the HDD tier.

Figure 7 Nutanix Write I/O

4 XC Appliances and vRealize Automation

Dell XC Series appliance enable you to run multiple workloads all on the same scalable converged infrastructure.

- Modular "incremental" scale: "With XC Series web-scale converged infrastructure, you can start small and scale. A single XC Series Node provides from 5 TB to 32 TB storage and up to 32 cores in a compact footprint. With the modularity of the solution, you can granularly scale per-node (ESXi host), giving you the ability to accurately match supply with demand and minimize the upfront capital expenditure (CapEx). The XC Series platform's scalability options are a perfect match for vRA, which in many situations, starts on a small scale and scales up based on public and/or internal demands."




Figure 8 Dell XC Series Conceptual Scalability

- Operating expense: "the Dell XC Series Solution helps organizations lower OpEx by making day-to-day operations simpler, and because of its converged architecture, Nutanix also makes it possible to overcome organizational silos."
- High performance: "it provides up to 100,000 random read IOPS and up to 3 GB/s of sequential throughput in a compact 2U footprint. ILM keeps indexes and heavily access data in the high performance SSD and cache tiers."
- Integrated: "the Dell XC Series Solution provides full support for VAAI, allowing you to leverage all the latest advancements from VMware and taking your solution to the next level. One of the out-of-the-box vRA capabilities includes cloning of existing vSphere-based virtual machines. The VAAI functionality will contribute to rapid virtual machine creation and provisioning."
- Data efficiency: "the Dell XC Series Solution is truly virtual machine-centric for all compression and deduplication policies. Unlike traditional solutions that perform these tasks mainly at the LUN level, the Nutanix software provides all of these capabilities at the virtual machine and file level, greatly increasing efficiency and simplicity."

By allowing for both inline and post process (MapReduce) compression and deduplication capabilities and cache, the Dell XC Series Solution breaks the restrictions set by traditional solutions.

- Effective "transparent" data tiering: "the XC solution incorporates heat-optimized or activity-based tiering, which leverages multiple tiers of storage and optimally places data on the tier that provides the best performance. The architecture was built to support local disks attached to the controller virtual machine (SSD, HDD) as well as cloud-based source targets. The tiering logic is fully extensible, allowing new tiers to be dynamically added and extended. The XC Series solution continuously monitors data-access patterns to determine whether access is random, sequential, or a mixed workload. Random I/O workloads are maintained in an SSD tier to minimize latencies. Sequential workloads can be automatically placed into HDD to improve SSD endurance.
- Business continuity and data protection: Native VMCaliber (VM-centric) snapshot and replication features provide an extensive disaster recovery (DR) and protection capability at the individual VM level. Windows volume shadow copy (VSS) provides integration for application consistent snapshots.
- Enterprise-grade system management: Through its Prism management framework, the XC solution delivers a simplified and intuitive way to manage large system, including a converged user interface (UI) that serves as a single pane of glass for servers, storage, alert notifications, and a bonjour mechanism to auto-detect new nodes in the cluster.
- New compute and storage: Capacity can be made available very simply through the Prism management interface and the vRA UI compared to other underlying physical infrastructures where integration with multiple UIs is often necessary. Additionally, multiple different systems can be managed through the Nutanix Prism Central tool, eliminating the need to have multiple sessions to manage multiple different systems.
- High-density architecture: the XC Solution uses an advanced server architecture in which 2 Intel CPUs (up to 32 cores) and up to 768GB of memory are integrated into a single 2U appliance, coupled with data archiving and compression.
- Automation: The ability to integrate vRA, vRealize Orchestrator (vRO) formerly vCenter Orchestrator, Nutanix Prism API and PowerShell makes it possible to create a fully automated environment including e.g.
 - Dell XC Series node/ESXi host provisioning
 - vRA capacity extension
 - Virtual machine backup and restore capabilities.
 - Third party integrations with e.g. Documentation system, IPAM, monitoring system and service ticket system,
- Multi-hypervisor support: The infrastructure as a service (IaaS) part of vRA can manage the lifecycle of virtual machines on multiple hypervisors. This means it is a perfect overlaying provisioning engine for the Dell XC solution infrastructure that also supports multiple hypervisors, such as Microsoft Hyper-V, KVM, and VMware ESXi.



5 Architecting vRealize Automation on XC Series appliances

This section describes the components, logical and physical, that will be included when building a vRA-based infrastructure with Dell XC Series appliances.

Architectural Elements

Table 1 Nutanix Design Elements

Item	Description
XC Block	The physical chassis where a Dell XC Series appliance resides.
XC System	A logical group of XC appliances grouped together based on availability, capacity, and performance requirements to provide compute (CPU and RAM) and storage capacity.
XC Appliance	Physical servers running the ESXi hypervisor.
Storage Pool	A group of physical storage devices, including SSD and HDD from the XC appliances within the cluster.
Container	A logical segmentation of the Storage Pool that contains one or more virtual machines and or files. This logical storage unit will be presented to the ESXi hosts as a NFS Datastore.

5.1



Figure 9 Dell XC Series block

The following table outlines components that are included in the vRA. Depending on deployment profile, some of the components can coexist on the same virtual machine.

Table 2 vRA elements

Item	Description
vRA Identity Appliance	A VMware virtual appliance providing end-user identity management and single sign on (SSO) functionality. The Windows-based vCenter SSO installation can be used in addition to the vRA Identity appliance.
vCenter Single Sign On (SSO)	vCenter SSO can be used in active-passive mode to enable high availability for the identity management component. Follow reference [8] to setup high availability for vCenter SSO. No session information is exchanged between the vCenter SSO nodes meaning the user might see a service interruption during a failover.
vRA Appliance	Provides the vRA portal functionality and governance functionality. In addition, it includes both a PostgreSQL database and a vRO. The PostgreSQL can be used in both a non-highly available implementation and in a



Item	Description
	highly available implementation according to the VMware knowledge base article [6].
Infrastructure Web Server	laaS component required for virtual machine provisioning.
Infrastructure Manager Server	laaS component required for virtual machine provisioning.
Distributed Execution Manager (DEM) Orchestrator	laaS component responsible for scheduling workflows and monitoring the state of the DEM workers.
Distributed Execution Manager (DEM) Worker	laaS component responsible for executing workflows.
Infrastructure Agent Server (Proxy Agent)	Connection to the vSphere cluster where virtual machines will be deployed.

