

# Dell EMC Ready Stack

## Design Considerations for Certified Reference Systems Built from Dell EMC Components

April 2019

H17501.1

### Reference Architecture Guide

#### Abstract

This reference architecture guide describes the architecture of Dell EMC Ready Stack and the options for building Ready Stack systems using Dell EMC hardware and software technologies.

Dell EMC Solutions

## Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2018–2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners.

Published in the USA 04/19 Reference Architecture Guide H17501.1.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.



# Contents

- Executive summary ..... 4**
- Architecture overview..... 8**
- Servers ..... 10**
- Storage..... 12**
- Networking..... 16**
- Management infrastructure..... 20**
- Virtualization..... 23**
- Data protection ..... 24**
- Dell Boomi: Application integration ..... 25**
- Big Cloud Fabric: Software-defined networking with Dell EMC Open Networking switches ..... 28**
- Summary ..... 35**
- References ..... 36**

## Executive summary

### Business challenge

As business processes become increasingly automated and software takes over previously manual operations, the role of applications and software in business becomes more widespread and far-reaching. As application ecosystems increase in complexity and become more intertwined with business operations, IT's relevance to the business increases as well. This represents a huge shift for many IT organizations, particularly for those that are accustomed to a back-office support function.

IT organizations need to re-examine every facet of what they do, how they do it, and how they prioritize their business objectives. In particular, IT organizations need to transform to deliver greater efficiency, predictability, and business agility.

Most IT organizations agree. According to 2017 research conducted by Enterprise Strategy Group (ESG) for Dell EMC, 71 percent agree that if they do not embrace IT transformation, they will no longer be competitive in the market. For details, see [How IT Transformation Maturity Drives IT Agility, Innovation, and Improved Business Outcomes](#).

A reference system such as Dell EMC Ready Stack, based on converged infrastructure (CI) methodologies, brings together the disparate infrastructure elements that power IT: servers, data storage devices, networking functions, virtualization, and management software. The value that CI brings to IT consumers is twofold—validated integration testing, which reduces risk and speeds deployment time, and leveraged management and automation tooling to provide efficiency and insight into the entire platform. Whether you are virtualizing workloads for private, public, or hybrid cloud deployments, CI systems expedite your journey to the cloud.

Underpinning the modern data center are multiple technology pillars—flash, scale-out architectures, software-defined Infrastructure, and cloud-enabled systems. CI helps make the adoption of these new technologies faster, simpler, more agile, more efficient, less risky, and less costly, which speeds overall IT transformation efforts.

Dell EMC Ready Stack is a certified reference system that simplifies and accelerates the process of deploying full-stack systems based on Dell EMC technologies.<sup>1</sup> By incorporating the core tenets of a CI system, Ready Stack helps organizations modernize faster, with less risk and less effort than through alternative solutions—whether the environment already includes Dell EMC or uses reference architectures from other vendors.

Not all customer objectives are the same, so Dell EMC offers a continuum of CI options. Customer preferences range from do-it-yourself (DIY) approaches to fully engineered, factory-built, enterprise-class architectures for global organizations with global data centers.

---

<sup>1</sup> Industry analyst firm IDC defines a certified reference system as a pre-integrated, vendor-certified system containing server, storage, networking, and basic component management software. Dell EMC Ready Stack combines all these elements and more, and includes validation by Dell EMC Engineering.

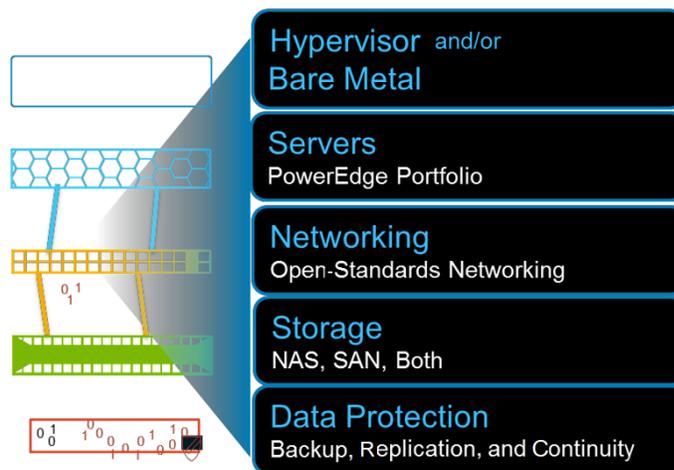
## Ready Stack overview

Ready Stack is a program that is directed primarily toward channel partners and their customers. The program enables them to easily build complete CI systems from combinations of Dell EMC servers, storage, networking, data protection, and other components. For networking, the use of Dell EMC open-standards-based networking technologies is preferred but not required.

Ready Stacks are compatible with hypervisors from VMware and Microsoft and with Microsoft operating systems. Ready Stacks can be managed by the Dell EMC portfolio of management products and protected by Dell EMC data protection solutions.

The following figure illustrates the Ready Stack reference system.

### BUILD WITH SIMPLICITY, SPEED AND CONFIDENCE



**Figure 1. Ready Stack reference system**

A Ready Stack reference architecture specifies all the hardware, software, resources, and services that are needed to run a scalable, highly available CI environment. This end-to-end solution approach means that a system can be operational in a shorter time than is typically possible with build-your-own solutions.

Designing, configuring, and running a production-ready CI involves multiple considerations, including:

- Selection of operating system and virtualization software distributions
- Selection of monitoring and management software
- Allocation of cluster services and data storage to physical nodes and arrays
- Selection of server hardware
- Design of the network fabric
- Sizing and scalability
- Performance

## Reference systems

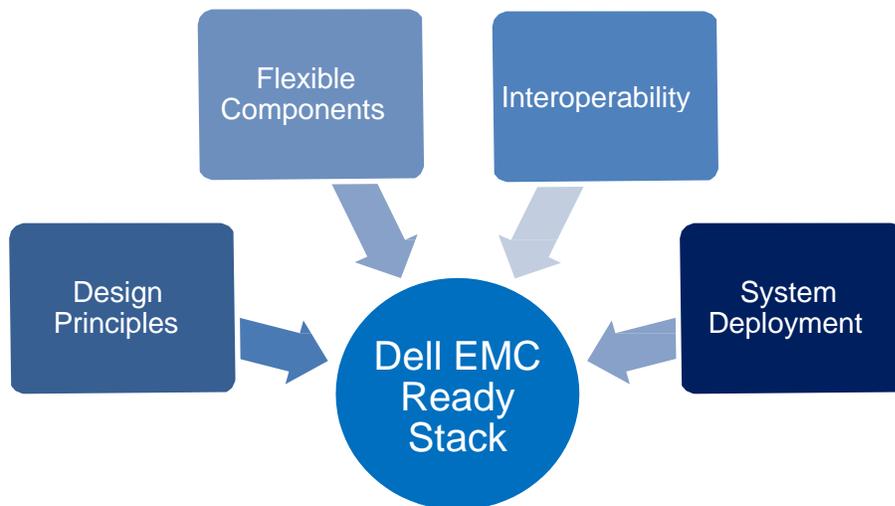
In addition to this guide, Dell EMC provides Ready Stack architecture designs and further guidance for specific hardware choices. For each Ready Stack architecture, the following documents provide those details:

- **Design guide**—Design, configuration, and sizing of the Ready Stack architecture
- **Deployment guide**—General guidance on deploying virtualization software on the complete Ready Stack platform, including systems management tools and plug-ins

The reference architectures that are described in the Ready Stack design and deployment guides are suggestions for specific Ready Stack designs. You can use other combinations of servers, storage, networking, and virtualization software when designing your own configurations provided they meet or exceed minimum hardware specifications. Ready Stack reference architectures are:

- Built on Dell EMC products that are designed for virtualization across the entire ecosystem
- Designed to help remove risk and accelerate time to value, yet flexible enough to be tailored for any organization
- Accompanied by software that makes Ready Stack easier to manage
- Delivered by Dell EMC trusted partners, who understand our joint customers' business needs and how to deliver comprehensive, low-risk solutions

A Ready Stack reference architecture removes a significant amount of complexity and risk from the DIY approach by providing guidance on many aspects of a modern data center, as illustrated in the following figure.



**Figure 2. Ready Stack elements**

The reference architecture includes:

- Component interoperability and connectivity requirements
- System design redundancy
- Operational management considerations
- Scale-up and scale-out methods

## Benefits of an all-Dell EMC solution

One of Dell EMC's strengths is its broad and complete product portfolio that can be tailored to meet customer requirements regardless of the customer's size, scale, or business model.

The breadth of the Dell EMC family of products affords the freedom to focus on rightsizing elements of the complete system, including compute, network, and storage.

Ready Stack systems are ideal for DIY customers who:

- Are considering a storage or compute refresh from any vendor
- Want to simplify their IT operations, reduce costs, and minimize support risks by using fewer vendors
- Have concerns about building a full infrastructure solution themselves and maintaining best-practice configurations
- Need to move quickly to deploy new applications or pursue new business initiatives
- Require a complete range of data protection options, including backup, replication, and business continuity, to protect data with various levels of business value/criticality
- Want to avoid "multivendor finger-pointing" or business-impacting delays in their support experience
- Need a repeatable platform for general-purpose virtualization

Dell EMC Services provides partners with choice and flexibility in how they provide Ready Stack implementation services. Gold tier partners and above can either resell Dell EMC services or deliver their own, provided they have earned the appropriate Services Delivery Competencies.

### Document purpose and scope

This reference architecture guide describes the component options for building the systems using Dell EMC hardware and software technologies.

The correct deployment methodology depends on identified business outcomes and a number of other variables that any consultative channel partner will identify and work through when developing solutions. This guide focuses on a single-vendor model using a Dell EMC Ready Stack architecture.

### Audience

This guide is for Dell EMC personnel, channel partners, and customers.

### We value your feedback

Dell EMC and the authors of this document welcome your feedback. Contact the Dell EMC Solutions team by [email](#) or provide your comments by completing our [documentation survey](#).

