

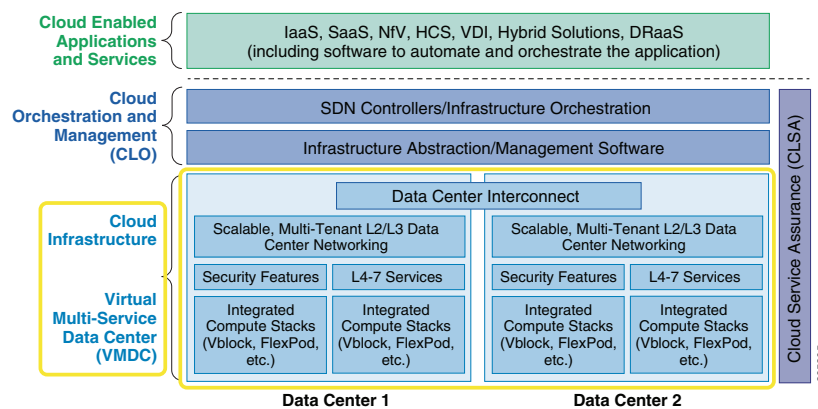


CHAPTER 1

Introduction

The Cisco Virtualized Multiservice Data Center (VMDC) system provides design and implementation guidance for enterprises deploying private cloud services, and for Service Providers (SPs) building public and virtual private cloud services. The Virtual Multi-Service Data Center (VMDC) is Cisco's reference architecture for cloud deployments and has been widely adopted by a large number of service providers and enterprises worldwide. VMDC integrates Cisco and third-party products across the cloud computing ecosystem into a validated end-to-end system that customers can deploy with confidence.

Figure 1-1 Cisco Cloud Systems Foundation



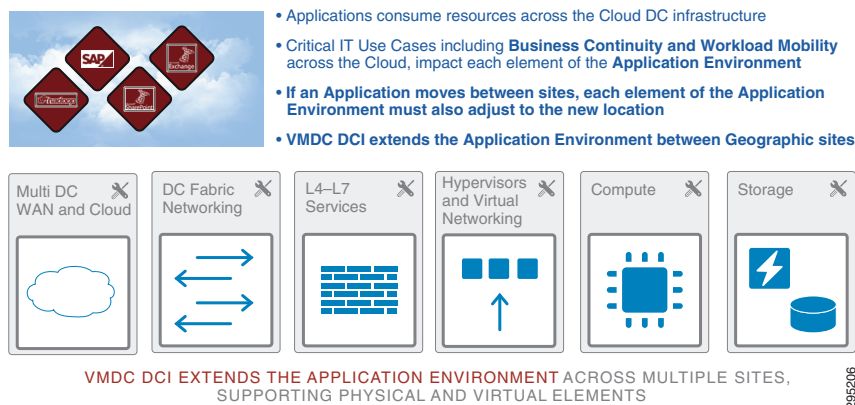
Data Center Interconnect (DCI) refers to underlying technologies used to connect geographically dispersed data centers to support Business Critical operations. This VMDC DCI solution provides validated guidelines for cloud data center connectivity across metro distances (less than 200 km) and geo distances (more than 200 km). This VMDC DCI solution enables critical business operations including:

- Application business continuity across multiple data center sites
- Application disaster recovery and avoidance across multiple data center sites
- Application geo-clustering and load balancing across multiple data center sites
- Complete workload mobility across multiple data center sites
- Operations functions across multiple data center sites including workload rebalancing, Maintenance operations, and consolidation of workloads

Solution Scope for VMDC DCI

The VMDC DCI solution provides metro and geo extensions that enable the interconnection of geographically diverse Cloud data centers. The VMDC DCI system enables elasticity, mobility, and recovery of applications and workloads from one physical data center to another with minimal disruption to the application. Application workloads consume a range of physical and virtual resources across the cloud, as described in Figure 1-2. If an application or workload moves between sites, the application environment must also adjust to the new location. VMDC DCI extends the application environment across sites to enable workload elasticity and more flexible deployment models. The application environment spans a number of critical elements including multi-site WAN connections, data center fabrics, L4-L7 services, hypervisors and virtual switching, compute resources, and storage resources. VMDC DCI extends these elements to unlock a range of business functions including business continuity, disaster recovery and avoidance, workload mobility, active-active data centers, and support of application geo-clusters. VMDC DCI also supports multi-site functions required by operations teams including workload rebalancing between sites, site migrations, and consolidation of workloads between sites.