VMware Architectural Elements

5.5.2.1 vCenter Server

vCenter Server comes in an appliance version and a Windows server installable version. vCenter Server is used to manage the ESXi hosts and the logical vSphere clusters. For availability, performance, and separation purposes, a minimum of two vCenter Servers should be implemented per deployment profile. One vCenter Server is used for the vRA management cluster, and at least one, depending on deployment size, vCenter Server is used for vRA resource workloads, also known as vRA-deployed virtual machines. The individual server VMs are described and depicted in Section 8 of this document.

The only exception would be the proof of concept (POC) / test environment deployment profile where one vCenter Server is enough..




Figure 10 Logical vCenter Server to vRA mapping

5.2.2 vRealize Application Services

vRealize Application Services (vRAS) is the component in the vRealize Suite that helps requestors visually map blueprints of applications that can be repeatedly deployed across various cloud environments.

5.2.3 vRealize Business

vRealize Business, formerly ITBM, is for presentation of costs and uses an industry reference database for calculating costs such as cloud total cost, average cost per virtual machine, and business unit demand.

5.2.4 vRealize Orchestrator

vRealize Orchestrator (vRO) is used for many cloud deployments to create custom workflows interacting with governance, provisioning, documentation, monitoring, service ticketing, IP address management (IPAM), backup and restore systems.

vRO can use the rich set of functionality exposed by the Nutanix Prism API to create end-user or administrator-specific vRA advanced services. vRO can also interact with other Nutanix products, including Foundation, meaning a wide variety of services can be created:

- Data protection as a service making it possible for end users to control backup and restore of their virtual machines
- XC Appliance setup and configuration
- Nutanix management
- Nutanix Controller Virtual Machine management

vRO is included in the vRA Appliance, available as a virtual appliance or available as a Windows server installable version. Depending on tenant separation requirements, one vRO instance can be setup per tenant.

5.2.5 vSphere Cluster

A vSphere cluster is a logical vCenter Server grouping of ESXi hosts. This logical component provides availability and physical load balancing capabilities based on CPU and RAM utilization. vSphere Cluster resource pools (RPs) are logical groupings of CPU and RAM resources that can be used to separate workloads and give priority to one type of workload over another.

5.3

Infrastructure

The components mentioned in this section can be kept in the vRA management cluster together with the vRA management components for administrative and availability reasons. Dell recommends that at least one LDAP server is placed in the vRA management cluster for the small, medium and large deployment profiles.

Table 3 vRA Infrastructure Items and Descriptions



Item	Description
Administration Server(s)	Used for cloud platform administrators
Domain Name Servers (DNS)	Provides name resolution functionality for vRA management cluster components. The vRA resource cluster components can also take advantage of the service if needed.
Lightweight Directory Access Protocol (LDAP) Server/s	Used by vRA management cluster components. The vRA resource cluster components can also take advantage of the service if needed
Network Time Protocol (NTP)	Used for appropriate time synchronization.
Simple Mail Transport Protocol Servers/s	Used by vRealize Automation management components. The vRA resource cluster components can also take advantage of the service if needed
Syslog	Standard for logging system messages



6 vRealize Automation Solution Architecture

This reference architecture will use the VMware vRA reference architecture [1] existing deployment profiles small, medium, and large. In addition, there will be an extra deployment profile, POC / test environment, for vRA on Nutanix evaluations.

The differences between the vRA outlined deployment profiles are deployed virtual machines, concurrent deployments, and service availability.

Dell advises following the outlined recommendations in [1] regarding all aspects where it is applicable to your organization's conditions, constraints, and requirements.

The deployment profiles assume that VMware vSphere will be used for both vRA management workloads and vRA resource workloads. Additional hypervisors are supported by vRA as described in section 4.1: What is vRealize Automation.

Deployment components

6.1.1 Dell XC System

The reference architecture strategy can contain one of two Dell XC system options: A single Xc appliance cluster, or multiple clusters. This section will help you decide which option is right for you.

6.1.1.1 Single Cluster Strategy

When using the one XC cluster strategy you will see increased Write Performance as write I/O is distributed across more XC Appliance Controller VMs. In addition to higher resiliency, one cluster can lower the impact of failures and also enable faster self-healing capabilities. A large cluster also enables increased utilization of available storage capacity and performance.

6.1.1.2 Multiple Cluster Strategy

A multiple XC cluster strategy can be used when existing requirements enforces physical separation between different environment such as management components and user components or based on availability zone, network segment or security zone like a Demilitarized zone (DMZ), internal development, internal production, or internal testing.

Separation between systems during system tuning and system upgrade also applies to the multiple XC cluster strategy.

6.1.2 XC Storage Pool

Dell recommends always using a one to one mapping between XC appliance clusters and XC Storage Pools.



6.1.3 Storage containers

Each XC container can be configured with compression or deduplication enabled to increase performance and or reduce disk space usage.

Compression of data occurs during or after write operations and tests by Nutanix indicate disk space usage reduction around 30%. However, test must be conducted per environment to get each environments specific space usage reduction figures.

Two types of compression are available.

- Inline"compression" – "Data" "is" "compressed" "as" "it" "is" "written."
- Post-process"compression" – "Data" "is" "compressed" "after" "it" "is" "written."

Both types of compression can be used for most workloads when deduplication is not applicable. However, in-line compression is available in all Nutanix software editions and is the recommended compression method. See application specific Nutanix best practice guides for recommendations when to use post-process compression.

Deduplication enables sharing of common blocks on performance storage tiers (RAM and Flash) and starting in Nutanix OS 4.0 on the capacity tier. If guest virtual machines are substantially similar, for example if the XC cluster is used to host numerous persistent Windows desktops or common server virtual machine images, enabling inline deduplication on the performance tier (SSDs) can substantially improves performance.

vRA does not support direct VMware Horizon View virtual desktop provisioning, meaning the reasons for creating a XC container with deduplication enabled are if vRA will:

- Deploy"identical"virtual"machines"from"templates.
- Initiate"deployment"of"virtual"machines"via"third"party"systems"using"a"similar"image-based" approach"of"creating"virtual"machines"as"vRA.

An XC container where compression and deduplication is not enabled will be referred to as a default container throughout the document.

Dell recommends using the following container setup for all deployment types.

Table 4 Container Recommendations

Container	Container feature enabled	Target workload
1	Compression (In-line)	Management cluster – All components
2	Compression (In-line)	Resource cluster – All components

Increase CVM RAM capacity when using compression as documented in the Nutanix administration guides.

Nutanix provides guidance for a number of application workloads. If an application is running in the vRA environment, make sure the container type recommended does not deviate from the general XC



container configuration above. If it does, create containers that match the Nutanix recommendation for the specific application.

Recommendations for containers can change based on the workload.