**Authors:** Robert Percy, Scott Stanford, Don Pisinski, David Hartman, Karen Johnson

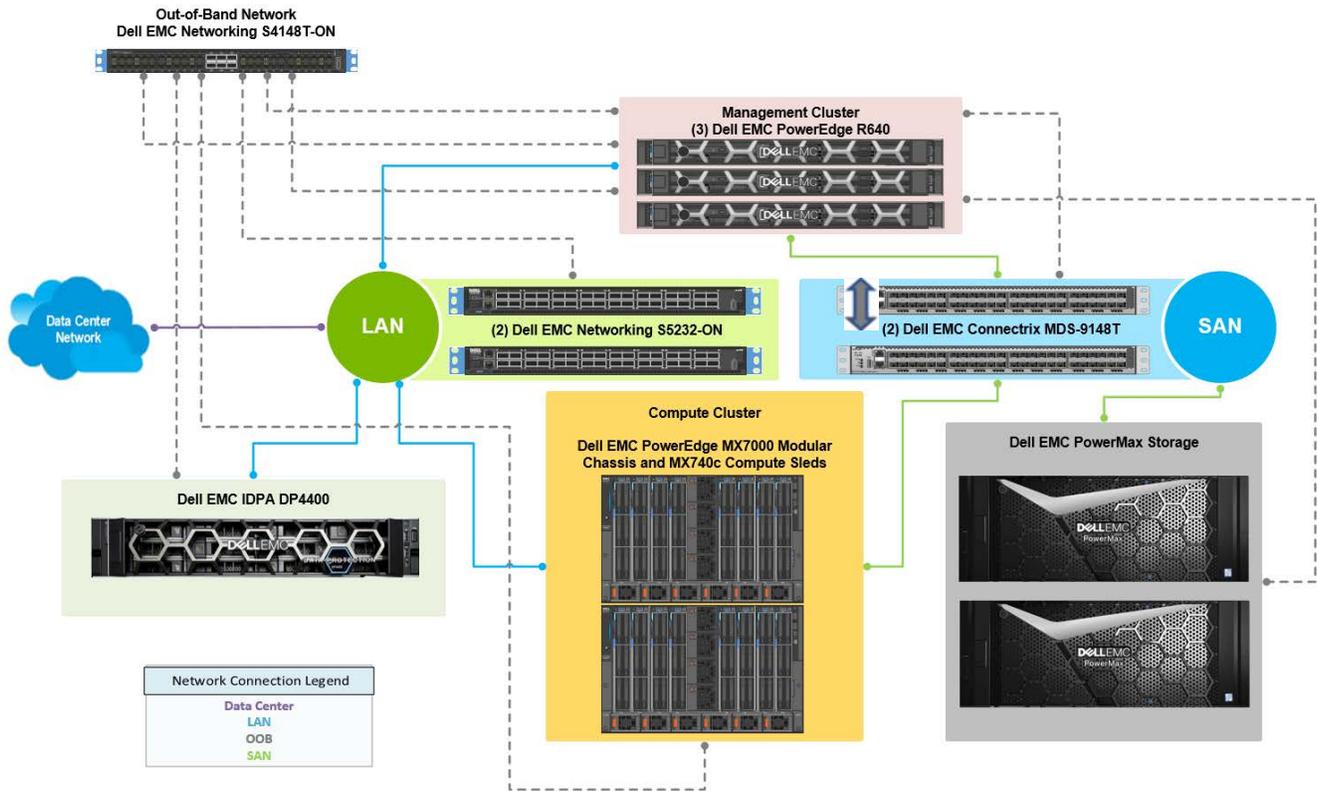
---

**Note:** The following page of the Ready Stack space on the Dell EMC Communities website provides links to additional documentation for this and other Ready Stack systems: [Ready Stack Info Hub](#)

---

# Architecture overview

The following figure provides a high-level overview of a sample Ready Stack architecture.



**Figure 3. Ready Stack architecture**

The following table outlines the Ready Stack component options.

**Table 1. Ready Stack component options**

| Component category                   | Options  |
|--------------------------------------|--|
| Server platforms                     | <ul style="list-style-type: none"> <li>Dell EMC PowerEdge rack servers</li> <li>Dell EMC PowerEdge modular infrastructure</li> </ul>                                       |
| LAN connectivity                     | Dell EMC Networking 25/40/50/100 GbE switches  |
| SAN connectivity                     | Dell EMC Connectrix 16 Gb and 32 Gb FC switches  |
| OOB connectivity (management switch) | Dell EMC Networking 10/100/1000 Mb switches  |
| Storage arrays                       | <ul style="list-style-type: none"> <li>Dell EMC Unity</li> <li>Dell EMC PowerMax</li> <li>Dell EMC SC Series</li> <li>Dell EMC XtremIO</li> <li>Dell EMC Isilon</li> </ul> |

| Component category                      | Options  |
|---|--|
| Management and compute server platforms | <ul style="list-style-type: none"> <li>• OOB management: Dell EMC PowerEdge R640 rack servers</li> <li>• Compute: Dell EMC MX740c compute sleds</li> </ul>   |
| Management software components          | <ul style="list-style-type: none"> <li>• For VMware vSphere: <ul style="list-style-type: none"> <li>▪ VMware vCenter Server Appliance (VCSA)</li> <li>▪ Dell EMC OpenManage Integration for VMware vCenter (OMIVV)</li> <li>▪ Dell EMC Virtual Storage Integrator (VSI)</li> </ul> </li> <li>• For Microsoft Hyper-V: <ul style="list-style-type: none"> <li>▪ Dell EMC OpenManage Integration for Microsoft System Center (OMIMSSC) for Microsoft System Center Configuration Manager (SCCM) and System Center Virtual Machine Manager (SCVMM)</li> <li>▪ Dell EMC Storage Integrator (ESI) for Windows Suite</li> <li>▪ Microsoft SCCM</li> <li>▪ Microsoft SCVMM</li> </ul> </li> </ul> |
| Data protection                         | <ul style="list-style-type: none"> <li>• Dell EMC Integrated Data Protection Appliance (IDPA)</li> <li>• Dell EMC Avamar</li> <li>• Dell EMC Data Domain</li> <li>• Dell EMC RecoverPoint</li> <li>• Dell EMC VPLEX</li> </ul>   |

## Design principles

Central to the design and architecture of a Dell EMC Ready Stack are the following principles:

- **No single point of failure**—The critical aspects of a Ready Stack system incorporate redundancy, including server high availability features, redundant networking, and multipath storage.
- **Integrated Management**—When the VMware hypervisor is deployed, Ready Stack includes integrated management by using Dell EMC OpenManage Integration for VMware vCenter (OMIVV) and Dell EMC Virtual Storage Integrator (VSI). Optionally, ESI for Windows Suite can be used to integrate with Microsoft System Center Operations Manager (SCOM) and SCVMM in Hyper-V environments.
- **Hardware configuration for virtualization**—Ready Stack is designed for general-purpose virtualized workloads. Each server is equipped with the processing power, memory, and converged network adapters that support Ethernet and Fibre Channel (FC), as required for virtualization.
- **Best practices adherence**—The Ready Stack design incorporates storage, networking, and hypervisor best practices to ensure availability and serviceability.
- **Flexible configurations**—Ready Stack can be configured to suit the requirements for a specific virtualized infrastructure. The system supports flexibility in the form of various options, such as rack and blade server models, number of compute servers, server processor models, server memory capacity, storage array types, and storage capacity based on customer needs.

## Servers

The Dell EMC server portfolio includes PowerEdge rack servers and the PowerEdge MX7000 chassis and blade servers.

PowerEdge rack and blade servers are designed to optimize application performance and ensure a stable, worry-free environment with intuitive tools that simplify and automate processes throughout the server lifecycle. Every customer has unique requirements, and PowerEdge servers provide the flexibility that is needed to build a scalable infrastructure that is tightly integrated with either VMware vSphere or Microsoft Hyper-V virtualization.

Embedded in every PowerEdge server, integrated Dell Remote Access Controller 9 (iDRAC 9) provides secure and remote server access for a multitude of common management functions. iDRAC with Lifecycle Controller operates regardless of the state of the operating system or the presence of a hypervisor and offers a complete set of server management features, including configuration, operating system deployment, firmware updates, health monitoring, and maintenance.

### PowerEdge server options

Ready Stack reference architectures offer a selection of server platforms for the compute infrastructure. A single Ready Stack configuration can contain any combination of those servers with resource configurations suitable for your workload and use case.

#### Rack servers versus blade enclosures

Ready Stack offers a choice between rack servers and blade enclosures. Both server types mount into standard data center racks, use floor space efficiently, and can co-exist with other components within the same rack. Dell EMC offers both rack and blade servers, providing flexibility to build the optimal CI stack.

The best server type for a particular situation depends on the system design criteria. Consider the following information:

- Rack servers scale up or out with the addition of one server at a time, with one, two, or four CPU sockets per server.
- Rack servers can have multiple disk drives onboard.
- Each rack server consumes between 1U and 4U of rack space.
- Each rack server has redundant power supplies.
- Blade enclosures scale with the addition of blade servers until the blade enclosure is filled to capacity, after which an additional enclosure is added.
- A typical blade enclosure requires 7U of rack space.
- Blade enclosures have n+1 power supplies for redundancy.
- Selecting the appropriate PowerEdge server can reduce CPU socket-based licensing costs and achieve greater virtual machine (VM) density.

#### Server memory and performance considerations

You can configure memory in various modes from within the BIOS. Optimizer Mode is the default mode and is recommended for most virtualization use cases to provide optimized

memory performance. For improved reliability and resiliency, other modes, such as mirror mode and spare mode, are available.

PowerEdge servers support various BIOS configuration profiles that control the processor, memory, and other configuration options. We recommend the following configuration settings:

- Select the Performance Optimized system profile and enable the Virtualization Technology option.
- In System Profile Settings, disable C1E and C States to ensure the highest performance in a virtualized environment.
- Enable Intel Hyper-threading and Virtualization features.

### **Rack servers**

PowerEdge rack servers help in building a modern infrastructure that minimizes IT challenges. Choose from a complete portfolio of 1-, 2-, and 4-socket rack servers to deliver high core density for traditional applications, virtualization, and cloud-native workloads. Enhanced memory speeds, faster NVMe storage options, and BIOS tuning help provide workloads with the performance they require. Intelligent automation throughout the IT lifecycle enables customers to maximize resource utilization and lower costs. Integrated security helps protect the data center from unauthorized changes and cyber-attacks.

### **Blade servers**

Dell EMC PowerEdge Modular Infrastructure with MX, FX, VRTX, and PowerEdge blade servers are built from the ground up for the software-defined data center. The PowerEdge platform is designed to support the right balance of density, capacity, and flexibility, and customizable modules of compute, storage, and networking are easily and rapidly scaled and managed. The OpenManage portfolio delivers innovative systems management that makes administrators more efficient and the infrastructure more productive and reliable.

## Storage

You can select a Ready Stack storage array from a broad range of systems in the Dell EMC portfolio to ensure that you get the storage technology that best fits your requirements. Storage arrays can be configured as part of the initial deployment and can be added later as well.

---

**Note:** This section includes an overview of recommended storage systems for Ready Stack. For links to more information about these systems, see [References](#) on page 36.

---

### Storage design considerations

Dell EMC provides a complete portfolio of storage array technologies, all of which are suitable for deployment in a Ready Stack architecture. Key considerations when choosing the appropriate storage include:

- Storage types
- Performance throughput
- Cost
- Reliability

The following table summarizes the recommended storage arrays.

**Table 2. Recommended Ready Stack storage arrays**

| Storage array      | Description   | Storage type   | Capacity                                | Use cases  |
|--------------------|---|----------------|---|--|
| Dell EMC PowerMax  | Hybrid or all-flash storage with enterprise data services for consolidation of mission-critical workloads | Block and file | Multi-petabyte                          | <ul style="list-style-type: none"> <li>• Consolidation of enterprise applications and databases</li> <li>• Open systems</li> </ul>   |
| Dell EMC Unity     | Simple, unified all-flash or hybrid storage with hybrid cloud   | Block and file | Multi-petabyte                          | <ul style="list-style-type: none"> <li>• SAN</li> <li>• NAS</li> <li>• Mixed workloads</li> <li>• Storage and server consolidation</li> <li>• Enterprise applications</li> </ul>                 |
| Dell EMC XtremIO   | All-flash storage with extreme performance and inline data reduction                                      | Block          | Multi-petabyte                          | <ul style="list-style-type: none"> <li>• Enterprise applications including virtual and database environments</li> <li>• Integrated copy data management</li> </ul>                               |
| Dell EMC SC Series | Affordable and efficient all-flash or hybrid storage  | Block and file | Multi-petabyte                          | <ul style="list-style-type: none"> <li>• General-purpose SAN and NAS workload consolidation</li> <li>• Virtual desktop infrastructure (VDI)</li> <li>• High-volume OLTP</li> </ul>               |
| Dell EMC Isilon    | Scale-out NAS solution  | File           | Over 90 petabytes in a single namespace | <ul style="list-style-type: none"> <li>• General-purpose file workloads</li> <li>• High-performance computing</li> <li>• Media and entertainment</li> <li>• Active and deep archiving</li> </ul> |

Consolidated block and file storage reduces costs and complexity while increasing business agility. You can use data services such as storage provisioning and dynamic host I/O limits across block and file storage.

---

**Note:** The Ready Stack design and deployment guides discuss specific Ready Stack designs; one design incorporates PowerMax storage while the other incorporates Unity storage. You can use other Dell EMC storage products as you design your own configuration.