Figure 1-2 Application Centric approach to Data Center Interconnect

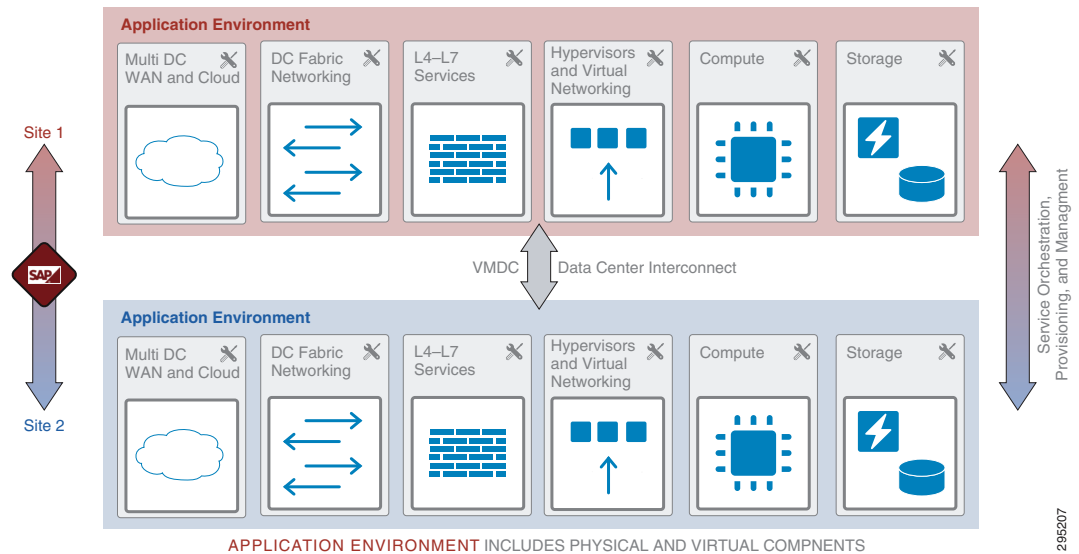


The VMDC DCI system provides design guidance on how the data center infrastructure can more easily support workload mobility and business continuity within Private and Public Clouds. This Cisco VMDC solution addresses how DCI extensions across metro/geo data centers directly impact each element of the application environment. The Application environment within Public and Private Cloud data centers includes many elements. Each element participates in the validated DCI design, providing much needed capabilities to support application mobility between geographic sites. VMDC DCI extends the application environment as described in Figure 1-3, across each element listed below:

- Redirection of external users to the appropriate site.
- L2 extensions between sites to enable workload mobility and the preservation the application's IP addressing.
- Extending data center fabric functions between sites including tenancy, network containers, traffic QoS, and bandwidth reservation.
- Extending L4-L7 Services between sites including service chaining for both physical or virtual services.
- Multi-site hypervisor features supporting workload migrations, extended clusters, and high availability for VMware and Microsoft Hyper-V environments.
- Distributed Virtual switching spanning multiple sites.

- Distributed Compute environment supporting integrated PoDs, with port and security profiles spanning multiple sites.
- Distributed Storage environment including NAS/SAN extensions, virtual volumes, storage fabric, and data replication across multiple sites.
- Service Orchestration, provisioning, and management of the application environment and infrastructure.