- Archive data and big data - A dedicated container should be used for this data type, separate to the Guest Operating system and other virtual machines.
- Do not use in-line compression for applications that natively compress data. For example, you would not compress Microsoft SQL server data.

There are however situations when a XC container with deduplication enabled should be considered for vRA resource workload. These include when small virtual machines, disk size equal to or smaller than 20 GB, are deployed across XC containers (not using VAAI) and are not using the vRA linked clone deployment feature.

A suggested tunable redundancy or replication factor (RF) will be included for each XC container per deployment profile. The available RFs are 2 and 3 where RF2 keeps two copies of the data and takes less disk space compared to RF3 that keeps three copies of the data. However, RF3 brings additional availability compared to RF2 since RF3 allows the XC cluster to survive two simultaneous XC appliance failures without impacting data availability.

VAAI works within one XC container and will not come into play when virtual machines are deployed across XC containers. It would be beneficial from a virtual machine provisioning and timing perspective to have one vSphere template per XC container if vRA uses vSphere-based templates when deploying virtual machines.

However, this can potentially add complexity to the vRA provisioning process.

6.1.4 vSphere Cluster Including XC appliances

The minimum number of vSphere clusters and XC appliances will be outlined per deployment type. The required number of vSphere clusters and XC appliances per deployment type or maximum number of vRA managed virtual machines will not be outlined since it will vary considerably depending on virtual machine hardware configuration, company rules and company policies.

6.1.5 Dell XC Platform and vSphere Cluster Relationship

There are pros and cons for running one or more XC clusters. There are two options in terms of XC cluster to vSphere cluster configuration.

A single XC cluster to many vSphere clusters relationship provides all the features mentioned in single XC cluster strategy in section “7.1: Deployment Components - XC cluster”. If the container is setup to use RF3, the system will ensure data availability if a XC appliance becomes unavailable in each of the two different vSphere HA clusters running on top of one XC cluster.

A single XC cluster to single vSphere cluster relationship can be used for the same reason as mentioned in multiple XC cluster strategy in section “7.1: Deployment Components – Multiple Cluster Strategy.” If using RF2 a one-to-one relationship will guarantee that data will always be available if one XC appliance (ESXi

host) is unavailable. If using RF3 you can guarantee data availability if two XC appliances (ESXi hosts) become unavailable in one vSphere HA cluster running on top of one XC cluster.

Deployment Logical Component Communication Structure

The following figure outlines the communication structure between the components included in the vRA setup.

6.2

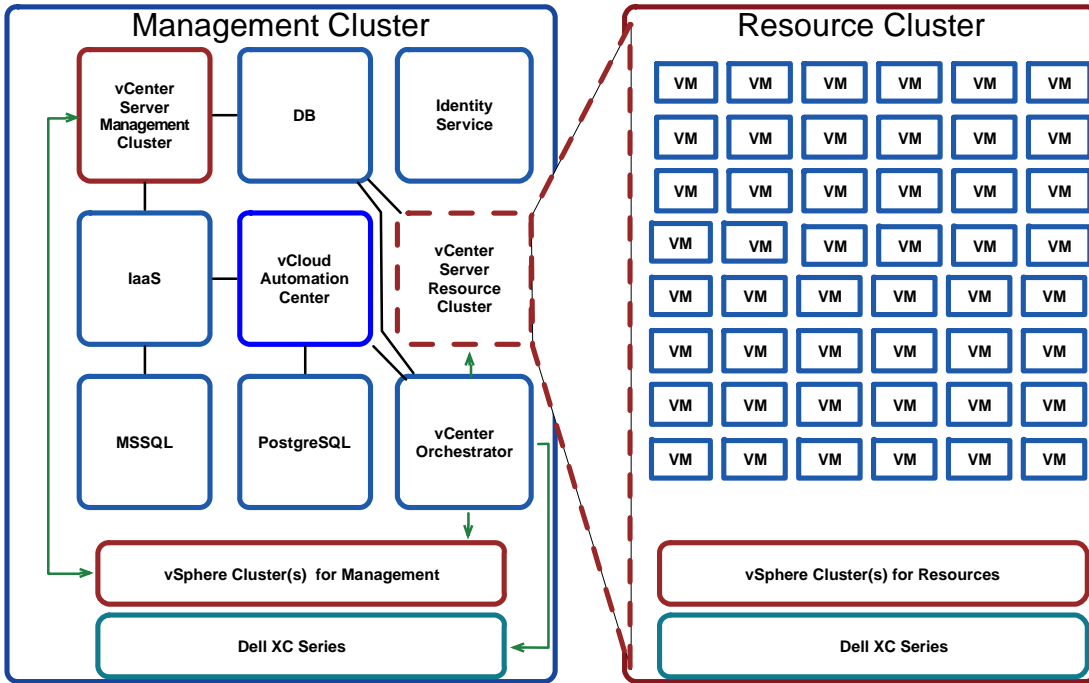


Figure 11 vRA separated into management and resource clusters

Important: A load balancer will be included in the medium and large deployment types in front of the vRA appliance and IaaS components.

All components will not be included in every deployment profile.

6.3

Note: Figure 11 includes VMware vSphere as the hypervisor and management solution for the resource workload. Additional hypervisors are supported by vRA as described in section 4.1: What is vRealize Automation.

vRealize Automation Center Component Placement

Dell recommends placing the vRA management components on the XC Series platform in a highly available configuration. Depending on the VMware license used for the vRA management cluster there are two different options: use vSphere Distributed Resource Scheduler (DRS) or manually place the virtual machines on the available XC appliances.

When vSphere DRS is available Dell recommends creating vSphere DRS rules separating the virtual machines delivering the same vRA component on different XC appliances. This applies to all components where applicable including:

- vRA"Appliance
- DEM"Orchestrator
- DEM"Worker
- IaaS"Manager
- IaaS"Web
- MSSQL
- PostgreSQL
- Proxy"Agent
- vRealize"Orchestrator
- vCenter"SSO
- Additional"DB

This will apply to the Medium and Large deployment profiles where vRA components are implemented in a highly available configuration meaning at least two virtual machines are available per component when applicable.

Depending on architectural and business requirements it might be possible to run one or more of the components on the same network. Virtual machines running on the same network (port group) will benefit in network communication performance from being placed on the same XC appliance.

DEM Orchestrator should have robust network connectivity to the IaaS manager. Proxy Agents will benefit from being deployed as close as possible to its associated endpoint. Consult reference [1] for potential placement recommendation changes.



7 Deployment profiles

This section includes the high level XC Series, vSphere and vRA components and options for successfully implement a vRA environment on XC Series appliances.

Additional information and recommendations can be found in the Nutanix vRA Best Practices guide.