---

### Dell EMC PowerMax storage

The Dell EMC PowerMax family of storage systems provides simplicity, modern design, and flexible deployments. PowerMax systems implement an integrated architecture for block, file, and VMware VVols, with concurrent support for native NAS, iSCSI, and FC protocols. PowerMax systems support file and block environments, point-in-time snapshots and thin clones, synchronous and asynchronous replication, built-in encryption, tiering to the cloud, and deep ecosystem integration with VMware, Microsoft, and OpenStack.

The PowerMax family consists of two models: The PowerMax 2000 is designed to provide customers with efficiency and maximum flexibility in half of a standard 19-in. rack; the PowerMax 8000 is designed for massive scale, performance, and I/O density, all within a 2-floor-tile footprint.

### Dell EMC Unity storage

Dell EMC Unity storage, powered by Intel Xeon processors, is well-suited for midsize deployments, remote office/branch office locations, and cost-sensitive mixed-workload environments. Unity storage is available in all-flash and hybrid flash arrays and in a software-defined virtual edition. With all-inclusive software, new differentiated features, and Internet-enabled management, Unity storage, which can require as little as 2U rack space, delivers simplicity and value while enabling organizations to speed deployment, streamline management, and seamlessly tier unified storage workload to the cloud.

### Dell EMC XtremIO storage

Dell EMC XtremIO storage delivers great efficiencies and business agility. The XtremIO array maximizes inline deduplication and compression, and provides application-integrated efficient copy services. It enables consistent performance with data services, providing speed, efficiency, and ease of use. The XtremIO X2 platform is ideal for virtual infrastructure use cases, especially those for VDI and large-scale copy data management. XtremIO software provides rapid in-memory computational mechanisms. The XtremIO architecture has a scaled-out, distributed key-value store that identifies data with a unique fingerprint and associates data content with data location for efficient data storage, movement, replication, and identification.

### Dell EMC SC Series storage

Dell EMC SC Series storage includes intelligent all-flash and hybrid solutions that are powered by Intel Xeon processors. SC Series storage provides a flexible tiered storage architecture, outstanding performance, and future-proof value.

End-to-end flash enables high levels of productivity. All-flash SC storage has an active/active, performance-centric design that keeps IOPS and throughput high during scale up and scale out.

## Dell EMC Isilon NAS

Dell EMC Isilon NAS enables you to store, manage, and protect unstructured data with efficiency and massive scalability. Isilon scale-out NAS systems are designed for demanding enterprise file workloads. Choose from all-flash, hybrid, and archive NAS platforms, all of which are powered by Intel Xeon processors, to:

- Simplify management of large data
- Cut costs with over 80 percent utilization and automated tiering
- Scale from tens of terabytes to tens of petabytes
- Increase operational flexibility with multiprotocol support

## SAN switches

For FC storage traffic, Ready Stack uses Dell EMC Connectrix B-Series and MDS Series FC switches. A wide range of Connectrix switches and director-class products, encompassing port densities from 24 to 768 ports, provides great flexibility in creating core/edge SAN topologies to handle any FC storage traffic design.

Brocade Web Tools, an embedded UI on the Connectrix switches, enables administrators to monitor and manage single or small fabrics, switches, and ports. Web Tools is launched directly from a web browser or from the Brocade Network Advisor. In addition, for Connectrix B-Series switches, [Connectrix Manager Converged Network Edition](#) is also available to help manage data center SAN fabrics. For Connectrix MDS Series switches, [optional software license key packages](#) provide additional SAN and traffic-management features.

## SAN design

For the Ready Stack SAN design, consider the following recommendations:

- For SAN arrays, always design for two FC fabrics to ensure high availability connectivity for the storage array.
- Base the number of port connections, whether SAN or IP, on throughput requirements and tolerance for degraded mode operation during a path failure.
- When using vSphere:
  - Multiple datastores within the vSphere cluster enable the use of vSphere HA Datastore Heartbeat. Enabling datastore heartbeating ensures that partitions or isolated host networks do not trigger VM movement within the cluster.
  - VMware currently supports a maximum datastore size of 64 TB and 2,048 powered-on VMs per VMFS datastore. However, in most circumstances and environments, a target of 15 to 25 VMs for each 500 to 750 GB datastore is the conservative recommendation. LUNs and VMware vSphere datastores can be easily expanded to address future growth. Maintaining a smaller number of VMs per datastore greatly reduces the potential for I/O contention, which results in more consistent performance across the Ready Stack environment.
  - Using thin provisioning within VMware on virtual disks does not initially result in additional space efficiency when thin provisioning is enabled on the array. However, the ability to reclaim space from within a compatible guest operating system requires that thin provisioning be used on both the storage and the virtual disks.

- When using Hyper-V:
  - Determine the appropriate disk size for your workload. The dynamically expanding disk type works well for most workloads on Unity All-Flash. Because Unity arrays use thin provisioning, only data that is actually written to a virtual hard disk consumes space on the array, regardless of the disk type—fixed, dynamic, or differencing. Therefore, the best disk type to choose depends more often on the workload rather than how it will impact storage utilization. For workloads generating I/O, such as Microsoft SQL Server databases, Microsoft recommends using the fixed-size virtual hard disk type for optimal performance.
  - Use separate storage LUNs for data and log files. To rule out disk contention, Dell EMC recommends using separate LUNs. For best performance, create distinct LUNs for a server's data files and log files.

## Networking

For Ethernet network traffic, the preferred choice for Ready Stack is Dell EMC Networking switches. Using Dell EMC Networking OS10 or select third-party network operating systems and tools, Dell EMC IP switches deliver nonblocking performance for workloads that are sensitive to packet loss.

Key features include scalable Layer 2 and Layer 3 Ethernet switching with Quality of Service (QoS) and a full complement of standards-based IPv4 and IPv6 features, including Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) routing support. Layer 2 multipathing is supported through Virtual Link Trunking (VLT) and multiple VLT (mVLT) multichassis link aggregation technology.

Dell EMC Networking OS10 has been tested and hardened to meet stringent requirements for reliability, scalability, and serviceability. Networking OS10 supports the full portfolio of Dell EMC Networking data center switch products and enables building cost-effective, end-to-end networks while reducing operational complexity. A fully compliant, industry-standard CLI enables certified engineers to be productive from day one. The CLI provides the primary method to configure, monitor, and administer Networking OS10 applications and switches.

### Dell EMC Networking best practices

The network architecture employs a VLT connection between the two top-of-rack (ToR) switches. The inherent redundancy of a non-VLT environment requires standby equipment, which drives up infrastructure costs and increases risks. In a VLT environment, all paths are active, adding immediate value and throughput while still protecting against hardware failures.

VLT technology enables a server or bridge to uplink a physical trunk into more than one Dell EMC Networking switch by treating the uplink as one logical trunk. A VLT-connected pair of switches acts as a single switch to a connecting bridge or server. Both links from the bridge network can actively forward and receive traffic. VLT provides a replacement for Spanning Tree Protocol (STP)-based networks by providing both redundancy and full bandwidth utilization using multiple active paths. Major benefits of VLT technology include:

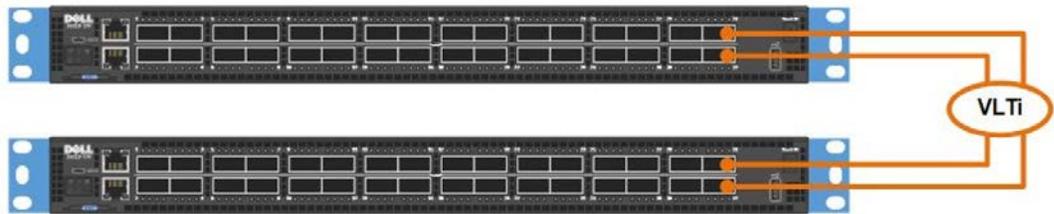
- Dual control plane for highly available resilient network services
- Full utilization of the active link aggregation group (LAG) interfaces
- Active/active design for seamless operations during maintenance events

### Network configurations

Ready Stacks can be built with Dell EMC Networking switches. Networking products for the data center include:

- Economical S-Series switches such as the S3048, S4048, and S4148 for OOB management
- High-performance S5058 and the new 25 GbE S5232 as ToR switches—leaf switches in a leaf/spine topology—for compute clusters
- Extreme-performance Z-Series switches, such as the Z9264F at the spine layer, and new multiprotocol embedded blade switches such as the MX9116n for the PowerEdge MX7000 series blade-server chassis

The following figure illustrates the Dell EMC Networking S5232-ON VLTi configuration.



**Figure 4. Networking S5232-ON VLTi configuration**

In a Ready Stack reference architecture that includes an MX7000 PowerEdge blade server, PowerEdge MX blade-server network traffic is supported by the MX7116n chassis switch's internal network ports that connect to the blade server's network ports.

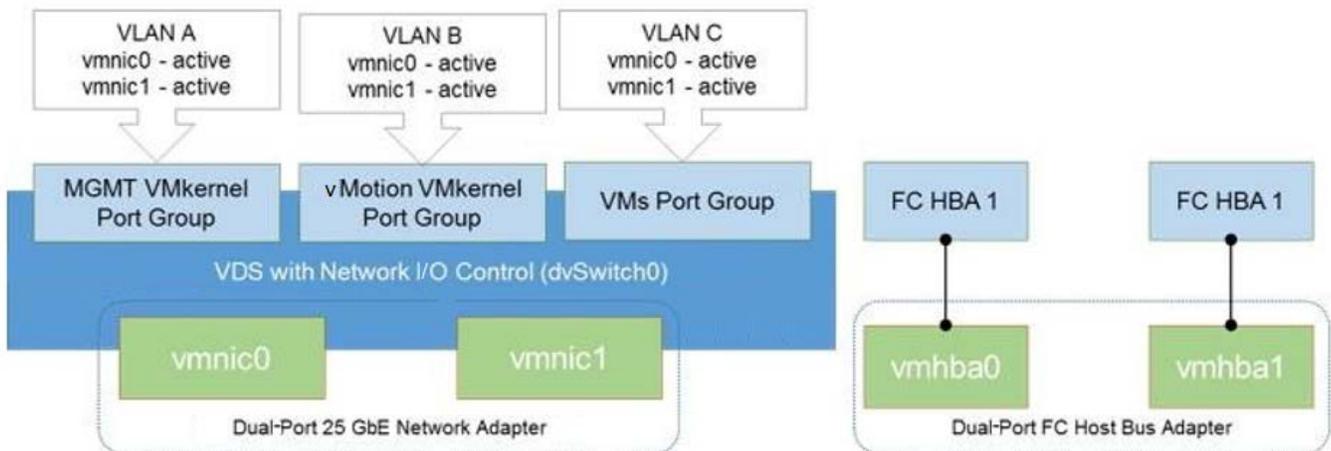
In a Ready Stack reference architecture that uses Dell EMC PowerEdge rack servers instead of PowerEdge blade servers, rack servers are used for compute and management.

### vSphere Distributed Switch LAN traffic network configuration

When using VMware as the hypervisor, customers can achieve bandwidth prioritization for different traffic classes such as host management, vSphere vMotion, and VM network using vSphere Distributed Switch (VDS) technology. The VDS, which can be configured, managed, and monitored from a central interface, provides:

- Simplified VM network configuration
- Enhanced network monitoring and troubleshooting capabilities
- Support for network bandwidth partitioning when NIC Partitioning (NPAR) is not available

The following diagram shows the VDS configuration. The dual-port FC host bus adapters are used for connecting to shared storage.