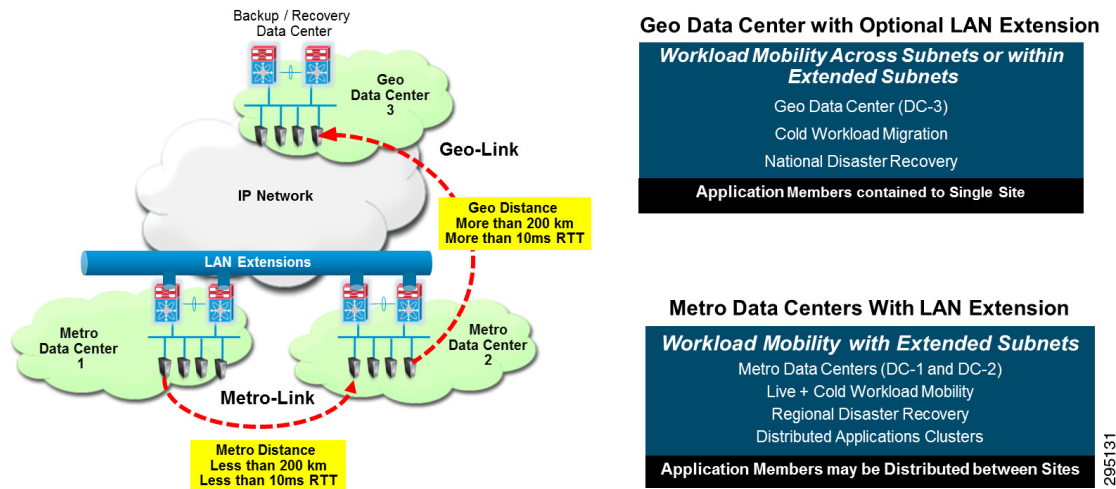
Figure 1-3 VMDC DCI Extends Application Environment Across Multiple Sites



Use Cases/Services/Deployment Models

The three site deployment model described in Figure 1-4 was used as a basis for the VMDC DCI design. This model integrates two data centers at a metro or regional distance of less than 200 km and 10 ms Round Trip Time (RTT). Metro data centers have the ability to operate as a single virtual data center spanning a metro distance, supporting active-active scenarios, live application migrations, and stretched cluster designs. The third data center is at a geo Distance of greater than 200 km and more than 10 ms RTT. The third data center provides a many-to-one recovery capability to a site spanning a much longer distance. The distance limitations to the third site typically force an active-standby operational model and cold workload migrations between independent sites.

Figure 1-4 VMDC DCI Validates a Three Site Data Center Model



The metro data centers support DCI features that enable the following business critical use cases:

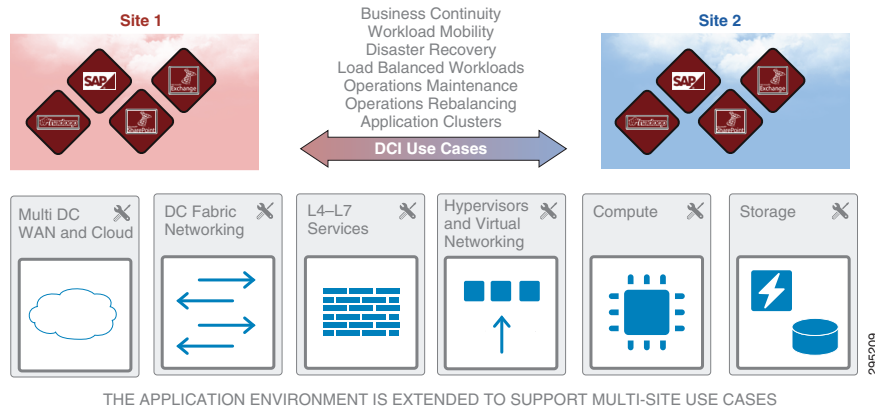
- Live workload mobility between metro data centers.
- Cold workload mobility between metro data centers.
- Stretched clusters and stateful services between metro data centers.
- Active-Active, Active-Standby, and load balanced applications designs between metro data centers.
- Regional disaster recovery between metro data centers.
- Workload rebalancing, operations maintenance activities, and consolidation of live or cold workloads between metro data centers.
- Site migrations of active workloads between metro data centers.