Deployment Profile: Proof of Concept or Test Environment

The purpose of this deployment profile is to provide a way of implementing a vRA environment on XC Series appliances quickly for evaluation purposes.

7.1 This deployment profile will use the vRA management components deployed according to VMware's vRA Reference Architectures [1]. In addition, one vCenter Server virtual appliance or Windows based machine is required to manage the XC appliances (ESXi hosts). This deployment includes the following vRA components:

- vRA"Identity"Appliance
- vRA"Virtual"Appliance
 - Internal"PostgreSQL" database
 - vRO
- vRA"laaS"
 - Running"all"laaS components"required
- vCenter"Server"
- Microsoft"SQL
 - Used"by"vRA"laaS" components

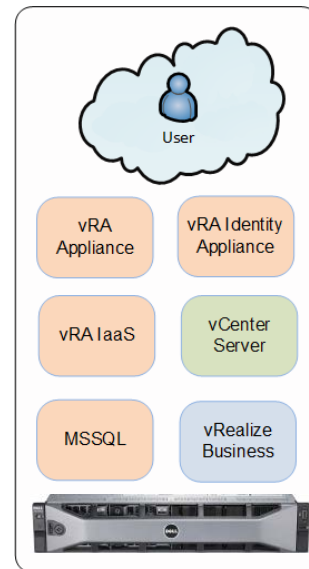


Figure 12 Logical overview of a POC / Test vRA deployment on XC Series appliances

No vRA management component redundancy is included and the XC container RF is set to its default value, 2.

Table 5 vRA Limits for Proof of Concept/ Test Deployments

Item	Maximum
vRA deployed virtual machines	1,000
vRA catalog items	2,500
vRA concurrent deployments	10



The following table outlines the Nutanix recommended setup for a POC / test environment deployment profile.

Table 6 Recommended Nutanix Setting for Proof of Concept/Test Deployments

Item	Quantity	Comment
XC clusters	1	Only one XC cluster deployed
XC Storage Pools	1	Use a one to one mapping between XC cluster and storage pools.
XC containers	2	1 container with in-line compression enabled for vRA management workload using RF2. 1 container with in-line compression enabled for vRA resource workload using RF2.
vCenter Servers	1	The vCenter Server manage both vRA management and resource workload.
vSphere Clusters	1	Two vSphere cluster resource pools should be created. 1 for vRA management components. 1 for vRA resource workload.
XC appliances / ESXi hosts	Minimum 3	Three XC appliances (ESXi hosts) are required to form a XC cluster.
XC Rack units	Minimum 6	
10 GbE Ports	Minimum 6	Two NICs per XC appliance.
100/1000 Ports (IPMI)	Minimum 3	One IPMI interface per XC appliance.
L2 Leaf Switches	2	

Important: Dell recommends an automatic process for keeping the resource prioritizing between the vSphere cluster resource pool for management components and the vSphere cluster resource pool for resource workloads at the desired level.



Deployment Profile: Small Production

The vRA management components are not implemented in a highly available configuration. The availability relies on the vSphere High availability (HA) feature included in the vSphere cluster where the management components run and the XC cluster out-of-the-box availability features.

7.2

This deployment profile will use the vRA management components deployed according to VMware's vRA Reference Architectures [1]. In addition, two vCenter Server virtual appliances or Windows based machine(s) vCenter Servers are required to manage the XC appliances.

This deployment includes the following vRA components:

- vRA Identity Appliance
- vRA Virtual Appliance
 - Internal PostgreSQL database
 - vRO¹
- vRA IaaS
 - Running all IaaS components required
- 2 x vCenter Server
- Microsoft SQL
 - Used by vRA IaaS components²

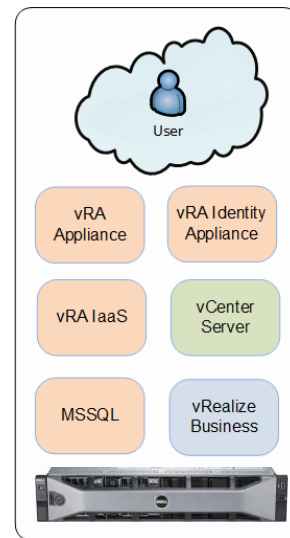


Figure 13 Logical overview of a small vRA production deployment on Nutanix

Notes:

1. The vRO appliance or the Windows based vRO can be used instead of the vRA built in vRO component based on vRO integration requirements. The Windows based vRO can be used when additional scripting integrations is being used such as PowerShell. This to make sure all scripting and workflows runs from the same source.
2. The Microsoft SQL database used for the vRO IaaS components can be used to host the vCenter Server database and the vRO database based on architectural, business and technical decisions.

vRA management components are implemented as single components for the small deployment type. If higher availability is required, based on business requirements, a XC container configured with RF3 can be used for the vRA management components.

The following VMware limitations apply to this deployment profile:

Table 7 vRA Limits for Small Deployments

Item	Maximum
vRA deployed virtual machines	1,000
vRA catalog items	2,500
vRA concurrent deployments	10

The following table outlines the Dell recommended setup for a small production deployment profile.

Table 8 vRA Recommended Setup for Small Deployments

Item	Quantity	Comment
XC cluster	1-2	Decision based on XC cluster to vSphere cluster configuration.
XC Storage Pools	1-2	Use a one to one mapping between XC cluster and Nutanix storage pools.
XC containers	Minimum 2	<ul style="list-style-type: none"> 1 container with in-line compression enabled for vRA management components using RF2. 1 container with in-line compression enabled for vRA resource workload using RF2.
vCenter Servers	Minimum 1	A two vCenter Server approach is recommended including: <ul style="list-style-type: none"> 1 vCenter Server for vRA workload 1 vCenter Server for vRA resource workload
vSphere Clusters	Minimum 2	
XC appliances / ESXi hosts	Minimum 4, 6	In a 1 vCenter Server and 1 XC cluster approach: <ul style="list-style-type: none"> 4 XC appliances. 2 for each vSphere cluster. In a 2 vCenter Server and 2 XC cluster approach: <ul style="list-style-type: none"> 6 XC appliances. 3 XC appliances for each vSphere cluster
XC Appliance Rack units	Minimum 8	
10 GbE Ports	Minimum 8	Two NIC per XC appliance.
100/1000 Ports (IPMI)	Minimum 4	One IPMI interface per XC appliance.
L2 Leaf Switches	Minimum 2	

Important: Even though a 1 vCenter Server, 1 XC cluster and 2 vSphere cluster implementation is possible based on software capabilities the approach must carefully consider operations management and a the possible lack of management and resource separation.