**Figure 5. VDS configuration**

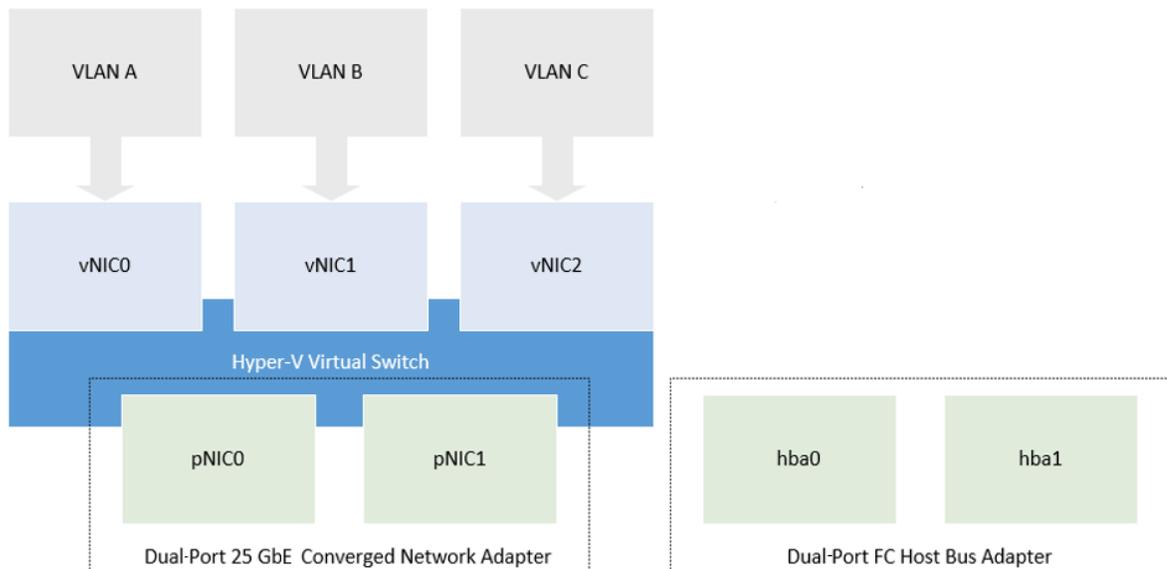
## Hyper-V Virtual Switch LAN traffic network configuration

Hyper-V Virtual Switch is a software-based layer-2 Ethernet network switch that is available in Hyper-V Manager when you install the Hyper-V server role. Hyper-V Virtual Switch includes programmatically managed and extensible capabilities to connect VMs to both virtual networks and the physical network.

Hyper-V Virtual Switch provides:

- Simplified VM network configuration
- Policy enforcement for security, isolation, and service levels
- Enhanced network monitoring and troubleshooting capabilities
- Support for network traffic segmentation

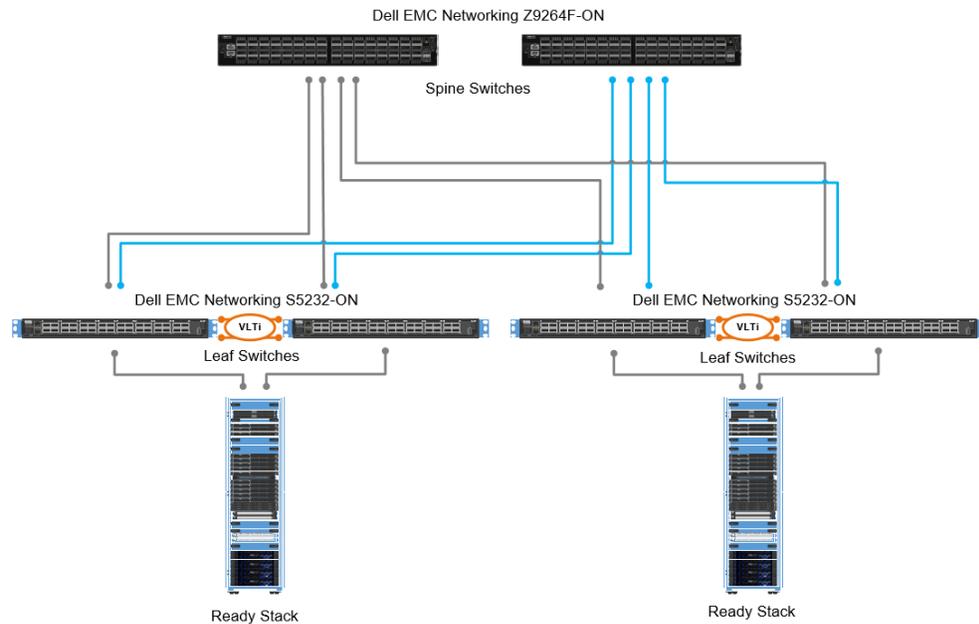
The following figure shows the Hyper-V Virtual Switch configuration for the management and compute servers. The dual-port FC host bus adapters are used for connecting to shared storage.



**Figure 6. Hyper-V Virtual Switch for dual port configuration**

## Network design for Ready Stack scaling

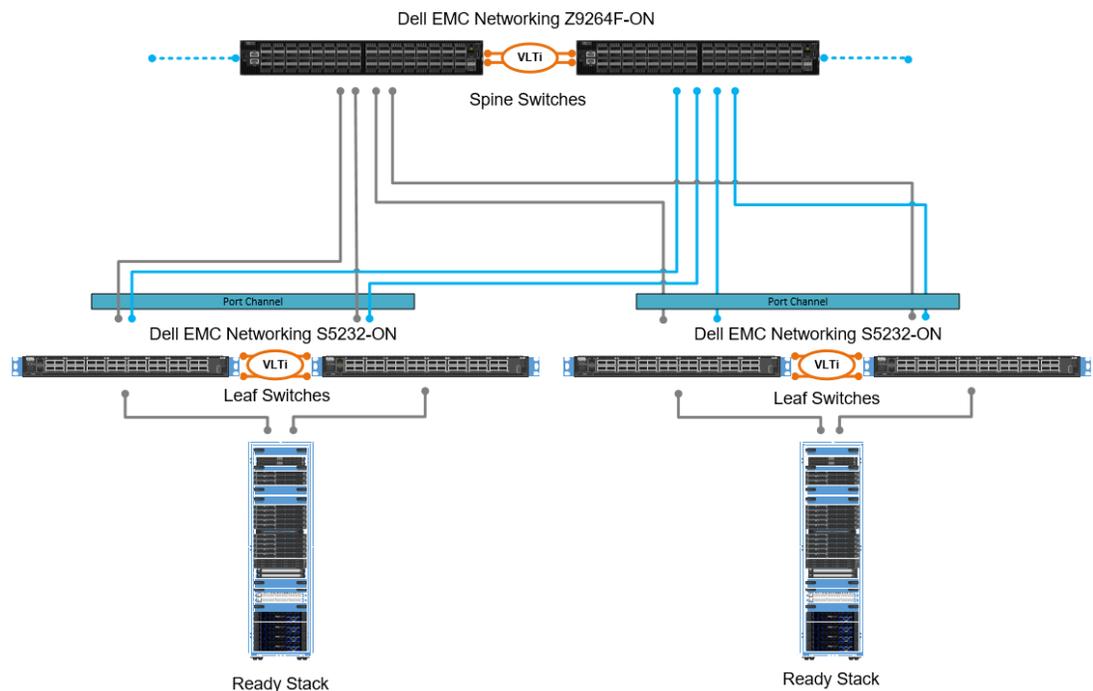
In a Ready Stack reference architecture that includes a PowerEdge MX7000 blade chassis, you can scale Ready Stack by adding multiple compute nodes (pods) in the customer data center. You can use the Dell EMC Networking Z9264 switch to create a simple yet scalable network, as shown in the following figure. The Z9264 switches serve as the spine switches in the leaf-spine architecture. The Z9264 is a multiline rate switch that supports 10/25/40/50/100 GbE connectivity and can aggregate multiple racks with little or no oversubscription.



**Figure 7. Multiple compute pods scaled out using leaf-spine architecture**

When connecting multiple racks, using the 40/100 GbE uplinks from the rack, you can build a large fabric that supports multiterabit clusters. The Z9264 density enables flattening the network tiers and creating an equal-cost fabric from any point to any other point in the network.

For large domain layer-2 requirements, extended VLT (eVLT) can be used on the Z9264, as shown in the following figure. The VLT pair that is formed can scale in terms of hundreds of servers inside multiple racks. Each rack has four 40/100 GbE links to the core network, providing bandwidth for traffic between racks.



**Figure 8. Multiple compute pods scaled out using eVLT**

## Management infrastructure

This section provides an overview of the Ready Stack management infrastructure as well as the software components that run on VMs within the management cluster.

The management infrastructure typically consists of at least two PowerEdge R440 or R640 servers that form a management cluster. Dell EMC recommends using three PowerEdge servers for HA clustering. Management components are virtualized to provide high availability. Redundant 10/25 GbE uplinks to the network infrastructure and redundant 16 Gb/s FC uplinks to the storage array combined with vSphere High Availability (HA) ensure that management components stay online. Two Dell EMC Networking S4148T switches provide OOB connectivity if HA is needed. iDRAC ports in each management and compute cluster connect to those switches.

For example, for a VMware environment, the management software components include:

- VMware vCenter Server Appliance (VCSA)
- Dell EMC OpenManage Integration for VMware vCenter (OMIVV)
- Dell EMC Virtual Storage Integrator (VSI)
- Appliance Configuration Manager for Integrated Data Protection Appliance (IDPA)

### Management servers

The management software components for Ready Stack require a nominal amount of virtual server resources. Although the components could reside on the compute server cluster, alongside the compute workload, you might choose to run these software components on a dedicated management cluster. Doing so results in less resource contention with the compute workloads and also simplifies system recovery in various failure scenarios. Server sizing is based on the specific workload to be run on the management cluster.

### Management software resources

The management software components run on VMs that reside in the management cluster. The following table lists the management components in the cluster and the recommended VM sizing of the components for vSphere and Hyper-V environments.

**Table 3. Ready Stack management components**

| Component                        | Number of VMs   | Number of CPU cores | Number of NICs |
|----------------------------------|---|---------------------|----------------|
| <b>For vSphere environments:</b> |   |                     |                |
| VMware VCSA                      | 1   | 4                   | 1              |
| Dell EMC OMIVV                   | 1   | 2                   | 1              |
| Dell EMC VSI                     | 1   | 2                   | 1              |
| Dell EMC IDPA DP4400             | For details, see: <ul style="list-style-type: none"> <li>• <a href="#">Dell EMC Integrated Data Protection Appliance (IDPA) DP4400 Spec Sheet</a></li> <li>• <a href="#">Dell EMC Integrated Data Protection Appliance DP4400 Installation Guide</a></li> </ul> |                     |                |

| Component   | Number of VMs   | Number of CPU cores | Number of NICs |
|---|---|---------------------|----------------|
| <b>For Hyper-V environments:</b>  |   |                     |                |
| Microsoft SCVMM<br>These values are for the SCVMM server only. For information about the other SCVMM and SQL Server component requirements, see <a href="#">System requirements for System Center Virtual Machine Manager</a> on the Microsoft website. | 1   | 16                  | 1              |
| Microsoft SCOM<br>These values are for the SCOM server only. For information about the other SCOM and SQL Server component requirements, see <a href="#">System requirements for System Center Operations Manager</a> on the Microsoft website.         | 1   | 4                   | 1              |
| Dell EMC OMIMSSC  | 1   | 4                   | 1              |
| Dell ESI for Windows Suite  | 1   | 2                   | 1              |
| Dell EMC IDPA DP4400  | For details, see: <ul style="list-style-type: none"> <li>• <a href="#">Dell EMC Integrated Data Protection Appliance (IDPA) DP4400 Spec Sheet</a></li> <li>• <a href="#">Dell EMC Integrated Data Protection Appliance DP4400 Installation Guide</a></li> </ul> |                     |                |

### OpenManage Essentials console

For administrators looking to reduce the complexity of hardware management, Dell EMC OpenManage Essentials provides an easy-to-use, customizable console for managing Dell EMC hardware, including servers, storage, and switches.

---

**Note:** For information about the next generation of the OpenManage Essentials console, see [Dell EMC OpenManage Enterprise v3.0](#).