The geo data center located at a further distance supports DCI features that enable the following business critical use cases:

- Cold workload mobility between metro/geo data centers.
- Certain Live Workload Mobility scenarios that support larger network latency between metro/geo data centers.
- Active-Standby and load balanced application designs between metro/geo data centers.
- National disaster recovery between metro/geo data centers.
- Workload rebalancing, operations maintenance activities, and consolidation of cold workloads between metro/geo data centers.
- Site migrations of cold workloads or halted workloads between metro/geo sites.

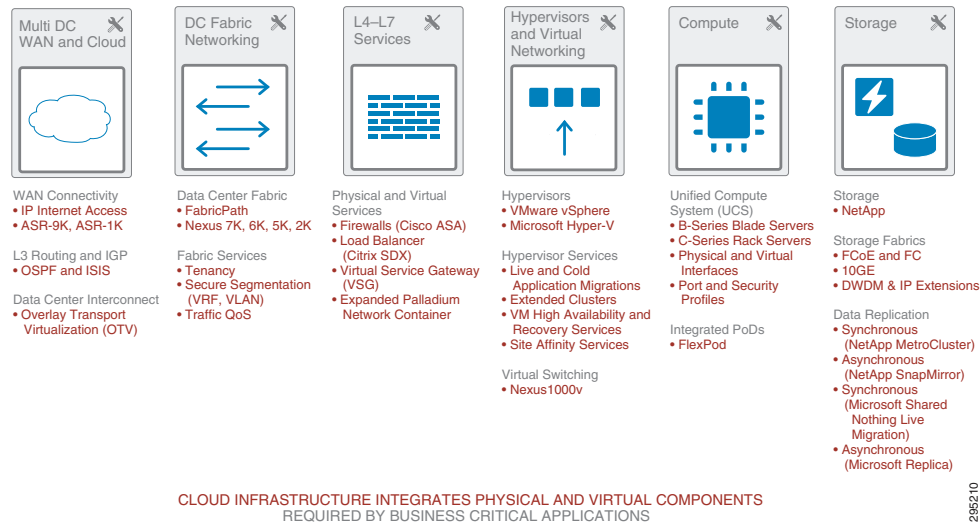
VMDC DCI enables a range of critical business functions and including business continuity and workload mobility (Figure 1-5).

Figure 1-5 VMDC DCI Enables a Range of Critical Business Functions



VMDC DCI infrastructure was validated with a range of products and features, needed to extend the application environment across multiple sites. The summary of infrastructure components is provided in Figure 1-6. Other product options are also available, and are described throughout this document.

Figure 1-6 Infrastructure Components Summary



Key Solution Benefits

The VMDC DCI solution incorporates a wide range of Cisco cloud innovations and Partner products. These products are integrated within the Cisco Validated Design (CVD). The DCI solution provides compelling benefits to Public and Private Cloud Providers:

- **Simplify the DCI Design Process for Operations Teams**—Interconnecting Cloud Data Centers involves many infrastructure elements and application components that provide critical business services. The VMDC DCI design provides a validated reference design that significantly reduces risk of implementation using Cisco’s latest product innovations and partner products. This VMDC DCI design builds upon previous VMDC releases that have been extensively validated and widely deployed by Enterprises and Service Providers worldwide. The validated VMDC DCI design enables Public and Private Cloud Providers to deploy DCI functions with confidence.