7.3

Deployment Profile: Medium Production

The vRA management components are implemented in a highly available configuration meaning at least two virtual machines are available per component. The XC container configuration uses RF2 since keeping additional copies, 3 instead of 2, per component will not substantially increase data availability but will use additional storage capacity. The only exception would be if either of the two databases used for the management components are implemented as a single point of failure.



This deployment profile will use the vRA management components deployed according to VMware's vRA reference architectures [1]. In addition, two vCenter Server(s) are required to manage the XC appliances. Additional load balancers are required.

This deployment includes the following vRA components:

- 2 x vCenter Single Sign On (SSO)
 - 2 x vRA Virtual Appliance
 - Internal PostgreSQL database².
- 2 x vRA IaaS Manager, Web and DEM Orchestrator.
- 2 x vRA DEM Worker & Proxy Agent
- 2 x vCenter Server
- 2 x vRO¹.
- 2 x Microsoft SQL used by vRA IaaS components³.
- 2 x PostgreSQL².
 - Used by vRA Appliances
- Load Balancer
 - Used by IaaS Mgr, IaaS Web, vRA Appliance and vRO

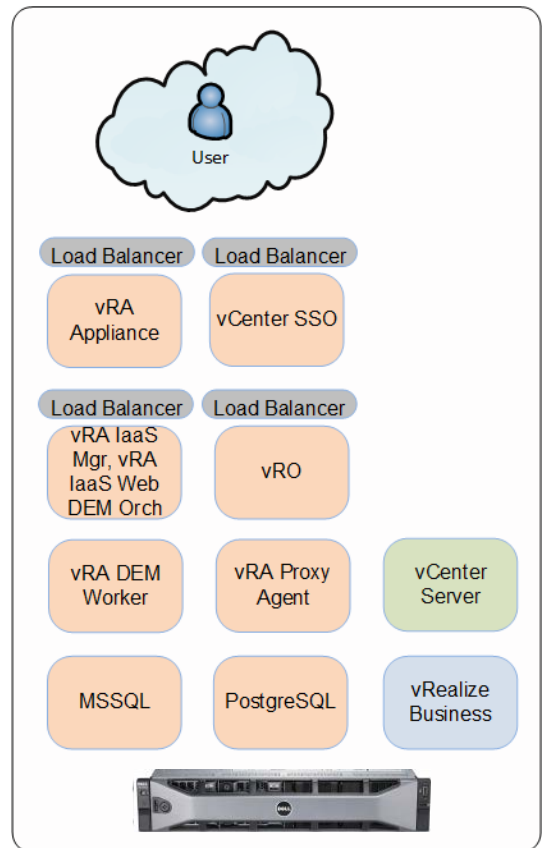


Figure 14 Logical overview of a medium vRA production deployment on Nutanix.

Notes:

1. Based on the vRO integration requirements the vRO appliance or the Windows based vRO can be used instead of the built-in vRA vRO component. A general recommendation includes an external vRO per vRA tenant. The Windows based vRO can be used when additional scripting integrations are being used such as PowerShell. This to make sure all scripting and workflows run from the same source.
2. The vRA virtual appliance internal PostgreSQL database can be implemented in a highly available configuration [6]. Architectural, business and technical decisions must determine if the vRA internal database or external databases must be used.
If using the internal vRA Appliance PostgreSQL database there are two options:
 - a. Run the databases on the same vRA virtual appliances where the vRA virtual appliance components run.
 - b. Run the databases on vRA virtual appliances where only the PostgreSQL databases run.
3. The Microsoft SQL database used for the vRO IaaS components can be used to host the vCenter Server database and the vRO database based on architectural, business and technical decisions. If another database engine is required there will be an additional two virtual machines required in the design.

Dell recommends using highly available deployed databases for medium production deployment profiles.

Table 9 vRA Limits for Medium Deployments

Item	Maximum
vRA deployed virtual machines	10,000
vRA catalog items	2,500
vRA concurrent deployments	50

The following table outlines the Nutanix recommended setup for a medium production deployment profile.

Table 10 vRA Recommended Setup for Medium Deployments

Item	Quantity	Comment
XC cluster	Minimum 2	For separation between management components and resource workload. Additional XC clusters will be required to support all potential, 10 000, vRA resource cluster deployed virtual machines.
XC Storage Pools	Minimum 2	Use a one to one mapping between XC cluster and Nutanix storage pools.
XC containers	Minimum 1	1 container with in-line compression enabled for vRA management workloads using RF2. 1 container with in-line compression enabled for vRA
vCenter Servers	Minimum 2	1 vCenter Server for vRA management workload. 1 vCenter Server for vRA resource workload.
vSphere Clusters	Minimum 2	1 per workload (vRA management and vRA resource) type.
XC appliances / ESXi hosts	Minimum 6	3 XC appliances per XC cluster.
Dell XC Rack units	Minimum 12	



Item	Quantity	Comment
10 GbE Ports	Minimum 12	Two NIC per XC appliance.
100/1000 Ports (IPMI)	Minimum 6	One IPMI interface per XC appliance.
L2 Leaf Switches	Minimum 2	



Deployment Profile: Large Production

The difference between the vRA management component layout for a medium and large deployment profile is the vRA IaaS component distribution. The vRA IaaS components Manager and Web are distributed on separate virtual machines compared to the medium production deployment profile where they coexist on the same virtual machines.

7.4

This deployment profile will use the vRA management components deployed according to VMware's vRA reference architectures [1]. In addition, two vCenter Server(s) are required to manage the XC appliances. Additional load balancers are required.

This deployment includes the following vRA components:

- 2 x vCenter Single Sign On (SSO)
- 2 x vRA Virtual Appliance
 - Internal PostgreSQL database ².
- 2 x vRA IaaS Manager and DEM Orchestrator.
- 2 x vRA IaaS Web.
- 2 x vRA DEM Worker + Proxy Agent
- 6 x vCenter Server
- 2 x vRO ¹
- 2 x Microsoft SQL used by vRA IaaS components ³.
- 2 x PostgreSQL ².
 - Used by vRA Appliances
- Load Balancer
 - Used by IaaS Mgr, IaaS Web, vRA Appliance and vRO