---

Basic, repetitive hardware management—updating enterprise servers, and discovering and monitoring servers, storage, and networking—is a reality of life for IT organizations. Often, such tasks are time-consuming and require cumbersome tools and a wide range of skill sets.

To help simplify hardware management and streamline IT operations, the OpenManage Essentials systems management console enables administrators to perform basic hardware management tasks from a single interface that is easy to install and use. OpenManage Essentials supports a range of management functions for Dell EMC servers, storage, and switches, including:

- Discovering and inventorying Dell EMC hardware assets
- Monitoring the health and status of discovered assets

## Management infrastructure

- Sending e-mail alerts for unattended monitoring
- Executing tasks such as remotely powering a server on and off
- Managing and automating server updates

### OpenManage Integration for VMware vCenter

OMIVV streamlines the management processes in data center environments by enabling the use of vCenter to manage the entire server infrastructure, both physical and virtual.

OMIVV functionality includes:

- Monitoring system-level information
- Generating system alerts for action in vCenter
- Rolling out cluster-level BIOS and firmware updates for a VMware ESXi cluster
- Deploying a bare metal infrastructure

### OpenManage Integration for Microsoft System Center

OMIMSSC provides agent-free, operating-system-agnostic, and hypervisor-agnostic configuration and deployment as well as uniform BIOS, driver, and firmware updates for 14th generation PowerEdge servers. Automated processes reduce the steps, time, and costs required to configure the PowerEdge servers from a bare-metal state. OMIMSSC also automates the preparation of remote one-to-many operating system deployments within multivendor operating system and hypervisor environments.

## Virtualization

Ready Stack supports both vSphere and Hyper-V virtualization.

### VMware vSphere

vSphere is highly configurable, which can make it an attractive hypervisor choice for organizations, whether they choose a private or public cloud, or opt for a hybrid cloud approach. Key components and features of vSphere include:

- **VMware ESXi**—Type 1 hypervisor that abstracts processors, memory, storage, and other resources into multiple VMs
- **VMware vCenter Server**—Central control point for data center management services, providing a single view of all ESXi hosts
- **VMware vSphere Client**—HTML5-based management interface that enables users to remotely connect to VMware vCenter Server
- **VMware vMotion**—Feature that enables live migration for powered-on VMs in the same data center
- **VMware Storage vMotion**—Feature that enables the live migration of virtual disks or configuration files to a new data store while a VM is running
- **VMware vSphere HA**—Services that restart failed VMs on other available servers

### Microsoft Hyper-V

Hyper-V is also designed to provide server virtualization for organizations that are implementing private, public, or hybrid clouds. Hyper-V is available as a role in Windows Server and can also be installed as a stand-alone server, known as Hyper-V Server. Both versions ease the learning curve for administrators who have experience with Microsoft products. Hyper-V features include:

- **Computing environment**—A Hyper-V VM includes the same basic parts as a physical computer, such as memory, processor, storage, and networking. All these parts have features and options that can be configured to meet different needs. Because of the many ways to configure them, storage and networking can be considered categories of their own.
- **Disaster recovery and backup**—For disaster recovery, Hyper-V Replica creates copies of VMs, which are intended to be stored in another physical location so that the VM can be restored from the copy. For backup, Hyper-V offers two types. One uses saved states and the other uses Volume Shadow Copy Service (VSS) to make application-consistent backups for programs that support VSS.
- **Optimization**—Each supported guest operating system has a customized set of services and drivers, called integration services. These services, while enhancing performance, make it easier to use the operating system in a Hyper-V VM.
- **Portability**—Features such as live migration, storage migration, and import/export make it easier to move or distribute a VM.
- **Remote connectivity**—Virtual Machine Connection is a remote connection tool for use with both Windows and Linux. Unlike Remote Desktop, this tool provides console access, so the administrator can see what is happening in the guest even when the operating system is not yet booted.
- **Security**—Secure boot and shielded VMs help protect against malware and other unauthorized access to a VM and its data.

## Data protection

All data centers need to protect their applications and data. Thus, it is always a good idea to have a fresh look at data protection requirements whenever a technology refresh is under consideration.

Daily backups ensure the ability to recover from corruptions, human error, malware, ransomware, and so on. Daily backups meet a 24-hour recovery point objective (RPO), limiting total data loss to one day or less.

Typically, organizations need an additional tier of data protection for an estimated 20 percent of their applications. For transactional data, zero data loss is the optimal RPO target. Synchronous local replication and asynchronous offsite replication are valuable tools for delivering a complete data protection solution.

### Dell EMC Integrated Data Protection Appliance

To provide simplicity and reduced risk of single-vendor support, the preconfigured IDPA DP4400, shown in the following figure, combines protection storage and software, search, and analytics in a single 2U appliance.



**Figure 9. Integrated Data Protection Appliance DP4400**

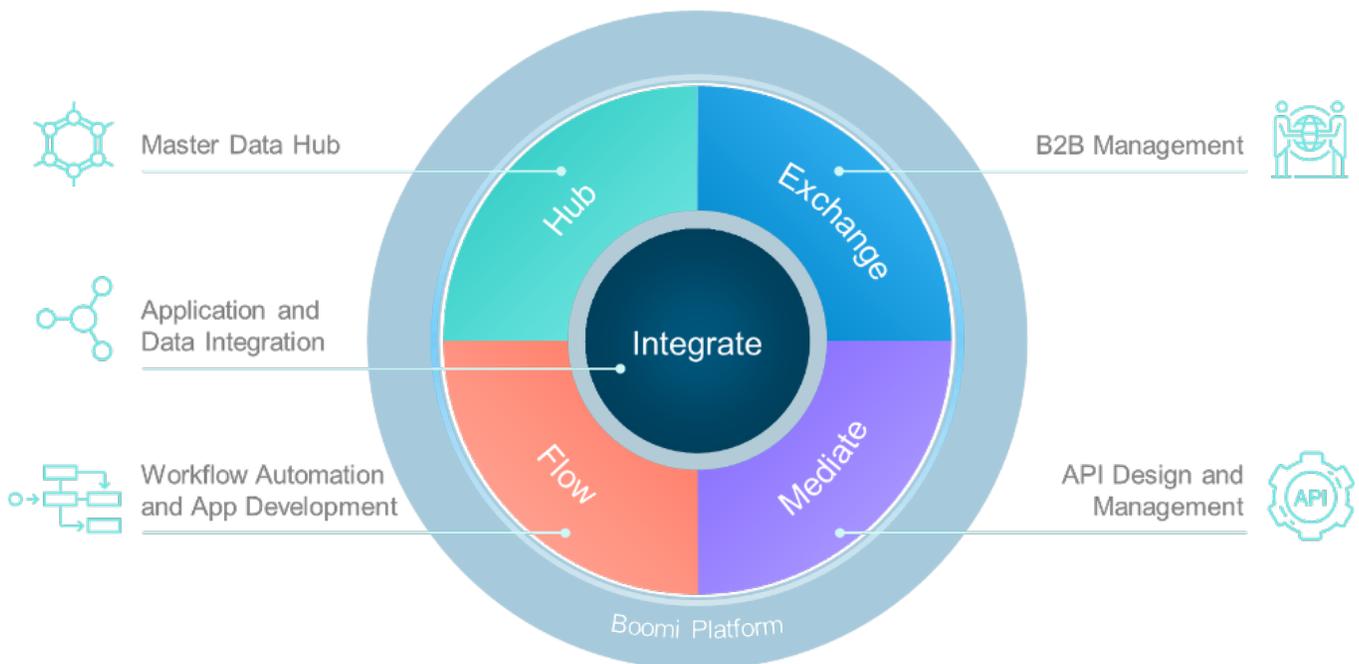
The IDPA DP4400 supports vSphere and Hyper-V environments. Combining power, simplicity, and simplified management through an HTML5-based UI, the IDPA DP4400 supports backup, replication, deduplication, instant access for application development and testing, and cloud disaster recovery.

Additional Dell EMC data protection products are available for Ready Stacks. For more information, see the [Data Protection web page](#) on DellEMC.com.

## Dell Boomi: Application integration

The goal of Ready Stack architectures—to speed overall IT transformation efforts—is not complete without connecting the business. Dell Boomi, from Dell Technologies, is a high-productivity data-management platform that integrates workflow automation and the cloud to build your connected business. Dell Boomi unlocks your data from the device to the edge, to the core, and to one or more clouds to support organizations in their workforce, digital, and IT transformations. Dell Boomi also enables organizations to connect public and private clouds with their on-premises data, facilitating data movement, governance, and management to reduce application sprawl and data silos.

**Unified platform** The Dell Boomi cloud-native, comprehensive, unified platform enables organizations to connect everything—people, processes, applications, data, devices, and things—so they can engage everywhere, across any channel, device, or platform. The low-code platform solves the hard-technical challenges, so you can focus on the “things” driving your business forward. The Dell Boomi platform, as illustrated in the following figure, provides the foundation on which your business can innovate.



**Figure 10. Dell Boomi platform**

The Dell Boomi platform incorporates the following elements:

- **Integrate: Application and data integration**—Connect all your applications and data sources across your hybrid IT landscape to unlock data silos and achieve pervasive integration. Connect public, private, and hybrid clouds and workloads.
- **Hub: Master data hub**—Synchronize and enrich your data through a data hub, so all applications—regardless of where they reside—have trusted data to support business decisions.

- **Exchange: Business-to-business (B2B) and electronic data interchange (EDI) management**—Achieve interoperability between your internal systems and your ecosystem of trading partners to extend the reach of your business operations.
- **Mediate: API design and management**—Modernize and expose underlying data as APIs to deliver scalable and secure real-time interactions for the new world of applications.
- **Flow: Workflow automation and application development**—Automate workflows on any channel, device, or platform to create new experiences for customers, partners, or employees.

**Faster transformation**

Dell Boomi provides a single platform to move, manage, and govern data across your enterprise. Implementations require very low code development, enabling you to build smarter and easier integrations with faster time-to-value. With Dell Boomi, you can easily create and manage access to APIs and data services. Accelerate your transition to a connected business with features such as Boomi Suggest for automated data mapping, Boomi Resolve for automated troubleshooting, and Boomi Assure for automatic updates with managed regression testing.