- **End-to-end Validation of the Application Environment**—The VMDC DCI solution delivers validated guidelines across the end-to-end layers of the cloud data center. The DCI design spans different sites and addresses each element of the Application environment including WAN connections, LAN extensions, tenancy, network containers, distributed virtual switching, and L4-L7 services; as well as traditional functions such as hypervisor vMotion and storage replication. This is a true DCI solution that directly addresses each element of the Application environment.
- **Validates 2 of the most used DCI Design Options**—VMDC DCI validates the most common design options to achieve 2 major Recovery Point Objective (RPO) and Recovery Time Objective (RTO) targets. The first design option enables the movement of applications, their data, their services, and network containers to support near zero RPO/RTO for the most business critical functions. Less business critical applications can be mapped to a second design option to achieve RPO/RTO targets of 15 minutes or more.
- **Minimal Disruption to the Application**—VMDC DCI allows operators to preserve IP addresses of moved applications, their services, and network container between sites.
- **Reduction in CAPEX/OPEX for DCI Deployments**—VMDC DCI helps customers align the correct DCI design to achieve the selected application RPO/RTO targets. The most stringent recovery targets typically require the highest CAPEX/OPEX. VMDC DCI provides a framework to map Applications to different Criticality Levels, and then select the most cost effective option that meets application requirements.
- **Planned Usage of Recovery Capacity**—Recovery capacity at remote sites can be used for other applications during “normal operations” and “reclaimed” as needed by Operations Teams during recovery events. This “Reuse-Reclaim” design strategy allows for planned utilization of extra capacity and many-to-one resource sharing, reducing CAPEX/OPEX.
- **DCI Use Cases Validated with Business Applications**—VMDC DCI utilized traditional business applications across each workload migration and business continuity use case. The test applications include Oracle database servers, Microsoft SharePoint and SQL, for single tier and multi-tier test applications.
- **Multiple Hypervisors supported**—Both VMware and Microsoft Hyper-V environments are supported. Microsoft Hyper-V design guidance is provided as a separate addendum.
- **Product Performance Measured across DCI Use Cases**—The performance of Cisco products and Partner Products used in VMDC DCI was measured and documented across metro/geo environments. Performance limitations, design recommendations, and configurations are provided for Cisco and Partner products.
- **Operational Simplicity**—This VMDC DCI release utilizes cloud service orchestration and resource provisioning products from Cisco and Cisco partners to support multi-site environments. Automated provisioning of cloud assets significantly simplifies operations, especially across multi-site designs.

Audience

This guide is intended for, but not limited to, system architects, network design engineers, system engineers, field consultants, advanced services specialists, and customers who want to understand how to deploy a public or private cloud data center infrastructure. This guide assumes that the reader has a basic understanding of enterprise and SP network designs and data center architectures.

Related CVD Guides

A brief description of previous VMDC System Releases is provided below for reference.

VMDC 2.x System Releases

In the data center portion of the architecture, VMDC 2.X designs were centered on traditional hierarchical infrastructure models incorporating leading Cisco platforms and Layer 2 (L2) resilience technologies such as Virtual Port Channel (vPC), providing network containers or tenancy models of different sizes and service profiles, with necessary network based services and orchestration and automation capabilities to accommodate the various needs of cloud providers and consumers.

VMDC 3.x System Releases

VMDC 3.X systems releases introduced Cisco FabricPath for intra-DC networks, as an optional L2 alternative to a hierarchical vPC-based design. FabricPath removes the complexities of Spanning Tree Protocol (STP) to enable more extensive, flexible, and scalable L2 designs. Customers leveraging VMDC reference architecture models can choose between vPC-based and FabricPath-based designs to meet their particular requirements.

VMDC Virtual Services Architecture (VSA) System Releases

VMDC VSA is the first VMDC release dealing specifically with the transition to NFV (Network Function Virtualization) of IaaS network services in the data center. Such services comprise virtual routers, virtual firewalls, load balancers, network analysis and WAN optimization virtual appliances.

The VMDC VSA release focuses mainly on public provider use cases, building a new logical topology model around the creation of virtual private cloud tenant containers in the shared data center infrastructure. Future releases will incorporate additional cloud consumer models specific to enterprise and private cloud use cases. In particular, future releases will address hybrid consumer models, comprising physical and virtual service appliances, used together as part of a per-consumer or per-tenant service set. These can be implemented on either a 2.X (classical Ethernet) or 3.X (FabricPath) VMDC infrastructure. However, the initial VMDC VSA release will focus on fundamental implications of an all-virtual approach, and a simple FabricPath data center topology previously validated in VMDC 3.0.