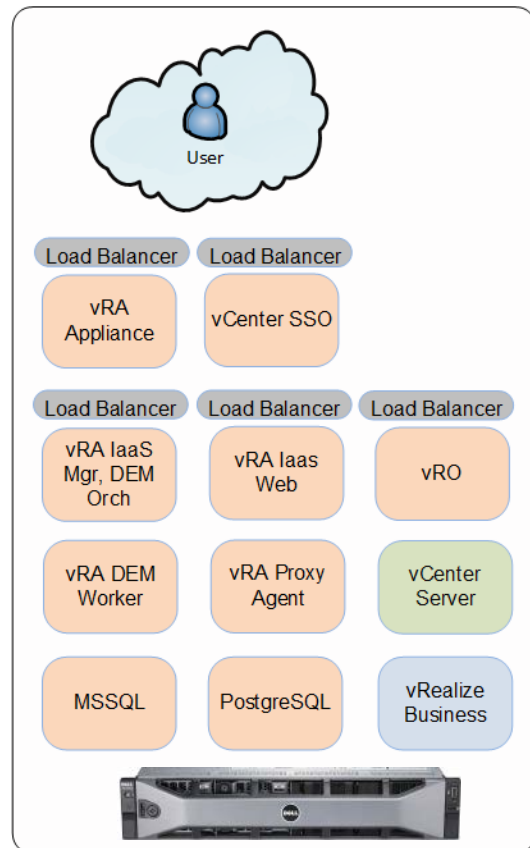


Figure 15 Logical overview of a large vRA production deployment on Nutanix.

Notes:

1. Based on the vRO integration requirements the vRO appliance or the Windows based vRO can be used instead of the built-in vRA vRO component. A general recommendation includes an external vRO per vRA tenant. The Windows based vRO can be used when additional scripting integrations are being used such as PowerShell. This to make sure all scripting and workflows run from the same source.
2. The vRA virtual appliance internal PostgreSQL database can be implemented in a highly available configuration [6]. Architectural, business and technical decisions must determine if the vRA internal database or external databases must be used.
If using the internal vRA Appliance PostgreSQL database there are two options:
 - a. Run the databases on the same vRA virtual appliances where the vRA virtual appliance components run.
 - b. Run the databases on vRA virtual appliances where only the PostgreSQL databases run.
3. The Microsoft SQL database used for the vRA IaaS components can be used to host the vCenter Server database and the vRO database based on architectural, business and technical decisions. If another database engine is required there will be an additional two virtual machines required in the design.

Important: vRealize Business supports up to 20,000 virtual machines deployed over 4 vCenter Servers.

Nutanix recommends using highly available deployed databases for large production deployment profiles.

Table 11 vRA Limits for Large Deployments

Item	Maximum
vRA deployed virtual machines	50,000
vRA Center catalog items	2,500
vRA concurrent deployments	100

The following table outlines the recommended setup for a large production deployment profile:

Table 12 vRA Recommended Setup for Large Deployments

Item	Quantity	Comment
XC cluster	Minimum 2	For separation between vRA management components and vRA resource workload. Additional XC clusters will be required to support all potential, 50 000, vRA resource cluster deployed virtual machines.
XC Storage Pools	Minimum 2	Use a one to one mapping between XC cluster and Nutanix storage pools.
XC containers	Minimum 4	1 container with in-line compression enabled for vRA management workloads using RF2. 1 container with in-line compression enabled for vRA resource workloads using RF2.
vCenter Servers	Minimum 2	1 vCenter Server for vRA management workload. Minimum 1 vCenter Server for vRA resource workload. Additional vCenter Servers, 4, are required to manage all potential, 50 000, vRA resource cluster deployed virtual machines.
vSphere Clusters	Minimum 2	1 per workload (vRA management and vRA resource) type.
XC appliances / ESXi hosts	Minimum 6	3 XC appliances per XC cluster.



Dell XC Rack units	Minimum 12	
10 GbE Ports	Minimum 12	Two NIC per XC appliance.
100/1000 Ports (IPMI)	Minimum 6	One IPMI interface per XC appliance.
L2 Leaf Switches	Minimum 2	



8 Scalability

This section will provide guidance regarding how to scale the Dell – Nutanix solution for the vRA resource workload.

The Nutanix web-scale converged infrastructure running the vRA management components can be scaled from one deployment profile to another within the existing environment if the XC appliances provides enough hardware. If not, the XC appliances can be scaled up with more RAM capacity to manage the larger vRA management component deployment profile requirements.

Scalability Example

The number of vRA deployed virtual machines per XC appliance will vary depending on workload type. Using a _____XC appliance with the below specification will enable approximately 40 vRA deployed virtual machines based on the following:

8.1

- node configuration:
- TBD
- Virtual machine configuration
 - 2 vCPU
 - 50% vCPU utilization
 - 6 GB RAM
 - 80 GB Disk

Consult reference [5] for future XC appliance hardware specification updates.

8.2

Scalability Input Parameters

Scalability input parameters can be divided into two sections, software-based and company-based.

8.2.1 Software-Based

The software-based scalability limitations and rules are listed in [1] and [4]. Some key software limitations to respect for scalability includes:

- Maximum number of vRA deployed virtual machines
- Maximum number of vRA catalog items
- Maximum number of vRA concurrent deployments
- vSphere configuration maximums [4] in terms of:
 - Maximum of 4,000 virtual machines per vSphere cluster
 - Maximum of 2,048 powered on virtual machine configuration files per vSphere HA datastore.
Note: One vSphere FT protected virtual machine counts for two.
- One vCenter Server can manage 10,000 powered on virtual machines. This applies to the Windows based vCenter Server and the vCenter Server Appliance when using an external Oracle database.

- One vCenter Server can manage 15,000 registered virtual machines. This applies to the Windows based vCenter Server and the vCenter Server Appliance when using an external Oracle database
- One vCenter Server can manage a maximum of 1,000 ESXi hosts. This applies to the Windows based vCenter Server and the vCenter Server Appliance when using an external Oracle database
- Maximum number of ESXi host per vCenter Server vSphere HA cluster is 32

8.2.2 Corporate Data center

The following bullets include examples of company-based scalability input parameters:

- Availability domains – A physical site, a computer room, and a computer rack can all be treated as availability domains. Dell – Nutanix cannot scale out between sites, but between computer rooms within the same physical location and between racks.
- Maximum number of XC appliances / ESXi hosts per XC cluster.
- Maximum number of XC appliances / ESXi hosts per vSphere cluster.
- Maximum number of virtual machines managed by same vCenter server.
- Maximum number of virtual machines included in same vSphere cluster.



9 vRealize Automation Deployment Tests

This section will describe what to expect in terms of provisioning times and potential disk savings when using the different vRA deployment types and the different XC container features plus VAAI.