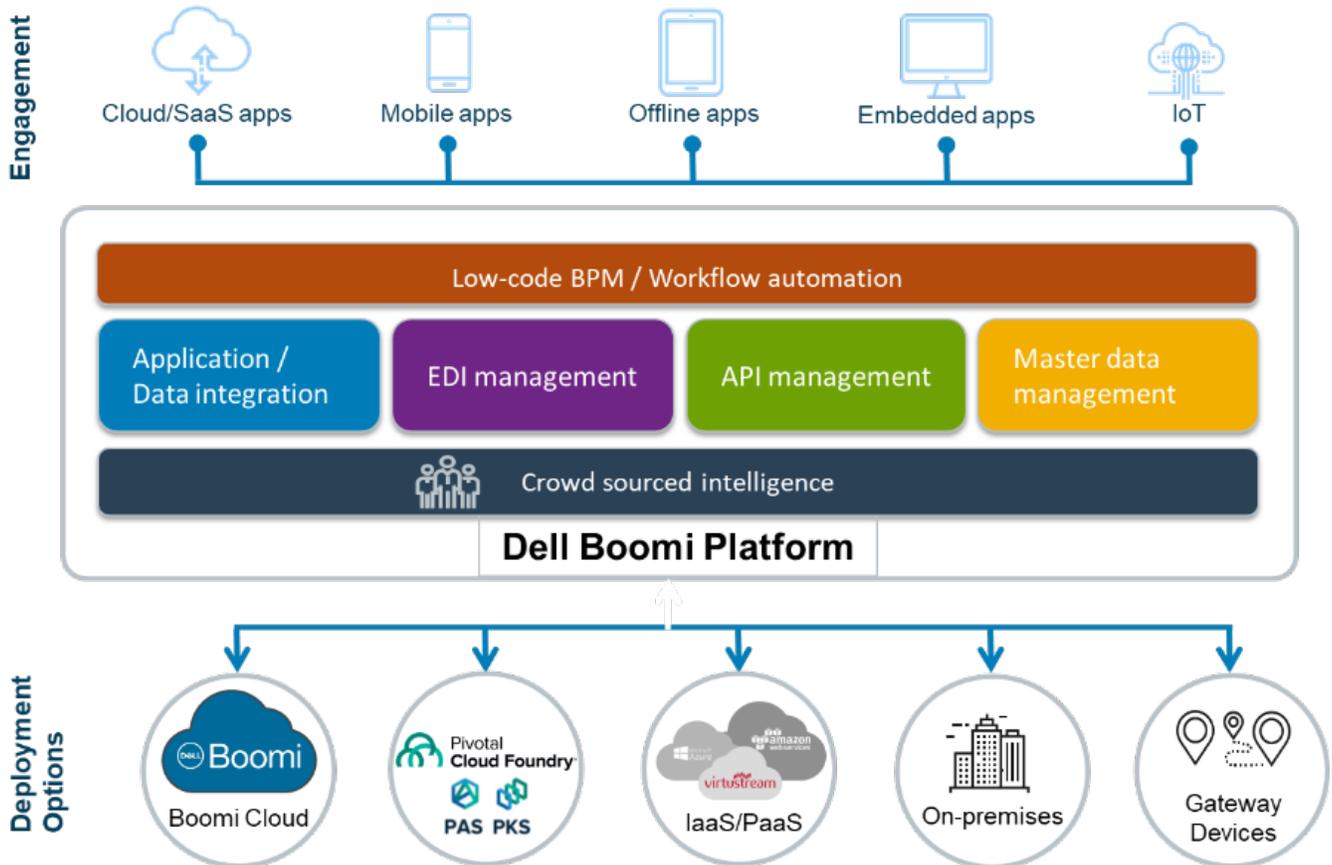


Figure 11. Dell Boomi core features

Running  
anywhere,  
connecting  
everything

The power of the platform is the Atom—the patented, lightweight runtime engine that enables you to deploy your integrations wherever needed: in the Dell Boomi Atom Cloud, in a public or private cloud, or on-premises, as illustrated in the following figure. The Boomi platform provides proven connectivity to more than 200 applications and 1,000 unique endpoints.

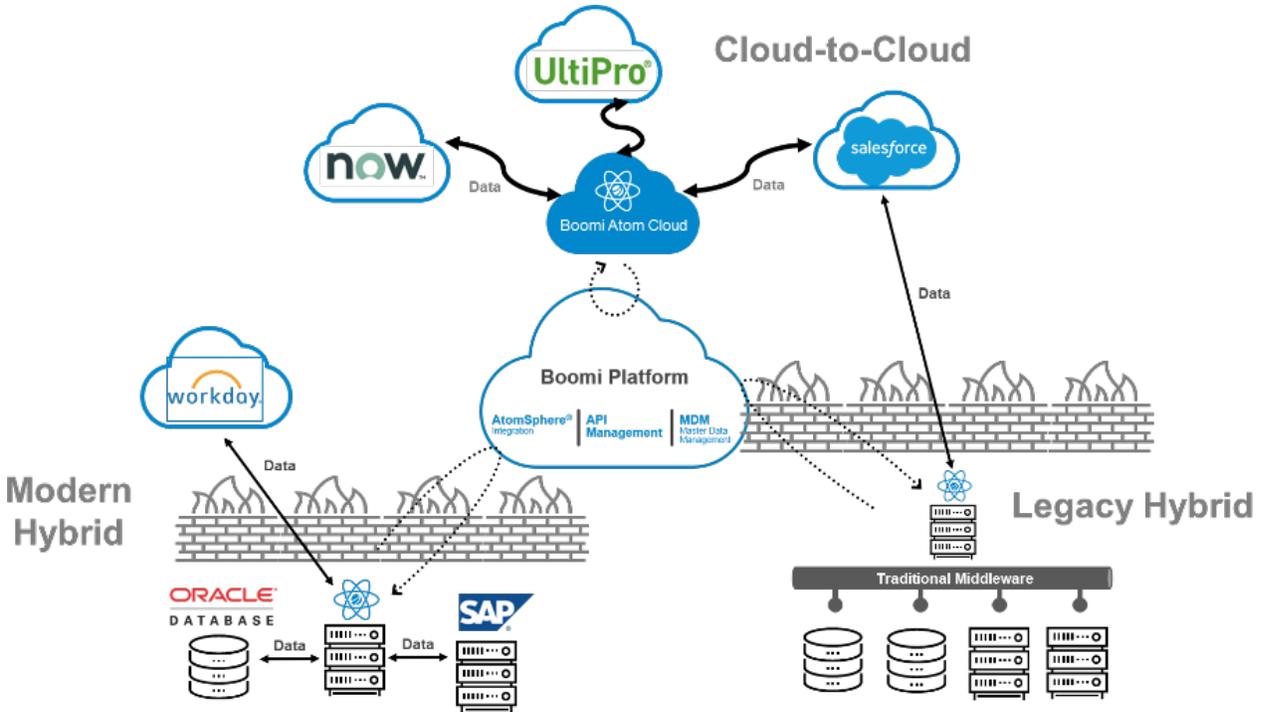


Figure 12. Dell Boomi application connectivity

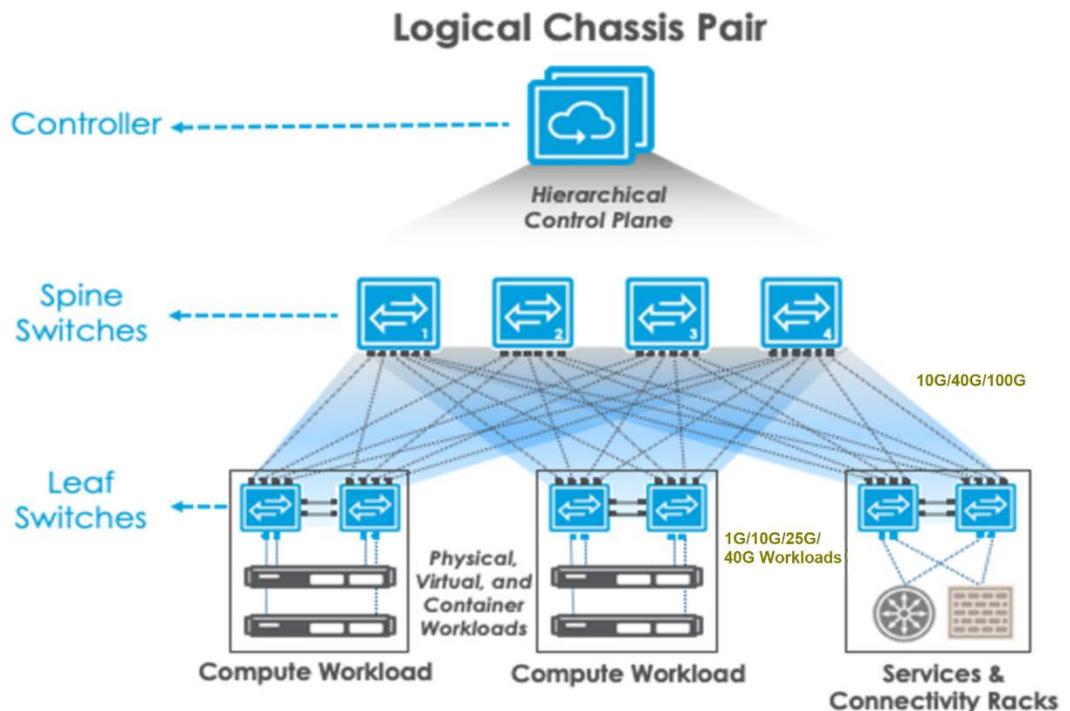
## Big Cloud Fabric: Software-defined networking with Dell EMC Open Networking switches

### Introduction

Big Cloud Fabric (BCF) from Big Switch Networks (BSN) brings hyperscale data center design principles to enterprise data centers. The industry's first software-defined networking (SDN) data center fabric, BCF is powered by a Clos-based fabric architecture that provides physical network automation, visibility, and troubleshooting for VMware environments.

BCF provides high east-west bisectional bandwidth, secure multi-tenancy (enterprise virtual private clouds), and workload elasticity. Customers benefit from application agility due to automation, massive operational simplification due to SDN, and dramatic cost reduction due to hardware/software disaggregation that is enabled by Dell Open Networking switches.

The following figure shows the network topology that supports various workload types.



**Figure 13. Big Cloud Fabric: Support for compute workloads**

BCF supports both physical and virtual (multi-hypervisor) workloads and a choice of orchestration software. It provides Layer 2 switching, Layer 3 routing, and Layer 4–7 service insertion and chaining while ensuring high bisectional bandwidth. The scalable fabric is fully resilient with no single point of failure and supports headless mode operations.

Combined with Dell high-performing Open Networking switches, BCF delivers an SDN solution that enables data centers to drive innovation in VMware SDDC environments. BCF provides network automation for multiple VMware products that include vSphere, NSX, vSAN, and Integrated OpenStack, while delivering new visibility and troubleshooting

capabilities for both VMware and network administrators with the aid of the BCF plug-in for vSphere Web Client and VMware vRealize Log Insight.

## VMware SDDC integration

Integrating BCF with the VMware software-defined data center (SDDC) provides:

- Physical network intelligence through a “one logical switch” operational model, deep fabric-wide visibility, and easy-to-configure service chaining
- Physical network agility through network automation, zero-touch fabric, controller-coordinated upgrading, and rapid VM-to-VM troubleshooting
- Deployment flexibility through Dell EMC Open Networking hardware and scale-as-you-grow options for all application workloads (physical, VM, and container)

BCF integration with VMware vSphere, NSX, vSAN, Integrated OpenStack, and vRealize Log Insight provides network automation and visibility. Network administrators gain visibility into the virtualization environment through the BCF controller and can use the BCF plug-in for vSphere Web Client to provide fabric visibility to the VMware administrator.

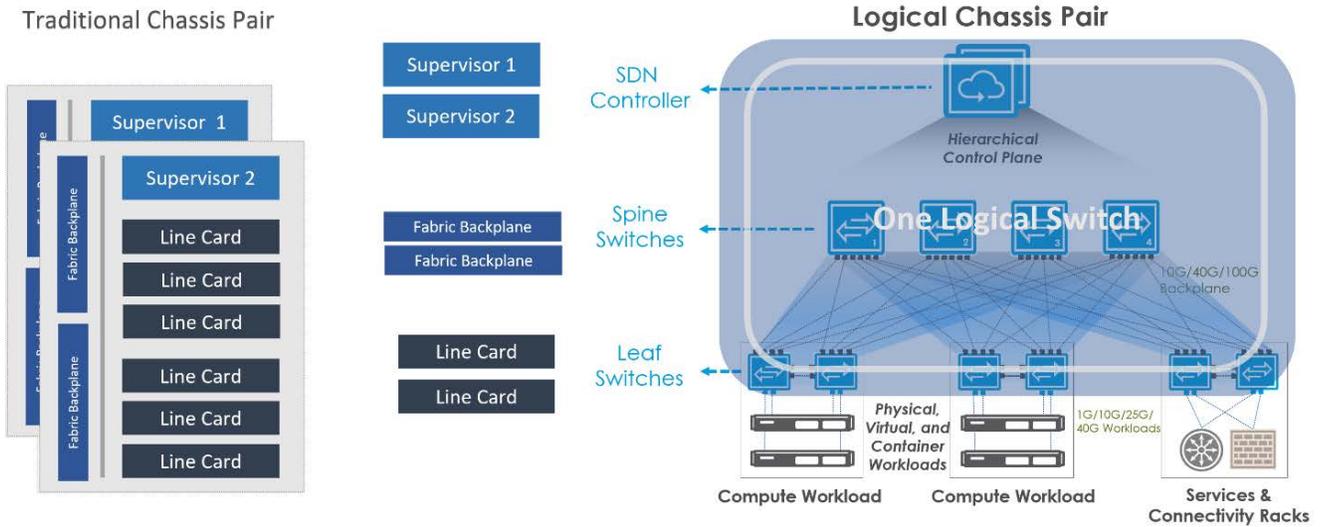
The following figure summarizes the benefits of BCF and VMware integration.



Figure 14. BCF and VMware integration benefits

## One logical switch

BCF is not bound by legacy operational paradigms. Instead, its software-defined fabric acts as one logical switch. This architecture, as depicted in the following figure, drives business velocity by simplifying operations, providing full visibility and telemetry, and enabling rapid service chaining at the network level.



**Figure 15. Logical chassis architecture: “One big switch”**

The network layer is often the least agile part of data center infrastructure to design, configure, and operate, especially when compared to compute infrastructure. Most data center networks are built using old network architecture, which inhibits the ability to meet the demands of modern applications and software-defined data centers. Network overlays consisting of virtual switches have emerged to provide agility but still lack the visibility required to effectively troubleshoot issues.

### BCF controller

Using the BCF controller to perform configuration, automation, and most troubleshooting tasks dramatically decreases the number of management consoles that are required to provision new physical capacity or new logical applications. For example, in a 16-rack pod with dual leaf switches and two spine switches, a traditional network design would have 34 management consoles. The BCF design has only one—the controller console—to perform the same functions. The result is massive time savings, reduced error rates, and simpler automation designs. As a powerful management tool, the controller console exposes a web-based UI, a traditional networking-style CLI, and REST APIs.

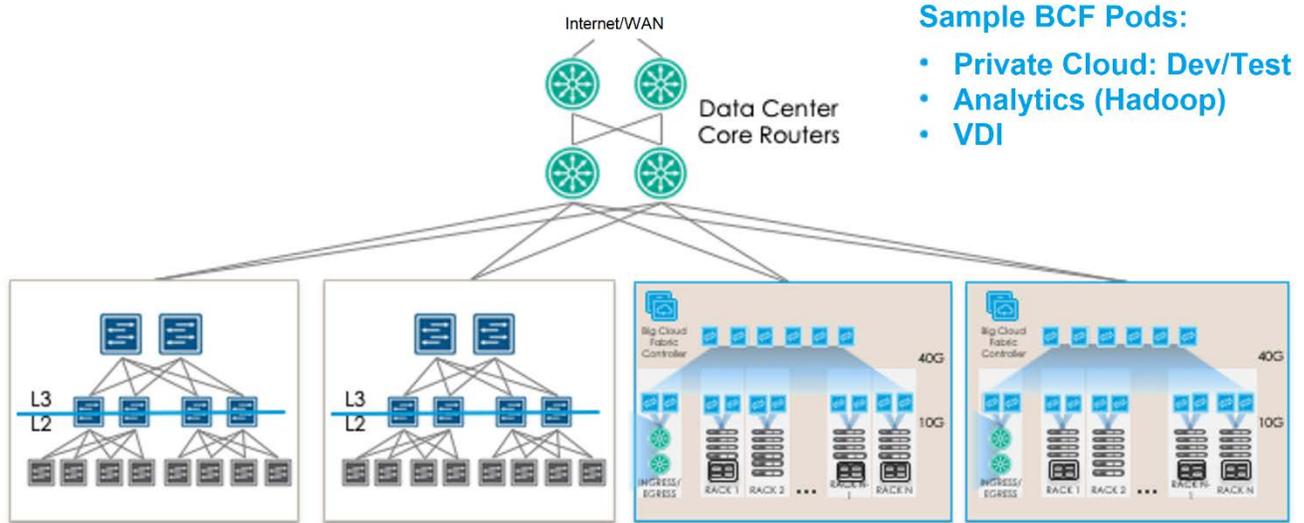
### Zero-touch workflows

In legacy networks, the operational workflows are mostly manual (sometimes scripted, but the scripts must be maintained), resulting in loss of agility. The BCF architecture has built-in network automation. Rapid network upgrades are up to 20 times faster, compared to legacy networks, resulting in lower network downtime, higher agility, and significant savings in operating expenses.

### Brownfield deployments

BCF has a modular architecture that enables enterprises to avoid a “rip and replace” approach when introducing next-generation data center networking technology into their environments. BCF can be introduced in brownfield networking deployments through pods, enabling inter-pod communication through core routers and enabling BCF pods that are positioned across data centers to co-exist with legacy networks. This capability supports a granular and methodical approach to expanding the BCF footprint without impact to existing network infrastructure.

The following figure illustrates the addition of pods to an existing data center.



**Sample BCF Pods:**

- Private Cloud: Dev/Test
- Analytics (Hadoop)
- VDI

Figure 16. Adding BCF pods to an existing data center

**Dell EMC Open Networking switches**

The Dell EMC Open Networking portfolio includes a variety of products for organizations from small businesses to large data centers. Dell EMC Open Networking switches that are coupled with BCF deliver high performance data center fabrics. BCF supports 1 GbE, 10 GbE, 25 GbE, and 40 GbE connectivity to Dell EMC Open Networking switches in a leaf role, and 10 GbE, 40 GbE, and 100 GbE for Open Networking switches in a spine role. BCF supports Dell EMC Networking data-center-class switches such as the S3000, S4000, S5000, S6000, and Z9000 series.

For information about specific switch models, see the [Dell EMC Data Center Networking Quick Reference Guide](#). For more information, contact your local Dell EMC Sales Representative.

**vSphere integration**

BCF integration with vSphere streamlines application deployment workflows by automating the physical network configuration for VMware virtual workloads. The BCF controller's vCenter API extension acts as a single point of integration with vCenter, as shown in the following figure.

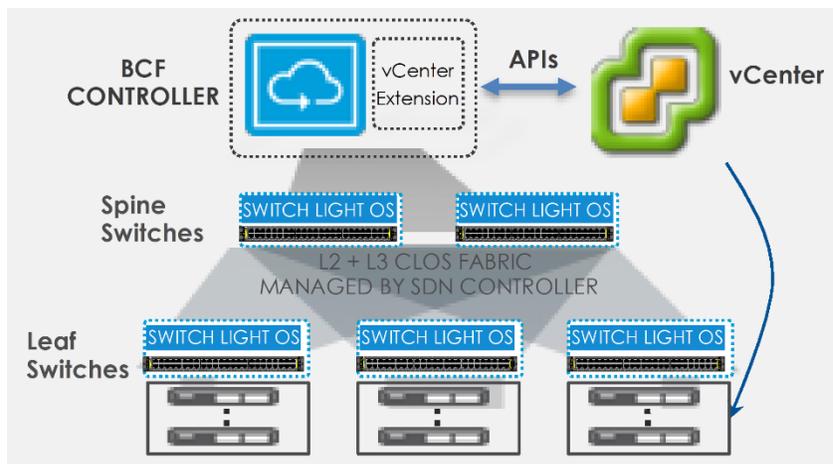


Figure 17. Integrating BCF with vSphere

When the controller is notified of events from vCenter, it performs the corresponding fabric operations, eliminating the need for any manual configuration of the physical network. The integration of BCF and VMware offers the following capabilities, which tremendously simplify network operations in VMware environments:

- Automatic host detection and link aggregation
- Automatic Layer 2 network creation and VM provisioning
- Network policy migration for vSphere vMotion
- Simplified VM-level visibility and troubleshooting
- Fabric analytics for VMware networking

### Networking at the speed of virtualization

SDDC systems might include one or more layers of compute, an overlay network, and multiple storage types. The physical network must be able to interact with all three components. In addition, SDDC systems deploy visibility and troubleshooting tools for the VMware administrator. This means that visibility of the physical network must extend into the SDDC tools for consistent operations and troubleshooting across both virtualization and network domains.

With virtualization becoming mainstream, networks must provide visibility into VMs and east-west traffic across VMs. They must also deliver network service connectivity easily. Emerging cloud-native applications require rapid application and services deployment, which requires network operations to be more automated and less dependent on manual CLI and limited UI workflows. Automation across physical and virtual networks, therefore, becomes a critical aspect of the SDDC automated infrastructure. In addition, gaining visibility across physical and virtual networks is becoming paramount for network and VMware administrators because traditional network troubleshooting is challenging.

### vSAN

The BCF “one big switch” architecture is similar to the vSAN concept of one logical datastore and shares the operational simplicity of vSAN. BCF enables rapid vSAN deployments with automatic configuration of cluster networking and one-click multicast deployment, as illustrated in the following figure. vSAN aware deployment tools in BCF simplify rollout and management, while a dedicated vSAN dashboard in BCF Fabric Analytics provides real-time and historical views of cluster events for rapid troubleshooting.

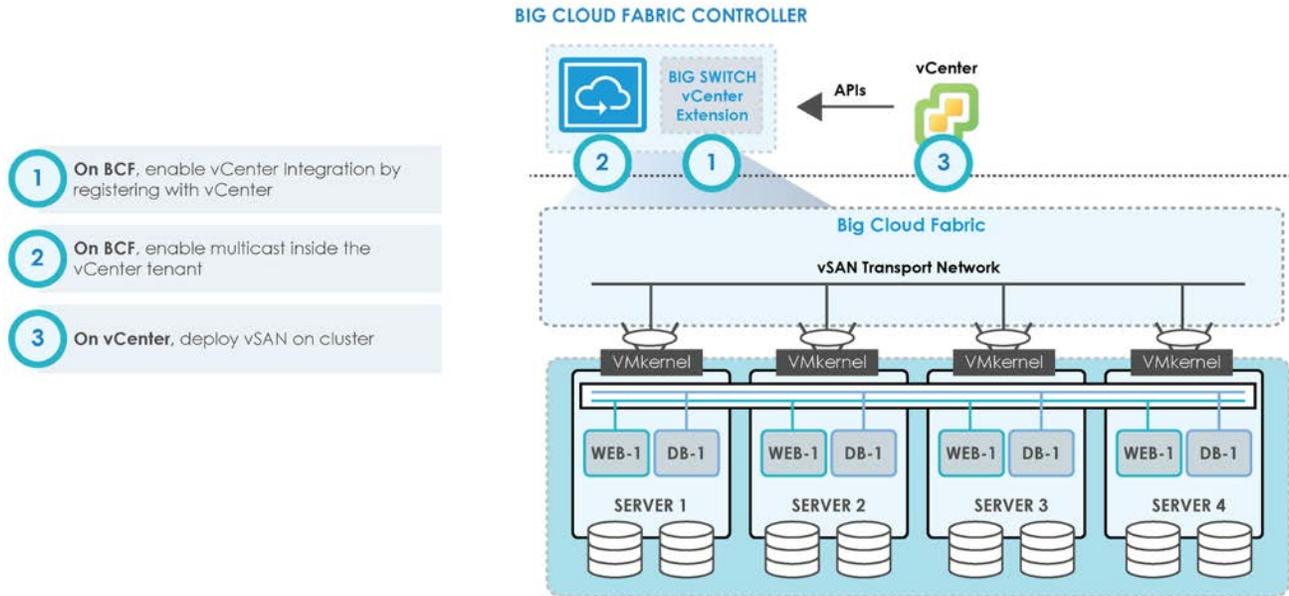


Figure 18. Deploying vSAN

### VM-to-VM troubleshooting

VM-to-VM traffic visibility across the virtual and physical network can be simulated using the BCF test path feature. This feature displays, on the controller, both the physical and logical paths that are taken by the traffic from one VM to another. This level of traffic visibility helps you rapidly determine if an application issue is network-related or compute-related without going through tedious trouble-ticket processes.