There are several ways of creating a vSphere vCenter Server-based virtual machine via vRA. Three options are relevant for a vRA deployment on Nutanix:

- Create – "A new virtual machine is created without an operating system. vRA can either start a virtual machine operating system installation or request a virtual machine operating system installation from a third-party system.
- Clone – "An existing vCenter Server virtual machine template is used to create new virtual machines.
- Linked Clone – "A snapshot of an existing vCenter Server managed virtual machine is used to create new virtual machines.

The clone and linked clone vRA deployment types have been tested and evaluated.

Important: The test results in this section should be used as guidelines when determining what vRA deployment types and which XC container features to use. Architectural and business requirements must be respected. The tests include only initial virtual machine deployments. No additional data is added to the virtual machines.

9.1 General Test Setup

This section outlines the configuration (hardware and software) used for the clone and linked clone tests.

9.1.1 Hardware

- Storage / Compute

9.1.2 Software

- Nutanix - "Version" "NOS" 4.0.1

9.1.3 Virtual machine

- Nutanix "Controller Virtual Machine" (CVM)
 - CPU: "8" vCPU
 - Memory: "32" GB
- vRA "deployed" virtual machine
 - Windows "Server" 2012 "R2"
 - 2 vCPU
 - 4 GB "RAM" (No reservation)
 - 40 GB "provisioned" disk
 - 9 GB VM disk size



Clone Test

The vRA clone virtual machine provisioning tests were carried out according to the following:

9.2

- 50 virtual machines deployed via vRA to the following XC containers:
 - Default XC container – Deploy VMs across XC containers without using VAAI
 - Default XC container – Deploy VMs within a single XC container using VAAI
 - XC container with inline compression enabled – Deploy VMs across XC containers
 - XC container with inline compression enabled – Deploy VMs within the same XC container
 - XC container with post compression (120 minutes) enabled – Deploy VMs across XC containers
 - XC container with post compression (120 minutes) enabled – Deploy VMs within the same XC container

9.2.1 Results

The clone deployment type using VAAI is approximately 10 times faster compared to the clone deployment type not using VAAI. The clone deployment type using VAAI uses approximately 8.5% of the disk space compared to the deployment type not using VAAI.

9.3

Linked Clones

The main reasons for using the linked clone feature compared to the clone feature includes:

- Faster provisioning of virtual machines
- Disk space savings

9.3.1 Linked Clone Test

The vRA linked clone virtual machine provisioning tests were carried out according to the following:

- 50 virtual machines deployed via vRA to the following XC containers
- Default XC container – VMs deployed across XC containers
- Default XC container – VMs deployed within the same XC container
- XC container with inline compression enabled - VMs deployed across XC containers
 - XC container with inline compression enabled - VMs deployed within the same XC container
 - XC container with post compression (120 minutes) enabled - VMs deployed across XC containers
 - XC container with post compression (120 minutes) enabled - VMs deployed within the same XC container

9.3.2 Linked Clone Test Results

9.3.2.1 Faster Provisioning

The scenarios for using the linked clone feature can be divided into two categories:



- Linked clone virtual machine deployment across and within a XC container compared to cloning virtual machines across XC containers (not using VAAI).
 - Tests showed up to 10 times faster virtual machine deployment using the linked clone feature.
- Linked clone virtual machine deployment across and within XC containers compared to cloning virtual machines within a XC container (using VAAI).
 - Tests showed equal deployment times for both features.

9.3.2.2 Disk Space Savings

The disk space savings for using the linked clone feature can be divided into three categories:

- Linked clone virtual machine deployment across and within XC containers compared to cloning virtual machines across XC containers (not using VAAI).
 - Tests showed that the linked clone feature used about 10% of the disk space compared to the clone feature.
- Linked clone virtual machine deployment across and within XC containers compared to cloning virtual machines within a XC container (using VAAI).
 - Tests showed that the linked clone feature used approximately 95% of the disk space compared to the clone feature.
- Linked clone virtual machines deployment across and within XC containers using the Nutanix compression (inline or post) compared to linked clone virtual machine deployment across and within XC containers.
 - Test shows additional disk space savings of 5% - 25%.

9.4 Test Results Summary

9.4.1 Provisioning Time

Both clone (using VAAI) and linked clone vRA deployment types deployed virtual machines equally in terms of provisioning time.

Note: The provisioning time was measured from a vRA perspective and not from either vCenter Server or the Nutanix platform.

The following table outlines the time difference between the different vRA deployment types and VAAI.

Table 13 Provisioning time test results

vRA deployment type	VAAI used	Provisioning time
Linked Clone	No	Baseline
Clone	Yes	Baseline
Clone	No	10x Longer



10 Deployment Profile Hardware Requirements

vRA Components

The vRA management components required hardware specifications are outlined in this section. The figures within parentheses, if any, are the recommended hardware specifications.

The figures are presented for vRA management components hardware calculating purposes. Refer to the [1] for potential hardware requirement changes.

10.1

- vRA"Identity"Appliance
 - 1"vCPU,"2"GB"RAM,"10"GB"disk,"1"GB/s"network
- vCenter"Single"Sign"On
 - 2"vCPU,"3"GB"RAM,"1GB
- vRA"Appliance"
 - 2"(4)"vCPU,"8"(16)"GB"RAM,"30"GB"Disk,"1"GB/s"network
- vRA"Appliance"- "only"running"PostgreSQL
 - 2"vCPU,"2"GB"RAM,"20"GB"Disk,"1"GB/s"network
- vRA"Infrastructure"Web"Server
 - 2"vCPU,"2"(4)"GB"RAM,"40"GB"Disk,"1"GB/s"network
- vRA"Infrastructure"Web"Manager"Server
 - 2"vCPU,"4"(8)"GB"RAM,"40"GB"Disk,"1"GB/s"network
- vRA"Distributed"Execution"Manager"(DEM)"Orchestrator
 - 2"vCPU,"2"(4)"GB"RAM,"40"GB"Disk,"1"GB/s"network
- vRA"Distributed"Execution"Manager"(DEM)"worker"
 - 2"vCPU,"2"(6)"GB"RAM,"60"GB"Disk,"1"GB/s"network
- vRA"Infrastructure"Agent"Server"(Proxy"Agent)"
 - 2"vCPU,"4"GB"RAM,"40"GB"Disk,"1"GB/s"network
- Microsoft"SQL"
 - 2"(8)"vCPU,"8"(16)"GB"RAM,"40"(80)"GB"Disk,"1"GB/s"network



Additional Components

10.2

- vCenter"Server"Appliance
 - 2-4"vCPU,"4-32"GB"RAM,"75-125"GB"Disk,"1"GB/s"network.
See"[2]"for"deployment"profile"hardware"requirements.
- Windows"based"vCenter"Server
 - 2"(4)"vCPU,"12"(16)"GB"RAM,"140"GB"Disk,"1"GB/s"network.
See"[3]"for"hardware"requirements"for"a"distributed"vCenter"Server"installation.
 - vRAS"Director:
 - Small"deployment"profile:
 - 2"vCPU,"4"GB"RAM,"16"GB"Disk,"1"GB/s"network.
 - Medium"deployment"profile:
 - 4"vCPU,"8"GB"RAM,"30"GB"Disk,"1"GB/s"network.
 - Large"deployment"profile:
 - 8"vCPU,"16"GB"RAM,"50"GB"Disk,"1"GB/s"network.
- vRealize"Business"Virtual"Appliance
 - 2"vCPU,"4"GB"RAM,"50"GB"Disk,"1"GB/s"network.
- vRO"Appliance
 - 2"vCPU,"3"GB"RAM,"12"GB"Disk,"1"GB/s"network.
- vRO"Windows"based
 - 2"vCPU,"4"GB"RAM,"10"GB"Disk,"1"GB/s"network.
- Administration"server/s" – "Windows"2012"or"2008"
 - 2"vCPU,"4"GB"RAM,"100"GB"Disk
- Domain"Name"Servers"(DNS)
 - 1"vCPU,"2"GB"RAM,"30"GB"Disk
- Lightweight"Directory"Access"Protocol"(LDAP)"server/s"
 - 2"vCPU,"4"GB"RAM,"40"GB"Disk
- Network"Time"Protocol"(NTP)
 - 1"vCPU,"2"GB"RAM,"20"GB"Disk
- Simple"Mail"Transport"Protocol"servers/s
 - X"vCPU,"X"GB"RAM,"X"GB"Disk
- Syslog
 - 2"vCPU,"4"GB"RAM,"100"GB"Disk



11 Conclusion

The Nutanix web-scale converged infrastructure offers the ability to run vRA management workloads for the four reference architecture-defined deployment types:

- Proof-of-concept/test environments
- Small deployments
- Medium deployments
- Large deployments

vRA deployment features clone and linked clone have been tested and detailed information is provided in terms of virtual machine provisioning times and disk space savings between the two features. The test results can be used as design architectural input parameters.

Configuration guidance is provided for the following Nutanix web-scale converged infrastructure and VMware vSphere components:

- Nutanix System
- Nutanix Storage Pool
- XC container
- XC appliance
- vCenter Server
- vSphere Cluster

The XC container offers features for disk saving. Detailed initial setup recommendations are included to give guidance if/when the features can/should be used.

In addition to the vRA management components, the vRA resource workload is discussed in sections of this reference architecture. VMware vSphere is outlined as the hypervisor platform for the resource workload, but both vRA and the Nutanix web-scale converged infrastructure support multiple hypervisors.

This makes vRA and the Dell – Nutanix web-scale converged infrastructure the perfect combination and building block for an enterprise, service provider, private, or public cloud infrastructure.

For any additional questions on Nutanix or vRA, please continue the conversation on our Nutanix Next online community: next.nutanix.com.



A Appendix

A.1 References

The following documents are used as references in the Nutanix vRA reference architecture.

1. VMware vCloud Automation Center 6.1 Reference Architecture:
<http://www.vmware.com/files/pdf/products/vCloud/VMware-vCloud-Automation-Center-61-Reference-Architecture.pdf>
2. vCenter Server Appliance hardware requirements:
http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2005086
3. vCenter Server 5.5 Hardware Requirements: <http://pubs.vmware.com/vsphere-55/index.jsp?topic=%2Fcom.vmware.vsphere.install.doc%2FGUID-67C4D2A0-10F7-4158-A249-D1B7D7B3BC99.html>
4. vSphere 5.5 Configuration Maximums: <http://www.vmware.se/pdf/vsphere5/r55/vsphere-55-configuration-maximums.pdf>
5. Nutanix Web-scale Converged Infrastructure Specifications:
http://go.nutanix.com/rs/nutanix/images/Nutanix_Spec_Sheet.pdf
6. Setting up PostgreSQL replication in the VMware vCloud Automation Center 6.0 virtual appliance (2083563): <http://kb.vmware.com/kb/2083563>
7. Microsoft SQL Server Best Practices for Nutanix: http://go.nutanix.com/TechGuide-Nutanix-SQLBestPractices_Asset.html
8. Using VMware® vCenter SSO 5.5 with VMware vCloud Automation Center 6.1:
<http://www.vmware.com/files/pdf/products/vCloud/VMW-vRealize-Automation-61-Deployment-Guide-HA-Configurations.pdf>

A.2 List of figures

Figure 1	Node Architecture	7
Figure 2	Dell – Nutanix Solution Architecture.....	7
Figure 3	Elastic Deduplication Engine.....	8
Figure 4	NDFS Shadow Clones.....	9
Figure 5	Logical overview of NDFS, VMware vSphere, vRA and VM / Virtual Disk.....	11



Figure 6	vRA Conceptual Architecture	12
Figure 7	Nutanix Write I/O	13
Figure 8	Dell XC Series Conceptual Scalability	14
Figure 9	Dell XC Series block	16
Figure 10	Logical vCenter Server to vRA mapping.....	17
Figure 11	vRA separated into management and resource clusters	23
Figure 12	Logical overview of a POC / Test vRA deployment on XC Series appliances.....	25
Figure 13	Logical overview of a small vRA production deployment on Nutanix	27
Figure 14	Logical overview of a medium vRA production deployment on Nutanix.....	29
Figure 15	Logical overview of a large vRA production deployment on Nutanix.....	32

A.3 List of tables

Table 1	Nutanix Design Elements	16
Table 2	vRA elements.....	16
Table 3	vRA Infrastructure Items and Descriptions	18
Table 4	Container Recommendations.....	21
Table 5	vRA Limits for Proof of Concept/ Test Deployments	25
Table 6	Recommended Nutanix Setting for Proof of Concept/Test Deployments.....	26
Table 7	vRA Limits for Small Deployments	28
Table 8	vRA Recommended Setup for Small Deployments.....	28
Table 9	vRA Limits for Medium Deployments	30
Table 10	vRA Recommended Setup for Medium Deployments	30
Table 11	vRA Limits for Large Deployments	33
Table 12	vRA Recommended Setup for Large Deployments	33
Table 13	Provisioning time test results.....	39



A.4 Further Research

As part of its continuous determination to deliver the best possible solutions, Nutanix plans to continue to conduct research into the following areas:

- Scalability and Sizing of vRA environments
- Hyper-V and KVM hypervisors in vRA environments
- The inclusion of Nutanix Metro Availability Clusters with vRA