The following table shows the various types of VM-to-VM traffic and where they are visible.

Table 4. VM-to-VM traffic

| VMware component     | Integration                            | Visibility                 | Troubleshooting |
|----------------------|--|----------------------------|-----------------|
| vSphere              | Fabric automation                      | VM host visibility         | VM-to-VM        |
| NSX                  | Hardware VXLAN tunnel end point (VTEP) | Overlay visibility         | VTEP-to-VTEP    |
| vSAN                 | vSAN transport network automation      | VMkernel (node) visibility | Node-to-node    |
| vSphere Web Client   | vCenter UI plug-in                     | Fabric visibility          | VM-to-VM        |
| vRealize Log Insight | Content pack                           | Fabric visibility          | Log correlation |

### VMware NSX

BCF seamlessly works with VMware NSX-V, providing automatic transport-network provisioning and enhanced overlay visibility to network administrators. Its one-logical-switch architecture with integrated Layer 2/Layer 3 Equal-Cost Multipath (ECMP) and distributed logical routing dramatically simplifies NSX deployments as compared to legacy networking. Additionally, a Fabric Analytics module provides the network administrator who is operating the BCF underlay with visibility to the NSX-V overlay (NSX-V logical switch and VXLAN virtual network identifier). With BCF eliminating the overlay-underlay visibility gap, NSX-V can be quickly adopted by IT organizations.

The following figure illustrates how BCF can help eliminate the overlay-underlay visibility gap.

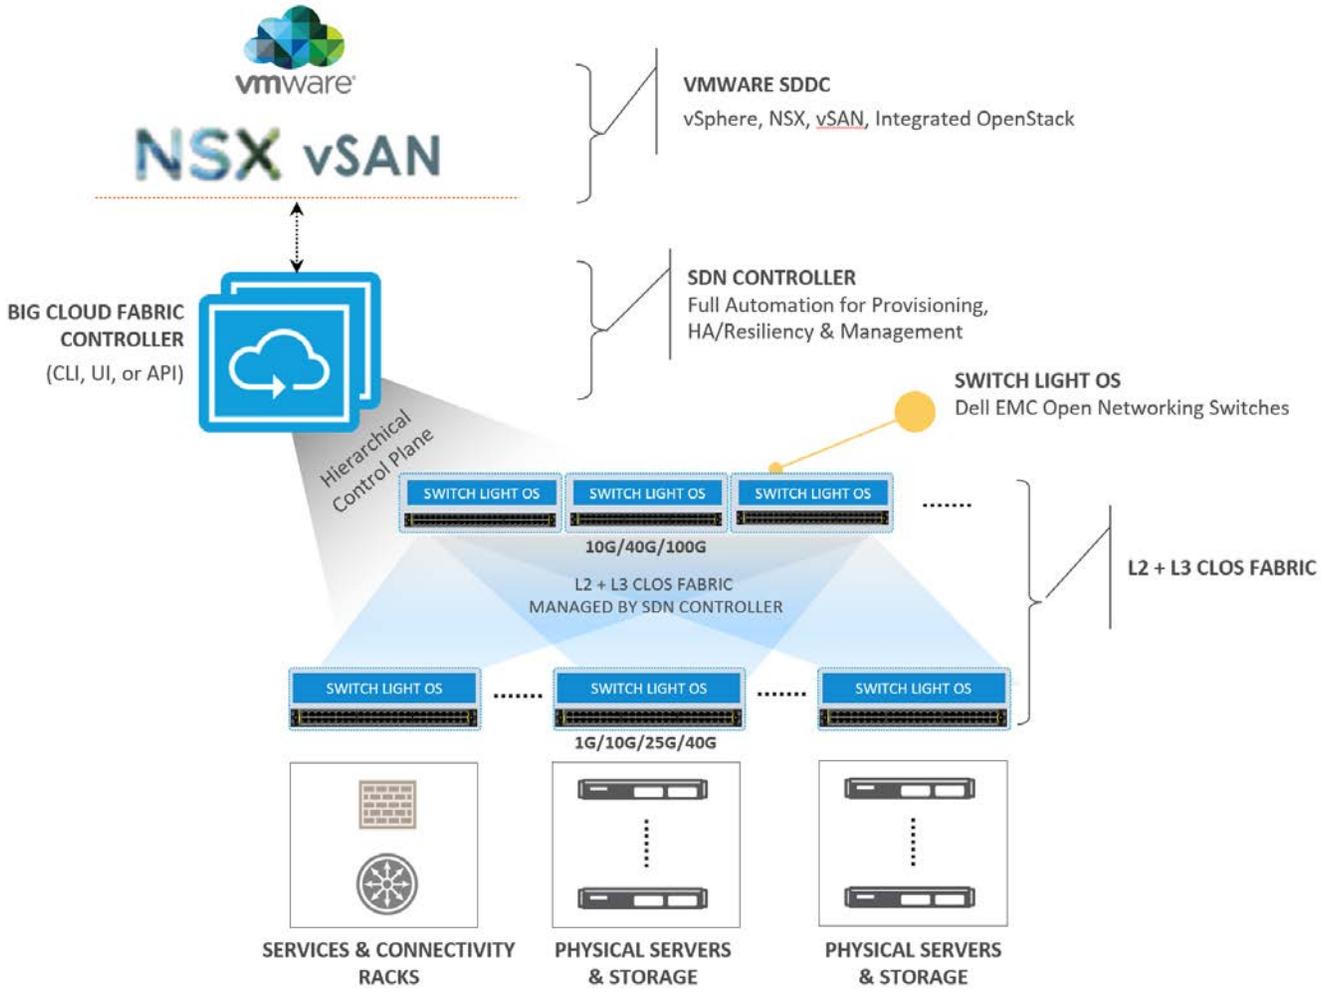


Figure 19. BCF network provisioning and visibility

## Summary

Building a CI platform from the ground up involves a substantial amount of highly technical skills, deep product knowledge across an IT ecosystem of diverse components, significant testing and validation resources, and a rather large risk factor. Cobbling together products from various vendors increases interoperability issues, as well as support complexities, by a factor that is anything but trivial.

Ready Stack from Dell EMC addresses these issues by:

- Using a product portfolio that has been extensively tested together
- Having architectures based on specific hardware and software configurations that complement each other and provide exceptions features, availability, security, and performance levels
- Providing end-to-end support and guidance across the entire stack

Ready Stack brings the benefits of CI without the needless complexities and risks, in an open and flexible architecture, enabling IT organizations to innovate rather than remediate. Ready Stack incorporates many Dell Technologies products and tools, such as iDRAC and VMware plug-ins, that are appropriate for the stack to improve ease of management and agility of the solution.

Add-ons such as Dell Boomi and Big Switch Networks' Big Cloud Fabric (BCF) enhance the value and ease of management. Dell Boomi data management is a comprehensive cloud-native platform that provides connectivity to more than 200 applications and 1,000 unique endpoints. It enables you to move, manage, and govern data across the enterprise. Big Switch Networks' BCF, the industry's first software-defined networking data center fabric, provides physical network automation and enhanced troubleshooting capabilities for VMware products such as vSphere, NSX, vSAN, and Integrated OpenStack.

## References

This section provides links to more information about Ready Stack and related Dell EMC products.

### Ready Stack

For a library of Ready Stack documentation, including design and deployment guides, see [Dell EMC Ready Stack](#) in the Ready Solutions Knowledge Center.

### Servers and server management

The following links provide more information about PowerEdge servers and server management:

- [Igniting Innovation with PowerEdge Server Solutions](#)
- [The Integrated Dell Remote Access Controller 9 \(iDRAC9\) with Lifecycle Controller](#)
- [Dell OpenManage Integration for VMware vCenter](#)
- [Dell EMC PowerEdge: Rack Servers Quick Reference Guide](#)
- [PowerEdge Modular Infrastructure Quick Reference Guide](#)
- [Security in Server Design](#)
- [PowerEdge SAS SSD Performance Specs](#)
- [PowerEdge C Series Quick Reference Guide](#)
- [PowerEdge FC640 Spec Sheet](#)
- [PowerEdge FX2 Enclosure Spec Sheet](#)
- [PowerEdge R440 Spec Sheet](#)
- [PowerEdge R640 Spec Sheet](#)
- [PowerEdge R740 Spec Sheet](#)
- [PowerEdge R740xd Spec Sheet](#)
- [PowerEdge R940 Spec Sheet](#)
- [PowerEdge M1000e Blade Chassis Spec Sheet](#)
- [PowerEdge R7425 Spec Sheet](#)
- [PowerEdge MX7000 Spec Sheet](#)
- [OpenManage Integration for VMware vCenter \(Support web page\)](#)

### Storage

The following links provide more information about Dell EMC storage:

- [Dell EMC Isilon Scale-Out NAS Product Family Data Sheet](#)
- [Dell EMC Isilon OneFS Operating System Data Sheet](#)
- [Isilon All-Flash Scale-Out NAS Storage Specification Sheet](#)
- [Isilon Hybrid Scale-Out NAS Storage Specifications Sheet](#)
- [Isilon Archive Scale-Out NAS Storage Specification Sheet](#)

- [\*Data Reduction with Dell EMC PowerMax: Inline Compression and Deduplication White Paper\*](#)
- [\*The Dell EMC PowerMax Family Overview\*](#)
- [\*Dell EMC SC All-Flash Storage Arrays Specification Sheet\*](#)
- [\*Dell Storage SC7020\*](#)
- [\*Dell EMC SC Series SC5020 Storage Array Specification Sheet\*](#)
- [\*Dell EMC Unity All-Flash Data Sheet\*](#)
- [\*Dell EMC Unity All-Flash Storage Specification Sheet\*](#)
- [\*Dell EMC Service Description: Dell Optimize\*](#)
- [\*Dell EMC XtremIO X2: Next-Generation All-Flash Array Data Sheet\*](#)
- [\*XtremIO X2 Specifications\*](#)

## Networking

The following links provide more information about Dell EMC Networking products:

- [\*Dell EMC Connectrix Storage Networking \(web page\)\*](#)
- [\*Open Networking Solutions for the Modern Enterprise \(web page\)\*](#)
- [\*Dell EMC Connectrix: Storage Networking for the Modern Data Center\*](#)
- [\*Dell EMC Networking Z9264F-ON Series Switch Spec Sheet\*](#)
- [\*Connectrix B-Series Management and Optional Features Solution Brief\*](#)
- [\*Perfect Your Network Modernization\*](#)
- [\*Network Automation with Dell Open Automation White Paper\*](#)
- [\*Data Center Switching Quick Reference Guide\*](#)
- [\*Dell EMC Networking Z9100-ON Series Switches\*](#)
- [\*Dell EMC Networking S5048F-ON\*](#)
- [\*Dell EMC Networking S4100-ON\*](#)
- [\*Dell EMC Networking S4200-ON\*](#)
- [\*Dell EMC Networking S4048-ON\*](#)
- [\*Dell Networking S4048T-ON\*](#)

## Data protection

The following links provide more information about IDPA DP4400 and information about additional Dell EMC data protection products:

- [\*Dell EMC Integrated Data Protection Appliance \(IDPA\) DP4400 Data Sheet\*](#)
- [\*Dell EMC Integrated Data Protection Appliance \(IDPA\) DP4400 Spec Sheet\*](#)
- [\*Dell EMC Data Protection \(web page\)\*](#)
- [\*Dell EMC Avamar Data Sheet\*](#)
- [\*Dell EMC Data Domain Deduplication Storage Systems Spec Sheet\*](#)
- [\*Dell EMC RecoverPoint Data Sheet\*](#)
- [\*Dell EMC RecoverPoint \(specification sheet\)\*](#)
- [\*EMC RecoverPoint for Virtual Machines Data Sheet\*](#)